Cloudera on Premises Data Services 1.5.4

Cloudera Data Services on premises Release Notes

Date published: 2023-12-16
Date modified: 2024-12-20



https://docs.cloudera.com/

Legal Notice

© Cloudera Inc. 2025. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 (“ASLv2”), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Cloudera software product page for more information on Cloudera software. For more information on Cloudera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.

Contents

Data Services Release Notes..... 4

What's new in Cloudera Data Services on premises 1.5.4.....7
Kubernetes Ingress NGINX Controller vulnerabilities.....7

Known issues for the Cloudera Data Services on premises 1.5.4.....9

Fixed Issues for the Cloudera Data Services on premises 1.5.4..... 23

Repository Locations for 1.5.4..... 25

Fixed CVEs.....25

Cumulative hotfixes..... 40
Cloudera Data Services on premises 1.5.4-CHF1..... 40
Whats new in Cloudera Private Cloud Data Services1.5.4-CHF1.....41
Known Issues in Cloudera Private Cloud Data Services 1.5.4-CHF1.....43
Fixed Issues in Cloudera Private Cloud Data Services 1.5.4-CHF1.....43
Repository Locations for 1.5.4-CHF1.....44
Fixed Common Vulnerabilities and Exposures in 1.5.4 CHF1..... 44
Cloudera Data Services on premises 1.5.4-CHF3..... 121
Whats new in Cloudera Data Services on premises 1.5.4-CHF3..... 122
Known Issues in Cloudera Data Services on premises 1.5.4-CHF3.....124
Repository Locations for 1.5.4-CHF3.....126
Fixed Common Vulnerabilities and Exposures in 1.5.4 CHF3.....126

Service packs..... 149
Cloudera Data Services on premises 1.5.4-SP1.....149
Certifications in 1.5.4-SP1.....150
What's new in Cloudera Data Services on premises 1.5.4-SP1.....150
Known Issues in Cloudera Data Services on premises 1.5.4-SP1..... 152
Fixed Issues in Cloudera Data Services on premises 1.5.4-SP1.....154
Repository Locations for 1.5.4-SP1..... 155
Fixed Common Vulnerabilities and Exposures in 1.5.4 SP1.....155
Cloudera Data Services on premises 1.5.4 SP2.....289
Certifications in 1.5.4 SP2..... 289
What's new in Cloudera Data Services on premises 1.5.4 SP2..... 289
Apache Parquet CVE-2025-30065..... 291
Known Issues in Cloudera Data Services on premises 1.5.4 SP2..... 292
Repository Locations for 1.5.4 SP2..... 295
Fixed Common Vulnerabilities and Exposures in 1.5.4 SP2.....295
Cumulative hotfixes: Cloudera Data Services on premises 1.5.4 SP2 CHF1..... 301

Data Services Release Notes

Cloudera Private Cloud Data Services includes the Management Console, Cloudera Data Warehouse (CDW), Cloudera Machine Learning (CML) and Cloudera Data Engineering (CDE). Learn about the new features and improvements in each of these services.

Cloudera Data Services on premises 1.5.4 support up to 7.1.9 SP1.

Certifications

- Cloudera Base on premises (7.1.9 CHF6, 7.1.8 CHF22 , 7.1.7 SP3)
- Cloudera Manager 7.11.3 CHF6
- Iceberg v2 GA on Cloudera Data Warehouse, Cloudera Data Engineering, & Cloudera AI with Ozone
- OEL (RHCK Kernel Only) 8.7, 8.8, 8.9, 9.1, 9.2, 9.3
- RHEL 8.7, 8.8, 8.9, 9.1, 9.2, 9.3
- K8s 1.27 and OCP 4.14
- Longhorn version 1.5.4



Note: Cloudera Manager 7.11.3 CHF6 support Cloudera Data Services on premises 1.5.4 release.

Cloudera Manager 7.11.3 CHF8 does not support any Cloudera Data Services on premises release.

New features in Platform

Stability and Resiliency: New prerequisite check in ECS Install Wizard

A new step is added in the ECS Install Wizard called Check Prerequisites. This ECS prerequisite checks fresh installations seamlessly and improves the overall installation experience for administrators. This step checks if the ECS hosts meet a list of minimum requirements before installation. For more information on this prerequisite check, see Installing CDP Private Cloud Data Services using ECS.

DRS automatic backups

Starting from CDP Private Cloud Data Services 1.5.4, DRS automatic backups for Control Plane, CDW, and CDE are enabled by default on ECS clusters for new installations or after cluster upgrade to version 1.5.4 or higher. You can disable this option, if required. You can also configure the external storage in Longhorn for ECS, and then initiate DRS automatic backups to it. Automatic backups (DRS) functionality is disabled by default on OCP clusters. For more information, see DRS automatic backups.

Authentication for ingress TLS/SSL

A new property (ssl_private_key_password) is added to the Cloudera Manager to specify the password for the private key in the Ingress Controller TLS/SSL Server Certificate and Private Key file.

Improved Diagnostics

The tez-site.xml file is now included in the Management Console diagnostic bundle download.

New features in Cloudera Data Engineering

Support for password protected private key to initialize the virtual cluster

You can now use the password protected private keys to initialize the virtual cluster. Currently, the password protected private keys are supported with RSA and EC algorithms only. For more information, see Initializing virtual clusters.

Support for Spark Connect session (Preview)

CDE supports Spark Connect sessions which are a type of CDE sessions that expose the Spark Connect interface and allows you to run spark commands from any remote Python environment.

New features in Cloudera Data Warehouse

Hue supports natural language query processing (Preview)

Hue leverages the power of Large Language Models (LLM) to help you generate SQL queries from natural language prompts and also provides options to optimize, explain, and fix queries, promoting efficient and accurate practices for accessing and manipulating data. You can use several AI services and models such as OpenAI's GPT service, Amazon Bedrock, and Azure's OpenAI service to run the Hue SQL AI assistant.

- To learn more about the supported models and services, limitations, and what data is shared with the LLMs, see About the SQL AI Assistant in CDW.
- To set up and enable the SQL AI Assistant, see About setting up the SQL AI Assistant in CDW.
- To see how to generate, edit, explain, optimize, and fix queries, see Starting the SQL AI Assistant in Hue.

Improvements to custom pod configuration (now known as Resource Templates)

Several improvements and changes have been made to the custom pod configuration functionality starting with CDP Private Cloud Data Services 1.5.4. The custom pod configuration has been renamed to "Resource Templates".

- A new menu option Resource Templates has been added to the left navigation pane on the CDW web interface.
- You can now configure the allocation of Kubernetes resources to the pods for Hive, Data Visualization, and Database Catalog in addition to Impala.
- The Impala pod configuration feature is moved from the **Environment Details** page to the **Resource Templates** page.

For more information, see Resource templates for CDW Private Cloud pods.

Flexibility to enable and disable quota management for CDW entities

Earlier, you were required to enable the quota management option before activating an environment to use quota-managed resource pools for environments, Data Catalogs, Virtual Warehouses, and Data Visualization instances in CDW. Starting with CDP Private Cloud Data Services 1.5.4, you can enable or disable the quota management option at any time during the lifecycle of the CDW entities. To learn more about the behavioral aspects, see Quota Management in CDW Private Cloud.

Added support for authentication between CDW and HMS database using mTLS

CDW and the Hive MetaStore (HMS) database on the base cluster can mutually authenticate each other during the SSL/TLS handshake using mTLS for all supported backend databases (Oracle, MySQL, MariaDB, and Postgres). To set up mTLS, you must upload the database client certificate and the private key files in PEM format while activating an environment in CDW. See Enabling mTLS between the HMS database and CDW on Private Cloud.

Ability to enable active-passive configuration for HiveServer2 pods

CDW provides an option to enable active-passive configuration for HiveServer2 (HS2) pods in Private Cloud. By enabling this feature, two HS2 pods run simultaneously–one active and the other inactive. When one pod terminates, the inactive pod becomes active, providing High Availability. See HiveServer2 High Availability in CDW Private Cloud. The most likely cause of a pod's termination is node failure.

CDW no longer has a dependency on YARN

Environment activation in CDW no longer depends on or fails if the YARN service is not installed on the CDP Base cluster.

Improvements to backup and restore CDW

There are two ways to create backups of the Data Warehosue service:

- Using Data Recovery Service (DRS)
- Using the CDW's CDP CLI cluster management commands

By default, CDW backs up namespace-related data before starting the upgrade process using the Data Recovery Service (DRS). A new option called Back up Virtual Warehouse namespaces before an upgrade has been added to disable the automatic backup process on the Advanced Configuration page in the CDW web interface.

New features in Cloudera Machine Learning

CML Service Accounts are available in CML Private Cloud

In CML, the Kerberos principal for the Service Account may not be the same as your login information. Therefore, ensure you provide the Kerberos identity when you sign in to the Service Account. For more information, see Authenticating Hadoop for CML Service Accounts.

Model Registry is available in CDP Private Cloud

Model Registry is now generally available (GA) in CDP Private Cloud. Model Registry in CDP Private Cloud uses Apache Ozone to store model artifacts. For creating a Model Registry you need the Ozone S3 gateway endpoint, the Ozone access key, and the Ozone secret key. If you deploy Model Registry in an environment that contains one or more CML workspaces, you must synchronize the Model Registry with the workspaces. For more information, see Prerequisites for creating Model Registry and Synchronizing the model registry with a workspace.

Heterogeneous GPU usage

When using heterogeneous GPU clusters to run sessions and jobs, the available GPU accelerator labels need to be selected during workload creation. For more information, see Heterogeneous GPU clusters.

Data connections without auto discovery

Cloudera Machine Learning is a flexible, open platform, supporting connections to many data sources. The provided code samples demonstrate how to access local data for CML workloads. For more information, see Connecting to CDW.

Spark Log4j Configuration

Cloudera Machine Learning allows you to update Spark's internal logging configuration on a per-project basis. Spark logging properties can be customized for every session, and job with a default file path found at the root of your project. You can also specify a custom location with a custom environment variable. For more information, see Spark Log4j Configuration.

ML Metrics Collector service

The Metrics Collector service gathers data about how users and groups use resource quota, like how much CPU, Memory and GPU capacity (if any) is allocated, and what the users or groups utilize from that. The Metrics Collector service is running by default, but to collect data about resource quota metrics, you need to enable the Quota Management feature. For more information, see ML Metrics Collector Service overview.

Quota Management for group level (Preview)

Quota Management Technical Preview (TP) release enables you to control how resources are allocated within your CML workspace on user and on group level. Yunikorn Gang Scheduling is also available, which is the default scheduling mechanism in Cloudera Machine Learning. For more information, see Quota Management overview and Yunikorn Gang Scheduling.

Restarting a failed AMP setup

You can now retry failed AMP deployment steps and continue the AMP setup to handle intermittent and configuration issues. For more information, see Restarting a failed AMP setup.

New Hadoop CLI Runtime Addon versions are available

The HadoopCLI 7.1.8.3-601 Runtime Addon is released for the Private Cloud.

Release notes for component services

- Data Catalog

- Management Console
- Cloudera Data Warehouse (CDW)
- Cloudera Machine Learning (CML)
- Cloudera Data Engineering (CDE)
- Cloudera Manager
- Replication Manager

What's new in Cloudera Data Services on premises 1.5.4

New features in the 1.5.4 release of the Cloudera Data Services on premises.

Stability and Resiliency: New prerequisite check in Cloudera Embedded Container Service Install Wizard

A new step is added in the Cloudera Embedded Container Service Install Wizard called Check Prerequisites. This Cloudera Embedded Container Service prerequisite checks fresh installations seamlessly and improves the overall installation experience for administrators. This step checks if the Cloudera Embedded Container Service hosts meet a list of minimum requirements before installation. For more information on this prerequisite check, see Installing CDP Private Cloud Data Services using ECS.

DRS automatic backups

Starting from CDP Private Cloud Data Services 1.5.4, DRS automatic backups for Control Plane, Cloudera Data Warehouse, and Cloudera Data Engineering are enabled by default on Cloudera Embedded Container Service clusters for new installations or after cluster upgrade to version 1.5.4 or higher. You can disable this option, if required. You can also configure the external storage in Longhorn for Cloudera Embedded Container Service, and then initiate DRS automatic backups to it.

Automatic backups (DRS) functionality is disabled by default on OCP clusters.

For more information, see DRS automatic backups.

Authentication for Ingress TLS/SSL

A new property (ssl_private_key_password) is added to the Cloudera Manager to specify the password for the private key in the Ingress Controller TLS/SSL Server Certificate and Private Key file.

Improved Diagnostics

The tez-site.xml file is now included in the Management Console diagnostic bundle download.

Kubernetes Ingress NGINX Controller vulnerabilities

Kubernetes Ingress NGINX Controller vulnerabilities.

Five vulnerabilities affecting the Ingress Nginx Controller for Kubernetes were publicly disclosed on March 24, 2025, and were given the nickname "IngressNightmare."

The ‘IngressNightmare’ vulnerabilities may allow Remote Code Execution (RCE) and potentially expose Kubernetes clusters to malicious configuration modifications. Exploitation requires specially crafted HTTP requests that bypass security measures, such as a Web Application Firewall (WAF). Successful exploitation may lead to complete cluster compromise, data exfiltration, and denial of service.

Details of the CVEs:

- CVE-2025-1974 (CVSS score: 9.8) – An unauthenticated attacker with access to the pod network can achieve arbitrary code execution in the context of the ingress-nginx controller under certain conditions
- CVE-2025-24514 (CVSS score: 8.8) – The auth-url Ingress annotation can be used to inject configuration into NGINX, resulting in arbitrary code execution in the context of the ingress-nginx controller and disclosure of secrets accessible to the controller
- CVE-2025-1097 (CVSS score: 8.8) – The auth-tls-match-cn Ingress annotation can be used to inject configuration into NGINX, resulting in arbitrary code execution in the context of the ingress-nginx controller and disclosure of secrets accessible to the controller
- CVE-2025-1098 (CVSS score: 8.8) – The mirror-target and mirror-host Ingress annotations can be used to inject arbitrary configuration into NGINX, resulting in arbitrary code execution in the context of the ingress-nginx controller and disclosure of secrets accessible to the controller
- CVE-2025-24513 (CVSS score: 4.8) – An improper input validation vulnerability that could result in directory traversal within the container, leading to denial-of-service (DoS) or limited disclosure of secret objects from the cluster when combined with other vulnerabilities



Note: Cloudera will not provide remediation options for unsupported versions, and has not tested mitigations on unsupported versions. Customers are advised to upgrade to a supported product version. For more information, refer to the Support Lifecycle Policy page.

Action required

Multiple mitigation scenarios are possible depending on the version of Cloudera Data Services deployed.

1. Check your Data Services version.

   To determine which version of Data Services is running, log into Cloudera Manager UI, click “Data Services” in the left sidebar, and find each Data Services cluster's version(s).
2. Check your Cloudera Manager version.

   To determine which version of Cloudera Manager is running, log into Cloudera Manager UI and click on the version number at the lower-left corner of the home page. The Cloudera Manager 7.11.3 CHF8 Build number shows 56304673.

Mitigation steps for Cloudera Data Services On Premises:

1. No action is needed in the following scenarios:
   - Running Data Services on OpenShift Container Platform (OCP) and using the default HAproxy ingress controller. If you have customized OCP to use NGINX as the ingress controller, please contact RedHat support for further information
   - Running Data Services 1.5.4 or later with Cloudera Manager 7.11.3 CHF8 or later. These versions already disable the affected admission webhook component.
2. Data Services before 1.5.4, or Cloudera Manager prior to 7.11.3 CHF8, perform one of the following:
   - (Recommended) Upgrade Cloudera Manager to version 7.13.1 CHF1 or higher and all Embedded Container Service clusters to 1.5.4 SP2 or higher. This will not upgrade the nginx version but will automatically disable the webhook via Cloudera Manager, which mitigates the CVEs.
   - If an upgrade is not possible, follow the below steps to disable the webhook to mitigate the CVEs manually.

     

     Note: These steps should be run on each ECS cluster. If you create a new ECS cluster while on an affected release, reapply these steps on the new cluster(s).

1. SSH into the ECS server host.
2. Create the following YAML file to patch the built-in Nginx Ingress Controller Helm Chart:

```
apiVersion: helm.cattle.io/v1
kind: HelmChartConfig
metadata:
```

```
 name: rke2-ingress-nginx
 namespace: kube-system
spec:
 valuesContent: |-
  controller:
   admissionWebhooks:
    enabled: false
```

3. Run the following commands to apply the patch:

```
alias kubectl=/var/lib/rancher/rke2/bin/kubectl
export KUBECONFIG=/etc/rancher/rke2/rke2.yaml
kubectl apply -f <PATCH_FILENAME>.yaml
```



Note: No cluster restart or ECS refresh is required.



Note: Cloudera will not provide remediation options for unsupported versions, and has not tested mitigations on unsupported versions. Customers are advised to upgrade to a supported product version. For more information, refer to the Support Lifecycle Policy page

Cloudera Data Services on premises version 1.5.4 SP2 CHF1 will contain fixes to upgrade the Ingress NGINX controller. Target release date: April 25. 2025.

For the latest update on this issue, see the corresponding Knowledge article: TSB 2025-839: Critical Kubernetes Ingress NGINX Controller Vulnerability Allows RCE Without Authentication

For the latest update on this issue, see the corresponding Data Services Knowledge article: https://my.cloudera.com/knowledge/TSB-2025-839-Mitigation-steps-for-Cloudera-Data-Services-on?id=405006

Known issues for the Cloudera Data Services on premises 1.5.4

This section lists known issues that you might run into while using the Cloudera Data Services on premises.

Known Issues in Cloudera Management Console 1.5.4

OPSX-5147: OOM when retrieving size of Binary File

Sometimes, diagnostics bundle collection fails to complete due to OOM issues.

Limit the time range for the diagnostics bundle.

DOCS-21833: Orphaned replicas/pods are not getting auto cleaned up leading to volume fill-up issues

By default, Longhorn will not automatically delete the orphaned replica directory. One can enable the automatic deletion by setting orphan-auto-deletion to true.

No workaround available.

OPSX-5310: Longhorn engine images were not deployed on Cloudera Embedded Container Service server nodes

Longhorn engine images were not deployed on Cloudera Embedded Container Service server nodes due to missing tolerations for Cloudera Control Plane taints. This caused the engine DaemonSet to schedule only on Cloudera Embedded Container Service agent nodes, preventing deployment on Cloudera Control Plane nodes.

1. Check the Engine `DaemonSet Status`. Run the following command to check if the Longhorn engine DaemonSet is missing on certain nodes:

```
kubectl get ds -n longhorn-system | grep engine
```

2. Identify Taints on Affected Nodes. Run the following command to check for taints on affected nodes:

```
kubectl describe node <node-name> | grep Taints
```



Note: If you see, `node-role.kubernetes.io/control-plane=true:NoSchedule`, this confirms the issue.

3. Manually Edit the DaemonSet to Add a Toleration. Edit the Longhorn engine DaemonSet YAML:

```
kubectl edit ds -n longhorn-system engine-image-ei-<your-engine-id>
```

4. Add the following under tolerations:

```
tolerations:
- effect: NoSchedule
  key: node-role.kubernetes.io/control-plane
  operator: Equal
  value: "true"
```

5. Apply the changes and verify deeployment. Save and exit the editor. Then, check if the DaemonSet is now running on all necessary nodes:

```
kubectl get pods -n longhorn-system -o wide | grep engine
```

Verify that the engine pods are successfully scheduled on the affected ECS server nodes.

OPSX-5148: Diagnostics Collection from UI w/ Default No Time Limit Should Not Invoke Timestamp Filtering

When the diagnostics collection is triggered through the UI, by default, "No Time Limit" is selected. Filtering of logs by timestamp is still observed.

No workaround available.

OPSX-4781: Vault pods may take long time to be ready during upgrades from 1.5.2 to 1.5.3

The 'vault-0' pod takes longer time to attach volume in some upgrade cases than usual. Due to the excess time taken the cluster upgrade may fail. But, usually in 15 minutes the volume can attach automatically and the pod would start running. In that case, the user can resume the upgrade.

No workaround available.

OPSX-5155: OS Upgrade | Pods are not starting after the OS upgrade from RHEL 8.6 to 8.8

After an OS upgrade and start of the Cloudera Embedded Container Service service, pods fail to come up due to stale state.

Restart the Cloudera Embedded Container Service cluster.

OPSX-5055: ECS upgrade failed at Unseal Vault step

During an Cloudera Embedded Container Service upgrade from 1.5.2 to 1.5.4 release, the vault pod fails to start due to an error caused by the Longhorn volume unable to attach to the host. The error is as below:

Warning FailedAttachVolume 3m16s (x166 over 5h26m) attachdetach-controller AttachVolume.Attach failed for volume "pvc-0ba86385-9064-4ef9-9019-71976b4902a5" :

rpc error: code = Internal desc = volume pvc-0ba86385-9064-4ef9-9019-71976b4902a5 failed to attach to node host-1.cloudera.com with attachmentID csi-7659ab0e6655d308d2316536269de47b4e66062539f135bf6012bfc8b41fc345: the volume is currently attached to different node host-2.cloudera.com

Follow below steps provided by SUSE to ensure the Longhorn volume is correctly attached to the node where the vault pod is running.

```
# Find out the volume name that is failing to attach to the vault
 pod.
For e.g. pvc-bc73e7d3-c7e7-468a-b8e0-afdb8033e40b from the pod
logs.
kubectl edit volumeattachments.longhorn.io -n longhorn-system
pvc-bc73e7d3-c7e7-468a-b8e0-afdb8033e40b

# Update the "spec:" section of the volumeattachment and replace
attachmentTickets section with {} as shown below and save.
spec:
 attachmentTickets: {}
 volume: pvc-bc73e7d3-c7e7-468a-b8e0-afdb8033e40b

# scale down the vault statefulset to 0 and scale it back up.
kubectl scale sts vault --replicas=0 -n vault-system
kubectl scale sts vault --replicas=1 -n vault-system
```

OPSX-4308: Display error in UI if listEnvironments failed

On the Environments page, if the `listEnvironments` API call fails, the error is hidden, and instead no environments are displayed, even though they do exist. This can be due to vault issues or connectivity issues.

No workaround available but the register environment page shows the error.

OPSX-4684: Start Cloudera Embedded Container Service command shows green(finished) even though start docker server failed on one of the hosts

The Docker service starts, but one or more Docker roles fail to start because the corresponding host is unhealthy.

Ensure the host is healthy. Start the the Docker role on the host.

OPSX-735: Kerberos service should handle Cloudera Manager downtime

The Cloudera Manager Server in the base cluster operates to generate Kerberos principals for Cloudera on premises If there is downtime, you may observe Kerberos-related errors.

Resolve downtime on Cloudera Manager. If you encounter Kerberos errors, you can retry the operation (such as retrying creation of the Virtual Warehouse).

Known issues from previous releases carried in 1.5.4

OPSX-4754 [Cloudera Embedded Container Service Restart Stability] DaemonSet rollout process is stuck post rolling restart where DaemonSet kube-system/rke2-canal has not finished or progressed for at least 15 minutes

On RHEL 9.x, an Cloudera Embedded Container Service service DaemonSet rollout health alert appears in the Cloudera Manager after an Cloudera Embedded Container Service installation and a rolling restart.

To fix the DaemonSet rollout issue:

1. Edit the DaemonSet rke2-canal configuration file by running the following command:

```
KUBECTL -n kube-system edit ds/rke2-canal
```

   Change the value of felixIptablesBackend from auto to Legacy and save the DaemonSet rke2-canal configuration file.
2. Reboot each node one-by-one.
3. Check to see if any of the nodes are cordoned off. If so, uncordon them:

```
[root@host-1 ~]# $KUBECTL get nodes
                  NAME STATUS ROLES AGE VERSION
                  host-1.ecs-restart1.kcloud.cloudera.com Re
ady,SchedulingDisabled control-plane,etcd,master 17h v1.26.1
0+rke2r1
                  host-2.ecs-restart1.kcloud.cloudera.com R
eady <none> 17h v1.26.10+rke2r1
                  host-3.ecs-restart1.kcloud.cloudera.com Re
ady <none> 17h v1.26.10+rke2r1
                  host-4.ecs-restart1.kcloud.cloudera.com Rea
dy <none> 17h v1.26.10+rke2r1
                  [root@host-1 ~]# $KUBECTL uncordon host-1.ec
s-restart1.kcloud.cloudera.com
                  node/host-1.ecs-restart1.kcloud.cloudera.com
 uncordoned
                  [root@host-1 ~]# $KUBECTL get nodes
                  NAME STATUS ROLES AGE VERSION
                  host-1.ecs-restart1.kcloud.cloudera.com Rea
dy control-plane,etcd,master 17h v1.26.10+rke2r1
                  host-2.ecs-restart1.kcloud.cloudera.com R
eady <none> 17h v1.26.10+rke2r1
                  host-3.ecs-restart1.kcloud.cloudera.com Re
ady <none> 17h v1.26.10+rke2r1
                  host-4.ecs-restart1.kcloud.cloudera.com Rea
dy <none> 17h v1.26.10+rke2r1
                  [root@host-1 ~]#
```

4. Ensure that the Vault is unsealed. , To unseal the vault in Cloudera Manager navigate to Clusters ECS <***ECS SERVICES***> such as ECS-1 or ECS-2 Actions Unseal Vault .
5. Wait for five to six minutes.
6. Check for longhorn pods that fail to come up on any of the hosts:

```
[root@host-1 ~]# kubectl -o wide get pods -n longhorn-system |
 grep -v "Running" | grep -v "Completed"
NAMESPACE                                NAME
                                             READY
STATUS                  RESTARTS         AGE     IP
  NODE
NOMINATED NODE   READINESS GATES
longhorn-system                          longhorn-csi-plugin-
frwnw                                         2/3
CrashLoopBackOff        14 (3m51s ago)   6h20m   10.x.x.x
 host-1.upgr-ecs-ext.kcloud.cloudera.com    <none>
          <none>
longhorn-system                          longhorn-manager-lgzm
b                                             0/1
CrashLoopBackOff        7 (97s ago)      6h24m   10.x.x.x
host-1.upgr-ecs-ext.kcloud.cloudera.com    <none>
          <none>
```

7. Reboot the host (In this case the host is: *host-1.upgr-ecs-ext.kcloud.cloudera.com*).
8. Wait for 15-30 minutes for pods to come up.

9. Post Cloudera Embedded Container Service reboot, if you notice buildkit pods in the following CrashLoopBackOff state, then delete those buildkit pods:

```
[root@host-1 ~]# kubectl -o wide get pods | grep -v "Running"
 | grep -v "Completed"
NAMESPACE                                NAME
                                             READY   STATUS
              RESTARTS         AGE     IP             NODE
                                   NOMINATED NODE   READINESS
 GATES
quasar-skl2-host-1                              buildkit-2jdmw
                                          2/3     CrashLoopBackOff
       14 (3m51s ago)   6h20m   10.x.x.x     host-1.upgr-ecs-
ext.kcloud.cloudera.com    <none>           <none>
quasar-skl2-host-1                             buildkit-k20smc
                                             0/1     CrashLoopBa
ckOff        7 (97s ago)      6h24m   10.x.x.x     host-2.up
gr-ecs-ext.kcloud.cloudera.com    <none>           <none>
```

You can delete the above buildkit pods by one of the following ways:

- On the Cloudera Manager UI, navigate to Clusters ECS <***ECS SERVICES***> such as ECS-1 or ECS-2 Web UI ECS Web UI Delete .
- Run the following command to delete all such buildkit pods:

```
[root@host-1 ~]# kubectl delete pod buildkit-2jdmw -n quasar
-skl2-host-1
```

Wait for the buildkit pods to start back up.

OPSAPS-69892: kube-proxy failure causing issues with cluster

After rebooting/restarting an Cloudera Embedded Container Service agent node, the kube-proxy Linux process may not start due to a race condition in the kubelet. When this happens, Cloudera Embedded Container Service cluster networking and other services – such as Vault, DNS, authentication, Longhorn storage, etc. – are affected. At the Kubernetes pod level, errors such as "connection refused", "connection timed out" and "i/o timeout" may be observed. If you suspect possible networking issues in your Cloudera Embedded Container Service cluster, checking kube-proxy is a good first step.

To fix this issue, perform the following steps on all of the affected nodes:

1. To identify which agent needs to be restarted, check the status of each kube-proxy pod to make sure it is in the "ready" state by running the following command on each host in the cluster.

```
kubectl describe pod [***POD-NAME***] -n kube-system
```

Here, *[***POD-NAME***]* should have a format such as: kube-proxy-<hostname>.

In the Conditions section of the describe pod output, confirm that the "ready" condition is "True".

```
Conditions:
  Type              Status
  Initialized       True
  Ready             True
  ContainersReady   True
```

```
  PodScheduled      True
```

Another option is to run the following command:

```
kubectl get pods -n kube-system -l component=kube-proxy -o go-
template='{{range .items}}
{{.metadata.name}}{{"\n"}}{{"  "}}{{range .status.conditions}}
{{ if eq .type "Ready" }}
Ready:{{.status}}{{"\n\n"}}{{end}}{{end}}{{end}}'
```

The sample output displays the status of all of the kube-proxy pods in the cluster:

```

kube-proxy-host-1.cloudera.com
  Ready:True

kube-proxy-host-2.cloudera.com
  Ready:True

kube-proxy-host-3.cloudera.com
  Ready:True
```

2. If the "ready" state is False, kube-proxy is not functioning properly, regardless of whether the kube-proxy process is running on that host or not. On each of the affected nodes, run the following command to delete the kube-proxy pod manifest:

   ```
   rm /var/lib/rancher/rke2/agent/pod-manifests/kube-proxy.yaml
   ```

3. Start the agent role.

   After the agent role is started, you may not immediately see the kube-proxy process running, but a new kube-proxy process should start shortly. Check the pod status to make sure it is ready. After all of the problem agents have been restarted, the cluster may complain that the vault is sealed – if so, unseal it. At this point, the Cloudera Control Plane should be functioning properly.

Additional details about this issue are available here: https://www.suse.com/support/kb/doc/?id=000021284

OPSX-4766: [Cloudera Embedded Container Service Restart]| Host Reboot | start command failed with error - "Timed out waiting for kube-apiserver to be ready"

In an Cloudera Embedded Container Service cluster with HA enabled, Cloudera Embedded Container Service Start fails with an error after stopping the cluster and rebooting the hosts.

Steps to reproduce:

1. Stop Cloudera Embedded Container Service.
2. Reboot hosts.
3. Start Cloudera Embedded Container Service.

The start command fails with the following error message:

"Timed out waiting for kube-apiserver to be ready"

Option 1:

Start each master role instance individually without waiting each node to be up and running.

Option 2:

If Option 1 does not work, follow the steps from SUSE to recover the cluster:https://docs.rke2.io/backup_restore#cluster-reset

Known Issues in Cloudera Management Console 1.5.2

OPSAPS-68923: After Cloudera Manager upgrade from 7.9.5 to 7.11.3.x Cloudera Embedded Container Service cluster showing stale config

After Cloudera Manager upgrade from 7.9.5 to 7.11.3.x, an Cloudera Embedded Container Service 1.5.0 cluster may show a stale config to add ""limit_fds": 1048576"

This can be ignored – no restart of the Cloudera Embedded Container Service cluster is necessary. When the Cloudera Embedded Container Service 1.5.0 cluster is upgraded to 1.5.2, the stale config will be resolved.

OPSX-4594: [Cloudera Embedded Container Service Restart Stability] Post rolling restart few volumes are in detached state (vault being one of them)

After rolling restart there may be some volumes in detached state.

1. Open the Longhorn UI to view the detached volumes.
2. Perform the following operations for each volume in a detached state:
   a. Identify the workload name and type from the volume details.
   b. Identify the workload and number of replicas using kubectl or the Kubernetes UI.
   c. Scale the workload down to 0.
   d. Wait for the pods associated with the workload to fully terminate.
   e. Scale up the workload up to the number of replicas it had originally.

To prevent this issue, use the Longhorn UI to set the number of replicas for the volume to at least 3.

OPSAPS-68558: [7.9.5->7.11.3.2] Cloudera Manager upgrade failed with BeanCreationException: Error creating bean with name 'com.cloudera.server.cmf.TrialState'

After upgrading the Cloudera Manager package, the Cloudera Manager Server does not start. An error about "Active Commands" is shown in the Cloudera Manager Server log.

This may happen when the Cloudera Data Services on premises is actively issuing requests to Cloudera Manager while an upgrade is being performed.

Before upgrading Cloudera Manager make sure there are no active commands. If there are any active commands, wait for them to complete before starting a Cloudera Manager upgrade.

If Cloudera Manager restart fails after upgrade due to an active getClientConfig command, check the Cloudera Manager server log for a "There are 1 active commands of type GetClientConfigFiles" error. This may block a Cloudera Manager restart after upgrade. Use the following steps to resolve this issue:

1. Login to Cloudera Manager database.
2. Search for any active GetClientConfigFiles command in the COMMANDS table.

```
UPDATE COMMANDS SET active=0,success=false,state='CANCELLED'
 where command_id=<command_id>;
```

3. Delete these entries, including foreign key dependencies, in the following tables:
   - PROCESSES
   - PROCESSES_DETAIL
   - COMMANDS_DETAIL

```
cm=> DELETE FROM COMMANDS where command_id=1546340765;
ERROR:  update or delete on table "commands" violates foreign
 key constraint "fk_process_command" on table "processes"
DETAIL:  Key (command_id)=(1546340765) is still referenced fro
m table "processes".
cm=>
cm=> DELETE FROM processes where command_id=1546340765;
```

```
ERROR:  update or delete on table "processes" violates foreign
 key constraint "fk_processes_detail_process" on table "proc
esses_detail"
DETAIL:  Key (process_id)=(1546340766) is still referenced fro
m table "processes_detail".
cm=>
cm=>                                          ^
cm=> DELETE FROM processes_detail where process_id=1546340766;
DELETE 1
cm=> DELETE FROM processes where command_id=1546340765;
DELETE 1
cm=> DELETE FROM COMMANDS where command_id=1546340765;
ERROR:  update or delete on table "commands" violates foreign
 key constraint "fk_commands_detail_command" on table "comma
nds_detail"
DETAIL:  Key (command_id)=(1546340765) is still referenced f
rom table "commands_detail".
cm=>
cm=> DELETE FROM commands_detail where command_id=1546340765;
DELETE 1
cm=> DELETE FROM COMMANDS where command_id=1546340765;
DELETE 1
```

4. Restart the Cloudera Manager server.

OPSX-4392: Getting the real client IP address in the application

Cloudera AI has a feature for adding the audit event for each user action (Monitoring User Events). In Cloudera on premises, instead of the client IP, we are getting the internal IP, which is logged into the internal DB.

In Cloudera Embedded Container Service, add the enable-real-ip configuration as true for the nginx ingress controller:

```
apiVersion: v1
data:
  allow-snippet-annotations: "true"
  enable-real-ip: "true"                    <<<<<<<<<< new config
kind: ConfigMap
metadata:
  annotations:
    meta.helm.sh/release-name: rke2-ingress-nginx
    meta.helm.sh/release-namespace: kube-system
  creationTimestamp: "2023-05-09T04:54:53Z"
  labels:
    app.kubernetes.io/component: controller
    app.kubernetes.io/instance: rke2-ingress-nginx
    app.kubernetes.io/managed-by: Helm
    app.kubernetes.io/name: rke2-ingress-nginx
    app.kubernetes.io/part-of: rke2-ingress-nginx
    app.kubernetes.io/version: 1.6.4
    helm.sh/chart: rke2-ingress-nginx-4.5.201
  name: rke2-ingress-nginx-controller
  namespace: kube-system
  resourceVersion: "162559439"
  uid: cca67b0c-bc05-4e1f-8439-7d44323f4624
```

In OCP, you may be able configure this using HAproxy with X-forward-for pass to OpenShift 4.

OPSX-4552: [Cloudera Embedded Container Service Restart] One of the docker servers failed to come up after starting the cluster post hosts reboot

At times the Docker server may fail to come up and return the following error message:

/var/run/docker.sock: Is a directory

On the Docker server role host, remove the /var/run/docker.sock directory, then restart the Docker server role.

CDPVC-1137, CDPAM-4388, COMPX-15083, and COMPX-15418: OpenShift Container Platform version upgrade from 4.10 to 4.11 fails due to a Pod Disruption Budget (PDB) issue

PDB can prevent a node from draining which makes the nodes to report the "Ready,SchedulingDisabled" state. As a result, the node is not updated to correct the Kubernetes version when you upgrade OCP from 4.10 to 4.11.

To resolve this issue, confirm that the upgrade has failed due to the PDB issue, and then manually delete the PDBs from the Cloudera on premises namespace.

1. Run the following command to check whether the nodes are stuck in the "Ready,SchedulingDisabled" state:

```
oc get nodes
```

2. Get the machine config daemon details of the particular pod as follows:

```
oc get po -n openshift-machine-config-operator -l 'k8s-app=m
achine-config-daemon' -o wide
```

3. Check the logs of the machine config operator of that particular node as follows:

```
oc logs -f -n openshift-machine-config-operator [***MACHINE-
CONFIG-DAEMON-NAME***] -c machine-config-daemon
```

Replace *[***MACHINE-CONFIG-DAEMON-NAME***]* with the actual machine config daemon name.

You may see one of the following errors in the node logs:

- error when evicting pods/cdp-release-cpx-liftie-****" -n "*[***PRIVATE-CLOUD-NAMESPACE***]* Cannot evict pod as it would violate the pod's disruption budget
- error when evicting pods/"cdp-release-cluster-proxy-[******]" -n "*[***PRIVATE-CLOUD-NAMESPACE***]* Cannot evict pod as it would violate the pod's disruption budget

Delete the PDB from the Cloudera on premises namespace as follows:

a. Obtain the PDB for the cdp-release-cluster-proxy namespace:

```
oc get pdb -n [***PRIVATE-CLOUD-NAMESPACE***] | grep cdp-
release-cluster-proxy
```

b. Back up the PDB:

```
oc get pdb [***PDB-NAME-OF-CLUSTER-PROXY***] -n [***PRIVATE-
CLOUD-NAMESPACE***] -o yaml >> [***BACKUP-FILE-NAME***].yaml
```

c. Delete the PDB:

```
oc delete pdb [***PDB-NAME-OF-CLUSTER-PROXY***] -
n [***PRIVATE-CLOUD-NAMESPACE***]
```

Repeat the steps to delete the cdp-release-cpx-liftie PDB as well.

PULSE-944 and PULSE-941 Cloudera Observability namespace not created after platform upgrade from 151 to 152

The Cloudera Observability namespace is not created after a platform upgrade from Cloudera Observability 1.5.1 to Cloudera Data Services on premises 1.5.2.

During the creation of the resource pool the Cloudera Observability namespace is provided by the Cloudera on premises. If the provisioning flow is not completed, such as due to a timing difference

between the start of the computeAPI pod and the call to the computeAPI pod by the service, the namespace is not created.

Trigger the Cloudera Observability namespace deployment by restarting the pvcservice pod.

PULSE-921 Cloudera Observability namespace has no pods

The Cloudera Observability namespace should have the same number of pods and nodes. When the Cloudera Observability namespace has no pods the prometheus-node-exporter-1.6.0 helm release state becomes invalid and the Cloudera Data Services on premises is unable to uninstall and reinstall the namespace. Also, as the Node Exporter is not installed into the Cloudera Observability namespace its metrics are unavailable when querying Prometheus in the control plane, for example the `node_cpu_seconds_total` metric.

Manually uninstall the invalid helm release with the `--debug` flag, verify that there are no helm releases listed by running `-n observability -a`, and then trigger the deployment process by restarting the pvcservice pod in the control plane.

PULSE-697 Add node-exporter to Cloudera Data Services on premises

When expanding a cluster with new nodes and there is insufficient CPU and memory resources, the Node Exporter will encounter difficulties deploying new pods on the additional nodes.

To ensure sufficient resource allocation, such as when the Cloudera Observability namespace requires adjustment, delete the existing namespace and restart the pvcservice pod. This automatically initiates the creation of the Cloudera Observability namespace with the appropriate resource allocation.



Note: During the namespace recreation process the Node Exporter metrics are temporarily unavailable.

PULSE-935 Longhorn volumes are over 90% of the capacity alerts on Prometheus volumes

Cloudera Manager displays the following alert about your Prometheus volumes: Concerning: Firing alerts for Longhorn: The actual used space of Longhorn volume is over 90% of the capacity.

Longhorn stores historical data as snapshots that are calculated with the active data for the volume's actual size. This size is therefore greater than the volume's nominal data value.

When the alert is displayed on the Cloudera Manager UI and it is related to Longhorn volumes used by Prometheus, ignore. For more information, see the Longhorn space consumption guidelines in the Longhorn documentation.

PULSE-937 Private-Key field change in Update Remote Write request does not reflect in enabling the metric flow

When using the Cloudera Management Console UI for Remote Storage the Disable option does not deactivate the remote write configuration, even when the action returns a positive result message. Therefore, when disabling a remote storage configuration use the CLI client to disable the remote storage configuration directly from the API.

At this time when a remote storage configuration is incorrect, do not use the Edit or Disable option from the configuration’s Actions menu (ellipsis icon) to change its configuration. Instead, delete the remote storage's configuration from the configuration’s Actions menu with the Remove Configuration action and then re-create the remote write configuration with the Delete and Create operations of the API, using the CLI client.

PULSE-841 Disabling the remote write configuration logs an error in both cp prometheus and env prometheus

When a metric replication is set up between the cluster and Cloudera Observability and the connection is disabled or deleted, Prometheus writes an error message that states that it cannot replicate the metrics.

No workaround is required. After a few minutes the errors are no longer logged and Prometheus no longer tries to replicate the metrics.

PULSE-895 Disabling the remote write config in the UI is broken in Cloudera Data Services on premises

The Remote Write Enable and Disable options in the Management Console's User Interface do not work when a Remote Storage configuration is created with a requestSignerAuth type from either the HTTP API or using the CDP-CLI tool.

At this time, do not use the Enable or Disable options from the Remote Storage configuration's Actions menu in the Cloudera Management Console UI. Instead, enable or disable the configuration from the HTTP API or using the CDP-CLI tool.

PULSE-936 No Alert to prompt the metric flow being affected b/c of wrong private key configuration

A remote write alert was not triggered when the wrong private key was used in a Remote Storage configuration.

No workaround. Incorrect configuration settings, such as in this case where a bad private key was used, may block the forwarding of metrics. When creating a Remote Storage configuration you must carefully verify each configuration setting.

Known Issues in Cloudera Management Console 1.5.1

External metadata databases are no longer supported on OCP

As of Cloudera Data Services on premises 1.5.1, external Cloudera Control Plane metadata databases are no longer supported. New installs require the use of an embedded Cloudera Control Plane database. Upgrades from Cloudera Data Services on premises 1.4.1 or 1.5.0 to 1.5.1 are supported, but there is currently no migration path from a previous external Cloudera Control Plane database to the embedded Cloudera Control Plane database. Upgrades from 1.4.0 or 1.5.0 with external Cloudera Control Plane metadata databases also require additional steps, which are described in the Cloudera Data Services on premises 1.5.1 upgrade topics.

DOCS-15855: Networking API is deprecated after upgrade to Cloudera Data Services on premises 1.5.1 (K8s 1.24)

When the control plane is upgraded from 1.4.1 to 1.5.1, the Kubernates version changes to 1.24. The Livy pods running in existing Virtual Clusters (VCs) use a deprecated networking API for creating ingress for Spark driver pods. Because the old networking API is deprecated and does not exist in Kubernates 1.24, any new job run will not work for the existing VCs.

OPSX-4266: Cloudera Embedded Container Service upgrade from 1.5.0 to 1.5.1 is failing in Cadence schema update job

When upgrading from Cloudera Embedded Container Service 1.5.0 to 1.5.1, the CONTROL_PLANE_CANARY fails with the following error:

```
Firing alerts for Control Plane: Job did not complete in time, J
ob failed to complete.
```

And the cdp-release-cdp-cadence-schema-update job fails.

Use the following steps to manually execute the job:

1. Export the job manifest into a file:

```
kubectl get job cdp-release-cdp-cadence-schema-update -n <cdp>
 -o yaml > job.yaml
```

2. Delete the cdp-release-cdp-cadence-schema-update job:

```
kubectl delete job cdp-release-cdp-cadence-schema-update -n
<cdp>
```

3. Remove runtime information from the manifest, such as:

```
resourceVersion
```

```
uid
selector
  matchLabels
    controller-uid
labels
  controller-uid
status section
```

4. Create the job:

```
kubectl apply -f job.yaml
```

OPSX-4076:

When you delete an environment after the backup event, the restore operation for the backup does not bring up the environment.

Create the environment manually.

OPSX-4024: Cloudera Manager truststore import into unified truststore should handle duplicate CommonNames

If multiple CA certificates with the exact same value for the Common Name field are present in the Cloudera Manager truststore when a Cloudera Data Services on premises cluster is installed, only one of them may be imported into the Data Services truststore. This may cause certificate errors if an incorrect/old certificate is imported.

Remove old certificates from the Cloudera Manager truststore, and ensure certificates have unique Common Names.

COMOPS-2822: OCP error x509: certificate signed by unknown authority

The error x509: certificate signed by unknown authority usually means that the Docker daemon that is used by Kubernetes on the managed cluster does not trust the self-signed certificate.

Usually the fix is to copy the certificate to the path below on all of the worker nodes in the cluster:

```
/etc/docker/certs.d/<your_registry_host_name>:<your_registry_hos
t_port>/ca.crt
```

OPSX-3073 First run command failed at setup storage step with error "Timed out waiting for local path storage to come up"

Pod stuck in pending state on host for a long time. Error in Role log related to CNI plugin:

Events:

```
   Type     Reason                  Age                    From
  Message
  ----     ------                  ----                   ----
 -------
  Warning  FailedCreatePodSandBox  3m5s (x269 over 61m)  kubelet
 (combined from similar events):
Failed to create pod sandbox: rpc error: code = Unknown desc = f
ailed to setup network for sandbox
"70427e9b26fb014750dfe4441fdfae96cb4d73e3256ff5673217602d503e806
f":
failed to find plugin "calico" in path [/opt/cni/bin]
```

Delete the cni directory on the host failing to launch pods:

ssh root@ecs-ha1-p-7.vpc.cloudera.com rm -rf /var/lib/cni

Restart the canal pod running on that host:

```
kubectl get pods -n kube-system -o wide | grep  ecs-ha1-p-7.vpc.
cloudera.com
kube-proxy-ecs-ha1-p-7.vpc.cloudera.com                1/1
 Running     0          11h   10.65.52.51    ecs-ha1-p-7.vpc.clo
udera.com   <none>           <none>
rke2-canal-llkc9                                       2/2
Running     0          11h   10.65.52.51    ecs-ha1-p-7.vpc.clou
dera.com   <none>           <none>
rke2-ingress-nginx-controller-dqtz8                    1/1     R
unning     0          11h   10.65.52.51    ecs-ha1-p-7.vpc.cloud
era.com   <none>           <none>
kubectl delete pod rke2-canal-llkc9 -n kube-system
```

OPSX-3528: [Pulse] Prometheus config reload fails if multiple remote storage configurations exist with the same name

It is possible to create multiple remote storage configurations with the same name. However, if such a situation occurs, the metrics will not flow to the remote storage as the config reload of the original prometheus will fail.

At any point in time, there should never be multiple remote storage configurations existing that have the same name.

OPSX-1405: Able to create multiple Cloudera Data Services on premises Environments with the same name

If two users try to create an environment with the same name at the same time, it might result in an unusable environment.

Delete the environment and try again with only one user trying to create the environment.

OPSX-1412: Creating a new environment through the CDP CLI reports intermittently that "Environment name is not unique" even though it is unique

When multiple users try to create the same environment at the same time or use automation to create an environment with retries, create environment may fail on collision with a previous request to create an environment.

Delete the existing environment, wait 5 minutes, and try again.

Known Issues in Cloudera Management Console 1.5.0

Somehow the Rebuilding field inside volume.meta is set to true causing the volume to get stuck in attaching/detaching loop

This is a condition that can occur in ECS Longhorn storage.

Since the volume has only 1 replica in this case, we can:

1. Scale down the workload. The Longhorn volume will be detached.

2. Wait for the Longhorn volume to be detached.

3. SSH into the node that has the replica.

4. cd into the replica folder (for example, /longhorn/replicas/pvc-126d40e2-7bff-4679-a310-e444e84df267-1a5dc941).

5. Change the"Rebuilding" field from true to false in the volume.meta file.

6. Scale up the workload to attach the volume.

Known Issues in Cloudera Management Console identified before 1.5.0

OPSX-5629: COE Insight from case 922848: Not able to connect to bit bucket

After installing Cloudera AI on an Cloudera Embedded Container Service cluster, users were not able to connect the internal bitbucket repo.

Workaround:

In this case the MTU of the Cloudera Embedded Container Service virtual network interfaces were larger than that of host external interface, which may cause the network requests from Cloudera Embedded Container Service containers to get truncated.

The Container Network Interface (CNI) is a framework for dynamically configuring networking resources. CNI integrates smoothly with Kubenetes to enable the use of an overlay or underlay network to automatically configure the network between pods. Cloudera Embedded Container Service uses Calico as the CNI network provider.

The MTU of the pods' virtual network interface can be seen by running the ifconfig command.

The default MTU of the virtual network interfaces is 1450.

The MTU setting of the virtual interfaces is stored as a configmap in the kube-system namespace. To modify the MTU, edit the rke2-canal-config configmap.

```
$ /var/lib/rancher/rke2/bin/kubectl --kubeconfig
 /etc/rancher/rke2/rke2.yaml --namespace kube-system
 edit cm rke2-canal-config
```

Find the veth_mtu parameter in the YAML content. Modify the default value of 1450 to the required MTU size.

Next, restart the rke2-canal pods from the kube-system namespace. There will be rke2-canal pods for each ECS node.

After the pods are restarted, all subsequent new pods will use the new MTU setting. However, existing pods that are already running will remain on the old MTU setting. Restart all of the pods to apply the new MTU setting.

OPSX-2484: FileAlreadyExistsException during timestamp filtering

The timestamp filtering may result in FileAlreadyExistsException when there is a file with same name already existing in the tmp directory.

None

OPSX-2772: For Account Administrator user, update roles functionality should be disabled

An Account Administrator user holds the biggest set of privileges, and is not allowed to modify via current UI, even user try to modify permissions system doesn't support changing for account administrator.

Cloudera Data Services on premises ECS Installation: Failed to perform First Run of services.

If an issue is encountered during the Install Cloudera Control Plane step of Containerized Cluster First Run, installation will be re-attempted infinitely rather than the command failing.

Since the control plane is installed and uninstalled in a continuous cycle, it is often possible to address the cause of the failure while the command is still running, at which point the next attempted installation should succeed. If this is not successful, abort the First Run command, delete the Containerized Cluster, address the cause of the failure, and retry from the beginning of the Add Cluster wizard. Any nodes that are re-used must be cleaned before re-attempting installation.

Environment creation through the CDP CLI fails when the base cluster includes Ozone

Problem: Attempt to create an environment using the CDP command-line interface fails in a Cloudera Data Services on premises deployment when the Cloudera Base on premises cluster is in a degraded state and includes Ozone service.

Workaround: Stopping the Ozone service temporarily in the Cloudera Base on premises cluster during environment creation prevents the control plane from using Ozone as a logging destination, and avoids this issue.

Filtering the diagnostic data by time range might result in a FileAlreadyExistsException

Problem:Filtering the collected diagnostic data might result in a FileAlreadyExistsException if the /tmp directory already contains a file by that name.

There is currently no workaround for this issue.

Kerberos service does not always handle Cloudera Manager downtime

Problem: The Cloudera Manager Server in the base cluster must be running to generate Kerberos principals for Cloudera on premises. If there is downtime, you might observe Kerberos-related errors.

Resolve downtime issues on Cloudera Manager. If you encounter Kerberos errors, you can retry the concerned operation such as creating Virtual Warehouses.

Updating user roles for the admin user does not update privileges

In the Cloudera Management Console, changing roles on the User Management page does not change privileges of the admin user.

None

Upgrade applies values that cannot be patched

If the size of a persistent volume claim in a Containerized Cluster is manually modified, subsequent upgrades of the cluster will fail.

None

Fixed Issues for the Cloudera Data Services on premises 1.5.4

This section lists the issues that have been fixed since the last release of the Cloudera Data Services on premises.

Fixed Issues in Cloudera Management Console 1.5.4

TSB 2024-746: Concurrent compactions from Spark and modify statements from Hive and Impala can corrupt Iceberg tables.

This issue has been fixed.

TSB 2024-745: Impala returns incorrect results for Iceberg V2 tables when optimized operator is being used in Cloudera Data Warehouse.

This issue has been fixed.

TSB 2024-758: Truncate command on Iceberg V2 branches cause unintentional data deletion.

This issue has been fixed.

OPSAPS-69250: Cloudera Embedded Container Service restart failure in an airgap environment due to "`yum install`' step in `rke.sh` script

In `rke.sh` script, "`yum install`" must not be in the Cloudera Embedded Container Service Server startup script, or any exceptions caught during the script running should be handled/propagated in the UI.

Add the following to /etc/yum.repos.d/cloudera-manager.repo

```
proxy=http://proxy-server-IP-address>:<proxy_port>
proxy_username=<proxy-user-name>
proxy_password=<proxy-password>
```

OPSX-4446: Duplicate entries in cdp-pvc-truststore

Duplicate certificates are no longer available in the unified truststore.

OPSX-4650: Cloudera Manager - OCP pvc install Wizard - fails if route name is too long

The kubernetes namespace field is limited to 30 characters. This does not affect existing installations.

OPSX-3666: mlx_crud_app DB connection fails with error "unable to create connection: x509: certificate relies on legacy Common Name field, use SANs instead"

If you are upgrading from Cloudera Data Services on premises 1.4.1 or 1.5.0 to 1.5.1 or higher versions, and you were previously using an external database, you must regenerate the DB certificate with SAN before upgrading to Cloudera Data Services on premises 1.5.1 or higher versions.

OPSX-4225: Upgrade failed as cadence pods are crashlooping post upgrade

When doing a fresh install of Cloudera Data Services on premises 1.5.1, external metadata databases are no longer supported. Instead, the Cloudera Data Services on premises installer will create an embedded database pod by default, which runs inside the Kubernetes cluster to host the databases required for installation.

If you are upgrading from Cloudera Data Services on premises 1.4.1 or 1.5.0 to 1.5.1 or higher versions, and you were previously using an external database, you must run the following psql commands to create the required databases. You should also ensure that the two new databases are owned by the common database users known by the control plane.

```
CREATE DATABASE db-cadence;
CREATE DATABASE db-cadence-visibility;
```

OPSX-4650 : OCP upgrade – OCP namespace name must be 29 characters or less

The kubernetes namespace field is limited to 30 characters in OCP. This does not affect existing installations.

COMPX-15475: Post upgrade prometheus-node-exporter-1.6.0 pod stuck in pending state

Applications, and their pods, that were running before an upgrade are no longer rejected. They get moved to a temporary queue during initialisation if they cannot be placed in the requested queue. This prevents a secondary issue, node rejections, from occurring which caused the pending pods.

OPSAPS-66166: FreeIPA cmadminrole needs more privileges for Cloudera Data Services on premises after upgrade

After upgrade, the Cloudera Manager admin role may be missing the Host Administrators privilege in an upgraded cluster.

The cluster administrator should run the following command to manually add this privilege to the role.

```
ipa role-add-privilege <cmadminrole> --privileges="Host Administ
rators"
```

For more information, see Upgrade from 1.5.2 or 1.5.3 to 1.5.4 (Cloudera Embedded Container Service).

Repository Locations for 1.5.4

The URLs for Cloudera Data Services on premises 1.5.4 are listed in the following table:

URL Type	Repository Location
Index	`https://username:password@archive.cloudera.com/p/cdp-pvc-ds/1.5.4/`
Manifest	Repository: `https://username:password@archive.cloudera.com/p/cdp-pvc-ds/1.5.4/manifest.json`
Parcels	Repository: `https://username:password@archive.cloudera.com/p/cdp-pvc-ds/1.5.4/parcels/`

List of fixed Common Vulnerabilities and Exposures in 1.5.4

Review the Common vulnerabilities and Exposures (CVEs) that were fixed in this release of Cloudera Data Services on premises.

- CVE-2023-27539: A denial of service vulnerability was found in rubygem-rack in how it parses headers. A carefully crafted input can cause header parsing to take an unexpected amount of time, possibly resulting in a denial of service.
- CVE-2024-7264: libcurl's ASN1 parser code has the `GTime2str()` function, used for parsing an ASN.1 Generalized Time field. If given an syntactically incorrect field, the parser might end up using -1 for the length of the *time fraction*, leading to a `strlen()` getting performed on a pointer to a heap buffer area that is not (purposely) null terminated. This flaw most likely leads to a crash, but can also lead to heap contents getting returned to the application when [CURLINFO_CERTINFO](https://curl.se/libcurl/c/CURLINFO_CERTINFO.html) is used.

  Note: The 'curl' binary has CVEs. However, this affected binary is not been used by any of the Cloudera Private Cloud Data Services applications and hence, cannot be exploited. This will be resolved in future release.
- CVE-2024-8096: When curl is told to use the Certificate Status Request TLS extension, often referred to as OCSP stapling, to verify that the server certificate is valid, it might fail to detect some OCSP problems and instead wrongly consider the response as fine. If the returned status reports another error than 'revoked' (like for example 'unauthorized') it is not treated as a bad certficate.

  Note: The 'curl' binary has CVEs. However, this affected binary is not been used by any of the Cloudera Private Cloud Data Services applications and hence, cannot be exploited. This will be resolved in future release.
- DSA-5692-1: ghostscript - security update
- CVE-2024-33871: An issue was discovered in Artifex Ghostscript before 10.03.1. contrib/opvp/gdevopvp.c allows arbitrary code execution via a custom Driver library, exploitable via a crafted PostScript document. This occurs because the Driver parameter for opvp (and oprp) devices can have an arbitrary name for a dynamic library; this library is then loaded.

- CVE-2024-33870: An issue was discovered in Artifex Ghostscript before 10.03.1. There is path traversal (via a crafted PostScript document) to arbitrary files if the current directory is in the permitted paths. For example, there can be a transformation of ../../foo to ./../../foo and this will grant access if ./ is permitted.
- CVE-2024-33869: An issue was discovered in Artifex Ghostscript before 10.03.1. Path traversal and command execution can occur (via a crafted PostScript document) because of path reduction in base/gpmisc.c. For example, restrictions on use of %pipe% can be bypassed via the aa/../%pipe%command# output filename.
- CVE-2024-29510: Artifex Ghostscript before 10.03.1 allows memory corruption, and SAFER sandbox bypass, via format string injection with a uniprint device.
- DSA-5679-1: less - security update
- DSA-5682-2: glib2.0 - regression update
- DSA-5682-1: glib2.0 - security update
- CVE-2024-23653: BuildKit is a toolkit for converting source code to build artifacts in an efficient, expressive and repeatable manner. In addition to running containers as build steps, BuildKit also provides APIs for running interactive containers based on built images. It was possible to use these APIs to ask BuildKit to run a container with elevated privileges. Normally, running such containers is only allowed if special `security.insecure` entitlement is enabled both by buildkitd configuration and allowed by the user initializing the build request. The issue has been fixed in v0.12.5. Avoid using BuildKit frontends from untrusted sources.
- CVE-2024-23652: BuildKit is a toolkit for converting source code to build artifacts in an efficient, expressive and repeatable manner. A malicious BuildKit frontend or Dockerfile using `RUN --mount` could trick the feature that removes empty files created for the mountpoints into removing a file outside the container, from the host system. The issue has been fixed in v0.12.5. Workarounds include avoiding using BuildKit frontends from an untrusted source or building an untrusted Dockerfile containing `RUN --mount` feature.
- CVE-2023-36665: protobuf.js (aka protobufjs) 6.10.0 through 7.x before 7.2.5 allows Prototype Pollution, a different vulnerability than CVE-2022-25878. A user-controlled protobuf message can be used by an attacker to pollute the prototype of Object.prototype by adding and overwriting its data and functions. Exploitation can involve: (1) using the function parse to parse protobuf messages on the fly, (2) loading .proto files by using load/loadSync functions, or (3) providing untrusted input to the functions ReflectionObject.setParsedOption and util.setProperty.
- CVE-2024-22682: DuckDB <=0.9.2 and DuckDB extension-template <=0.9.2 are vulnerable to malicious extension injection via the custom extension feature.
- CVE-2022-30123: A sequence injection vulnerability exists in Rack <2.0.9.1, <2.1.4.1 and <2.2.3.1 which could allow is a possible shell escape in the Lint and CommonLogger components of Rack.
- CVE-2023-38545: This flaw makes curl overflow a heap based buffer in the SOCKS5 proxy handshake. When curl is asked to pass along the hostname to the SOCKS5 proxy to allow that to resolve the address instead of it getting done by curl itself, the maximum length that hostname can be is 255 bytes. If the hostname is detected to be longer than 255 bytes, curl switches to local name resolving and instead passes on the resolved address only to the proxy. Due to a bug, the local variable that means 'let the host resolve the name' could get the wrong value during a slow SOCKS5 handshake, and contrary to the intention, copy the too long hostname to the target buffer instead of copying just the resolved address there.
- CVE-2023-32002: The use of `Module._load()` can bypass the policy mechanism and require modules outside of the policy.json definition for a given module. This vulnerability affects all users using the experimental policy mechanism in all active release lines: 16.x, 18.x and, 20.x. Please note that at the time this CVE was issued, the policy is an experimental feature of Node.js.
- CVE-2016-5397: The Apache Thrift Go client library exposed the potential during code generation for command injection due to using an external formatting tool. Affected Apache Thrift 0.9.3 and older, Fixed in Apache Thrift 0.10.0.
- CVE-2022-3294: Users may have access to secure endpoints in the control plane network. Kubernetes clusters are only affected if an untrusted user can modify Node objects and send proxy requests to them. Kubernetes supports node proxying, which allows clients of kube-apiserver to access endpoints of a Kubelet to establish connections to Pods, retrieve container logs, and more. While Kubernetes already validates the proxying address for Nodes, a bug in kube-apiserver made it possible to bypass this validation. Bypassing this validation could allow authenticated requests destined for Nodes to to the API server's private network.
- CVE-2023-46402: git-urls 1.0.0 allows ReDOS (Regular Expression Denial of Service) in urls.go.

- RHSA-2024:2098: The container-tools module contains tools for working with containers, notably podman, buildah, skopeo, and runc.
- RHSA-2024:0752: The container-tools module contains tools for working with containers, notably podman, buildah, skopeo, and runc.
- CVE-2024-23651: BuildKit is a toolkit for converting source code to build artifacts in an efficient, expressive and repeatable manner. Two malicious build steps running in parallel sharing the same cache mounts with subpaths could cause a race condition that can lead to files from the host system being accessible to the build container. The issue has been fixed in v0.12.5. Workarounds include, avoiding using BuildKit frontend from an untrusted source or building an untrusted Dockerfile containing cache mounts with --mount=type=cache,source=... options.
- RHSA-2023:4419: OpenSSH is an SSH protocol implementation supported by a number of Linux, UNIX, and similar operating systems. It includes the core files necessary for both the OpenSSH client and server.
- RHSA-2024:2699: Git Large File Storage (LFS) replaces large files such as audio samples, videos, datasets, and graphics with text pointers inside Git, while storing the file contents on a remote server.
- RHSA-2024:1444: Node.js is a software development platform for building fast and scalable network applications in the JavaScript programming language.
- RHSA-2023:5360: Node.js is a software development platform for building fast and scalable network applications in the JavaScript programming language.
- RHSA-2023:5850: Node.js is a software development platform for building fast and scalable network applications in the JavaScript programming language.
- CVE-2023-43646: get-func-name is a module to retrieve a function's name securely and consistently both in NodeJS and the browser. Versions prior to 2.0.1 are subject to a regular expression denial of service (redos) vulnerability which may lead to a denial of service when parsing malicious input. This vulnerability can be exploited when there is an imbalance in parentheses, which results in excessive backtracking and subsequently increases the CPU load and processing time significantly. This vulnerability can be triggered using the following input: '\t'.repeat(54773) + '\t/function/i'. This issue has been addressed in commit `f934b228b` which has been included in releases from 2.0.1. Users are advised to upgrade. There are no known workarounds for this vulnerability.
- CVE-2023-45133: Babel is a compiler for writingJavaScript. In `@babel/traverse` prior to versions 7.23.2 and 8.0.0-alpha.4 and all versions of `babel-traverse`, using Babel to compile code that was specifically crafted by an attacker can lead to arbitrary code execution during compilation, when using plugins that rely on the `path.evaluate()`or `path.evaluateTruthy()` internal Babel methods. Known affected plugins are `@babel/plugin-transform-runtime`; `@babel/preset-env` when using its `useBuiltIns` option; and any "polyfill provider" plugin that depends on `@babel/helper-define-polyfill-provider`, such as `babel-plugin-polyfill-corejs3`, `babel-plugin-polyfill-corejs2`, `babel-plugin-polyfill-es-shims`, `babel-plugin-polyfill-regenerator`. No other plugins under the `@babel/` namespace are impacted, but third-party plugins might be. Users that only compile trusted code are not impacted. The vulnerability has been fixed in `@babel/traverse@7.23.2` and `@babel/traverse@8.0.0-alpha.4`. Those who cannot upgrade `@babel/traverse` and are using one of the affected packages mentioned above should upgrade them to their latest version to avoid triggering the vulnerable code path in affected `@babel/traverse` versions: `@babel/plugin-transform-runtime` v7.23.2, `@babel/preset-env` v7.23.2, `@babel/helper-define-polyfill-provider` v0.4.3, `babel-plugin-polyfill-corejs2` v0.4.6, `babel-plugin-polyfill-corejs3` v0.8.5, `babel-plugin-polyfill-es-shims` v0.10.0, `babel-plugin-polyfill-regenerator` v0.5.3.
- CVE-2024-27983: An attacker can make the Node.js HTTP/2 server completely unavailable by sending a small amount of HTTP/2 frames packets with a few HTTP/2 frames inside. It is possible to leave some data in nghttp2 memory after reset when headers with HTTP/2 CONTINUATION frame are sent to the server and then a TCP connection is abruptly closed by the client triggering the Http2Session destructor while header frames are still being processed (and stored in memory) causing a race condition.
- CVE-2021-33910: basic/unit-name.c in systemd prior to 246.15, 247.8, 248.5, and 249.1 has a Memory Allocation with an Excessive Size Value (involving strdupa and alloca for a pathname controlled by a local attacker) that results in an operating system crash.
- CVE-2023-43665: In Django 3.2 before 3.2.22, 4.1 before 4.1.12, and 4.2 before 4.2.6, the django.utils.text.Truncator chars() and words() methods (when used with html=True) are subject to a potential DoS (denial of service) attack via certain inputs with very long, potentially malformed HTML text. The chars() and words() methods are used to implement the truncatechars_html and truncatewords_html template filters, which are thus also vulnerable. NOTE: this issue exists because of an incomplete fix for CVE-2019-14232.

- CVE-2023-46695: An issue was discovered in Django 3.2 before 3.2.23, 4.1 before 4.1.13, and 4.2 before 4.2.7. The NFKC normalization is slow on Windows. As a consequence, django.contrib.auth.forms.UsernameField is subject to a potential DoS (denial of service) attack via certain inputs with a very large number of Unicode characters.
- CVE-2023-41164: In Django 3.2 before 3.2.21, 4.1 before 4.1.11, and 4.2 before 4.2.5, django.utils.encoding.uri_to_iri() is subject to a potential DoS (denial of service) attack via certain inputs with a very large number of Unicode characters.
- CVE-2024-24680: An issue was discovered in Django 3.2 before 3.2.24, 4.2 before 4.2.10, and Django 5.0 before 5.0.2. The intcomma template filter was subject to a potential denial-of-service attack when used with very long strings.
- CVE-2022-44570: A denial of service vulnerability in the Range header parsing component of Rack >= 1.5.0. A Carefully crafted input can cause the Range header parsing component in Rack to take an unexpected amount of time, possibly resulting in a denial of service attack vector. Any applications that deal with Range requests (such as streaming applications, or applications that serve files) may be impacted.
- CVE-2023-27530: A DoS vulnerability exists in Rack <v3.0.4.2, <v2.2.6.3, <v2.1.4.3 and <v2.0.9.3 within in the Multipart MIME parsing code in which could allow an attacker to craft requests that can be abuse to cause multipart parsing to take longer than expected.
- CVE-2022-44571: There is a denial of service vulnerability in the Content-Disposition parsingcomponent of Rack fixed in 2.0.9.2, 2.1.4.2, 2.2.4.1, 3.0.0.1. This could allow an attacker to craft an input that can cause Content-Disposition header parsing in Rackto take an unexpected amount of time, possibly resulting in a denial ofservice attack vector. This header is used typically used in multipartparsing. Any applications that parse multipart posts using Rack (virtuallyall Rails applications) are impacted.
- CVE-2020-8184: A reliance on cookies without validation/integrity check security vulnerability exists in rack < 2.2.3, rack < 2.1.4 that makes it is possible for an attacker to forge a secure or host-only cookie prefix.
- CVE-2022-44572: A denial of service vulnerability in the multipart parsing component of Rack fixed in 2.0.9.2, 2.1.4.2, 2.2.4.1 and 3.0.0.1 could allow an attacker tocraft input that can cause RFC2183 multipart boundary parsing in Rack to take an unexpected amount of time, possibly resulting in a denial of service attack vector. Any applications that parse multipart posts using Rack (virtually all Rails applications) are impacted.
- CVE-2022-30122: A possible denial of service vulnerability exists in Rack <2.0.9.1, <2.1.4.1 and <2.2.3.1 in the multipart parsing component of Rack.
- CVE-2023-28319: A use after free vulnerability exists in curl <v8.1.0 in the way libcurl offers a feature to verify an SSH server's public key using a SHA 256 hash. When this check fails, libcurl would free the memory for the fingerprint before it returns an error message containing the (now freed) hash. This flaw risks inserting sensitive heap-based data into the error message that might be shown to users or otherwise get leaked and revealed.
- CVE-2023-35945: Envoy is a cloud-native high-performance edge/middle/service proxy. Envoy’s HTTP/2 codec may leak a header map and bookkeeping structures upon receiving `RST_STREAM` immediately followed by the `GOAWAY` frames from an upstream server. In nghttp2, cleanup of pending requests due to receipt of the `GOAWAY` frame skips de-allocation of the bookkeeping structure and pending compressed header. The error return [code path] is taken if connection is already marked for not sending more requests due to `GOAWAY` frame. The clean-up code is right after the return statement, causing memory leak. Denial of service through memory exhaustion. This vulnerability was patched in versions(s) 1.26.3, 1.25.8, 1.24.9, 1.23.11.
- RHSA-2023:4035: Node.js is a software development platform for building fast and scalable network applications in the JavaScript programming language.
- RHSA-2023:5362: Node.js is a software development platform for building fast and scalable network applications in the JavaScript programming language.
- RHSA-2023:5869: Node.js is a software development platform for building fast and scalable network applications in the JavaScript programming language.
- RHSA-2024:1435: PostgreSQL is an advanced object-relational database management system. The postgresql-jdbc package includes the .jar files needed for Java programs to access a PostgreSQL database.
- CVE-2024-23226: The issue was addressed with improved memory handling. This issue is fixed in macOS Sonoma 14.4, visionOS 1.1, iOS 17.4 and iPadOS 17.4, watchOS 10.4, tvOS 17.4. Processing web content may lead to arbitrary code execution.

- CVE-2023-42950: A use after free issue was addressed with improved memory management. This issue is fixed in Safari 17.2, iOS 17.2 and iPadOS 17.2, tvOS 17.2, watchOS 10.2, macOS Sonoma 14.2. Processing maliciously crafted web content may lead to arbitrary code execution.
- RHSA-2024:2126: WebKitGTK is the port of the portable web rendering engine WebKit to the GTK platform.
- CVE-2023-30608: sqlparse is a non-validating SQL parser module for Python. In affected versions the SQL parser contains a regular expression that is vulnerable to ReDoS (Regular Expression Denial of Service). This issue was introduced by commit `e75e358`. The vulnerability may lead to Denial of Service (DoS). This issues has been fixed in sqlparse 0.4.4 by commit `c457abd5f`. Users are advised to upgrade. There are no known workarounds for this issue.
- CVE-2023-6932: A use-after-free vulnerability in the Linux kernel's ipv4: igmp component can be exploited to achieve local privilege escalation. A race condition can be exploited to cause a timer be mistakenly registered on a RCU read locked object which is freed by another thread. We recommend upgrading past commit e2b706c691905fe78468c361aaabc719d0a496f1.
- CVE-2023-6931: A heap out-of-bounds write vulnerability in the Linux kernel's Performance Events system component can be exploited to achieve local privilege escalation. A perf_event's read_size can overflow, leading to an heap out-of-bounds increment or write in perf_read_group(). We recommend upgrading past commit 382c27f4ed28f803b1f1473ac2d8db0afc795a1b.
- CVE-2023-20588: A division-by-zero error on some AMD processors can potentially return speculative data resulting in loss of confidentiality.
- CVE-2023-40590: GitPython is a python library used to interact with Git repositories. When resolving a program, Python/Windows look for the current working directory, and after that the PATH environment. GitPython defaults to use the `git` command, if a user runs GitPython from a repo has a `git.exe` or `git` executable, that program will be run instead of the one in the user's `PATH`. This is more of a problem on how Python interacts with Windows systems, Linux and any other OS aren't affected by this. But probably people using GitPython usually run it from the CWD of a repo. An attacker can trick a user to download a repository with a malicious `git` executable, if the user runs/imports GitPython from that directory, it allows the attacker to run any arbitrary commands. There is no fix currently available for windows users, however there are a few mitigations. 1: Default to an absolute path for the git program on Windows, like `C:\\Program Files\\Git\\cmd\\git.EXE` (default git path installation). 2: Require users to set the `GIT_PYTHON_GIT_EXECUTABLE` environment variable on Windows systems. 3: Make this problem prominent in the documentation and advise users to never run GitPython from an untrusted repo, or set the `GIT_PYTHON_GIT_EXECUTABLE` env var to an absolute path. 4: Resolve the executable manually by only looking into the `PATH` environment variable.
- CVE-2023-32559: A privilege escalation vulnerability exists in the experimental policy mechanism in all active release lines: 16.x, 18.x and, 20.x. The use of the deprecated API `process.binding()` can bypass the policy mechanism by requiring internal modules and eventually take advantage of `process.binding('spawn_sync')` run arbitrary code, outside of the limits defined in a `policy.json` file. Please note that at the time this CVE was issued, the policy is an experimental feature of Node.js.
- CVE-2023-32006: The use of `module.constructor.createRequire()` can bypass the policy mechanism and require modules outside of the policy.json definition for a given module. This vulnerability affects all users using the experimental policy mechanism in all active release lines: 16.x, 18.x, and, 20.x. Please note that at the time this CVE was issued, the policy is an experimental feature of Node.js.
- CVE-2023-30585: A vulnerability has been identified in the Node.js (.msi version) installation process, specifically affecting Windows users who install Node.js using the .msi installer. This vulnerability emerges during the repair operation, where the ""msiexec.exe"" process, running under the NT AUTHORITY\SYSTEM context, attempts to read the %USERPROFILE% environment variable from the current user's registry.

  The issue arises when the path referenced by the %USERPROFILE% environment variable does not exist. In such cases, the ""msiexec.exe"" process attempts to create the specified path in an unsafe manner, potentially leading to the creation of arbitrary folders in arbitrary locations.

  The severity of this vulnerability is heightened by the fact that the %USERPROFILE% environment variable in the Windows registry can be modified by standard (or ""non-privileged"") users. Consequently, unprivileged actors, including malicious entities or trojans, can manipulate the environment variable key to deceive the

privileged ""msiexec.exe"" process. This manipulation can result in the creation of folders in unintended and potentially malicious locations.

It is important to note that this vulnerability is specific to Windows users who install Node.js using the .msi installer. Users who opt for other installation methods are not affected by this particular issue.

- CVE-2023-4807: Issue summary: The POLY1305 MAC (message authentication code) implementationcontains a bug that might corrupt the internal state of applications on theWindows 64 platform when running on newer X86_64 processors supporting theAVX512-IFMA instructions.Impact summary: If in an application that uses the OpenSSL library an attackercan influence whether the POLY1305 MAC algorithm is used, the applicationstate might be corrupted with various application dependent consequences.

  The POLY1305 MAC (message authentication code) implementation in OpenSSL doesnot save the contents of non-volatile XMM registers on Windows 64 platformwhen calculating the MAC of data larger than 64 bytes. Before returning tothe caller all the XMM registers are set to zero rather than restoring theirprevious content. The vulnerable code is used only on newer x86_64 processorssupporting the AVX512-IFMA instructions.

  The consequences of this kind of internal application state corruption canbe various - from no consequences, if the calling application does notdepend on the contents of non-volatile XMM registers at all, to the worstconsequences, where the attacker could get complete control of the applicationprocess. However given the contents of the registers are just zeroized sothe attacker cannot put arbitrary values inside, the most likely consequence,if any, would be an incorrect result of some application dependentcalculations or a crash leading to a denial of service.

  The POLY1305 MAC algorithm is most frequently used as part of theCHACHA20-POLY1305 AEAD (authenticated encryption with associated data)algorithm. The most common usage of this AEAD cipher is with TLS protocolversions 1.2 and 1.3 and a malicious client can influence whether this AEADcipher is used by the server. This implies that server applications usingOpenSSL can be potentially impacted. However we are currently not aware ofany concrete application that would be affected by this issue therefore weconsider this a Low severity security issue.

  As a workaround the AVX512-IFMA instructions support can be disabled atruntime by setting the environment variable OPENSSL_ia32cap: OPENSSL_ia32cap=:~0x200000The FIPS provider is not affected by this issue.
- CVE-2023-40283: An issue was discovered in l2cap_sock_release in net/bluetooth/l2cap_sock.c in the Linux kernel before 6.4.10. There is a use-after-free because the children of an sk are mishandled.
- CVE-2023-42752: An integer overflow flaw was found in the Linux kernel. This issue leads to the kernel allocating `skb_shared_info` in the userspace, which is exploitable in systems without SMAP protection since `skb_shared_info` contains references to function pointers.
- CVE-2023-1436: An infinite recursion is triggered in Jettison when constructing a JSONArray from a Collection that contains a self-reference in one of its elements. This leads to a StackOverflowError exception being thrown.
- CVE-2022-40149: Those using Jettison to parse untrusted XML or JSON data may be vulnerable to Denial of Service attacks (DOS). If the parser is running on user supplied input, an attacker may supply content that causes the parser to crash by stackoverflow. This effect may support a denial of service attack.
- CVE-2022-40150: Those using Jettison to parse untrusted XML or JSON data may be vulnerable to Denial of Service attacks (DOS). If the parser is running on user supplied input, an attacker may supply content that causes the parser to crash by Out of memory. This effect may support a denial of service attack.
- CVE-2022-45685: A stack overflow in Jettison before v1.5.2 allows attackers to cause a Denial of Service (DoS) via crafted JSON data.
- CVE-2022-45693: Jettison before v1.5.2 was discovered to contain a stack overflow via the map parameter. This vulnerability allows attackers to cause a Denial of Service (DoS) via a crafted string.
- RHSA-2024:2447: OpenSSL is a toolkit that implements the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols, as well as a full-strength general-purpose cryptography library.
- CVE-2020-29562: The iconv function in the GNU C Library (aka glibc or libc6) 2.30 to 2.32, when converting UCS4 text containing an irreversible character, fails an assertion in the code path and aborts the program, potentially resulting in a denial of service.
- CVE-2021-27645: The nameserver caching daemon (nscd) in the GNU C Library (aka glibc or libc6) 2.29 through 2.33, when processing a request for netgroup lookup, may crash due to a double-free, potentially resulting in degraded service or Denial of Service on the local system. This is related to netgroupcache.c.

- CVE-2020-12723: regcomp.c in Perl before 5.30.3 allows a buffer overflow via a crafted regular expression because of recursive S_study_chunk calls.
- CVE-2020-10878: Perl before 5.30.3 has an integer overflow related to mishandling of a "PL_regkind[OP(n)] == NOTHING" situation. A crafted regular expression could lead to malformed bytecode with a possibility of instruction injection.
- CVE-2020-10543: Perl before 5.30.3 on 32-bit platforms allows a heap-based buffer overflow because nested regular expression quantifiers have an integer overflow.
- CVE-2021-20232: A flaw was found in gnutls. A use after free issue in client_send_params in lib/ext/pre_shared_key.c may lead to memory corruption and other potential consequences.
- CVE-2021-20231: A flaw was found in gnutls. A use after free issue in client sending key_share extension may lead to memory corruption and other consequences.
- CVE-2023-38546: This flaw allows an attacker to insert cookies at will into a running program using libcurl, if the specific series of conditions are met. libcurl performs transfers. In its API, an application creates 'easy handles' that are the individual handles for single transfers. libcurl provides a function call that duplicates an easy handle called curl_easy_duphandle. If a transfer has cookies enabled when the handle is duplicated, the cookie-enable state is also cloned - but without cloning the actual cookies. If the source handle did not read any cookies from a specific file on disk, the cloned version of the handle would instead store the file name as none (using the four ASCII letters, no quotes).Subsequent use of the cloned handle that does not explicitly set a source to load cookies from would then inadvertently load cookies from a file named none - if such a file exists and is readable in the current directory of the program using libcurl. And if using the correct file format of course.
- CVE-2017-7244: The _pcre32_xclass function in pcre_xclass.c in libpcre1 in PCRE 8.40 allows remote attackers to cause a denial of service (invalid memory read) via a crafted file.
- CVE-2018-16429: GNOME GLib 2.56.1 has an out-of-bounds read vulnerability in g_markup_parse_context_parse() in gmarkup.c, related to utf8_str().
- CVE-2019-13012: The keyfile settings backend in GNOME GLib (aka glib2.0) before 2.60.0 creates directories using g_file_make_directory_with_parents (kfsb->dir, NULL, NULL) and files using g_file_replace_contents (kfsb->file, contents, length, NULL, FALSE, G_FILE_CREATE_REPLACE_DESTINATION, NULL, NULL, NULL). Consequently, it does not properly restrict directory (and file) permissions. Instead, for directories, 0777 permissions are used; for files, default file permissions are used. This is similar to CVE-2019-12450.
- CVE-2021-28153: An issue was discovered in GNOME GLib before 2.66.8. When g_file_replace() is used with G_FILE_CREATE_REPLACE_DESTINATION to replace a path that is a dangling symlink, it incorrectly also creates the target of the symlink as an empty file, which could conceivably have security relevance if the symlink is attacker-controlled. (If the path is a symlink to a file that already exists, then the contents of that file correctly remain unchanged.)
- CVE-2023-2602: A vulnerability was found in the pthread_create() function in libcap. This issue may allow a malicious actor to use cause __real_pthread_create() to return an error, which can exhaust the process memory.
- CVE-2015-2059: The stringprep_utf8_to_ucs4 function in libin before 1.31, as used in jabberd2, allows context-dependent attackers to read system memory and possibly have other unspecified impact via invalid UTF-8 characters in a string, which triggers an out-of-bounds read.
- CVE-2015-8948: idn in GNU libidn before 1.33 might allow remote attackers to obtain sensitive memory information by reading a zero byte as input, which triggers an out-of-bounds read.
- CVE-2017-5969: libxml2 2.9.4, when used in recover mode, allows remote attackers to cause a denial of service (NULL pointer dereference) via a crafted XML document. NOTE: The maintainer states "I would disagree of a CVE with the Recover parsing option which should only be used for manual recovery at least for XML parser.
- CVE-2017-8872: The htmlParseTryOrFinish function in HTMLparser.c in libxml2 2.9.4 allows attackers to cause a denial of service (buffer over-read) or information disclosure.
- CVE-2017-9048: libxml2 20904-GITv2.9.4-16-g0741801 is vulnerable to a stack-based buffer overflow. The function xmlSnprintfElementContent in valid.c is supposed to recursively dump the element content definition into a char buffer 'buf' of size 'size'. At the end of the routine, the function may strcat two more characters without checking whether the current strlen(buf) + 2 < size. This vulnerability causes programs that use libxml2, such as PHP, to crash.
- CVE-2016-4984: /usr/libexec/openldap/generate-server-cert.sh in openldap-servers sets weak permissions for the TLS certificate, which allows local users to obtain the TLS certificate by leveraging a race condition between the creation of the certificate, and the chmod to protect it.

- CVE-2017-11462: Double free vulnerability in MIT Kerberos 5 (aka krb5) allows attackers to have unspecified impact via vectors involving automatic deletion of security contexts on error.
- CVE-2016-8621: The `curl_getdate` function in curl before version 7.51.0 is vulnerable to an out of bounds read if it receives an input with one digit short.
- CVE-2016-8622: The URL percent-encoding decode function in libcurl before 7.51.0 is called `curl_easy_unescape`. Internally, even if this function would be made to allocate a unscape destination buffer larger than 2GB, it would return that new length in a signed 32 bit integer variable, thus the length would get either just truncated or both truncated and turned negative. That could then lead to libcurl writing outside of its heap based buffer.
- CVE-2016-8623: A flaw was found in curl before version 7.51.0. The way curl handles cookies permits other threads to trigger a use-after-free leading to information disclosure.
- CVE-2021-3200: Buffer overflow vulnerability in libsolv 2020-12-13 via the Solver * testcase_read(Pool *pool, FILE *fp, const char *testcase, Queue *job, char **resultp, int *resultflagsp function at src/testcase.c: line 2334, which could cause a denial of service
- CVE-2016-9586: curl before version 7.52.0 is vulnerable to a buffer overflow when doing a large floating point output in libcurl's implementation of the printf() functions. If there are any application that accepts a format string from the outside without necessary input filtering, it could allow remote attacks.
- CVE-2017-1000100: When doing a TFTP transfer and curl/libcurl is given a URL that contains a very long file name (longer than about 515 bytes), the file name is truncated to fit within the buffer boundaries, but the buffer size is still wrongly updated to use the untruncated length. This too large value is then used in the sendto() call, making curl attempt to send more data than what is actually put into the buffer. The endto() function will then read beyond the end of the heap based buffer. A malicious HTTP(S) server could redirect a vulnerable libcurl-using client to a crafted TFTP URL (if the client hasn't restricted which protocols it allows redirects to) and trick it to send private memory contents to a remote server over UDP. Limit curl's redirect protocols with --proto-redir and libcurl's with CURLOPT_REDIR_PROTOCOLS.
- CVE-2021-37621: Exiv2 is a command-line utility and C++ library for reading, writing, deleting, and modifying the metadata of image files. An infinite loop was found in Exiv2 versions v0.27.4 and earlier. The infinite loop is triggered when Exiv2 is used to print the metadata of a crafted image file. An attacker could potentially exploit the vulnerability to cause a denial of service, if they can trick the victim into running Exiv2 on a crafted image file. Note that this bug is only triggered when printing the image ICC profile, which is a less frequently used Exiv2 operation that requires an extra command line option (`-p C`). The bug is fixed in version v0.27.5.
- CVE-2021-37620: Exiv2 is a command-line utility and C++ library for reading, writing, deleting, and modifying the metadata of image files. An out-of-bounds read was found in Exiv2 versions v0.27.4 and earlier. The out-of-bounds read is triggered when Exiv2 is used to read the metadata of a crafted image file. An attacker could potentially exploit the vulnerability to cause a denial of service, if they can trick the victim into running Exiv2 on a crafted image file. The bug is fixed in version v0.27.5.
- CVE-2021-37616: Exiv2 is a command-line utility and C++ library for reading, writing, deleting, and modifying the metadata of image files. A null pointer dereference was found in Exiv2 versions v0.27.4 and earlier. The null pointer dereference is triggered when Exiv2 is used to print the metadata of a crafted image file. An attacker could potentially exploit the vulnerability to cause a denial of service, if they can trick the victim into running Exiv2 on a crafted image file. Note that this bug is only triggered when printing the interpreted (translated) data, which is a less frequently used Exiv2 operation that requires an extra command line option (`-p t` or `-P t`). The bug is fixed in version v0.27.5.
- CVE-2021-34335: Exiv2 is a command-line utility and C++ library for reading, writing, deleting, and modifying the metadata of image files. A floating point exception (FPE) due to an integer divide by zero was found in Exiv2 versions v0.27.4 and earlier. The FPE is triggered when Exiv2 is used to print the metadata of a crafted image file. An attacker could potentially exploit the vulnerability to cause a denial of service, if they can trick the victim into running Exiv2 on a crafted image file. Note that this bug is only triggered when printing the interpreted (translated) data, which is a less frequently used Exiv2 operation that requires an extra command line option (`-p t` or `-P t`). The bug is fixed in version v0.27.5.
- CVE-2021-37623: Exiv2 is a command-line utility and C++ library for reading, writing, deleting, and modifying the metadata of image files. An infinite loop was found in Exiv2 versions v0.27.4 and earlier. The infinite loop is triggered when Exiv2 is used to modify the metadata of a crafted image file. An attacker could potentially exploit the vulnerability to cause a denial of service, if they can trick the victim into running Exiv2 on a crafted

image file. Note that this bug is only triggered when deleting the IPTC data, which is a less frequently used Exiv2 operation that requires an extra command line option (`-d I rm`). The bug is fixed in version v0.27.5.

- CVE-2021-34334: Exiv2 is a command-line utility and C++ library for reading, writing, deleting, and modifying the metadata of image files. An infinite loop is triggered when Exiv2 is used to read the metadata of a crafted image file. An attacker could potentially exploit the vulnerability to cause a denial of service, if they can trick the victim into running Exiv2 on a crafted image file. The bug is fixed in version v0.27.5.
- CVE-2021-32815: Exiv2 is a command-line utility and C++ library for reading, writing, deleting, and modifying the metadata of image files. The assertion failure is triggered when Exiv2 is used to modify the metadata of a crafted image file. An attacker could potentially exploit the vulnerability to cause a denial of service, if they can trick the victim into running Exiv2 on a crafted image file. Note that this bug is only triggered when modifying the metadata, which is a less frequently used Exiv2 operation than reading the metadata. For example, to trigger the bug in the Exiv2 command-line application, you need to add an extra command-line argument such as `fi`. ### Patches The bug is fixed in version v0.27.5. ### References Regression test and bug fix: #1739 ### For more information Please see our security policy for information about Exiv2 security.
- CVE-2021-37622: Exiv2 is a command-line utility and C++ library for reading, writing, deleting, and modifying the metadata of image files. An infinite loop was found in Exiv2 versions v0.27.4 and earlier. The infinite loop is triggered when Exiv2 is used to modify the metadata of a crafted image file. An attacker could potentially exploit the vulnerability to cause a denial of service, if they can trick the victim into running Exiv2 on a crafted image file. Note that this bug is only triggered when deleting the IPTC data, which is a less frequently used Exiv2 operation that requires an extra command line option (`-d I rm`). The bug is fixed in version v0.27.5.
- CVE-2020-18771: Exiv2 0.27.99.0 has a global buffer over-read in Exiv2::Internal::Nikon1MakerNote::print0x0088 in nikonmn_int.cpp which can result in an information leak.
- CVE-2021-37615: Exiv2 is a command-line utility and C++ library for reading, writing, deleting, and modifying the metadata of image files. A null pointer dereference was found in Exiv2 versions v0.27.4 and earlier. The null pointer dereference is triggered when Exiv2 is used to print the metadata of a crafted image file. An attacker could potentially exploit the vulnerability to cause a denial of service, if they can trick the victim into running Exiv2 on a crafted image file. Note that this bug is only triggered when printing the interpreted (translated) data, which is a less frequently used Exiv2 operation that requires an extra command line option (`-p t` or `-P t`). The bug is fixed in version v0.27.5.
- CVE-2018-13419: An issue has been found in libsndfile 1.0.28. There is a memory leak in psf_allocate in common.c, as demonstrated by sndfile-convert. NOTE: The maintainer and third parties were unable to reproduce and closed the issue
- CVE-2023-4132: A use-after-free vulnerability was found in the siano smsusb module in the Linux kernel. The bug occurs during device initialization when the siano device is plugged in. This flaw allows a local user to crash the system, causing a denial of service condition.
- CVE-2021-41617: sshd in OpenSSH 6.2 through 8.x before 8.8, when certain non-default configurations are used, allows privilege escalation because supplemental groups are not initialized as expected. Helper programs for AuthorizedKeysCommand and AuthorizedPrincipalsCommand may run with privileges associated with group memberships of the sshd process, if the configuration specifies running the command as a different user.
- CVE-2023-35827: An issue was discovered in the Linux kernel through 6.3.8. A use-after-free was found in ravb_remove in drivers/net/ethernet/renesas/ravb_main.c.
- CVE-2023-3212: A NULL pointer dereference issue was found in the gfs2 file system in the Linux kernel. It occurs on corrupt gfs2 file systems when the evict code tries to reference the journal descriptor structure after it has been freed and set to NULL. A privileged local user could use this flaw to cause a kernel panic.
- CVE-2022-3162: Users authorized to list or watch one type of namespaced custom resource cluster-wide can read custom resources of a different type in the same API group without authorization. Clusters are impacted by this vulnerability if all of the following are true: 1. There are 2+ CustomResourceDefinitions sharing the same API group 2. Users have cluster-wide list or watch authorization on one of those custom resources. 3. The same users are not authorized to read another custom resource in the same API group.
- RHSA-2023:3042: GNU Emacs is a powerful, customizable, self-documenting text editor. It provides special code editing features, a scripting language (elisp), and the capability to read e-mail and news.
- RHSA-2024:0606: OpenSSH is an SSH protocol implementation supported by a number of Linux, UNIX, and similar operating systems. It includes the core files necessary for both the OpenSSH client and server.

- CVE-2024-23650: BuildKit is a toolkit for converting source code to build artifacts in an efficient, expressive and repeatable manner. A malicious BuildKit client or frontend could craft a request that could lead to BuildKit daemon crashing with a panic. The issue has been fixed in v0.12.5. As a workaround, avoid using BuildKit frontends from untrusted sources.
- RHSA-2023:2758: The container-tools module contains tools for working with containers, notably podman, buildah, skopeo, and runc.
- RHSA-2023:6939: The container-tools module contains tools for working with containers, notably podman, buildah, skopeo, and runc.
- RHSA-2023:2866: Git Large File Storage (LFS) replaces large files such as audio samples, videos, datasets, and graphics with text pointers inside Git, while storing the file contents on a remote server.
- CVE-2024-22025: A vulnerability in Node.js has been identified, allowing for a Denial of Service (DoS) attack through resource exhaustion when using the fetch() function to retrieve content from an untrusted URL.

  The vulnerability stems from the fact that the fetch() function in Node.js always decodes Brotli, making it possible for an attacker to cause resource exhaustion when fetching content from an untrusted URL.

  An attacker controlling the URL passed into fetch() can exploit this vulnerability to exhaust memory, potentially leading to process termination, depending on the system configuration.
- CVE-2022-29244: npm pack ignores root-level .gitignore and .npmignore file exclusion directives when run in a workspace or with a workspace flag (ie. `--workspaces`, `--workspace=<name>`). Anyone who has run `npm pack` or `npm publish` inside a workspace, as of v7.9.0 and v7.13.0 respectively, may be affected and have published files into the npm registry they did not intend to include. Users should upgrade to the latest, patched version of npm v8.11.0, run: npm i -g npm@latest . Node.js versions v16.15.1, v17.19.1, and v18.3.0 include the patched v8.11.0 version of npm.
- CVE-2023-46809: A flaw was found in Node.js. The privateDecrypt() API of the crypto library may allow a covert timing side-channel during PKCS#1 v1.5 padding error handling. This issue revealed significant timing differences in decryption for valid and invalid ciphertexts, which may allow a remote attacker to decrypt captured RSA ciphertexts or forge signatures, especially in scenarios involving API endpoints processing JSON Web Encryption messages.
- CVE-2024-27982: The team has identified a critical vulnerability in the http server of the most recent version of Node, where malformed headers can lead to HTTP request smuggling. Specifically, if a space is placed before a content-length header, it is not interpreted correctly, enabling attackers to smuggle in a second request within the body of the first.
- CVE-2024-29041: Express.js minimalist web framework for node. Versions of Express.js prior to 4.19.0 and all pre-release alpha and beta versions of 5.0 are affected by an open redirect vulnerability using malformed URLs. When a user of Express performs a redirect using a user-provided URL Express performs an encode [using `encodeurl`](https://github.com/pillarjs/encodeurl) on the contents before passing it to the `location` header. This can cause malformed URLs to be evaluated in unexpected ways by common redirect allow list implementations in Express applications, leading to an Open Redirect via bypass of a properly implemented allow list. The main method impacted is `res.location()` but this is also called from within `res.redirect()`. The vulnerability is fixed in 4.19.2 and 5.0.0-beta.3.
- RHSA-2023:7747: The libxml2 library is a development toolbox providing the implementation of various XML standards.
- RHSA-2024:0463: The RPM Package Manager (RPM) is a command-line driven package management system capable of installing, uninstalling, verifying, querying, and updating software packages.
- RHSA-2024:0465: SQLite is a C library that implements an SQL database engine. A large subset of SQL92 is supported. A complete database is stored in a single disk file. The API is designed for convenience and ease of use. Applications that link against SQLite can enjoy the power and flexibility of an SQL database without the administrative hassles of supporting a separate database server.
- RHSA-2024:2438: Pluggable Authentication Modules (PAM) provide a system to set up authentication policies without the need to recompile programs to handle authentication.
- CVE-2020-29363: An issue was discovered in p11-kit 0.23.6 through 0.23.21. A heap-based buffer overflow has been discovered in the RPC protocol used by p11-kit server/remote commands and the client library. When the remote entity supplies a serialized byte array in a CK_ATTRIBUTE, the receiving entity may not allocate sufficient length for the buffer to store the deserialized value.

- CVE-2020-27350: APT had several integer overflows and underflows while parsing .deb packages, aka GHSL-2020-168 GHSL-2020-169, in files apt-pkg/contrib/extracttar.cc, apt-pkg/deb/debfile.cc, and apt-pkg/contrib/arfile.cc. This issue affects: apt 1.2.32ubuntu0 versions prior to 1.2.32ubuntu0.2; 1.6.12ubuntu0 versions prior to 1.6.12ubuntu0.2; 2.0.2ubuntu0 versions prior to 2.0.2ubuntu0.2; 2.1.10ubuntu0 versions prior to 2.1.10ubuntu0.1;
- CVE-2020-24659: An issue was discovered in GnuTLS before 3.6.15. A server can trigger a NULL pointer dereference in a TLS 1.3 client if a no_renegotiation alert is sent with unexpected timing, and then an invalid second handshake occurs. The crash happens in the application's error handling path, where the gnutls_deinit function is called after detecting a handshake failure.
- CVE-2023-32360: An authentication issue was addressed with improved state management. This issue is fixed in macOS Big Sur 11.7.7, macOS Monterey 12.6.6, macOS Ventura 13.4. An unauthenticated user may be able to access recently printed documents.
- CVE-2023-34241: OpenPrinting CUPS is a standards-based, open source printing system for Linux and other Unix-like operating systems. Starting in version 2.0.0 and prior to version 2.4.6, CUPS logs data of free memory to the logging service AFTER the connection has been closed, when it should have logged the data right before. This is a use-after-free bug that impacts the entire cupsd process.

  The exact cause of this issue is the function `httpClose(con->http)` being called in `scheduler/client.c`. The problem is that httpClose always, provided its argument is not null, frees the pointer at the end of the call, only for cupsdLogClient to pass the pointer to httpGetHostname. This issue happens in function `cupsdAcceptClient` if LogLevel is warn or higher and in two scenarios: there is a double-lookup for the IP Address (HostNameLookups Double is set in `cupsd.conf`) which fails to resolve, or if CUPS is compiled with TCP wrappers and the connection is refused by rules from `/etc/hosts.allow` and `/etc/hosts.deny`.

  Version 2.4.6 has a patch for this issue.
- CVE-2021-3995: A logic error was found in the libmount library of util-linux in the function that allows an unprivileged user to unmount a FUSE filesystem. This flaw allows an unprivileged local attacker to unmount FUSE filesystems that belong to certain other users who have a UID that is a prefix of the UID of the attacker in its string form. An attacker may use this flaw to cause a denial of service to applications that use the affected filesystems.
- CVE-2021-3996: A logic error was found in the libmount library of util-linux in the function that allows an unprivileged user to unmount a FUSE filesystem. This flaw allows a local user on a vulnerable system to unmount other users' filesystems that are either world-writable themselves (like /tmp) or mounted in a world-writable directory. An attacker may use this flaw to cause a denial of service to applications that use the affected filesystems.
- CVE-2023-3138: A vulnerability was found in libX11. The security flaw occurs because the functions in src/InitExt.c in libX11 do not check that the values provided for the Request, Event, or Error IDs are within the bounds of the arrays that those functions write to, using those IDs as array indexes. They trust that they were called with values provided by an Xserver adhering to the bounds specified in the X11 protocol, as all X servers provided by X.Org do. As the protocol only specifies a single byte for these values, an out-of-bounds value provided by a malicious server (or a malicious proxy-in-the-middle) can only overwrite other portions of the Display structure and not write outside the bounds of the Display structure itself, possibly causing the client to crash with this memory corruption.
- CVE-2021-20305: A flaw was found in Nettle in versions before 3.7.2, where several Nettle signature verification functions (GOST DSA, EDDSA & ECDSA) result in the Elliptic Curve Cryptography point (ECC) multiply function being called with out-of-range scalers, possibly resulting in incorrect results. This flaw allows an attacker to force an invalid signature, causing an assertion failure or possible validation. The highest threat to this vulnerability is to confidentiality, integrity, as well as system availability.
- CVE-2021-3580: A flaw was found in the way nettle's RSA decryption functions handled specially crafted ciphertext. An attacker could use this flaw to provide a manipulated ciphertext leading to application crash and denial of service.
- CVE-2021-24031: In the Zstandard command-line utility prior to v1.4.1, output files were created with default permissions. Correct file permissions (matching the input) would only be set at completion time. Output files could therefore be readable or writable to unintended parties.
- CVE-2023-22045: Vulnerability in the Oracle Java SE, Oracle GraalVM Enterprise Edition, Oracle GraalVM for JDK product of Oracle Java SE (component: Hotspot). Supported versions that are affected are Oracle Java

SE: 8u371, 8u371-perf, 11.0.19, 17.0.7, 20.0.1; Oracle GraalVM Enterprise Edition: 20.3.10, 21.3.6, 22.3.2; Oracle GraalVM for JDK: 17.0.7 and 20.0.1. Difficult to exploit vulnerability allows unauthenticated attacker with network access via multiple protocols to compromise Oracle Java SE, Oracle GraalVM Enterprise Edition, Oracle GraalVM for JDK. Successful attacks of this vulnerability can result in unauthorized read access to a subset of Oracle Java SE, Oracle GraalVM Enterprise Edition, Oracle GraalVM for JDK accessible data. Note: This vulnerability can be exploited by using APIs in the specified Component, e.g., through a web service which supplies data to the APIs. This vulnerability also applies to Java deployments, typically in clients running sandboxed Java Web Start applications or sandboxed Java applets, that load and run untrusted code (e.g., code that comes from the internet) and rely on the Java sandbox for security. CVSS 3.1 Base Score 3.7 (Confidentiality impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:H/PR:N/UI:N/S:U/C:L/I:N/A:N).

- CVE-2023-22049: Vulnerability in the Oracle Java SE, Oracle GraalVM Enterprise Edition, Oracle GraalVM for JDK product of Oracle Java SE (component: Libraries). Supported versions that are affected are Oracle Java SE: 8u371, 8u371-perf, 11.0.19, 17.0.7, 20.0.1; Oracle GraalVM Enterprise Edition: 20.3.10, 21.3.6, 22.3.2; Oracle GraalVM for JDK: 17.0.7 and 20.0.1. Difficult to exploit vulnerability allows unauthenticated attacker with network access via multiple protocols to compromise Oracle Java SE, Oracle GraalVM Enterprise Edition, Oracle GraalVM for JDK. Successful attacks of this vulnerability can result in unauthorized update, insert or delete access to some of Oracle Java SE, Oracle GraalVM Enterprise Edition, Oracle GraalVM for JDK accessible data. Note: This vulnerability can be exploited by using APIs in the specified Component, e.g., through a web service which supplies data to the APIs. This vulnerability also applies to Java deployments, typically in clients running sandboxed Java Web Start applications or sandboxed Java applets, that load and run untrusted code (e.g., code that comes from the internet) and rely on the Java sandbox for security. CVSS 3.1 Base Score 3.7 (Integrity impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:H/PR:N/UI:N/S:U/C:N/I:L/A:N).
- RHSA-2023:7034: Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exceptions, very high level dynamic data types and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing systems.
- CVE-2023-49081: aiohttp is an asynchronous HTTP client/server framework for asyncio and Python. Improper validation made it possible for an attacker to modify the HTTP request (e.g. to insert a new header) or create a new HTTP request if the attacker controls the HTTP version. The vulnerability only occurs if the attacker can control the HTTP version of the request. This issue has been patched in version 3.9.0.
- CVE-2024-23829: aiohttp is an asynchronous HTTP client/server framework for asyncio and Python. Security-sensitive parts of the Python HTTP parser retained minor differences in allowable character sets, that must trigger error handling to robustly match frame boundaries of proxies in order to protect against injection of additional requests. Additionally, validation could trigger exceptions that were not handled consistently with processing of other malformed input. Being more lenient than internet standards require could, depending on deployment environment, assist in request smuggling. The unhandled exception could cause excessive resource consumption on the application server and/or its logging facilities. This vulnerability exists due to an incomplete fix for CVE-2023-47627. Version 3.9.2 fixes this vulnerability.
- CVE-2023-49082: aiohttp is an asynchronous HTTP client/server framework for asyncio and Python. Improper validation makes it possible for an attacker to modify the HTTP request (e.g. insert a new header) or even create a new HTTP request if the attacker controls the HTTP method. The vulnerability occurs only if the attacker can control the HTTP method (GET, POST etc.) of the request. If the attacker can control the HTTP version of the request it will be able to modify the request (request smuggling). This issue has been patched in version 3.9.0.
- CVE-2024-25629: c-ares is a C library for asynchronous DNS requests. `ares__read_line()` is used to parse local configuration files such as `/etc/resolv.conf`, `/etc/nsswitch.conf`, the `HOSTALIASES` file, and if using a c-ares version prior to 1.27.0, the `/etc/hosts` file. If any of these configuration files has an embedded `NULL` character as the first character in a new line, it can lead to attempting to read memory prior to the start of the given buffer which may result in a crash. This issue is fixed in c-ares 1.27.0. No known workarounds exist.
- CVE-2023-23916: An allocation of resources without limits or throttling vulnerability exists in curl <v7.88.0 based on the "chained" HTTP compression algorithms, meaning that a server response can be compressed multiple times and potentially with differentalgorithms. The number of acceptable "links" in this "decompression chain" wascapped, but the cap was implemented on a per-header basis allowing a maliciousserver to insert a virtually unlimited number of compression steps simply byusing many headers. The use of such a decompression chain could result in a "malloc bomb", making curl end up spending enormous amounts of allocated heap memory, or trying to and returning out of memory errors.

- CVE-2023-27537: A double free vulnerability exists in libcurl <8.0.0 when sharing HSTS data between separate "handles". This sharing was introduced without considerations for do this sharing across separate threads but there was no indication of this fact in the documentation. Due to missing mutexes or thread locks, two threads sharing the same HSTS data could end up doing a double-free or use-after-free.
- CVE-2018-1002104: Versions < 1.5 of the Kubernetes ingress default backend, which handles invalid ingress traffic, exposed prometheus metrics publicly.
- DSA-5686-1: dav1d - security update
- RHSA-2024:1530: Expat is a C library for parsing XML documents.
- RHSA-2023:1583: Node.js is a software development platform for building fast and scalable network applications in the JavaScript programming language.
- RHSA-2023:4536: Node.js is a software development platform for building fast and scalable network applications in the JavaScript programming language.
- RHSA-2022:1830: PostgreSQL is an advanced object-relational database management system (DBMS).
- CVE-2021-3782: An internal reference count is held on the buffer pool, incremented every time a new buffer is created from the pool. The reference count is maintained as an int; on LP64 systems this can cause the reference count to overflow if the client creates a large number of wl_shm buffer objects, or if it can coerce the server to create a large number of external references to the buffer storage. With the reference count overflowing, a use-after-free can be constructed on the wl_shm_pool tracking structure, where values may be incremented or decremented; it may also be possible to construct a limited oracle to leak 4 bytes of server-side memory to the attacking client at a time.
- CVE-2020-36023: An issue was discovered in freedesktop poppler version 20.12.1, allows remote attackers to cause a denial of service (DoS) via crafted .pdf file to FoFiType1C::cvtGlyph function.
- CVE-2020-36024: An issue was discovered in freedesktop poppler version 20.12.1, allows remote attackers to cause a denial of service (DoS) via crafted .pdf file to FoFiType1C::convertToType1 function.
- CVE-2022-37050: In Poppler 22.07.0, PDFDoc::savePageAs in PDFDoc.c callows attackers to cause a denial-of-service (application crashes with SIGABRT) by crafting a PDF file in which the xref data structure is mishandled in getCatalog processing. Note that this vulnerability is caused by the incomplete patch of CVE-2018-20662.
- CVE-2022-37051: An issue was discovered in Poppler 22.07.0. There is a reachable abort which leads to denial of service because the main function in pdfunite.cc lacks a stream check before saving an embedded file.
- CVE-2022-37052: A reachable Object::getString assertion in Poppler 22.07.0 allows attackers to cause a denial of service due to a failure in markObject.
- RHSA-2024:2302: GStreamer is a streaming media framework based on graphs of filters which operate on media data. The gstreamer1-plugins-base packages contain a collection of well-maintained base plug-ins.
- RHSA-2024:2295: The libjpeg-turbo packages contain a library of functions for manipulating JPEG images. They also contain simple client programs for accessing the libjpeg functions. These packages provide the same functionality and API as libjpeg but with better performance.
- RHSA-2024:2184: libsndfile is a C library for reading and writing files containing sampled sound, such as AIFF, AU, or WAV.
- RHSA-2024:2410: HarfBuzz is an implementation of the OpenType Layout engine.
- CVE-2023-42843: An inconsistent user interface issue was addressed with improved state management. This issue is fixed in iOS 16.7.2 and iPadOS 16.7.2, iOS 17.1 and iPadOS 17.1, Safari 17.1, macOS Sonoma 14.1. Visiting a malicious website may lead to address bar spoofing.
- CVE-2023-42956: The issue was addressed with improved memory handling. This issue is fixed in Safari 17.2, iOS 17.2 and iPadOS 17.2, macOS Sonoma 14.2. Processing web content may lead to a denial-of-service.
- CVE-2024-23252: Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.
- CVE-2024-23254: The issue was addressed with improved UI handling. This issue is fixed in tvOS 17.4, macOS Sonoma 14.4, visionOS 1.1, iOS 17.4 and iPadOS 17.4, watchOS 10.4, Safari 17.4. A malicious website may exfiltrate audio data cross-origin.
- CVE-2024-23263: A logic issue was addressed with improved validation. This issue is fixed in tvOS 17.4, macOS Sonoma 14.4, visionOS 1.1, iOS 17.4 and iPadOS 17.4, watchOS 10.4, iOS 16.7.6 and iPadOS 16.7.6, Safari 17.4. Processing maliciously crafted web content may prevent Content Security Policy from being enforced.

- CVE-2024-23280: An injection issue was addressed with improved validation. This issue is fixed in Safari 17.4, macOS Sonoma 14.4, iOS 17.4 and iPadOS 17.4, watchOS 10.4, tvOS 17.4. A maliciously crafted webpage may be able to fingerprint the user.
- CVE-2024-23284: A logic issue was addressed with improved state management. This issue is fixed in tvOS 17.4, macOS Sonoma 14.4, visionOS 1.1, iOS 17.4 and iPadOS 17.4, watchOS 10.4, iOS 16.7.6 and iPadOS 16.7.6, Safari 17.4. Processing maliciously crafted web content may prevent Content Security Policy from being enforced.
- RHSA-2024:2145: The libX11 packages contain the core X11 protocol client library.
- RHSA-2024:2433: Avahi is an implementation of the DNS Service Discovery and Multicast DNS specifications for Zero Configuration Networking. It facilitates service discovery on a local network. Avahi and Avahi-aware applications allow you to plug your computer into a network and, with no configuration, view other people to chat with, view printers to print with, and find shared files on other computers.
- RHSA-2024:2289: The libtiff packages contain a library of functions for manipulating Tagged Image File Format (TIFF) files.
- RHSA-2023:2867: PostgreSQL is an advanced object-relational database management system. The postgresql-jdbc package includes the .jar files needed for Java programs to access a PostgreSQL database.
- CVE-2022-21724: pgjdbc is the offical PostgreSQL JDBC Driver. A security hole was found in the jdbc driver for postgresql database while doing security research. The system using the postgresql library will be attacked when attacker control the jdbc url or properties. pgjdbc instantiates plugin instances based on class names provided via `authenticationPluginClassName`, `sslhostnameverifier`, `socketFactory`, `sslfactory`, `sslpasswordcallback` connection properties. However, the driver did not verify if the class implements the expected interface before instantiating the class. This can lead to code execution loaded via arbitrary classes. Users using plugins are advised to upgrade. There are no known workarounds for this issue.
- CVE-2023-1206: A hash collision flaw was found in the IPv6 connection lookup table in the Linux kernel’s IPv6 functionality when a user makes a new kind of SYN flood attack. A user located in the local network or with a high bandwidth connection can increase the CPU usage of the server that accepts IPV6 connections up to 95%.
- CVE-2023-3338: A null pointer dereference flaw was found in the Linux kernel's DECnet networking protocol. This issue could allow a remote user to crash the system.
- CVE-2023-34319: The fix for XSA-423 added logic to Linux'es netback driver to deal with a frontend splitting a packet in a way such that not all of the headers would come in one piece. Unfortunately the logic introduced there didn't account for the extreme case of the entire packet being split into as many pieces as permitted by the protocol, yet still being smaller than the area that's specially dealt with to keep all (possible) headers together. Such an unusual packet would therefore trigger a buffer overrun in the driver.
- CVE-2023-34324: Closing of an event channel in the Linux kernel can result in a deadlock. This happens when the close is being performed in parallel to an unrelated Xen console action and the handling of a Xen console interrupt in an unprivileged guest.

  The closing of an event channel is e.g. triggered by removal of a paravirtual device on the other side. As this action will cause console messages to be issued on the other side quite often, the chance of triggering the deadlock is not neglectable.

  Note that 32-bit Arm-guests are not affected, as the 32-bit Linux kernel on Arm doesn't use queued-RW-locks, which are required to trigger the issue (on Arm32 a waiting writer doesn't block further readers to get the lock).
- CVE-2023-3863: A use-after-free flaw was found in nfc_llcp_find_local in net/nfc/llcp_core.c in NFC in the Linux kernel. This flaw allows a local user with special privileges to impact a kernel information leak issue.
- CVE-2023-4194: A flaw was found in the Linux kernel's TUN/TAP functionality. This issue could allow a local user to bypass network filters and gain unauthorized access to some resources. The original patches fixing CVE-2023-1076 are incorrect or incomplete. The problem is that the following upstream commits - a096ccca6e50 ("tun: tun_chr_open(): correctly initialize socket uid"), - 66b2c338adce ("tap: tap_open(): correctly initialize socket uid"), pass "inode->i_uid" to sock_init_data_uid() as the last parameter and that turns out to not be accurate.
- CVE-2023-3341: The code that processes control channel messages sent to `named` calls certain functions recursively during packet parsing. Recursion depth is only limited by the maximum accepted packet size; depending on the environment, this may cause the packet-parsing code to run out of available stack memory, causing `named` to terminate unexpectedly. Since each incoming control channel message is fully parsed before

its contents are authenticated, exploiting this flaw does not require the attacker to hold a valid RNDC key; only network access to the control channel's configured TCP port is necessary.

This issue affects BIND 9 versions 9.2.0 through 9.16.43, 9.18.0 through 9.18.18, 9.19.0 through 9.19.16, 9.9.3-S1 through 9.16.43-S1, and 9.18.0-S1 through 9.18.18-S1.

- CVE-2021-4001: A race condition was found in the Linux kernel's ebpf verifier between bpf_map_update_elem and bpf_map_freeze due to a missing lock in kernel/bpf/syscall.c. In this flaw, a local user with a special privilege (cap_sys_admin or cap_bpf) can modify the frozen mapped address space. This flaw affects kernel versions prior to 5.16 rc2.
- CVE-2021-46174: Heap-based Buffer Overflow in function bfd_getl32 in Binutils objdump 3.37.
- CVE-2022-35205: An issue was discovered in Binutils readelf 2.38.50, reachable assertion failure in function display_debug_names allows attackers to cause a denial of service.
- CVE-2022-44840: Heap buffer overflow vulnerability in binutils readelf before 2.40 via function find_section_in_set in file readelf.c.
- CVE-2022-45703: Heap buffer overflow vulnerability in binutils readelf before 2.40 via function display_debug_section in file readelf.c.
- CVE-2022-47008: An issue was discovered function make_tempdir, and make_tempname in bucomm.c in Binutils 2.34 thru 2.38, allows attackers to cause a denial of service due to memory leaks.
- CVE-2020-19726: An issue was discovered in binutils libbfd.c 2.36 relating to the auxiliary symbol data allows attackers to read or write to system memory or cause a denial of service.
- CVE-2023-51385: In ssh in OpenSSH before 9.6, OS command injection might occur if a user name or host name has shell metacharacters, and this name is referenced by an expansion token in certain situations. For example, an untrusted Git repository can have a submodule with shell metacharacters in a user name or host name.
- CVE-2023-41040: GitPython is a python library used to interact with Git repositories. In order to resolve some git references, GitPython reads files from the `.git` directory, in some places the name of the file being read is provided by the user, GitPython doesn't check if this file is located outside the `.git` directory. This allows an attacker to make GitPython read any file from the system. This vulnerability is present in https://github.com/gitpython-developers/GitPython/blob/1c8310d7cae144f74a671cbe17e51f63a830adbf/git/refs/symbolic.py#L174-L175. That code joins the base directory with a user given string without checking if the final path is located outside the base directory. This vulnerability cannot be used to read the contents of files but could in theory be used to trigger a denial of service for the program. This issue has not yet been addressed.
- CVE-2023-5178: A use-after-free vulnerability was found in drivers/nvme/target/tcp.c` in `nvmet_tcp_free_crypto` due to a logical bug in the NVMe/TCP subsystem in the Linux kernel. This issue may allow a malicious user to cause a use-after-free and double-free problem, which may permit remote code execution or lead to local privilege escalation.
- CVE-2023-5717: A heap out-of-bounds write vulnerability in the Linux kernel's Linux Kernel Performance Events (perf) component can be exploited to achieve local privilege escalation.

  If perf_read_group() is called while an event's sibling_list is smaller than its child's sibling_list, it can increment or write to memory locations outside of the allocated buffer.

  We recommend upgrading past commit 32671e3799ca2e4590773fd0e63aaa4229e50c06.
- CVE-2018-25091: urllib3 before 1.24.2 does not remove the authorization HTTP header when following a cross-origin redirect (i.e., a redirect that differs in host, port, or scheme). This can allow for credentials in the authorization header to be exposed to unintended hosts or transmitted in cleartext. Note: his issue exists because of an incomplete fix for CVE-2018-20060 (which was case-sensitive).
- CVE-2023-38552: When the Node.js policy feature checks the integrity of a resource against a trusted manifest, the application can intercept the operation and return a forged checksum to the node's policy implementation, thus effectively disabling the integrity check.
- Impacts: This vulnerability affects all users using the experimental policy mechanism in all active release lines: 18.x and, 20.x. Note that at the time this CVE was issued, the policy mechanism is an experimental feature of Node.js.
- CVE-2019-15847: The POWER9 backend in GNU Compiler Collection (GCC) before version 10 could optimize multiple calls of the __builtin_darn intrinsic into a single call, thus reducing the entropy of the random number generator. This occurred because a volatile operation was not specified. For example, within a single execution of a program, the output of every __builtin_darn() call may be the same.

- CVE-2018-13410: An issue was discovered IW44Image.cpp in djvulibre 3.5.28 in allows attackers to cause a denial of service via divide by zero.
- CVE-2021-46312: An issue was discovered IW44EncodeCodec.cpp in djvulibre 3.5.28 in allows attackers to cause a denial of service via divide by zero.
- CVE-2021-31239: An issue found in SQLite SQLite3 v.3.35.4 that allows a remote attacker to cause a denial of service via the appendvfs.c function.
- CVE-2021-45346: A Memory Leak vulnerability exists in SQLite Project SQLite3 3.35.1 and 3.37.0 via maliciously crafted SQL Queries (made via editing the Database File), it is possible to query a record, and leak subsequent bytes of memory that extend beyond the record, which could let a malicious user obtain sensitive information. NOTE: The developer disputes this as a vulnerability stating that If you give SQLite a corrupted database file and submit a query against the database, it might read parts of the database that you did not intend or expect.
- CVE-2023-32570: VideoLAN dav1d before 1.2.0 has a thread_task.c race condition that can lead to an application crash, related to dav1d_decode_frame_exit.
- TEMP-0841856-B18BAF
- CVE-2018-13410: Info-ZIP Zip 3.0, when the -T and -TT command-line options are used, allows attackers to cause a denial of service (invalid free and application crash) or possibly have unspecified other impact because of an off-by-one error. NOTE: it is unclear whether there are realistic scenarios in which an untrusted party controls the -TT value, given that the entire purpose of -TT is execution of arbitrary commands
- CVE-2024-28757: libexpat through 2.6.1 allows an XML Entity Expansion attack when there is isolated use of external parsers (created via XML_ExternalEntityParserCreate).
- CVE-2012-0039: GLib 2.31.8 and earlier, when the g_str_hash function is used, computes hash values without restricting the ability to trigger hash collisions predictably, which allows context-dependent attackers to cause a denial of service (CPU consumption) via crafted input to an application that maintains a hash table. NOTE: this issue may be disputed by the vendor; the existence of the g_str_hash function is not a vulnerability in the library, because callers of g_hash_table_new and g_hash_table_new_full can specify an arbitrary hash function that is appropriate for the application.
- CVE-2022-2817: Use After Free in GitHub repository vim/vim prior to 9.0.0213.
- CVE-2022-2862: Use After Free in GitHub repository vim/vim prior to 9.0.0221.
- CVE-2022-2874: NULL Pointer Dereference in GitHub repository vim/vim prior to 9.0.0224.
- CVE-2022-2889: Use After Free in GitHub repository vim/vim prior to 9.0.0225.
- CVE-2022-2982: Use After Free in GitHub repository vim/vim prior to 9.0.0260.
- CVE-2022-3016: Use After Free in GitHub repository vim/vim prior to 9.0.0286.
- CVE-2022-3099: Use After Free in GitHub repository vim/vim prior to 9.0.0360.
- CVE-2022-3134: Use After Free in GitHub repository vim/vim prior to 9.0.0389.
- CVE-2014-8166: The browsing feature in the server in CUPS does not filter ANSI escape sequences from shared printer names, which might allow remote attackers to execute arbitrary code via a crafted printer name.

Cumulative hotfixes

The cumulative hotfixes that have been shipped for Cloudera Data Services on premises 1.5.4-CHF1.

Cloudera Data Services on premises 1.5.4-CHF1

The cumulative hotfixes for new features, known issues, and fixed issues for 1.5.4-CHF1.



Note: ECS Customers: Direct upgrade path is not available. Customers must upgrade to Cloudera Data Services on premises 1.5.4 prior to consuming any CHFs built on top of 1.5.4.

Note:

OCP Customers: Direct upgrade path is available. Customers can directly upgrade from Cloudera Data Services on premises 1.5.2 to any 1.5.4 CHFs.

Whats new in Cloudera Private Cloud Data Services1.5.4-CHF1

New features introduced in this cumulative hotfix release of CDP Private Cloud Data Services 1.5.4-CHF1.

Note: Cloudera Manager 7.11.3 CHF7 Data Services (version: 7.11.3.14) support Cloudera Private Cloud Data Services 1.5.4 CHF1 release.

Note: Cloudera Manager 7.11.3 CHF8 does not support any Cloudera Private Cloud Data Services release.

Restore CP namespaces independently from system-generated DRS backups

Enhancements to the system-generated DRS backups:

- System-generated backups in DRS are automatic, periodic backups that include control plane and data services' namespaces.
- Through Private Cloud Data Services 1.5.4 release, restoring a system-generated backup from the private cloud management console UI restores all the namespaces present in the backup.
- With this change, such a restore action independently restores only control plane namespaces that are present in the backup.

Kubernetes Ingress NGINX Controller vulnerabilities

Kubernetes Ingress NGINX Controller vulnerabilities.

Five vulnerabilities affecting the Ingress Nginx Controller for Kubernetes were publicly disclosed on March 24, 2025, and were given the nickname "IngressNightmare."

The ‘IngressNightmare’ vulnerabilities may allow Remote Code Execution (RCE) and potentially expose Kubernetes clusters to malicious configuration modifications. Exploitation requires specially crafted HTTP requests that bypass security measures, such as a Web Application Firewall (WAF). Successful exploitation may lead to complete cluster compromise, data exfiltration, and denial of service.

Details of the CVEs:

- CVE-2025-1974 (CVSS score: 9.8) – An unauthenticated attacker with access to the pod network can achieve arbitrary code execution in the context of the ingress-nginx controller under certain conditions
- CVE-2025-24514 (CVSS score: 8.8) – The auth-url Ingress annotation can be used to inject configuration into NGINX, resulting in arbitrary code execution in the context of the ingress-nginx controller and disclosure of secrets accessible to the controller
- CVE-2025-1097 (CVSS score: 8.8) – The auth-tls-match-cn Ingress annotation can be used to inject configuration into NGINX, resulting in arbitrary code execution in the context of the ingress-nginx controller and disclosure of secrets accessible to the controller
- CVE-2025-1098 (CVSS score: 8.8) – The mirror-target and mirror-host Ingress annotations can be used to inject arbitrary configuration into NGINX, resulting in arbitrary code execution in the context of the ingress-nginx controller and disclosure of secrets accessible to the controller
- CVE-2025-24513 (CVSS score: 4.8) – An improper input validation vulnerability that could result in directory traversal within the container, leading to denial-of-service (DoS) or limited disclosure of secret objects from the cluster when combined with other vulnerabilities

Note: Cloudera will not provide remediation options for unsupported versions, and has not tested mitigations on unsupported versions. Customers are advised to upgrade to a supported product version. For more information, refer to the Support Lifecycle Policy page.

Action required

Multiple mitigation scenarios are possible depending on the version of Cloudera Data Services deployed.

1. Check your Data Services version.

   To determine which version of Data Services is running, log into Cloudera Manager UI, click “Data Services” in the left sidebar, and find each Data Services cluster's version(s).
2. Check your Cloudera Manager version.

   To determine which version of Cloudera Manager is running, log into Cloudera Manager UI and click on the version number at the lower-left corner of the home page. The Cloudera Manager 7.11.3 CHF8 Build number shows 56304673.

Mitigation steps for Cloudera Data Services On Premises:

1. No action is needed in the following scenarios:
   - Running Data Services on OpenShift Container Platform (OCP) and using the default HAproxy ingress controller. If you have customized OCP to use NGINX as the ingress controller, please contact RedHat support for further information
   - Running Data Services 1.5.4 or later with Cloudera Manager 7.11.3 CHF8 or later. These versions already disable the affected admission webhook component.
2. Data Services before 1.5.4, or Cloudera Manager prior to 7.11.3 CHF8, perform one of the following:
   - (Recommended) Upgrade Cloudera Manager to version 7.13.1 CHF1 or higher and all Embedded Container Service clusters to 1.5.4 SP2 or higher. This will not upgrade the nginx version but will automatically disable the webhook via Cloudera Manager, which mitigates the CVEs.
   - If an upgrade is not possible, follow the below steps to disable the webhook to mitigate the CVEs manually.

     Note: These steps should be run on each ECS cluster. If you create a new ECS cluster while on an affected release, reapply these steps on the new cluster(s).

1. SSH into the ECS server host.
2. Create the following YAML file to patch the built-in Nginx Ingress Controller Helm Chart:

```
apiVersion: helm.cattle.io/v1
kind: HelmChartConfig
metadata:
 name: rke2-ingress-nginx
 namespace: kube-system
spec:
 valuesContent: |-
  controller:
   admissionWebhooks:
    enabled: false
```

3. Run the following commands to apply the patch:

```
alias kubectl=/var/lib/rancher/rke2/bin/kubectl
export KUBECONFIG=/etc/rancher/rke2/rke2.yaml
kubectl apply -f <PATCH_FILENAME>.yaml
```

Note: No cluster restart or ECS refresh is required.

Note: Cloudera will not provide remediation options for unsupported versions, and has not tested mitigations on unsupported versions. Customers are advised to upgrade to a supported product version. For more information, refer to the Support Lifecycle Policy page

Cloudera Data Services on premises version 1.5.4 SP2 CHF1 will contain fixes to upgrade the Ingress NGINX controller. Target release date: April 25. 2025.

For the latest update on this issue, see the corresponding Knowledge article: TSB 2025-839: Critical Kubernetes Ingress NGINX Controller Vulnerability Allows RCE Without Authentication

For the latest update on this issue, see the corresponding Data Services Knowledge article: https://my.cloudera.com/knowledge/TSB-2025-839-Mitigation-steps-for-Cloudera-Data-Services-on?id=405006

Known Issues in Cloudera Private Cloud Data Services 1.5.4-CHF1

Following are new known issues in the 1.5.4 cumulative hotfix CHF1 release of CDP Private Cloud Data Services.

DSE-36967 - Namespace Termination issue when using Portworx storage

There is an issue with Portworx version lower than 3.1.1, as the namespace deletion gets stuck in terminating state. Portworx is not able to cleanly unmount and clean up the underlying resources.

The issues is fixed with Portworx version 3.1.1. Upgrade to Portworx version 3.1.1 or to later versions.

OPSX-5413 - Implement Rke2 fix to ECS for CoreDNS unreachable issue

ECS uses the Canal CNI plugin by default, meaning OPSX-5413 does not resolve issues with Canal.

Users with the default Canal CNI plugin on ECS should switch to the Calico CNI plugin. Refer Flannel Troubleshooting Guide.

Fixed Issues in Cloudera Private Cloud Data Services 1.5.4-CHF1

The fixes in this cumulative hotfix release of CDP Private Cloud Data Services 1.5.4-CHF1.

OPSX-5104: Cluster Ingress Controller improvements for CML workloads scaling

Customers may experience performance degradation when a large number of CML workload sessions are launched simultaneously, which may result in session timeouts. Customer creates a large amount of CML sessions in a short period of time.

The fix does not need user intervention. The fix added cluster ingress improvements for CML workload scaling.

OPSX-5147: OOM when retrieving size of Binary File

Diagnostics bundle collection no longer fails due to OOM errors.

OPSX-5148: Diagnostics Collection from UI w/ Default No Time Limit Should Not Invoke Timestamp Filtering

The diagnostics collection triggered through the UI, with the default "No Time Limit selected", no longer filters logs by timestamp.

Repository Locations for 1.5.4-CHF1

The URLs for Cloudera Private Cloud Data Services 1.5.4-CHF1 are listed in the following table:

URL Type	Repository Location
Index	`https://username:password@archive.cloudera.com/p/cdp-pvc-ds/1.5.4-h2/`
Manifest	Repository: `https://username:password@archive.cloudera.com/p/cdp-pvc-ds/1.5.4-h2/manifest.json`
Parcels	Repository: `https://username:password@archive.cloudera.com/p/cdp-pvc-ds/1.5.4-h2/parcels/`

Fixed Common Vulnerabilities and Exposures in 1.5.4 CHF1

Review the Common vulnerabilities and Exposures (CVEs) that were fixed in 1.5.4 CHF1 release of Cloudera Private Cloud Data Services.

Issue ID	Description
CVE-2004-0230	TCP, when using a large Window Size, makes it easier for remote attackers to guess sequence numbers and cause a persistent TCP connections by repeatedly injecting a TCP RST packet, especially in protocols that use long-lived c
CVE-2005-3660	Linux kernel 2.4 and 2.6 allows attackers to cause a denial of service (memory exhaustion and panic) by creating a socketpairs and setting a large data transfer buffer, then preventing Linux from being able to finish the transfer by c closing the file descriptor without closing an associated reference.
CVE-2007-3719	The process scheduler in the Linux kernel 2.6.16 gives preference to "interactive" processes that perform voluntary denial of service (CPU consumption), as described in "Secretly Monopolizing the CPU Without Superuser Privileg
CVE-2008-2544	Mounting /proc filesystem via chroot command silently mounts it in read-write mode. The user could bypass the ch files, he would never have otherwise.
CVE-2008-4609	The TCP implementation in (1) Linux, (2) platforms based on BSD Unix, (3) Microsoft Windows, (4) Cisco produ allows remote attackers to cause a denial of service (connection queue exhaustion) via multiple vectors that manipu demonstrated by sockstress.
CVE-2009-5155	In the GNU C Library (aka glibc or libc6) before 2.28, parse_reg_exp in posix/regcomp.c misparses alternatives, w service (assertion failure and application exit) or trigger an incorrect result by attempting a regular-expression matc
CVE-2010-4563	The Linux kernel, when using IPv6, allows remote attackers to determine whether a host is sniffing the network by multicast address and determining whether an Echo Reply is sent, as demonstrated by thcping.
CVE-2010-5321	Memory leak in drivers/media/video/videobuf-core.c in the videobuf subsystem in the Linux kernel 2.6.x through 4 of service (memory consumption) by leveraging /dev/video access for a series of mmap calls that require new alloc CVE-2007-6761. NOTE: as of 2016-06-18, this affects only 11 drivers that have not been updated to use videobuf2
CVE-2011-4915	fs/proc/base.c in the Linux kernel through 3.1 allows local users to obtain sensitive keystroke information via acces
CVE-2011-4916	Linux kernel through 3.1 allows local users to obtain sensitive keystroke information via access to /dev/pts/ and /de
CVE-2011-4917	In the Linux kernel through 3.1 there is an information disclosure issue via /proc/stat.
CVE-2012-4542	block/scsi_ioctl.c in the Linux kernel through 3.8 does not properly consider the SCSI device class during authoriz users to bypass intended access restrictions via an SG_IO ioctl call that leverages overlapping opcodes.
CVE-2012-6702	Expat, when used in a parser that has not called XML_SetHashSalt or passed it a seed of 0, makes it easier for cont cryptographic protection mechanisms via vectors involving use of the srand function.
CVE-2012-6711	A heap-based buffer overflow exists in GNU Bash before 4.3 when wide characters, not supported by the current lo variable, are printed through the echo built-in function. A local attacker, who can provide data to print through the to crash a script or execute code with the privileges of the bash process. This occurs because ansicstr() in lib/sh/strt

CVE-2013-0341	Rejected reason: DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: none. Reason: This candidate was w showed that it was not a security issue. Notes: none.
CVE-2013-1664	The XML libraries for Python 3.4, 3.3, 3.2, 3.1, 2.7, and 2.6, as used in OpenStack Keystone Essex, Folsom, and G Cinder Folsom; Django; and possibly other products allow remote attackers to cause a denial of service (resource c Expansion (XEE) attack.
CVE-2013-1665	The XML libraries for Python 3.4, 3.3, 3.2, 3.1, 2.7, and 2.6, as used in OpenStack Keystone Essex and Folsom, Dj remote attackers to read arbitrary files via an XML external entity declaration in conjunction with an entity referen attack.
CVE-2013-7040	Python 2.7 before 3.4 only uses the last eight bits of the prefix to randomize hash values, which causes it to compu to trigger hash collisions predictably and makes it easier for context-dependent attackers to cause a denial of servic application that maintains a hash table.  **Note:** This vulnerability exists because of an incomplete fix for CVE-2012-1150.
CVE-2014-3477	The dbus-daemon in D-Bus 1.2.x through 1.4.x, 1.6.x before 1.6.20, and 1.8.x before 1.8.4, sends an AccessDenied the client is prohibited from accessing the service, which allows local users to cause a denial of service (initializatio channel attack via a D-Bus message to an inactive service.
CVE-2014-3532	dbus 1.3.0 before 1.6.22 and 1.8.x before 1.8.6, when running on Linux 2.6.37-rc4 or later, allows local users to ca of other services or applications) by sending a message containing a file descriptor, then exceeding the maximum r forwarded.
CVE-2014-3533	dbus 1.3.0 before 1.6.22 and 1.8.x before 1.8.6 allows local users to cause a denial of service (disconnect) via a cer the dbus-daemon to forward a message containing an invalid file descriptor.
CVE-2014-3564	Multiple heap-based buffer overflows in the status_handler function in (1) engine-gpgsm.c and (2) engine-uiserver. attackers to cause a denial of service (crash) and possibly execute arbitrary code via vectors related to "different lin
CVE-2014-3566	The SSL protocol 3.0, as used in OpenSSL through 1.0.1i and other products, uses nondeterministic CBC padding, attackers to obtain cleartext data via a padding-oracle attack, aka the "POODLE" issue.
CVE-2014-3591	Libgcrypt before 1.6.3 and GnuPG before 1.4.19 does not implement ciphertext blinding for Elgamal decryption, w obtain the server's private key by determining factors using crafted ciphertext and the fluctuations in the electromag
CVE-2014-3636	D-Bus 1.3.0 through 1.6.x before 1.6.24 and 1.8.x before 1.8.8 allows local users to (1) cause a denial of service (p drop) by queuing the maximum number of file descriptors or (2) cause a denial of service (disconnect) via multiple allowed number of file descriptors for a single sendmsg call.
CVE-2014-3637	D-Bus 1.3.0 through 1.6.x before 1.6.24 and 1.8.x before 1.8.8 does not properly close connections for processes th to cause a denial of service via a D-bus message containing a D-Bus connection file descriptor.
CVE-2014-3638	The bus_connections_check_reply function in config-parser.c in D-Bus before 1.6.24 and 1.8.x before 1.8.8 allows consumption) via a large number of method calls.
CVE-2014-3639	The dbus-daemon in D-Bus before 1.6.24 and 1.8.x before 1.8.8 does not properly close old connections, which all (incomplete connection consumption and prevention of new connections) via a large number of incomplete connec
CVE-2014-4043	The posix_spawn_file_actions_addopen function in glibc before 2.20 does not copy its path argument in accordanc context-dependent attackers to trigger use-after-free vulnerabilities.
CVE-2014-4617	The do_uncompress function in g10/compress.c in GnuPG 1.x before 1.4.17 and 2.x before 2.0.24 allows context-d service (infinite loop) via malformed compressed packets, as demonstrated by an a3 01 5b ff byte sequence.
CVE-2014-5044	Multiple integer overflows in libgfortran might allow remote attackers to execute arbitrary code or cause a denial o vectors related to array allocation.
CVE-2014-5270	Libgcrypt before 1.5.4, as used in GnuPG and other products, does not properly perform ciphertext normalization a easier for physically proximate attackers to conduct key-extraction attacks by leveraging the ability to collect volta than CVE-2013-4576.
CVE-2014-5351	The kadm5_randkey_principal_3 function in lib/kadm5/srv/svr_principal.c in kadmind in MIT Kerberos 5 (aka krb a -randkey -keepold request, which allows remote authenticated users to forge tickets by leveraging administrative
CVE-2014-5461	Buffer overflow in the vararg functions in ldo.c in Lua 5.1 through 5.2.x before 5.2.3 allows context-dependent atta a small number of arguments to a function with a large number of fixed arguments.
CVE-2014-9114	Blkid in util-linux before 2.26rc-1 allows local users to execute arbitrary code.
CVE-2014-9620	The ELF parser in file 5.08 through 5.21 allows remote attackers to cause a denial of service via a large number of

CVE-2014-9892	The snd_compr_tstamp function in sound/core/compress_offload.c in the Linux kernel through 4.7, as used in And (2013) devices, does not properly initialize a timestamp data structure, which allows attackers to obtain sensitive in Android internal bug 28770164 and Qualcomm internal bug CR568717.
CVE-2014-9900	The ethtool_get_wol function in net/core/ethtool.c in the Linux kernel through 4.7, as used in Android before 2016 not initialize a certain data structure, which allows local users to obtain sensitive information via a crafted applicati Qualcomm internal bug CR570754.
CVE-2014-9939	ihex.c in GNU Binutils before 2.26 contains a stack buffer overflow when printing bad bytes in Intel Hex objects.
CVE-2015-0245	D-Bus 1.4.x through 1.6.x before 1.6.30, 1.8.x before 1.8.16, and 1.9.x before 1.9.10 does not validate the source o local users to cause a denial of service (activation failure error returned) by leveraging a race condition involving s systemd responds.
CVE-2015-0247	Heap-based buffer overflow in openfs.c in the libext2fs library in e2fsprogs before 1.42.12 allows local users to ex descriptor data in a filesystem image.
CVE-2015-0837	The mpi_powm function in Libgcrypt before 1.6.3 and GnuPG before 1.4.19 allows attackers to obtain sensitive in when accessing a pre-computed table during modular exponentiation, related to a "Last-Level Cache Side-Channel
CVE-2015-1197	cpio 2.11, when using the --no-absolute-filenames option, allows local users to write to arbitrary files via a symlink
CVE-2015-1572	Heap-based buffer overflow in closefs.c in the libext2fs library in e2fsprogs before 1.42.12 allows local users to ex block group descriptor to be marked as dirty. NOTE: this vulnerability exists because of an incomplete fix for CVE
CVE-2015-1606	The keyring DB in GnuPG before 2.1.2 does not properly handle invalid packets, which allows remote attackers to use-after-free) via a crafted keyring file.
CVE-2015-1607	kbx/keybox-search.c in GnuPG before 1.4.19, 2.0.x before 2.0.27, and 2.1.x before 2.1.2 does not properly handle attackers to cause a denial of service (invalid read operation) via a crafted keyring file, related to sign extensions ar
CVE-2015-2059	The stringprep_utf8_to_ucs4 function in libin before 1.31, as used in jabberd2, allows context-dependent attackers other unspecified impact via invalid UTF-8 characters in a string, which triggers an out-of-bounds read.
CVE-2015-2327	PCRE before 8.36 mishandles the /(((a\2)|(a*)\g<-1>))*/ pattern and related patterns with certain internal recursive attackers to cause a denial of service (segmentation fault) or possibly have unspecified other impact via a crafted re JavaScript RegExp object encountered by Konqueror.
CVE-2015-2328	PCRE before 8.36 mishandles the /((?(R)a|(?1)))+/ pattern and related patterns with certain recursion, which allows (segmentation fault) or possibly have unspecified other impact via a crafted regular expression, as demonstrated by Konqueror.
CVE-2015-2613	Unspecified vulnerability in Oracle Java SE 7u80 and 8u45, and Java SE Embedded 7u75 and 8u33 allows remote related to JCE.
CVE-2015-2695	lib/gssapi/spnego/spnego_mech.c in MIT Kerberos 5 (aka krb5) before 1.14 relies on an inappropriate context hand denial of service (incorrect pointer read and process crash) via a crafted SPNEGO packet that is mishandled during
CVE-2015-2696	lib/gssapi/krb5/iakerb.c in MIT Kerberos 5 (aka krb5) before 1.14 relies on an inappropriate context handle, which service (incorrect pointer read and process crash) via a crafted IAKERB packet that is mishandled during a gss_inq
CVE-2015-2697	The build_principal_va function in lib/krb5/krb/bld_princ.c in MIT Kerberos 5 (aka krb5) before 1.14 allows remo service (out-of-bounds read and KDC crash) via an initial '\0' character in a long realm field within a TGS request.
CVE-2015-2808	The RC4 algorithm, as used in the TLS protocol and SSL protocol, does not properly combine state data with key d makes it easier for remote attackers to conduct plaintext-recovery attacks against the initial bytes of a stream by sni on keys affected by the Invariance Weakness, and then using a brute-force approach involving LSB values, aka the
CVE-2015-2877	Kernel Samepage Merging (KSM) in the Linux kernel 2.6.32 through 4.x does not prevent use of a write-timing sid defeat the ASLR protection mechanism on other guest OS instances via a Cross-VM ASL INtrospection (CAIN) at if you care about this attack vector, disable deduplication." Share-until-written approaches for memory conservatio inherently detectable for information disclosure, and can be classified as potentially misunderstood behaviors rathe
CVE-2015-3153	The default configuration for cURL and libcurl before 7.42.1 sends custom HTTP headers to both the proxy and de proxy servers to obtain sensitive information by reading the header contents.
CVE-2015-3217	PCRE 7.8 and 8.32 through 8.37, and PCRE2 10.10 mishandle group empty matches, which might allow remote at based buffer overflow) via a crafted regular expression, as demonstrated by /^(?:(?(1)\\.|([^\\\\W_])?)+)+$/.
CVE-2015-5073	Heap-based buffer overflow in the find_fixedlength function in pcre_compile.c in PCRE before 8.38 allows remote or obtain sensitive information from heap memory and possibly bypass the ASLR protection mechanism via a craft parenthesis.
CVE-2015-5186	Audit before 2.4.4 in Linux does not sanitize escape characters in filenames.

CVE-2015-5218	Buffer overflow in text-utils/colcrt.c in colcrt in util-linux before 2.27 allows local users to cause a denial of servic global variable.
CVE-2015-5276	The std::random_device class in libstdc++ in the GNU Compiler Collection (aka GCC) before 4.9.4 does not prope which makes it easier for context-dependent attackers to predict the random values via unspecified vectors.
CVE-2015-7036	The fts3_tokenizer function in SQLite, as used in Apple iOS before 8.4 and OS X before 10.10.4, allows remote at denial of service (application crash) via a SQL command that triggers an API call with a crafted pointer value in th
CVE-2015-8382	The match function in pcre_exec.c in PCRE before 8.37 mishandles the /(?:((abcd))|(((?:(?:(?:(?:abc|(?:abcdef))))b) and related patterns involving (*ACCEPT), which allows remote attackers to obtain sensitive information from pro (partially initialized memory and application crash) via a crafted regular expression, as demonstrated by a JavaScri aka ZDI-CAN-2547.
CVE-2015-8386	PCRE before 8.38 mishandles the interaction of lookbehind assertions and mutually recursive subpatterns, which a service (buffer overflow) or possibly have unspecified other impact via a crafted regular expression, as demonstrat by Konqueror.
CVE-2015-8388	PCRE before 8.38 mishandles the /(?=di(?<=(?1))|(?=(.))))/ pattern and related patterns with an unmatched closing cause a denial of service (buffer overflow) or possibly have unspecified other impact via a crafted regular expressio object encountered by Konqueror.
CVE-2015-8391	The pcre_compile function in pcre_compile.c in PCRE before 8.38 mishandles certain [: nesting, which allows rem (CPU consumption) or possibly have unspecified other impact via a crafted regular expression, as demonstrated by Konqueror.
CVE-2015-8538	dwarf_leb.c in libdwarf allows attackers to cause a denial of service (SIGSEGV).
CVE-2015-8865	The file_check_mem function in funcs.c in file before 5.23, as used in the Fileinfo component in PHP before 5.5.34 mishandles continuation-level jumps, which allows context-dependent attackers to cause a denial of service (buffer execute arbitrary code via a crafted magic file.
CVE-2015-8948	idn in GNU libidn before 1.33 might allow remote attackers to obtain sensitive memory information by reading a z bounds read.
CVE-2015-8982	Integer overflow in the strxfrm function in the GNU C Library (aka glibc or libc6) before 2.21 allows context-depe (crash) or possibly execute arbitrary code via a long string, which triggers a stack-based buffer overflow.
CVE-2015-8982	Integer overflow in the strxfrm function in the GNU C Library (aka glibc or libc6) before 2.21 allows context-depe (crash) or possibly execute arbitrary code via a long string, which triggers a stack-based buffer overflow.
CVE-2015-8983	Integer overflow in the _IO_wstr_overflow function in libio/wstrops.c in the GNU C Library (aka glibc or libc6) be to cause a denial of service (application crash) or possibly execute arbitrary code via vectors related to computing a buffer overflow.
CVE-2015-8983	Integer overflow in the _IO_wstr_overflow function in libio/wstrops.c in the GNU C Library (aka glibc or libc6) be to cause a denial of service (application crash) or possibly execute arbitrary code via vectors related to computing a buffer overflow.
CVE-2015-8984	The fnmatch function in the GNU C Library (aka glibc or libc6) before 2.22 might allow context-dependent attacke crash) via a malformed pattern, which triggers an out-of-bounds read.
CVE-2015-8985	The pop_fail_stack function in the GNU C Library (aka glibc or libc6) allows context-dependent attackers to cause application crash) via vectors related to extended regular expression processing.
CVE-2016-0755	The ConnectionExists function in lib/url.c in libcurl before 7.47.0 does not properly re-use NTLM-authenticated pr attackers to authenticate as other users via a request, a similar issue to CVE-2014-0015.
CVE-2016-10228	The iconv program in the GNU C Library (aka glibc or libc6) 2.31 and earlier, when invoked with multiple suffixes or IGNORE) along with the -c option, enters an infinite loop when processing invalid multi-byte input sequences, l
CVE-2016-10254	The allocate_elf function in common.h in elfutils before 0.168 allows remote attackers to cause a denial of service memory allocation failure.
CVE-2016-10255	The __libelf_set_rawdata_wrlock function in elf_getdata.c in elfutils before 0.168 allows remote attackers to cause sh_off or (2) sh_size ELF header value, which triggers a memory allocation failure.
CVE-2016-10255	The __libelf_set_rawdata_wrlock function in elf_getdata.c in elfutils before 0.168 allows remote attackers to cause sh_off or (2) sh_size ELF header value, which triggers a memory allocation failure.
CVE-2016-10255	The __libelf_set_rawdata_wrlock function in elf_getdata.c in elfutils before 0.168 allows remote attackers to cause sh_off or (2) sh_size ELF header value, which triggers a memory allocation failure.
CVE-2016-10505	NULL pointer dereference vulnerabilities in the imagetopnm function in convert.c, sycc444_to_rgb function in col and sycc422_to_rgb function in color.c in OpenJPEG before 2.2.0 allow remote attackers to cause a denial of servi

CVE-2016-10506	Division-by-zero vulnerabilities in the functions opj_pi_next_cprl, opj_pi_next_pcrl, and opj_pi_next_rpcl in pi.c i attackers to cause a denial of service (application crash) via crafted j2k files.
CVE-2016-10723	An issue was discovered in the Linux kernel through 4.17.2. Since the page allocator does not yield CPU resources unprivileged user can trivially lock up the system forever by wasting CPU resources from the page allocator (e.g., v global OOM killer is invoked. NOTE: the software maintainer has not accepted certain proposed patches, in part be problem is non-trivial to handle.
CVE-2016-1234	Stack-based buffer overflow in the glob implementation in GNU C Library (aka glibc) before 2.24, when GLOB_A dependent attackers to cause a denial of service (crash) via a long name.
CVE-2016-1938	The s_mp_div function in lib/freebl/mpi/mpi.c in Mozilla Network Security Services (NSS) before 3.21, as used in divides numbers, which might make it easier for remote attackers to defeat cryptographic protection mechanisms b mp_exptmod function.
CVE-2016-1951	Multiple integer overflows in io/prprf.c in Mozilla Netscape Portable Runtime (NSPR) before 4.12 allow remote at overflow) or possibly have unspecified other impact via a long string to a PR_*printf function.
CVE-2016-2037	The cpio_safer_name_suffix function in util.c in cpio 2.11 allows remote attackers to cause a denial of service (out
CVE-2016-2183	The DES and Triple DES ciphers, as used in the TLS, SSH, and IPSec protocols and other protocols and products, billion blocks, which makes it easier for remote attackers to obtain cleartext data via a birthday attack against a lon by an HTTPS session using Triple DES in CBC mode, aka a "Sweet32" attack.
CVE-2016-2226	Integer overflow in the string_appends function in cplus-dem.c in libiberty allows remote attackers to execute arbit triggers a buffer overflow.
CVE-2016-2779	runuser in util-linux allows local users to escape to the parent session via a crafted TIOCSTI ioctl call, which pushe
CVE-2016-3189	Use-after-free vulnerability in bzip2recover in bzip2 1.0.6 allows remote attackers to cause a denial of service (cras ends set to before the start of the block.
CVE-2016-4008	The _asn1_extract_der_octet function in lib/decoding.c in GNU Libtasn1 before 4.8, when used without the ASN1_ allows remote attackers to cause a denial of service (infinite recursion) via a crafted certificate.
CVE-2016-4429	Stack-based buffer overflow in the clntudp_call function in sunrpc/clnt_udp.c in the GNU C Library (aka glibc or l of service (crash) or possibly unspecified other impact via a flood of crafted ICMP and UDP packets.
CVE-2016-4472	The overflow protection in Expat is removed by compilers with certain optimization settings, which allows remote or possibly execute arbitrary code via crafted XML data. NOTE: this vulnerability exists because of an incomplete
CVE-2016-4483	The xmlBufAttrSerializeTxtContent function in xmlsave.c in libxml2 allows context-dependent attackers to cause a application crash) via a non-UTF-8 attribute value, related to serialization. NOTE: this vulnerability may be a dupli
CVE-2016-4484	The Debian initrd script for the cryptsetup package 2:1.7.3-2 and earlier allows physically proximate attackers to g an invalid password.
CVE-2016-4487	Use-after-free vulnerability in libiberty allows remote attackers to cause a denial of service (segmentation fault and "btypevec."
CVE-2016-4488	Use-after-free vulnerability in libiberty allows remote attackers to cause a denial of service (segmentation fault and "ktypevec."
CVE-2016-4489	Integer overflow in the gnu_special function in libiberty allows remote attackers to cause a denial of service (segm related to the "demangling of virtual tables."
CVE-2016-4490	Integer overflow in cp-demangle.c in libiberty allows remote attackers to cause a denial of service (segmentation fa inconsistent use of the long and int types for lengths.
CVE-2016-4491	The d_print_comp function in cp-demangle.c in libiberty allows remote attackers to cause a denial of service (segm which triggers infinite recursion and a buffer overflow, related to a node having "itself as ancestor more than once."
CVE-2016-4492	Buffer overflow in the do_type function in cplus-dem.c in libiberty allows remote attackers to cause a denial of ser crafted binary.
CVE-2016-4493	The demangle_template_value_parm and do_hpacc_template_literal functions in cplus-dem.c in libiberty allow rem (out-of-bounds read and crash) via a crafted binary.
CVE-2016-4984	/usr/libexec/openldap/generate-server-cert.sh in openldap-servers sets weak permissions for the TLS certificate, wh certificate by leveraging a race condition between the creation of the certificate, and the chmod to protect it.
CVE-2016-5300	The XML parser in Expat does not use sufficient entropy for hash initialization, which allows context-dependent at consumption) via crafted identifiers in an XML document. NOTE: this vulnerability exists because of an incomplet

CVE-2016-6153	os_unix.c in SQLite before 3.13.0 improperly implements the temporary directory search algorithm, which might a information, cause a denial of service (application crash), or have unspecified other impact by leveraging use of the files.
CVE-2016-6261	The idna_to_ascii_4i function in lib/idna.c in libidn before 1.33 allows context-dependent attackers to cause a denia via 64 bytes of input.
CVE-2016-6262	idn in libidn before 1.33 might allow remote attackers to obtain sensitive memory information by reading a zero by read, a different vulnerability than CVE-2015-8948.
CVE-2016-6263	The stringprep_utf8_nfkc_normalize function in lib/nfkc.c in libidn before 1.33 allows context-dependent attackers read and crash) via crafted UTF-8 data.
CVE-2016-6318	Stack-based buffer overflow in the FascistGecosUser function in lib/fascist.c in cracklib allows local users to cause privileges via a long GECOS field, involving longbuffer.
CVE-2016-6321	Directory traversal vulnerability in the safer_name_suffix function in GNU tar 1.14 through 1.29 might allow remo mechanism and write to arbitrary files via vectors related to improper sanitization of the file_name parameter, aka
CVE-2016-6349	The machinectl command in oci-register-machine allows local users to list running containers and possibly obtain s command.
CVE-2016-7091	sudo: It was discovered that the default sudo configuration on Red Hat Enterprise Linux and possibly other Linux i INPUTRC which could lead to information disclosure. A local user with sudo access to a restricted program that us from specially formatted files with elevated privileges provided by sudo.
CVE-2016-8615	A flaw was found in curl before version 7.51. If cookie state is written into a cookie jar file that is later read back a HTTP server can inject new cookies for arbitrary domains into said cookie jar.
CVE-2016-8616	A flaw was found in curl before version 7.51.0 When re-using a connection, curl was doing case insensitive compa existing connections. This means that if an unused connection with proper credentials exists for a protocol that has can cause that connection to be reused if s/he knows the case-insensitive version of the correct password.
CVE-2016-8617	The base64 encode function in curl before version 7.51.0 is prone to a buffer being under allocated in 32bit system `CURLOPT_USERNAME`.
CVE-2016-8618	The libcurl API function called `curl_maprintf()` before version 7.51.0 can be tricked into doing a double-free due systems using 32 bit `size_t` variables.
CVE-2016-8619	The function `read_data()` in security.c in curl before version 7.51.0 is vulnerable to memory double free.
CVE-2016-8621	The `curl_getdate` function in curl before version 7.51.0 is vulnerable to an out of bounds read if it receives an inpu
CVE-2016-8622	The URL percent-encoding decode function in libcurl before 7.51.0 is called `curl_easy_unescape`. Internally, ever unscape destination buffer larger than 2GB, it would return that new length in a signed 32 bit integer variable, thus both truncated and turned negative. That could then lead to libcurl writing outside of its heap based buffer.
CVE-2016-8623	A flaw was found in curl before version 7.51.0. The way curl handles cookies permits other threads to trigger a use
CVE-2016-8624	curl before version 7.51.0 doesn't parse the authority component of the URL correctly when the host name part end tricked into connecting to a different host. This may have security implications if you for example use an URL pars domains before using curl to request them.
CVE-2016-8625	curl before version 7.51.0 uses outdated IDNA 2003 standard to handle International Domain Names and this may issue network transfer requests to the wrong host.
CVE-2016-9063	An integer overflow during the parsing of XML using the Expat library. This vulnerability affects Firefox < 50.
CVE-2016-9074	An existing mitigation of timing side-channel attacks is insufficient in some circumstances. This issue is addressed This vulnerability affects Thunderbird < 45.5, Firefox ESR < 45.5, and Firefox < 50.
CVE-2016-9113	There is a NULL pointer dereference in function imagetobmp of convertbmp.c:980 of OpenJPEG 2.1.2. image->co initialization(NULL). Impact is Denial of Service.
CVE-2016-9114	There is a NULL Pointer Access in function imagetopnm of convert.c:1943(jp2) of OpenJPEG 2.1.2. image->comp initialization(NULL). Impact is Denial of Service.
CVE-2016-9115	Heap Buffer Over-read in function imagetotga of convert.c(jp2):942 in OpenJPEG 2.1.2. Impact is Denial of Servi
CVE-2016-9116	NULL Pointer Access in function imagetopnm of convert.c:2226(jp2) in OpenJPEG 2.1.2. Impact is Denial of Serv
CVE-2016-9117	NULL Pointer Access in function imagetopnm of convert.c(jp2):1289 in OpenJPEG 2.1.2. Impact is Denial of Serv
CVE-2016-9318	libxml2 2.9.4 and earlier, as used in XMLSec 1.2.23 and earlier and other products, does not offer a flag directly in read but other files may not be opened, which makes it easier for remote attackers to conduct XML External Entity
CVE-2016-9574	nss before version 3.30 is vulnerable to a remote denial of service during the session handshake when using Sessio

CVE-2016-9580	An integer overflow vulnerability was found in tiftoimage function in openjpeg 2.1.2, resulting in heap buffer over
CVE-2016-9581	An infinite loop vulnerability in tiftoimage that results in heap buffer overflow in convert_32s_C1P1 was found in
CVE-2016-9586	curl before version 7.52.0 is vulnerable to a buffer overflow when doing a large floating point output in libcurl's im are any application that accepts a format string from the outside without necessary input filtering, it could allow rer
CVE-2017-0630	An information disclosure vulnerability in the kernel trace subsystem could enable a local malicious application to This issue is rated as Moderate because it first requires compromising a privileged process. Product: Android. Vers A-34277115.
CVE-2017-0663	A remote code execution vulnerability in libxml2 could enable an attacker using a specially crafted file to execute a unprivileged process. This issue is rated as High due to the possibility of remote code execution in an application th Versions: 4.4.4, 5.0.2, 5.1.1, 6.0, 6.0.1, 7.0, 7.1.1, 7.1.2. Android ID: A-37104170.
CVE-2017-1000100	When doing a TFTP transfer and curl/libcurl is given a URL that contains a very long file name (longer than about within the buffer boundaries, but the buffer size is still wrongly updated to use the untruncated length. This too larg making curl attempt to send more data than what is actually put into the buffer. The endto() function will then read A malicious HTTP(S) server could redirect a vulnerable libcurl-using client to a crafted TFTP URL (if the client ha redirects to) and trick it to send private memory contents to a remote server over UDP. Limit curl's redirect protoco CURLOPT_REDIR_PROTOCOLS.
CVE-2017-1000158	CPython (aka Python) up to 2.7.13 is vulnerable to an integer overflow in the PyString_DecodeEscape function in overflow (and possible arbitrary code execution)
CVE-2017-1000254	libcurl may read outside of a heap allocated buffer when doing FTP. When libcurl connects to an FTP server and su asks the server for the current directory with the `PWD` command. The server then responds with a 257 response c The returned path name is then kept by libcurl for subsequent uses. Due to a flaw in the string parser for this direct but without a closing double quote would lead to libcurl not adding a trailing NUL byte to the buffer holding the na the string, it could read beyond the allocated heap buffer and crash or wrongly access data beyond the buffer, think server could abuse this fact and effectively prevent libcurl-based clients to work with it - the PWD command is alw mistake has a high chance of causing a segfault. The simple fact that this has issue remained undiscovered for this responses are rare in benign servers. We are not aware of any exploit of this flaw. This bug was introduced in comm commit/415d2e7cb7), March 2005. In libcurl version 7.56.0, the parser always zero terminates the string but also r double quote.
CVE-2017-10140	Postfix before 2.11.10, 3.0.x before 3.0.10, 3.1.x before 3.1.6, and 3.2.x before 3.2.2 might allow local users to gai functionality in Berkeley DB 2.x and later, related to reading settings from DB_CONFIG in the current directory.
CVE-2017-10684	In ncurses 6.0, there is a stack-based buffer overflow in the fmt_entry function. A crafted input will lead to a remot
CVE-2017-10685	In ncurses 6.0, there is a format string vulnerability in the fmt_entry function. A crafted input will lead to a remote
CVE-2017-10790	The _asn1_check_identifier function in GNU Libtasn1 through 4.12 causes a NULL pointer dereference and crash assignment of a NULL value within an asn1_node structure. It may lead to a remote denial of service attack.
CVE-2017-10989	The getNodeSize function in ext/rtree/rtree.c in SQLite through 3.19.3, as used in GDAL and other products, misha database, leading to a heap-based buffer over-read or possibly unspecified other impact.
CVE-2017-11112	In ncurses 6.0, there is an attempted 0xffffffffffffffff access in the append_acs function of tinfo/parse_entry.c. It co the terminfo library code is used to process untrusted terminfo data.
CVE-2017-11113	In ncurses 6.0, there is a NULL Pointer Dereference in the _nc_parse_entry function of tinfo/parse_entry.c. It could the terminfo library code is used to process untrusted terminfo data.
CVE-2017-11462	Double free vulnerability in MIT Kerberos 5 (aka krb5) allows attackers to have unspecified impact via vectors inv on error.
CVE-2017-12449	The _bfd_vms_save_sized_string function in vms-misc.c in the Binary File Descriptor (BFD) library (aka libbfd), a earlier, allows remote attackers to cause an out of bounds heap read via a crafted vms file.
CVE-2017-12451	The _bfd_xcoff_read_ar_hdr function in bfd/coff-rs6000.c and bfd/coff64-rs6000.c in the Binary File Descriptor (B GNU Binutils 2.29 and earlier, allows remote attackers to cause an out of bounds stack read via a crafted COFF im
CVE-2017-12452	The bfd_mach_o_i386_canonicalize_one_reloc function in bfd/mach-o-i386.c in the Binary File Descriptor (BFD) Binutils 2.29 and earlier, allows remote attackers to cause an out of bounds heap read via a crafted mach-o file.
CVE-2017-12453	The _bfd_vms_slurp_eeom function in libbfd.c in the Binary File Descriptor (BFD) library (aka libbfd), as distribu remote attackers to cause an out of bounds heap read via a crafted vms alpha file.
CVE-2017-12454	The _bfd_vms_slurp_egsd function in bfd/vms-alpha.c in the Binary File Descriptor (BFD) library (aka libbfd), as allows remote attackers to cause an arbitrary memory read via a crafted vms alpha file.
CVE-2017-12455	The evax_bfd_print_emh function in vms-alpha.c in the Binary File Descriptor (BFD) library (aka libbfd), as distri allows remote attackers to cause an out of bounds heap read via a crafted vms alpha file.

CVE-2017-12456	The read_symbol_stabs_debugging_info function in rddbg.c in GNU Binutils 2.29 and earlier allows remote attack crafted binary file.
CVE-2017-12457	The bfd_make_section_with_flags function in section.c in the Binary File Descriptor (BFD) library (aka libbfd), as allows remote attackers to cause a NULL dereference via a crafted file.
CVE-2017-12458	The nlm_swap_auxiliary_headers_in function in bfd/nlmcode.h in the Binary File Descriptor (BFD) library (aka lib earlier, allows remote attackers to cause an out of bounds heap read via a crafted nlm file.
CVE-2017-12799	The elf_read_notesfunction in bfd/elf.c in GNU Binutils 2.29 allows remote attackers to cause a denial of service (possibly have unspecified other impact via a crafted binary file.
CVE-2017-12967	The getsym function in tekhex.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Bi denial of service (stack-based buffer over-read and application crash) via a malformed tekhex binary.
CVE-2017-13685	The dump_callback function in SQLite 3.20.0 allows remote attackers to cause a denial of service (EXC_BAD_AC file.
CVE-2017-13694	The acpi_ps_complete_final_op() function in drivers/acpi/acpica/psobject.c in the Linux kernel through 4.12.9 doe causes a kernel stack dump, which allows local users to obtain sensitive information from kernel memory and bypa kernel through 4.9) via a crafted ACPI table.
CVE-2017-13710	The setup_group function in elf.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU B a denial of service (NULL pointer dereference and application crash) via a group section that is too small.
CVE-2017-13728	There is an infinite loop in the next_char function in comp_scan.c in ncurses 6.0, related to libtic. A crafted input w
CVE-2017-13729	There is an illegal address access in the _nc_save_str function in alloc_entry.c in ncurses 6.0. It will lead to a remo
CVE-2017-13730	There is an illegal address access in the function _nc_read_entry_source() in progs/tic.c in ncurses 6.0 that might le
CVE-2017-13731	There is an illegal address access in the function postprocess_termcap() in parse_entry.c in ncurses 6.0 that will lea
CVE-2017-13732	There is an illegal address access in the function dump_uses() in progs/dump_entry.c in ncurses 6.0 that might lead
CVE-2017-13733	There is an illegal address access in the fmt_entry function in progs/dump_entry.c in ncurses 6.0 that might lead to
CVE-2017-13734	There is an illegal address access in the _nc_safe_strcat function in strings.c in ncurses 6.0 that will lead to a remot
CVE-2017-13757	The Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2.29, does not validate the Pl attackers to cause a denial of service (heap-based buffer over-read and application crash) via a crafted ELF file, rel elf32-i386.c and elf_x86_64_get_synthetic_symtab in elf64-x86-64.c.
CVE-2017-14062	Integer overflow in the decode_digit function in puny_decode.c in Libidn2 before 2.0.4 allows remote attackers to unspecified other impact.
CVE-2017-14128	The decode_line_info function in dwarf2.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed i to cause a denial of service (read_1_byte heap-based buffer over-read and application crash) via a crafted ELF file.
CVE-2017-14129	The read_section function in dwarf2.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GN cause a denial of service (parse_comp_unit heap-based buffer over-read and application crash) via a crafted ELF fi
CVE-2017-14130	The _bfd_elf_parse_attributes function in elf-attrs.c in the Binary File Descriptor (BFD) library (aka libbfd), as dis attackers to cause a denial of service (_bfd_elf_attr_strdup heap-based buffer over-read and application crash) via a
CVE-2017-14529	The pe_print_idata function in peXXigen.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed i vector entries, which allows remote attackers to cause a denial of service (heap-based buffer over-read and applicat bfd_getl16 function.
CVE-2017-14729	The *_get_synthetic_symtab functions in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GI entry for a symbol, which allows remote attackers to cause a denial of service (heap-based buffer overflow and app other impact via a crafted ELF file, related to elf32-i386.c and elf64-x86-64.c.
CVE-2017-14745	The *_get_synthetic_symtab functions in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GI a sorting count instead of an error flag, which allows remote attackers to cause a denial of service (integer overflow unspecified other impact via a crafted ELF file, related to elf32-i386.c and elf64-x86-64.c.
CVE-2017-14930	Memory leak in decode_line_info in dwarf2.c in the Binary File Descriptor (BFD) library (aka libbfd), as distribute attackers to cause a denial of service (memory consumption) via a crafted ELF file.
CVE-2017-14932	decode_line_info in dwarf2.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binut denial of service (infinite loop) via a crafted ELF file.
CVE-2017-14933	read_formatted_entries in dwarf2.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU cause a denial of service (infinite loop) via a crafted ELF file.

CVE-2017-14934	process_debug_info in dwarf.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binu denial of service (infinite loop) via a crafted ELF file that contains a negative size value in a CU structure.
CVE-2017-14938	_bfd_elf_slurp_version_tables in elf.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GN cause a denial of service (excessive memory allocation and application crash) via a crafted ELF file.
CVE-2017-14939	decode_line_info in dwarf2.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binut which allows remote attackers to cause a denial of service (heap-based buffer over-read and application crash) via a
CVE-2017-14940	scan_unit_for_symbols in dwarf2.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU cause a denial of service (NULL pointer dereference and application crash) via a crafted ELF file.
CVE-2017-14974	The *_get_synthetic_symtab functions in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GI certain canonicalization step, which allows remote attackers to cause a denial of service (NULL pointer dereferenc related to elf32-i386.c and elf64-x86-64.c.
CVE-2017-15020	dwarf1.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2.29, mishandles to cause a denial of service (application crash) or possibly have unspecified other impact via a crafted ELF file, rela demonstrated by a parse_die heap-based buffer over-read.
CVE-2017-15021	bfd_get_debug_link_info_1 in opncls.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in G cause a denial of service (heap-based buffer over-read and application crash) via a crafted ELF file, related to bfd_g
CVE-2017-15022	dwarf2.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2.29, does not val allows remote attackers to cause a denial of service (bfd_hash_hash NULL pointer dereference, or out-of-bounds a ELF file, related to scan_unit_for_symbols and parse_comp_unit.
CVE-2017-15023	read_formatted_entries in dwarf2.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU format count, which allows remote attackers to cause a denial of service (NULL pointer dereference and applicatio concat_filename.
CVE-2017-15024	find_abstract_instance_name in dwarf2.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in to cause a denial of service (infinite recursion and application crash) via a crafted ELF file.
CVE-2017-15025	decode_line_info in dwarf2.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binut denial of service (divide-by-zero error and application crash) via a crafted ELF file.
CVE-2017-15088	plugins/preauth/pkinit/pkinit_crypto_openssl.c in MIT Kerberos 5 (aka krb5) through 1.15.2 mishandles Distinguis attackers to execute arbitrary code or cause a denial of service (buffer overflow and application crash) in situations get_matching_data and X509_NAME_oneline_ex functions. NOTE: this has security relevance only in use cases o the use of get_matching_data in KDC certauth plugin code that is specific to Red Hat.
CVE-2017-15225	_bfd_dwarf2_cleanup_debug_info in dwarf2.c in the Binary File Descriptor (BFD) library (aka libbfd), as distribut attackers to cause a denial of service (memory leak) via a crafted ELF file.
CVE-2017-15286	SQLite 3.20.1 has a NULL pointer dereference in tableColumnList in shell.c because it fails to consider certain cas `sqlite3_step(pStmt)==SQLITE_ROW` is false and a data structure is never initialized.
CVE-2017-15671	The glob function in glob.c in the GNU C Library (aka glibc or libc6) before 2.27, when invoked with GLOB_TIL when processing the ~ operator with a long user name, potentially leading to a denial of service (memory leak).
CVE-2017-15938	dwarf2.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2.29, miscalculate of a relocatable object file, which allows remote attackers to cause a denial of service (find_abstract_instance_nam and application crash).
CVE-2017-15939	dwarf2.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2.29, mishandles allows remote attackers to cause a denial of service (NULL pointer dereference and application crash) via a crafted **Note:** This issue is caused by an incomplete fix for CVE-2017-15023.
CVE-2017-15996	elfcomm.c in readelf in GNU Binutils 2.29 allows remote attackers to cause a denial of service (excessive memory impact via a crafted ELF file that triggers a "buffer overflow on fuzzed archive header," related to an uninitialized v the get_archive_member_name, process_archive_index_and_symbols, and setup_archive functions.
CVE-2017-16826	The coff_slurp_line_table function in coffcode.h in the Binary File Descriptor (BFD) library (aka libbfd), as distrib attackers to cause a denial of service (invalid memory access and application crash) or possibly have unspecified ot
CVE-2017-16827	The aout_get_external_symbols function in aoutx.h in the Binary File Descriptor (BFD) library (aka libbfd), as dist remote attackers to cause a denial of service (slurp_symtab invalid free and application crash) or possibly have uns
CVE-2017-16828	The display_debug_frames function in dwarf.c in GNU Binutils 2.29.1 allows remote attackers to cause a denial of buffer over-read, and application crash) or possibly have unspecified other impact via a crafted ELF file, related to

CVE-2017-16829	The _bfd_elf_parse_gnu_properties function in elf-properties.c in the Binary File Descriptor (BFD) library (aka lib does not prevent negative pointers, which allows remote attackers to cause a denial of service (out-of-bounds read unspecified other impact via a crafted ELF file.
CVE-2017-16830	The print_gnu_property_note function in readelf.c in GNU Binutils 2.29.1 does not have integer-overflow protecti attackers to cause a denial of service (segmentation violation and application crash) or possibly have unspecified ot
CVE-2017-16831	coffgen.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2.29.1, does not remote attackers to cause a denial of service (integer overflow and application crash, or excessive memory allocati via a crafted PE file.
CVE-2017-16832	The pe_bfd_read_buildid function in peicode.h in the Binary File Descriptor (BFD) library (aka libbfd), as distribu size and offset values in the data dictionary, which allows remote attackers to cause a denial of service (segmentati have unspecified other impact via a crafted PE file.
CVE-2017-16879	Stack-based buffer overflow in the _nc_write_entry function in tinfo/write_entry.c in ncurses 6.0 allows attackers t or possibly execute arbitrary code via a crafted terminfo file, as demonstrated by tic.
CVE-2017-16931	parser.c in libxml2 before 2.9.5 mishandles parameter-entity references because the NEXTL macro calls the xmlPa of a '%' character in a DTD name.
CVE-2017-16932	parser.c in libxml2 before 2.9.5 does not prevent infinite recursion in parameter entities.
CVE-2017-17080	elf.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2.29.1, does not valida allows remote attackers to cause a denial of service (bfd_getl32 heap-based buffer over-read and application crash) elfcore_grok_netbsd_procinfo, elfcore_grok_openbsd_procinfo, and elfcore_grok_nto_status.
CVE-2017-17121	The Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2.29.1, allows remote attacke access violation) or possibly have unspecified other impact via a COFF binary in which a relocation refers to a loca section.
CVE-2017-17122	The dump_relocs_in_section function in objdump.c in GNU Binutils 2.29.1 does not check for reloc count integer cause a denial of service (excessive memory allocation, or heap-based buffer overflow and application crash) or po crafted PE file.
CVE-2017-17123	The coff_slurp_reloc_table function in coffcode.h in the Binary File Descriptor (BFD) library (aka libbfd), as distri attackers to cause a denial of service (NULL pointer dereference and application crash) via a crafted COFF based f
CVE-2017-17124	The _bfd_coff_read_string_table function in coffgen.c in the Binary File Descriptor (BFD) library (aka libbfd), as properly validate the size of the external string table, which allows remote attackers to cause a denial of service (ex buffer overflow and application crash) or possibly have unspecified other impact via a crafted COFF binary.
CVE-2017-17125	nm.c and objdump.c in GNU Binutils 2.29.1 mishandle certain global symbols, which allows remote attackers to ca (_bfd_elf_get_symbol_version_string buffer over-read and application crash) or possibly have unspecified other im
CVE-2017-17126	The load_debug_section function in readelf.c in GNU Binutils 2.29.1 allows remote attackers to cause a denial of s application crash) or possibly have unspecified other impact via an ELF file that lacks section headers.
CVE-2017-17479	In OpenJPEG 2.3.0, a stack-based buffer overflow was discovered in the pgxtoimage function in jpwl/convert.c. Th write, which may lead to remote denial of service or possibly remote code execution.
CVE-2017-18078	systemd-tmpfiles in systemd before 237 attempts to support ownership/permission changes on hardlinked files ever turned off, which allows local users to bypass intended access restrictions via vectors involving a hard link to a file demonstrated by changing the ownership of the /etc/passwd file.
CVE-2017-18640	The Alias feature in SnakeYAML before 1.26 allows entity expansion during a load operation, a related issue to CV
CVE-2017-5969	libxml2 2.9.4, when used in recover mode, allows remote attackers to cause a denial of service (NULL pointer dere NOTE: The maintainer states "I would disagree of a CVE with the Recover parsing option which should only be us parser.
CVE-2017-6004	The compile_bracket_matchingpath function in pcre_jit_compile.c in PCRE through 8.x before revision 1680 (e.g. remote attackers to cause a denial of service (out-of-bounds read and application crash) via a crafted regular expres
CVE-2017-6891	Two errors in the "asn1_find_node()" function (lib/parser_aux.c) within GnuTLS libtasn1 version 4.10 can be expl overflow by tricking a user into processing a specially crafted assignments file via the e.g. asn1Coding utility.
CVE-2017-6965	readelf in GNU Binutils 2.28 writes to illegal addresses while processing corrupt input files containing symbol-diff buffer overflow.
CVE-2017-6966	readelf in GNU Binutils 2.28 has a use-after-free (specifically read-after-free) error while processing multiple, relo caused by mishandling of an invalid symbol index, and mishandling of state across invocations.
CVE-2017-6969	readelf in GNU Binutils 2.28 is vulnerable to a heap-based buffer over-read while processing corrupt RL78 binarie crashes. It may lead to an information leak as well.

CVE-2017-7000	An issue was discovered in certain Apple products. iOS before 10.3.2 is affected. macOS before 10.12.5 is affected It allows remote attackers to execute arbitrary code or cause a denial of service (memory corruption and applicatio
CVE-2017-7186	libpcre1 in PCRE 8.40 and libpcre2 in PCRE2 10.23 allow remote attackers to cause a denial of service (segmentat crash) by triggering an invalid Unicode property lookup.
CVE-2017-7209	The dump_section_as_bytes function in readelf in GNU Binutils 2.28 accesses a NULL pointer while reading secti program crash.
CVE-2017-7210	objdump in GNU Binutils 2.28 is vulnerable to multiple heap-based buffer over-reads (of size 1 and size 8) while h a crafted object file, leading to program crash.
CVE-2017-7223	GNU assembler in GNU Binutils 2.28 is vulnerable to a global buffer overflow (of size 1) while attempting to unge potentially leading to a program crash.
CVE-2017-7224	The find_nearest_line function in objdump in GNU Binutils 2.28 is vulnerable to an invalid write (of size 1) while an empty function name, leading to a program crash.
CVE-2017-7225	The find_nearest_line function in addr2line in GNU Binutils 2.28 does not handle the case where the main file nam triggering a NULL pointer dereference and an invalid write, and leading to a program crash.
CVE-2017-7226	The pe_ILF_object_p function in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binu over-read of size 4049 because it uses the strlen function instead of strnlen, leading to program crashes in several u could lead to information disclosure as well.
CVE-2017-7227	GNU linker (ld) in GNU Binutils 2.28 is vulnerable to a heap-based buffer overflow while processing a bogus inpu relates to lack of '\0' termination of a name field in ldlex.l.
CVE-2017-7244	The _pcre32_xclass function in pcre_xclass.c in libpcre1 in PCRE 8.40 allows remote attackers to cause a denial o file.
CVE-2017-7299	The Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2.28, has an invalid read (of (bfd_elf_final_link function in bfd/elflink.c) does not check the format of the input file before trying to read the EL leads to a GNU linker (ld) program crash.
CVE-2017-7300	The Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2.28, has an aout_link_add_s vulnerable to a heap-based buffer over-read (off-by-one) because of an incomplete check for invalid string offsets v linker (ld) program crash.
CVE-2017-7301	The Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2.28, has an aout_link_add_s off-by-one vulnerability because it does not carefully check the string offset. The vulnerability could lead to a GNU
CVE-2017-7302	The Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2.28, has a swap_std_reloc_ to an invalid read (of size 4) because of missing checks for relocs that could not be recognised. This vulnerability c
CVE-2017-7303	The Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2.28, is vulnerable to an inva check (in the find_link function) for null headers before attempting to match them. This vulnerability causes Binuti
CVE-2017-7304	The Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2.28, is vulnerable to an inva check (in the copy_special_section_fields function) for an invalid sh_link field before attempting to follow it. This to crash.
CVE-2017-7375	A flaw in libxml2 allows remote XML entity inclusion with default parser flags (i.e., when the caller did not reques DTD subset loading, or default DTD attributes). Depending on the context, this may expose a higher-risk attack su default parser flags, and expose content from local files, HTTP, or FTP servers (which might be otherwise unreach
CVE-2017-7407	The ourWriteOut function in tool_writeout.c in curl 7.53.1 might allow physically proximate attackers to obtain se opportunistic circumstances by reading a workstation screen during use of a --write-out argument ending in a '%' cl over-read.
CVE-2017-7500	It was found that rpm did not properly handle RPM installations when a destination path was a symbolic link to a d and permissions of an arbitrary directory, and RPM files being placed in an arbitrary destination. An attacker, with subdirectory will be installed, could redirect that directory to an arbitrary location and gain root privilege.
CVE-2017-7501	It was found that versions of rpm before 4.13.0.2 use temporary files with predictable names when installing an RP directory where files will be installed could create symbolic links to an arbitrary location and modify content, and p could be used for denial of service or possibly privilege escalation.
CVE-2017-7526	libgcrypt before version 1.7.8 is vulnerable to a cache side-channel attack resulting into a complete break of RSA-1 computing the sliding-window expansion. The same attack is believed to work on RSA-2048 with moderately mor attacker can run arbitrary software on the hardware where the private RSA key is used.
CVE-2017-7607	The handle_gnu_hash function in readelf.c in elfutils 0.168 allows remote attackers to cause a denial of service (he crash) via a crafted ELF file.

CVE-2017-7608	The ebl_object_note_type_name function in eblobjnotetypename.c in elfutils 0.168 allows remote attackers to caus read and application crash) via a crafted ELF file.
CVE-2017-7609	elf_compress.c in elfutils 0.168 does not validate the zlib compression factor, which allows remote attackers to cau via a crafted ELF file.
CVE-2017-7610	The check_group function in elflint.c in elfutils 0.168 allows remote attackers to cause a denial of service (heap-ba via a crafted ELF file.
CVE-2017-7611	The check_symtab_shndx function in elflint.c in elfutils 0.168 allows remote attackers to cause a denial of service crash) via a crafted ELF file.
CVE-2017-7612	The check_sysv_hash function in elflint.c in elfutils 0.168 allows remote attackers to cause a denial of service (hea crash) via a crafted ELF file.
CVE-2017-7613	elflint.c in elfutils 0.168 does not validate the number of sections and the number of segments, which allows remot (memory consumption) via a crafted ELF file.
CVE-2017-7614	elflink.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2.28, has a "memb behavior issue, which might allow remote attackers to cause a denial of service (application crash) or possibly have {return 0;}" program.
CVE-2017-7781	An error occurs in the elliptic curve point addition algorithm that uses mixed Jacobian-affine coordinates where it c when it should not. A man-in-the-middle attacker could use this to interfere with a connection, resulting in an attac secret. This vulnerability affects Firefox < 55.
CVE-2017-7781	An error occurs in the elliptic curve point addition algorithm that uses mixed Jacobian-affine coordinates where it c when it should not. A man-in-the-middle attacker could use this to interfere with a connection, resulting in an attac secret. This vulnerability affects Firefox < 55.
CVE-2017-8392	The Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2.28, is vulnerable to an inva to determine whether symbols are NULL in the _bfd_dwarf2_find_nearest_line function. This vulnerability causes programs using the libbfd library, such as objdump, to crash.
CVE-2017-8393	The Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2.28, is vulnerable to a globa assumption made by code that runs for objcopy and strip, that SHT_REL/SHR_RELA sections are always named s vulnerability causes programs that conduct an analysis of binary programs using the libbfd library, such as objcopy
CVE-2017-8394	The Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2.28, is vulnerable to an inva dereferencing of _bfd_elf_large_com_section. This vulnerability causes programs that conduct an analysis of binar objcopy, to crash.
CVE-2017-8395	The Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2.28, is vulnerable to an inva malloc() return-value check to see if memory had actually been allocated in the _bfd_generic_get_section_contents that conduct an analysis of binary programs using the libbfd library, such as objcopy, to crash.
CVE-2017-8396	The Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2.28, is vulnerable to an inva offset range tests didn't catch small negative offsets less than the size of the reloc field. This vulnerability causes pr programs using the libbfd library, such as objdump, to crash.
CVE-2017-8397	The Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2.28, is vulnerable to an inva size 1 during processing of a corrupt binary containing reloc(s) with negative addresses. This vulnerability causes p programs using the libbfd library, such as objdump, to crash.
CVE-2017-8398	dwarf.c in GNU Binutils 2.28 is vulnerable to an invalid read of size 1 during dumping of debug information from programs that conduct an analysis of binary programs, such as objdump and readelf, to crash.
CVE-2017-8421	The function coff_set_alignment_hook in coffcode.h in Binary File Descriptor (BFD) library (aka libbfd), as distrib leak vulnerability which can cause memory exhaustion in objdump via a crafted PE file. Additional validation in du resolve this.
CVE-2017-8804	The xdr_bytes and xdr_string functions in the GNU C Library (aka glibc or libc6) 2.25 mishandle failures of buffer attackers to cause a denial of service (virtual memory allocation, or memory consumption if an overcommit setting 111, a related issue to CVE-2017-8779. NOTE: [Information provided from upstream and references
CVE-2017-8817	The FTP wildcard function in curl and libcurl before 7.57.0 allows remote attackers to cause a denial of service (ou possibly have unspecified other impact via a string that ends with an '[' character.
CVE-2017-8872	The htmlParseTryOrFinish function in HTMLparser.c in libxml2 2.9.4 allows attackers to cause a denial of service
CVE-2017-9038	GNU Binutils 2.28 allows remote attackers to cause a denial of service (heap-based buffer over-read and applicatio byte_get_little_endian function in elfcomm.c, the get_unwind_section_word function in readelf.c, and ARM unwin offsets.

CVE-2017-9039	GNU Binutils 2.28 allows remote attackers to cause a denial of service (memory consumption) via a crafted ELF fi get_program_headers function in readelf.c.
CVE-2017-9040	GNU Binutils 2017-04-03 allows remote attackers to cause a denial of service (NULL pointer dereference and appl process_mips_specific function in readelf.c, via a crafted ELF file that triggers a large memory-allocation attempt.
CVE-2017-9041	GNU Binutils 2.28 allows remote attackers to cause a denial of service (heap-based buffer over-read and applicatio MIPS GOT mishandling in the process_mips_specific function in readelf.c.
CVE-2017-9042	readelf.c in GNU Binutils 2017-04-12 has a "cannot be represented in type long" issue, which might allow remote a (application crash) or possibly have unspecified other impact via a crafted ELF file.
CVE-2017-9043	readelf.c in GNU Binutils 2017-04-12 has a "shift exponent too large for type unsigned long" issue, which might al service (application crash) or possibly have unspecified other impact via a crafted ELF file.
CVE-2017-9044	The print_symbol_for_build_attribute function in readelf.c in GNU Binutils 2017-04-12 allows remote attackers to SEGV) via a crafted ELF file.
CVE-2017-9047	A buffer overflow was discovered in libxml2 20904-GITv2.9.4-16-g0741801. The function xmlSnprintfElementCo recursively dump the element content definition into a char buffer 'buf' of size 'size'. The variable len is assigned st XML_ELEMENT_CONTENT_ELEMENT, then (i) the content->prefix is appended to buf (if it actually fits) whe buffer. However, the check for whether the content->name actually fits also uses 'len' rather than the updated buffe about "size" many bytes beyond the allocated memory. This vulnerability causes programs that use libxml2, such a
CVE-2017-9048	libxml2 20904-GITv2.9.4-16-g0741801 is vulnerable to a stack-based buffer overflow. The function xmlSnprintfE recursively dump the element content definition into a char buffer 'buf' of size 'size'. At the end of the routine, the f without checking whether the current strlen(buf) + 2 < size. This vulnerability causes programs that use libxml2, su
CVE-2017-9049	libxml2 20904-GITv2.9.4-16-g0741801 is vulnerable to a heap-based buffer over-read in the xmlDictComputeFast causes programs that use libxml2, such as PHP, to crash. This vulnerability exists because of an incomplete fix for
CVE-2017-9050	libxml2 20904-GITv2.9.4-16-g0741801 is vulnerable to a heap-based buffer over-read in the xmlDictAddString fu programs that use libxml2, such as PHP, to crash. This vulnerability exists because of an incomplete fix for CVE-2
CVE-2017-9233	XML External Entity vulnerability in libexpat 2.2.0 and earlier (Expat XML Parser Library) allows attackers to put malformed external entity definition from an external DTD.
CVE-2017-9742	The score_opcodes function in opcodes/score7-dis.c in GNU Binutils 2.28 allows remote attackers to cause a denia crash) or possibly have unspecified other impact via a crafted binary file, as demonstrated by mishandling of this fi
CVE-2017-9743	The print_insn_score32 function in opcodes/score7-dis.c:552 in GNU Binutils 2.28 allows remote attackers to caus and application crash) or possibly have unspecified other impact via a crafted binary file, as demonstrated by misha execution.
CVE-2017-9744	The sh_elf_set_mach_from_flags function in bfd/elf32-sh.c in the Binary File Descriptor (BFD) library (aka libbfd remote attackers to cause a denial of service (buffer overflow and application crash) or possibly have unspecified o demonstrated by mishandling of this file during "objdump -D" execution.
CVE-2017-9745	The _bfd_vms_slurp_etir function in bfd/vms-alpha.c in the Binary File Descriptor (BFD) library (aka libbfd), as d remote attackers to cause a denial of service (buffer overflow and application crash) or possibly have unspecified o demonstrated by mishandling of this file during "objdump -D" execution.
CVE-2017-9746	The disassemble_bytes function in objdump.c in GNU Binutils 2.28 allows remote attackers to cause a denial of se or possibly have unspecified other impact via a crafted binary file, as demonstrated by mishandling of rae insns pri execution.
CVE-2017-9747	The ieee_archive_p function in bfd/ieee.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in attackers to cause a denial of service (buffer overflow and application crash) or possibly have unspecified other imp by mishandling of this file during "objdump -D" execution. NOTE: this may be related to a compiler bug.
CVE-2017-9748	The ieee_object_p function in bfd/ieee.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in attackers to cause a denial of service (buffer overflow and application crash) or possibly have unspecified other imp by mishandling of this file during "objdump -D" execution. NOTE: this may be related to a compiler bug.
CVE-2017-9749	The *regs* macros in opcodes/bfin-dis.c in GNU Binutils 2.28 allow remote attackers to cause a denial of service (possibly have unspecified other impact via a crafted binary file, as demonstrated by mishandling of this file during
CVE-2017-9750	opcodes/rx-decode.opc in GNU Binutils 2.28 lacks bounds checks for certain scale arrays, which allows remote att overflow and application crash) or possibly have unspecified other impact via a crafted binary file, as demonstrated -D" execution.
CVE-2017-9751	opcodes/rl78-decode.opc in GNU Binutils 2.28 has an unbounded GETBYTE macro, which allows remote attacker and application crash) or possibly have unspecified other impact via a crafted binary file, as demonstrated by misha execution.

CVE-2017-9752	bfd/vms-alpha.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2.28, allov service (buffer overflow and application crash) or possibly have unspecified other impact via a crafted binary file, a the _bfd_vms_get_value and _bfd_vms_slurp_etir functions during "objdump -D" execution.
CVE-2017-9753	The versados_mkobject function in bfd/versados.c in the Binary File Descriptor (BFD) library (aka libbfd), as distr initialize a certain data structure, which allows remote attackers to cause a denial of service (buffer overflow and ap other impact via a crafted binary file, as demonstrated by mishandling of this file during "objdump -D" execution.
CVE-2017-9754	The process_otr function in bfd/versados.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed i certain offset, which allows remote attackers to cause a denial of service (buffer overflow and application crash) or crafted binary file, as demonstrated by mishandling of this file during "objdump -D" execution.
CVE-2017-9755	opcodes/i386-dis.c in GNU Binutils 2.28 does not consider the number of registers for bnd mode, which allows ren (buffer overflow and application crash) or possibly have unspecified other impact via a crafted binary file, as demo "objdump -D" execution.
CVE-2017-9756	The aarch64_ext_ldst_reglist function in opcodes/aarch64-dis.c in GNU Binutils 2.28 allows remote attackers to ca and application crash) or possibly have unspecified other impact via a crafted binary file, as demonstrated by misha execution.
CVE-2017-9954	The getvalue function in tekhex.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU B a denial of service (stack-based buffer over-read and application crash) via a crafted tekhex file, as demonstrated by
CVE-2017-9955	The get_build_id function in opncls.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GN cause a denial of service (heap-based buffer over-read and application crash) via a crafted file in which a certain siz field, as demonstrated by mishandling within the objdump program.
CVE-2018-1000030	Python 2.7.14 is vulnerable to a Heap-Buffer-Overflow as well as a Heap-Use-After-Free. Python versions prior to that Python 2.7.17 and prior may also be vulnerable however this has not been confirmed. The vulnerability lies wh amounts of data. In both cases there is essentially a race condition that occurs. For the Heap-Buffer-Overflow, Thre Thread1 is already writing to the buffer without knowing how much to write. So when a large amount of data is bei corruption using a Heap-Buffer-Overflow. As for the Use-After-Free, Thread3->Malloc->Thread1->Free's->Thread stated that this is not a security vulnerability due to the fact that the attacker must be able to run code, however in s this vulnerability can potentially be used by an attacker to violate a trust boundary, as such the DWF feels this issue
CVE-2018-1058	A flaw was found in the way Postgresql allowed a user to modify the behavior of a query for other users. An attack execute code with the permissions of superuser in the database. Versions 9.3 through 10 are affected.
CVE-2018-10754	Rejected reason: DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: none. Reason: This candidate was w showed that it was not a security issue.  **Note:** None.
CVE-2018-11212	An issue was discovered in libjpeg 9a and 9d. The alloc_sarray function in jmemmgr.c allows remote attackers to c via a crafted file.
CVE-2018-1123	procps-ng before version 3.3.15 is vulnerable to a denial of service in ps via mmap buffer overflow. Inbuilt protecti overflowed buffer, ensuring that the impact of this flaw is limited to a crash (temporary denial of service).
CVE-2018-1125	procps-ng before version 3.3.15 is vulnerable to a stack buffer overflow in pgrep. This vulnerability is mitigated by allocated string. When pgrep is compiled with FORTIFY (as on Red Hat Enterprise Linux and Fedora), the impact
CVE-2018-13785	In libpng 1.6.34, a wrong calculation of row_factor in the png_check_chunk_length function (pngrutil.c) may trigg by-zero while processing a crafted PNG file, leading to a denial of service.
CVE-2018-16375	An issue was discovered in OpenJPEG 2.3.0. Missing checks for header_info.height and header_info.width in the f can lead to a heap-based buffer overflow.
CVE-2018-16429	GNOME GLib 2.56.1 has an out-of-bounds read vulnerability in g_markup_parse_context_parse() in gmarkup.c, re
CVE-2018-18508	In Network Security Services (NSS) before 3.36.7 and before 3.41.1, a malformed signature can cause a crash due Service.
CVE-2018-18508	In Network Security Services (NSS) before 3.36.7 and before 3.41.1, a malformed signature can cause a crash due Service.
CVE-2018-20482	GNU Tar through 1.30, when --sparse is used, mishandles file shrinkage during read access, which allows local use loop in sparse_dump_region in sparse.c) by modifying a file that is supposed to be archived by a different user's pr
CVE-2018-25032	zlib before 1.2.12 allows memory corruption when deflating (i.e., when compressing) if the input has many distant

CVE-2018-2938	Vulnerability in the Java SE component of Oracle Java SE (subcomponent: Java DB). Supported versions that are a 8u172. Difficult to exploit vulnerability allows unauthenticated attacker with network access via multiple protocols vulnerability is in Java SE, attacks may significantly impact additional products. Successful attacks of this vulnerab This vulnerability can only be exploited by supplying data to APIs in the specified Component without using Untru Java applets, such as through a web service. CVE-2018-2938 addresses CVE-2018-1313. CVSS 3.0 Base Score 9.0 impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:C/C:H/I:H/A:H).
CVE-2018-2940	Vulnerability in the Java SE, Java SE Embedded component of Oracle Java SE (subcomponent: Libraries). Suppor 6u191, 7u181, 8u172 and 10.0.1; Java SE Embedded: 8u171. Easily exploitable vulnerability allows unauthenticate protocols to compromise Java SE, Java SE Embedded. Successful attacks require human interaction from a person of this vulnerability can result in unauthorized read access to a subset of Java SE, Java SE Embedded accessible da deployments, typically in clients running sandboxed Java Web Start applications or sandboxed Java applets, that lo comes from the internet) and rely on the Java sandbox for security. This vulnerability does not apply to Java deploy only trusted code (e.g., code installed by an administrator). CVSS 3.0 Base Score 4.3 (Confidentiality impacts). CV UI:R/S:U/C:L/I:N/A:N).
CVE-2018-2952	Vulnerability in the Java SE, Java SE Embedded, JRockit component of Oracle Java SE (subcomponent: Concurrer Java SE: 6u191, 7u181, 8u172 and 10.0.1; Java SE Embedded: 8u171; JRockit: R28.3.18. Difficult to exploit vulne network access via multiple protocols to compromise Java SE, Java SE Embedded, JRockit. Successful attacks of t ability to cause a partial denial of service (partial DOS) of Java SE, Java SE Embedded, JRockit. Note: Applies to c vulnerability can be exploited through sandboxed Java Web Start applications and sandboxed Java applets. It can a the specified Component without using sandboxed Java Web Start applications or sandboxed Java applets, such as 3.7 (Availability impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:U/C:N/I:N/A:L).
CVE-2018-2973	Vulnerability in the Java SE, Java SE Embedded component of Oracle Java SE (subcomponent: JSSE). Supported v 7u181, 8u172 and 10.0.1; Java SE Embedded: 8u171. Difficult to exploit vulnerability allows unauthenticated attac compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized creatio data or all Java SE, Java SE Embedded accessible data. Note: This vulnerability applies to Java deployments, typic Start applications or sandboxed Java applets, that load and run untrusted code (e.g., code that comes from the intern This vulnerability does not apply to Java deployments, typically in servers, that load and run only trusted code (e.g 3.0 Base Score 5.9 (Integrity impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:U/C:N/I:H/A:N).
CVE-2018-3136	Vulnerability in the Java SE, Java SE Embedded component of Oracle Java SE (subcomponent: Security). Supporte 6u201, 7u191, 8u182 and 11; Java SE Embedded: 8u181. Difficult to exploit vulnerability allows unauthenticated a protocols to compromise Java SE, Java SE Embedded. Successful attacks require human interaction from a person vulnerability is in Java SE, Java SE Embedded, attacks may significantly impact additional products. Successful at unauthorized update, insert or delete access to some of Java SE, Java SE Embedded accessible data. Note: This vul typically in clients running sandboxed Java Web Start applications or sandboxed Java applets (in Java SE 8), that lo comes from the internet) and rely on the Java sandbox for security. This vulnerability does not apply to Java deploy only trusted code (e.g. code installed by an administrator). CVSS 3.0 Base Score 3.4 (Integrity impacts). CVSS Ve S:C/C:N/I:L/A:N).
CVE-2018-3139	Vulnerability in the Java SE, Java SE Embedded component of Oracle Java SE (subcomponent: Networking). Supp 6u201, 7u191, 8u182 and 11; Java SE Embedded: 8u181. Difficult to exploit vulnerability allows unauthenticated a protocols to compromise Java SE, Java SE Embedded. Successful attacks require human interaction from a person of this vulnerability can result in unauthorized read access to a subset of Java SE, Java SE Embedded accessible da deployments, typically in clients running sandboxed Java Web Start applications or sandboxed Java applets (in Jav code that comes from the internet) and rely on the Java sandbox for security. This vulnerability does not apply to Ja load and run only trusted code (e.g. code installed by an administrator). CVSS 3.0 Base Score 3.1 (Confidentiality AC:H/PR:N/UI:R/S:U/C:L/I:N/A:N).
CVE-2018-3149	Vulnerability in the Java SE, Java SE Embedded, JRockit component of Oracle Java SE (subcomponent: JNDI). Su SE: 6u201, 7u191, 8u182 and 11; Java SE Embedded: 8u181; JRockit: R28.3.19. Difficult to exploit vulnerability a access via multiple protocols to compromise Java SE, Java SE Embedded, JRockit. Successful attacks require hum attacker and while the vulnerability is in Java SE, Java SE Embedded, JRockit, attacks may significantly impact ad vulnerability can result in takeover of Java SE, Java SE Embedded, JRockit. Note: This vulnerability applies to Jav sandboxed Java Web Start applications or sandboxed Java applets (in Java SE 8), that load and run untrusted code rely on the Java sandbox for security. This vulnerability can also be exploited by using APIs in the specified Comp supplies data to the APIs. CVSS 3.0 Base Score 8.3 (Confidentiality, Integrity and Availability impacts). CVSS Ve S:C/C:H/I:H/A:H).
CVE-2018-3169	Vulnerability in the Java SE, Java SE Embedded component of Oracle Java SE (subcomponent: Hotspot). Supporte 7u191, 8u182 and 11; Java SE Embedded: 8u181. Difficult to exploit vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. Successful attacks require human interaction from a person other than in Java SE, Java SE Embedded, attacks may significantly impact additional products. Successful attacks of this vul Java SE Embedded. Note: This vulnerability applies to Java deployments, typically in clients running sandboxed Ja applets (in Java SE 8), that load and run untrusted code (e.g. code that comes from the internet) and rely on the Jav does not apply to Java deployments, typically in servers, that load and run only trusted code (e.g. code installed by (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:R/S:C/C:H/I

CVE-2018-3180	Vulnerability in the Java SE, Java SE Embedded, JRockit component of Oracle Java SE (subcomponent: JSSE). Su Java SE: 6u201, 7u191, 8u182 and 11; Java SE Embedded: 8u181; JRockit: R28.3.19. Difficult to exploit vulnerab network access via SSL/TLS to compromise Java SE, Java SE Embedded, JRockit. Successful attacks of this vulne insert or delete access to some of Java SE, Java SE Embedded, JRockit accessible data as well as unauthorized reac Embedded, JRockit accessible data and unauthorized ability to cause a partial denial of service (partial DOS) of Jav This vulnerability applies to Java deployments, typically in clients running sandboxed Java Web Start applications that load and run untrusted code (e.g. code that comes from the internet) and rely on the Java sandbox for security. using APIs in the specified Component, e.g. through a web service which supplies data to the APIs. CVSS 3.0 Base Availability impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:U/C:L/I:L/A:L).
CVE-2018-3183	Vulnerability in the Java SE, Java SE Embedded, JRockit component of Oracle Java SE (subcomponent: Scripting) Java SE: 8u182 and 11; Java SE Embedded: 8u181; JRockit: R28.3.19. Difficult to exploit vulnerability allows una via multiple protocols to compromise Java SE, Java SE Embedded, JRockit. While the vulnerability is in Java SE, significantly impact additional products. Successful attacks of this vulnerability can result in takeover of Java SE, J vulnerability applies to Java deployments, typically in clients running sandboxed Java Web Start applications or sa and run untrusted code (e.g. code that comes from the internet) and rely on the Java sandbox for security. This vuln in the specified Component, e.g. through a web service which supplies data to the APIs. CVSS 3.0 Base Score 9.0 impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:C/C:H/I:H/A:H).
CVE-2018-3214	Vulnerability in the Java SE, Java SE Embedded, JRockit component of Oracle Java SE (subcomponent: Sound). S SE: 6u201, 7u191 and 8u182; Java SE Embedded: 8u181; JRockit: R28.3.19. Easily exploitable vulnerability allow access via multiple protocols to compromise Java SE, Java SE Embedded, JRockit. Successful attacks of this vulne to cause a partial denial of service (partial DOS) of Java SE, Java SE Embedded, JRockit. Note: This vulnerability clients running sandboxed Java Web Start applications or sandboxed Java applets (in Java SE 8), that load and run the internet) and rely on the Java sandbox for security. This vulnerability can also be exploited by using APIs in the service which supplies data to the APIs. CVSS 3.0 Base Score 5.3 (Availability impacts). CVSS Vector: (CVSS:3.0
CVE-2018-3639	Systems with microprocessors utilizing speculative execution and speculative execution of memory reads before th known may allow unauthorized disclosure of information to an attacker with local user access via a side-channel an Variant 4.
CVE-2018-6003	An issue was discovered in the _asn1_decode_simple_ber function in decoding.c in GNU Libtasn1 before 4.13. Un stack exhaustion and DoS.
CVE-2018-6323	The elf_object_p function in elfcode.h in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GN overflow because bfd_size_type multiplication is not used. A crafted ELF file allows remote attackers to cause a de have unspecified other impact.
CVE-2018-6759	The bfd_get_debug_link_info_1 function in opncls.c in the Binary File Descriptor (BFD) library (aka libbfd), as di unchecked strnlen operation. Remote attackers could leverage this vulnerability to cause a denial of service (segme
CVE-2018-6829	cipher/elgamal.c in Libgcrypt through 1.8.2, when used to encrypt messages directly, improperly encodes plaintext information by reading ciphertext data (i.e., it does not have semantic security in face of a ciphertext-only attack). T assumption does not hold for Libgcrypt's ElGamal implementation.
CVE-2018-6954	systemd-tmpfiles in systemd through 237 mishandles symlinks present in non-terminal path components, which all arbitrary files via vectors involving creation of a directory and a file under that directory, and later replacing that di the fs.protected_symlinks sysctl is turned on.
CVE-2018-8740	In SQLite through 3.22.0, databases whose schema is corrupted using a CREATE TABLE AS statement could caus build.c and prepare.c.
CVE-2018-9234	GnuPG 2.2.4 and 2.2.5 does not enforce a configuration in which key certification requires an offline master Certif certifications that occurred only with access to a signing subkey.
CVE-2019-11191	The Linux kernel through 5.0.7, when CONFIG_IA32_AOUT is enabled and ia32_aout is loaded, allows local use (if any exist) because install_exec_creds() is called too late in load_aout_binary() in fs/binfmt_aout.c, and thus the condition when reading /proc/pid/stat. NOTE: the software maintainer disputes that this is a vulnerability because A been supported
CVE-2019-12378	An issue was discovered in ip6_ra_control in net/ipv6/ipv6_sockglue.c in the Linux kernel through 5.1.5. There is might allow an attacker to cause a denial of service (NULL pointer dereference and system crash). NOTE: This has
CVE-2019-12379	An issue was discovered in con_insert_unipair in drivers/tty/vt/consolemap.c in the Linux kernel through 5.1.5. Th ENOMEM outcome of kmalloc. NOTE: This id is disputed as not being an issue
CVE-2019-12381	An issue was discovered in ip_ra_control in net/ipv4/ip_sockglue.c in the Linux kernel through 5.1.5. There is an u allow an attacker to cause a denial of service (NULL pointer dereference and system crash). NOTE: this is disputed
CVE-2019-12382	An issue was discovered in drm_load_edid_firmware in drivers/gpu/drm/drm_edid_load.c in the Linux kernel throu fwstr, which might allow an attacker to cause a denial of service (NULL pointer dereference and system crash). NO being a vulnerability because kstrdup() returning NULL is handled sufficiently and there is no chance for a NULL

CVE-2019-12455	An issue was discovered in sunxi_divs_clk_setup in drivers/clk/sunxi/clk-sunxi.c in the Linux kernel through 5.1.5 derived_name, which might allow an attacker to cause a denial of service (NULL pointer dereference and system c an issue because ,ÄúThe memory allocation that was not checked is part of a code that only runs at boot time, befo is no possibility for an unprivileged user to control it, and no denial of service.,Äù
CVE-2019-12456	An issue was discovered in the MPT3COMMAND case in _ctl_ioctl_main in drivers/scsi/mpt3sas/mpt3sas_ctl.c in local users to cause a denial of service or possibly have unspecified other impact by changing the value of ioc_num aka a "double fetch" vulnerability. NOTE: a third party reports that this is unexploitable because the doubly fetched
CVE-2019-13012	The keyfile settings backend in GNOME GLib (aka glib2.0) before 2.60.0 creates directories using g_file_make_di NULL) and files using g_file_replace_contents (kfsb->file, contents, length, NULL, FALSE, G_FILE_CREATE_P NULL). Consequently, it does not properly restrict directory (and file) permissions. Instead, for directories, 0777 p permissions are used. This is similar to CVE-2019-12450.
CVE-2019-13050	Interaction between the sks-keyserver code through 1.2.0 of the SKS keyserver network, and GnuPG through 2.2.1 configuration line referring to a host on the SKS keyserver network. Retrieving data from this network may cause a Certificate Spamming Attack.
CVE-2019-13057	An issue was discovered in the server in OpenLDAP before 2.4.48. When the server administrator delegates rootDI databases but wants to maintain isolation (e.g., for multi-tenant deployments), slapd does not properly stop a rootD from another database during a SASL bind or with a proxyAuthz (RFC 4370) control. (It is not a common configur administrator and a DB administrator enjoy different levels of trust.)
CVE-2019-13565	An issue was discovered in OpenLDAP 2.x before 2.4.48. When using SASL authentication and session encryption in slapd access controls, it is possible to obtain access that would otherwise be denied via a simple bind for any ide SASL bind is completed, the sasl_ssf value is retained for all new non-SASL connections. Depending on the ACL of operations (searches, modifications, etc.). In other words, a successful authorization step completed by one user different user.
CVE-2019-13751	Uninitialized data in SQLite in Google Chrome prior to 79.0.3945.79 allowed a remote attacker to obtain potentiall via a crafted HTML page.
CVE-2019-13752	Out of bounds read in SQLite in Google Chrome prior to 79.0.3945.79 allowed a remote attacker to obtain potentia memory via a crafted HTML page.
CVE-2019-13753	Out of bounds read in SQLite in Google Chrome prior to 79.0.3945.79 allowed a remote attacker to obtain potentia memory via a crafted HTML page.
CVE-2019-16231	drivers/net/fjes/fjes_main.c in the Linux kernel 5.2.14 does not check the alloc_workqueue return value, leading to
CVE-2019-16232	drivers/net/wireless/marvell/libertas/if_sdio.c in the Linux kernel 5.2.14 does not check the alloc_workqueue return dereference.
CVE-2019-16233	drivers/scsi/qla2xxx/qla_os.c in the Linux kernel 5.2.14 does not check the alloc_workqueue return value, leading t
CVE-2019-16234	drivers/net/wireless/intel/iwlwifi/pcie/trans.c in the Linux kernel 5.2.14 does not check the alloc_workqueue return dereference.
CVE-2019-16276	Go before 1.12.10 and 1.13.x before 1.13.1 allow HTTP Request Smuggling.
CVE-2019-16276	Go before 1.12.10 and 1.13.x before 1.13.1 allow HTTP Request Smuggling.
CVE-2019-17450	find_abstract_instance in dwarf2.c in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU a denial of service (infinite recursion and application crash) via a crafted ELF file.
CVE-2019-17451	An issue was discovered in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2. SEGV in _bfd_dwarf2_find_nearest_line in dwarf2.c, as demonstrated by nm.
CVE-2019-17594	There is a heap-based buffer over-read in the _nc_find_entry function in tinfo/comp_hash.c in the terminfo library
CVE-2019-17595	There is a heap-based buffer over-read in the fmt_entry function in tinfo/comp_hash.c in the terminfo library in ncu
CVE-2019-18276	An issue was discovered in disable_priv_mode in shell.c in GNU Bash through 5.0 patch 11. By default, if Bash is UID, it will drop privileges by setting its effective UID to its real UID. However, it does so incorrectly. On Linux a functionality, the saved UID is not dropped. An attacker with command execution in the shell can use "enable -f" fo be a shared object that calls setuid() and therefore regains privileges. However, binaries running with an effective U
CVE-2019-18348	An issue was discovered in urllib2 in Python 2.x through 2.7.17 and urllib in Python 3.x through 3.8.0. CRLF injec parameter, as demonstrated by the first argument to urllib.request.urlopen with \r\n (specifically in the host compon This is similar to the CVE-2019-9740 query string issue and the CVE-2019-9947 path string issue. (This is not exp fixed.). This is fixed in: v2.7.18, v2.7.18rc1; v3.5.10, v3.5.10rc1; v3.6.11, v3.6.11rc1, v3.6.12; v3.7.8, v3.7.8rc1, v3 v3.8.5, v3.8.6, v3.8.6rc1.

CVE-2019-19070	A memory leak in the spi_gpio_probe() function in drivers/spi/spi-gpio.c in the Linux kernel through 5.3.11 allows (memory consumption) by triggering devm_add_action_or_reset() failures, aka CID-d3b0ffa1d75d. NOTE: third p the system must have already been out of memory before the probe began
CVE-2019-19449	In the Linux kernel 5.0.21, mounting a crafted f2fs filesystem image can lead to slab-out-of-bounds read access in segment.c, related to init_min_max_mtime in fs/f2fs/segment.c (because the second argument to get_seg_entry is
CVE-2019-19603	SQLite 3.30.1 mishandles certain SELECT statements with a nonexistent VIEW, leading to an application crash.
CVE-2019-19645	alter.c in SQLite through 3.30.1 allows attackers to trigger infinite recursion via certain types of self-referential vie statements.
CVE-2019-19880	exprListAppendList in window.c in SQLite 3.30.1 allows attackers to trigger an invalid pointer dereference becaus clauses of window definitions are mishandled.
CVE-2019-19906	cyrus-sasl (aka Cyrus SASL) 2.1.27 has an out-of-bounds write leading to unauthenticated remote denial-of-service packet. The OpenLDAP crash is ultimately caused by an off-by-one error in _sasl_add_string in common.c in cyru
CVE-2019-19924	SQLite 3.30.1 mishandles certain parser-tree rewriting, related to expr.c, vdbeaux.c, and window.c. This is caused handling.
CVE-2019-20218	selectExpander in select.c in SQLite 3.30.1 proceeds with WITH stack unwinding even after a parsing error.
CVE-2019-20387	repodata_schema2id in repodata.c in libsolv before 0.7.6 has a heap-based buffer over-read via a last schema whos schema.
CVE-2019-2422	Vulnerability in the Java SE component of Oracle Java SE (subcomponent: Libraries). Supported versions that are 11.0.1; Java SE Embedded: 8u191. Difficult to exploit vulnerability allows unauthenticated attacker with network a Java SE. Successful attacks require human interaction from a person other than the attacker. Successful attacks of t read access to a subset of Java SE accessible data. Note: This vulnerability applies to Java deployments, typically i applications or sandboxed Java applets (in Java SE 8), that load and run untrusted code (e.g., code that comes from security. This vulnerability does not apply to Java deployments, typically in servers, that load and run only trusted CVSS 3.0 Base Score 3.1 (Confidentiality impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:R/S:U/C:L/I:
CVE-2019-2426	Vulnerability in the Java SE component of Oracle Java SE (subcomponent: Networking). Supported versions that a 11.0.1; Java SE Embedded: 8u191. Difficult to exploit vulnerability allows unauthenticated attacker with network a Java SE. Successful attacks of this vulnerability can result in unauthorized read access to a subset of Java SE acces to Java deployments, typically in clients running sandboxed Java Web Start applications or sandboxed Java applets code (e.g., code that comes from the internet) and rely on the Java sandbox for security. This vulnerability can also Component, e.g., through a web service which supplies data to the APIs. CVSS 3.0 Base Score 3.7 (Confidentiality AC:H/PR:N/UI:N/S:U/C:L/I:N/A:N).
CVE-2019-2602	Vulnerability in the Java SE, Java SE Embedded component of Oracle Java SE (subcomponent: Libraries). Suppor SE: 7u211, 8u202, 11.0.2 and 12; Java SE Embedded: 8u201. Easily exploitable vulnerability allows unauthenticat multiple protocols to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result i frequently repeatable crash (complete DOS) of Java SE, Java SE Embedded. Note: This vulnerability can only be e specified Component without using Untrusted Java Web Start applications or Untrusted Java applets, such as throu (Availability impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:L/PR:N/UI:N/S:U/C:N/I:N/A:H).
CVE-2019-2684	Vulnerability in the Java SE, Java SE Embedded component of Oracle Java SE (subcomponent: RMI). Supported v 8u202, 11.0.2 and 12; Java SE Embedded: 8u201. Difficult to exploit vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized crea critical data or all Java SE, Java SE Embedded accessible data. Note: This vulnerability applies to Java deployment Web Start applications or sandboxed Java applets (in Java SE 8), that load and run untrusted code (e.g., code that c sandbox for security. This vulnerability can also be exploited by using APIs in the specified Component, e.g., throu APIs. CVSS 3.0 Base Score 5.9 (Integrity impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:U/C:N/I
CVE-2019-2698	Vulnerability in the Java SE component of Oracle Java SE (subcomponent: 2D). Supported versions that are affect exploit vulnerability allows unauthenticated attacker with network access via multiple protocols to compromise Jav can result in takeover of Java SE. Note: This vulnerability applies to Java deployments, typically in clients running sandboxed Java applets (in Java SE 8), that load and run untrusted code (e.g., code that comes from the internet) an vulnerability does not apply to Java deployments, typically in servers, that load and run only trusted code (e.g., cod Base Score 8.1 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/
CVE-2019-2708	Vulnerability in the Data Store component of Oracle Berkeley DB. Supported versions that are affected are Prior to Easily exploitable vulnerability allows low privileged attacker having Local Logon privilege with logon to the infra compromise Data Store. Successful attacks of this vulnerability can result in unauthorized ability to cause a partial CVSS 3.0 Base Score 3.3 (Availability impacts). CVSS Vector: (CVSS:3.0/AV:L/AC:L/PR:L/UI:N/S:U/C:N/I:N/A

CVE-2019-2745	Vulnerability in the Java SE component of Oracle Java SE (subcomponent: Security). Supported versions that are a and 11.0.3. Difficult to exploit vulnerability allows unauthenticated attacker with logon to the infrastructure where SE. Successful attacks of this vulnerability can result in unauthorized access to critical data or complete access to a vulnerability applies to Java deployments, typically in clients running sandboxed Java Web Start applications or sa and run untrusted code (e.g., code that comes from the internet) and rely on the Java sandbox for security. This vul in the specified Component, e.g., through a web service which supplies data to the APIs. CVSS 3.0 Base Score 5.1 (CVSS:3.0/AV:L/AC:H/PR:N/UI:N/S:U/C:H/I:N/A:N).
CVE-2019-2762	Vulnerability in the Java SE, Java SE Embedded component of Oracle Java SE (subcomponent: Utilities). Supporte 7u221, 8u212, 11.0.3 and 12.0.1; Java SE Embedded: 8u211. Easily exploitable vulnerability allows unauthenticate protocols to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unautho service (partial DOS) of Java SE, Java SE Embedded. Note: This vulnerability applies to Java deployments, typical Start applications or sandboxed Java applets (in Java SE 8), that load and run untrusted code (e.g., code that comes sandbox for security. This vulnerability can also be exploited by using APIs in the specified Component, e.g., throu APIs. CVSS 3.0 Base Score 5.3 (Availability impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:L/PR:N/UI:N/S:U/C:
CVE-2019-2766	Vulnerability in the Java SE, Java SE Embedded component of Oracle Java SE (subcomponent: Networking). Supp SE: 7u221, 8u212, 11.0.3 and 12.0.1; Java SE Embedded: 8u211. Difficult to exploit vulnerability allows unauthen multiple protocols to compromise Java SE, Java SE Embedded. Successful attacks require human interaction from attacks of this vulnerability can result in unauthorized read access to a subset of Java SE, Java SE Embedded acces to Java deployments, typically in clients running sandboxed Java Web Start applications or sandboxed Java applets code (e.g., code that comes from the internet) and rely on the Java sandbox for security. This vulnerability can also Component, e.g., through a web service which supplies data to the APIs. CVSS 3.0 Base Score 3.1 (Confidentiality AC:H/PR:N/UI:R/S:U/C:L/I:N/A:N).
CVE-2019-2769	Vulnerability in the Java SE, Java SE Embedded component of Oracle Java SE (subcomponent: Utilities). Supporte 7u221, 8u212, 11.0.3 and 12.0.1; Java SE Embedded: 8u211. Easily exploitable vulnerability allows unauthenticate protocols to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unautho service (partial DOS) of Java SE, Java SE Embedded. Note: This vulnerability applies to Java deployments, typical Start applications or sandboxed Java applets (in Java SE 8), that load and run untrusted code (e.g., code that comes sandbox for security. This vulnerability can also be exploited by using APIs in the specified Component, e.g., throu APIs. CVSS 3.0 Base Score 5.3 (Availability impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:L/PR:N/UI:N/S:U/C:
CVE-2019-2786	Vulnerability in the Java SE, Java SE Embedded component of Oracle Java SE (subcomponent: Security). Supporte 8u212, 11.0.3 and 12.0.1; Java SE Embedded: 8u211. Difficult to exploit vulnerability allows unauthenticated attac protocols to compromise Java SE, Java SE Embedded. Successful attacks require human interaction from a person vulnerability is in Java SE, Java SE Embedded, attacks may significantly impact additional products. Successful at unauthorized read access to a subset of Java SE, Java SE Embedded accessible data. Note: This vulnerability applie running sandboxed Java Web Start applications or sandboxed Java applets (in Java SE 8), that load and run untruste internet) and rely on the Java sandbox for security. This vulnerability can also be exploited by using APIs in the sp service which supplies data to the APIs. CVSS 3.0 Base Score 3.4 (Confidentiality impacts). CVSS Vector: (CVSS A:N).
CVE-2019-2816	Vulnerability in the Java SE, Java SE Embedded component of Oracle Java SE (subcomponent: Networking). Supp SE: 7u221, 8u212, 11.0.3 and 12.0.1; Java SE Embedded: 8u211. Difficult to exploit vulnerability allows unauthen multiple protocols to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in to some of Java SE, Java SE Embedded accessible data as well as unauthorized read access to a subset of Java SE, This vulnerability applies to Java deployments, typically in clients running sandboxed Java Web Start applications load and run untrusted code (e.g., code that comes from the internet) and rely on the Java sandbox for security. This APIs in the specified Component, e.g., through a web service which supplies data to the APIs. CVSS 3.0 Base Scor CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:U/C:L/I:L/A:N).
CVE-2019-2842	Vulnerability in the Java SE component of Oracle Java SE (subcomponent: JCE). The supported version that is affe vulnerability allows unauthenticated attacker with network access via multiple protocols to compromise Java SE. S result in unauthorized ability to cause a partial denial of service (partial DOS) of Java SE. Note: This vulnerability clients running sandboxed Java Web Start applications or sandboxed Java applets (in Java SE 8), that load and run the internet) and rely on the Java sandbox for security. This vulnerability can also be exploited by using APIs in the service which supplies data to the APIs. CVSS 3.0 Base Score 3.7 (Availability impacts). CVSS Vector: (CVSS:3.0
CVE-2019-2894	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Security). Supported vers 8u221, 11.0.4 and 13; Java SE Embedded: 8u221. Difficult to exploit vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized read Embedded accessible data. Note: This vulnerability applies to Java deployments, typically in clients running sandb sandboxed Java applets (in Java SE 8), that load and run untrusted code (e.g., code that comes from the internet) an vulnerability can also be exploited by using APIs in the specified Component, e.g., through a web service which su 3.7 (Confidentiality impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:U/C:L/I:N/A:N).

CVE-2019-2933	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Libraries). Supported vers 8u221, 11.0.4 and 13; Java SE Embedded: 8u221. Difficult to exploit vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. Successful attacks require human interaction from a person other than vulnerability can result in unauthorized read access to a subset of Java SE, Java SE Embedded accessible data. Not deployments, typically in clients running sandboxed Java Web Start applications or sandboxed Java applets (in Jav (e.g., code that comes from the internet) and rely on the Java sandbox for security. This vulnerability can also be ex Component, e.g., through a web service which supplies data to the APIs. CVSS 3.0 Base Score 3.1 (Confidentiality AC:H/PR:N/UI:R/S:U/C:L/I:N/A:N).
CVE-2019-2945	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Networking). Supported v 7u231, 8u221, 11.0.4 and 13; Java SE Embedded: 8u221. Difficult to exploit vulnerability allows unauthenticated a protocols to compromise Java SE, Java SE Embedded. Successful attacks require human interaction from a person of this vulnerability can result in unauthorized ability to cause a partial denial of service (partial DOS) of Java SE, applies to Java deployments, typically in clients running sandboxed Java Web Start applications or sandboxed Java untrusted code (e.g., code that comes from the internet) and rely on the Java sandbox for security. This vulnerabilit typically in servers, that load and run only trusted code (e.g., code installed by an administrator). CVSS 3.0 Base S (CVSS:3.0/AV:N/AC:H/PR:N/UI:R/S:U/C:N/I:N/A:L).
CVE-2019-2949	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Kerberos). Supported ver 8u221, 11.0.4 and 13; Java SE Embedded: 8u221. Difficult to exploit vulnerability allows unauthenticated attacker compromise Java SE, Java SE Embedded. While the vulnerability is in Java SE, Java SE Embedded, attacks may s Successful attacks of this vulnerability can result in unauthorized access to critical data or complete access to all Ja Note: This vulnerability applies to Java deployments, typically in clients running sandboxed Java Web Start applic 8), that load and run untrusted code (e.g., code that comes from the internet) and rely on the Java sandbox for secur by using APIs in the specified Component, e.g., through a web service which supplies data to the APIs. CVSS 3.0 CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:C/C:H/I:N/A:N).
CVE-2019-2958	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Libraries). Supported vers 8u221, 11.0.4 and 13; Java SE Embedded: 8u221. Difficult to exploit vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized crea critical data or all Java SE, Java SE Embedded accessible data. Note: This vulnerability applies to Java deployment Web Start applications or sandboxed Java applets (in Java SE 8), that load and run untrusted code (e.g., code that c sandbox for security. This vulnerability can also be exploited by using APIs in the specified Component, e.g., throu APIs. CVSS 3.0 Base Score 5.9 (Integrity impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:U/C:N/I
CVE-2019-2962	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: 2D). Supported versions t 11.0.4 and 13; Java SE Embedded: 8u221. Difficult to exploit vulnerability allows unauthenticated attacker with ne compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized ability DOS) of Java SE, Java SE Embedded. Note: This vulnerability applies to Java deployments, typically in clients run or sandboxed Java applets (in Java SE 8), that load and run untrusted code (e.g., code that comes from the internet) This vulnerability can also be exploited by using APIs in the specified Component, e.g., through a web service whi Score 3.7 (Availability impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:U/C:N/I:N/A:L).
CVE-2019-2964	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Concurrency). Supported 8u221, 11.0.4 and 13; Java SE Embedded: 8u221. Difficult to exploit vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized abil DOS) of Java SE, Java SE Embedded. Note: This vulnerability can only be exploited by supplying data to APIs in Untrusted Java Web Start applications or Untrusted Java applets, such as through a web service. CVSS 3.0 Base Sc (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:U/C:N/I:N/A:L).
CVE-2019-2973	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: JAXP). Supported versio 8u221, 11.0.4 and 13; Java SE Embedded: 8u221. Difficult to exploit vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized abil DOS) of Java SE, Java SE Embedded. Note: This vulnerability applies to Java deployments, typically in clients run or sandboxed Java applets (in Java SE 8), that load and run untrusted code (e.g., code that comes from the internet) This vulnerability can also be exploited by using APIs in the specified Component, e.g., through a web service whi Score 3.7 (Availability impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:U/C:N/I:N/A:L).
CVE-2019-2975	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Scripting). Supported ver 11.0.4 and 13; Java SE Embedded: 8u221. Difficult to exploit vulnerability allows unauthenticated attacker with ne compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized update SE, Java SE Embedded accessible data and unauthorized ability to cause a partial denial of service (partial DOS) o vulnerability applies to Java deployments, typically in clients running sandboxed Java Web Start applications or sa and run untrusted code (e.g., code that comes from the internet) and rely on the Java sandbox for security. This vul in the specified Component, e.g., through a web service which supplies data to the APIs. CVSS 3.0 Base Score 4.8 Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:U/C:N/I:L/A:L).

CVE-2019-2978	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Networking). Supported v 8u221, 11.0.4 and 13; Java SE Embedded: 8u221. Difficult to exploit vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized abil DOS) of Java SE, Java SE Embedded. Note: This vulnerability applies to Java deployments, typically in clients run or sandboxed Java applets (in Java SE 8), that load and run untrusted code (e.g., code that comes from the internet) This vulnerability can also be exploited by using APIs in the specified Component, e.g., through a web service whi Score 3.7 (Availability impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:U/C:N/I:N/A:L).
CVE-2019-2981	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: JAXP). Supported versio 8u221, 11.0.4 and 13; Java SE Embedded: 8u221. Difficult to exploit vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized abil DOS) of Java SE, Java SE Embedded. Note: This vulnerability applies to Java deployments, typically in clients run or sandboxed Java applets (in Java SE 8), that load and run untrusted code (e.g., code that comes from the internet) This vulnerability can also be exploited by using APIs in the specified Component, e.g., through a web service whi Score 3.7 (Availability impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:U/C:N/I:N/A:L).
CVE-2019-2983	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Serialization). Supported 8u221, 11.0.4 and 13; Java SE Embedded: 8u221. Difficult to exploit vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized abil DOS) of Java SE, Java SE Embedded. Note: This vulnerability applies to Java deployments, typically in clients run or sandboxed Java applets (in Java SE 8), that load and run untrusted code (e.g., code that comes from the internet) This vulnerability can also be exploited by using APIs in the specified Component, e.g., through a web service whi Score 3.7 (Availability impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:U/C:N/I:N/A:L).
CVE-2019-2987	Vulnerability in the Java SE product of Oracle Java SE (component: 2D). Supported versions that are affected are J vulnerability allows unauthenticated attacker with network access via multiple protocols to compromise Java SE. S result in unauthorized ability to cause a partial denial of service (partial DOS) of Java SE. Note: This vulnerability clients running sandboxed Java Web Start applications or sandboxed Java applets (in Java SE 8), that load and run the internet) and rely on the Java sandbox for security. This vulnerability can also be exploited by using APIs in the service which supplies data to the APIs. CVSS 3.0 Base Score 3.7 (Availability impacts). CVSS Vector: (CVSS:3.0
CVE-2019-2988	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: 2D). Supported versions t 11.0.4 and 13; Java SE Embedded: 8u221. Difficult to exploit vulnerability allows unauthenticated attacker with ne compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized ability DOS) of Java SE, Java SE Embedded. Note: This vulnerability applies to Java deployments, typically in clients run or sandboxed Java applets (in Java SE 8), that load and run untrusted code (e.g., code that comes from the internet) This vulnerability does not apply to Java deployments, typically in servers, that load and run only trusted code (e.g 3.0 Base Score 3.7 (Availability impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:U/C:N/I:N/A:L).
CVE-2019-2989	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Networking). Supported v 8u221, 11.0.4 and 13; Java SE Embedded: 8u221. Difficult to exploit vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. While the vulnerability is in Java SE, Java SE Embedded, attacks may Successful attacks of this vulnerability can result in unauthorized creation, deletion or modification access to critica accessible data. Note: This vulnerability applies to Java deployments, typically in clients running sandboxed Java V applets (in Java SE 8), that load and run untrusted code (e.g., code that comes from the internet) and rely on the Jav also be exploited by using APIs in the specified Component, e.g., through a web service which supplies data to the impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:C/C:N/I:H/A:N).
CVE-2019-2992	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: 2D). Supported versions t 11.0.4 and 13; Java SE Embedded: 8u221. Difficult to exploit vulnerability allows unauthenticated attacker with ne compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized ability DOS) of Java SE, Java SE Embedded. Note: This vulnerability applies to Java deployments, typically in clients run or sandboxed Java applets (in Java SE 8), that load and run untrusted code (e.g., code that comes from the internet) This vulnerability does not apply to Java deployments, typically in servers, that load and run only trusted code (e.g 3.0 Base Score 3.7 (Availability impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:U/C:N/I:N/A:L).
CVE-2019-2999	Vulnerability in the Java SE product of Oracle Java SE (component: Javadoc). Supported versions that are affected 13. Difficult to exploit vulnerability allows unauthenticated attacker with network access via multiple protocols to require human interaction from a person other than the attacker and while the vulnerability is in Java SE, attacks m Successful attacks of this vulnerability can result in unauthorized update, insert or delete access to some of Java SE read access to a subset of Java SE accessible data. Note: This vulnerability applies to Java deployments, typically i applications or sandboxed Java applets (in Java SE 8), that load and run untrusted code (e.g., code that comes from security. This vulnerability does not apply to Java deployments, typically in servers, that load and run only trusted CVSS 3.0 Base Score 4.7 (Confidentiality and Integrity impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI
CVE-2019-3842	In systemd before v242-rc4, it was discovered that pam_systemd does not properly sanitize the environment before for an attacker, in some particular configurations, to set a XDG_SEAT environment variable which allows for com using the "allow_active" element rather than "allow_any".
CVE-2019-3859	An out of bounds read flaw was discovered in libssh2 before 1.8.1 in the _libssh2_packet_require and _libssh2_pac who compromises a SSH server may be able to cause a Denial of Service or read data in the client memory.

CVE-2019-3860	An out of bounds read flaw was discovered in libssh2 before 1.8.1 in the way SFTP packets with empty payloads a compromises a SSH server may be able to cause a Denial of Service or read data in the client memory.
CVE-2019-5827	Integer overflow in SQLite via WebSQL in Google Chrome prior to 74.0.3729.131 allowed a remote attacker to po HTML page.
CVE-2019-9074	An issue was discovered in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2. SEGV in bfd_getl32 in libbfd.c, when called from pex64_get_runtime_function in pei-x86_64.c.
CVE-2020-11725	snd_ctl_elem_add in sound/core/control.c in the Linux kernel through 5.6.3 has a count=info->owner line, which la multiplication for unspecified "interesting side effects." NOTE: kernel engineers dispute this finding, because it co were added that were unfamiliar with the misuse of the info->owner field to represent data unrelated to the "owner" SNDRV_CTL_IOCTL_ELEM_ADD and SNDRV_CTL_IOCTL_ELEM_REPLACE, have been designed to misu
CVE-2020-12762	json-c through 0.14 has an integer overflow and out-of-bounds write via a large JSON file, as demonstrated by prin
CVE-2020-13435	SQLite through 3.32.0 has a segmentation fault in sqlite3ExprCodeTarget in expr.c.
CVE-2020-13631	SQLite before 3.32.0 allows a virtual table to be renamed to the name of one of its shadow tables, related to alter.c
CVE-2020-13776	systemd through v245 mishandles numerical usernames such as ones composed of decimal digits or 0x followed by privileges when privileges of the 0x0 user account were intended.  **Note:** This issue exists because of an incomplete fix for CVE-2017-1000082.
CVE-2020-14155	libpcre in PCRE before 8.44 allows an integer overflow via a large number after a (?C substring.
CVE-2020-14350	It was found that some PostgreSQL extensions did not use search_path safely in their installation script. An attacke to trick an administrator into executing a specially crafted script, during the installation or update of such extension 12.4, before 11.9, before 10.14, before 9.6.19, and before 9.5.23.
CVE-2020-14556	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Libraries). Supported vers 11.0.7 and 14.0.1; Java SE Embedded: 8u251. Difficult to exploit vulnerability allows unauthenticated attacker wit to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized upd SE, Java SE Embedded accessible data as well as unauthorized read access to a subset of Java SE, Java SE Embedd and server deployment of Java. This vulnerability can be exploited through sandboxed Java Web Start applications be exploited by supplying data to APIs in the specified Component without using sandboxed Java Web Start applic through a web service. CVSS 3.1 Base Score 4.8 (Confidentiality and Integrity impacts). CVSS Vector: (CVSS:3.1
CVE-2020-14577	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: JSSE). Supported version 11.0.7 and 14.0.1; Java SE Embedded: 8u251. Difficult to exploit vulnerability allows unauthenticated attacker wit Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized read access to a sub data. Note: Applies to client and server deployment of Java. This vulnerability can be exploited through sandboxed Java applets. It can also be exploited by supplying data to APIs in the specified Component without using sandbox Java applets, such as through a web service. CVSS 3.1 Base Score 3.7 (Confidentiality impacts). CVSS Vector: (C I:N/A:N).
CVE-2020-14578	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Libraries). Supported vers 8u251; Java SE Embedded: 8u251. Difficult to exploit vulnerability allows unauthenticated attacker with network a Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized ability to cause a pa Java SE, Java SE Embedded. Note: Applies to client and server deployment of Java. This vulnerability can be expl applications and sandboxed Java applets. It can also be exploited by supplying data to APIs in the specified Compo applications or sandboxed Java applets, such as through a web service. CVSS 3.1 Base Score 3.7 (Availability impa PR:N/UI:N/S:U/C:N/I:N/A:L).
CVE-2020-14579	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Libraries). Supported vers 8u251; Java SE Embedded: 8u251. Difficult to exploit vulnerability allows unauthenticated attacker with network a Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized ability to cause a pa Java SE, Java SE Embedded. Note: Applies to client and server deployment of Java. This vulnerability can be expl applications and sandboxed Java applets. It can also be exploited by supplying data to APIs in the specified Compo applications or sandboxed Java applets, such as through a web service. CVSS 3.1 Base Score 3.7 (Availability impa PR:N/UI:N/S:U/C:N/I:N/A:L).
CVE-2020-14581	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: 2D). Supported versions t 11.0.7 and 14.0.1; Java SE Embedded: 8u251. Difficult to exploit vulnerability allows unauthenticated attacker wit to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized read Embedded accessible data. Note: Applies to client and server deployment of Java. This vulnerability can be exploit applications and sandboxed Java applets. It can also be exploited by supplying data to APIs in the specified Compo applications or sandboxed Java applets, such as through a web service. CVSS 3.1 Base Score 3.7 (Confidentiality i AC:H/PR:N/UI:N/S:U/C:L/I:N/A:N).

CVE-2020-14583	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Libraries). Supported vers 8u251, 11.0.7 and 14.0.1; Java SE Embedded: 8u251. Difficult to exploit vulnerability allows unauthenticated attac protocols to compromise Java SE, Java SE Embedded. Successful attacks require human interaction from a person vulnerability is in Java SE, Java SE Embedded, attacks may significantly impact additional products. Successful at in takeover of Java SE, Java SE Embedded. Note: This vulnerability applies to Java deployments, typically in clien applications or sandboxed Java applets, that load and run untrusted code (e.g., code that comes from the internet) a vulnerability does not apply to Java deployments, typically in servers, that load and run only trusted code (e.g., cod Base Score 8.3 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:H/PR:N/
CVE-2020-14593	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: 2D). Supported versions t 11.0.7 and 14.0.1; Java SE Embedded: 8u251. Easily exploitable vulnerability allows unauthenticated attacker with compromise Java SE, Java SE Embedded. Successful attacks require human interaction from a person other than th Java SE, Java SE Embedded, attacks may significantly impact additional products. Successful attacks of this vulne deletion or modification access to critical data or all Java SE, Java SE Embedded accessible data. Note: This vulne in clients running sandboxed Java Web Start applications or sandboxed Java applets, that load and run untrusted co and rely on the Java sandbox for security. This vulnerability does not apply to Java deployments, typically in serve code installed by an administrator). CVSS 3.1 Base Score 7.4 (Integrity impacts). CVSS Vector: (CVSS:3.1/AV:N/
CVE-2020-14621	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: JAXP). Supported versio 8u251, 11.0.7 and 14.0.1; Java SE Embedded: 8u251. Easily exploitable vulnerability allows unauthenticated attac protocols to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unautho of Java SE, Java SE Embedded accessible data. Note: This vulnerability can only be exploited by supplying data to using Untrusted Java Web Start applications or Untrusted Java applets, such as through a web service. CVSS 3.1 B Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:N/I:L/A:N).
CVE-2020-14779	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Serialization). Supported 8u261, 11.0.8 and 15; Java SE Embedded: 8u261. Difficult to exploit vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized abil DOS) of Java SE, Java SE Embedded. Note: Applies to client and server deployment of Java. This vulnerability ca Start applications and sandboxed Java applets. It can also be exploited by supplying data to APIs in the specified C Start applications or sandboxed Java applets, such as through a web service. CVSS 3.1 Base Score 3.7 (Availability AC:H/PR:N/UI:N/S:U/C:N/I:N/A:L).
CVE-2020-14781	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: JNDI). Supported version 8u261, 11.0.8 and 15; Java SE Embedded: 8u261. Difficult to exploit vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized read Embedded accessible data. Note: Applies to client and server deployment of Java. This vulnerability can be exploit applications and sandboxed Java applets. It can also be exploited by supplying data to APIs in the specified Compo applications or sandboxed Java applets, such as through a web service. CVSS 3.1 Base Score 3.7 (Confidentiality i AC:H/PR:N/UI:N/S:U/C:L/I:N/A:N).
CVE-2020-14782	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Libraries). Supported vers 8u261, 11.0.8 and 15; Java SE Embedded: 8u261. Difficult to exploit vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized upd SE, Java SE Embedded accessible data. Note: Applies to client and server deployment of Java. This vulnerability c Start applications and sandboxed Java applets. It can also be exploited by supplying data to APIs in the specified C Start applications or sandboxed Java applets, such as through a web service. CVSS 3.1 Base Score 3.7 (Integrity im AC:H/PR:N/UI:N/S:U/C:N/I:L/A:N).
CVE-2020-14792	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Hotspot). Supported versi 8u261, 11.0.8 and 15; Java SE Embedded: 8u261. Difficult to exploit vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. Successful attacks require human interaction from a person other than vulnerability can result in unauthorized update, insert or delete access to some of Java SE, Java SE Embedded acce access to a subset of Java SE, Java SE Embedded accessible data. Note: Applies to client and server deployment of through sandboxed Java Web Start applications and sandboxed Java applets. It can also be exploited by supplying without using sandboxed Java Web Start applications or sandboxed Java applets, such as through a web service. C Integrity impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:H/PR:N/UI:R/S:U/C:L/I:L/A:N).
CVE-2020-14796	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Libraries). Supported vers 8u261, 11.0.8 and 15; Java SE Embedded: 8u261. Difficult to exploit vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. Successful attacks require human interaction from a person other than vulnerability can result in unauthorized read access to a subset of Java SE, Java SE Embedded accessible data. Not deployments, typically in clients running sandboxed Java Web Start applications or sandboxed Java applets, that lo comes from the internet) and rely on the Java sandbox for security. This vulnerability does not apply to Java deploy only trusted code (e.g., code installed by an administrator). CVSS 3.1 Base Score 3.1 (Confidentiality impacts). CV UI:R/S:U/C:L/I:N/A:N).

CVE-2020-14797	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Libraries). Supported vers 8u261, 11.0.8 and 15; Java SE Embedded: 8u261. Difficult to exploit vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized upd SE, Java SE Embedded accessible data. Note: Applies to client and server deployment of Java. This vulnerability c Start applications and sandboxed Java applets. It can also be exploited by supplying data to APIs in the specified C Start applications or sandboxed Java applets, such as through a web service. CVSS 3.1 Base Score 3.7 (Integrity im AC:H/PR:N/UI:N/S:U/C:N/I:L/A:N).
CVE-2020-14798	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Libraries). Supported vers 8u261, 11.0.8 and 15; Java SE Embedded: 8u261. Difficult to exploit vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. Successful attacks require human interaction from a person other than vulnerability can result in unauthorized update, insert or delete access to some of Java SE, Java SE Embedded acce to Java deployments, typically in clients running sandboxed Java Web Start applications or sandboxed Java applets that comes from the internet) and rely on the Java sandbox for security. This vulnerability does not apply to Java de and run only trusted code (e.g., code installed by an administrator). CVSS 3.1 Base Score 3.1 (Integrity impacts). C UI:R/S:U/C:N/I:L/A:N).
CVE-2020-14803	Vulnerability in the Java SE product of Oracle Java SE (component: Libraries). Supported versions that are affecte exploitable vulnerability allows unauthenticated attacker with network access via multiple protocols to compromise vulnerability can result in unauthorized read access to a subset of Java SE accessible data. Note: This vulnerability clients running sandboxed Java Web Start applications or sandboxed Java applets, that load and run untrusted code rely on the Java sandbox for security. This vulnerability does not apply to Java deployments, typically in servers, th installed by an administrator). CVSS 3.1 Base Score 5.3 (Confidentiality impacts). CVSS Vector: (CVSS:3.1/AV:N
CVE-2020-16590	A double free vulnerability exists in the Binary File Descriptor (BFD) (aka libbrd) in GNU Binutils 2.35 in the proc readelf, via a crafted file.
CVE-2020-16591	A Denial of Service vulnerability exists in the Binary File Descriptor (BFD) in GNU Binutils 2.35 due to an invalid demonstrated in readeif.
CVE-2020-16592	A use after free issue exists in the Binary File Descriptor (BFD) library (aka libbfd) in GNU Binutils 2.34 in bfd_ha can cause a denial of service via a crafted file.
CVE-2020-16593	A Null Pointer Dereference vulnerability exists in the Binary File Descriptor (BFD) library (aka libbfd), as distribu scan_unit_for_symbols, as demonstrated in addr2line, that can cause a denial of service via a crafted file.
CVE-2020-16599	A Null Pointer Dereference vulnerability exists in the Binary File Descriptor (BFD) library (aka libbfd), as distribu _bfd_elf_get_symbol_version_string, as demonstrated in nm-new, that can cause a denial of service via a crafted fi
CVE-2020-16845	Go before 1.13.15 and 14.x before 1.14.7 can have an infinite read loop in ReadUvarint and ReadVarint in encodin
CVE-2020-16845	Go before 1.13.15 and 14.x before 1.14.7 can have an infinite read loop in ReadUvarint and ReadVarint in encodin
CVE-2020-21583	An issue was discovered in hwclock.13-v2.27 allows attackers to gain escalated privlidges or execute arbitrary com the date.
CVE-2020-22218	An issue was discovered in function _libssh2_packet_add in libssh2 1.10.0 allows attackers to access out of bounds
CVE-2020-24553	Go before 1.14.8 and 1.15.x before 1.15.1 allows XSS because text/html is the default for CGI/FCGI handlers that
CVE-2020-24736	Buffer Overflow vulnerability found in SQLite3 v.3.27.1 and before allows a local attacker to cause a denial of serv
CVE-2020-24977	GNOME project libxml2 v2.9.10 has a global buffer over-read vulnerability in xmlEncodeEntitiesInternal at libxm commit 50f06b3e.
CVE-2020-25696	A flaw was found in the psql interactive terminal of PostgreSQL in versions before 13.1, before 12.5, before 11.10, 9.5.24. If an interactive psql session uses \gset when querying a compromised server, the attacker can execute arbit running psql. The highest threat from this vulnerability is to data confidentiality and integrity as well as system ava
CVE-2020-2583	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Serialization). Supported 7u241, 8u231, 11.0.5 and 13.0.1; Java SE Embedded: 8u231. Difficult to exploit vulnerability allows unauthenticat protocols to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unautho service (partial DOS) of Java SE, Java SE Embedded. Note: This vulnerability applies to Java deployments, typical Start applications or sandboxed Java applets (in Java SE 8), that load and run untrusted code (e.g., code that comes sandbox for security. This vulnerability can also be exploited by using APIs in the specified Component, e.g., throu APIs. CVSS 3.0 Base Score 3.7 (Availability impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:U/C:
CVE-2020-2590	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Security). Supported vers 8u231, 11.0.5 and 13.0.1; Java SE Embedded: 8u231. Difficult to exploit vulnerability allows unauthenticated attac compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized update Java SE Embedded accessible data. Note: This vulnerability applies to Java deployments, typically in clients runni sandboxed Java applets (in Java SE 8), that load and run untrusted code (e.g., code that comes from the internet) an vulnerability can also be exploited by using APIs in the specified Component, e.g., through a web service which su 3.7 (Integrity impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:U/C:N/I:L/A:N).

CVE-2020-2593	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Networking). Supported v 7u241, 8u231, 11.0.5 and 13.0.1; Java SE Embedded: 8u231. Difficult to exploit vulnerability allows unauthenticat protocols to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unautho some of Java SE, Java SE Embedded accessible data as well as unauthorized read access to a subset of Java SE, Jav vulnerability applies to Java deployments, typically in clients running sandboxed Java Web Start applications or sa load and run untrusted code (e.g., code that comes from the internet) and rely on the Java sandbox for security. Thi APIs in the specified Component, e.g., through a web service which supplies data to the APIs. CVSS 3.0 Base Scor CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:U/C:L/I:L/A:N).
CVE-2020-2601	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Security). Supported vers 8u231, 11.0.5 and 13.0.1; Java SE Embedded: 8u231. Difficult to exploit vulnerability allows unauthenticated attac compromise Java SE, Java SE Embedded. While the vulnerability is in Java SE, Java SE Embedded, attacks may s Successful attacks of this vulnerability can result in unauthorized access to critical data or complete access to all Ja Note: This vulnerability applies to Java deployments, typically in clients running sandboxed Java Web Start applica 8), that load and run untrusted code (e.g., code that comes from the internet) and rely on the Java sandbox for secur by using APIs in the specified Component, e.g., through a web service which supplies data to the APIs. CVSS 3.0 B CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:C/C:H/I:N/A:N).
CVE-2020-2604	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Serialization). Supported 7u241, 8u231, 11.0.5 and 13.0.1; Java SE Embedded: 8u231. Difficult to exploit vulnerability allows unauthenticat protocols to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in takeove vulnerability applies to Java deployments, typically in clients running sandboxed Java Web Start applications or sa and run untrusted code (e.g., code that comes from the internet) and rely on the Java sandbox for security. This vul in the specified Component, e.g., through a web service which supplies data to the APIs. CVSS v3.0 Base Score 8. impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:U/C:H/I:H/A:H).
CVE-2020-2654	Vulnerability in the Java SE product of Oracle Java SE (component: Libraries). Supported versions that are affecte 13.0.1. Difficult to exploit vulnerability allows unauthenticated attacker with network access via multiple protocols of this vulnerability can result in unauthorized ability to cause a partial denial of service (partial DOS) of Java SE. by supplying data to APIs in the specified Component without using Untrusted Java Web Start applications or Untr service. CVSS 3.0 Base Score 3.7 (Availability impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:U/(
CVE-2020-2659	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Networking). Supported v 7u241 and 8u231; Java SE Embedded: 8u231. Difficult to exploit vulnerability allows unauthenticated attacker wit compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized ability DOS) of Java SE, Java SE Embedded. Note: This vulnerability applies to Java deployments, typically in clients run or sandboxed Java applets (in Java SE 8), that load and run untrusted code (e.g., code that comes from the internet) This vulnerability can also be exploited by using APIs in the specified Component, e.g., through a web service whi Score 3.7 (Availability impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:N/S:U/C:N/I:N/A:L).
CVE-2020-27216	In Eclipse Jetty versions 1.0 thru 9.4.32.v20200930, 10.0.0.alpha1 thru 10.0.0.beta2, and 11.0.0.alpha1 thru 11.0.0. temporary directory is shared between all users on that system. A collocated user can observe the process of creatin temporary directory and race to complete the creation of the temporary subdirectory. If the attacker wins the race th to the subdirectory used to unpack web applications, including their WEB-INF/lib jar files and JSP files. If any cod directory, this can lead to a local privilege escalation vulnerability.
CVE-2020-27218	In Eclipse Jetty version 9.4.0.RC0 to 9.4.34.v20201102, 10.0.0.alpha0 to 10.0.0.beta2, and 11.0.0.alpha0 to 11.0.0. enabled and requests from different clients are multiplexed onto a single connection, and if an attacker can send a r but not consumed by the application, then a subsequent request on the same connection will see that body prepende data but may inject data into the body of the subsequent request.
CVE-2020-27223	In Eclipse Jetty 9.4.6.v20170531 to 9.4.36.v20210114 (inclusive), 10.0.0, and 11.0.0 when Jetty handles a request large number of ,Äúquality,Äù (i.e. q) parameters, the server may enter a denial of service (DoS) state due to high resulting in minutes of CPU time exhausted processing those quality values.
CVE-2020-2754	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Scripting). Supported ver 11.0.6 and 14; Java SE Embedded: 8u241. Difficult to exploit vulnerability allows unauthenticated attacker with ne compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized ability DOS) of Java SE, Java SE Embedded. Note: Applies to client and server deployment of Java. This vulnerability ca Start applications and sandboxed Java applets. It can also be exploited by supplying data to APIs in the specified C Start applications or sandboxed Java applets, such as through a web service. CVSS 3.0 Base Score 3.7 (Availability AC:H/PR:N/UI:N/S:U/C:N/I:N/A:L).
CVE-2020-2755	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Scripting). Supported ver 11.0.6 and 14; Java SE Embedded: 8u241. Difficult to exploit vulnerability allows unauthenticated attacker with ne compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized ability DOS) of Java SE, Java SE Embedded. Note: Applies to client and server deployment of Java. This vulnerability ca Start applications and sandboxed Java applets. It can also be exploited by supplying data to APIs in the specified C Start applications or sandboxed Java applets, such as through a web service. CVSS 3.0 Base Score 3.7 (Availability AC:H/PR:N/UI:N/S:U/C:N/I:N/A:L).

CVE-2020-2756	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Serialization). Supported 8u241, 11.0.6 and 14; Java SE Embedded: 8u241. Difficult to exploit vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized abil DOS) of Java SE, Java SE Embedded. Note: Applies to client and server deployment of Java. This vulnerability ca Start applications and sandboxed Java applets. It can also be exploited by supplying data to APIs in the specified C Start applications or sandboxed Java applets, such as through a web service. CVSS 3.0 Base Score 3.7 (Availability AC:H/PR:N/UI:N/S:U/C:N/I:N/A:L).
CVE-2020-2757	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Serialization). Supported 8u241, 11.0.6 and 14; Java SE Embedded: 8u241. Difficult to exploit vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized abil DOS) of Java SE, Java SE Embedded. Note: Applies to client and server deployment of Java. This vulnerability ca Start applications and sandboxed Java applets. It can also be exploited by supplying data to APIs in the specified C Start applications or sandboxed Java applets, such as through a web service. CVSS 3.0 Base Score 3.7 (Availability AC:H/PR:N/UI:N/S:U/C:N/I:N/A:L).
CVE-2020-2773	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Security). Supported vers 8u241, 11.0.6 and 14; Java SE Embedded: 8u241. Difficult to exploit vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized abil DOS) of Java SE, Java SE Embedded. Note: Applies to client and server deployment of Java. This vulnerability ca Start applications and sandboxed Java applets. It can also be exploited by supplying data to APIs in the specified C Start applications or sandboxed Java applets, such as through a web service. CVSS 3.0 Base Score 3.7 (Availability AC:H/PR:N/UI:N/S:U/C:N/I:N/A:L).
CVE-2020-2781	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: JSSE). Supported version 11.0.6 and 14; Java SE Embedded: 8u241. Easily exploitable vulnerability allows unauthenticated attacker with ne SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized ability to cause a partial Java SE Embedded. Note: Applies to client and server deployment of Java. This vulnerability can be exploited thro and sandboxed Java applets. It can also be exploited by supplying data to APIs in the specified Component without or sandboxed Java applets, such as through a web service. CVSS 3.0 Base Score 5.3 (Availability impacts). CVSS S:U/C:N/I:N/A:L).
CVE-2020-2800	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Lightweight HTTP Serve Java SE: 7u251, 8u241, 11.0.6 and 14; Java SE Embedded: 8u241. Difficult to exploit vulnerability allows unauthe multiple protocols to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in to some of Java SE, Java SE Embedded accessible data as well as unauthorized read access to a subset of Java SE, This vulnerability can only be exploited by supplying data to APIs in the specified Component without using Untru Java applets, such as through a web service. CVSS 3.0 Base Score 4.8 (Confidentiality and Integrity impacts). CVS UI:N/S:U/C:L/I:L/A:N).
CVE-2020-2803	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Libraries). Supported vers 8u241, 11.0.6 and 14; Java SE Embedded: 8u241. Difficult to exploit vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. Successful attacks require human interaction from a person other than in Java SE, Java SE Embedded, attacks may significantly impact additional products. Successful attacks of this vul Java SE Embedded. Note: This vulnerability applies to Java deployments, typically in clients running sandboxed Ja applets, that load and run untrusted code (e.g., code that comes from the internet) and rely on the Java sandbox for Java deployments, typically in servers, that load and run only trusted code (e.g., code installed by an administrator) Integrity and Availability impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:R/S:C/C:H/I:H/A:H).
CVE-2020-2805	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Libraries). Supported vers 8u241, 11.0.6 and 14; Java SE Embedded: 8u241. Difficult to exploit vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. Successful attacks require human interaction from a person other than in Java SE, Java SE Embedded, attacks may significantly impact additional products. Successful attacks of this vul Java SE Embedded. Note: This vulnerability applies to Java deployments, typically in clients running sandboxed Ja applets, that load and run untrusted code (e.g., code that comes from the internet) and rely on the Java sandbox for Java deployments, typically in servers, that load and run only trusted code (e.g., code installed by an administrator) Integrity and Availability impacts). CVSS Vector: (CVSS:3.0/AV:N/AC:H/PR:N/UI:R/S:C/C:H/I:H/A:H).
CVE-2020-28196	MIT Kerberos 5 (aka krb5) before 1.17.2 and 1.18.x before 1.18.3 allows unbounded recursion via an ASN.1-encod asn.1/asn1_encode.c support for BER indefinite lengths lacks a recursion limit.
CVE-2020-2830	Vulnerability in the Java SE, Java SE Embedded product of Oracle Java SE (component: Concurrency). Supported 8u241, 11.0.6 and 14; Java SE Embedded: 8u241. Easily exploitable vulnerability allows unauthenticated attacker to compromise Java SE, Java SE Embedded. Successful attacks of this vulnerability can result in unauthorized abil DOS) of Java SE, Java SE Embedded. Note: Applies to client and server deployment of Java. This vulnerability ca Start applications and sandboxed Java applets. It can also be exploited by supplying data to APIs in the specified C Start applications or sandboxed Java applets, such as through a web service. CVSS 3.0 Base Score 5.3 (Availability AC:L/PR:N/UI:N/S:U/C:N/I:N/A:L).
CVE-2020-28362	Go before 1.14.12 and 1.15.x before 1.15.4 allows Denial of Service.

CVE-2020-28366	Code injection in the go command with cgo before Go 1.14.12 and Go 1.15.5 allows arbitrary code execution at bu name in a linked object file.
CVE-2020-28367	Code injection in the go command with cgo before Go 1.14.12 and Go 1.15.5 allows arbitrary code execution at bu a #cgo directive.
CVE-2020-29361	An issue was discovered in p11-kit 0.21.1 through 0.23.21. Multiple integer overflows have been discovered in the the p11-kit list command, where overflow checks are missing before calling realloc or calloc.
CVE-2020-29362	An issue was discovered in p11-kit 0.21.1 through 0.23.21. A heap-based buffer over-read has been discovered in t remote commands and the client library. When the remote entity supplies a byte array through a serialized PKCS#1 allow the reading of up to 4 bytes of memory past the heap allocation.
CVE-2020-35448	An issue was discovered in the Binary File Descriptor (BFD) library (aka libbfd), as distributed in GNU Binutils 2. occur in bfd_getl_signed_32 in libbfd.c because sh_entsize is not validated in _bfd_elf_slurp_secondary_reloc_sec
CVE-2020-35525	In SQlite 3.31.1, a potential null pointer derreference was found in the INTERSEC query processing.
CVE-2020-35527	In SQLite 3.31.1, there is an out of bounds access problem through ALTER TABLE for views that have a nested F
CVE-2020-36221	An integer underflow was discovered in OpenLDAP before 2.4.57 leading to slapd crashes in the Certificate Exact service (schema_init.c serialNumberAndIssuerCheck).
CVE-2020-36222	A flaw was discovered in OpenLDAP before 2.4.57 leading to an assertion failure in slapd in the saslAuthzTo valic
CVE-2020-36223	A flaw was discovered in OpenLDAP before 2.4.57 leading to a slapd crash in the Values Return Filter control han free and out-of-bounds read).
CVE-2020-36224	A flaw was discovered in OpenLDAP before 2.4.57 leading to an invalid pointer free and slapd crash in the saslAu service.
CVE-2020-36225	A flaw was discovered in OpenLDAP before 2.4.57 leading to a double free and slapd crash in the saslAuthzTo pro
CVE-2020-36226	A flaw was discovered in OpenLDAP before 2.4.57 leading to a memch->bv_len miscalculation and slapd crash in denial of service.
CVE-2020-36227	A flaw was discovered in OpenLDAP before 2.4.57 leading to an infinite loop in slapd with the cancel_extop Canc
CVE-2020-36228	An integer underflow was discovered in OpenLDAP before 2.4.57 leading to a slapd crash in the Certificate List E of service.
CVE-2020-36229	A flaw was discovered in ldap_X509dn2bv in OpenLDAP before 2.4.57 leading to a slapd crash in the X.509 DN p service.
CVE-2020-36230	A flaw was discovered in OpenLDAP before 2.4.57 leading in an assertion failure in slapd in the X.509 DN parsing denial of service.
CVE-2020-36694	An issue was discovered in netfilter in the Linux kernel before 5.10. There can be a use-after-free in the packet pro sequence count is mishandled during concurrent iptables rules replacement. This could be exploited with the CAP_ namespace. NOTE: cc00bca was reverted in 5.12.
CVE-2020-8231	Due to use of a dangling pointer, libcurl 7.29.0 through 7.71.1 can use the wrong connection when sending data.
CVE-2020-8284	A malicious server can use the FTP PASV response to trick curl 7.73.0 and earlier into connecting back to a given make curl extract information about services that are otherwise private and not disclosed, for example doing port sc
CVE-2020-8285	curl 7.21.0 to and including 7.73.0 is vulnerable to uncontrolled recursion due to a stack overflow issue in FTP wil
CVE-2020-8991	vg_lookup in daemons/lvmetad/lvmetad-core.c in LVM2 2.02 mismanages memory, leading to an lvmetad memory NOTE: RedHat disputes CVE-2020-8991 as not being a vulnerability since there,Âôs no apparent route to either pr through the bug
CVE-2021-20197	There is an open race window when writing output in the following utilities in GNU binutils version 2.35 and earli utilities are run as a privileged user (presumably as part of a script updating binaries across different users), an unpr getting ownership of arbitrary files through a symlink.
CVE-2021-20229	A flaw was found in PostgreSQL in versions before 13.2. This flaw allows a user with SELECT privilege on one c columns of the table. The highest threat from this vulnerability is to confidentiality.
CVE-2021-20266	A flaw was found in RPM's hdrblobInit() in lib/header.c. This flaw allows an attacker who can modify the rpmdb t threat from this vulnerability is to system availability.
CVE-2021-20294	A flaw was found in binutils readelf 2.35 program. An attacker who is able to convince a victim using readelf to re overflow, out-of-bounds write of arbitrary data supplied by the attacker. The highest impact of this flaw is to confi

CVE-2021-22876	curl 7.1.1 to and including 7.75.0 is vulnerable to an "Exposure of Private Personal Information to an Unauthorized Referer: header. libcurl does not strip off user credentials from the URL when automatically populating the Referer HTTP requests, and therefore risks leaking sensitive data to the server that is the target of the second HTTP request
CVE-2021-22898	curl 7.7 through 7.76.1 suffers from an information disclosure when the `-t` command line option, known as `CURL used to send variable=content pairs to TELNET servers. Due to a flaw in the option parser for sending NEW_ENV uninitialized data from a stack based buffer to the server, resulting in potentially revealing sensitive internal inform protocol.
CVE-2021-22924	libcurl keeps previously used connections in a connection pool for subsequenttransfers to reuse, if one of them mat config matching function did not take 'issuercert' into account and it compared the involved paths *case insensitive connections.File paths are, or can be, case sensitive on many systems but not all, and caneven vary depending on us include the 'issuer cert' which a transfer can setto qualify how to verify the server certificate.
CVE-2021-22925	curl supports the `-t` command line option, known as `CURLOPT_TELNETOPTIONS`in libcurl. This rarely used toTELNET servers.Due to flaw in the option parser for sending `NEW_ENV` variables, libcurlcould be made to pa buffer to theserver. Therefore potentially revealing sensitive internal information to theserver using a clear-text net did not call and use sscanf() correctly whenparsing the string provided by the application.
CVE-2021-22946	A user can tell curl >= 7.20.0 and <= 7.78.0 to require a successful upgrade to TLS when speaking to an IMAP, PO command line or`CURLOPT_USE_SSL` set to `CURLUSESSL_CONTROL` or `CURLUSESSL_ALL` withlibcu the server would return a properly crafted but perfectly legitimate response.This flaw would then make curl silently contrary to the instructions and expectations, exposing possibly sensitive data in clear text over the network.
CVE-2021-22947	When curl >= 7.20.0 and <= 7.78.0 connects to an IMAP or POP3 server to retrieve data using STARTTLS to upgr and send back multiple responses at once that curl caches. curl would then upgrade to TLS but not flush the in-queu using and trustingthe responses it got *before* the TLS handshake as if they were authenticated.Using this flaw, it inject the fake responses, then pass-through the TLS traffic from the legitimate server and trick curl into sending da injected data comes from the TLS-protected server.
CVE-2021-23214	When the server is configured to use trust authentication with a clientcert requirement or to use cert authentication, arbitrary SQL queries when a connection is first established, despite the use of SSL certificate verification and encr
CVE-2021-23222	A man-in-the-middle attacker can inject false responses to the client's first few queries, despite the use of SSL certi
CVE-2021-27212	In OpenLDAP through 2.4.57 and 2.5.x through 2.5.1alpha, an assertion failure in slapd can occur in the issuerAnd packet, resulting in a denial of service (daemon exit) via a short timestamp. This is related to schema_init.c and che
CVE-2021-27218	An issue was discovered in GNOME GLib before 2.66.7 and 2.67.x before 2.67.4. If g_byte_array_new_take() was bit platform, the length would be truncated modulo 2**32, causing unintended length truncation.
CVE-2021-28153	An issue was discovered in GNOME GLib before 2.66.8. When g_file_replace() is used with G_FILE_CREATE_R that is a dangling symlink, it incorrectly also creates the target of the symlink as an empty file, which could conceiv attacker-controlled. (If the path is a symlink to a file that already exists, then the contents of that file correctly rema
CVE-2021-28165	In Eclipse Jetty 7.2.2 to 9.4.38, 10.0.0.alpha0 to 10.0.1, and 11.0.0.alpha0 to 11.0.1, CPU usage can reach 100% up
CVE-2021-28831	decompress_gunzip.c in BusyBox through 1.32.1 mishandles the error bit on the huft_build result pointer, with a re malformed gzip data.
CVE-2021-3200	Buffer overflow vulnerability in libsolv 2020-12-13 via the Solver * testcase_read(Pool *pool, FILE *fp, const cha *resultflagsp function at src/testcase.c: line 2334, which could cause a denial of service
CVE-2021-32028	A flaw was found in postgresql. Using an INSERT ... ON CONFLICT ... DO UPDATE command on a purpose-cra could read arbitrary bytes of server memory. The highest threat from this vulnerability is to data confidentiality.
CVE-2021-32029	A flaw was found in postgresql. Using an UPDATE ... RETURNING command on a purpose-crafted table, an auth bytes of server memory. The highest threat from this vulnerability is to data confidentiality.
CVE-2021-33061	Insufficient control flow management for the Intel(R) 82599 Ethernet Controllers and Adapters may allow an authe service via local access.
CVE-2021-33560	Libgcrypt before 1.8.8 and 1.9.x before 1.9.3 mishandles ElGamal encryption because it lacks exponent blinding to mpi_powm, and the window size is not chosen appropriately. This, for example, affects use of ElGamal in OpenPG
CVE-2021-33574	The mq_notify function in the GNU C Library (aka glibc) versions 2.32 and 2.33 has a use-after-free. It may use th through its struct sigevent parameter) after it has been freed by the caller, leading to a denial of service (application
CVE-2021-33621	The cgi gem before 0.1.0.2, 0.2.x before 0.2.2, and 0.3.x before 0.3.5 for Ruby allows HTTP response splitting. Th user input either to generate an HTTP response or to create a CGI::Cookie object.
CVE-2021-33631	Integer Overflow or Wraparound vulnerability in openEuler kernel on Linux (filesystem modules) allows Forced In kernel: from 4.19.90 before 4.19.90-2401.3, from 5.10.0-60.18.0 before 5.10.0-183.0.0.
CVE-2021-33928	Buffer overflow vulnerability in function pool_installable in src/repo.h in libsolv before 0.7.17 allows attackers to

CVE-2021-33929	Buffer overflow vulnerability in function pool_disabled_solvable in src/repo.h in libsolv before 0.7.17 allows attac
CVE-2021-33930	Buffer overflow vulnerability in function pool_installable_whatprovides in src/repo.h in libsolv before 0.7.17 allow
CVE-2021-33938	Buffer overflow vulnerability in function prune_to_recommended in src/policy.c in libsolv before 0.7.17 allows att
CVE-2021-3421	A flaw was found in the RPM package in the read functionality. This flaw allows an attacker who can convince a v or compromise an RPM repository, to cause RPM database corruption. The highest threat from this vulnerability is versions before 4.17.0-alpha.
CVE-2021-3516	There's a flaw in libxml2's xmllint in versions before 2.9.11. An attacker who is able to submit a crafted file to be p free. The greatest impact of this flaw is to confidentiality, integrity, and availability.
CVE-2021-3517	There is a flaw in the xml entity encoding functionality of libxml2 in versions before 2.9.11. An attacker who is ab by an application linked with the affected functionality of libxml2 could trigger an out-of-bounds read. The most li availability, with some potential impact to confidentiality and integrity if an attacker is able to use memory informa
CVE-2021-3518	There's a flaw in libxml2 in versions before 2.9.11. An attacker who is able to submit a crafted file to be processed trigger a use-after-free. The greatest impact from this flaw is to confidentiality, integrity, and availability.
CVE-2021-3520	There's a flaw in lz4. An attacker who submits a crafted file to an application linked with lz4 may be able to trigger memmove() on a negative size argument, causing an out-of-bounds write and/or a crash. The greatest impact of thi impact to confidentiality and integrity as well.
CVE-2021-3521	There is a flaw in RPM's signature functionality. OpenPGP subkeys are associated with a primary key via a "bindi binding signature of subkeys prior to importing them. If an attacker is able to add or socially engineer another party public key, RPM could wrongly trust a malicious signature. The greatest impact of this flaw is to data integrity. To compromise an RPM repository or convince an administrator to install an untrusted RPM or public key. It is strong public keys from trusted sources.
CVE-2021-3537	A vulnerability found in libxml2 in versions before 2.9.11 shows that it did not propagate errors while parsing XM dereference. If an untrusted XML document was parsed in recovery mode and post-validated, the flaw could be use from this vulnerability is to system availability.
CVE-2021-3541	A flaw was found in libxml2. Exponential entity expansion attack its possible bypassing all existing protection mec
CVE-2021-35937	A race condition vulnerability was found in rpm. A local unprivileged user could use this flaw to bypass the checks CVE-2017-7500 and CVE-2017-7501, potentially gaining root privileges. The highest threat from this vulnerabilit well as system availability.
CVE-2021-35938	A symbolic link issue was found in rpm. It occurs when rpm sets the desired permissions and credentials after insta use this flaw to exchange the original file with a symbolic link to a security-critical file and escalate their privileges vulnerability is to data confidentiality and integrity as well as system availability.
CVE-2021-35939	It was found that the fix for CVE-2017-7500 and CVE-2017-7501 was incomplete: the check was only implemente created. A local unprivileged user who owns another ancestor directory could potentially use this flaw to gain root vulnerability is to data confidentiality and integrity as well as system availability.
CVE-2021-3601	Rejected reason: DO NOT USE THIS CANDIDATE NUMBER. ConsultIDs: none. Reason: This candidate was w class this issue as a security vulnerability. The trusted CA store should not contain anything that the user does not t github.com/openssl/openssl/issues/5236#issuecomment-119646061
CVE-2021-36222	ec_verify in kdc/kdc_preauth_ec.c in the Key Distribution Center (KDC) in MIT Kerberos 5 (aka krb5) before 1.18 attackers to cause a NULL pointer dereference and daemon crash. This occurs because a return value is not properl
CVE-2021-3669	A flaw was found in the Linux kernel. Measuring usage of the shared memory does not scale with large shared mer resource exhaustion and DoS.
CVE-2021-3671	A null pointer de-reference was found in the way samba kerberos server handled missing sname in TGS-REQ (Tic authenticated user could use this flaw to crash the samba server.
CVE-2021-37322	GCC c++filt v2.26 was discovered to contain a use-after-free vulnerability via the component cplus-dem.c.
CVE-2021-37600	An integer overflow in util-linux through 2.37.1 can potentially cause a buffer overflow if an attacker were able to large number in the /proc/sysvipc/sem file.  **Note:** This is unexploitable in GNU C Library environments, and possibly in all realistic environments.
CVE-2021-3800	A flaw was found in glib before version 2.63.6. Due to random charset alias, pkexec can leak content from files ow under the right condition.
CVE-2021-38185	GNU cpio through 2.13 allows attackers to execute arbitrary code via a crafted pattern file, because of a dstring.c d of-bounds heap write. NOTE: it is unclear whether there are common cases where the pattern file, associated with t

CVE-2021-3847	An unauthorized access to the execution of the setuid file with capabilities flaw in the Linux kernel OverlayFS subs capable file from a nosuid mount into another mount. A local user could use this flaw to escalate their privileges or
CVE-2021-39686	In several functions of binder.c, there is a possible way to represent the wrong domain to SELinux due to a race co privilege with no additional execution privileges needed. User interaction is not needed for exploitation.Product: A A-200688826References: Upstream kernel
CVE-2021-4023	A flaw was found in the io-workqueue implementation in the Linux kernel versions prior to 5.15-rc1. The kernel ca operation triggers the submission of new io-uring operations during a shortage of free space. This flaw allows a loc requests to possibly crash the system.
CVE-2021-40528	The ElGamal implementation in Libgcrypt before 1.9.4 allows plaintext recovery because, during interaction betwe dangerous combination of the prime defined by the receiver's public key, the generator defined by the receiver's pu exponents can lead to a cross-configuration attack against OpenPGP.
CVE-2021-4149	A vulnerability was found in btrfs_alloc_tree_b in fs/btrfs/extent-tree.c in the Linux kernel due to an improper lock local privilege may cause a denial of service (DOS) due to a deadlock problem.
CVE-2021-4204	An out-of-bounds (OOB) memory access flaw was found in the Linux kernel's eBPF due to an Improper Input Vali a special privilege to crash the system or leak internal information.
CVE-2021-42374	An out-of-bounds heap read in Busybox's unlzma applet leads to information leak and denial of service when craft This can be triggered by any applet/format that
CVE-2021-42376	A NULL pointer dereference in Busybox's hush applet leads to denial of service when processing a crafted shell co delimiter character. This may be used for DoS under very rare conditions of filtered command input.
CVE-2021-42378	A use-after-free in Busybox's awk applet leads to denial of service and possibly code execution when processing a
CVE-2021-42379	A use-after-free in Busybox's awk applet leads to denial of service and possibly code execution when processing a function
CVE-2021-42380	A use-after-free in Busybox's awk applet leads to denial of service and possibly code execution when processing a
CVE-2021-42381	A use-after-free in Busybox's awk applet leads to denial of service and possibly code execution when processing a
CVE-2021-42382	A use-after-free in Busybox's awk applet leads to denial of service and possibly code execution when processing a
CVE-2021-42384	A use-after-free in Busybox's awk applet leads to denial of service and possibly code execution when processing a function
CVE-2021-42385	A use-after-free in Busybox's awk applet leads to denial of service and possibly code execution when processing a
CVE-2021-42386	A use-after-free in Busybox's awk applet leads to denial of service and possibly code execution when processing a
CVE-2021-42574	An issue was discovered in the Bidirectional Algorithm in the Unicode Specification through 14.0. It permits the vi sequences, which can be used to craft source code that renders different logic than the logical ordering of tokens in Adversaries can leverage this to encode source code for compilers accepting Unicode such that targeted vulnerabili reviewers. **Note:** The Unicode Consortium offers the following alternative approach to presenting this concern. An text that can affect applications that implement support for The Unicode Standard and the Unicode Bidire to text display behavior when text includes left-to-right and right-to-left characters, the visual order of tol order. Additionally, control characters needed to fully support the requirements of bidirectional text can f Unless mitigated, an adversary could craft source code such that the ordering of tokens perceived by hum processed by a compiler/interpreter/etc. The Unicode Consortium has documented this class of vulnerabi Report #36, Unicode Security Considerations. The Unicode Consortium also provides guidance on mitig Technical Standard #39, Unicode Security Mechanisms, and in Unicode Standard Annex #31, Unicode Id specification allows applications to tailor the implementation in ways that can mitigate misleading visual Unicode Standard Annex #9, Unicode Bidirectional Algorithm.
CVE-2021-43519	Stack overflow in lua_resume of ldo.c in Lua Interpreter 5.1.0~5.4.4 allows attackers to perform a Denial of Servic
CVE-2021-43618	GNU Multiple Precision Arithmetic Library (GMP) through 6.2.1 has an mpz/inp_raw.c integer overflow and resul to a segmentation fault on 32-bit platforms.
CVE-2021-43797	Netty is an asynchronous event-driven network application framework for rapid development of maintainable high Netty prior to version 4.1.71.Final skips control chars when they are present at the beginning / end of the header na allowed by the spec and could lead to HTTP request smuggling. Failing to do the validation might cause netty to "s to another remote system when used as proxy. This remote system can't see the invalid usage anymore, and therefo should upgrade to version 4.1.71.Final.
CVE-2021-44568	Two heap-overflow vulnerabilities exist in openSUSE/libsolv libsolv through 13 Dec 2020 in the decisionmap vari src/solver.c (line 1940 & line 1995), which could cause a remote Denial of Service.

CVE-2021-44758	Heimdal before 7.7.1 allows attackers to cause a NULL pointer dereference in a SPNEGO acceptor via a preferred_ nonzero initial_response value to send_accept.
CVE-2021-44758	Heimdal before 7.7.1 allows attackers to cause a NULL pointer dereference in a SPNEGO acceptor via a preferred_ nonzero initial_response value to send_accept.
CVE-2021-46195	GCC v12.0 was discovered to contain an uncontrolled recursion via the component libiberty/rust-demangle.c. This Denial of Service (DoS) by consuming excessive CPU and memory resources.
CVE-2021-46822	The PPM reader in libjpeg-turbo through 2.0.90 mishandles use of tjLoadImage for loading a 16-bit binary PPM fil binary PGM file into an RGB buffer. This is related to a heap-based buffer overflow in the get_word_rgb_row func
CVE-2021-46848	GNU Libtasn1 before 4.19.0 has an ETYPE_OK off-by-one array size check that affects asn1_encode_simple_der.
CVE-2021-47028	In the Linux kernel, the following vulnerability has been resolved: mt76: mt7915: fix txrate reporting Properly chec reporting. [1215.161863] Call trace: [1215.164307] cfg80211_calculate_bitrate+0x124/0x200 [cfg80211] [1215. +0x80/0xc0 [mac80211] [1215.175624] ieee80211_tx_status_ext+0x508/0x838 [mac80211] [1215.181190] mt79 [mt7915e] [1215.186580] mt7915_mac_tx_free+0x324/0x7c0 [mt7915e] [1215.191623] mt7915_queue_rx_skb+ mt76_dma_cleanup+0x7b0/0x11d0 [mt76] [1215.201276] __napi_poll+0x38/0xf8 [1215.204668] napi_workfn+0 +0x1fc/0x390 [1215.212062] worker_thread+0x48/0x4d0 [1215.215715] kthread+0x120/0x128 [1215.218935] r
CVE-2021-47037	In the Linux kernel, the following vulnerability has been resolved: ASoC: q6afe-clocks: fix reprobing of the driver example if the APR services are restarted after the firmware crash. However currently Q6afe-clocks driver will oop _probe call. Rewrite the driver to fill the clock data at runtime rather than using big static array of clocks.
CVE-2021-47105	In the Linux kernel, the following vulnerability has been resolved: ice: xsk: return xsk buffers back to pool when cl xdp_buff pointer in the internal SW ring but we never give it back to the xsk buffer pool. This means that buffers c used again. Add missing xsk_buff_free() call to the routine that is supposed to clean the entries that are left in the r used by other sockets. Also, only go through the space that is actually left to be cleaned instead of a whole ring.
CVE-2021-47131	In the Linux kernel, the following vulnerability has been resolved: net/tls: Fix use-after-free after the TLS device g active TLS offload goes down, tls_device_down is called to stop the offload and tear down the TLS context. Howe points to the TLS context, which is now deallocated. If a netdev goes up, while the connection is still active, and th retransmissions, it will lead to a use-after-free of the TLS context. This commit addresses this bug by keeping the c and implements the necessary fallbacks, so that the connection can resume in software (non-offloaded) kTLS mode encrypt all packets. The RX side already has all the necessary fallbacks, because receiving non-decrypted packets i is to block resync requests, which are normally produced after receiving non-decrypted packets. The necessary syn teardown: first the fallbacks are deployed, then the driver resources are released (it used to be possible to have a tls called TLS_RX_DEV_DEGRADED is added to indicate the fallback mode. It's used to skip the RX resync logic c objects may be released (for example, resync_async, which is allocated and freed by the driver).
CVE-2021-47193	In the Linux kernel, the following vulnerability has been resolved: scsi: pm80xx: Fix memory leak during rmmod I This would lead to memory leak during driver removal. Properly free memory when the module is removed.
CVE-2021-47200	In the Linux kernel, the following vulnerability has been resolved: drm/prime: Fix use after free in mmap with drm drops a reference to the gem object on success. If the gem object's refcount == 1 on entry to drm_gem_prime_mma and the subsequent drm_gem_object_get() will be a UAF. Fix by grabbing a reference before calling the mmap hel reference dropping was adding in commit 9786b65bc61ac ("drm/ttm: fix mmap refcounting"): "For that to work pr drm_gem_ttm_mmap() must be moved so it happens before calling obj->funcs->mmap(), otherwise the gem refcou
CVE-2021-47382	In the Linux kernel, the following vulnerability has been resolved: s390/qeth: fix deadlock during failing recovery deadlock during recovery") removed taking discipline_mutex inside qeth_do_reset(), fixing potential deadlocks. A still takes discipline_mutex and thus has the original deadlock potential. Intermittent deadlocks were seen when a q causing a race between qeth_do_reset and ccwgroup_remove. Call qeth_set_offline() directly in the qeth_do_reset(ccwgroup_set_offline(), without taking discipline_mutex.
CVE-2021-47484	In the Linux kernel, the following vulnerability has been resolved: octeontx2-af: Fix possible null pointer dereferen dereference in files "rvu_debugfs.c" and "rvu_nix.c"
CVE-2021-47489	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: Fix even more out of bound writes fixed by: commit f23750b5b3d98653b31d4469592935ef6364ad67 Author: Thelford Williams <tdwilliamsiv@gma -0400 drm/amdgpu: fix out of bounds write but amdgpu_dm_debugfs.c contains more of the same issue so fix the r dp_max_bpc_write (Harry Wentland)
CVE-2021-47622	In the Linux kernel, the following vulnerability has been resolved: scsi: ufs: Fix a deadlock in the error handler The on a test setup: - All tags allocated - The SCSI error handler calls ufshcd_eh_host_reset_handler() - ufshcd_eh_hos ufshcd_err_handler() - ufshcd_err_handler() locks up as follows: Workqueue: ufs_eh_wq_0 ufshcd_err_handler.cfi __schedule+0x6cc/0xa94 schedule+0x12c/0x298 blk_mq_get_tag+0x210/0x480 __blk_mq_alloc_request+0x1c8/0 ufshcd_exec_dev_cmd+0x68/0x640 ufshcd_verify_dev_init+0x68/0x35c ufshcd_probe_hba+0x12c/0x1cb8 ufshcd ufshcd_reset_and_restore+0xd0/0x354 ufshcd_err_handler+0x408/0xc58 process_one_work+0x24c/0x66c worker_ ret_from_fork+0x10/0x30 Fix this lockup by making ufshcd_exec_dev_cmd() allocate a reserved request.
CVE-2022-0500	A flaw was found in unrestricted eBPF usage by the BPF_BTF_LOAD, leading to a possible out-of-bounds memor subsystem due to the way a user loads BTF. This flaw allows a local user to crash or escalate their privileges on the

CVE-2022-0536	Improper Removal of Sensitive Information Before Storage or Transfer in NPM follow-redirects prior to 1.14.8.
CVE-2022-1205	A NULL pointer dereference flaw was found in the Linux kernel,Äôs Amateur Radio AX.25 protocol functionality This flaw allows a local user to crash the system.
CVE-2022-1271	An arbitrary file write vulnerability was found in GNU gzip's zgrep utility. When zgrep is applied on the attacker's name), this can overwrite an attacker's content to an arbitrary attacker-selected file. This flaw occurs due to insuffic with two or more newlines where selected content and the target file names are embedded in crafted multi-line file privileged attacker to force zgrep to write arbitrary files on the system.
CVE-2022-1586	An out-of-bounds read vulnerability was discovered in the PCRE2 library in the compile_xclass_matchingpath() fu involves a unicode property matching issue in JIT-compiled regular expressions. The issue occurs because the char within JIT.
CVE-2022-1587	An out-of-bounds read vulnerability was discovered in the PCRE2 library in the get_recurse_data_length() functio affects recursions in JIT-compiled regular expressions caused by duplicate data transfers.
CVE-2022-1664	Dpkg::Source::Archive in dpkg, the Debian package management system, before version 1.21.8, 1.20.10, 1.19.8, 1. vulnerability. When extracting untrusted source packages in v2 and v3 source package formats that include a debia directory traversal situations on specially crafted orig.tar and debian.tar tarballs.
CVE-2022-2196	A regression exists in the Linux Kernel within KVM: nVMX that allowed for speculative execution attacks.¬†L2 c L1 thinking it doesn't need retpolines or IBPB¬†after running L2 due to KVM (L0) advertising eIBRS support to L execute code on an indirect branch on the host machine. We recommend upgrading to Kernel 6.2 or past commit¬†
CVE-2022-23476	Nokogiri is an open source XML and HTML library for the Ruby programming language. Nokogiri `1.13.8` and `1 `xmlTextReaderExpand` in the method `Nokogiri::XML::Reader#attribute_hash`. This can lead to a null pointer ex For applications using `XML::Reader` to parse untrusted inputs, this may potentially be a vector for a denial of serv Nokogiri `>= 1.13.10`. Users may be able to search their code for calls to either `XML::Reader#attributes` or `XM they are affected.
CVE-2022-23491	Certifi is a curated collection of Root Certificates for validating the trustworthiness of SSL certificates while verify 2022.12.07 removes root certificates from "TrustCor" from the root store. These are in the process of being remove root certificates are being removed pursuant to an investigation prompted by media reporting that TrustCor's owner spyware. Conclusions of Mozilla's investigation can be found in the linked google group discussion.
CVE-2022-24823	Netty is an open-source, asynchronous event-driven network application framework. The package `io.netty:netty-c contains an insufficient fix for CVE-2021-21290. When Netty's multipart decoders are used local information discl temporary directory if temporary storing uploads on the disk is enabled. This only impacts applications running on vulnerability impacts code running on Unix-like systems, and very old versions of Mac OSX and Windows as they between all users. Version 4.1.77.Final contains a patch for this vulnerability. As a workaround, specify one's own DefaultHttpDataFactory.setBaseDir(...) to set the directory to something that is only readable by the current user.
CVE-2022-2509	A vulnerability found in gnutls. This security flaw happens because of a double free error occurs during verificatio function.
CVE-2022-25265	In the Linux kernel through 5.16.10, certain binary files may have the exec-all attribute if they were built in approx kernel 2.4.20). This can cause execution of bytes located in supposedly non-executable regions of a file.
CVE-2022-25313	In Expat (aka libexpat) before 2.4.5, an attacker can trigger stack exhaustion in build_model via a large nesting dep
CVE-2022-2625	A vulnerability was found in PostgreSQL. This attack requires permission to create non-temporary objects in at lea an administrator to create or update an affected extension in that schema, and the ability to lure or wait for a victim REPLACE or CREATE IF NOT EXISTS. Given all three prerequisites, this flaw allows an attacker to run arbitrary superuser.
CVE-2022-27672	When SMT is enabled, certain AMD processors may speculatively execute instructions using a target from the sibl potentially resulting in information disclosure.
CVE-2022-27774	An insufficiently protected credentials vulnerability exists in curl 4.9 to and include curl 7.82.0 are affected that co when follows HTTP(S) redirects is used with authentication could leak credentials to other services that exist on di
CVE-2022-27776	A insufficiently protected credentials vulnerability in fixed in curl 7.83.0 might leak authentication or cookie heade another port number.
CVE-2022-27778	A use of incorrectly resolved name vulnerability fixed in 7.83.1 might remove the wrong file when `--no-clobber` i
CVE-2022-27779	libcurl wrongly allows cookies to be set for Top Level Domains (TLDs) if thehost name is provided with a trailing cookies. curl's "cookie engine" can bebuilt with or without [Public Suffix List](https://publicsuffix.org/)awareness. rudimentary check exists to atleast prevent cookies from being set on TLDs. This check was broken if thehost nam allow arbitrary sites to set cookies that then would get sent to adifferent and unrelated site or domain.
CVE-2022-27780	The curl URL parser wrongly accepts percent-encoded URL separators like '/'when decoding the host name part of wrong host name when it is later retrieved.For example, a URL like `http://example.com%2F127.0.0.1/`, would be `http://example.com/127.0.0.1/`. This flawcan be used to circumvent filters, checks and more.

CVE-2022-27781	libcurl provides the `CURLOPT_CERTINFO` option to allow applications torequest details to be returned about a function, a malicious server could make libcurl built withNSS get stuck in a never-ending busy-loop when trying to
CVE-2022-27782	libcurl would reuse a previously created connection even when a TLS or SSHrelated option had been changed that previously used connections in a connection pool for subsequenttransfers to reuse if one of them matches the setup left out from the configuration match checks, making themmatch too easily.
CVE-2022-28321	The Linux-PAM package before 1.5.2-6.1 for openSUSE Tumbleweed allows authentication bypass for SSH login correctly restrict login if a user tries to connect from an IP address that is not resolvable via DNS. In such condition still get access. NOTE: the relevance of this issue is largely limited to openSUSE Tumbleweed and openSUSE Fac
CVE-2022-28391	BusyBox through 1.35.0 allows remote attackers to execute arbitrary code if netstat is used to print a DNS PTR rec Alternatively, the attacker could choose to change the terminal's colors.
CVE-2022-28948	An issue in the Unmarshal function in Go-Yaml v3 causes the program to crash when attempting to deserialize inva
CVE-2022-29155	In OpenLDAP 2.x before 2.5.12 and 2.6.x before 2.6.2, a SQL injection vulnerability exists in the experimental bac within an LDAP query. This can occur during an LDAP search operation when the search filter is processed, due to
CVE-2022-29824	In libxml2 before 2.9.14, several buffer handling functions in buf.c (xmlBuf*) and tree.c (xmlBuffer*) don't check out-of-bounds memory writes. Exploitation requires a victim to open a crafted, multi-gigabyte XML file. Other sof example libxslt through 1.1.35, is affected as well.
CVE-2022-30115	Using its HSTS support, curl can be instructed to use HTTPS directly insteadof using an insecure clear-text HTTP This mechanism could be bypassed if the host name in the given URL used atrailing dot while not using one when around - by having the trailing dot in the HSTS cache and *not* using thetrailing dot in the URL.
CVE-2022-30115	Using its HSTS support, curl can be instructed to use HTTPS directly insteadof using an insecure clear-text HTTP This mechanism could be bypassed if the host name in the given URL used atrailing dot while not using one when around - by having the trailing dot in the HSTS cache and *not* using thetrailing dot in the URL.
CVE-2022-3108	An issue was discovered in the Linux kernel through 5.16-rc6. kfd_parse_subtype_iolink in drivers/gpu/drm/amd/a value of kmemdup().
CVE-2022-3114	An issue was discovered in the Linux kernel through 5.16-rc6. imx_register_uart_clocks in drivers/clk/imx/clk.c la will cause the null pointer dereference.
CVE-2022-31159	The AWS SDK for Java enables Java developers to work with Amazon Web Services. A partial-path traversal issu method in the AWS S3 TransferManager component of the AWS SDK for Java v1 prior to version 1.12.261. Appli `destinationDirectory` argument, but S3 object keys are determined by the application that uploaded the objects. Th allows the caller to pass a filesystem object in the object key but contained an issue in the validation logic for the k could bypass the validation logic by including a UNIX double-dot in the bucket key. Under certain conditions, this directory from their S3 bucket that is one level up in the filesystem from their working directory. This issue,Äôs sc name prefix matches the destinationDirectory. E.g. for destination directory`/tmp/foo`, the actor can cause a downl bar`. If `com.amazonaws.services.s3.transfer.TransferManager::downloadDirectory` is used to download an untrus that bucket can be written outside of the intended destination directory. Version 1.12.261 contains a patch for this i `com.amazonaws.services.s3.transfer.TransferManager::downloadDirectory`, pass a `KeyFilter` that forbids `S3Ob return a string containing the substring `..` .
CVE-2022-3116	The Heimdal Software Kerberos 5 implementation is vulnerable to a null pointer dereferance. An attacker with net the vulnerable code path can cause the application to crash.
CVE-2022-31197	PostgreSQL JDBC Driver (PgJDBC for short) allows Java programs to connect to a PostgreSQL database using sta The PGJDBC implementation of the `java.sql.ResultRow.refreshRow()` method is not performing escaping of colu contains a statement terminator, e.g. `;`, could lead to SQL injection. This could lead to executing additional SQL c User applications that do not invoke the `ResultSet.refreshRow()` method are not impacted. User application that d underlying database that they are querying via their JDBC application may be under the control of an attacker. The into executing SQL against a table name who's column names would contain the malicious SQL and subsequently ResultSet. Note that the application's JDBC user and the schema owner need not be the same. A JDBC application database schemas owned by potentially malicious less-privileged users would be vulnerable. In that situation it may schema that causes the application to execute commands as the privileged user. Patched versions will be released a upgrade. There are no known workarounds for this issue.
CVE-2022-32208	When curl < 7.84.0 does FTP transfers secured by krb5, it handles message verification failures wrongly. This flaw attack to go unnoticed and even allows it to inject data to the client.

CVE-2022-3358	OpenSSL supports creating a custom cipher via the legacy EVP_CIPHER_meth_new() function and associated fun in OpenSSL 3.0 and application authors are instead encouraged to use the new provider mechanism in order to imp 3.0.0 to 3.0.5 incorrectly handle legacy custom ciphers passed to the EVP_EncryptInit_ex2(), EVP_DecryptInit_ex (as well as other similarly named encryption and decryption initialisation functions). Instead of using the custom ci an equivalent cipher from the available providers. An equivalent cipher is found based on the NID passed to EVP_ supposed to represent the unique NID for a given cipher. However it is possible for an application to incorrectly pa EVP_CIPHER_meth_new(). When NID_undef is used in this way the OpenSSL encryption/decryption initialisatio being equivalent and will fetch this from the available providers. This will succeed if the default provider has been loaded that offers this cipher). Using the NULL cipher means that the plaintext is emitted as the ciphertext. Applica call EVP_CIPHER_meth_new() using NID_undef and subsequently use it in a call to an encryption/decryption init SSL/TLS are not impacted by this issue. Fixed in OpenSSL 3.0.6 (Affected 3.0.0-3.0.5).
CVE-2022-3437	A heap-based buffer overflow vulnerability was found in Samba within the GSSAPI unwrap_des() and unwrap_des Triple-DES decryption routines in the Heimdal GSSAPI library allow a length-limited write buffer overflow on ma a maliciously small packet. This flaw allows a remote user to send specially crafted malicious data to the applicatio (DoS) attack.
CVE-2022-34903	GnuPG through 2.3.6, in unusual situations where an attacker possesses any secret-key information from a victim's GPGME) are met, allows signature forgery via injection into the status line.
CVE-2022-3515	A vulnerability was found in the Libksba library due to an integer overflow within the CRL parser. The vulnerabili execution on the target system by passing specially crafted data to the application, for example, a malicious S/MIM
CVE-2022-35252	When curl is used to retrieve and parse cookies from a HTTP(S) server, itaccepts cookies using control codes that might make the server return 400 responses. Effectively allowing a"sister site" to deny service to all siblings.
CVE-2022-3566	A vulnerability, which was classified as problematic, was found in Linux Kernel. This affects the function tcp_gets TCP Handler. The manipulation leads to race condition. It is recommended to apply a patch to fix this issue. The id vulnerability.
CVE-2022-3567	A vulnerability has been found in Linux Kernel and classified as problematic. This vulnerability affects the functio component IPv6 Handler. The manipulation leads to race condition. It is recommended to apply a patch to fix this i to this vulnerability.
CVE-2022-36944	Scala 2.13.x before 2.13.9 has a Java deserialization chain in its JAR file. On its own, it cannot be exploited. There object deserialization within an application. In such situations, it allows attackers to erase contents of arbitrary files arbitrary code (specifically, Function0 functions) via a gadget chain.
CVE-2022-3707	A double-free memory flaw was found in the Linux kernel. The Intel GVT-g graphics driver triggers VGA card sys intel_gvt_dma_map_guest_page function. This issue could allow a local user to crash the system.
CVE-2022-37434	zlib through 1.2.12 has a heap-based buffer over-read or buffer overflow in inflate in inflate.c via a large gzip head call inflateGetHeader are affected. Some common applications bundle the affected zlib source code but may be una nodejs/node reference).
CVE-2022-40152	Those using Woodstox to parse XML data may be vulnerable to Denial of Service attacks (DOS) if DTD support is supplied input, an attacker may supply content that causes the parser to crash by stackoverflow. This effect may su
CVE-2022-40303	An issue was discovered in libxml2 before 2.10.3. When parsing a multi-gigabyte XML document with the XML_ integer counters can overflow. This results in an attempt to access an array at a negative 2GB offset, typically leadi
CVE-2022-40304	An issue was discovered in libxml2 before 2.10.3. Certain invalid XML entity definitions can corrupt a hash table l errors. In one case, a double-free can be provoked.
CVE-2022-4129	A flaw was found in the Linux kernel's Layer 2 Tunneling Protocol (L2TP). A missing lock when clearing sk_user pointer dereference. A local user could use this flaw to potentially crash the system causing a denial of service.
CVE-2022-41916	Heimdal is an implementation of ASN.1/DER, PKIX, and Kerberos. Versions prior to 7.7.1 are vulnerable to a den certificate validation library, affecting the KDC (via PKINIT) and kinit (via PKINIT), as well as any third-party ap should upgrade to Heimdal 7.7.1 or 7.8. There are no known workarounds for this issue.
CVE-2022-42010	An issue was discovered in D-Bus before 1.12.24, 1.13.x and 1.14.x before 1.14.4, and 1.15.x before 1.15.2. An au and other programs that use libdbus to crash when receiving a message with certain invalid type signatures.
CVE-2022-42011	An issue was discovered in D-Bus before 1.12.24, 1.13.x and 1.14.x before 1.14.4, and 1.15.x before 1.15.2. An au and other programs that use libdbus to crash when receiving a message where an array length is inconsistent with t
CVE-2022-42012	An issue was discovered in D-Bus before 1.12.24, 1.13.x and 1.14.x before 1.14.4, and 1.15.x before 1.15.2. An au and other programs that use libdbus to crash by sending a message with attached file descriptors in an unexpected f
CVE-2022-4285	An illegal memory access flaw was found in the binutils package. Parsing an ELF file containing corrupt symbol v service. This issue is the result of an incomplete fix for CVE-2020-16599.

CVE-2022-42898	PAC parsing in MIT Kerberos 5 (aka krb5) before 1.19.4 and 1.20.x before 1.20.1 has integer overflows that may l kadmind, or a GSS or Kerberos application server) on 32-bit platforms (which have a resultant heap-based buffer o other platforms. This occurs in krb5_pac_parse in lib/krb5/krb/pac.c. Heimdal before 7.7.1 has "a similar bug."
CVE-2022-43680	In libexpat through 2.4.9, there is a use-after free caused by overeager destruction of a shared DTD in XML_Exter situations.
CVE-2022-4379	A use-after-free vulnerability was found in __nfs42_ssc_open() in fs/nfs/nfs4file.c in the Linux kernel. This flaw al
CVE-2022-4382	A use-after-free flaw caused by a race among the superblock operations in the gadgetfs Linux driver was found. It that is running the gadgetfs side.
CVE-2022-44640	Heimdal before 7.7.1 allows remote attackers to execute arbitrary code because of an invalid free in the ASN.1 cod (KDC).
CVE-2022-45142	The fix for CVE-2022-3437 included changing memcmp to be constant time and a workaround for a compiler bug of memcmp. When these patches were backported to the heimdal-7.7.1 and heimdal-7.8.0 branches (and possibly o causing the validation of message integrity codes in gssapi/arcfour to be inverted.
CVE-2022-45868	The web-based admin console in H2 Database Engine before 2.2.220 can be started via the CLI with the argument user to specify the password in cleartext for the web admin console. Consequently, a local user (or an attacker that means) would be able to discover the password by listing processes and their arguments. NOTE: the vendor states Passwords should never be passed on the command line and every qualified DBA or system administrator is expec fixed in 2.2.220.
CVE-2022-45873	systemd 250 and 251 allows local users to achieve a systemd-coredump deadlock by triggering a crash that has a lo in shared/elf-util.c. The exploitation methodology is to crash a binary calling the same function recursively, and pu backtrace large enough to cause the deadlock. This must be done 16 times when MaxConnections=16 is set for the
CVE-2022-45886	An issue was discovered in the Linux kernel through 6.0.9. drivers/media/dvb-core/dvb_net.c has a .disconnect ver leads to a use-after-free.
CVE-2022-45887	An issue was discovered in the Linux kernel through 6.0.9. drivers/media/usb/ttusb-dec/ttusb_dec.c has a memory dvb_frontend_detach call.
CVE-2022-45919	An issue was discovered in the Linux kernel through 6.0.10. In drivers/media/dvb-core/dvb_ca_en50221.c, a use-a after an open, because of the lack of a wait_event.
CVE-2022-47629	Libksba before 1.6.3 is prone to an integer overflow vulnerability in the CRL signature parser.
CVE-2022-48174	There is a stack overflow vulnerability in ash.c:6030 in busybox before 1.35. In the environment of Internet of Veh command to arbitrary code execution.
CVE-2022-48554	File before 5.43 has an stack-based buffer over-read in file_copystr in funcs.c. NOTE: "File" is the name of an Ope
CVE-2022-48566	An issue was discovered in compare_digest in Lib/hmac.py in Python through 3.9.1. Constant-time-defeating optin variable in hmac.compare_digest.
CVE-2022-48566	An issue was discovered in compare_digest in Lib/hmac.py in Python through 3.9.1. Constant-time-defeating optin variable in hmac.compare_digest.
CVE-2022-48626	In the Linux kernel, the following vulnerability has been resolved: moxart: fix potential use-after-free on remove p structure could be accessed after it was freed in moxart_remove(), so fix this by saving the base register of the devi dereference.
CVE-2022-48645	In the Linux kernel, the following vulnerability has been resolved: net: enetc: deny offload of tc-based TSN feature ENETC (taprio, cbs, gate, police) are configured through a mix of command BD ring messages and port registers: registers are a region of the ENETC memory map which are only accessible from the PCIe Physical Function. The Functions. Moreover, attempting to access these registers crashes the kernel: $ echo 1 > /sys/bus/pci/devices/0000\ [1957:ef00] type 00 class 0x020001 fsl_enetc_vf 0000:00:01.0: Adding to iommu group 15 fsl_enetc_vf 0000:00:0 fsl_enetc_vf 0000:00:01.0 eno0vf0: renamed from eth0 $ tc qdisc replace dev eno0vf0 root taprio num_tc 8 map 0 1@4 1@5 1@6 1@7 base-time 0 \ sched-entry S 0x7f 900000 sched-entry S 0x80 100000 flags 0x2 Unable to han ffff800009551a08 Internal error: Oops: 96000007 [#1] PREEMPT SMP pc : enetc_setup_tc_taprio+0x170/0x47c l Call trace: enetc_setup_tc_taprio+0x170/0x47c enetc_setup_tc+0x38/0x2dc taprio_change+0x43c/0x970 taprio_in tc_modify_qdisc+0x1fc/0x6c0 rtnetlink_rcv_msg+0x12c/0x390 Split enetc_setup_tc() into separate functions for t enetc_qos.o from being included into enetc-vf.ko, since it serves absolutely no purpose there.
CVE-2022-48655	In the Linux kernel, the following vulnerability has been resolved: firmware: arm_scmi: Harden accesses to the res descriptors by the index upon the SCMI drivers requests through the SCMI reset operations interface can potentiall driver misbehave. Add an internal consistency check before any such domains descriptors accesses.

CVE-2022-48666	In the Linux kernel, the following vulnerability has been resolved: scsi: core: Fix a use-after-free There are two .ex implementations. Both implementations use resources associated with the SCSI host. Make sure that these resource when .exit_cmd_priv is called by waiting inside scsi_remove_host() until the tag set has been freed. This commit fi ========== BUG: KASAN: use-a +0x27/0xd0 [ib_srp] Read of size 8 at addr ffff888100337000 by task multipathd/16727 Call Trace: <TASK> dum +0x5e/0x5db kasan_report+0xab/0x120 srp_exit_cmd_priv+0x27/0xd0 [ib_srp] scsi_mq_exit_request+0x4d/0x70 __blk_mq_free_map_and_rqs+0x6e/0x100 blk_mq_free_tag_set+0x2b/0x160 scsi_host_dev_release+0xf3/0x1a0 +0xa5/0x120 device_release+0x54/0xe0 kobject_put+0xa5/0x120 scsi_device_dev_release_usercontext+0x4c1/0x device_release+0x54/0xe0 kobject_put+0xa5/0x120 scsi_disk_release+0x3f/0x50 device_release+0x54/0xe0 kobje +0x17f/0x1b0 device_release+0x54/0xe0 kobject_put+0xa5/0x120 dm_put_table_device+0xa3/0x160 [dm_mod] free_priority_group+0xd8/0x110 [dm_multipath] free_multipath+0x94/0xe0 [dm_multipath] dm_table_destroy+0x +0x196/0x350 [dm_mod] dev_remove+0x10c/0x160 [dm_mod] ctl_ioctl+0x2c2/0x590 [dm_mod] dm_ctl_ioctl+0 +0xb4/0xf0 dm_ctl_ioctl+0x5/0x10 [dm_mod] __x64_sys_ioctl+0xb4/0xf0 do_syscall_64+0x3b/0x90 entry_SYS
CVE-2022-48674	In the Linux kernel, the following vulnerability has been resolved: erofs: fix pcluster use-after-free on UP platform disabled, KASAN reports as below: ========== free in __mutex_lock+0xe5/0xc30 Read of size 8 at addr ffff8881094223f8 by task stress/7789 CPU: 0 PID: 7789 g0d53d2e882f9 #3 Hardware name: Red Hat KVM, BIOS 0.5.1 01/01/2011 Call Trace: <TASK> .. __mutex_lock +0x8ce/0x1560 .. z_erofs_readahead+0x31c/0x580 .. Freed by task 7787 kasan_save_stack+0x1e/0x40 kasan_set_ +0x20/0x40 __kasan_slab_free+0x10c/0x190 kmem_cache_free+0xed/0x380 rcu_core+0x3d5/0xc90 __do_softirq creation: kasan_save_stack+0x1e/0x40 __kasan_record_aux_stack+0x97/0xb0 call_rcu+0x3d/0x3f0 erofs_shrink_ +0xdc/0x170 shrink_slab.constprop.0+0x296/0x530 drop_slab+0x1c/0x70 drop_caches_sysctl_handler+0x70/0x8 vfs_write+0x555/0x6c0 ksys_write+0xbe/0x160 do_syscall_64+0x3b/0x90 The root cause is that erofs_workgroup it causes a race that the pcluster reuses unexpectedly before freeing. Since UP platforms are quite rare now, such pa specific-designed path directly instead.
CVE-2022-48708	In the Linux kernel, the following vulnerability has been resolved: pinctrl: single: fix potential NULL dereference A pcs_set_mux(). pinmux_generic_get_function() can return NULL and the pointer "function" was dereferenced with Verification Center (linuxtesting.org) with SVACE.
CVE-2022-48781	In the Linux kernel, the following vulnerability has been resolved: crypto: af_alg - get rid of alg_memory_allocated not seem to be really used. alg_proto does have a .memory_allocated field, but no corresponding .sysctl_mem. This true, but all sk_prot_mem_limits() users will trigger a NULL dereference [1]. THis was not a problem until SO_RE protection fault, probably for non-canonical address 0xdffffc0000000001: 0000 [#1] PREEMPT SMP KASAN KA [0x0000000000000008-0x000000000000000f] CPU: 1 PID: 3591 Comm: syz-executor153 Not tainted 5.17.0-rc3- Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google 01/01/2011 RIP: 0010:s [inline] RIP: 0010:sock_reserve_memory+0x1d7/0x330 net/core/sock.c:1000 Code: 08 00 74 08 48 89 ef e8 27 20 c5 08 48 89 e8 48 c1 e8 03 48 b9 00 00 00 00 00 fc ff df <80> 3c 08 00 74 08 48 89 ef e8 fb 1f bb f9 48 8b 6d 00 4 EFLAGS: 00010202 RAX: 0000000000000001 RBX: ffff88814aabc000 RCX: dffffc0000000000 RDX: 0000000 RDI: ffffffff90e18120 RBP: 0000000000000008 R08: dffffc0000000000 R09: fffffbfff21c3025 R10: fffffbfff21c3 ffffffff8d109840 R13: 0000000000001002 R14: 0000000000000001 R15: 0000000000000001 FS: 0000555556e08 knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 00007fc74416f130 CR3: 00000000003506e0 DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 DR3: 0000000 DR7: 0000000000000400 Call Trace: <TASK> sock_setsockopt+0x14a9/0x3a30 net/core/sock.c:1446 __sys_setso __do_sys_setsockopt net/socket.c:2191 [inline] __se_sys_setsockopt net/socket.c:2188 [inline] __x64_sys_setsock do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x44/0xd0 arch/x86/entry/common.c:80 entry RIP: 0033:0x7fc7440fddc9 Code: 28 00 00 00 75 05 48 83 c4 28 c3 e8 51 15 00 00 90 48 89 f8 48 89 f7 48 89 d6 4 0f 05 <48> 3d 01 f0 ff ff 73 01 c3 48 c7 c1 c0 ff ff ff f7 d8 64 89 01 48 RSP: 002b:00007ffe98f07968 EFLAGS: 00 RAX: ffffffffffffffda RBX: 0000000000000003 RCX: 00007fc7440fddc9 RDX: 0000000000000049 RSI: 0000000 RBP: 0000000000000000 R08: 0000000000000004 R09: 00007ffe98f07990 R10: 0000000020000000 R11: 00000 R13: 00007ffe98f079a0 R14: 00007ffe98f079e0 R15: 0000000000000000 </TASK> Modules linked in: ---[end tr 0010:sk_prot_mem_limits include/net/sock.h:1523 [inline] RIP: 0010:sock_reserve_memory+0x1d7/0x330 net/cor ef e8 27 20 bb f9 4c 03 7c 24 10 48 8b 6d 00 48 83 c5 08 48 89 e8 48 c1 e8 03 48 b9 00 00 00 00 00 fc ff df <80> 48 8b 6d 00 4c 89 ff 48 RSP: 0018:ffffc90001f1fb68 EFLAGS: 00010202 RAX: 0000000000000001 RBX: ffff888 RDX: 0000000000000001 RSI: 0000000000000008 RDI: ffffffff90e18120 RBP: 0000000000000008 R08: dffffc0 R10: fffffbfff21c3025 R11: 0000000000000000 R12: ffffffff8d109840 R13: 0000000000001002 R14: 0000000000 FS: 0000555556e08300(0000) GS:ffff8880b9b00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 00 00007fc74416f130 CR3: 0000000073d9e000 CR4: 00000000003506e0 DR0: 0000000000000000 DR1: 0000000
CVE-2022-48791	In the Linux kernel, the following vulnerability has been resolved: scsi: pm8001: Fix use-after-free for aborted TM occur if a TMF sas_task is aborted before we handle the IO completion in mpi_ssp_completion(). The abort occurs SAS_TASK_STATE_ABORTED flag is set and the sas_task is freed in pm8001_exec_internal_tmf_task(). Howe O completion still thinks that the sas_task is available. Fix this by clearing the ccb->task if the TMF times out - the pointer is cleared.
CVE-2022-48792	In the Linux kernel, the following vulnerability has been resolved: scsi: pm8001: Fix use-after-free for aborted SSP may occur if a sas_task is aborted by the upper layer before we handle the I/O completion in mpi_ssp_completion() following are the two steps in handling those I/O completions: - Call complete() to inform the upper layer handler o resources associated with the sas_task in pm8001_ccb_task_free() call. When complete() is called, the upper layer touch the associated sas_task afterwards, but we do so in the pm8001_ccb_task_free() call. Fix by swapping the co ordering.

CVE-2022-48814	In the Linux kernel, the following vulnerability has been resolved: net: dsa: seville: register the mdiobus under dev ("net: dsa: realtek: register the MDIO bus under devres") 5135e96a3dd2 ("net: dsa: don't allocate the slave_mii_bu when called from devm_mdiobus_free() <- devres_release_all() <- __device_release_driver(), and that mdiobus wa VSC9959 switch is a platform device, so the initial set of constraints that I thought would cause this (I2C or SPI bu do not apply. But there is one more which applies here. If the DSA master itself is on a bus that calls ->remove fro the fsl-mc bus), there is a device link between the switch and the DSA master, and device_links_unbind_consumer shutdown. So the same treatment must be applied to all DSA switch drivers, which is: either use devres for both th use devres at all. The seville driver has a code structure that could accommodate both the mdiobus_unregister and dependency upon mscc_miim_setup() from mdio-mscc-miim.c, which calls devm_mdiobus_alloc_size() on its beh exporting yet one more symbol mscc_miim_teardown(), let's work with devres and replace of_mdiobus_register w devres, we can ensure that devres doesn't free a still-registered bus (it either runs both callbacks, or none).
CVE-2022-48816	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: lock against ->sock changing during asynchronously unless ->recv_mutex is held. So it is important to hold that mutex. Otherwise a sysfs read can trigg ("SUNRPC: Check if the xprt is connected before handling sysfs reads") appears to attempt to fix this problem, but
CVE-2022-48818	In the Linux kernel, the following vulnerability has been resolved: net: dsa: mv88e6xxx: don't use devres for mdiob ("net: dsa: realtek: register the MDIO bus under devres") 5135e96a3dd2 ("net: dsa: don't allocate the slave_mii_bu panic when called from devm_mdiobus_free() <- devres_release_all() <- __device_release_driver(), and that mdiob mv88e6xxx is an MDIO device, so the initial set of constraints that I thought would cause this (I2C or SPI buses w not apply. But there is one more which applies here. If the DSA master itself is on a bus that calls ->remove from - the fsl-mc bus), there is a device link between the switch and the DSA master, and device_links_unbind_consumer on shutdown. systemd-shutdown[1]: Powering off. mv88e6085 0x0000000008b96000:00 sw_gl0: Link is Down fs group 7 fsl-mc dpbp.8: Removing from iommu group 7 ----------[cut here]---------- kernel BUG at drivers/net/p BUG: 0 [#1] PREEMPT SMP Modules linked in: CPU: 0 PID: 1 Comm: systemd-shutdow Not tainted 5.16.5-0004 +0x44/0x50 lr : devm_mdiobus_free+0x10/0x20 Call trace: mdiobus_free+0x44/0x50 devm_mdiobus_free+0x10/0 __device_release_driver+0x190/0x220 device_release_driver_internal+0xac/0xb0 device_links_unbind_consumers +0x4c/0x220 device_release_driver_internal+0xac/0xb0 device_links_unbind_consumers+0xd4/0x100 __device_r device_release_driver+0x28/0x40 bus_remove_device+0x118/0x124 device_del+0x174/0x420 fsl_mc_device_rem +0xc/0x20 device_for_each_child+0x58/0xa0 dprc_remove+0x90/0xb0 fsl_mc_driver_remove+0x20/0x5c __devi device_release_driver+0x28/0x40 bus_remove_device+0x118/0x124 device_del+0x174/0x420 fsl_mc_bus_remov +0xc/0x1c platform_shutdown+0x20/0x30 device_shutdown+0x154/0x330 kernel_power_off+0x34/0x6c __do_sy +0x20/0x30 invoke_syscall.constprop.0+0x4c/0xe0 do_el0_svc+0x4c/0x150 el0_svc+0x24/0xb0 el0t_64_sync_ha So the same treatment must be applied to all DSA switch drivers, which is: either use devres for both the mdiobus a at all. The Marvell driver already has a good structure for mdiobus removal, so just plug in mdiobus_free and get ri
CVE-2022-48826	In the Linux kernel, the following vulnerability has been resolved: drm/vc4: Fix deadlock on DSI device attach erro with host device's lock held. Un-registering host in "device attach" error path (ex: probe retry) will result in deadloc DSI display. Startup Call trace: [35.043036] rt_mutex_slowlock.constprop.21+0x184/0x1b8 [35.043048] mutex_l device_del+0x4c/0x3e8 [35.043075] device_unregister+0x20/0x40 [35.043082] mipi_dsi_remove_device_fn+0x +0x68/0xb0 [35.043105] mipi_dsi_host_unregister+0x40/0x90 [35.043115] vc4_dsi_host_attach+0xf0/0x120 [vc +0x30/0x48 [35.043209] tc358762_probe+0x128/0x164 [tc358762] [35.043225] mipi_dsi_drv_probe+0x28/0x38 [35.043244] __driver_probe_device+0x80/0xe8 [35.043254] driver_probe_device+0xb8/0x118 [35.043263] __d bus_for_each_drv+0x84/0xd8 [35.043281] __device_attach+0xf0/0x150 [35.043290] device_initial_probe+0x1c/ +0xa4/0xb0 [35.043308] deferred_probe_work_func+0xa0/0xe0 [35.043318] process_one_work+0x254/0x700 [[35.043339] kthread+0x19c/0x1a8 [35.043348] ret_from_fork+0x10/0x20 Shutdown Call trace: [365.565417] Ca +0x148/0x200 [365.565452] __schedule+0x340/0x9c8 [365.565467] schedule+0x48/0x110 [365.565479] schedu wait_for_completion+0xac/0x138 [365.565509] __flush_work+0x218/0x4e0 [365.565523] flush_work+0x1c/0x2 +0x68/0x158 [365.565550] device_shutdown+0x24/0x348 [365.565561] kernel_restart_prepare+0x40/0x50 [365 [365.565591] __do_sys_reboot+0x10c/0x220 [365.565605] __arm64_sys_reboot+0x2c/0x38 [365.565619] invol el0_svc_common.constprop.3+0xfc/0x120 [365.565648] do_el0_svc+0x2c/0x90 [365.565661] el0_svc+0x4c/0xf +0x90/0xb8 [365.565682] el0t_64_sync+0x180/0x184

CVE-2022-48833	In the Linux kernel, the following vulnerability has been resolved: btrfs: skip reserved bytes warning on unmount a changes made by commit c2e39305299f01 ("btrfs: clear extent buffer uptodate when we fail to write it") and its fol ("btrfs: check WRITE_ERR when trying to read an extent buffer"), we can now end up not cleaning up space reser transaction abort happens, as well as not cleaning up still dirty extent buffers. This happens because if writeback fo have cleared the bit EXTENT_BUFFER_UPTODATE from the extent buffer and we have also set the bit EXTENT when trying to free the log tree with free_log_tree(), which iterates over the tree, we can end up getting an -EIO err read_extent_buffer_pages() returns -EIO if an extent buffer does not have EXTENT_BUFFER_UPTODATE set an bit set. Getting that -EIO means that we return immediately as we can not iterate over the entire tree. In that case w extent buffer in the respective block group and space_info object. When this happens we get the following traces w BTRFS: error (device dm-0) in cleanup_transaction:1913: errno=-5 IO failure [174957.286497] BTRFS: error (dev errno=-5 IO failure [174957.399379] ------------[cut here]------------ [174957.402497] WARNING: CPU: 2 PID: 3 btrfs_put_block_group+0x77/0xb0 [btrfs] [174957.407523] Modules linked in: btrfs overlay dm_zero (...) [174957 umount Tainted: G W 5.16.0-rc5-btrfs-next-109 #1 [174957.426689] Hardware name: QEMU Standard PC (i440F g155821a1990b-prebuilt.qemu.org 04/01/2014 [174957.428716] RIP: 0010:btrfs_put_block_group+0x77/0xb0 [bt bd (...) [174957.432867] RSP: 0018:ffffb70d41cffdd0 EFLAGS: 00010206 [174957.433632] RAX: 000000000000 ffff8b0758edd1c8 [174957.434689] RDX: 0000000000000001 RSI: ffffffffc0b467e7 RDI: ffff8b0758edd000 [174 0000000000000000 R09: 0000000000000000 [174957.437114] R10: 0000000000000246 R11: 000000000000000 R13: ffff8b09c3848198 R14: ffff8b0758edd188 R15: dead000000000100 [174957.439317] FS: 00007f328fb8280 knlGS:0000000000000000 [174957.440402] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [174957.441 CR3: 0000000404f4e005 CR4: 0000000000370ee0 [174957.442117] DR0: 0000000000000000 DR1: 000000000 [174957.443076] DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 [174957.443948] C [174957.444538] btrfs_free_block_groups+0x255/0x3c0 [btrfs] [174957.445238] close_ctree+0x301/0x357 [btrfs] +0x16c/0x290 [174957.446250] generic_shutdown_super+0x74/0x120 [174957.446832] kill_anon_super+0x14/0x +0x12/0x20 [btrfs] [174957.447890] deactivate_locked_super+0x31/0xa0 [174957.448440] cleanup_mnt+0x147/0 +0x5c/0xa0 [174957.449336] exit_to_user_mode_prepare+0x1e5/0x1f0 [174957.449934] syscall_exit_to_user_mo do_syscall_64+0x48/0xc0 [174957.450980] entry_SYSCALL_64_after_hwframe+0x44/0xae [174957.451605] RI Code: 03 0c 00 f7 (...) [174957.454320] RSP: 002b:00007fff13564ec8 EFLAGS: 00000246 ORIG_RAX: 000000 0000000000000000 RBX: 00007f328feea264 RCX: 00007f328fdc4a97 [174957.456131] RDX: 000000000000000
CVE-2022-48852	In the Linux kernel, the following vulnerability has been resolved: drm/vc4: hdmi: Unregister codec device on unbi device but we don't unregister it on unbind, leading to a device leakage. Unregister our device at unbind.
CVE-2022-48859	In the Linux kernel, the following vulnerability has been resolved: net: marvell: prestera: Add missing of_node_pu This node pointer is returned by of_find_compatible_node() with refcount incremented. Calling of_node_put() to a
CVE-2022-4886	Ingress-nginx `path` sanitization can be bypassed with `log_format` directive.
CVE-2023-0361	A timing side-channel in the handling of RSA ClientKeyExchange messages was discovered in GnuTLS. This side encrypted in the RSA ciphertext across a network in a Bleichenbacher style attack. To achieve a successful decrypt amount of specially crafted messages to the vulnerable server. By recovering the secret from the ClientKeyExchang decrypt the application data exchanged over that connection.
CVE-2023-0458	A speculative pointer dereference problem exists in the Linux Kernel on the do_prlimit() function. The resource arg and is used in pointer arithmetic for the 'rlim' variable and can be used to leak the contents. We recommend upgrad commit¬†739790605705ddcf18f21782b9c99ad7d53a8c11
CVE-2023-0459	Copy_from_user on 64-bit versions of the Linux kernel does not implement the __uaccess_begin_nospec allowing check and pass a kernel pointer to copy_from_user(). This would allow an attacker to leak information. We recomm commit¬†74e19ef0ff8061ef55957c3abd71614ef0f42f47
CVE-2023-0461	There is a use-after-free vulnerability in the Linux Kernel which can be exploited to achieve local privilege escalati configuration flag CONFIG_TLS¬†or CONFIG_XFRM_ESPINTCP¬†has to be configured, but the operation doe is a use-after-free bug of icsk_ulp_data¬†of a struct inet_connection_sock. When CONFIG_TLS¬†is enabled, user tls_context) on a connected tcp socket. The context is not cleared if this socket is disconnected and reused as a liste listener, the context is inherited and vulnerable. The setsockopt¬†TCP_ULP¬†operation does not require any privi commit¬†2c02d41d71f90a5168391b6a5f2954112ba2307c
CVE-2023-0597	A flaw possibility of memory leak in the Linux kernel cpu_entry_area mapping of X86 CPU data to memory was f exception stack(s) or other important data. A local user could use this flaw to get access to some important data wit
CVE-2023-1073	A memory corruption flaw was found in the Linux kernel,Äôs human interface device (HID) subsystem in how a u allows a local user to crash or potentially escalate their privileges on the system.
CVE-2023-1074	A memory leak flaw was found in the Linux kernel's Stream Control Transmission Protocol. This issue may occur service and someone connects to this service. This could allow a local user to starve resources, causing a denial of s
CVE-2023-1075	A flaw was found in the Linux Kernel. The tls_is_tx_ready() incorrectly checks for list emptiness, potentially acces leaking the last byte of the confused field that overlaps with rec->tx_ready.
CVE-2023-1078	A flaw was found in the Linux Kernel in RDS (Reliable Datagram Sockets) protocol. The rds_rm_zerocopy_callba causing a type confusion. Local user can trigger this with rds_message_put(). Type confusion leads to `struct rds_m something else that is potentially controlled by local user. It is known how to trigger this, which causes an out of bo

CVE-2023-1079	A flaw was found in the Linux kernel. A use-after-free may be triggered in asus_kbd_backlight_set when plugging which advertises itself as an Asus device. Similarly to the previous known CVE-2023-25012, but in asus devices, t controller while the device is disconnecting, triggering a use-after-free on the struct asus_kbd_leds *led structure. A to cause memory corruption with controlled data.
CVE-2023-1118	A flaw use after free in the Linux kernel integrated infrared receiver/transceiver driver was found in the way user d flaw to crash the system or potentially escalate their privileges on the system.
CVE-2023-1192	A use-after-free flaw was found in smb2_is_status_io_timeout() in CIFS in the Linux Kernel. After CIFS transfers local variable points to the memory region, and if the system call frees it faster than CIFS uses it, CIFS will access service.
CVE-2023-1281	Use After Free vulnerability in Linux kernel traffic control index filter (tcindex) allows Privilege Escalation.¬†The while packets are traversing, which will cause a use-after-free when 'tcf_exts_exec()' is called with the destroyed tc vulnerability to elevate its privileges to root. This issue affects Linux Kernel: from 4.14 before git commit ee05917
CVE-2023-1513	A flaw was found in KVM. When calling the KVM_GET_DEBUGREGS ioctl, on 32-bit systems, there might be s kvm_debugregs structure that could be copied to userspace, causing an information leak.
CVE-2023-1579	Heap based buffer overflow in binutils-gdb/bfd/libbfd.c in bfd_getl64.
CVE-2023-1611	A use-after-free flaw was found in btrfs_search_slot in fs/btrfs/ctree.c in btrfs in the Linux Kernel.This flaw allows cause a kernel information lea
CVE-2023-1670	A flaw use after free in the Linux kernel Xircom 16-bit PCMCIA (PC-card) Ethernet driver was found.A local user potentially escalate their privileges on the system.
CVE-2023-1829	A use-after-free vulnerability in the Linux Kernel traffic control index filter (tcindex) can be exploited to achieve lo tcindex_delete function which does not properly deactivate filters in case of a perfect hashes while deleting the und to double freeing the structure.¬†A local attacker user can use this vulnerability to elevate its privileges to root. We 8c710f75256bb3cf05ac7b1672c82b92c43f3d28.
CVE-2023-1855	A use-after-free flaw was found in xgene_hwmon_remove in drivers/hwmon/xgene-hwmon.c in the Hardware Mon This flaw could allow a local attacker to crash the system due to a race problem. This vulnerability could even lead
CVE-2023-1859	A use-after-free flaw was found in xen_9pfs_front_removet in net/9p/trans_xen.c in Xen transport for 9pfs in the L attacker to crash the system due to a race problem, possibly leading to a kernel information leak.
CVE-2023-1872	A use-after-free vulnerability in the Linux Kernel io_uring system can be exploited to achieve local privilege escal the presence of ctx->uring_lock which can lead to a Use-After-Free vulnerability due a race condition with fixed fi upgrading past commit da24142b1ef9fd5d36b76e36bab328a5b27523e8.
CVE-2023-1989	A use-after-free flaw was found in btsdio_remove in drivers\bluetooth\btsdio.c in the Linux Kernel. In this flaw, a may cause a race problem leading to a UAF on hdev devices.
CVE-2023-1990	A use-after-free flaw was found in ndlc_remove in drivers/nfc/st-nci/ndlc.c in the Linux Kernel. This flaw could al race problem.
CVE-2023-1998	The Linux kernel allows userspace processes to enable mitigations by calling prctl with PR_SET_SPECULATION feature as well as by using seccomp. We had noticed that on VMs of at least one major cloud provider, the kernel s attacks in some cases even after enabling the spectre-BTI mitigation with prctl. The same behavior can be observed the mitigation to IBRS on boot command line. This happened because when plain IBRS was enabled (not enhanced determined that STIBP was not needed. The IBRS bit implicitly protects against cross-thread branch target injectio bit was cleared on returning to userspace, due to performance reasons, which disabled the implicit STIBP and left u branch target injection against which STIBP protects.
CVE-2023-20860	Spring Framework running version 6.0.0 - 6.0.6 or 5.3.0 - 5.3.25 using "**" as a pattern in Spring Security configu mismatch in pattern matching between Spring Security and Spring MVC, and the potential for a security bypass.
CVE-2023-20861	In Spring Framework versions 6.0.0 - 6.0.6, 5.3.0 - 5.3.25, 5.2.0.RELEASE - 5.2.22.RELEASE, and older unsuppo provide a specially crafted SpEL expression that may cause a denial-of-service (DoS) condition.
CVE-2023-21102	In __efi_rt_asm_wrapper of efi-rt-wrapper.S, there is a possible bypass of shadow stack protection due to a logic er escalation of privilege with no additional execution privileges needed. User interaction is not needed for exploitatio kernelAndroid ID: A-260821414References: Upstream kernel
CVE-2023-2162	A use-after-free vulnerability was found in iscsi_sw_tcp_session_create in drivers/scsi/iscsi_tcp.c in SCSI sub-com attacker could leak kernel internal information.
CVE-2023-2163	Incorrect verifier pruning¬†in BPF in Linux Kernel¬†>=5.4¬†leads to unsafe code paths being incorrectly marked kernel memory, lateral privilege escalation, and container escape.
CVE-2023-2194	An out-of-bounds write vulnerability was found in the Linux kernel's SLIMpro I2C device driver. The userspace "c number between 0-255 and was used as the size of a memcpy, possibly writing beyond the end of dma_buffer. This crash the system or potentially achieve code execution.

CVE-2023-22025	Vulnerability in the Oracle Java SE, Oracle GraalVM for JDK, Oracle GraalVM Enterprise Edition, product of Ora versions that are affected are Oracle Java SE: 8u381-perf, 17.0.8, 21; Oracle GraalVM for JDK: 17.0.8, 21; Oracle 22.3.3. Difficult to exploit vulnerability allows unauthenticated attacker with network access via multiple protocols GraalVM for JDK, Oracle GraalVM Enterprise Edition,. Successful attacks of this vulnerability can result in unaut of Oracle Java SE, Oracle GraalVM for JDK, Oracle GraalVM Enterprise Edition, accessible data. Note: This vuln the specified Component, e.g., through a web service which supplies data to the APIs. This vulnerability also appli running sandboxed Java Web Start applications or sandboxed Java applets, that load and run untrusted code (e.g., c the Java sandbox for security. CVSS 3.1 Base Score 3.7 (Integrity impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:H
CVE-2023-22067	Vulnerability in the Oracle Java SE, Oracle GraalVM Enterprise Edition product of Oracle Java SE (component: C are Oracle Java SE: 8u381, 8u381-perf; Oracle GraalVM Enterprise Edition: 20.3.11 and 21.3.7. Easily exploitable with network access via CORBA to compromise Oracle Java SE, Oracle GraalVM Enterprise Edition. Successful a unauthorized update, insert or delete access to some of Oracle Java SE, Oracle GraalVM Enterprise Edition accessi exploited by supplying data to APIs in the specified Component without using Untrusted Java Web Start applicatio a web service. CVSS 3.1 Base Score 5.3 (Integrity impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:
CVE-2023-22081	Vulnerability in the Oracle Java SE, Oracle GraalVM for JDK, Oracle GraalVM Enterprise Edition product of Orac versions that are affected are Oracle Java SE: 8u381, 8u381-perf, 11.0.20, 17.0.8, 21; Oracle GraalVM for JDK: 17 20.3.11, 21.3.7 and 22.3.3. Easily exploitable vulnerability allows unauthenticated attacker with network access via Oracle GraalVM for JDK, Oracle GraalVM Enterprise Edition. Successful attacks of this vulnerability can result in of service (partial DOS) of Oracle Java SE, Oracle GraalVM for JDK, Oracle GraalVM Enterprise Edition. Note: T typically in clients running sandboxed Java Web Start applications or sandboxed Java applets, that load and run unt internet) and rely on the Java sandbox for security. This vulnerability does not apply to Java deployments, typically code (e.g., code installed by an administrator). CVSS 3.1 Base Score 5.3 (Availability impacts). CVSS Vector: (CV I:N/A:L).
CVE-2023-2283	A vulnerability was found in libssh, where the authentication check of the connecting client can be bypassed in the memory allocation problems. This issue may happen if there is insufficient memory or the memory usage is limited `rc,` which is initialized to SSH_ERROR and later rewritten to save the return value of the function call `pki_key_c variable is not changed between this point and the cryptographic verification. Therefore any error between them ca
CVE-2023-22998	In the Linux kernel before 6.0.3, drivers/gpu/drm/virtio/virtgpu_object.c misinterprets the drm_gem_shmem_get_s in the error case, whereas it is actually an error pointer).
CVE-2023-23003	In the Linux kernel before 5.16, tools/perf/util/expr.c lacks a check for the hashmap__new return value.
CVE-2023-23039	An issue was discovered in the Linux kernel through 6.2.0-rc2. drivers/tty/vcc.c has a race condition and resultant u attacker removes a VCC device while calling open(), aka a race condition between vcc_open() and vcc_remove().
CVE-2023-23559	In rndis_query_oid in drivers/net/wireless/rndis_wlan.c in the Linux kernel through 6.1.5, there is an integer overfl
CVE-2023-2454	schema_element defeats protective search_path changes; It was found that certain database calls in PostgreSQL co database-level privileges to execute arbitrary code.
CVE-2023-2455	Row security policies disregard user ID changes after inlining; PostgreSQL could permit incorrect policies to be ap policies are used and a given query is planned under one role and then executed under other roles. This scenario ca when a common user and query is planned initially and then re-used across multiple SET ROLEs. Applying an inc otherwise-forbidden reads and modifications. This affects only databases that have used CREATE POLICY to defi
CVE-2023-24998	Apache Commons FileUpload before 1.5 does not limit the number of request parts to be processed resulting in the with a malicious upload or series of uploads. Note that, like all of the file upload limits, the new configuration optio enabled by default and must be explicitly configured.
CVE-2023-25012	The Linux kernel through 6.1.9 has a Use-After-Free in bigben_remove in drivers/hid/hid-bigbenff.c via a crafted U remain registered for too long.
CVE-2023-25613	An LDAP Injection vulnerability exists in the¬†LdapIdentityBackend of Apache Kerby before 2.0.3.¬†
CVE-2023-2602	A vulnerability was found in the pthread_create() function in libcap. This issue may allow a malicious actor to use error, which can exhaust the process memory.
CVE-2023-2603	A vulnerability was found in libcap. This issue occurs in the _libcap_strdup() function and can lead to an integer ov
CVE-2023-26159	Versions of the package follow-redirects before 1.15.4 are vulnerable to Improper Input Validation due to the impr function. When new URL() throws an error, it can be manipulated to misinterpret the hostname. An attacker could malicious site, potentially leading to information disclosure, phishing attacks, or other security breaches.
CVE-2023-26545	In the Linux kernel before 6.1.13, there is a double free in net/mpls/af_mpls.c upon an allocation failure (for registe during the renaming of a device.
CVE-2023-27533	A vulnerability in input validation exists in curl <8.0 during communication using the TELNET protocol may allow user name and "telnet options" during server negotiation. The lack of proper input scrubbing allows an attacker to s without the application's intent. This vulnerability could be exploited if an application allows user input, thereby en the system.

CVE-2023-27535	An authentication bypass vulnerability exists in libcurl <8.0.0 in the FTP connection reuse feature that can result in subsequent transfers. Previously created connections are kept in a connection pool for reuse if they match the curre such as CURLOPT_FTP_ACCOUNT, CURLOPT_FTP_ALTERNATIVE_TO_USER, CURLOPT_FTP_SSL_CC included in the configuration match checks, causing them to match too easily. This could lead to libcurl using the w potentially allowing unauthorized access to sensitive information.
CVE-2023-27536	An authentication bypass vulnerability exists libcurl <8.0.0 in the connection reuse feature which can reuse previou incorrect user permissions due to a failure to check for changes in the CURLOPT_GSSAPI_DELEGATION option negotiate/GSSAPI transfers and could potentially result in unauthorized access to sensitive information. The safest CURLOPT_GSSAPI_DELEGATION option has been changed.
CVE-2023-27538	An authentication bypass vulnerability exists in libcurl prior to v8.0.0 where it reuses a previously established SSH option was modified, which should have prevented reuse. libcurl maintains a pool of previously used connections t configurations match. However, two SSH settings were omitted from the configuration check, allowing them to ma an inappropriate connection.
CVE-2023-28321	An improper certificate validation vulnerability exists in curl <v8.1.0 in the way it supports matching of wildcard p Name" in TLS server certificates. curl can be built to use its own name matching function for TLS rather than one p wildcard matching function would match IDN (International Domain Name) hosts incorrectly and could as a result mismatch. IDN hostnames are converted to puny code before used for certificate checks. Puny coded names always to pattern match, but the wildcard check in curl could still check for `x*`, which would match even though the IDN resembling an `x`.
CVE-2023-28322	An information disclosure vulnerability exists in curl <v8.1.0 when doing HTTP(S) transfers, libcurl might erroneo (`CURLOPT_READFUNCTION`) to ask for data to send, even when the `CURLOPT_POSTFIELDS` option has wasused to issue a `PUT` request which used that callback. This flaw may surprise the application and cause it to m or use memory after free or similar in the second transfer. The problem exists in the logic for a reused handle when a POST.
CVE-2023-28328	A NULL pointer dereference flaw was found in the az6027 driver in drivers/media/usb/dev-usb/az6027.c in the Lin not checked properly before transferring into the device. This flaw allows a local user to crash the system or potent
CVE-2023-28466	do_tls_getsockopt in net/tls/tls_main.c in the Linux kernel through 6.2.6 lacks a lock_sock call, leading to a race co NULL pointer dereference).
CVE-2023-28484	In libxml2 before 2.10.4, parsing of certain invalid XSD schemas can lead to a NULL pointer dereference and subs xmlSchemaFixupComplexType in xmlschemas.c.
CVE-2023-29469	An issue was discovered in libxml2 before 2.10.4. When hashing empty dict strings in a crafted XML document, xm non-deterministic values, leading to various logic and memory errors, such as a double free. This behavior occurs b of an empty string, and any value is possible (not solely the '\0' value).
CVE-2023-2985	A use after free flaw was found in hfsplus_put_super in fs/hfsplus/super.c in the Linux Kernel. This flaw could allo problem.
CVE-2023-30456	An issue was discovered in arch/x86/kvm/vmx/nested.c in the Linux kernel before 6.2.8. nVMX on x86_64 lacks c
CVE-2023-30772	The Linux kernel before 6.2.9 has a race condition and resultant use-after-free in drivers/power/supply/da9150-cha unplugs a device.
CVE-2023-31081	An issue was discovered in drivers/media/test-drivers/vidtv/vidtv_bridge.c in the Linux kernel 6.2. There is a NUL vidtv_mux_stop_thread. In vidtv_stop_streaming, after dvb->mux=NULL occurs, it executes vidtv_mux_stop_thre
CVE-2023-31083	An issue was discovered in drivers/bluetooth/hci_ldisc.c in the Linux kernel 6.2. In hci_uart_tty_ioctl, there is a rac and HCIUARTGETPROTO. HCI_UART_PROTO_SET is set before hu->proto is set. A NULL pointer dereferenc
CVE-2023-3141	A use-after-free flaw was found in r592_remove in drivers/memstick/host/r592.c in media access in the Linux Kern the system at device disconnect, possibly leading to a kernel information leak.
CVE-2023-3161	A flaw was found in the Framebuffer Console (fbcon) in the Linux Kernel. When providing font->width and font->h since there are no checks in place, a shift-out-of-bounds occurs leading to undefined behavior and possible denial o
CVE-2023-3220	An issue was discovered in the Linux kernel through 6.1-rc8. dpu_crtc_atomic_check in drivers/gpu/drm/msm/disp value of kzalloc() and will cause the NULL Pointer Dereference.
CVE-2023-32269	An issue was discovered in the Linux kernel before 6.1.11. In net/netrom/af_netrom.c, there is a use-after-free beca connected AF_NETROM socket. However, in order for an attacker to exploit this, the system must have netrom ro CAP_NET_ADMIN capability.
CVE-2023-33203	The Linux kernel before 6.2.9 has a race condition and resultant use-after-free in drivers/net/ethernet/qualcomm/em unplugs an emac based device.
CVE-2023-33460	There's a memory leak in yajl 2.1.0 with use of yajl_tree_parse function. which will cause out-of-memory in server

CVE-2023-33953	gRPC contains a vulnerability that allows hpack table accounting errors could lead to unwanted disconnects betwee ¬†Three vectors were found that allow the following DOS attacks: - Unbounded memory buffering in the HPACK the HPACK parser The unbounded CPU consumption is down to a copy that occurred per-input-block in the parser to the memory copy bug we end up with an O(n^2) parsing loop, with n selected by the client. The unbounded mem check was behind the string reading code, so we needed to first buffer up to a 4 gigabyte string before rejecting it a have an encoding quirk whereby an infinite number of 0,Äôs can be added at the start of an integer. gRPC,Äôs hpa concluding a parse. - gRPC,Äôs metadata overflow check was performed per frame, so that the following sequence HEADERS: containing a: 1 CONTINUATION: containing a: 2 CONTINUATION: containing a: 3 etc,Ä¶
CVE-2023-3397	A race condition occurred between the functions lmLogClose and txEnd in JFS, in the Linux Kernel, executed in d attacker with normal user privileges to crash the system or leak internal kernel information.
CVE-2023-34256	An issue was discovered in the Linux kernel before 6.3.3. There is an out-of-bounds read in crc16 in lib/crc16.c wh ext4_group_desc_csum does not properly check an offset. NOTE: this is disputed by third parties because the kern with the stated "When modifying the block device while it is mounted by the filesystem" access.
CVE-2023-3567	A use-after-free flaw was found in vcs_read in drivers/tty/vt/vc_screen.c in vc_screen in the Linux Kernel. This iss access to cause a system crash or leak internal kernel information.
CVE-2023-35823	An issue was discovered in the Linux kernel before 6.3.2. A use-after-free was found in saa7134_finidev in drivers
CVE-2023-35824	An issue was discovered in the Linux kernel before 6.3.2. A use-after-free was found in dm1105_remove in drivers
CVE-2023-35828	An issue was discovered in the Linux kernel before 6.3.2. A use-after-free was found in renesas_usb3_remove in d
CVE-2023-35829	An issue was discovered in the Linux kernel before 6.3.2. A use-after-free was found in rkvdec_remove in drivers/
CVE-2023-37453	An issue was discovered in the USB subsystem in the Linux kernel through 6.4.2. There is an out-of-bounds and cr sysfs.c.
CVE-2023-3772	A flaw was found in the Linux kernel,Äôs IP framework for transforming packets (XFRM subsystem). This issue r CAP_NET_ADMIN privileges to directly dereference a NULL pointer in xfrm_update_ae_params(), leading to a p
CVE-2023-3773	A flaw was found in the Linux kernel,Äôs IP framework for transforming packets (XFRM subsystem). This issue r CAP_NET_ADMIN privileges to cause a 4 byte out-of-bounds read of XFRMA_MTIMER_THRESH when parsin leakage of sensitive heap data to userspace.
CVE-2023-3777	A use-after-free vulnerability in the Linux kernel's netfilter: nf_tables component can be exploited to achieve local nf_tables_delrule() is flushing table rules, it is not checked whether the chain is bound and the chain's owner rule ca circumstances. We recommend upgrading past commit 6eaf41e87a223ae6f8e7a28d6e78384ad7e407f8.
CVE-2023-37788	goproxy v1.1 was discovered to contain an issue which can lead to a Denial of service (DoS) via unspecified vecto
CVE-2023-38546	This flaw allows an attacker to insert cookies at will into a running program using libcurl, if the specific series of c In its API, an application creates 'easy handles' that are the individual handles for single transfers. libcurl provides a called curl_easy_duphandle. If a transfer has cookies enabled when the handle is duplicated, the cookie-enable stat actual cookies. If the source handle did not read any cookies from a specific file on disk, the cloned version of the h as none (using the four ASCII letters, no quotes).Subsequent use of the cloned handle that does not explicitly set a inadvertently load cookies from a file named none - if such a file exists and is readable in the current directory of th correct file format of course.
CVE-2023-39189	A flaw was found in the Netfilter subsystem in the Linux kernel. The nfnl_osf_add_callback function did not valida This flaw allows a local privileged (CAP_NET_ADMIN) attacker to trigger an out-of-bounds read, leading to a cra
CVE-2023-39192	A flaw was found in the Netfilter subsystem in the Linux kernel. The xt_u32 module did not validate the fields in th local privileged attacker to trigger an out-of-bounds read by setting the size fields with a value beyond the array bo disclosure.
CVE-2023-39193	A flaw was found in the Netfilter subsystem in the Linux kernel. The sctp_mt_check did not validate the flag_coun (CAP_NET_ADMIN) attacker to trigger an out-of-bounds read, leading to a crash or information disclosure.
CVE-2023-39194	A flaw was found in the XFRM subsystem in the Linux kernel. The specific flaw exists within the processing of sta the end of an allocated buffer. This flaw allows a local privileged (CAP_NET_ADMIN) attacker to trigger an out-o information disclosure.
CVE-2023-39197	An out-of-bounds read vulnerability was found in Netfilter Connection Tracking (conntrack) in the Linux kernel. T sensitive information via the DCCP protocol.
CVE-2023-39417	IN THE EXTENSION SCRIPT, a SQL Injection vulnerability was found in PostgreSQL if it uses @extowner@, @ quoting construct (dollar quoting, ", or ""). If an administrator has installed files of a vulnerable, trusted, non-bundl CREATE privilege can execute arbitrary code as the bootstrap superuser.
CVE-2023-40577	Alertmanager handles alerts sent by client applications such as the Prometheus server. An attacker with the permiss alerts endpoint could be able to execute arbitrary JavaScript code on the users of Prometheus Alertmanager. This is 0.2.51.

CVE-2023-4206	A use-after-free vulnerability in the Linux kernel's net/sched: cls_route component can be exploited to achieve loca is called on an existing filter, the whole tcf_result struct is always copied into the new instance of the filter. This ca to a class, as tcf_unbind_filter() is always called on the old instance in the success path, decreasing filter_cnt of the deleted, leading to a use-after-free. We recommend upgrading past commit b80b829e9e2c1b3f7aae34855e04d8f6e
CVE-2023-4207	A use-after-free vulnerability in the Linux kernel's net/sched: cls_fw component can be exploited to achieve local p called on an existing filter, the whole tcf_result struct is always copied into the new instance of the filter. This caus to a class, as tcf_unbind_filter() is always called on the old instance in the success path, decreasing filter_cnt of the deleted, leading to a use-after-free. We recommend upgrading past commit 76e42ae831991c828cffa8c37736ebfb83
CVE-2023-4208	A use-after-free vulnerability in the Linux kernel's net/sched: cls_u32 component can be exploited to achieve local called on an existing filter, the whole tcf_result struct is always copied into the new instance of the filter. This caus to a class, as tcf_unbind_filter() is always called on the old instance in the success path, decreasing filter_cnt of the deleted, leading to a use-after-free. We recommend upgrading past commit 3044b16e7c6fe5d24b1cdbcf1bd0a9d92
CVE-2023-42753	An array indexing vulnerability was found in the netfilter subsystem of the Linux kernel. A missing macro could le offset, providing attackers with the primitive to arbitrarily increment/decrement a memory buffer out-of-bound. Th system or potentially escalate their privileges on the system.
CVE-2023-42754	A NULL pointer dereference flaw was found in the Linux kernel ipv4 stack. The socket buffer (skb) was assumed t __ip_options_compile, which is not always the case if the skb is re-routed by ipvs. This issue may allow a local use crash the system.
CVE-2023-42755	A flaw was found in the IPv4 Resource Reservation Protocol (RSVP) classifier in the Linux kernel. The xprt pointe leading to an out-of-bounds read in the `rsvp_classify` function. This issue may allow a local user to crash the syste
CVE-2023-42756	A flaw was found in the Netfilter subsystem of the Linux kernel. A race condition between IPSET_CMD_ADD an panic due to the invocation of `__ip_set_put` on a wrong `set`. This issue may allow a local user to crash the syster
CVE-2023-43785	A vulnerability was found in libX11 due to a boundary condition within the _XkbReadKeySyms() function. This fl bounds read error and read the contents of memory on the system.
CVE-2023-43786	A vulnerability was found in libX11 due to an infinite loop within the PutSubImage() function. This flaw allows a l resources and cause a denial of service condition.
CVE-2023-43787	A vulnerability was found in libX11 due to an integer overflow within the XCreateImage() function. This flaw allo and execute arbitrary code with elevated privileges.
CVE-2023-44981	Authorization Bypass Through User-Controlled Key vulnerability in Apache ZooKeeper. If SASL Quorum Peer au (quorum.auth.enableSasl=true), the authorization is done by verifying that the instance part in SASL authentication The instance part in SASL auth ID is optional and if it's missing, like 'eve@EXAMPLE.COM', the authorization ch arbitrary endpoint could join the cluster and begin propagating counterfeit changes to the leader, essentially giving tree.¬†Quorum Peer authentication is not enabled by default.Users are recommended to upgrade to version 3.9.1, 3 ensure the ensemble election/quorum communication is protected by a firewall as this will mitigate the issue.See th cluster administration.
CVE-2023-4569	A memory leak flaw was found in nft_set_catchall_flush in net/netfilter/nf_tables_api.c in the Linux Kernel. This i double-deactivations of catchall elements, which can result in a memory leak.
CVE-2023-45862	An issue was discovered in drivers/usb/storage/ene_ub6250.c for the ENE UB6250 reader driver in the Linux kern extend beyond the end of an allocation.
CVE-2023-45871	An issue was discovered in drivers/net/ethernet/intel/igb/igb_main.c in the IGB driver in the Linux kernel before 6. frames larger than the MTU.
CVE-2023-46120	The RabbitMQ Java client library allows Java and JVM-based applications to connect to and interact with RabbitM when receiving Message objects. Attackers could send a very large Message causing a memory overflow and trigg may suffer from DoS attacks from RabbitMQ Java client which will ultimately exhaust the memory of the consum 5.18.0.
CVE-2023-46218	This flaw allows a malicious HTTP server to set "super cookies" in curl that are then passed back to more origins th This allows a site to set cookies that then would get sent to different and unrelated sites and domains. It could do th function that verifies a given cookie domain against the Public Suffix List (PSL). For example a cookie could be se lower case hostname `curl.co.uk`, even though `co.uk` is listed as a PSL domain.
CVE-2023-4622	A use-after-free vulnerability in the Linux kernel's af_unix component can be exploited to achieve local privilege e function tries to add data to the last skb in the peer's recv queue without locking the queue. Thus there is a race whe could access an skb locklessly that is being released by garbage collection, resulting in use-after-free. We recomme 790c2f9d15b594350ae9bca7b236f2b1859de02c.
CVE-2023-4623	A use-after-free vulnerability in the Linux kernel's net/sched: sch_hfsc (HFSC qdisc traffic control) component can escalation. If a class with a link-sharing curve (i.e. with the HFSC_FSC flag set) has a parent without a link-sharing on the parent, but vttree_remove() will be skipped in update_vf(). This leaves a dangling pointer that can cause a us commit b3d26c5702c7d6c45456326e56d2ccf3f103e60f.

CVE-2023-46343	In the Linux kernel before 6.5.9, there is a NULL pointer dereference in send_acknowledge in net/nfc/nci/spi.c.
CVE-2023-4921	A use-after-free vulnerability in the Linux kernel's net/sched: sch_qfq component can be exploited to achieve local used as a class of the qfq qdisc, sending network packets triggers use-after-free in qfq_dequeue() due to the incorre checking in agg_dequeue(). We recommend upgrading past commit 8fc134fee27f2263988ae38920bc03da416b03d
CVE-2023-49295	quic-go is an implementation of the QUIC protocol (RFC 9000, RFC 9001, RFC 9002) in Go. An attacker can caus large number of PATH_CHALLENGE frames. The receiver is supposed to respond to each PATH_CHALLENGE The attacker can prevent the receiver from sending out (the vast majority of) these PATH_RESPONSE frames by c selectively acknowledging received packets) and by manipulating the peer's RTT estimate. This vulnerability has b 0.39.4.
CVE-2023-49568	A denial of service (DoS) vulnerability was discovered in go-git versions prior to v5.11. This vulnerability allows a attacks by providing specially crafted responses from a Git server which triggers resource exhaustion in go-git¬†cl filesystem supported by go-git¬†are not affected by this vulnerability. This is a go-git¬†implementation issue and
CVE-2023-49569	A path traversal vulnerability was discovered in go-git versions prior to v5.11. This vulnerability allows an attacke filesystem. In the worse case scenario, remote code execution could be achieved. Applications are only affected if t pkg.go.dev/github.com/go-git/go-billy/v5/osfs#ChrootOS , which is the default when using "Plain" versions of Ope Applications using BoundOS https://pkg.go.dev/github.com/go-git/go-billy/v5/osfs#BoundOS ¬†or in-memory file a go-git¬†implementation issue and does not affect the upstream git¬†cli.
CVE-2023-5043	Ingress nginx annotation injection causes arbitrary command execution.
CVE-2023-5044	Code injection via nginx.ingress.kubernetes.io/permanent-redirect annotation.
CVE-2023-51042	In the Linux kernel before 6.4.12, amdgpu_cs_wait_all_fences in drivers/gpu/drm/amd/amdgpu/amdgpu_cs.c has a
CVE-2023-51043	In the Linux kernel before 6.4.5, drivers/gpu/drm/drm_atomic.c has a use-after-free during a race condition betwee unload.
CVE-2023-52438	In the Linux kernel, the following vulnerability has been resolved: binder: fix use-after-free in shinker's callback Th the shrinker's callback, which means that using alloc->vma pointer isn't safe as it can race with munmap(). As of co zap pages with read mmap_sem in munmap") the mmap lock is downgraded after the vma has been isolated. I was manually adding some delays and triggering page reclaiming through the shrinker's debug sysfs. The following KA == BUG: KASAN: slab- +0x470/0x4b8 Read of size 8 at addr ffff356ed50e50f0 by task bash/478 CPU: 1 PID: 478 Comm: bash Not tainted #70 Hardware name: linux,dummy-virt (DT) Call trace: zap_page_range_single+0x470/0x4b8 binder_alloc_free_p +0x130/0x3b0 list_lru_walk_node+0xc4/0x22c binder_shrink_scan+0x108/0x1dc shrinker_debugfs_scan_write+0 vfs_write+0x1ac/0x758 ksys_write+0xf0/0x1dc __arm64_sys_write+0x6c/0x9c Allocated by task 492: kmem_cac +0x2c/0x190 mmap_region+0x258/0x18bc do_mmap+0x694/0xa60 vm_mmap_pgoff+0x170/0x29c ksys_mmap_ +0xcc/0x144 Freed by task 491: kmem_cache_free+0x17c/0x3c8 vm_area_free_rcu_cb+0x74/0x98 rcu_core+0xa3 __do_softirq+0x2fc/0xd24 Last potentially related work creation: __call_rcu_common.constprop.0+0x6c/0xba0 ca remove_vma+0xe4/0x118 do_vmi_align_munmap.isra.0+0x718/0xb5c do_vmi_munmap+0xdc/0x1fc __vm_munm +0x58/0x7c Fix this issue by performing instead a vma_lookup() which will fail to find the vma that was isolated b that this option has better performance than upgrading to a mmap write lock which would increase contention. Plus removed anyway.
CVE-2023-52439	In the Linux kernel, the following vulnerability has been resolved: uio: Fix use-after-free in uio_open core-1 core-2 -- uio_unregister_device uio_open idev = idr_find() device_unregister(>dev) uio_device_release get_device(&idev->dev) kfree(idev) uio_free_minor(minor) uio_release put_device(&id -- In the core-1 uio_unregister_device(), the device_unregister will kfre But after core-1 device_unregister, put_device and before doing kfree, the core-2 may get_device. Then: 1. After c free for idev. 2. When core-2 do uio_release and put_device, the idev will be double freed. To address this issue, w with minor_lock.
CVE-2023-52456	In the Linux kernel, the following vulnerability has been resolved: serial: imx: fix tx statemachine deadlock When tx statemachine is used to control the RTS pin to drive the RS485 transceiver TX_EN pin. When the TTY port is cl (for instance during userland application crash), imx_uart_shutdown disables the interface and disables the Transm imx_uart_stop_tx bails on an incomplete transmission, to be retriggered by the TC interrupt. This interrupt is disab transitions out of SEND. The statemachine is in deadlock now, and the TX_EN remains low, making the interface incomplete transmission AND whether TC interrupts are enabled before bailing to be retriggered. This makes sure properly set to WAIT_AFTER_SEND.
CVE-2023-52462	In the Linux kernel, the following vulnerability has been resolved: bpf: fix check for attempt to corrupt spilled poin a 1/2/4-byte register, we set slot_type[BPF_REG_SIZE - 1] (plus potentially few more below it, depending on actu has spilled register we need to consult slot_type[7], not slot_type[0]. To avoid the need to remember and double-ch helper.
CVE-2023-52467	In the Linux kernel, the following vulnerability has been resolved: mfd: syscon: Fix null pointer dereference in of_ to dynamically allocated memory which can be NULL upon failure.

CVE-2023-52477	In the Linux kernel, the following vulnerability has been resolved: usb: hub: Guard against accesses to uninitialized usb/core/hub.c and drivers/usb/core/hub.h access fields inside udev->bos without checking if it was allocated and i for whatever reason, udev->bos will be NULL and those accesses will result in a crash: BUG: kernel NULL pointe PGD 0 P4D 0 Oops: 0000 [#1] PREEMPT SMP NOPTI CPU: 5 PID: 17818 Comm: kworker/5:1 Tainted: G W 5.1 <HASH:1f9e 1> Hardware name: Google Kindred/Kindred, BIOS Google_Kindred.12672.413.0 02/03/2021 Work 0010:hub_port_reset+0x193/0x788 Code: 89 f7 e8 20 f7 15 00 48 8b 43 08 80 b8 96 03 00 00 03 75 36 0f b7 88 92 a8 03 00 00 <48> 83 78 18 00 74 19 48 89 df 48 8b 75 b0 ba 02 00 00 00 4c 89 e9 RSP: 0018:ffffab740c53fcf8 EF RBX: ffffa1bc5f678000 RCX: 0000000000000310 RDX: ffffffffffffdff RSI: 0000000000000286 RDI: ffffa1be96 00001b7d5edaa20c R09: ffffffffb005e060 R10: 0000000000000001 R11: 0000000000000000 R12: 000000000000 0000000000000032 R15: 0000000000000000 FS: 0000000000000000(0000) GS:ffffa1be96540000(0000) knlGS:0 ES: 0000 CR0: 0000000080050033 CR2: 0000000000000018 CR3: 000000022e80c005 CR4: 00000000003706e0 hub_activate+0x5b7/0x68f process_one_work+0x1a2/0x487 worker_thread+0x11a/0x288 kthread+0x13a/0x152 ? kthread_associate_blkcg+0x70/0x70 ret_from_fork+0x1f/0x30 Fall back to a default behavior if the BOS descripto functionalities that depend on it: LPM support checks, Super Speed capabilitiy checks, U1/U2 states setup.
CVE-2023-52480	In the Linux kernel, the following vulnerability has been resolved: ksmbd: fix race condition between session looku ksmbd_session_lookup | smb2_sess_setup sess = xa_load | | | xa_erase(&conn->sessions, sess->id); | | ksmbd_sessi sess->last_active = jiffies | + This patch add rwsem to fix race condition between ksmbd_session_lookup and ksmb
CVE-2023-52484	In the Linux kernel, the following vulnerability has been resolved: iommu/arm-smmu-v3: Fix soft lockup triggered When running an SVA case, the following soft lockup is triggered: ---------- - CPU#244 stuck for 26s! pstate: 83400009 (Nzcv daif +PAN -UAO +TCO +DIT -SSBS BTYPE=--) pc : arm_smm lr : arm_smmu_cmdq_issue_cmdlist+0x150/0xa50 sp : ffff8000d83ef290 x29: ffff8000d83ef290 x28: 000000003b x26: ffff8000d83ef3c0 x25: da86c0812194a0e8 x24: 0000000000000000 x23: 0000000000000040 x22: ffff8000d8 x20: 0000000000000001 x19: ffff0000c6398080 x18: 0000000000000000 x17: 0000000000000000 x16: 0000000 x14: ffff3000b4a30888 x13: ffff3000b4a3cf60 x12: 0000000000000000 x11: 0000000000000000 x10: 0000000000 0000000000000000 x7 : 0000000000000000 x6 : 0000000000048cfa x5 : 0000000000000000 x4 : 0000000000000 0000000080000000 x1 : 0000000000000000 x0 : 0000000000000001 Call trace: arm_smmu_cmdq_issue_cmdlist +0x118/0x254 arm_smmu_tlb_inv_range_asid+0x6c/0x130 arm_smmu_mm_invalidate_range+0xa0/0xa4 __mmu +0x88/0x120 unmap_vmas+0x194/0x1e0 unmap_region+0xb4/0x144 do_mas_align_munmap+0x290/0x490 do_m +0xa8/0x19c __arm64_sys_munmap+0x28/0x50 invoke_syscall+0x78/0x11c el0_svc_common.constprop.0+0x58 +0x2c/0xd4 el0t_64_sync_handler+0x114/0x140 el0t_64_sync+0x1a4/0x1a8 ---------- rc1 the arm_smmu_mm_invalidate_range above is renamed to "arm_smmu_mm_arch_invalidate_secondary_tlbs", 06ff87bae8d3 ("arm64: mm: remove unused functions and variable protoypes") fixed a similar lockup on the CPU too, since arm_smmu_mm_arch_invalidate_secondary_tlbs() is called typically next to MMU tlb flush function, e.g { __flush_tlb_range { // check MAX_TLBI_OPS } } mmu_notifier_arch_invalidate_secondary_tlbs { arm_smmu_ not check MAX_TLBI_OPS } } } Clone a CMDQ_MAX_TLBI_OPS from the MAX_TLBI_OPS in tlbflush.h, sin page table, so it makes sense to align with the tlbflush code. Then, replace per-page TLBI commands with a single hits this threshold.
CVE-2023-52494	In the Linux kernel, the following vulnerability has been resolved: bus: mhi: host: Add alignment check for event r event ring read pointer by "is_valid_ring_ptr" to make sure it is in the buffer range, but there is another risk the poi expecting event ring elements are 128 bits(struct mhi_ring_element) aligned, an unaligned read pointer could lead t memory corruption. So add a alignment check for event ring read pointer.
CVE-2023-52502	In the Linux kernel, the following vulnerability has been resolved: net: nfc: fix races in nfc_llcp_sock_get() and nfc race in nfc_llcp_sock_get(), leading to UAF. Getting a reference on the socket found in a lookup while holding a lo nfc_llcp_sock_get_sn() has a similar problem. Finally nfc_llcp_recv_snl() needs to make sure the socket found by
CVE-2023-52503	In the Linux kernel, the following vulnerability has been resolved: tee: amdtee: fix use-after-free vulnerability in ar condition in amdtee_close_session that may cause use-after-free in amdtee_open_session. For instance, if a session free this session via: kref_put(&sess->refcount, destroy_session); the reference count will get decremented, and the However, if in another thread, amdtee_open_session() is called before destroy_session() has completed execution, freed up later in destroy_session() leading to use-after-free in amdtee_open_session. To fix this issue, treat decreme from session list in destroy_session() as a critical section, so that it is executed atomically.
CVE-2023-52507	In the Linux kernel, the following vulnerability has been resolved: nfc: nci: assert requested protocol is valid The p the protocol is supported. Assert the provided protocol is less than the maximum defined so it doesn't potentially pe clearer error for undefined protocols vs unsupported ones.
CVE-2023-52509	In the Linux kernel, the following vulnerability has been resolved: ravb: Fix use-after-free issue in ravb_tx_timeou call cancel_work_sync(). Otherwise, ravb_tx_timeout_work() is possible to use the freed priv after ravb_remove() ravb_tx_timeout() ravb_remove() unregister_netdev() free_netdev(ndev) // free priv ravb_tx_timeout_work() // use so that ravb_stop() is called. And, after phy_stop() is called, netif_carrier_off() is also called. So that .ndo_tx_timeo
CVE-2023-52510	In the Linux kernel, the following vulnerability has been resolved: ieee802154: ca8210: Fix a potential UAF in ca8 ca8210_register_ext_clock(), it calls clk_unregister() to release priv->clk and returns an error. However, the caller where priv->clk is freed again in ca8210_unregister_ext_clock(). In this case, a use-after-free may happen in the se by removing the first clk_unregister(). Also, priv->clk could be an error code on failure of clk_register_fixed_rate(in ca8210_unregister_ext_clock().

CVE-2023-52513	In the Linux kernel, the following vulnerability has been resolved: RDMA/siw: Fix connection failure handling In the newly created endpoint unlinks the listening endpoint and is ready to be dropped. This special case was not han TCP socket close, causing a NULL dereference crash in siw_cm_work_handler() when dereferencing a NULL liste timeout, if immediate MPA request processing fails. This patch furthermore simplifies MPA processing in general: sk_data_ready() upcall is now surpressed, if the socket is already moved out of TCP_ESTABLISHED state.
CVE-2023-52524	In the Linux kernel, the following vulnerability has been resolved: net: nfc: llcp: Add lock when modifying device held when modifying it, or the list could become corrupted, as syzbot discovered.
CVE-2023-52525	In the Linux kernel, the following vulnerability has been resolved: wifi: mwifiex: Fix oob check condition in mwifi path trying to access the rfc1042 headers when the buffer is too small, so the driver can still process packets withou
CVE-2023-52527	In the Linux kernel, the following vulnerability has been resolved: ipv4, ipv6: Fix handling of transhdrlen in __ip{, transhdrlen in length is a problem when the packet is partially filled (e.g. something like send(MSG_MORE) happe to an IPv4 or IPv6 packet as we don't want to repeat the transport header or account for it twice. This can happen u splicing into an L2TP socket. The symptom observed is a warning in __ip6_append_data(): WARNING: CPU: 1 P __ip6_append_data.isra.0+0x1be8/0x47f0 net/ipv6/ip6_output.c:1800 that occurs when MSG_SPLICE_PAGES is partially occupied skbuff. The warning occurs when 'copy' is larger than the amount of data in the message iterator includes the transport header length when it shouldn't. This can be triggered by, for example: sfd = socket(AF_INE bind(sfd, ...); // ::1 connect(sfd, ...); // ::1 port 7 send(sfd, buffer, 4100, MSG_MORE); sendfile(sfd, dfd, NULL, 10 into the length if the write queue is empty in l2tp_ip6_sendmsg(), analogously to how UDP does things. l2tp_ip_se problem as it builds the UDP packet itself.
CVE-2023-52528	In the Linux kernel, the following vulnerability has been resolved: net: usb: smsc75xx: Fix uninit-value access in _ the following uninit-value access issue: ========== in smsc75xx_wait_ready drivers/net/usb/smsc75xx.c:975 [inline] BUG: KMSAN: uninit-value in smsc75xx_bind+ smsc75xx.c:1482 CPU: 0 PID: 8696 Comm: kworker/0:3 Not tainted 5.8.0-rc5-syzkaller #0 Hardware name: Goog Compute Engine, BIOS Google 01/01/2011 Workqueue: usb_hub_wq hub_event Call Trace: __dump_stack lib/du +0x21c/0x280 lib/dump_stack.c:118 kmsan_report+0xf7/0x1e0 mm/kmsan/kmsan_report.c:121 __msan_warning+ smsc75xx_wait_ready drivers/net/usb/smsc75xx.c:975 [inline] smsc75xx_bind+0x5c9/0x11e0 drivers/net/usb/sms drivers/net/usb/usbnet.c:1737 usb_probe_interface+0xece/0x1550 drivers/usb/core/driver.c:374 really_probe+0xf2 driver_probe_device+0x293/0x390 drivers/base/dd.c:701 __device_attach_driver+0x63f/0x830 drivers/base/dd.c:8 base/bus.c:431 __device_attach+0x4e2/0x7f0 drivers/base/dd.c:873 device_initial_probe+0x4a/0x60 drivers/base/c drivers/base/bus.c:491 device_add+0x3b0e/0x40d0 drivers/base/core.c:2680 usb_set_configuration+0x380f/0x3f1 usb_generic_driver_probe+0x138/0x300 drivers/usb/core/generic.c:241 usb_probe_device+0x311/0x490 drivers/u +0xf20/0x20b0 drivers/base/dd.c:529 driver_probe_device+0x293/0x390 drivers/base/dd.c:701 __device_attach_d bus_for_each_drv+0x2ca/0x3f0 drivers/base/bus.c:431 __device_attach+0x4e2/0x7f0 drivers/base/dd.c:873 device dd.c:920 bus_probe_device+0x177/0x3d0 drivers/base/bus.c:491 device_add+0x3b0e/0x40d0 drivers/base/core.c:2 drivers/usb/core/hub.c:2554 hub_port_connect drivers/usb/core/hub.c:5208 [inline] hub_port_connect_change driv drivers/usb/core/hub.c:5494 [inline] hub_event+0x5e7b/0x8a70 drivers/usb/core/hub.c:5576 process_one_work+0x worker_thread+0x10bc/0x2730 kernel/workqueue.c:2415 kthread+0x551/0x590 kernel/kthread.c:292 ret_from_for Local variable ----buf.i87@smsc75xx_bind created at: __smsc75xx_read_reg drivers/net/usb/smsc75xx.c:83 [inlin smsc75xx.c:968 [inline] smsc75xx_bind+0x485/0x11e0 drivers/net/usb/smsc75xx.c:1482 __smsc75xx_read_reg d smsc75xx_wait_ready drivers/net/usb/smsc75xx.c:968 [inline] smsc75xx_bind+0x485/0x11e0 drivers/net/usb/sms usbnet_read_cmd() reads less bytes than requested (zero byte in the reproducer). In this case, 'buf' is not properly fi ENODATA if usbnet_read_cmd() reads less bytes than requested.
CVE-2023-52564	In the Linux kernel, the following vulnerability has been resolved: Revert "tty: n_gsm: fix UAF in gsm_cleanup_m 9b9c8195f3f0d74a826077fc1c01b9ee74907239. The commit above is reverted as it did not solve the original issue virtual ttys by calling gsm_dlci_release() for each available DLCI. There, dlci_put() is called to decrease the refere which finally calls gsm_dlci_free(). This already clears the pointer which is being checked in gsm_cleanup_mux() it is not necessary to clear this pointer in gsm_cleanup_mux() as done in the reverted commit. The commit introduc __die+0x1f/0x70 ? page_fault_oops+0x156/0x420 ? search_exception_tables+0x37/0x50 ? fixup_exception+0x21 asm_exc_page_fault+0x26/0x30 ? tty_port_put+0x19/0xa0 gsmtty_cleanup+0x29/0x80 [n_gsm] release_one_tty+0 worker_thread+0x4c/0x3d0 ? __pfx_worker_thread+0x10/0x10 kthread+0xe1/0x110 ? __pfx_kthread+0x10/0x10 +0x10/0x10 ret_from_fork_asm+0x1b/0x30 </TASK> The actual issue is that nothing guards dlci_put() from bein was triggered but did not yet finished calling gsm_dlci_free().
CVE-2023-52566	In the Linux kernel, the following vulnerability has been resolved: nilfs2: fix potential use after free in nilfs_gccach nilfs_gccache_submit_read_data(), brelse(bh) is called to drop the reference count of bh when the call to nilfs_dat_ 0 and its owner page gets unlocked, bh may be freed. However, bh->b_page is dereferenced to put the page after th This patch moves the release operation after unlocking and putting the page. NOTE: The function in question is on current userland tools, address translation using DAT does not occur in that function, so the code path that causes t possible to run that code path by intentionally modifying the userland GC library or by calling the GC ioctl directly to the commit log]
CVE-2023-52567	In the Linux kernel, the following vulnerability has been resolved: serial: 8250_port: Check IRQ data before use In (irq = 0) and IIR register shows that an interrupt happened in the 8250 hardware the IRQ data can be NULL. In suc came to this path from the timer interrupt and quite likely system is already awake. Without this fix we have got an base_baud = 115200) is a 16550A ... BUG: kernel NULL pointer dereference, address: 0000000000000010 RIP: 0 Trace: ? serial8250_handle_irq+0x7c/0x240 ? __pfx_serial8250_timeout+0x10/0x10

CVE-2023-52573	In the Linux kernel, the following vulnerability has been resolved: net: rds: Fix possible NULL-pointer dereference check, if conn pointer exists before dereferencing it as rdma_set_service_type() argument Found by Linux Verifica
CVE-2023-52574	In the Linux kernel, the following vulnerability has been resolved: team: fix null-ptr-deref when team device type i with reproducer [1]. BUG: kernel NULL pointer dereference, address: 0000000000000228 ... RIP: 0010:vlan_dev_ Trace: <TASK> ? __die+0x24/0x70 ? page_fault_oops+0x82/0x150 ? exc_page_fault+0x69/0x150 ? asm_exc_pag +0x35/0x140 [8021q] ? vlan_dev_hard_header+0x8e/0x140 [8021q] neigh_connected_output+0xb2/0x100 ip6_fin +0x43/0xc0 ? ip6_mtu+0x46/0x80 ip6_finish_output+0x2a/0xb0 mld_sendpack+0x18f/0x250 mld_ifc_work+0x3 worker_thread+0x4d/0x2f0 ? __pfx_worker_thread+0x10/0x10 kthread+0xe5/0x120 ? __pfx_kthread+0x10/0x10 +0x10/0x10 ret_from_fork_asm+0x1b/0x30 [1] $ teamd -t team0 -d -c '{"runner": {"name": "loadbalance"}}' $ ip add link t-dummy name t-dummy.100 type vlan id 100 $ ip link add name t-nlmon type nlmon $ ip link set t-nlmon $ ip link set t-dummy up $ ip link set team0 up $ ip link set t-dummy.100 down $ ip link set t-dummy.100 master t device and team device type is changed from non-ether to ether, header_ops of team device is changed to vlan_hea null-ptr-deref for vlan->real_dev in vlan_dev_hard_header() because team device is not a vlan device. Cache eth_h cached header_ops to header_ops of team net device when its type is changed from non-ether to ether to fix the bug
CVE-2023-52577	In the Linux kernel, the following vulnerability has been resolved: dccp: fix dccp_v4_err()/dccp_v6_err() again dh- "struct dccp_hdr", not in the "byte 7" as Jann claimed. We need to make sure the ICMP messages are big enough, u assumptions). syzbot reported: BUG: KMSAN: uninit-value in pskb_may_pull_reason include/linux/skbuff.h:2667 pskb_may_pull include/linux/skbuff.h:2681 [inline] BUG: KMSAN: uninit-value in dccp_v6_err+0x426/0x1aa0 n include/linux/skbuff.h:2667 [inline] pskb_may_pull include/linux/skbuff.h:2681 [inline] dccp_v6_err+0x426/0x1a +0x4c7/0x880 net/ipv6/icmp.c:867 icmpv6_rcv+0x19d5/0x30d0 ip6_protocol_deliver_rcu+0xda6/0x2a60 net/ipv6 ip6_input.c:483 [inline] NF_HOOK include/linux/netfilter.h:304 [inline] ip6_input+0x15d/0x430 net/ipv6/ip6_inp ipv6/ip6_input.c:586 dst_input include/net/dst.h:468 [inline] ip6_rcv_finish+0x5db/0x870 net/ipv6/ip6_input.c:79 [inline] ipv6_rcv+0xda/0x390 net/ipv6/ip6_input.c:310 __netif_receive_skb_one_core net/core/dev.c:5523 [inline] core/dev.c:5637 netif_receive_skb_internal net/core/dev.c:5723 [inline] netif_receive_skb+0x58/0x660 net/core/de tun_get_user+0x564c/0x6940 drivers/net/tun.c:2002 tun_chr_write_iter+0x3af/0x5d0 drivers/net/tun.c:2048 call_w new_sync_write fs/read_write.c:491 [inline] vfs_write+0x8ef/0x15c0 fs/read_write.c:584 ksys_write+0x20f/0x4c0 fs/read_write.c:649 [inline] __se_sys_write fs/read_write.c:646 [inline] __x64_sys_write+0x93/0xd0 fs/read_write common.c:50 [inline] do_syscall_64+0x41/0xc0 arch/x86/entry/common.c:80 entry_SYSCALL_64_after_hwfram slab_post_alloc_hook+0x12f/0xb70 mm/slab.h:767 slab_alloc_node mm/slub.c:3478 [inline] kmem_cache_alloc_ kmalloc_reserve+0x13d/0x4a0 net/core/skbuff.c:559 __alloc_skb+0x318/0x740 net/core/skbuff.c:650 alloc_skb in alloc_skb_with_frags+0xc8/0xbd0 net/core/skbuff.c:6313 sock_alloc_send_pskb+0xa80/0xbf0 net/core/sock.c:279 [inline] tun_get_user+0x23cf/0x6940 drivers/net/tun.c:1846 tun_chr_write_iter+0x3af/0x5d0 drivers/net/tun.c:204 [inline] new_sync_write fs/read_write.c:491 [inline] vfs_write+0x8ef/0x15c0 fs/read_write.c:584 ksys_write+0x20 fs/read_write.c:649 [inline] __se_sys_write fs/read_write.c:646 [inline] __x64_sys_write+0x93/0xd0 fs/read_write common.c:50 [inline] do_syscall_64+0x41/0xc0 arch/x86/entry/common.c:80 entry_SYSCALL_64_after_hwfram syz-executor153 Not tainted 6.6.0-rc1-syzkaller-00014-ga747acc0b752 #0 Hardware name: Google Google Compu Google 08/04/2023
CVE-2023-52578	In the Linux kernel, the following vulnerability has been resolved: net: bridge: use DEV_STATS_INC() syzbot/KC br_handle_frame_finish() [1] This function can run from multiple cpus without mutual exclusion. Adopt SMP safe >stats fields. Handles updates to dev->stats.tx_dropped while we are at it. [1] BUG: KCSAN: data-race in br_hand read-write to 0xffff8881374b2178 of 8 bytes by interrupt on cpu 1: br_handle_frame_finish+0xd4f/0xef0 net/bridg +0x1ed/0x220 br_nf_pre_routing_finish_ipv6+0x50f/0x540 NF_HOOK include/linux/netfilter.h:304 [inline] br_nf bridge/br_netfilter_ipv6.c:178 br_nf_pre_routing+0x526/0xba0 net/bridge/br_netfilter_hooks.c:508 nf_hook_entry [inline] nf_hook_bridge_pre net/bridge/br_input.c:272 [inline] br_handle_frame+0x4c9/0x940 net/bridge/br_input. +0xa8a/0x21e0 net/core/dev.c:5417 __netif_receive_skb_one_core net/core/dev.c:5521 [inline] __netif_receive_sk process_backlog+0x21f/0x380 net/core/dev.c:5965 __napi_poll+0x60/0x3b0 net/core/dev.c:6527 napi_poll net/cor +0x32b/0x750 net/core/dev.c:6727 __do_softirq+0xc1/0x265 kernel/softirq.c:553 run_ksoftirqd+0x17/0x20 kernel +0x30a/0x4a0 kernel/smpboot.c:164 kthread+0x1d7/0x210 kernel/kthread.c:388 ret_from_fork+0x48/0x60 arch/x8 +0x11/0x20 arch/x86/entry/entry_64.S:304 read-write to 0xffff8881374b2178 of 8 bytes by interrupt on cpu 0: br_ net/bridge/br_input.c:189 br_nf_hook_thresh+0x1ed/0x220 br_nf_pre_routing_finish_ipv6+0x50f/0x540 NF_HOO [inline] br_nf_pre_routing_ipv6+0x1e3/0x2a0 net/bridge/br_netfilter_ipv6.c:178 br_nf_pre_routing+0x526/0xba0 nf_hook_entry_hookfn include/linux/netfilter.h:144 [inline] nf_hook_bridge_pre net/bridge/br_input.c:272 [inline] br_input.c:417 __netif_receive_skb_core+0xa8a/0x21e0 net/core/dev.c:5417 __netif_receive_skb_one_core net/co +0x57/0x1b0 net/core/dev.c:5637 process_backlog+0x21f/0x380 net/core/dev.c:5965 __napi_poll+0x60/0x3b0 net dev.c:6594 [inline] net_rx_action+0x32b/0x750 net/core/dev.c:6727 __do_softirq+0xc1/0x265 kernel/softirq.c:553 __local_bh_enable_ip+0x64/0x70 kernel/softirq.c:381 __raw_spin_unlock_bh include/linux/spinlock_api_smp.h:1 +0x36/0x40 kernel/locking/spinlock.c:210 spin_unlock_bh include/linux/spinlock.h:396 [inline] batadv_tt_local_p translation-table.c:1356 batadv_tt_purge+0x2b/0x630 net/batman-adv/translation-table.c:3560 process_one_work process_scheduled_works+0x5b8/0xa30 kernel/workqueue.c:2703 worker_thread+0x525/0x730 kernel/workqueue kthread.c:388 ret_from_fork+0x48/0x60 arch/x86/kernel/process.c:147 ret_from_fork_asm+0x11/0x20 arch/x86/e 0x00000000000d7190 -> 0x00000000000d7191 Reported by Kernel Concurrency Sanitizer on: CPU: 0 PID: 1484 rc1-syzkaller-00236-gad8a69f361b9 #0

CVE-2023-52608	In the Linux kernel, the following vulnerability has been resolved: firmware: arm_scmi: Check mailbox/SMT chan completion interrupt the shared memory area is accessed to retrieve the message header at first and then, if the mes which is still pending, the related payload is fetched too. When an SCMI command times out the channel ownershi a late reply is received and, as a consequence, any further transmission attempt remains pending, waiting for the ch Once that late reply is received the channel ownership is given back to the agent and any pending request is then al area of the just delivered late reply; then the wait for the reply to the new request starts. It has been observed that th can be wrongly associated with the freshly enqueued request: when that happens the SCMI stack in-flight lookup p message header now present in the SMT area is related to the new pending transaction, even though the real reply h the A2P channel can be detected by looking at the channel status bits: a genuine reply from the platform will have s completion IRQ. Add a consistency check to validate such condition in the A2P ISR.
CVE-2023-52628	In the Linux kernel, the following vulnerability has been resolved: netfilter: nftables: exthdr: fix 4-byte stack OOB dst[len / 4] can write past the destination array which leads to stack corruption. This construct is necessary to clean NOT a multiple of the register size, so make it conditional just like nft_payload.c does. The bug was added in 4.1 c and ip option support was added. Bug reported by Zero Day Initiative project (ZDI-CAN-21950, ZDI-CAN-21951,
CVE-2023-52654	In the Linux kernel, the following vulnerability has been resolved: io_uring/af_unix: disable sending io_uring over lots of problems for io_uring in the past, and it still doesn't work exactly right and races with unix_stream_read_ge disallow sending io_uring files via sockets via SCM_RIGHT, so there are no possible cycles invloving registered f the io_uring side unnecessary.
CVE-2023-52655	In the Linux kernel, the following vulnerability has been resolved: usb: aqc111: check packet for fixup for true lim 0 and sizeof(u64) the value passed to skb_trim() as length will wrap around ending up as some very large value. Th located at that position, which will either oops or process some random value. The fix is to check against sizeof(u6 does. The issue exists since the introduction of the driver.
CVE-2023-52670	In the Linux kernel, the following vulnerability has been resolved: rpmsg: virtio: Free driver_override when rpmsg when rpmsg_remove(), otherwise the following memory leak will occur: unreferenced object 0xffff0000d55d7080 pid 56, jiffies 4294893188 (age 214.272s) hex dump (first 32 bytes): 72 70 6d 73 67 5f 6e 73 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ..... backtrace: [<000000009c94c9c1>] __kmem_cache_alloc_node+0x1f8 __kmalloc_node_track_caller+0x44/0x70 [<00000000228a60c3>] kstrndup+0x4c/0x90 [<0000000077158695>] d [<000000003e9c4ea5>] rpmsg_register_device_override+0x98/0x170 [<000000001c0c89a8>] rpmsg_ns_register_ rpmsg_probe+0x2e0/0x3ec [<00000000e65a68df>] virtio_dev_probe+0x1c0/0x280 [<00000000443331cc>] really __driver_probe_device+0x78/0xe0 [<00000000a41c9a5b>] driver_probe_device+0xd8/0x160 [<000000009c3bd5 [<0000000043cd7614>] bus_for_each_drv+0x7c/0xd4 [<000000003b929a36>] __device_attach+0x9c/0x19c [<00 +0x14/0x20 [<000000003c999637>] bus_probe_device+0xa0/0xac
CVE-2023-52672	In the Linux kernel, the following vulnerability has been resolved: pipe: wakeup wr_wait after setting max_usage C notification queue support") a regression was introduced that would lock up resized pipes under certain conditions. resizing the pipe ring size was moved to a different function, doing that moved the wakeup for pipe->wr_wait befo pipe was full before the resize occured it would result in the wakeup never actually triggering pipe_write. Set @ma writers if this isn't a watch queue. [Christian Brauner <brauner@kernel.org>: rewrite to account for watch queues]
CVE-2023-52675	In the Linux kernel, the following vulnerability has been resolved: powerpc/imc-pmu: Add a null pointer check in u a pointer to dynamically allocated memory which can be NULL upon failure.
CVE-2023-52677	In the Linux kernel, the following vulnerability has been resolved: riscv: Check if the code to patch lies in the exit s vmalloc_to_page() which panics since the address does not lie in the vmalloc region.
CVE-2023-52682	In the Linux kernel, the following vulnerability has been resolved: f2fs: fix to wait on block writeback for post_rea not encrypted, it missed to call f2fs_wait_on_block_writeback() to wait for GCed page writeback in IPU write path - do_garbage_collect - gc_data_segment - move_data_block - f2fs_submit_page_write migrate normal cluster's blo f2fs_write_single_data_page - f2fs_do_write_data_page - f2fs_inplace_write_data - f2fs_submit_page_bio IRQ - f data due to out-of-order GC and common IO. - f2fs_read_end_io
CVE-2023-52686	In the Linux kernel, the following vulnerability has been resolved: powerpc/powernv: Add a null pointer check in c to dynamically allocated memory which can be NULL upon failure.
CVE-2023-52690	In the Linux kernel, the following vulnerability has been resolved: powerpc/powernv: Add a null pointer check to s pointer to dynamically allocated memory which can be NULL upon failure. Add a null pointer check, and release '
CVE-2023-52691	In the Linux kernel, the following vulnerability has been resolved: drm/amd/pm: fix a double-free in si_dpm_init V >pm.dpm.dyn_state.vddc_dependency_on_dispclk.entries fails, amdgpu_free_extended_power_table is called to fr control flow returns to si_dpm_sw_init, it goes to label dpm_failed and calls si_dpm_fini, which calls amdgpu_free fields again. Thus a double-free is triggered.
CVE-2023-52693	In the Linux kernel, the following vulnerability has been resolved: ACPI: video: check for error while searching fo called in acpi_video_dev_register_backlight() fails, for example, because acpi_ut_acquire_mutex() fails inside acpi (uninitialized) acpi_parent handle being passed to acpi_get_pci_dev() for detecting the parent pci device. Check ac only in case of success. Found by Linux Verification Center (linuxtesting.org) with SVACE.

CVE-2023-52694	In the Linux kernel, the following vulnerability has been resolved: drm/bridge: tpd12s015: Drop buggy __exit anno tpd12s015_remove() marked with __exit this function is discarded when the driver is compiled as a built-in. The re no cleanup done which results in resource leakage or worse.
CVE-2023-52696	In the Linux kernel, the following vulnerability has been resolved: powerpc/powernv: Add a null pointer check in c pointer to dynamically allocated memory which can be NULL upon failure.
CVE-2023-52703	In the Linux kernel, the following vulnerability has been resolved: net/usb: kalmia: Don't pass act_len in usb_bulk_ kalmia_send_init_packet() is uninitialized when passing it to the first usb_bulk_msg error path. Jiri Pirko noted tha that the value that would be printed in the second error path would be the value of act_len from the first call to usb_ pass act_len to the usb_bulk_msg error paths. 1: https://lore.kernel.org/lkml/Y9pY61y1nwTuzMOa@nanopsycho/
CVE-2023-52705	In the Linux kernel, the following vulnerability has been resolved: nilfs2: fix underflow in second superblock posit NILFS_SB2_OFFSET_BYTES, which computes the position of the second superblock, underflows when the argu Therefore, when using this macro, it is necessary to check in advance that the device size is not less than a lower li occur. The current nilfs2 implementation lacks this check, causing out-of-bound block access when mounting devic loop0, sector 36028797018963960 op 0x0:(READ) flags 0x0 phys_seg 1 prio class 2 NILFS (loop0): unable to rea In addition, when trying to resize the filesystem to a size below 4096 bytes, this underflow occurs in nilfs_resize_fs to nilfs_sufile_resize(), corrupting parameters such as the number of segments in superblocks. This causes excessiv during a subsequent resize ioctl, causing semaphore ns_segctor_sem to block for a long time and hang the writer th more than 143 seconds. Not tainted 6.2.0-rc8-syzkaller-00015-gf6feea56f66d #0 "echo 0 > /proc/sys/kernel/hung_t task:segctord state:D stack:23456 pid:5067 ppid:2 flags:0x00004000 Call Trace: <TASK> context_switch kernel/s +0x1409/0x43f0 kernel/sched/core.c:6606 schedule+0xc3/0x190 kernel/sched/core.c:6682 rwsem_down_write_slo rwsem.c:1190 nilfs_transaction_lock+0x25c/0x4f0 fs/nilfs2/segment.c:357 nilfs_segctor_thread_construct fs/nilfs2 +0x52f/0x1140 fs/nilfs2/segment.c:2570 kthread+0x270/0x300 kernel/kthread.c:376 ret_from_fork+0x1f/0x30 arc Call Trace: <TASK> folio_mark_accessed+0x51c/0xf00 mm/swap.c:515 __nilfs_get_page_block fs/nilfs2/page.c: fs/nilfs2/page.c:61 nilfs_mdt_submit_block+0xd7/0x8f0 fs/nilfs2/mdt.c:121 nilfs_mdt_read_block+0xeb/0x430 fs/ +0x12d/0xbb0 fs/nilfs2/mdt.c:251 nilfs_sufile_get_segment_usage_block fs/nilfs2/sufile.c:92 [inline] nilfs_sufile_ nilfs_sufile_resize+0x7a3/0x12b0 fs/nilfs2/sufile.c:777 nilfs_resize_fs+0x20c/0xed0 fs/nilfs2/super.c:422 nilfs_ioc nilfs_ioctl+0x137c/0x2440 fs/nilfs2/ioctl.c:1301 ... This fixes these issues by inserting appropriate minimum devic depending on where the macro is used.
CVE-2023-52708	In the Linux kernel, the following vulnerability has been resolved: mmc: mmc_spi: fix error handling in mmc_spi_ doesn't need to call mmc_remove_host(), or it will cause null-ptr-deref, because of deleting a not added device in m 'fail_glue_init', if mmc_add_host() fails, and change the label 'fail_add_host' to 'fail_gpiod_request'.
CVE-2023-52733	In the Linux kernel, the following vulnerability has been resolved: s390/decompressor: specify __decompress() buf to __decompress() didn't specify "out_len" parameter on many architectures including s390, expecting that no write performed. This has changed since commit 2aa14b1ab2c4 ("zstd: import usptream v1.5.2") which includes zstd libr of dctx by reutilizing dst buffer (#2751)"). Now zstd decompression code might store literal buffer in the unwritten "out_len" is not set, it is considered to be unlimited and hence free to use for optimization needs. On s390 this migh placed right after the decompressor buffer. Luckily the size of uncompressed kernel image is already known to the specify it in the "out_len" parameter.
CVE-2023-52742	In the Linux kernel, the following vulnerability has been resolved: net: USB: Fix wrong-direction WARNING in pl the plusb network driver: A zero-length control-OUT transfer was treated as a read instead of a write. In modern ke usb 1-1: BOGUS control dir, pipe 80000280 doesn't match bRequestType c0 WARNING: CPU: 0 PID: 4645 at dri +0x14a7/0x1880 drivers/usb/core/urb.c:411 Modules linked in: CPU: 1 PID: 4645 Comm: dhcpcd Not tainted 6.2.0 Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google 01/12/2023 RIP: 0010:u usb/core/urb.c:411 ... Call Trace: <TASK> usb_start_wait_urb+0x101/0x4b0 drivers/usb/core/message.c:58 usb_in message.c:102 [inline] usb_control_msg+0x320/0x4a0 drivers/usb/core/message.c:153 __usbnet_read_cmd+0xb9/ usbnet_read_cmd+0x96/0xf0 drivers/net/usb/usbnet.c:2068 pl_vendor_req drivers/net/usb/plusb.c:60 [inline] pl_se plusb.c:75 [inline] pl_reset+0x2f/0xf0 drivers/net/usb/plusb.c:85 usbnet_open+0xcc/0x5d0 drivers/net/usb/usbnet.c dev.c:1417 __dev_change_flags+0x587/0x750 net/core/dev.c:8530 dev_change_flags+0x97/0x170 net/core/dev.c:8 devinet.c:1147 inet_ioctl+0x33f/0x380 net/ipv4/af_inet.c:979 sock_do_ioctl+0xcc/0x230 net/socket.c:1169 sock_io fs/ioctl.c:51 [inline] __do_sys_ioctl fs/ioctl.c:870 [inline] __se_sys_ioctl fs/ioctl.c:856 [inline] __x64_sys_ioctl+0x arch/x86/entry/common.c:50 [inline] do_syscall_64+0x39/0xb0 arch/x86/entry/common.c:80 entry_SYSCALL_64 usbnet_write_cmd() instead of usbnet_read_cmd() and remove the USB_DIR_IN flag.
CVE-2023-52746	In the Linux kernel, the following vulnerability has been resolved: xfrm/compat: prevent potential spectre v1 gadge nla_type(nla); if (type > XFRMA_MAX) { return -EOPNOTSUPP; } @type is then used as an array index and can (nla_len(nla) < compat_policy[type].len) { array_index_nospec() can be used to prevent leaking content of kernel m
CVE-2023-52752	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix use-after-free bug in cifs_debug_ that are being teared down (e.g. @ses->ses_status == SES_EXITING) in cifs_debug_data_proc_show() to avoid us following GPF when reading from /proc/fs/cifs/DebugData while mounting and umounting [816.251274] general p canonical address 0x6b6b6b6b6b6b6d81: 0000 [#1] PREEMPT SMP NOPTI ... [816.260138] Call Trace: [816.26 die_addr+0x36/0x90 [816.260762] ? exc_general_protection+0x1b3/0x410 [816.261126] ? asm_exc_general_pro cifs_debug_tcon+0xbd/0x240 [cifs] [816.261878] ? cifs_debug_tcon+0xab/0x240 [cifs] [816.262249] cifs_debug [816.262689] ? seq_read_iter+0x379/0x470 [816.262995] seq_read_iter+0x118/0x470 [816.263291] proc_reg_re srso_alias_return_thunk+0x5/0x7f [816.263945] vfs_read+0x201/0x350 [816.264211] ksys_read+0x75/0x100 [8 [816.264750] entry_SYSCALL_64_after_hwframe+0x6e/0xd8 [816.265135] RIP: 0033:0x7fd5e669d381

CVE-2023-52753	In the Linux kernel, the following vulnerability has been resolved: drm/amd/display: Avoid NULL dereference of t whether assigned timing generator is NULL or not before accessing its funcs to prevent NULL dereference.
CVE-2023-52810	In the Linux kernel, the following vulnerability has been resolved: fs/jfs: Add check for negative db_l2nbperpage l and the minimum legal value should be 0, not negative. In the case of l2nbperpage being negative, an error will occ Syzbot reported this bug: UBSAN: shift-out-of-bounds in fs/jfs/jfs_dmap.c:799:12 shift exponent -16777216 is neg
CVE-2023-52827	In the Linux kernel, the following vulnerability has been resolved: wifi: ath12k: fix possible out-of-bound read in a from HTT message and could be an unexpected value in case errors happen, so add validation before using to avoi message iteration and parsing. The same issue also applies to ppdu_info->ppdu_stats.common.num_users, so valid code review. Compile test only.
CVE-2023-52844	In the Linux kernel, the following vulnerability has been resolved: media: vidtv: psi: Add check for kstrdup Add ch return the error if it fails in order to avoid NULL pointer dereference.
CVE-2023-52858	In the Linux kernel, the following vulnerability has been resolved: clk: mediatek: clk-mt7629: Add check for mtk_ value of mtk_alloc_clk_data() in order to avoid NULL pointer dereference.
CVE-2023-52869	In the Linux kernel, the following vulnerability has been resolved: pstore/platform: Add check for kstrdup Add che the error if it fails in order to avoid NULL pointer dereference.
CVE-2023-6176	A null pointer dereference flaw was found in the Linux kernel API for the cryptographic algorithm scatterwalk fun constructs a malicious packet with specific socket configuration, which could allow a local user to crash the system
CVE-2023-6240	A Marvin vulnerability side-channel leakage was found in the RSA decryption operation in the Linux Kernel. This decrypt ciphertexts or forge signatures, limiting the services that use that private key.
CVE-2023-6356	A flaw was found in the Linux kernel's NVMe driver. This issue may allow an unauthenticated malicious actor to s using NVMe over TCP, leading the NVMe driver to a NULL pointer dereference in the NVMe driver and causing
CVE-2023-6535	A flaw was found in the Linux kernel's NVMe driver. This issue may allow an unauthenticated malicious actor to s using NVMe over TCP, leading the NVMe driver to a NULL pointer dereference in the NVMe driver, causing kern
CVE-2023-6536	A flaw was found in the Linux kernel's NVMe driver. This issue may allow an unauthenticated malicious actor to s using NVMe over TCP, leading the NVMe driver to a NULL pointer dereference in the NVMe driver, causing kern
CVE-2023-6546	A race condition was found in the GSM 0710 tty multiplexor in the Linux kernel. This issue occurs when two threa the same tty file descriptor with the gsm line discipline enabled, and can lead to a use-after-free problem on a struc could allow a local unprivileged user to escalate their privileges on the system.
CVE-2023-6606	An out-of-bounds read vulnerability was found in smbCalcSize in fs/smb/client/netmisc.c in the Linux Kernel. This the system or leak internal kernel information.
CVE-2023-6915	A Null pointer dereference problem was found in ida_free in lib/idr.c in the Linux Kernel. This issue may allow an service problem due to a missing check at a function return.
CVE-2023-6918	A flaw was found in the libssh implements abstract layer for message digest (MD) operations implemented by diffe values from these were not properly checked, which could cause low-memory situations failures, NULL dereferenc memory as an input for the KDF. In this case, non-matching keys will result in decryption/integrity failures, termin
CVE-2023-7192	A memory leak problem was found in ctnetlink_create_conntrack in net/netfilter/nf_conntrack_netlink.c in the Lin attacker with CAP_NET_ADMIN privileges to cause a denial of service (DoS) attack due to a refcount overflow.
CVE-2024-0193	A use-after-free flaw was found in the netfilter subsystem of the Linux kernel. If the catchall element is garbage-co element can be deactivated twice. This can cause a use-after-free issue on an NFT_CHAIN object or NFT_OBJEC with CAP_NET_ADMIN capability to escalate their privileges on the system.
CVE-2024-0775	A use-after-free flaw was found in the __ext4_remount in fs/ext4/super.c in ext4 in the Linux kernel. This flaw allo problem while freeing the old quota file names before a potential failure, leading to a use-after-free.
CVE-2024-21803	Use After Free vulnerability in Linux Linux kernel kernel on Linux, x86, ARM (bluetooth modules) allows Local E associated with program files https://gitee.Com/anolis/cloud-kernel/blob/devel-5.10/net/bluetooth/af_bluetooth.C. T rc2 before v6.8-rc1.
CVE-2024-22189	quic-go is an implementation of the QUIC protocol in Go. Prior to version 0.42.0, an attacker can cause its peer to number of `NEW_CONNECTION_ID` frames that retire old connection IDs. The receiver is supposed to respond `RETIRE_CONNECTION_ID` frame. The attacker can prevent the receiver from sending out (the vast majority of frames by collapsing the peers congestion window (by selectively acknowledging received packets) and by manipu 0.42.0 contains a patch for the issue. No known workarounds are available.
CVE-2024-22257	In Spring Security, versions 5.7.x prior to 5.7.12, 5.8.x prior to 5.8.11, versions 6.0.x prior to 6.0.9, versions 6.1.x p application is possible vulnerable to broken access control when it directly uses the AuthenticatedVoter#vote passi
CVE-2024-22386	A race condition was found in the Linux kernel's drm/exynos device driver in¬†exynos_drm_crtc_atomic_disable(dereference issue, possibly leading to a kernel panic or denial of service issue.

CVE-2024-23342	The `ecdsa` PyPI package is a pure Python implementation of ECC (Elliptic Curve Cryptography) with support for Algorithm), EdDSA (Edwards-curve Digital Signature Algorithm) and ECDH (Elliptic Curve Diffie-Hellman). Ver Minerva attack. As of time of publication, no known patched version exists.
CVE-2024-24864	A race condition was found in the Linux kernel's media/dvb-core in dvbdmx_write()¬†function. This can result in a leading to a kernel panic or denial of service issue.
CVE-2024-26583	In the Linux kernel, the following vulnerability has been resolved: tls: fix race between async notify and socket clo recvmsg/sendmsg) may exit as soon as the async crypto handler calls complete() so any code past that point risks to locking and extra flags altogether. Have the main thread hold an extra reference, this way we can depend solely on Don't futz with reiniting the completion, either, we are now tightly controlling when completion fires.
CVE-2024-26583	In the Linux kernel, the following vulnerability has been resolved: tls: fix race between async notify and socket clo recvmsg/sendmsg) may exit as soon as the async crypto handler calls complete() so any code past that point risks to locking and extra flags altogether. Have the main thread hold an extra reference, this way we can depend solely on Don't futz with reiniting the completion, either, we are now tightly controlling when completion fires.
CVE-2024-26584	In the Linux kernel, the following vulnerability has been resolved: net: tls: handle backlogging of crypto requests S CRYPTO_TFM_REQ_MAY_BACKLOG flag on our requests to the crypto API, crypto_aead_{encrypt,decrypt} c instead of -EINPROGRESS in valid situations. For example, when the cryptd queue for AESNI is full (easy to trig cryptd.cryptd_max_cpu_qlen), requests will be enqueued to the backlog but still processed. In that case, the async c first with err == -EINPROGRESS, which it seems we can just ignore, then with err == 0. Compared to Sabrina's or tls_*crypt_async_wait() helpers and converts the EBUSY to EINPROGRESS to avoid having to modify all the erro
CVE-2024-26585	In the Linux kernel, the following vulnerability has been resolved: tls: fix race between tx work scheduling and soc submitting thread (recvmsg/sendmsg) may exit as soon as the async crypto handler calls complete(). Reorder sched This seems more logical in the first place, as it's the inverse order of what the submitting thread will do.
CVE-2024-26800	In the Linux kernel, the following vulnerability has been resolved: tls: fix use-after-free on failed backlog decryptio to the backlog and crypto_aead_decrypt returns -EBUSY, tls_do_decryption will wait until all async decryptions ha tls_do_decryption will return -EBADMSG and tls_decrypt_sg jumps to the error path, releasing all the pages. But t callback, and have already been released by tls_decrypt_done. The only true async case is when crypto_aead_decry we already waited so we can tell tls_sw_recvmsg that the data is available for immediate copy, but we need to notif flag) that the memory has already been released.
CVE-2024-26811	In the Linux kernel, the following vulnerability has been resolved: ksmbd: validate payload size in ipc response If i ksmbd.mountd can return invalid ipc response to ksmbd kernel server. ksmbd should validate payload size of ipc re overrun or slab-out-of-bounds. This patch validate 3 ipc response that has payload.
CVE-2024-26841	In the Linux kernel, the following vulnerability has been resolved: LoongArch: Update cpu_sibling_map when disa cpu_sibling_map when disabling nonboot CPUs by defining & calling clear_cpu_sibling_map(), otherwise we get s label: negative count! WARNING: CPU: 6 PID: 45 at kernel/jump_label.c:263 __static_key_slow_dec_cpuslocked cpuhp/6 Not tainted 6.8.0-rc5+ #1340 pc 90000000004c302c ra 90000000004c302c tp 90000001005bc000 sp 9000 900000000224c278 a2 90000001005bfb58 a3 900000000224c280 a4 900000000224c278 a5 90000001005bfb50 a t0 ce87a4763eb5234a t1 ce87a4763eb5234a t2 0000000000000000 t3 0000000000000000 t4 0000000000000006 t t7 0000000000001964 t8 000000000009ebf6 u0 9000000001f2a068 s9 0000000000000000 s0 900000000246a2d8 90000000021518c0 s4 0000000000000040 s5 9000000002151058 s6 9000000009828e40 s7 00000000000000b4 s __static_key_slow_dec_cpuslocked+0xec/0x100 ERA: 90000000004c302c __static_key_slow_dec_cpuslocked+0x -IE -DA +PG DACF=CC DACM=CC -WE) PRMD: 00000004 (PPLV0 +PIE -PWE) EUEN: 00000000 (-FPE -SX (LIE=2-4,10-12 VS=7) ESTAT: 000c0000 [BRK] (IS= ECode=12 EsubCode=0) PRID: 0014d000 (Loongson-64b 45 Comm: cpuhp/6 Not tainted 6.8.0-rc5+ #1340 Stack : 0000000000000000 900000000203f258 900000000179af 0000000000000000 90000001005bf988 9000000001fe0be0 900000000224c280 900000000224c278 90000001005 0000000000000001 ce87a4763eb5234a 0000000007f38000 90000001003f8cc0 0000000000000000 00000000000 4c206e6f73676e6f 6f4c203a656d616e 000000000009ec99 0000000007f38000 0000000000000000 900000000214 0000000000000000 0000000000000107 0000000000000009 ffffffffffafdabe 00000000000000b4 00000000000000 00005555939a0c7c 00000000000000b0 0000000000000004 0000000000000000 0000000000071c1c ... Call Trace +0x48/0x1a0 [<900000000179afc8>] dump_stack_lvl+0x78/0xa0 [<9000000000263ed0>] __warn+0x90/0x1a0 [< +0x1b8/0x280 [<900000000179c564>] do_bp+0x264/0x420 [<90000000004c302c>] __static_key_slow_dec_cpus sched_cpu_deactivate+0x2fc/0x300 [<9000000000266498>] cpuhp_invoke_callback+0x178/0x8a0 [<900000000 [<90000000002a117c>] smpboot_thread_fn+0x1dc/0x2e0 [<900000000029a720>] kthread+0x140/0x160 [<90000 +0xc/0xa4

CVE-2024-26893	In the Linux kernel, the following vulnerability has been resolved: firmware: arm_scmi: Fix double free in SMC tra code tears down a channel, it calls the chan_free callback function, defined by each transport. Since multiple protoc member, chan_free() might want to clean up the same member multiple times within the given SCMI transport imp transport. This will lead to a NULL pointer dereference at the second time: | scmi_protocol scmi_dev.1: Enabled po | arm-scmi firmware:scmi: SCMI Notifications - Core Enabled. | arm-scmi firmware:scmi: unable to communicate NULL pointer dereference at virtual address 0000000000000000 | Mem abort info: | ESR = 0x0000000096000004 32 bits | SET = 0, FnV = 0 | EA = 0, S1PTW = 0 | FSC = 0x04: level 0 translation fault | Data abort info: | ISV = 0, | CM = 0, WnR = 0, TnD = 0, TagAccess = 0 | GCS = 0, Overlay = 0, DirtyBit = 0, Xs = 0 | user pgtable: 4k pages, [0000000000000000] pgd=0000000000000000, p4d=0000000000000000 | Internal error: Oops: 000000009600000 | CPU: 4 PID: 1 Comm: swapper/0 Not tainted 6.7.0-rc2-00124-g455ef3d016c9-dirty #793 | Hardware name: FVP daif +PAN -UAO -TCO +DIT -SSBS BTYPE=--) | pc : smc_chan_free+0x3c/0x6c | lr : smc_chan_free+0x3c/0x6c | idr_for_each+0x68/0xf8 | scmi_cleanup_channels.isra.0+0x2c/0x58 | scmi_probe+0x434/0x734 | platform_probe | __driver_probe_device+0x78/0x12c | driver_probe_device+0x3c/0x118 | __driver_attach+0x74/0x128 | bus_for_ +0x24/0x30 | bus_add_driver+0xe4/0x1e8 | driver_register+0x60/0x128 | __platform_driver_register+0x28/0x34 | +0x78/0x33c | kernel_init_freeable+0x2b8/0x51c | kernel_init+0x24/0x130 | ret_from_fork+0x10/0x20 | Code: f00 (b9400280) | ---[end trace 0000000000000000]--- Simply check for the struct pointer being NULL before trying t This was found when a transport doesn't really work (for instance no SMC service), the probe routines then tries to
CVE-2024-26896	In the Linux kernel, the following vulnerability has been resolved: wifi: wfx: fix memory leak when starting AP Kr object 0xd73d1180 (size 184): comm "wpa_supplicant", pid 1559, jiffies 13006305 (age 964.245s) hex dump (first 00 00 00 00 00 ..... 00 00 00 00 00 00 00 00 1e 00 01 00 00 00 00 00 ..... backtrace: [<5ca11420>] km __alloc_skb+0x144/0x170 [<fb8a5e38>] __netdev_alloc_skb+0x50/0x180 [<0f9fa1d5>] __ieee80211_beacon_get ieee80211_beacon_get_tim+0x54/0x18c [mac80211] [<41e25cc3>] wfx_start_ap+0xc8/0x234 [wfx] [<93a70356> [mac80211] [<a4a661cd>] nl80211_start_ap+0x76c/0x9e0 [cfg80211] [<47bd8b68>] genl_rcv_msg+0x198/0x378 +0xd0/0x130 [<6b7c977a>] genl_rcv+0x34/0x44 [<66b2d04d>] netlink_unicast+0x1b4/0x258 [<f965b9b6>] netli ____sys_sendmsg+0x1e0/0x274 [<d2b5212d>] ___sys_sendmsg+0x80/0xb4 [<69954f45>] __sys_sendmsg+0x64 (size 1024): comm "wpa_supplicant", pid 1559, jiffies 13006305 (age 964.246s) hex dump (first 32 bytes): 00 00 0 00 00 00 ..... 10 00 07 40 00 00 00 00 00 00 00 00 00 00 00 00 ...@..... backtrace: [<9a993714>] __kma [<f83ea192>] kmalloc_reserve.constprop.0+0x30/0x74 [<a2c61343>] __alloc_skb+0xa0/0x170 [<fb8a5e38>] __n __ieee80211_beacon_get+0x290/0x4d4 [mac80211] [<7accd02d>] ieee80211_beacon_get_tim+0x54/0x18c [mac8 +0xc8/0x234 [wfx] [<93a70356>] ieee80211_start_ap+0x404/0x6b4 [mac80211] [<a4a661cd>] nl80211_start_ap genl_rcv_msg+0x198/0x378 [<453ef796>] netlink_rcv_skb+0xd0/0x130 [<6b7c977a>] genl_rcv+0x34/0x44 [<66 [<f965b9b6>] netlink_sendmsg+0x1e8/0x428 [<aadb8231>] ____sys_sendmsg+0x1e0/0x274 [<d2b5212d>] ___s kernel is build optimized, it seems the stack is not accurate. It appears the issue is related to wfx_set_mfp_ap(). The allocated by ieee80211_beacon_get() is never released. Fixing this leak makes kmemleak happy.
CVE-2024-26907	In the Linux kernel, the following vulnerability has been resolved: RDMA/mlx5: Fix fortify source warning while [cut here]------------ memcpy: detected field-spanning write (size 56) of single field "eseg->inline_hdr.start" at /va build/drivers/infiniband/hw/mlx5/wr.c:131 (size 2) WARNING: CPU: 0 PID: 293779 at /var/lib/dkms/mlnx-ofed-k hw/mlx5/wr.c:131 mlx5_ib_post_send+0x191b/0x1a60 [mlx5_ib] Modules linked in: 8021q garp mrp stp llc rdma_ ib_ipoib(OE) ib_cm(OE) ib_umad(OE) mlx5_ib(OE) ib_uverbs(OE) ib_core(OE) mlx5_core(OE) pci_hyperv_intf mlxfw(OE) psample nft_fib_inet nft_fib_ipv4 nft_fib_ipv6 nft_fib nft_reject_inet nf_reject_ipv4 nf_reject_ipv6 nf nf_conntrack nf_defrag_ipv6 nf_defrag_ipv4 ip_set nf_tables libcrc32c nfnetlink mst_pciconf(OE) knem(OE) vfio vfio iommufd irqbypass cuse nfsv3 nfs fscache netfs xfrm_user xfrm_algo ipmi_devintf ipmi_msghandler binfmt_ polyval_clmulni polyval_generic ghash_clmulni_intel sha512_ssse3 snd_pcsp aesni_intel crypto_simd cryptd snd_ input_leds serio_raw evbug nfsd auth_rpcgss nfs_acl lockd grace sch_fq_codel sunrpc drm efi_pstore ip_tables x_t net_failover failover floppy [last unloaded: mlx_compat(OE)] CPU: 0 PID: 293779 Comm: ssh Tainted: G OE 6.2. Hardware name: Red Hat KVM, BIOS 0.5.1 01/01/2011 RIP: 0010:mlx5_ib_post_send+0x191b/0x1a60 [mlx5_ib] a0 b9 02 00 00 00 48 c7 c2 10 5b fd c0 48 c7 c7 80 5b fd c0 c6 05 57 0c 03 00 01 e8 95 4d 93 da <0f> 0b 44 8b 4d ff ff 41 0f b7 RSP: 0018:ffffb5b48478b570 EFLAGS: 00010046 RAX: 0000000000000000 RBX: 0000000000000 0000000000000000 RSI: 0000000000000000 RDI: 0000000000000000 RBP: ffffb5b48478b628 R08: 0000000000 R10: 0000000000000000 R11: 0000000000000000 R12: ffffb5b48478b5e8 R13: ffff963a3c609b5e R14: ffff9639c FS: 00007fc03b444c80(0000) GS:ffff963a3dc00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 000 0000556f46bdf000 CR3: 0000000006ac6003 CR4: 0000000003706f0 DR0: 0000000000000000 DR1: 00000000 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 Call Trace: <TASK> ? show_regs+0x72/0x [mlx5_ib] ? __warn+0x8d/0x160 ? mlx5_ib_post_send+0x191b/0x1a60 [mlx5_ib] ? report_bug+0x1bb/0x1d0 ? ha +0x19/0x80 ? asm_exc_invalid_op+0x1b/0x20 ? mlx5_ib_post_send+0x191b/0x1a60 [mlx5_ib] mlx5_ib_post_sen +0x2ec/0x770 [ib_ipoib] ipoib_start_xmit+0x5a0/0x770 [ib_ipoib] dev_hard_start_xmit+0x8e/0x1e0 ? validate_xm +0x116/0x3a0 __dev_xmit_skb+0x1fd/0x580 __dev_queue_xmit+0x284/0x6b0 ? _raw_spin_unlock_irq+0xe/0x5 push_pseudo_header+0x17/0x40 [ib_ipoib] neigh_connected_output+0xcd/0x110 ip_finish_output2+0x179/0x480 __ip_finish_output+0xc3/0x190 ip_finish_output+0x2e/0xf0 ip_output+0x78/0x110 ? __pfx_ip_finish_output+0x1 __ip_queue_xmit+0x18a/0x460 ip_queue_xmit+0x15/0x30 __tcp_transmit_skb+0x914/0x9c0 tcp_write_xmit+0x3 tcp_sendmsg_locked+0x2e1/0xac0 tcp_sendmsg+0x2d/0x50 inet_sendmsg+0x43/0x90 sock_sendmsg+0x68/0x80 +0x326/0x3c0 ksys_write+0xbd/0xf0 ? do_syscall_64+0x69/0x90 __x64_sys_write+0x19/0x30 do_syscall_ ---trun
CVE-2024-26913	In the Linux kernel, the following vulnerability has been resolved: drm/amd/display: Fix dcn35 8k30 Underflow/C missing for pipe split policy determination and cause Underflow/Corruption issue. [how] Add the odm calculation.
CVE-2024-26914	In the Linux kernel, the following vulnerability has been resolved: drm/amd/display: fix incorrect mpc_combine ar stream, while MAX_PLANES is per asic. The mpc_combine is an array that records all the planes per asic. Therefo array size. Using MAX_SURFACES causes array overflow when there are more than 3 planes. [how] Use the MA

CVE-2024-26952	In the Linux kernel, the following vulnerability has been resolved: ksmbd: fix potencial out-of-bounds when buffer bounds when buffer offset fields of a few requests is invalid. This patch set the minimum value of buffer offset fiel
CVE-2024-27034	In the Linux kernel, the following vulnerability has been resolved: f2fs: compress: fix to cover normal cluster write compressed cluster w/ normal cluster, we should not unlock cp_rwsem during f2fs_write_raw_pages(), otherwise d persisted before CP & SPOR, due to cluster metadata wasn't updated atomically.
CVE-2024-27035	In the Linux kernel, the following vulnerability has been resolved: f2fs: compress: fix to guarantee persisting comp compressed cluster is not persisted with metadata during checkpoint, after SPOR, the data may be corrupted, let's g checkpoint.
CVE-2024-27389	In the Linux kernel, the following vulnerability has been resolved: pstore: inode: Only d_invalidate() is needed Unl records in pstorefs would trigger the dput() double-drop warning: WARNING: CPU: 0 PID: 2569 at fs/dcache.c:76 of d_drop()/dput() (as mentioned in Documentation/filesystems/vfs.rst) isn't the right approach here, and leads to th Use d_invalidate() and update the code to not bother checking for error codes that can never happen. ---
CVE-2024-27393	In the Linux kernel, the following vulnerability has been resolved: xen-netfront: Add missing skb_mark_for_recycl introduced later than fixes tag in commit 6a5bcd84e886 ("page_pool: Allow drivers to hint on SKB recycling"). It to page_pool_release_page() between v5.9 to v5.14, after which is should have used skb_mark_for_recycle(). Since were removed (in commit 535b9c61bdef ("net: page_pool: hide page_pool_release_page()") and remaining callers branch 'net-page_pool-remove-page_pool_release_page'")). This leak became visible in v6.8 via commit dba1b8a7 memory leaks").
CVE-2024-28849	follow-redirects is an open source, drop-in replacement for Node's `http` and `https` modules that automatically fol follow-redirects only clears authorization header during cross-domain redirect, but keep the proxy-authentication h vulnerability may lead to credentials leak, but has been addressed in version 1.15.6. Users are advised to upgrade. T vulnerability.
CVE-2024-35785	In the Linux kernel, the following vulnerability has been resolved: tee: optee: Fix kernel panic caused by incorrect to register devices on the TEE bus has a bug leading to kernel panic as follows: [15.398930] Unable to handle kern ffff07ed00626d7c [15.406913] Mem abort info: [15.409722] ESR = 0x0000000096000005 [15.413490] EC = 0x [15.418814] SET = 0, FnV = 0 [15.421878] EA = 0, S1PTW = 0 [15.425031] FSC = 0x05: level 1 translation fau ISV = 0, ISS = 0x00000005, ISS2 = 0x00000000 [15.438310] CM = 0, WnR = 0, TnD = 0, TagAccess = 0 [15.44 = 0, Xs = 0 [15.448697] swapper pgtable: 4k pages, 48-bit VAs, pgdp=00000000d9e3e000 [15.455413] [ffff07ed p4d=1800000bffdf9003, pud=0000000000000000 [15.464146] Internal error: Oops: 0000000096000005 [#1] PRE optee: Fix supplicant based device enumeration") lead to the introduction of this bug. So fix it appropriately.
CVE-2024-35796	In the Linux kernel, the following vulnerability has been resolved: net: ll_temac: platform_get_resource replaced b platform_get_resource was replaced with devm_platform_ioremap_resource_byname and is called using 0 as name platform_get_resource_byname in the call stack, where it causes a null pointer in strcmp. if (type == resource_type have been replaced with devm_platform_ioremap_resource.
CVE-2024-35811	In the Linux kernel, the following vulnerability has been resolved: wifi: brcmfmac: Fix use-after-free bug in brcmf patch of CVE-2023-47233 : https://nvd.nist.gov/vuln/detail/CVE-2023-47233 In brcm80211 driver,it starts with th timeout worker: ->brcmf_usb_probe ->brcmf_usb_probe_cb ->brcmf_attach ->brcmf_bus_started ->brcmf_cfg802 ->INIT_WORK(&cfg->escan_timeout_work, brcmf_cfg80211_escan_timeout_worker); If we disconnect the USB to make cleanup. The invoking chain is : brcmf_usb_disconnect ->brcmf_usb_disconnect_cb ->brcmf_detach ->br the timeout woker may still be running. This will cause a use-after-free bug on cfg in brcmf_cfg80211_escan_time canceling the worker in brcmf_cfg80211_detach. [arend.vanspriel@broadcom.com: keep timer delete as is and can
CVE-2024-35818	In the Linux kernel, the following vulnerability has been resolved: LoongArch: Define the __io_aw() hook as mmi ("drivers: Remove explicit invocations of mmiowb()") remove all mmiowb() in drivers, but it says: "NOTE: mmiow in conjunction with spin_unlock(). However, pairing each mmiowb() removal in this patch with the corresponding so there is a small chance that this change may regress any drivers incorrectly relying on mmiowb() to order MMIO synchronisation." The mmio in radeon_ring_commit() is protected by a mutex rather than a spinlock, but in the mu We can add mmiowb() calls in the radeon driver but the maintainer says he doesn't like such a workaround, and rad protected mmio. So we should extend the mmiowb tracking system from spinlock to mutex, and maybe other locki prone, so we solve it in the architectural code, by simply defining the __io_aw() hook as mmiowb(). And we no lon so use the generic definition. Without this, we get such an error when run 'glxgears' on weak ordering architectures ring 0 stalled for more than 10324msec radeon 0000:04:00.0: ring 3 stalled for more than 10240msec radeon 0000: 0x000000000001f412 last fence id 0x000000000001f414 on ring 3) radeon 0000:04:00.0: GPU lockup (current fen 0x000000000000f941 on ring 0) radeon 0000:04:00.0: scheduling IB failed (-35). [drm:radeon_gem_va_ioctl [rade (-35) radeon 0000:04:00.0: scheduling IB failed (-35). [drm:radeon_gem_va_ioctl [radeon]] *ERROR* Couldn't up scheduling IB failed (-35). [drm:radeon_gem_va_ioctl [radeon]] *ERROR* Couldn't update BO_VA (-35) radeon [drm:radeon_gem_va_ioctl [radeon]] *ERROR* Couldn't update BO_VA (-35) radeon 0000:04:00.0: scheduling IB [radeon]] *ERROR* Couldn't update BO_VA (-35) radeon 0000:04:00.0: scheduling IB failed (-35). [drm:radeon_ update BO_VA (-35) radeon 0000:04:00.0: scheduling IB failed (-35). [drm:radeon_gem_va_ioctl [radeon]] *ERR
CVE-2024-35829	In the Linux kernel, the following vulnerability has been resolved: drm/lima: fix a memleak in lima_heap_alloc Wh need to be deallocated, or there will be memleaks.

CVE-2024-35833	In the Linux kernel, the following vulnerability has been resolved: dmaengine: fsl-qdma: Fix a memory leak relate dma_alloc_coherent() is undone neither in the remove function, nor in the error handling path of fsl_qdma_probe() issues.
CVE-2024-35835	In the Linux kernel, the following vulnerability has been resolved: net/mlx5e: fix a double-free in arfs_create_grou arfs_create_groups will free ft->g and return an error. However, arfs_create_table, the only caller of arfs_create_gr mlx5e_destroy_flow_table, in which the ft->g will be freed again.
CVE-2024-35844	In the Linux kernel, the following vulnerability has been resolved: f2fs: compress: fix reserve_cblocks counting err one direct_node, performing the following operations will cause the file to be unrepairable: unisoc # ./f2fs_io comp dev/block/dm-48 112G 112G 1.2M 100% /data unisoc # ./f2fs_io release_cblocks test.apk 924 unisoc # df -h | grep 100% /data unisoc # dd if=/dev/random of=file4 bs=1M count=3 3145728 bytes (3.0 M) copied, 0.025 s, 120 M/s u 112G 112G 1.8M 100% /data unisoc # ./f2fs_io reserve_cblocks test.apk F2FS_IOC_RESERVE_COMPRESS_BL reboot unisoc # df -h | grep dm-48 /dev/block/dm-48 112G 112G 11M 100% /data unisoc # ./f2fs_io reserve_cbloc one direct_node. After returning to -ENOSPC, reserved_blocks += ret will not be executed. As a result, the reserve the real number of reserved blocks. Therefore, fsck cannot be set to repair the file. After this patch, the fsck flag wi grep dm-48 /dev/block/dm-48 112G 112G 1.8M 100% /data unisoc # ./f2fs_io reserve_cblocks test.apk F2FS_IOC No space left on device adb reboot then fsck will be executed unisoc # df -h | grep dm-48 /dev/block/dm-48 112G reserve_cblocks test.apk 924
CVE-2024-35845	In the Linux kernel, the following vulnerability has been resolved: wifi: iwlwifi: dbg-tlv: ensure NUL termination string, so we must ensure the string is terminated correctly before using it.
CVE-2024-35879	In the Linux kernel, the following vulnerability has been resolved: of: dynamic: Synchronize of_changeset_destroy sequence: 1) of_platform_depopulate() 2) of_overlay_remove() During the step 1, devices are destroyed and devlin are destroyed but __of_changeset_entry_destroy() can raise warnings related to missing of_node_put(): ERROR: n 2 ... Indeed, during the devlink removals performed at step 1, the removal itself releasing the device (and the attach workqueue and so, it is done asynchronously with respect to function calls. When the warning is present, of_node_ the workqueue job. In order to be sure that any ongoing devlink removals are done before the of_node destruction, the devlink removals.
CVE-2024-35902	In the Linux kernel, the following vulnerability has been resolved: net/rds: fix possible cp null dereference cp migh null dereference [Simon Horman adds:] Analysis: * cp is a parameter of __rds_rdma_map and is not reassigned. * cp argument to __rds_rdma_map() - rds_get_mr() - rds_get_mr_for_dest * Prior to the code above, the following a indicative, but could itself be unnecessary) trans_private = rs->rs_transport->get_mr(sg, nents, rs, &mr->r_key, cp args->vec.bytes, need_odp ? ODP_ZEROBASED : ODP_NOT_NEEDED); * The code modified by this patch is g trans_private is assigned as per the previous point in this analysis. The only implementation of get_mr that I could an ERR_PTR if the conn (4th) argument is NULL. * ret is set to PTR_ERR(trans_private). rds_ib_get_mr can retu argument is NULL. Thus ret may be -ENODEV in which case the code in question will execute. Conclusion: * cp adds a check; this patch does seem to address a possible bug
CVE-2024-35956	In the Linux kernel, the following vulnerability has been resolved: btrfs: qgroup: fix qgroup prealloc rsv leak in sub create snapshot and delete subvolume all use btrfs_subvolume_reserve_metadata() to reserve metadata for the chan tree, which cannot be mediated in the normal way via start_transaction. When quota groups (squota or qgroups) are of type PREALLOC. Once the operation is associated to a transaction, we convert PREALLOC to PERTRANS, wi transaction. However, the error paths of these three operations were not implementing this lifecycle correctly. They to PERTRANS in a generic cleanup step regardless of errors or whether the operation was fully associated to a tran occasionally converting this rsv to PERTRANS without calling record_root_in_trans successfully, which meant tha transaction by some other thread, the end of the transaction would not free that root's PERTRANS, leaking it. Ultim CONFIG_BTRFS_DEBUG builds at unmount for the leaked reservation. The fix is to ensure that every qgroup PR properties: 1. any failure before record_root_in_trans is called successfully results in freeing the PREALLOC reser we convert to PERTRANS, and now the transaction owns freeing the reservation. This patch enforces those proper generic/269 with squotas enabled at mkfs time would fail in ~5-10 runs on my system. With this patch, it ran succe
CVE-2024-35971	In the Linux kernel, the following vulnerability has been resolved: net: ks8851: Handle softirqs at the end of IRQ th thread may call ks8851_rx_pkts() in case there are any packets in the MAC FIFO, which calls netif_rx(). This netif local_bh_disable() and local_bh_enable(). The local_bh_enable() may call do_softirq() to run softirqs in case any a net_rx_action, which ultimately reaches the driver .start_xmit callback. If that happens, the system hangs. The entir from netdev_start_xmit netdev_start_xmit from dev_hard_start_xmit dev_hard_start_xmit from sch_direct_xmit sc __dev_queue_xmit from __neigh_update __neigh_update from neigh_update neigh_update from arp_process.const __netif_receive_skb_one_core __netif_receive_skb_one_core from process_backlog process_backlog from __napi from net_rx_action net_rx_action from __do_softirq __do_softirq from call_with_stack call_with_stack from do_s __local_bh_enable_ip from netif_rx netif_rx from ks8851_irq ks8851_irq from irq_thread_fn irq_thread_fn from ir from ret_from_fork The hang happens because ks8851_irq() first locks a spinlock in ks8851_par.c ks8851_lock_pa with that spinlock locked, calls netif_rx(). Once the execution reaches ks8851_start_xmit_par(), it calls ks8851_loc already locked spinlock again, and the hang happens. Move the do_softirq() call outside of the spinlock protected s around the entire spinlock protected section of ks8851_irq() handler. Place local_bh_enable() outside of the spinloc do_softirq() without the ks8851_par.c ks8851_lock_par() spinlock being held, and safely call ks8851_start_xmit_p locked spinlock. Since ks8851_irq() is protected by local_bh_disable()/local_bh_enable() now, replace netif_rx() w local_bh_disable()/local_bh_enable() calls.

CVE-2024-35988	In the Linux kernel, the following vulnerability has been resolved: riscv: Fix TASK_SIZE on 64-bit NOMMU On anywhere in physical RAM. The current definition of TASK_SIZE is wrong if any RAM exists above 4G, causing routines.
CVE-2024-35990	In the Linux kernel, the following vulnerability has been resolved: dma: xilinx_dpdma: Fix locking There are sever chan->vchan.lock was not held. Add appropriate locking. This fixes lockdep warnings like [31.077578] ---------- WARNING: CPU: 2 PID: 40 at drivers/dma/xilinx/xilinx_dpdma.c:834 xilinx_dpdma_chan_queue_transfer+0x27 linked in: [31.078019] CPU: 2 PID: 40 Comm: kworker/u12:1 Not tainted 6.6.20+ #98 [31.078102] Hardware nar Workqueue: events_unbound deferred_probe_work_func [31.078272] pstate: 600000c5 (nZCv daIF -PAN -UAO [31.078377] pc : xilinx_dpdma_chan_queue_transfer+0x274/0x5e0 [31.078473] lr : xilinx_dpdma_chan_queue_t sp : ffffffc083bb2e10 [31.078590] x29: ffffffc083bb2e10 x28: 0000000000000000 x27: ffffff880165a168 [31.078 ffffff880164eab8 x24: ffffff880164d480 [31.078920] x23: ffffff880165a148 x22: ffffff880164e988 x21: 0000000 ffffffc082aa3000 x19: ffffff880164e880 x18: 0000000000000000 [31.079295] x17: 0000000000000000 x16: 0000 [31.079453] x14: 0000000000000000 x13: ffffff8802263dc0 x12: 0000000000000001 [31.079613] x11: 0001ffc0 x9 : 0001ffc082aa3def [31.079824] x8 : 0001ffc082aa3dec x7 : 0000000000000000 x6 : 0000000000000516 [31. ffffff88003c9c40 x3 : ffffffffffffffff [31.080147] x2 : ffffffc7f8d43000 x1 : 00000000000000c0 x0 : 0000000000 xilinx_dpdma_chan_queue_transfer+0x274/0x5e0 [31.080518] xilinx_dpdma_issue_pending+0x11c/0x120 [31.0 +0x180/0x3ac [31.080712] zynqmp_dpsub_plane_atomic_update+0x11c/0x21c [31.080825] drm_atomic_helper_ drm_atomic_helper_commit_tail+0x5c/0xb0 [31.081139] commit_tail+0x234/0x294 [31.081246] drm_atomic_he drm_atomic_commit+0x100/0x140 [31.081477] drm_client_modeset_commit_atomic+0x318/0x384 [31.081634] +0x8c/0x24c [31.081725] drm_client_modeset_commit+0x34/0x5c [31.081812] __drm_fb_helper_restore_fbdev [31.081899] drm_fb_helper_set_par+0x50/0x70 [31.081971] fbcon_init+0x538/0xc48 [31.082047] visual_init+0 do_bind_con_driver.isra.0+0x2d0/0x634 [31.082320] do_take_over_console+0x24c/0x33c [31.082429] do_fbcon fbcon_fb_registered+0x2d0/0x34c [31.082663] register_framebuffer+0x27c/0x38c [31.082767] __drm_fb_helper [31.082939] drm_fb_helper_initial_config+0x50/0x74 [31.083012] drm_fbdev_dma_client_hotplug+0xb8/0x108 +0xa0/0xf4 [31.083195] drm_fbdev_dma_setup+0xb0/0x1cc [31.083293] zynqmp_dpsub_drm_init+0x45c/0x4e0 +0x444/0x5e0 [31.083616] platform_probe+0x8c/0x13c [31.083713] really_probe+0x258/0x59c [31.083793] __ [31.083878] driver_probe_device+0x70/0x1c0 [31.083961] __device_attach_driver+0x108/0x1e0 [31.084052] b [31.084125] __device_attach+0x100/0x298 [31.084207] device_initial_probe+0x14/0x20 [31.084292] bus_probe deferred_probe_work_func+0x11c/0x180 [31.084451] process_one_work+0x3ac/0x988 [31.084643] worker_thre +0x1bc/0x1c0 [31.084848] ret_from_fork+0x10/0x20 [31.084932] irq event stamp: 64549 [31.084970] hardirqs _raw_spin_unlock_irqrestore+0x80/0x90 [31.085157] ---truncated---
CVE-2024-36008	In the Linux kernel, the following vulnerability has been resolved: ipv4: check for NULL idev in ip_route_use_hin deref in fib_validate_source() in an old tree [1]. It appears the bug exists in latest trees. All calls to __in_dev_get_r result. [1] general protection fault, probably for non-canonical address 0xdffffc0000000000: 0000 [#1] SMP KASA [0x0000000000000000-0x0000000000000007] CPU: 2 PID: 3257 Comm: syz-executor.3 Not tainted 5.10.0-syzka (Q35 + ICH9, 2009), BIOS 1.16.3-debian-1.16.3-2~bpo12+1 04/01/2014 RIP: 0010:fib_validate_source+0xbf/0x1 f2 f2 f2 42 c7 44 20 23 f3 f3 f3 f3 48 89 44 24 78 42 c6 44 20 27 f3 e8 5d 88 48 fc 4c 89 e8 48 c1 e8 03 48 89 44 2 15 98 fc 48 89 5c 24 10 41 bf RSP: 0018:ffffc900015fee40 EFLAGS: 00010246 RAX: 0000000000000000 RBX: RDX: 0000000000000000 RSI: 0000000004001eac RDI: ffff8880160c64c0 RBP: ffffc900015ff060 R08: 0000000 R10: 0000000000000002 R11: ffff88800f4f90c0 R12: dfffffc0000000000 R13: 0000000000000000 R14: 0000000 FS: 00007f938acfe6c0(0000) GS:ffff888058c00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 000 00007f938acddd58 CR3: 000000001248e000 CR4: 0000000000352ef0 DR0: 0000000000000000 DR1: 0000000 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 Call Trace: ip_route_use_hint+0x410/0x9b +0x2c4/0x1a30 net/ipv4/ip_input.c:327 ip_list_rcv_finish net/ipv4/ip_input.c:612 [inline] ip_sublist_rcv+0x3ed/0x +0x422/0x470 net/ipv4/ip_input.c:673 __netif_receive_skb_list_ptype net/core/dev.c:5572 [inline] __netif_receive dev.c:5620 __netif_receive_skb_list net/core/dev.c:5672 [inline] netif_receive_skb_list_internal+0x9f9/0xdc0 net/ +0x55/0x3e0 net/core/dev.c:5816 xdp_recv_frames net/bpf/test_run.c:257 [inline] xdp_test_run_batch net/bpf/test_ +0x1818/0x1d00 net/bpf/test_run.c:363 bpf_prog_test_run_xdp+0x81f/0x1170 net/bpf/test_run.c:1376 bpf_prog_t syscall.c:3736 __sys_bpf+0x45c/0x710 kernel/bpf/syscall.c:5115 __do_sys_bpf kernel/bpf/syscall.c:5201 [inline] _ [inline] __x64_sys_bpf+0x7c/0x90 kernel/bpf/syscall.c:5199
CVE-2024-36898	In the Linux kernel, the following vulnerability has been resolved: gpiolib: cdev: fix uninitialised kfifo If a line is r debouncing in software, and the line is subsequently reconfigured to enable edge detection then the allocation of th This results in events being written to and read from an uninitialised kfifo. Read events are returned to userspace. I software debounce is already active.
CVE-2024-36899	In the Linux kernel, the following vulnerability has been resolved: gpiolib: cdev: Fix use after free in lineinfo_chan as follows: when the GPIO chip device file is being closed by invoking gpio_chrdev_release(), watched_lines is fre of lineinfo_changed_nb notifier chain failed due to waiting write rwsem. Additionally, one of the GPIO chip's lines notifier chain's read rwsem. Consequently, a race condition leads to the use-after-free of watched_lines. Here is the gpio_chrdev_release() --> bitmap_free(cdev->watched_lines) <-- freed --> blocking_notifier_chain_unregister() --> rwsem --> __down_write_common() --> rwsem_down_write_slowpath() --> schedule_preempt_disabled() --> sche gpio_free() --> gpiod_free() --> gpiod_free_commit() --> gpiod_line_state_notify() --> blocking_notifier_call_cha rwsem --> notifier_call_chain() --> lineinfo_changed_notify() --> test_bit(xxxx, cdev->watched_lines) <-- use afte issue is that a GPIO line event is being generated for userspace where it shouldn't. However, since the chrdev is bei read that event anyway. To fix the issue, call the bitmap_free() function after the unregistration of lineinfo_changed

CVE-2024-36942	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: qca: fix firmware check error path A firmware files before downloading them to the controller but introduced a memory leak in case the sanity checks ev before returning on errors.
CVE-2024-36964	In the Linux kernel, the following vulnerability has been resolved: fs/9p: only translate RWX permissions for plain bits is allowed through, which causes it to be able to set (among others) the suid bit. This was presumably not the in explicitly and conditionally on .u.
CVE-2024-38577	In the Linux kernel, the following vulnerability has been resolved: rcu-tasks: Fix show_rcu_tasks_trace_gp_kthread buffer overflow in show_rcu_tasks_trace_gp_kthread() if counters, passed to sprintf() are huge. Counter numbers, r buffer overflow is still possible. Use snprintf() with buffer size instead of sprintf(). Found by Linux Verification Ce
CVE-2024-38620	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: HCI: Remove HCI_AMP support Sin controllers no longer has any use so remove it along with the capability of creating AMP controllers. Since we no l Primary controllers, as only HCI_PRIMARY is left, this also remove hdev->dev_type altogether.
CVE-2024-38667	In the Linux kernel, the following vulnerability has been resolved: riscv: prevent pt_regs corruption for secondary i should be reserved for pt_regs. However this is not the case for the idle threads of the secondary boot harts. Their s may get corrupted. Similar issue has been fixed for the primary hart, see c7cdd96eca28 ("riscv: prevent stack corru However that fix was not propagated to the secondary harts. The problem has been noticed in some CPU hotplug te stored several registers on stack, corrupting top of pt_regs structure including status field. As a result, kernel attemp
CVE-2024-39496	In the Linux kernel, the following vulnerability has been resolved: btrfs: zoned: fix use-after-free due to race with c creation of a block group, we can race with a device replace operation and then trigger a use-after-free on the devic the replace operation). This happens because at btrfs_load_zone_info() we extract a device from the chunk map int while not under the protection of the device replace rwsem. So if there's a device replace operation happening when source of the replace operation, we will trigger a use-after-free if before we finish using the device the replace oper enlarging the critical section under the protection of the device replace rwsem so that all uses of the device are done
CVE-2024-39496	In the Linux kernel, the following vulnerability has been resolved: btrfs: zoned: fix use-after-free due to race with c creation of a block group, we can race with a device replace operation and then trigger a use-after-free on the devic the replace operation). This happens because at btrfs_load_zone_info() we extract a device from the chunk map int while not under the protection of the device replace rwsem. So if there's a device replace operation happening when source of the replace operation, we will trigger a use-after-free if before we finish using the device the replace oper enlarging the critical section under the protection of the device replace rwsem so that all uses of the device are done
CVE-2024-39497	In the Linux kernel, the following vulnerability has been resolved: drm/shmem-helper: Fix BUG_ON() on mmap(P of check for copy-on-write (COW) mapping in drm_gem_shmem_mmap allows users to call mmap with PROT_W a kernel panic due to BUG_ON in vmf_insert_pfn_prot: BUG_ON((vma->vm_flags & VM_PFNMAP) && is_cow EINVAL early if COW mapping is detected. This bug affects all drm drivers using default shmem helpers. It can b *ptr = mmap(0, size, PROT_WRITE, MAP_PRIVATE, fd, mmap_offset); ptr[0] = 0;
CVE-2024-39507	In the Linux kernel, the following vulnerability has been resolved: net: hns3: fix kernel crash problem in concurren driver need to notify the roce driver to handle this event, but at this time, the roce driver may uninit, then cause ker status change, need to check whether the roce registered, and when uninit, need to wait link update finish.
CVE-2024-39508	In the Linux kernel, the following vulnerability has been resolved: io_uring/io-wq: Use set_bit() and test_bit() at w on worker->flags within io_uring/io-wq to address potential data races. The structure io_worker->flags may be acc to concurrency issues. When KCSAN is enabled, it reveals data races occurring in io_worker_handle_work and io_ BUG: KCSAN: data-race in io_worker_handle_work / io_wq_activate_free_worker write to 0xffff8885c4246404 o io_worker_handle_work (io_uring/io-wq.c:434 io_uring/io-wq.c:569) io_wq_worker (io_uring/io-wq.c:?) <snip> r by task 49024 on cpu 5: io_wq_activate_free_worker (io_uring/io-wq.c:? io_uring/io-wq.c:285) io_wq_enqueue (io (io_uring/io_uring.c:524) io_req_task_submit (io_uring/io_uring.c:1511) io_handle_tw_list (io_uring/io_uring.c:11 18daea77cca6 ("Merge tag 'for-linus' of git://git.kernel.org/pub/scm/virt/kvm/kvm"). These races involve writes an different tasks running on different CPUs. To mitigate this, refactor the code to use atomic operations such as set_b "and" and "or" operations. This ensures thread-safe manipulation of worker flags. Also, move `create_index` to avo

CVE-2024-39510	In the Linux kernel, the following vulnerability has been resolved: cachefiles: fix slab-use-after-free in cachefiles_ondemand_daemon_read() We got the following issue in a fuzz test of randomly issuing the restore com ========== BUG: KASAN: slab- cachefiles_ondemand_daemon_read+0xb41/0xb60 Read of size 8 at addr ffff888122e84088 by task ondemand-04- Comm: ondemand-04-dae Not tainted 6.8.0-dirty #564 Call Trace: kasan_report+0x93/0xc0 cachefiles_ondemand_ vfs_read+0x169/0xb50 ksys_read+0xf5/0x1e0 Allocated by task 116: kmem_cache_alloc+0x140/0x3a0 cachefiles fscache_cookie_state_machine+0x43c/0x1230 [...] Freed by task 792: kmem_cache_free+0xfe/0x390 cachefiles_p fscache_cookie_state_machine+0x5c8/0x1230 [...] ========== Following is the process that triggers the issue: mount | daemon_thread1 | daemon_thread2 ---------- cachefiles_withdraw_cookie cachefiles_ondemand_clean_object(object) cachefiles_ondemand_send_req REQ_A = wait_for_completion(&REQ_A->done) cachefiles_daemon_read cachefiles_ondemand_daemon_read REQ_A = ca msg->object_id = req->object->ondemand->ondemand_id ------ restore ------ cachefiles_ondemand_restore xas_fo xas_set_mark(&xas, CACHEFILES_REQ_NEW) cachefiles_daemon_read cachefiles_ondemand_daemon_read R copy_to_user(_buffer, msg, n) xa_erase(&cache->reqs, id) complete(&REQ_A->done) ------ close(fd) ------ cachef cachefiles_put_object cachefiles_put_object kmem_cache_free(cachefiles_object_jar, object) REQ_A->object->on When we see the request within xa_lock, req->object must not have been freed yet, so grab the reference count of c issue.
CVE-2024-40899	In the Linux kernel, the following vulnerability has been resolved: cachefiles: fix slab-use-after-free in cachefiles_c issue in a fuzz test of randomly issuing the restore command: ========== BUG: KASAN: slab-use-after-free in cachefiles_ondemand_daemon_read+0x609/0xab0 Write of size 4 at addr fff ondemand-04-dae/4962 CPU: 11 PID: 4962 Comm: ondemand-04-dae Not tainted 6.8.0-rc7-dirty #542 Call Trace: cachefiles_ondemand_daemon_read+0x609/0xab0 vfs_read+0x169/0xb50 ksys_read+0xf5/0x1e0 Allocated by tas +0x1df/0x4b0 cachefiles_ondemand_send_req+0x24d/0x690 cachefiles_create_tmpfile+0x249/0xb30 cachefiles_c cachefiles_look_up_object+0x29c/0xa60 cachefiles_lookup_cookie+0x37d/0xca0 fscache_cookie_state_machine+ by task 626: kfree+0xf1/0x2c0 cachefiles_ondemand_send_req+0x568/0x690 cachefiles_create_tmpfile+0x249/0x +0x6f/0x140 cachefiles_look_up_object+0x29c/0xa60 cachefiles_lookup_cookie+0x37d/0xca0 fscache_cookie_st ========== Following is the proc daemon_thread1 | daemon_thread2 ---------- cachefiles_ondemand_init_o REQ_A = kzalloc(sizeof(*req) + data_len) wait_for_completion(&REQ_A->done) cachefiles_daemon_read cache REQ_A = cachefiles_ondemand_select_req cachefiles_ondemand_get_fd copy_to_user(_buffer, msg, n) process_o restore ------ cachefiles_ondemand_restore xas_for_each(&xas, req, ULONG_MAX) xas_set_mark(&xas, CACHE cachefiles_daemon_read cachefiles_ondemand_daemon_read REQ_A = cachefiles_ondemand_select_req write(de size)); cachefiles_ondemand_copen xa_erase(&cache->reqs, id) complete(&REQ_A->done) kfree(REQ_A) cachef get_unused_fd_flags file = anon_inode_getfile fd_install(fd, file) load = (void *)REQ_A->msg.data; load->fd = fd; issuing a restore command when the daemon is still alive, which results in a request being processed multiple times problem, add an additional reference count to cachefiles_req, which is held while waiting and reading, and then rel Note that since there is only one reference count for waiting, we need to avoid the same request being completed m request if it is successfully removed from the xarray.
CVE-2024-40900	In the Linux kernel, the following vulnerability has been resolved: cachefiles: remove requests from xarray during CACHEFILES_DEAD set, we can still read the requests, so in the following concurrency the request may be used daemon_thread1 | daemon_thread2 ---------- cachefiles_ondemand_init_o REQ_A = kzalloc(sizeof(*req) + data_len) wait_for_completion(&REQ_A->done) cachefiles_daemon_read cache dev fd cachefiles_flush_reqs complete(&REQ_A->done) kfree(REQ_A) xa_lock(&cache->reqs); cachefiles_onder CACHEFILES_OP_READ // req use-after-free !!! xa_unlock(&cache->reqs); xa_destroy(&cache->reqs) Hence re flushing them to avoid accessing freed requests.
CVE-2024-40910	In the Linux kernel, the following vulnerability has been resolved: ax25: Fix refcount imbalance on inbound conne ax25_release(), we call netdev_put() to decrease the refcount on the associated ax.25 device. However, the executio never calls netdev_hold(). This imbalance leads to refcount errors, and ultimately to kernel crashes. A typical call t one of the following errors: refcount_t: decrement hit 0; leaking memory. refcount_t: underflow; use-after-free. An <TASK> ? show_regs+0x64/0x70 ? __warn+0x83/0x120 ? refcount_warn_saturate+0xb2/0x100 ? report_bug+0x1 handle_bug+0x3e/0x70 ? exc_invalid_op+0x1c/0x70 ? asm_exc_invalid_op+0x1f/0x30 ? refcount_warn_saturate+ +0xb2/0x100 ax25_release+0x2ad/0x360 __sock_release+0x35/0xa0 sock_close+0x19/0x20 [...] On reboot (or any gets stuck in an infinite loop: unregister_netdevice: waiting for ax0 to become free. Usage count = 0 This patch cor netdev_hold() and ax25_dev_hold() for new connections in ax25_accept(). This makes the logic leading to ax25_a both cases we increment the refcount, which is ultimately decremented in ax25_release().

CVE-2024-40915	In the Linux kernel, the following vulnerability has been resolved: riscv: rewrite __kernel_map_pages() to fix sleep is a debug function which clears the valid bit in page table entry for deallocated pages to detect illegal memory acc set/clear the valid bit using __set_memory(). __set_memory() acquires init_mm's semaphore, and this operation ma __kernel_map_pages() can be called in atomic context, and thus is illegal to sleep. An example warning that this ca invalid context at kernel/locking/rwsem.c:1578 in_atomic(): 1, irqs_disabled(): 0, non_block: 0, pid: 2, name: kthre 0 PID: 2 Comm: kthreadd Not tainted 6.9.0-g1d4c6d784ef6 #37 Hardware name: riscv-virtio,qemu (DT) Call Trac +0x1c/0x24 [<ffffffff8091ef6e>] show_stack+0x2c/0x38 [<ffffffff8092baf8>] dump_stack_lvl+0x5a/0x72 [<ffffff [<ffffffff8003b7ac>] __might_resched+0x104/0x10e [<ffffffff8003b7f4>] __might_sleep+0x3e/0x62 [<ffffffff809 [<ffffffff8000cf00>] __set_memory+0x82/0x2fa [<ffffffff8000d324>] __kernel_map_pages+0x5a/0xd4 [<ffffffff8 [<ffffffff8018ee82>] __vmalloc_node_range+0x196/0x6ba [<ffffffff80011904>] copy_process+0x72c/0x17ec [<ff [<ffffffff80012f62>] kernel_thread+0x82/0xa0 [<ffffffff8003552c>] kthreadd+0x14a/0x1be [<ffffffff809357de>] with apply_to_existing_page_range(). It is fine to not have any locking, because __kernel_map_pages() works with pages are not changed by anyone else in the meantime.
CVE-2024-40918	In the Linux kernel, the following vulnerability has been resolved: parisc: Try to fix random segmentation faults in PA8800 and PA8900 processors have had problems with random segmentation faults for many years. Systems with Systems with PA8800 and PA8900 processors have a large L2 cache which needs per page flushing for decent perf The combined cache in these systems is also more sensitive to non-equivalent aliases than the caches in earlier syst faults that I have looked at appear to be memory corruption in memory allocated using mmap and malloc. My first work. On reviewing the cache code, I realized that there were two issues which the existing code didn't handle corr issue is that the present bit in PTEs is racy. 1) PA-RISC caches have a mind of their own and they can speculatively long as there is a entry in the TLB for the page which allows move-in. TLBs are local to each CPU. Thus, the TLB flushing the page. This is particularly important on SMP systems. In some of the flush routines, the flush routine w would be purged. This was because the flush routine needed the TLB entry to do the flush. 2) My initial approach t try and use flush_cache_page_if_present for all flush operations. This actually made things worse and led to a coup on me that some lines weren't being flushed because the pte check code was racy. This resulted in random inequiva __flush_cache_page tmpalias flush sets up its own TLB entry and it doesn't need the existing TLB entry. As long a we can get the pfn and physical address of the page. We can also purge the TLB entry for the page before doing the special TLB entry that inhibits cache move-in. When switching page mappings, we need to ensure that lines are rer just flush the lines to memory as they may come back. This made it clear that we needed to implement all the requi This includes flushes for user and kernel pages. After modifying the code to use tmpalias flushes, it became clear th fully resolved. The frequency of faults was worse on systems with a 64 MB L2 (PA8900) and systems with more C to flush_cache_page_if_present to detect pages that couldn't be flushed triggered frequently on some systems. Helg flushed and found that the PTE was either cleared or for a swap page. Ignoring pages that were swapped out seeme problematic. I looked at routines related to pte_clear and noticed ptep_clear_flush. The default implementation just obvious that on parisc we need to flush the cache page as well. If we don't flush the cache page, stale lines will be l Once a PTE is cleared, there is no way to find the physical address associated with the PTE and flush the associated updated change with a parisc specific version of ptep_clear_flush. It fixed the random data corruption on Helge's rp this point, I realized that I could restore the code where we only flush in flush_cache_page_if_present if the page h need to flush the cache when the accessed bit is cleared in ---truncated---
CVE-2024-40927	In the Linux kernel, the following vulnerability has been resolved: xhci: Handle TD clearing for multiple streams c multiple TDs might be in flight when an endpoint is stopped. We need to issue a Set TR Dequeue Pointer for each, and the caches cleared. Change the logic so that any N>1 TDs found active for different streams are deferred until a xhci_invalidate_cancelled_tds() again from xhci_handle_cmd_set_deq() to queue another command until we are do error/"should never happen" paths to ensure we at least clear any affected TDs, even if we can't issue a command to loudly with an xhci_warn() if this ever happens. This problem case dates back to commit e9df17eb1408 ("USB: xh rings per endpoint.") early on in the XHCI driver's life, when stream support was first added. It was then identified in commit 674f8438c121 ("xhci: split handling halted endpoints into two steps"), which added a FIXME comment changing the behavior as far as I can tell, though the new logic made the problem more obvious). Then later, in cor give back some cached cancelled URBs."), it was acknowledged again. [Mathias: commit 94f339147fc3 ("xhci: Fi URBs.") was a targeted regression fix to the previously mentioned patch. Users reported issues with usb stuck after This rolled back the TD clearing of multiple streams to its original state.] Apparently the commit author was aware It was still mentioned as a FIXME, an xhci_dbg() was added to log the problem condition, and the remaining issue The choice of making the log type xhci_dbg() for what is, at this point, a completely unhandled and known broken it guarantees that no actual users would see the log in production, thereby making it nigh undebuggable (indeed, ev doesn't really hint at there being a problem at all). It took me *months* of random xHC crashes to finally find a rel debug session, which could all have been avoided had this unhandled, broken condition been actually reported with intentionally left in unfixed (never mind that it shouldn't have been left in at all). > Another fix to solve clearing the TDs is needed, but not as urgent. 3 years after that statement and 14 years after the original bug was introduced, I th next time let's not leave bugs unfixed (that are actually worse than the original bug), and let's actually get people to crashes and IOMMU faults with UAS devices when handling errors/faults. Easiest repro is to use `hdparm` to mark bad, then `cat /dev/sdX > /dev/null` in a loop. At least in the case of JMicron controllers, the read errors end up hav requests to different streams) and the one that didn't get cleared properly ends up faulting the xHC entirely when it been unmapped, referred to by the stale TDs. This normally happens quickly (after two or three loops). After this fi and experienced no xHC failures, with all read errors recovered properly. Repro'd and tested on an Apple M1 Mac IOMMU, this bug would instead silently corrupt freed memory, making this a ---truncated---

CVE-2024-40947	In the Linux kernel, the following vulnerability has been resolved: ima: Avoid blocking in RCU read-side critical s BUG: unable to handle kernel NULL pointer dereference at 0000000000000010 PGD 42f873067 P4D 0 Oops: 000 1286325 Comm: kubeletmonit.sh Kdump: loaded Tainted: P Hardware name: QEMU Standard PC (i440FX + PIIX 0010:ima_match_policy+0x84/0x450 Code: 49 89 fc 41 89 cf 31 ed 89 44 24 14 eb 1c 44 39 7b 18 74 26 41 83 ff 0 9c 01 00 00 <44> 85 73 10 74 ea 44 8b 6b 14 41 f6 c5 01 75 d4 41 f6 c5 02 74 0f RSP: 0018:ff71570009e07a80 E RBX: 0000000000000000 RCX: 0000000000000200 RDX: ffffffffad8dc7c0 RSI: 0000000024924925 RDI: ff3e27 R08: 0000000000000000 R09: ffffffffabfce739 R10: ff3e27810cc42400 R11: 0000000000000000 R12: ff3e27818 000000000000000c R15: 0000000000000001 FS: 00007f5195b51740(0000) GS:ff3e278b12d40000(0000) knlGS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 0000000000000010 CR3: 0000000626d24002 CR4: 00000000003 0000000000000000 DR2: 0000000000000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 000000000 +0x22/0x30 process_measurement+0xb0/0x830 ? page_add_file_rmap+0x15/0x170 ? alloc_set_pte+0x269/0x4c0 simple_xattr_get+0x75/0xa0 ? selinux_file_open+0x9d/0xf0 ima_file_check+0x64/0x90 path_openat+0x571/0x17 page_counter_try_charge+0x57/0xc0 ? files_cgroup_alloc_fd+0x38/0x60 ? __alloc_fd+0xd4/0x250 ? do_sys_oper +0x1bd/0x250 do_syscall_64+0x5d/0x1d0 entry_SYSCALL_64_after_hwframe+0x65/0xca Commit c7423dbdbc9 ima_filter_rule_match()") introduced call to ima_lsm_copy_rule within a RCU read-side critical section which con implies a possible sleep and violates limitations of RCU read-side critical sections on non-PREEMPT systems. Sle might cause synchronize_rcu() returning early and break RCU protection, allowing a UAF to happen. The root cau follows: | Thread A | Thread B | |ima_match_policy | | | rcu_read_lock | |ima_lsm_update_rule | | | synchronize_rcu ==> synchronize_rcu returns early | kfree(entry) | | | entry = entry->next| ==> UAF happens and entry now becom >action | ==> Accessing entry might cause panic. To fix this issue, we are converting all kmalloc that is called with GFP_ATOMIC. [PM: fixed missing comment, long lines, !CONFIG_IMA_LSM_RULES case]
CVE-2024-40953	In the Linux kernel, the following vulnerability has been resolved: KVM: Fix a data race on last_boosted_vcpu in k {READ,WRITE}_ONCE() to access kvm->last_boosted_vcpu to ensure the loads and stores are atomic. In the ext tears the stores, it's theoretically possible for KVM to attempt to get a vCPU using an out-of-bounds index, e.g. if t and is paired with a 32-bit load on a VM with 257 vCPUs: CPU0 CPU1 last_boosted_vcpu = 0xff; (last_boosted_v 0x01; i = (last_boosted_vcpu = 0x1ff) last_boosted_vcpu[7:0] = 0x00; vcpu = kvm->vcpu_array[0x1ff]; As detecte in kvm_vcpu_on_spin [kvm] / kvm_vcpu_on_spin [kvm] write to 0xffffc90025a92344 of 4 bytes by task 4340 on c kvm/../../../virt/kvm/kvm_main.c:4112) kvm handle_pause (arch/x86/kvm/vmx/vmx.c:5929) kvm_intel vmx_handl x86/kvm/vmx/vmx.c:6606) kvm_intel vcpu_run (arch/x86/kvm/x86.c:11107 arch/x86/kvm/x86.c:11211) kvm kvm x86.c:?) kvm kvm_vcpu_ioctl (arch/x86/kvm/../../../virt/kvm/kvm_main.c:?) kvm __se_sys_ioctl (fs/ioctl.c:52 fs/io (fs/ioctl.c:890) x64_sys_call (arch/x86/entry/syscall_64.c:33) do_syscall_64 (arch/x86/entry/common.c:?) entry_S entry/entry_64.S:130) read to 0xffffc90025a92344 of 4 bytes by task 4342 on cpu 4: kvm_vcpu_on_spin (arch/x86 kvm handle_pause (arch/x86/kvm/vmx/vmx.c:5929) kvm_intel vmx_handle_exit (arch/x86/kvm/vmx/vmx.c:? arch vcpu_run (arch/x86/kvm/x86.c:11107 arch/x86/kvm/x86.c:11211) kvm kvm_arch_vcpu_ioctl_run (arch/x86/kvm/x kvm/../../../virt/kvm/kvm_main.c:?) kvm __se_sys_ioctl (fs/ioctl.c:52 fs/ioctl.c:904 fs/ioctl.c:890) __x64_sys_ioctl syscall_64.c:33) do_syscall_64 (arch/x86/entry/common.c:?) entry_SYSCALL_64_after_hwframe (arch/x86/entry -> 0x00000000
CVE-2024-40965	In the Linux kernel, the following vulnerability has been resolved: i2c: lpi2c: Avoid calling clk_get_rate during tra clk_get_rate for each transfer, lock the clock rate and cache the value. A deadlock has been observed while adding When this clock provider adds its clock, the clk mutex is locked already, it needs to access i2c, which in return nee
CVE-2024-40967	In the Linux kernel, the following vulnerability has been resolved: serial: imx: Introduce timeout when waiting on t second for USR2_TXDC to be set, we avoid a potential deadlock. In case of the timeout, there is not much we can and optimistically try to continue.
CVE-2024-40969	In the Linux kernel, the following vulnerability has been resolved: f2fs: don't set RO when shutting down f2fs Shut due to readonly, which causes a deadlock like below. f2fs_ioc_shutdown(F2FS_GOING_DOWN_FULLSYNC) is freeze_super - f2fs_stop_checkpoint() - f2fs_handle_critical_error - sb_start_write - set RO - waiting - bdev_thaw sb_rdonly() - f2fs_stop_discard_thread -> wait for kthread_stop(discard_thread);
CVE-2024-40970	In the Linux kernel, the following vulnerability has been resolved: Avoid hw_desc array overrun in dw-axi-dmac I and in which each descriptor is composed by 3 segments, resulting in the DMA channel descs_allocated to be 9. Si considering the descs_allocated, this scenario would result in a kernel panic (hw_desc array will be overrun). To fi to the axi_dma_desc structure, where we keep the number of allocated hw_descs (axi_desc_alloc()) and use it in ax correctly. Additionally I propose to remove the axi_chan_start_first_queued() call after completing the transfer, sin (started descriptors can be interrupted and transfer ignored due to DMA channel not being enabled).
CVE-2024-40971	In the Linux kernel, the following vulnerability has been resolved: f2fs: remove clear SB_INLINECRYPT flag in c SB_INLINECRYPT flag will be clear and re-set. If create new file or open file during this gap, these files will not lead to data corruption if wrappedkey_v0 is enable. Thread A: Thread B: -f2fs_remount -f2fs_file_open or f2fs_ne SB_INLINECRYPT flag -fscrypt_select_encryption_impl -parse_options <- set SB_INLINECRYPT again
CVE-2024-40973	In the Linux kernel, the following vulnerability has been resolved: media: mtk-vcodec: potential null pointer defere devm_kzalloc() needs to be checked to avoid NULL pointer deference. This is similar to CVE-2022-3113.

CVE-2024-40975	In the Linux kernel, the following vulnerability has been resolved: platform/x86: x86-android-tablets: Unregister d subsystems support a device getting removed while there are still consumers of the device with a reference to the d regulator subsystem. If a regulator gets unregistered while there are still drivers holding a reference a WARN() at d e.g.: WARNING: CPU: 1 PID: 1587 at drivers/regulator/core.c:5829 regulator_unregister Hardware name: Intel Co BYT-T FFD8, BIOS BLADE_21.X64.0005.R00.1504101516 FFD8_X64_R_2015_04_10_1516 04/10/2015 RIP: <TASK> regulator_unregister devres_release_group i2c_device_remove device_release_driver_internal bus_remo x86_android_tablet_remove On the Lenovo Yoga Tablet 2 series the bq24190 charger chip also provides a 5V boo devices connected to the micro USB port, the bq24190-charger driver exports this as a Vbus regulator. On the 830 is controlled by a platform_device and x86_android_tablet_remove() removes platform_device-s before i2c_clients on the 1380 (13") model there is a lc824206xa micro-USB switch connected over I2C and the extcon driver for tha client *must* be registered first, because that creates the regulator with the lc824206xa listed as its consumer. If th lc824206xa driver will end up getting a dummy regulator. Since in this case both the regulator provider and consu that the consumer is unregistered first is to unregister the I2C devices in reverse order of in which they were create problems in the future change x86_android_tablet_remove() to unregister all device types in reverse order.
CVE-2024-40976	In the Linux kernel, the following vulnerability has been resolved: drm/lima: mask irqs in timeout path before hard a rendering job might take just long enough to trigger the drm sched job timeout handler but also still complete bef handler. This runs into race conditions not expected by the timeout handler. In some very specific cases it currently lima_pm_idle, with a stack dump such as: [10136.669170] WARNING: CPU: 0 PID: 0 at drivers/gpu/drm/lima/lim +0xa0/0xb0 ... [10136.669459] pc : lima_devfreq_record_idle+0xa0/0xb0 ... [10136.669628] Call trace: [10136.66 [10136.669646] lima_sched_pipe_task_done+0x5c/0xb0 [10136.669656] lima_gp_irq_handler+0xa8/0x120 [1013 +0x48/0x160 [10136.669679] handle_irq_event+0x4c/0xc0 We can prevent that race condition entirely by maskin handler, at which point we give up on waiting for that job entirely. The irqs will be enabled again at the next hard r the timeout handler.
CVE-2024-40979	In the Linux kernel, the following vulnerability has been resolved: wifi: ath12k: fix kernel crash during resume Cu not properly handled, resulting in kernel crash in case DMA remap is not supported: BUG: Bad page state in proce mapcount:0 mapping:0000000000000000 index:0x0 pfn:0x36e80 page dumped because: nonzero _refcount Call T __free_pages_ok __free_pages dma_direct_free dma_free_attrs ath12k_qmi_free_target_mem_chunk ath12k_qmi_ Once ath12k module is loaded, firmware sends memory request to host. In case DMA remap not supported, ath12k in allocating with large segment size: ath12k_pci 0000:04:00.0: qmi firmware request memory request ath12k_pci 7077888 ath12k_pci 0000:04:00.0: qmi mem seg type 4 size 8454144 ath12k_pci 0000:04:00.0: qmi dma allocatio with small size ath12k_pci 0000:04:00.0: qmi delays mem_request 2 ath12k_pci 0000:04:00.0: qmi firmware reque back with more but small segments and allocation succeeds: ath12k_pci 0000:04:00.0: qmi mem seg type 1 size 52 seg type 1 size 524288 ath12k_pci 0000:04:00.0: qmi mem seg type 1 size 524288 ath12k_pci 0000:04:00.0: qmi m 0000:04:00.0: qmi mem seg type 1 size 524288 ath12k_pci 0000:04:00.0: qmi mem seg type 1 size 524288 ath12k 524288 ath12k_pci 0000:04:00.0: qmi mem seg type 1 size 262144 ath12k_pci 0000:04:00.0: qmi mem seg type 1 mem seg type 1 size 524288 ath12k_pci 0000:04:00.0: qmi mem seg type 1 size 524288 ath12k_pci 0000:04:00.0: 0000:04:00.0: qmi mem seg type 1 size 524288 ath12k_pci 0000:04:00.0: qmi mem seg type 4 size 524288 ath12k 524288 ath12k_pci 0000:04:00.0: qmi mem seg type 4 size 524288 ath12k_pci 0000:04:00.0: qmi mem seg type 4 mem seg type 4 size 524288 ath12k_pci 0000:04:00.0: qmi mem seg type 4 size 524288 ath12k_pci 0000:04:00.0: 0000:04:00.0: qmi mem seg type 4 size 524288 ath12k_pci 0000:04:00.0: qmi mem seg type 4 size 524288 ath12k 524288 ath12k_pci 0000:04:00.0: qmi mem seg type 4 size 524288 ath12k_pci 0000:04:00.0: qmi mem seg type 4 mem seg type 4 size 524288 ath12k_pci 0000:04:00.0: qmi mem seg type 4 size 524288 ath12k_pci 0000:04:00.0: 0000:04:00.0: qmi mem seg type 4 size 524288 ath12k_pci 0000:04:00.0: qmi mem seg type 4 size 65536 ath12k_ 524288 Now ath12k is working. If suspend is triggered, firmware will be reloaded during resume. As same as befo first. In ath12k_qmi_msg_mem_request_cb() segment count and size are assigned: ab->qmi.mem_seg_count == 2 >qmi.target_mem[1].size == 8454144 Then allocation failed like before and ath12k_qmi_free_target_mem_chunk(the first segment is skipped because its v.addr is cleared due to allocation failure: chunk->v.addr = dma_alloc_cohe because it has not been freed. While freeing the second segment, a size of 8454144 is passed to dma_free_coherent allocated at the first time firmware is loaded, before suspend. So its real size is 524288, much smaller than 845414 some memory which is in use and thus cras ---truncated---
CVE-2024-40982	In the Linux kernel, the following vulnerability has been resolved: ssb: Fix potential NULL pointer dereference in s function first attempts to convert the 'dev' pointer to 'struct ssb_device *'. However, it mistakenly dereferences 'dev potentially leading to a NULL pointer dereference if 'dev' is NULL. To fix this issue, move the NULL check before the pointer is valid before attempting to use it. Found by Linux Verification Center (linuxtesting.org) with SVACE
CVE-2024-41001	In the Linux kernel, the following vulnerability has been resolved: io_uring/sqpoll: work around a potential audit m there's a memory leak related to connect handling: unreferenced object 0xffff0001093bdf00 (size 128): comm "iou- hex dump (first 32 bytes): 02 00 fa ea 7f 00 00 01 00 00 00 00 00 00 00 00 ..... 00 00 00 00 00 00 00 00 00 00 0 (crc 2e481b1a): [<00000000c0a26af4>] kmemleak_alloc+0x30/0x38 [<000000009c30bb45>] kmalloc_trace+0x22 __audit_sockaddr+0xd0/0x138 [<0000000089a93e34>] move_addr_to_kernel+0x1a0/0x1f8 [<000000000b4e80e6 [<00000000abfbcd99>] io_submit_sqes+0x588/0x1e48 [<00000000e7c25e07>] io_sq_thread+0x8a4/0x10e4 [<00 +0x10/0x20 which can can happen if: 1) The command type does something on the prep side that triggers an audit before this that triggered an audit call inside ->issue(), where we have audit_uring_entry() and audit_uring_exit(). W operation before the SQPOLL does anything.
CVE-2024-41002	In the Linux kernel, the following vulnerability has been resolved: crypto: hisilicon/sec - Fix memory leak for sec r resources. When releasing resources, it need to release the AIV resources at the same time. Otherwise, memory lea to the sec resource release function.

CVE-2024-41011	In the Linux kernel, the following vulnerability has been resolved: drm/amdkfd: don't allow mapping the MMIO H right offset in that case. The GPU has an unused 4K area of the register BAR space into which you can remap regis this space to allow userspace (CPU or GPU) to flush the HDP when it updates VRAM. However, on systems with of MMIO space.
CVE-2024-41013	In the Linux kernel, the following vulnerability has been resolved: xfs: don't walk off the end of a directory data bl xfs_dir2_data_unused and xfs_dir2_data_entry to make sure don't stray beyond valid memory region. Before patch offset of the dup and dep is within the range. So in a crafted image, if last entry is xfs_dir2_data_unused, we can ch 1 byte of space. In the next traversal, this space will be considered as dup or dep. We may encounter an out of bour In the patch, we make sure that the remaining bytes large enough to hold an unused entry before accessing xfs_dir2 XFS_DIR2_DATA_ALIGN byte aligned. We also make sure that the remaining bytes large enough to hold a diren xfs_dir2_data_entry.
CVE-2024-41014	In the Linux kernel, the following vulnerability has been resolved: xfs: add bounds checking to xlog_recover_proce of the space occupied by fixed members of xlog_op_header in the xlog_recover_process_data. We can create a cra bounds read by following these steps: 1) Mount an image of xfs, and do some file operations to leave records 2) Be subsequent steps to simulate abnormal exit. Because umount will ensure that tail_blk and head_blk are the same, w xlog_recover_process_data 3) Write a tool to parse and modify the copied image in step 2 4) Make the end of the x xlog_rec_header->h_size 5) xlog_rec_header->h_num_logops++ 6) Modify xlog_rec_header->h_crc Fix: Add a ch access fixed members of xlog_op_header.
CVE-2024-41015	In the Linux kernel, the following vulnerability has been resolved: ocfs2: add bounds checking to ocfs2_check_dir_ ocfs2_dir_entry to make sure all members of ocfs2_dir_entry don't stray beyond valid memory region.
CVE-2024-41016	In the Linux kernel, the following vulnerability has been resolved: ocfs2: strict bound check before memcmp in ocf 'non-indexed', which saved with additional space requested. It's better to check if the memory is out of bound befor comes from crafted poisonous images.
CVE-2024-41017	In the Linux kernel, the following vulnerability has been resolved: jfs: don't walk off the end of ealist Add a check sure each ea stays within the ealist.
CVE-2024-41019	In the Linux kernel, the following vulnerability has been resolved: fs/ntfs3: Validate ff offset This adds sanity chec >first_free at first, but walking through by ff without any check. If the second ff is a large offset. We may encounte
CVE-2024-41020	In the Linux kernel, the following vulnerability has been resolved: filelock: Fix fcntl/close race recovery compat pa ("filelock: Remove locks reliably when fcntl/close race is detected"), I missed that there are two copies of the code the version for 64-bit offsets on 32-bit kernels. Thanks to Greg KH for stumbling over this while doing the stable b compat path for 32-bit kernels.
CVE-2024-41022	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: Fix signedness bug in sdma_v4_0 needs to be signed for the error handling to work.
CVE-2024-41023	In the Linux kernel, the following vulnerability has been resolved: sched/deadline: Fix task_struct reference leak D test with linux-rt: stress-ng --cyclic 30 --timeout 30 --minimize --quiet kmemleak frequently reported a memory lea object 0xffff8881305b8000 (size 16136): comm "stress-ng", pid 614, jiffies 4294883961 (age 286.412s) object hex 00 00 00 00 00 00 00 00 00 00 .@..... 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ..... debug hex 00 00 00 00 00 00 00 00 00 00 S..... backtrace: [<00000000046b6790>] dup_task_struct+0x30/0x540 [<000000 [<00000000ced59777>] kernel_clone+0xb0/0x770 [<00000000a50befdc>] __do_sys_clone+0xb6/0xf0 [<000000 [<00000000552900ff>] entry_SYSCALL_64_after_hwframe+0x6e/0x76 The issue occurs in start_dl_timer(), whi and sets a timer. The timer callback, dl_task_timer, is supposed to decrement the reference count upon expiration. I before the timer expires and cancels it, the reference count is not decremented, leading to the leak. This patch fixes reference count is properly decremented when the timer is canceled.
CVE-2024-41027	In the Linux kernel, the following vulnerability has been resolved: Fix userfaultfd_api to return EINVAL as expect not set in the Kernel config we fail silently and return all the available features. However, the man page indicates w fix this issue since we can end up with a Kernel warning should a program request the feature UFFD_FEATURE_ config not set with this feature. [200.812896] WARNING: CPU: 91 PID: 13634 at mm/memory.c:1660 zap_pte_r linked in: [200.869387] CPU: 91 PID: 13634 Comm: userfaultfd Kdump: loaded Not tainted 6.9.0-rc5+ #8 [200.8 R6525/0N7YGH, BIOS 2.7.3 03/30/2022 [200.885052] RIP: 0010:zap_pte_range+0x43d/0x660
CVE-2024-41030	In the Linux kernel, the following vulnerability has been resolved: ksmbd: discard write access to the directory ope to be opened with the write access. However, some writing flags set by client result in adding write access on serve FUSE file system. Simply, let's discard the write access when opening a directory. list_add corruption. next is NUL BUG at lib/list_debug.c:26! pc : __list_add_valid+0x88/0xbc lr : __list_add_valid+0x88/0xbc Call trace: __list_ad +0x11c/0x170 fuse_open_common+0x284/0x5e8 fuse_dir_open+0x14/0x24 do_dentry_open+0x2a4/0x4e0 dentry handle_ksmbd_work+0x478/0x5ec process_one_work+0x1b4/0x448 worker_thread+0x25c/0x430 kthread+0x104
CVE-2024-41034	In the Linux kernel, the following vulnerability has been resolved: nilfs2: fix kernel bug on rename operation of br directory operation on broken directory on nilfs2, __block_write_begin_int() called to prepare block write may fail the folio/page size. This is because nilfs_dotdot(), which gets parent directory reference entry ("..") of the directory consistency enough, and may return location exceeding folio/page size for broken directories. Fix this issue by che in the first chunk of the directory in nilfs_dotdot().

CVE-2024-41035	In the Linux kernel, the following vulnerability has been resolved: USB: core: Fix duplicate endpoint bug by cleari identified a bug in usbcore (see the Closes: tag below) caused by our assumption that the reserved bits in an endpoi always be 0. As a result of the bug, the endpoint_is_duplicate() routine in config.c (and possibly other routines as v for distinct endpoints, even though they have the same direction and endpoint number. This can lead to confusion, descriptors with matching endpoint numbers and directions, where one was interrupt and the other was bulk). To fi bEndpointAddress when we parse the descriptor. (Note that both the USB-2.0 and USB-3.1 specs say these bits are to make a copy of the descriptor earlier in usb_parse_endpoint() and use the copy instead of the original when chec
CVE-2024-41036	In the Linux kernel, the following vulnerability has been resolved: net: ks8851: Fix deadlock with the SPI chip vari are actually functional then there is a deadlock with the 'statelock' spinlock between ks8851_start_xmit_spi and ks8 lockup - CPU#0 stuck for 27s! call trace: queued_spin_lock_slowpath+0x100/0x284 do_raw_spin_lock+0x34/0x4 ks8851_start_xmit+0x14/0x20 netdev_start_xmit+0x40/0x6c dev_hard_start_xmit+0x6c/0xbc sch_direct_xmit+0x qdisc_run+0x24/0x3c net_tx_action+0xf8/0x130 handle_softirqs+0x1ac/0x1f0 __do_softirq+0x14/0x20 ____do_s +0x3c/0x58 do_softirq_own_stack+0x1c/0x28 __irq_exit_rcu+0x54/0x9c irq_exit_rcu+0x10/0x1c el1_interrupt+0 el1h_64_irq+0x64/0x68 __netif_schedule+0x6c/0x80 netif_tx_wake_queue+0x38/0x48 ks8851_irq+0xb8/0x2c8 i +0x10c/0x1b0 kthread+0xc8/0xd8 ret_from_fork+0x10/0x20 This issue has not been identified earlier because test and so spinlocks were actually NOPs. Now use spin_(un)lock_bh for TX queue related locking to avoid execution to a deadlock.
CVE-2024-41040	In the Linux kernel, the following vulnerability has been resolved: net/sched: Fix UAF when resolving a clash KAS BUG: KASAN: slab-use-after-free in tcf_ct_flow_table_process_conn+0x12b/0x380 [act_ct] Read of size 1 at add handler130/6469 Call Trace: <IRQ> dump_stack_lvl+0x48/0x70 print_address_description.constprop.0+0x33/0x3 +0xd0/0x120 __asan_load1+0x6c/0x80 tcf_ct_flow_table_process_conn+0x12b/0x380 [act_ct] tcf_ct_act+0x886/(+0xf8/0x1f0 fl_classify+0x355/0x360 [cls_flower] __tcf_classify+0x1fd/0x330 tcf_classify+0x21c/0x3c0 sch_han __netif_receive_skb_core.constprop.0+0xb25/0x1510 __netif_receive_skb_list_core+0x220/0x4c0 netif_receive_s napi_complete_done+0x157/0x3d0 gro_cell_poll+0xcf/0x100 __napi_poll+0x65/0x310 net_rx_action+0x30c/0x5c +0x82/0xc0 irq_exit_rcu+0xe/0x20 common_interrupt+0xa1/0xb0 </IRQ> <TASK> asm_common_interrupt+0x2 kasan_save_stack+0x38/0x70 kasan_set_track+0x25/0x40 kasan_save_alloc_info+0x1e/0x40 __kasan_krealloc+0 nf_ct_ext_add+0xed/0x230 [nf_conntrack] tcf_ct_act+0x1095/0x1350 [act_ct] tcf_action_exec+0xf8/0x1f0 fl_clas __tcf_classify+0x1fd/0x330 tcf_classify+0x21c/0x3c0 sch_handle_ingress.constprop.0+0x2c5/0x500 __netif_rece __netif_receive_skb_list_core+0x220/0x4c0 netif_receive_skb_list_internal+0x446/0x620 napi_complete_done+0 __napi_poll+0x65/0x310 net_rx_action+0x30c/0x5c0 __do_softirq+0x14f/0x491 Freed by task 6469: kasan_save_ +0x25/0x40 kasan_save_free_info+0x2b/0x60 ____kasan_slab_free+0x180/0x1f0 __kasan_slab_free+0x12/0x30 s __kmem_cache_free+0x1a2/0x2f0 kfree+0x78/0x120 nf_conntrack_free+0x74/0x130 [nf_conntrack] nf_ct_destro __nf_ct_resolve_clash+0x529/0x5d0 [nf_conntrack] nf_ct_resolve_clash+0xf6/0x490 [nf_conntrack] __nf_conntra tcf_ct_act+0x12ad/0x1350 [act_ct] tcf_action_exec+0xf8/0x1f0 fl_classify+0x355/0x360 [cls_flower] __tcf_classi sch_handle_ingress.constprop.0+0x2c5/0x500 __netif_receive_skb_core.constprop.0+0xb25/0x1510 __netif_recei netif_receive_skb_list_internal+0x446/0x620 napi_complete_done+0x157/0x3d0 gro_cell_poll+0xcf/0x100 __nap +0x30c/0x5c0 __do_softirq+0x14f/0x491 The ct may be dropped if a clash has been resolved but is still passed to t for further usage. This issue can be fixed by retrieving ct from skb again after confirming conntrack.
CVE-2024-41041	In the Linux kernel, the following vulnerability has been resolved: udp: Set SOCK_RCU_FREE earlier in udp_lib_ [0] in udp_v4_early_demux(). In udp_v[46]_early_demux() and sk_lookup(), we do not touch the refcount of the lo skb->destructor, so we check SOCK_RCU_FREE to ensure that the sk is safe to access during the RCU grace perio flagged for a bound socket after being put into the hash table. Moreover, the SOCK_RCU_FREE check is done too and sk_lookup(), so there could be a small race window: CPU1 CPU2 ---- ---- udp_v4_early_demux() udp_lib_get_ __udp4_lib_demux_lookup() | |- DEBUG_NET_WARN_ON_ONCE(sk_is_refcounted(sk)); `- sock_set_flag(sk, S bug in TCP and fixed it in commit 871019b22d1b ("net: set SOCK_RCU_FREE before inserting socket into hashta [0]: WARNING: CPU: 0 PID: 11198 at net/ipv4/udp.c:2599 udp_v4_early_demux+0x481/0xb70 net/ipv4/udp.c:25 11198 Comm: syz-executor.1 Not tainted 6.9.0-g93bda33046e7 #13 Hardware name: QEMU Standard PC (i440FX gd239552ce722-prebuilt.qemu.org 04/01/2014 RIP: 0010:udp_v4_early_demux+0x481/0xb70 net/ipv4/udp.c:2599 e9 31 ff d3 e3 81 e3 bf ef ff ff 89 de e8 2c 74 15 fe 85 db 0f 85 02 06 00 00 e8 9f 7a 15 fe <0f> 0b e8 98 7a 15 fe 4 52 RSP: 0018:ffffc9000ce3fa58 EFLAGS: 00010293 RAX: 0000000000000000 RBX: 0000000000000000 RCX: RSI: ffffffff8318c2f1 RDI: 0000000000000001 RBP: ffff88805a2dd6e0 R08: 0000000000000001 R09: 0000000(R11: 0001ffffffffffff R12: ffff88805a2dd680 R13: 0000000000000007 R14: ffff88800923f900 R15: ffff888054560 GS:ffff88807dc00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 C 000000003de4b002 CR4: 0000000000770ef0 DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000(00000000fffe0ff0 DR7: 0000000000000600 PKRU: 55555554 Call Trace: <TASK> ip_rcv_finish_core.constprop ip_rcv_finish+0xda/0x150 net/ipv4/ip_input.c:447 NF_HOOK include/linux/netfilter.h:314 [inline] NF_HOOK inc +0x16c/0x180 net/ipv4/ip_input.c:569 __netif_receive_skb_one_core+0xb3/0xe0 net/core/dev.c:5624 __netif_rece netif_receive_skb_internal net/core/dev.c:5824 [inline] netif_receive_skb+0x271/0x300 net/core/dev.c:5884 tun_rx tun_get_user+0x24db/0x2c50 drivers/net/tun.c:2002 tun_chr_write_iter+0x107/0x1a0 drivers/net/tun.c:2048 new_ vfs_write+0x76f/0x8d0 fs/read_write.c:590 ksys_write+0xbf/0x190 fs/read_write.c:643 __do_sys_write fs/read_w read_write.c:652 [inline] __x64_sys_write+0x41/0x50 fs/read_write.c:652 x64_sys_call+0xe66/0x1990 arch/x86/i do_syscall_x64 arch/x86/entry/common.c:52 [inline] do_syscall_64+0x4b/0x110 arch/x86/entry/common.c:83 entr RIP: 0033:0x7fc44a68bc1f Code: 89 54 24 18 48 89 74 24 10 89 7c 24 08 e8 e9 cf f5 ff 48 8b 54 24 18 48 8b 74 24 05 <48> 3d 00 f0 ff ff 77 31 44 89 c7 48 89 44 24 08 e8 3c d0 f5 ff 48 RSP: 002b:00007fc449126c90 EFLAGS: 00 RAX: ffffffffffffffda RBX: 00000000004bc050 RCX: 00007fc44a68bc1f R ---truncated---

CVE-2024-41042	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_tables: prefer nft_chain_validate nf detection because a cycle will result in a call stack overflow (ctx->level >= NFT_JUMP_STACK_SIZE). It also fo nft_lookup, so there appears no reason to iterate the maps again. nf_tables_check_loops() and all its helper functio time significantly, from 23s down to 12s. This also fixes a crash bug. Old loop detection code can result in unboun was hit at .... Oops: stack guard page: 0000 [#1] PREEMPT SMP KASAN CPU: 4 PID: 1539 Comm: nft Not tainte during validation of register stores. I can't see any actual reason to attempt to check for this from nft_validate_regis in progress, so we don't have a full picture of the rule graph. For nf-next it might make sense to either remove it or case we could catch an error earlier (for improved error reporting to userspace).
CVE-2024-41044	In the Linux kernel, the following vulnerability has been resolved: ppp: reject claimed-as-LCP but actually malform assumes valid LCP packets (with code from 1 to 7 inclusive), add 'ppp_check_packet()' to ensure that LCP packet h bytes, and reject claimed-as-LCP but actually malformed data otherwise.
CVE-2024-41046	In the Linux kernel, the following vulnerability has been resolved: net: ethernet: lantiq_etop: fix double free in deta descriptor is never incremented which results in the same skb being released multiple times.
CVE-2024-41048	In the Linux kernel, the following vulnerability has been resolved: skmsg: Skip zero length skb in sk_msg_recvmsg t sockmap_basic) on a Loongarch platform, the following kernel panic occurs: [...] Oops[#1]: CPU: 22 PID: 2824 C #18 Hardware name: LOONGSON Dabieshan/Loongson-TC542F0, BIOS Loongson-UDK2018 ... ... ra: 90000000 ERA: 9000000004162774 copy_page_to_iter+0x74/0x1c0 CRMD: 000000b0 (PLV0 -IE -DA +PG DACF=CC DA (PPLV0 +PIE +PWE) EUEN: 00000007 (+FPE +SXE +ASXE -BTE) ECFG: 00071c1d (LIE=0,2-4,10-12 VS=7) EsubCode=0) BADV: 0000000000000040 PRID: 0014c011 (Loongson-64bit, Loongson-3C5000) Modules linked xt_MASQUERADE xt_conntrack Process test_progs (pid: 2824, threadinfo=0000000000863a31, task=...) Stack : copy_page_to_iter+0x74/0x1c0 [<90000000048bf6c0>] sk_msg_recvmsg+0x120/0x560 [<90000000049f2b90>] t [<90000000049aae34>] inet_recvmsg+0x54/0x100 [<900000000481ad5c>] sock_recvmsg+0x7c/0xe0 [<900000000 [<900000000481e27c>] sys_recvfrom+0x1c/0x40 [<9000000004c076ec>] do_syscall+0x8c/0xc0 [<900000000037 ---[end trace 0000000000000000]--- Kernel panic - not syncing: Fatal exception Kernel relocated by 0x3510000 . 0x9000000004d70000 .bss @ 0x9000000006469400 ---[end Kernel panic - not syncing: Fatal exception]--- [...] T sockmap_skb_verdict_shutdown subtest in sockmap_basic. This crash is because a NULL pointer is passed to page to the different implementations depending on the architecture, page_address(NULL) will trigger a panic on Loong this bug was hidden on x86 platform for a while, but now it is exposed on Loongarch platform. The root cause is th put on the queue. This zero length skb is a TCP FIN packet, which was sent by shutdown(), invoked in test_sockma SHUT_WR); In this case, in sk_psock_skb_ingress_enqueue(), num_sge is zero, and no page is put to this sge (see empty sge is queued into ingress_msg list. And in sk_msg_recvmsg(), this empty sge is used, and a NULL page is to copy_page_to_iter(), which passes it to kmap_local_page() and to page_address(), then kernel panics. To solve t So in sk_msg_recvmsg(), if copy is zero, that means it's a zero length skb, skip invoking copy_page_to_iter(). We a copy_page_to_iter to check for is_fin in tcp_bpf.c.
CVE-2024-41049	In the Linux kernel, the following vulnerability has been resolved: filelock: fix potential use-after-free in posix_loc warning in trace_posix_lock_inode(). The request pointer had been changed earlier to point to a lock entry that was the tracepoint could fire, another task raced in and freed that lock. Fix this by moving the tracepoint inside the spin happen.
CVE-2024-41055	In the Linux kernel, the following vulnerability has been resolved: mm: prevent derefencing NULL ptr in pfn_secti sparsemem: fix race in accessing memory_section->usage") changed pfn_section_valid() to add a READ_ONCE() section_deactivate() where ms->usage can be cleared. The READ_ONCE() call, by itself, is not enough to prevent its value before dereferencing it.
CVE-2024-41059	In the Linux kernel, the following vulnerability has been resolved: hfsplus: fix uninit-value in copy_name [syzbot r value in sized_strscpy+0xc4/0x160 sized_strscpy+0xc4/0x160 copy_name+0x2af/0x320 fs/hfsplus/xattr.c:411 hfsp hfsplus/xattr.c:750 vfs_listxattr fs/xattr.c:493 [inline] listxattr+0x1f3/0x6b0 fs/xattr.c:840 path_listxattr fs/xattr.c:8 xattr.c:876 [inline] __se_sys_listxattr fs/xattr.c:873 [inline] __x64_sys_listxattr+0x16b/0x2f0 fs/xattr.c:873 x64_sy include/generated/asm/syscalls_64.h:195 do_syscall_x64 arch/x86/entry/common.c:52 [inline] do_syscall_64+0xc entry_SYSCALL_64_after_hwframe+0x77/0x7f Uninit was created at: slab_post_alloc_hook mm/slub.c:3877 [inl kmalloc_trace+0x57b/0xbe0 mm/slub.c:4065 kmalloc include/linux/slab.h:628 [inline] hfsplus_listxattr+0x4cc/0x1 fs/xattr.c:493 [inline] listxattr+0x1f3/0x6b0 fs/xattr.c:840 path_listxattr fs/xattr.c:864 [inline] __do_sys_listxattr fs/ fs/xattr.c:873 [inline] __x64_sys_listxattr+0x16b/0x2f0 fs/xattr.c:873 x64_sys_call+0x2ba0/0x3b50 arch/x86/inclu do_syscall_x64 arch/x86/entry/common.c:52 [inline] do_syscall_64+0xcf/0x1e0 arch/x86/entry/common.c:83 entr [Fix] When allocating memory to strbuf, initialize memory to 0.
CVE-2024-41060	In the Linux kernel, the following vulnerability has been resolved: drm/radeon: check bo_va->bo is non-NULL bef radeon_vm_clear_freed might clear bo_va->bo, so we have to check it before dereferencing it.
CVE-2024-41061	In the Linux kernel, the following vulnerability has been resolved: drm/amd/display: Fix array-index-out-of-bounds Potential out of bounds access in dml2_calculate_rq_and_dlg_params() because the value of out_lowest_state_idx array can be greater than 1. [How] Currently dml2 core specifies identical values for all FCLKChangeSupport elem avoid out of bounds access.

CVE-2024-41062	In the Linux kernel, the following vulnerability has been resolved: bluetooth/l2cap: sync sock recv cb and release T call to close the sock and hci_rx_work, where the former releases the sock and the latter accesses it without lock pr hci_rx_work l2cap_sock_release hci_acldata_packet l2cap_sock_kill l2cap_recv_frame sk_free l2cap_conless_cha processes the data that needs to be received before the sock is closed, then everything is normal; Otherwise, the wo receiving data. Add a chan mutex in the rx callback of the sock to achieve synchronization between the sock releas to NULL, avoid others use invalid sock pointer.
CVE-2024-41063	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: hci_core: cancel all works upon hci_ calling hci_release_dev() from hci_error_reset() due to hci_dev_put() from hci_error_reset() can cause deadlock at is called from hdev->req_workqueue which destroy_workqueue() needs to flush. We need to make sure that hdev-> are queued into hdev->workqueue and hdev->{power_on,error_reset} which are queued into hdev->req_workqueu destroy_workqueue(hdev->workqueue); destroy_workqueue(hdev->req_workqueue); are called from hci_release_ items from hci_unregister_dev() as soon as hdev->list is removed from hci_dev_list.
CVE-2024-41064	In the Linux kernel, the following vulnerability has been resolved: powerpc/eeh: avoid possible crash when edev-> during eeh_pe_report_edev(), edev->pdev will change and can cause a crash, hold the PCI rescan/remove lock whi
CVE-2024-41065	In the Linux kernel, the following vulnerability has been resolved: powerpc/pseries: Whitelist dtl slub object for co trace log from /sys/kernel/debug/powerpc/dtl/cpu-* results in a BUG() when the config CONFIG_HARDENED_U kernel BUG at mm/usercopy.c:102! Oops: Exception in kernel mode, sig: 5 [#1] LE PAGE_SIZE=64K MMU=Rad Modules linked in: xfs libcrc32c dm_service_time sd_mod t10_pi sg ibmvfc scsi_transport_fc ibmveth pseries_wd dm_log dm_mod fuse CPU: 27 PID: 1815 Comm: python3 Not tainted 6.10.0-rc3 #85 Hardware name: IBM,9040- of:IBM,FW1060.00 (NM1060_042) hv:phyp pSeries NIP: c0000000005d23d4 LR: c0000000005d23d0 CTR: 0000 TRAP: 0700 Not tainted (6.10.0-rc3) MSR: 8000000000029033 <SF,EE,ME,IR,DR,RI,LE> CR: 2828220f XER: 0 IRQMASK: 0 [... GPRs omitted ...] NIP [c0000000005d23d4] usercopy_abort+0x78/0xb0 LR [c0000000005d23d usercopy_abort+0x74/0xb0 (unreliable) __check_heap_object+0xf8/0x120 check_heap_object+0x218/0x240 __ch +0x17c/0x2c4 full_proxy_read+0x8c/0x110 vfs_read+0xdc/0x3a0 ksys_read+0x84/0x144 system_call_exception+ +0x15c/0x2ec --- interrupt: 3000 at 0x7fff81f3ab34 Commit 6d07d1cd300f ("usercopy: Restrict non-usercopy cach whitelisted areas in slab/slub objects can be copied to userspace when usercopy hardening is enabled using CONFI contains hypervisor dispatch events which are expected to be read by privileged users. Hence mark this safe for use usersize=DISPATCH_LOG_BYTES to whitelist the entire object.
CVE-2024-41066	In the Linux kernel, the following vulnerability has been resolved: ibmvnic: Add tx check to prevent skb leak Belo a reference to an skb during transmit: tx_buff[free_map[consumer_index]]->skb = new_skb; free_map[consumer_i consumer_index ++; Where variable data looks like this: free_map == [4, IBMVNIC_INVALID_MAP, IBMVNIC tx_buff == [skb=null, skb=<ptr>, skb=<ptr>, skb=null, skb=null] The driver has checks to ensure that free_map[co there was no check to ensure that this index pointed to an unused/null skb address. So, if, by some chance, our free then we were previously risking an skb memory leak. This could then cause tcp congestion control to stop sending Therefore, add a conditional to ensure that the skb address is null. If not then warn the user (because this is still a b pointer to prevent memleak/tcp problems.
CVE-2024-41067	In the Linux kernel, the following vulnerability has been resolved: btrfs: scrub: handle RST lookup error correctly RST feature, it would crash the following ASSERT() inside scrub_read_endio(): ASSERT(sector_nr < stripe->nr_s dump from btrfs_get_raid_extent_offset(), as we failed to find the RST entry for the range. [CAUSE] Inside scrub_ allocated a new bbio we immediately called btrfs_map_block() to make sure there was some RST range covering th we immediately call endio for the bbio, while the bbio is newly allocated, it's completely empty. Then inside scrub_ find the sector number (as bi_sector is no longer reliable if the bio is submitted to lower layers). And since the bio i any sector matching the sector, and return sector_nr == stripe->nr_sectors, triggering the ASSERT(). [FIX] Instead a new bbio, call btrfs_map_block() first. Since our only objective of calling btrfs_map_block() is only to update str btrfs_alloc_bio(). This new timing would avoid the problem of handling empty bbio completely, and in fact fixes a if the submission thread is the only owner of the pending_io, the scrub would never finish (since we didn't decrease cause of RST lookup failure still needs to be addressed.
CVE-2024-41068	In the Linux kernel, the following vulnerability has been resolved: s390/sclp: Fix sclp_init() cleanup on failure If sc up: if there are multiple failing calls to sclp_init() sclp_state_change_event will be added several times to sclp_reg_ warning: ------------[cut here]------------ list_add double add: new=000003ffe1598c10, prev=000003ffe1598bf0, ne CPU: 0 PID: 1 at lib/list_debug.c:35 __list_add_valid_or_report+0xde/0xf8 CPU: 0 PID: 1 Comm: swapper/0 Not 0404c00180000000 000003ffe0d6076a (__list_add_valid_or_report+0xe2/0xf8) R:0 T:1 IO:0 EX:0 Key:0 M:1 W: Trace: [<000003ffe0d6076a>] __list_add_valid_or_report+0xe2/0xf8 ([<000003ffe0d60766>] __list_add_valid_or sclp_init+0x40e/0x450 [<000003ffe00009f2>] do_one_initcall+0x42/0x1e0 [<000003ffe15b77a6>] do_initcalls+0 kernel_init_freeable+0x1ba/0x1f8 [<000003ffe0d6650e>] kernel_init+0x2e/0x180 [<000003ffe000301c>] __ret_fr ret_from_fork+0xa/0x30 Fix this by removing sclp_state_change_event from sclp_reg_list when sclp_init() fails.
CVE-2024-41069	In the Linux kernel, the following vulnerability has been resolved: ASoC: topology: Fix references to freed memor release memory used by it, so having pointer references directly into topology file contents is wrong. Use devm_km

CVE-2024-41070	In the Linux kernel, the following vulnerability has been resolved: KVM: PPC: Book3S HV: Prevent UAF in kvm_ Al reported a possible use-after-free (UAF) in kvm_spapr_tce_attach_iommu_group(). It looks up `stt` from tablefd doing fdput() on the returned fd. After the fdput() the tablefd is free to be closed by another thread. The close calls release_spapr_tce_table() (via call_rcu()) which frees `stt`. Although there are calls to rcu_read_lock() in kvm_spa not sufficient to prevent the UAF, because `stt` is used outside the locked regions. With an artifcial delay after the triggers the race, KASAN detects the UAF: BUG: KASAN: slab-use-after-free in kvm_spapr_tce_attach_iommu_ addr c000200027552c30 by task kvm-vfio/2505 CPU: 54 PID: 2505 Comm: kvm-vfio Not tainted 6.10.0-rc3-next- GTH POWER9 0x4e1202 opal:skiboot-v6.5.3-35-g1851b2a06 PowerNV Call Trace: dump_stack_lvl+0xb4/0x108 kasan_report+0x118/0x2b0 __asan_load4+0xb8/0xd0 kvm_spapr_tce_attach_iommu_group+0x298/0x720 [kvm] kvm_device_ioctl+0x144/0x240 [kvm] sys_ioctl+0x62c/0x1810 system_call_exception+0x190/0x440 system_call by task 0: ... kfree+0xec/0x3e0 release_spapr_tce_table+0xd4/0x11c [kvm] rcu_core+0x568/0x16a0 handle_softirq +0x6c/0x90 do_softirq_own_stack+0x58/0x90 __irq_exit_rcu+0x218/0x2d0 irq_exit+0x30/0x80 arch_local_irq_re +0x1c/0x30 cpuidle_enter_state+0x134/0x5cc cpuidle_enter+0x6c/0xb0 call_cpuidle+0x7c/0x100 do_idle+0x394/ start_secondary+0x3fc/0x410 start_secondary_prolog+0x10/0x14 Fix it by delaying the fdput() until `stt` is no long function. To keep the patch minimal add a call to fdput() at each of the existing return paths. Future work can conv cleanup. With the fix in place the test case no longer triggers the UAF.
CVE-2024-41071	In the Linux kernel, the following vulnerability has been resolved: wifi: mac80211: Avoid address calculations via >n_channels must be set before req->channels[] can be used. This patch fixes one of the issues encountered in [1]. bounds in net/mac80211/scan.c:364:4 [83.964258] index 0 is out of range for type 'struct ieee80211_channel *[]' [<TASK> [83.964269] dump_stack_lvl+0x3f/0xc0 [83.964274] __ubsan_handle_out_of_bounds+0xec/0x110 [83 +0x2db/0x4b0 [83.964281] __ieee80211_start_scan+0x601/0x990 [83.964291] nl80211_trigger_scan+0x874/0x9 +0xe8/0x160 [83.964298] genl_rcv_msg+0x240/0x270 [...] [1] https://bugzilla.kernel.org/show_bug.cgi?id=21881
CVE-2024-41072	In the Linux kernel, the following vulnerability has been resolved: wifi: cfg80211: wext: add extra SIOCSIWSCAN add extra check whether number of channels passed via 'ioctl(sock, SIOCSIWSCAN, ...)' doesn't exceed IW_MAX with -EINVAL otherwise.
CVE-2024-41073	In the Linux kernel, the following vulnerability has been resolved: nvme: avoid double free special payload If a dis may fail before a new special payload is added, a double free will result. Clear the RQF_SPECIAL_LOAD when th
CVE-2024-41074	In the Linux kernel, the following vulnerability has been resolved: cachefiles: Set object to close if ondemand_id < the user mode, it may delete the request corresponding to the random id. And the request may have not been read y the open request will be done with the still reopen state in above case. As a result, the request corresponding to this function, so the read request is never completed and blocks other process. Fix this issue by simply set object to clos
CVE-2024-41075	In the Linux kernel, the following vulnerability has been resolved: cachefiles: add consistency check for copen/crea completing random copen/cread requests and crashing the system. Added checks are listed below: * Generic, copen can only complete read requests. * For copen, ondemand_id must not be 0, because this indicates that the request h the object corresponding to fd and req should be the same.
CVE-2024-41076	In the Linux kernel, the following vulnerability has been resolved: NFSv4: Fix memory leak in nfs4_set_security_l time we set a security xattr.
CVE-2024-41077	In the Linux kernel, the following vulnerability has been resolved: null_blk: fix validation of block size Block size be a power of 2. The current check does not validate this, so update the check. Without this patch, null_blk would with bs=1536 [1]. [axboe: remove unnecessary braces and != 0 check]
CVE-2024-41078	In the Linux kernel, the following vulnerability has been resolved: btrfs: qgroup: fix quota root leak after quota disa fail when cleaning the quota tree or when deleting the root from the root tree, we jump to the 'out' label without eve resulting in a leak of the root since fs_info->quota_root is no longer pointing to the root (we have set it to NULL ju doing a btrfs_put_root() call under the 'out' label. This is a problem that exists since qgroups were first added in 20 implementation and prototypes"), but back then we missed a kfree on the quota root and free_extent_buffer() calls then roots were not yet reference counted.
CVE-2024-41079	In the Linux kernel, the following vulnerability has been resolved: nvmet: always initialize cqe.result The spec doe (aka results) for the command queue entry need to be set to 0 when they are not used (not specified). Though, the ta but not for RDMA. Let's make RDMA behave the same and thus explicitly initializing the result field. This preven
CVE-2024-41080	In the Linux kernel, the following vulnerability has been resolved: io_uring: fix possible deadlock in io_register_io io_register_iowq_max_workers() function calls io_put_sq_data(), which acquires the sqd->lock without releasing t 009ad9f0c6ee ("io_uring: drop ctx->uring_lock before acquiring sqd->lock"), this can lead to a potential deadlock is released before calling io_put_sq_data(), and then it is re-acquired after the function call. This change ensures tha preventing the possibility of a deadlock.
CVE-2024-41081	In the Linux kernel, the following vulnerability has been resolved: ila: block BH in ila_output() As explained in co before using dst_cache"), net/core/dst_cache.c helpers need to be called with BH disabled. ila_output() is called fro context, and under rcu_read_lock(). We might be interrupted by a softirq, re-enter ila_output() and corrupt dst_cach local_bh_disable().

CVE-2024-41082	In the Linux kernel, the following vulnerability has been resolved: nvme-fabrics: use reserved tag for reg read/writ commands are issued by nvme command in the same time by user tasks, this may exhaust all tags of admin_q. If a before these commands finish, reconnect routine may fail to update nvme regs due to insufficient tags, which will workaround this issue, maybe we can let reg_read32()/reg_read64()/reg_write32() use reserved tags. This maybe sa will not issue connect command 2. For the enable ctrl / fw activate path, since connect and reg_xx() are called seria while reg_xx() use reserved tags.
CVE-2024-41087	In the Linux kernel, the following vulnerability has been resolved: ata: libata-core: Fix double free on error If e.g. t fails, we will jump to the err_out label, which will call devres_release_group(). devres_release_group() will trigger ata_host_release() calls kfree(host), so executing the kfree(host) in ata_host_alloc() will lead to a double free: kern opcode: 0000 [#1] PREEMPT SMP NOPTI CPU: 11 PID: 599 Comm: (udev-worker) Not tainted 6.10.0-rc5 #47 H + PIIX, 1996), BIOS 1.16.3-2.fc40 04/01/2014 RIP: 0010:kfree+0x2cf/0x2f0 Code: 5d 41 5e 41 5f 5d e9 80 d6 ff f 89 da RSP: 0018:ffffc90000f377f0 EFLAGS: 00010246 RAX: ffff888112b1f2c0 RBX: ffff888112b1f2c0 RCX: ff RSI: ffffffffc02c9de5 RDI: ffff888112b1f2c0 RBP: ffffc90000f37830 R08: 0000000000000000 R09: 00000000000 617461203a736b6e R12: ffffea00044ac780 R13: ffff888100046400 R14: ffffffffc02c9de5 R15: 000000000000000 GS:ffff88813b380000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 C 0000000111724000 CR4: 0000000000750ef0 PKRU: 55555554 Call Trace: <TASK> ? __die_body.cold+0x19/0x do_error_trap+0x6a/0x90 ? kfree+0x2cf/0x2f0 ? exc_invalid_op+0x50/0x70 ? kfree+0x2cf/0x2f0 ? asm_exc_inval +0xf5/0x120 [libata] ? ata_host_alloc+0xf5/0x120 [libata] ? kfree+0x2cf/0x2f0 ata_host_alloc+0xf5/0x120 [libata] ahci_init_one+0x6c9/0xd20 [ahci] Ensure that we will not call kfree(host) twice, by performing the kfree() only if
CVE-2024-41088	In the Linux kernel, the following vulnerability has been resolved: can: mcp251xfd: fix infinite loop when xmit fai function fails, the driver stops processing messages, and the interrupt routine does not return, running indefinitely e Error messages: [441.298819] mcp251xfd spi2.0 can0: ERROR in mcp251xfd_start_xmit: -16 [441.306498] mcp buffer not empty. (seq=0x000017c7, tef_tail=0x000017cf, tef_head=0x000017d0, tx_head=0x000017d3). ... and re when multiple devices share the same SPI interface. And there is concurrent access to the bus. The problem occurs mcp251xfd_start_xmit() fails. Consequently, the driver skips one TX package while still expecting a response in m issue by starting a workqueue to write the tx obj synchronously if err = -EBUSY. In case of another error, decreme stack, and drop the message. [mkl: use more imperative wording in patch description]
CVE-2024-41089	In the Linux kernel, the following vulnerability has been resolved: drm/nouveau/dispnv04: fix null pointer derefere nv17_tv_get_hd_modes(), the return value of drm_mode_duplicate() is assigned to mode, which will lead to a poss drm_mode_duplicate(). The same applies to drm_cvt_mode(). Add a check to avoid null pointer dereference.
CVE-2024-41092	In the Linux kernel, the following vulnerability has been resolved: drm/i915/gt: Fix potential UAF by revoke of fer reporting the following issue triggered by igt@i915_selftest@live@hangcheck on ADL-P and similar machines: <6 intel_hangcheck_live_selftests/igt_reset_evict_fence ... <6> [414.068804] i915 0000:00:02.0: [drm] GT0: GUC: su i915 0000:00:02.0: [drm] GT0: GUC: SLPC enabled <3> [414.070354] Unable to pin Y-tiled fence; err:-4 <3> [41 GEM_BUG_ON(!i915_active_is_idle(&fence->active)) ... <4>[609.603992] ------------[cut here]------------ <2>[gpu/drm/i915/gt/intel_ggtt_fencing.c:301! <4>[609.604003] invalid opcode: 0000 [#1] PREEMPT SMP NOPTI < Comm: kworker/u64:3 Tainted: G U W 6.9.0-CI_DRM_14785-g1ba62f8cea9c+ #1 <4>[609.604008] Hardware na Platform/AlderLake-P DDR4 RVP, BIOS RPLPFWI1.R00.4035.A00.2301200723 01/20/2023 <4>[609.604010] V [i915] <4>[609.604149] RIP: 0010:i915_vma_revoke_fence+0x187/0x1f0 [i915] ... <4>[609.604271] Call Trace: <4>[609.604716] __i915_vma_evict+0x2e9/0x550 [i915] <4>[609.604852] __i915_vma_unbind+0x7c/0x160 [i9 +0x24/0xa0 [i915] <4>[609.605098] i915_vma_destroy+0x2f/0xa0 [i915] <4>[609.605210] __i915_gem_object_ <4>[609.605330] __i915_gem_free_objects.isra.0+0x6a/0xc0 [i915] <4>[609.605440] process_scheduled_works similar failures reported by CI from other IGT tests, observed on other platforms. Before commit 63baf4f3d587 ("d before unbinding a GGTT fence"), i915_vma_revoke_fence() was waiting for idleness of vma->active via fence_up >active in order for the fence_update() to be able to wait selectively on that one instead of vma->active since only i then, another commit 0d86ee35097a ("drm/i915/gt: Make fence revocation unequivocal") replaced the call to fence with only fence_write(), and also added that GEM_BUG_ON(!i915_active_is_idle(&fence->active)) in front. No ju might then expect idleness of vma->fence->active without first waiting on it. The issue can be potentially caused by registers on one side and sequential execution of signal callbacks invoked on completion of a request that was using parallel to revocation of those fence registers. Fix it by waiting for idleness of vma->fence->active in i915_vma_re 24bb052d3dd499c5956abad5f7d8e4fd07da7fb1)
CVE-2024-41093	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: avoid using null object of framebu directly, get object from framebuffer by calling drm_gem_fb_get_obj() and return error code when object is null to
CVE-2024-41095	In the Linux kernel, the following vulnerability has been resolved: drm/nouveau/dispnv04: fix null pointer derefere nv17_tv_get_ld_modes(), the return value of drm_mode_duplicate() is assigned to mode, which will lead to a possi drm_mode_duplicate(). Add a check to avoid npd.

CVE-2024-41097	In the Linux kernel, the following vulnerability has been resolved: usb: atm: cxacru: fix endpoint checking in cxacr an old issue [1] that occurs due to incomplete checking of present usb endpoints. As such, wrong endpoints types n in turn triggers a warning in usb_submit_urb(). Fix the issue by verifying that required endpoint types are present f account cmd endpoint type. Unfortunately, this patch has not been tested on real hardware. [1] Syzbot report: usb 1 WARNING: CPU: 0 PID: 8667 at drivers/usb/core/urb.c:502 usb_submit_urb+0xed2/0x18a0 drivers/usb/core/urb. Comm: kworker/0:4 Not tainted 5.14.0-rc4-syzkaller #0 Hardware name: Google Google Compute Engine/Google Workqueue: usb_hub_wq hub_event RIP: 0010:usb_submit_urb+0xed2/0x18a0 drivers/usb/core/urb.c:502 ... Call atm/cxacru.c:649 cxacru_card_status+0x22/0xd0 drivers/usb/atm/cxacru.c:760 cxacru_bind+0x7ac/0x11a0 drivers +0x321/0x1ae0 drivers/usb/atm/usbatm.c:1055 cxacru_usb_probe+0xdf/0x1e0 drivers/usb/atm/cxacru.c:1363 usb_ core/driver.c:396 call_driver_probe drivers/base/dd.c:517 [inline] really_probe+0x23c/0xcd0 drivers/base/dd.c:595 drivers/base/dd.c:747 driver_probe_device+0x4c/0x1a0 drivers/base/dd.c:777 __device_attach_driver+0x20b/0x2f +0x15f/0x1e0 drivers/base/bus.c:427 __device_attach+0x228/0x4a0 drivers/base/dd.c:965 bus_probe_device+0x1e +0xc2f/0x2180 drivers/base/core.c:3354 usb_set_configuration+0x113a/0x1910 drivers/usb/core/message.c:2170 u drivers/usb/core/generic.c:238 usb_probe_device+0xd9/0x2c0 drivers/usb/core/driver.c:293
CVE-2024-41098	In the Linux kernel, the following vulnerability has been resolved: ata: libata-core: Fix null pointer dereference on ata_host_alloc() fails, ata_host_release() will get called. However, the code in ata_host_release() tries to free ata_p can lead to the following: BUG: unable to handle page fault for address: 0000000000003990 PGD 0 P4D 0 Oops: 0 CPU: 10 PID: 594 Comm: (udev-worker) Not tainted 6.10.0-rc5 #44 Hardware name: QEMU Standard PC (i440FX 04/01/2014 RIP: 0010:ata_host_release.cold+0x2f/0x6e [libata] Code: e4 4d 63 f4 44 89 e2 48 c7 c6 90 ad 32 c0 48 0018:ffffc90000ebb968 EFLAGS: 00010246 RAX: 0000000000000041 RBX: ffff88810fb52e78 RCX: 0000000000 RSI: ffff88813b3218c0 RDI: ffff88813b3218c0 RBP: ffff88810fb52e40 R08: 0000000000000000 R09: 6c65725f7 73692033203a746e R12: 0000000000000004 R13: 0000000000000000 R14: 0000000000000011 R15: 0000000000 GS:ffff88813b300000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 C 00000001122a2000 CR4: 0000000000750ef0 PKRU: 55555554 Call Trace: <TASK> ? __die_body.cold+0x19/0x exc_page_fault+0x7e/0x180 ? asm_exc_page_fault+0x26/0x30 ? ata_host_release.cold+0x2f/0x6e [libata] ? ata_ho release_nodes+0x35/0xb0 devres_release_group+0x113/0x140 ata_host_alloc+0xed/0x120 [libata] ata_host_alloc_ +0x6c9/0xd20 [ahci] Do not access ata_port struct members unconditionally.
CVE-2024-42063	In the Linux kernel, the following vulnerability has been resolved: bpf: Mark bpf prog stack with kmsan_unposion_ reported uninit memory usages during map_{lookup,delete}_elem. ========== BUG: KMSAN: uninit-value in _ devmap.c:441 [inline] BUG: KMSAN: uninit-value in dev_map_lookup_elem+0xf3/0x170 kernel/bpf/devmap.c:79 bpf/devmap.c:441 [inline] dev_map_lookup_elem+0xf3/0x170 kernel/bpf/devmap.c:796 ____bpf_map_lookup_ele bpf_map_lookup_elem+0x5c/0x80 kernel/bpf/helpers.c:38 ___bpf_prog_run+0x13fe/0xe0f0 kernel/bpf/core.c:199 bpf/core.c:2237 ========== The reproducer should be in the interpreter mode. The C reproducer is trying to run (18) r1 = map[id:49] 4: (b7) r8 = 16777216 5: (7b) *(u64 *)(r10 -8) = r8 6: (bf) r2 = r10 7: (07) r2 += -229 ^^^^^^ dev_map_lookup_elem#1543472 11: (95) exit It is due to the "void *key" (r2) passed to the helper. bpf allows unin the right privileges. This patch uses kmsan_unpoison_memory() to mark the stack as initialized. This should addres *key" argument during map_{lookup,delete}_elem.
CVE-2024-42064	In the Linux kernel, the following vulnerability has been resolved: drm/amd/display: Skip pipe if the pipe idx not s idx not set properly [how] Add code to skip the pipe that idx not set properly
CVE-2024-42065	In the Linux kernel, the following vulnerability has been resolved: drm/xe: Add a NULL check in xe_ttm_stolen_m the mgr is not NULL.
CVE-2024-42066	In the Linux kernel, the following vulnerability has been resolved: drm/xe: Fix potential integer overflow in page s >page_alignment to u64 before bit-shifting to prevent overflow when assigning to min_page_size.
CVE-2024-42067	In the Linux kernel, the following vulnerability has been resolved: bpf: Take return from set_memory_rox() into ac set_memory_rox() can fail, leaving memory unprotected. Check return and bail out when bpf_jit_binary_lock_ro()
CVE-2024-42068	In the Linux kernel, the following vulnerability has been resolved: bpf: Take return from set_memory_ro() into acc set_memory_ro() can fail, leaving memory unprotected. Check its return and take it into account as an error.
CVE-2024-42070	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_tables: fully validate NFT_DATA_ store validation for NFT_DATA_VALUE is conditional, however, the datatype is always either NFT_DATA_VAL requires a new helper function to infer the register type from the set datatype so this conditional check can be remo leaked through the registers.

CVE-2024-42076	In the Linux kernel, the following vulnerability has been resolved: net: can: j1939: Initialize unused data in j1939_ in raw_recvmsg() [1]. j1939_send_one() creates full frame including unused data, but it doesn't initialize it. This ca by initializing unused data. [1] BUG: KMSAN: kernel-infoleak in instrument_copy_to_user include/linux/instrume kernel-infoleak in copy_to_user_iter lib/iov_iter.c:24 [inline] BUG: KMSAN: kernel-infoleak in iterate_ubuf inclu KMSAN: kernel-infoleak in iterate_and_advance2 include/linux/iov_iter.h:245 [inline] BUG: KMSAN: kernel-inf linux/iov_iter.h:271 [inline] BUG: KMSAN: kernel-infoleak in _copy_to_iter+0x366/0x2520 lib/iov_iter.c:185 ins instrumented.h:114 [inline] copy_to_user_iter lib/iov_iter.c:24 [inline] iterate_ubuf include/linux/iov_iter.h:29 [inl iov_iter.h:245 [inline] iterate_and_advance include/linux/iov_iter.h:271 [inline] _copy_to_iter+0x366/0x2520 lib/i linux/uio.h:196 [inline] memcpy_to_msg include/linux/skbuff.h:4113 [inline] raw_recvmsg+0x2b8/0x9e0 net/can/ socket.c:1046 [inline] sock_recvmsg+0x2c4/0x340 net/socket.c:1068 ____sys_recvmsg+0x18a/0x620 net/socket.c net/socket.c:2845 do_recvmmsg+0x4fc/0xfd0 net/socket.c:2939 __sys_recvmmsg net/socket.c:3018 [inline] __do_ [inline] __se_sys_recvmmsg net/socket.c:3034 [inline] __x64_sys_recvmmsg+0x397/0x490 net/socket.c:3034 x64 include/generated/asm/syscalls_64.h:300 do_syscall_x64 arch/x86/entry/common.c:52 [inline] do_syscall_64+0xc entry_SYSCALL_64_after_hwframe+0x77/0x7f Uninit was created at: slab_post_alloc_hook mm/slub.c:3804 [inl [inline] kmem_cache_alloc_node+0x613/0xc50 mm/slub.c:3888 kmalloc_reserve+0x13d/0x4a0 net/core/skbuff.c: skbuff.c:668 alloc_skb include/linux/skbuff.h:1313 [inline] alloc_skb_with_frags+0xc8/0xbf0 net/core/skbuff.c:65 core/sock.c:2795 sock_alloc_send_skb include/net/sock.h:1842 [inline] j1939_sk_alloc_skb net/can/j1939/socket.c j1939/socket.c:1142 [inline] j1939_sk_sendmsg+0xc0a/0x2730 net/can/j1939/socket.c:1277 sock_sendmsg_nosec +0x30f/0x380 net/socket.c:745 ____sys_sendmsg+0x877/0xb60 net/socket.c:2584 ___sys_sendmsg+0x28d/0x3c0 socket.c:2667 [inline] __do_sys_sendmsg net/socket.c:2676 [inline] __se_sys_sendmsg net/socket.c:2674 [inline] net/socket.c:2674 x64_sys_call+0xc4b/0x3b50 arch/x86/include/generated/asm/syscalls_64.h:47 do_syscall_x64 a do_syscall_64+0xcf/0x1e0 arch/x86/entry/common.c:83 entry_SYSCALL_64_after_hwframe+0x77/0x7f Bytes 12 of size 16 starts at ffff888120969690 Data copied to user address 00000000200017c0 CPU: 1 PID: 5050 Comm: sy syzkaller-00031-g71b1543c83d6 #0 Hardware name: Google Google Compute Engine/Google Compute Engine, B
CVE-2024-42077	In the Linux kernel, the following vulnerability has been resolved: ocfs2: fix DIO failure due to insufficient transac ocfs2_dio_end_io_write() estimates number of necessary transaction credits using ocfs2_calc_extend_credits(). Th that the IO could be arbitrarily large and can contain arbitrary number of extents. Extent tree manipulations do ofte in all of the cases. For example if we have only single block extents in the tree, ocfs2_mark_extent_written() will e all the time and we will never extend the current transaction and eventually exhaust all the transaction credits if the Once that happens a WARN_ON(jbd2_handle_buffer_credits(handle) <= 0) is triggered in jbd2_journal_dirty_met in response to this error. This was actually triggered by one of our customers on a heavily fragmented OCFS2 files transaction always has enough credits for one extent insert before each call of ocfs2_mark_extent_written(). Hemir - not syncing: OCFS2: (device dm-1): panic forced after error" PID: xxx TASK: xxxx CPU: 5 COMMAND: "Subr ffffffff8c069932 #1 __crash_kexec at ffffffff8c1338fa #2 panic at ffffffff8c1d69b9 #3 ocfs2_handle_error at ffffffff ffffffffc0c88387 [ocfs2] #5 ocfs2_journal_dirty at ffffffffc0c51e98 [ocfs2] #6 ocfs2_split_extent at ffffffffc0c27ea3 ffffffffc0c28053 [ocfs2] #8 ocfs2_mark_extent_written at ffffffffc0c28347 [ocfs2] #9 ocfs2_dio_end_io_write at ff at ffffffffc0c2c0f5 [ocfs2] #11 dio_complete at ffffffff8c2b9fa7 #12 do_blockdev_direct_IO at ffffffff8c2bc09f #1 #14 generic_file_direct_write at ffffffff8c1dcf14 #15 __generic_file_write_iter at ffffffff8c1dd07b #16 ocfs2_file_ aio_write at ffffffff8c2cc72e #18 kmem_cache_alloc at ffffffff8c248dde #19 do_io_submit at ffffffff8c2ccada #20 entry_SYSCALL_64_after_hwframe at ffffffff8c8000ba
CVE-2024-42079	In the Linux kernel, the following vulnerability has been resolved: gfs2: Fix NULL pointer dereference in gfs2_log >sd_jdesc to NULL under the log flush lock to provide exclusion against gfs2_log_flush(). In gfs2_log_flush(), che dereferencing it. Otherwise, we could run into a NULL pointer dereference when outstanding glock work races wit run_queue -> do_xmote -> inode_go_sync -> gfs2_log_flush).
CVE-2024-42080	In the Linux kernel, the following vulnerability has been resolved: RDMA/restrack: Fix potential invalid address ac was set to KBUILD_MODNAME in ib_create_cq(), while if the module exited but forgot del this rdma_restrack_e in rdma_restrack_clean() when print the owner of this rdma_restrack_entry. These code is used to help find one for is not needed anymore, so delete them.
CVE-2024-42081	In the Linux kernel, the following vulnerability has been resolved: drm/xe/xe_devcoredump: Check NULL before a 'xe_devcoredump_snapshot *' and 'xe_device *' only if 'coredump' is not NULL. v2 - Fix commit messages. v3 - D v4 - Drop return check for coredump_to_xe. (Jose/Rodrigo) v5 - Modify misleading commit message. (Matt)
CVE-2024-42082	In the Linux kernel, the following vulnerability has been resolved: xdp: Remove WARN() from __xdp_reg_mem_ __xdp_reg_mem_model(). The warning occurs only if __mem_id_init_hash_table() returns an error. It returns the fails; 2. rhashtable_init() fails when some fields of rhashtable_params struct are not initialized properly. The secon static const rhashtable_params struct with valid fields. So, warning is only triggered when there is a problem with sense in using WARN() to handle this error and it can be safely removed. WARNING: CPU: 0 PID: 5065 at net/co +0x2d9/0x650 net/core/xdp.c:299 CPU: 0 PID: 5065 Comm: syz-executor883 Not tainted 6.8.0-syzkaller-05271-g Google Compute Engine/Google Compute Engine, BIOS Google 03/27/2024 RIP: 0010:__xdp_reg_mem_model+ Trace: xdp_reg_mem_model+0x22/0x40 net/core/xdp.c:344 xdp_test_run_setup net/bpf/test_run.c:188 [inline] bpf bpf/test_run.c:377 bpf_prog_test_run_xdp+0x813/0x11b0 net/bpf/test_run.c:1267 bpf_prog_test_run+0x33a/0x3b +0x48d/0x810 kernel/bpf/syscall.c:5649 __do_sys_bpf kernel/bpf/syscall.c:5738 [inline] __se_sys_bpf kernel/bpf/ +0x7c/0x90 kernel/bpf/syscall.c:5736 do_syscall_64+0xfb/0x240 entry_SYSCALL_64_after_hwframe+0x6d/0x7 (linuxtesting.org) with syzkaller.

CVE-2024-42084	In the Linux kernel, the following vulnerability has been resolved: ftruncate: pass a signed offset The old ftruncate(extension when called in compat mode on 64-bit architectures. As a result, passing a negative length accidentally s 2GiB and 4GiB. Changing the type of the compat syscall to the signed compat_off_t changes the behavior so it ins the truncate() syscall and the corresponding loff_t based variants are all correct already and do not suffer from this
CVE-2024-42086	In the Linux kernel, the following vulnerability has been resolved: iio: chemical: bme680: Fix overflows in compe compensate functions of the driver that there could be overflows of variables due to bit shifting ops. These implica they were mentioned in log message of Commit 1b3bd8592780 ("iio: chemical: Add support for Bosch BME680 s iio/20180728114028.3c1bbe81@archlinux/
CVE-2024-42087	In the Linux kernel, the following vulnerability has been resolved: drm/panel: ilitek-ili9881c: Fix warning with GP controls the reset GPIO using the non-sleeping gpiod_set_value() function. This complains loudly when the GPIO sleep, use gpiod_set_value_cansleep() to fix the issue.
CVE-2024-42089	In the Linux kernel, the following vulnerability has been resolved: ASoC: fsl-asoc-card: set priv->pdev before usin being used in fsl_asoc_card_audmux_init(). Move this assignment at the start of the probe function, so sub-functio fsl_asoc_card_audmux_init() dereferences priv->pdev to get access to the dev struct, used with dev_err macros. As NULL pointer dereference. Note that if priv->dev is dereferenced before assignment but never used, for example if won't crash probably due to compiler optimisations.
CVE-2024-42090	In the Linux kernel, the following vulnerability has been resolved: pinctrl: fix deadlock in create_pinctrl() when ha pinctrl_maps_mutex is acquired before calling add_setting(). If add_setting() returns -EPROBE_DEFER, create_pi pinctrl_free() attempts to acquire pinctrl_maps_mutex, which is already held by create_pinctrl(), leading to a poten by releasing pinctrl_maps_mutex before calling pinctrl_free(), preventing the deadlock. This bug was discovered a Security Testing (SAST) by Synopsys, Inc.
CVE-2024-42091	In the Linux kernel, the following vulnerability has been resolved: drm/xe: Check pat.ops before dumping PAT set running on brand new platform or when running as a VF. While the former is unlikely, the latter is valid (future) us will try to dump PAT settings by debugfs. It's better to check pointer to pat.ops instead of specific .dump hook, as every .ops variant.
CVE-2024-42092	In the Linux kernel, the following vulnerability has been resolved: gpio: davinci: Validate the obtained number of I from Device Tree. In case of broken DT due to any error this value can be any. Without this value validation there access in davinci_gpio_probe(). Validate the obtained nirq value so that it won't exceed the maximum number of II Center (linuxtesting.org) with SVACE.
CVE-2024-42093	In the Linux kernel, the following vulnerability has been resolved: net/dpaa2: Avoid explicit cpumask var allocatio CONFIG_CPUMASK_OFFSTACK=y kernel, explicit allocation of cpumask variable on stack is not recommende overflow. Instead, kernel code should always use *cpumask_var API(s) to allocate cpumask var in config-neutral v CONFIG_CPUMASK_OFFSTACK. Use *cpumask_var API(s) to address it.
CVE-2024-42094	In the Linux kernel, the following vulnerability has been resolved: net/iucv: Avoid explicit cpumask var allocation CONFIG_CPUMASK_OFFSTACK=y kernel, explicit allocation of cpumask variable on stack is not recommende overflow. Instead, kernel code should always use *cpumask_var API(s) to allocate cpumask var in config-neutral v CONFIG_CPUMASK_OFFSTACK. Use *cpumask_var API(s) to address it.
CVE-2024-42095	In the Linux kernel, the following vulnerability has been resolved: serial: 8250_omap: Implementation of Errata i2 timeout can be triggered, if this Erroneous interrupt is not cleared then it may leads to storm of interrupts, therefore www.ti.com/lit/pdf/sprz536 page 23
CVE-2024-42096	In the Linux kernel, the following vulnerability has been resolved: x86: stop playing stack games in profile_pc() Th timer-based profiling, which isn't really all that relevant any more to begin with, but it also ends up making assump necessarily valid. Basically, the code tries to account the time spent in spinlocks to the caller rather than the spinlo not worth the code complexity or the KASAN warnings when no serious profiling is done using timers anyway the stack layout that is only true in the simplest of cases. We've lost the comment at some point (I think when the 32-bi to say: Assume the lock function has either no stack frame or a copy of eflags from PUSHF. which explains why it off the stack pointer and then takes a minimal look at the values to just check if they might be eflags or the return p unlike kernel addresses but that basic stack layout assumption assumes that there isn't any lock debugging etc going a stack frame. It causes KASAN unhappiness reported for years by syzkaller [1] and others [2]. With no real practi the code. Just for historical interest, here's some background commits relating to this code from 2006: 0cb91a2293 during profiling for !FP kernels") 31679f38d886 ("Simplify profile_pc on x86-64") and a code unification from 20 profile_pc") but the basics of this thing actually goes back to before the git tree.
CVE-2024-42097	In the Linux kernel, the following vulnerability has been resolved: ALSA: emux: improve patch ioctl data validatio skipping over the main info block match that in load_guspatch(). In load_guspatch(), add checking that the specifie data, like load_data() already did.
CVE-2024-42098	In the Linux kernel, the following vulnerability has been resolved: crypto: ecdh - explicitly zeroize private_key pri parameter passed in by the caller (if present), or alternatively a newly generated private key. However, it is possibl generated key) which is shorter than the previous key. In that scenario, some key material from the previous key w is to explicitly zeroize the entire private_key array first. Note that this patch slightly changes the behavior of this fu failed, the old private_key would remain. Now, the private_key is always zeroized. This behavior is consistent with ecc_is_key_valid fails.

CVE-2024-42101	In the Linux kernel, the following vulnerability has been resolved: drm/nouveau: fix null pointer dereference in nou nouveau_connector_get_modes(), the return value of drm_mode_duplicate() is assigned to mode, which will lead t failure of drm_mode_duplicate(). Add a check to avoid npd.
CVE-2024-42102	In the Linux kernel, the following vulnerability has been resolved: Revert "mm/writeback: fix possible divide-by-z "mm: Avoid possible overflows in dirty throttling". Dirty throttling logic assumes dirty limits in page units fit into true (see patch 2/2 for more details). This patch (of 2): This reverts commit 9319b647902cbd5cc884ac08a8a6d54c ways. Firstly, the removed (u64) cast from the multiplication will introduce a multiplication overflow on 32-bit arc is actually common - the default settings with 4GB of RAM will trigger this). Secondly, the div64_u64() is unneces div64_ul() in case we want to be safe & cheap. Thirdly, if dirty thresholds are larger than 1<<32 pages, then dirty b spectacular ways anyway so trying to fix one possible overflow is just moot.
CVE-2024-42104	In the Linux kernel, the following vulnerability has been resolved: nilfs2: add missing check for inode numbers on mounting and unmounting a specific pattern of corrupted nilfs2 filesystem images causes a use-after-free of metada lru_add_fn(). As Jan Kara pointed out, this is because the link count of a metadata file gets corrupted to 0, and nilfs tries to delete that inode (ifile inode in this case). The inconsistency occurs because directories containing the inode not be visible in the namespace are read without checking. Fix this issue by treating the inode numbers of these inte when reading directory folios/pages. Also thanks to Hillf Danton and Matthew Wilcox for their initial mm-layer an
CVE-2024-42105	In the Linux kernel, the following vulnerability has been resolved: nilfs2: fix inode number range checks Patch seri to reserved inodes". This series fixes one use-after-free issue reported by syzbot, caused by nilfs2's internal inode b corrupted filesystem, and a couple of flaws that cause problems if the starting number of non-reserved inodes writt (or corruptly) changed from its default value. This patch (of 3): In the current implementation of nilfs2, "nilfs->ns_ reserved inode number, is read from the superblock, but its lower limit is not checked. As a result, if a number that of reserved inodes such as the root directory or metadata files is set in the super block parameter, the inode number NILFS_VALID_INODE) will not function properly. In addition, these test macros use left bit-shift calculations usi via the BIT macro, but the result of a shift calculation that exceeds the bit width of an integer is undefined in the C a large value other than the default value NILFS_USER_INO (=11), the macros may potentially malfunction deper by checking the lower bound of "nilfs->ns_first_ino" and by preventing bit shifts equal to or greater than the NILF test macros. Also, change the type of "ns_first_ino" from signed integer to unsigned integer to avoid the need for ty bound check introduced this time.
CVE-2024-42106	In the Linux kernel, the following vulnerability has been resolved: inet_diag: Initialize pad field in struct inet_diag value access in raw_lookup() [1]. Diag for raw sockets uses the pad field in struct inet_diag_req_v2 for the underly to the sdiag_raw_protocol field in struct inet_diag_req_raw. inet_diag_get_exact_compat() converts inet_diag_req the pad field uninitialized. So the issue occurs when raw_lookup() accesses the sdiag_raw_protocol field. Fix this b inet_diag_get_exact_compat(). Also, do the same fix in inet_diag_dump_compat() to avoid the similar issue in the in raw_lookup net/ipv4/raw_diag.c:49 [inline] BUG: KMSAN: uninit-value in raw_sock_get+0x657/0x800 net/ipv raw_diag.c:49 [inline] raw_sock_get+0x657/0x800 net/ipv4/raw_diag.c:71 raw_diag_dump_one+0xa1/0x660 net/i +0x7d9/0x980 inet_diag_get_exact_compat net/ipv4/inet_diag.c:1404 [inline] inet_diag_rcv_msg_compat+0x469/ sock_diag_rcv_msg+0x23d/0x740 net/core/sock_diag.c:282 netlink_rcv_skb+0x537/0x670 net/netlink/af_netlink.c net/core/sock_diag.c:297 netlink_unicast_kernel net/netlink/af_netlink.c:1335 [inline] netlink_unicast+0xe74/0x12 netlink_sendmsg+0x10c6/0x1260 net/netlink/af_netlink.c:1905 sock_sendmsg_nosec net/socket.c:730 [inline] __s socket.c:745 ____sys_sendmsg+0x7f0/0xb70 net/socket.c:2585 ___sys_sendmsg+0x271/0x3b0 net/socket.c:2639 __do_sys_sendmsg net/socket.c:2677 [inline] __se_sys_sendmsg net/socket.c:2675 [inline] __x64_sys_sendmsg+0 +0x135e/0x3ce0 arch/x86/include/generated/asm/syscalls_64.h:47 do_syscall_x64 arch/x86/entry/common.c:52 [in entry/common.c:83 entry_SYSCALL_64_after_hwframe+0x77/0x7f Uninit was stored to memory at: raw_sock_g raw_diag_dump_one+0xa1/0x660 net/ipv4/raw_diag.c:99 inet_diag_cmd_exact+0x7d9/0x980 inet_diag_get_exac [inline] inet_diag_rcv_msg_compat+0x469/0x530 net/ipv4/inet_diag.c:1426 sock_diag_rcv_msg+0x23d/0x740 ne +0x537/0x670 net/netlink/af_netlink.c:2564 sock_diag_rcv+0x35/0x40 net/core/sock_diag.c:297 netlink_unicast_k netlink_unicast+0xe74/0x1240 net/netlink/af_netlink.c:1361 netlink_sendmsg+0x10c6/0x1260 net/netlink/af_netli socket.c:730 [inline] __sock_sendmsg+0x332/0x3d0 net/socket.c:745 ____sys_sendmsg+0x7f0/0xb70 net/socket.c net/socket.c:2639 __sys_sendmsg net/socket.c:2668 [inline] __do_sys_sendmsg net/socket.c:2677 [inline] __se_sy __x64_sys_sendmsg+0x27e/0x4a0 net/socket.c:2675 x64_sys_call+0x135e/0x3ce0 arch/x86/include/generated/asm arch/x86/entry/common.c:52 [inline] do_syscall_64+0xd9/0x1e0 arch/x86/entry/common.c:83 entry_SYSCALL_6 variable req.i created at: inet_diag_get_exact_compat net/ipv4/inet_diag.c:1396 [inline] inet_diag_rcv_msg_compa sock_diag_rcv_msg+0x23d/0x740 net/core/sock_diag.c:282 CPU: 1 PID: 8888 Comm: syz-executor.6 Not tainted Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 1.16.3-2.fc40 04/01/2014

CVE-2024-42110	In the Linux kernel, the following vulnerability has been resolved: net: ntb_netdev: Move ntb_netdev_rx_handler() The following is emitted when using idxd (DSA) dmanegine as the data mover for ntb_transport that ntb_netdev us smp_processor_id() in preemptible [00000000] code: irq/52-idxd-por/14526 [74412.556784] caller is netif_rx_inte CPU: 6 PID: 14526 Comm: irq/52-idxd-por Not tainted 6.9.5 #5 [74412.569870] Hardware name: Intel Corporatio EGSDCRB1.E9I.1752.P05.2402080856 02/08/2024 [74412.581699] Call Trace: [74412.584514] <TASK> [74412 [74412.591129] check_preemption_disabled+0xc8/0xf0 [74412.596374] netif_rx_internal+0x42/0x130 [74412.60 ntb_netdev_rx_handler+0x66/0x150 [ntb_netdev] [74412.610985] ntb_complete_rxc+0xed/0x140 [ntb_transport] +0x53/0x80 [ntb_transport] [74412.623332] idxd_dma_complete_txd+0xe3/0x160 [idxd] [74412.628963] idxd_w irq_thread_fn+0x21/0x60 [74412.638134] ? irq_thread+0xa8/0x290 [74412.642218] irq_thread+0x1a0/0x290 [744 +0x10/0x10 [74412.651071] ? __pfx_irq_thread_dtor+0x10/0x10 [74412.656117] ? __pfx_irq_thread+0x10/0x10 [74412.664384] ? __pfx_kthread+0x10/0x10 [74412.668639] ret_from_fork+0x31/0x50 [74412.672716] ? __pfx_ ret_from_fork_asm+0x1a/0x30 [74412.681457] </TASK> The cause is due to the idxd driver interrupt completion threaded handler is not hard or soft interrupt context. However __netif_rx() can only be called from interrupt conte to allow completion via normal context for dmaengine drivers that utilize threaded irq handling. While the followir __netif_rx(), baebdf48c360 ("net: dev: Makes sure netif_rx() can be invoked in any context."), the change should've precedes this fix should've been using netif_rx_ni() or netif_rx_any_context().
CVE-2024-42114	In the Linux kernel, the following vulnerability has been resolved: wifi: cfg80211: restrict NL80211_ATTR_TXQ_ trigger softlockups, setting NL80211_ATTR_TXQ_QUANTUM to 2^31. We had a similar issue in sch_fq, fixed v fq: do not accept silly TCA_FQ_QUANTUM") watchdog: BUG: soft lockup - CPU#1 stuck for 26s! [kworker/1:0: 131135 hardirqs last enabled at (131134): [<ffff80008ae8778c>] __exit_to_kernel_mode arch/arm64/kernel/entry- enabled at (131134): [<ffff80008ae8778c>] exit_to_kernel_mode+0xdc/0x10c arch/arm64/kernel/entry-common.c [<ffff80008ae85378>] __el1_irq arch/arm64/kernel/entry-common.c:533 [inline] hardirqs last disabled at (131135 +0x24/0x68 arch/arm64/kernel/entry-common.c:551 softirqs last enabled at (125892): [<ffff80008907e82c>] neigh [inline] softirqs last enabled at (125892): [<ffff80008907e82c>] neigh_resolve_output+0x268/0x658 net/core/neigh (125896): [<ffff80008904166c>] local_bh_disable+0x10/0x34 include/linux/bottom_half.h:19 CPU: 1 PID: 24 Co rc7-syzkaller-gfda5695d692c #0 Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS mld mld_ifc_work pstate: 80400005 (Nzcv daif +PAN -UAO -TCO -DIT -SSBS BTYPE=--) pc : __list_del includ __list_del_entry include/linux/list.h:218 [inline] pc : list_move_tail include/linux/list.h:310 [inline] pc : fq_tin_deq ieee80211_tx_dequeue+0x6b8/0x3b4c net/mac80211/tx.c:3854 lr : __list_del_entry include/linux/list.h:218 [inline [inline] lr : fq_tin_dequeue include/net/fq_impl.h:112 [inline] lr : ieee80211_tx_dequeue+0x67c/0x3b4c net/mac80 x29: ffff800093d36a60 x28: ffff800093d36960 x27: dfff800000000000 x26: ffff0000d800ad50 x25: ffff0000d800a x23: ffff0000e0032468 x22: ffff0000e00324d4 x21: ffff0000d800abf0 x20: ffff0000d800abf8 x19: ffff0000d800ab 000000000000d476 x16: ffff8000805519dc x15: ffff7000127a6cc8 x14: 1ffff000127a6cc8 x13: 00000000000000 x10: 0000000000ff0100 x9 : 0000000000000000 x8 : 0000000000000000 x7 : 0000000000000000 x6 : 00000000 0000000000000008 x3 : ffff80008034c7fc x2 : ffff0000e0032468 x1 : 00000000da0e46b8 x0 : ffff0000e0032470 [inline] __list_del_entry include/linux/list.h:218 [inline] list_move_tail include/linux/list.h:310 [inline] fq_tin_dequ ieee80211_tx_dequeue+0x6b8/0x3b4c net/mac80211/tx.c:3854 wake_tx_push_queue net/mac80211/util.c:294 [inl +0x118/0x274 net/mac80211/util.c:315 drv_wake_tx_queue net/mac80211/driver-ops.h:1350 [inline] schedule_an [inline] ieee80211_queue_skb+0x18e8/0x2244 net/mac80211/tx.c:1664 ieee80211_tx+0x260/0x400 net/mac80211 net/mac80211/tx.c:2062 __ieee80211_subif_start_xmit+0xab8/0x122c net/mac80211/tx.c:4338 ieee80211_subif_s tx.c:4532 __netdev_start_xmit include/linux/netdevice.h:4903 [inline] netdev_start_xmit include/linux/netdevice.h [inline] dev_hard_start_xmit+0x27c/0x938 net/core/dev.c:3547 __dev_queue_xmit+0x1678/0x33fc net/core/dev.c: netdevice.h:3091 [inline] neigh_resolve_output+0x558/0x658 net/core/neighbour.c:1563 neigh_output include/net/ truncated---
CVE-2024-42115	In the Linux kernel, the following vulnerability has been resolved: jffs2: Fix potential illegal address access in jffs2 jffs2 file system,the following abnormal printouts were found: [2430.649000] Unable to handle kernel paging requ [2430.649622] Mem abort info: [2430.649829] ESR = 0x96000004 [2430.650115] EC = 0x25: DABT (current E FnV = 0 [2430.650795] EA = 0, S1PTW = 0 [2430.651032] FSC = 0x04: level 0 translation fault [2430.651446] = 0x00000004 [2430.652001] CM = 0, WnR = 0 [2430.652558] [0069696969696948] address between user and k error: Oops: 96000004 [#1] PREEMPT SMP [2430.654512] CPU: 2 PID: 20919 Comm: cat Not tainted 5.15.25-g name: linux,dummy-virt (DT) [2430.655517] pstate: 20000005 (nzCv daif -PAN -UAO -TCO -DIT -SSBS BTYP [2430.656630] lr : jffs2_free_inode+0x24/0x48 [2430.657051] sp : ffff800009eebd10 [2430.657355] x29: ffff80 0000000000000000 [2430.658327] x26: ffff000038f09d80 x25: 0080000000000000 x24: ffff800009d38000 [243 ffff000038f09d80 x21: ffff8000084f0d14 [2430.659434] x20: ffff0000bf9a6ac0 x19: 0169696969696940 x18: 00 ffff8000b6506000 x16: ffff800009eec000 x15: 0000000000004000 [2430.660637] x14: 0000000000000000 x13: [2430.661345] x11: 0004000800000000 x10: 0000000000000001 x9 : ffff8000084f0d14 [2430.662025] x8 : ffff0 x6 : 0000000003470302 [2430.662695] x5 : ffff00002e41dcc0 x4 : ffff0000bf9aa3b0 x3 : 0000000003470342 [24 x1 : ffff8000084f0d14 x0 : fffffc0000000000 [2430.664217] Call trace: [2430.664528] kfree+0x78/0x348 [2430. [2430.665233] i_callback+0x24/0x50 [2430.665528] rcu_do_batch+0x1ac/0x448 [2430.665892] rcu_core+0x28 +0x18/0x28 [2430.666473] __do_softirq+0x138/0x3cc [2430.666781] irq_exit+0xf0/0x110 [2430.667065] handl gic_handle_irq+0xac/0xe8 [2430.667739] call_on_irq_stack+0x28/0x54 The parameter passed to kfree was 5a5a5 the jffs_inode_info structure. It was found that all variables in the jffs_inode_info structure were 5a5a5a5a, except these variables are not initialized because they were set to 5a5a5a5a during memory testing, which is meant to dete is initialized in the function jffs2_i_init_once, while other members are initialized in the function jffs2_init_inode_ is called after iget_locked, but in the iget_locked function, the destroy_inode process is triggered, which releases th member of the inode is not initialized.In concurrent high pressure scenarios, iget_locked may enter the destroy_ino destroy_inode functionality of jffs2 only releases the target, the fix method is to set target to NULL in jffs2_i_init_

CVE-2024-42116	In the Linux kernel, the following vulnerability has been resolved: igc: fix a log entry using uninitialized netdev Du [5.133667] igc 0000:01:00.0 (unnamed net_device) (uninitialized): PHC added ^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^ is called very early, even before register_netdev() has been called. So the netdev_info() call works on a partially un igc_ptp_init() after register_netdev(), right after the media autosense check, just as in igb. Add a comment, just as i [5.200987] igc 0000:01:00.0 eth0: PHC added
CVE-2024-42117	In the Linux kernel, the following vulnerability has been resolved: drm/amd/display: ASSERT when failing to find find_disp_cfg_idx_by_plane_id and find_disp_cfg_idx_by_stream_id returns an array index and they return -1 whe number. [HOW] When this happens, call ASSERT(), and return a positive number (which is fewer than callers' arr 2 NEGATIVE_RETURNS issues reported by Coverity.
CVE-2024-42118	In the Linux kernel, the following vulnerability has been resolved: drm/amd/display: Do not return negative stream resource_stream_to_stream_idx returns an array index and it return -1 when not found; however, -1 is not a valid a happens, call ASSERT(), and return a zero instead. This fixes an OVERRUN and an NEGATIVE_RETURNS issu
CVE-2024-42119	In the Linux kernel, the following vulnerability has been resolved: drm/amd/display: Skip finding free audio for un ENGINE_ID_UNKNOWN = -1 and can not be used as an array index. Plus, it also means it is uninitialized and do return NULL. This fixes 2 OVERRUN issues reported by Coverity.
CVE-2024-42120	In the Linux kernel, the following vulnerability has been resolved: drm/amd/display: Check pipe offset before setti so checking its index before accessing the array. This fixes an OVERRUN issue reported by Coverity.
CVE-2024-42121	In the Linux kernel, the following vulnerability has been resolved: drm/amd/display: Check index msg_id before re array index and it cannot be a negative value, and therefore cannot be equal to MOD_HDCP_MESSAGE_ID_INV valid before reading and setting. This fixes 4 OVERRUN issues reported by Coverity.
CVE-2024-42122	In the Linux kernel, the following vulnerability has been resolved: drm/amd/display: Add NULL pointer check for of kzalloc before using it.
CVE-2024-42123	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: fix double free err_addr pointer w amdgpu_umc_bad_page_polling_timeout, the amdgpu_umc_handle_bad_pages will be run many times so that dou the err_addr to NULL to avoid the warnings.
CVE-2024-42124	In the Linux kernel, the following vulnerability has been resolved: scsi: qedf: Make qedf_execute_tmf() non-preem preemptible code in qedf_execute_tmf90. This results in BUG_ON() when running an RT kernel. [659.343280] B [00000000] code: sg_reset/3646 [659.343282] caller is qedf_execute_tmf+0x8b/0x360 [qedf]
CVE-2024-42125	In the Linux kernel, the following vulnerability has been resolved: wifi: rtw89: fw: scan offload prohibit all 6 GHz policy via BIOS to block uses of 6 GHz. In this case, 6 GHz sband will be NULL even if it is WiFi 7 chip. So, add
CVE-2024-42126	In the Linux kernel, the following vulnerability has been resolved: powerpc: Avoid nmi_enter/nmi_exit in real moc per cpu variables which can lead to kernel crash when invoked during real mode interrupt handling (e.g. early HMI comes from vmalloc area. Early HMI/MCE handlers are called through DEFINE_INTERRUPT_HANDLER_NMI calls. We don't see any issue when percpu allocation is from the embedded first chunk. However with CONFIG_NI enabled there are chances where percpu allocation can come from the vmalloc area. With kernel command line "pe percpu allocation to come from vmalloc area and can see kernel crash in machine_check_early: [1.215714] NIP [c +0x24/0x110 [1.215717] LR [c0000000000461a0] machine_check_early+0xf0/0x2c0 [1.215719] --- interrupt: 20 [0000000000000000] 0x0 (unreliable) [1.215722] [c000000fffd731b0] [0000000000000000] 0x0 [1.215724] [c00 machine_check_early_common+0x134/0x1f8 Fix this by avoiding use of nmi_enter()/nmi_exit() in real mode if pe
CVE-2024-42127	In the Linux kernel, the following vulnerability has been resolved: drm/lima: fix shared irq handling on driver remo interrupt handlers must be prepared to be called at any time. At driver removal time, the clocks are disabled early a very end of the remove process due to the devm usage. This is potentially a bug as the interrupts access device regi crash can be triggered by removing the driver in a kernel with CONFIG_DEBUG_SHIRQ enabled. This patch free callback so that the handlers are already unregistered by the time we fully disable clocks.
CVE-2024-42128	In the Linux kernel, the following vulnerability has been resolved: leds: an30259a: Use devm_mutex_init() for mut registered using devm_led_classdev_register() so they are automatically unregistered after module's remove() is do led_set_brightness() to turn off the LEDs and that callback uses mutex which was destroyed already in module's re
CVE-2024-42129	In the Linux kernel, the following vulnerability has been resolved: leds: mlxreg: Use devm_mutex_init() for mutex registered using devm_led_classdev_register() so they are automatically unregistered after module's remove() is do led_set_brightness() to turn off the LEDs and that callback uses mutex which was destroyed already in module's re
CVE-2024-42130	In the Linux kernel, the following vulnerability has been resolved: nfc/nci: Add the inconsistency check between th &(0x7f0000000740)=ANY=[@ANYBLOB="610501"], 0xf) Syzbot constructed a write() call with a data length o passed too little data to meet the basic requirements of the function nci_rf_intf_activated_ntf_packet(). Therefore, i length and count value to avoid problems caused by inconsistent data length and count.

CVE-2024-42131	In the Linux kernel, the following vulnerability has been resolved: mm: avoid overflows in dirty throttling logic Th assumptions that dirty limits in PAGE_SIZE units fit into 32-bit (so that various multiplications fit into 64-bits). If overflows, possible divisions by 0 etc. Fix these problems by never allowing so large dirty limits as they have dubi dirty_background_bytes interfaces we can just refuse to set so large limits. For dirty_ratio / dirty_background_ratio computed from the amount of available memory which can change due to memory hotplug etc. So when convertin we just don't allow the result to exceed UINT_MAX. This is root-only triggerable problem which occurs when the
CVE-2024-42134	In the Linux kernel, the following vulnerability has been resolved: virtio-pci: Check if is_avq is NULL [bug] In the vp_dev->is_avq is involved to determine whether it is admin virtqueue, but this function vp_dev->is_avq may be e does not assign a value to vp_dev->is_avq. [fix] Check whether it is vp_dev->is_avq before use. [test] Test with vi following command would crash the guest system After applying the patch, everything seems to be working fine.
CVE-2024-42135	In the Linux kernel, the following vulnerability has been resolved: vhost_task: Handle SIGKILL by flushing work device is closed, this has us handle SIGKILL by: 1. marking the worker as killed so we no longer try to use it with setting the virtqueue to worker mapping so no new works are queued. 3. running all the exiting works.
CVE-2024-42136	In the Linux kernel, the following vulnerability has been resolved: cdrom: rearrange last_media_change check to a syzkaller with the newly reintroduced signed integer wrap sanitizer we encounter this splat: [366.015950] UBSAN cdrom/cdrom.c:2361:33 [366.021089] -9223372036854775808 - 346321 cannot be represented in type '__s64' (ak executor.4 is using a deprecated SCSI ioctl, please convert it to SG_IO [366.027502] CPU: 5 PID: 28472 Comm: gb3ef86b5a957 #1 [366.027512] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 1.16.3-debia Call Trace: [366.027523] <TASK> [366.027533] dump_stack_lvl+0x93/0xd0 [366.027899] handle_overflow+0x multi_count 32 ignored [366.043924] cdrom_ioctl+0x2c3f/0x2d10 [366.063932] ? __pm_runtime_resume+0xe6/ +0x15d/0x1d0 [366.074624] ? __pfx_sr_block_ioctl+0x10/0x10 [366.077642] blkdev_ioctl+0x419/0x500 [366.0 Historically, the signed integer overflow sanitizer did not work in the kernel due to its interaction with `-fwrapv` bu newest version of Clang. It was re-enabled in the kernel with Commit 557f8c582a9ba8ab ("ubsan: Reintroduce sig check to not perform any arithmetic, thus not tripping the sanitizer.
CVE-2024-42137	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: qca: Fix BT enable failure again for 272970be3dab ("Bluetooth: hci_qca: Fix driver shutdown on closed serdev") will cause below regression issue: BT boot -> enable BT -> disable BT -> warm reboot -> BT enable failure if property enable-gpios is not configured wi is to fix a use-after-free issue within qca_serdev_shutdown() by adding condition to avoid the serdev is flushed or regression issue regarding above steps since the VSC is not sent to reset controller during warm reboot. Fixed by se qca_serdev_shutdown() once BT was ever enabled, and the use-after-free issue is also fixed by this change since th or wrote. Verified by the reported machine Dell XPS 13 9310 laptop over below two kernel commits: commit e00f coredump implementation for QCA") of bluetooth-next tree. commit b23d98d46d28 ("Bluetooth: btusb: Fix trigger linus mainline tree.
CVE-2024-42143	In the Linux kernel, the following vulnerability has been resolved: orangefs: fix out-of-bounds fsid access Arnd Be "orangefs_statfs() copies two consecutive fields of the superblock into the statfs structure, which triggers a warning Kara suggested an alternate way to do the patch to make it more readable. I ran both ideas through xfstests and both suggestion.
CVE-2024-42144	In the Linux kernel, the following vulnerability has been resolved: thermal/drivers/mediatek/lvts_thermal: Check N not NULL before using it.
CVE-2024-42145	In the Linux kernel, the following vulnerability has been resolved: IB/core: Implement a limit on UMAD receive L maintains received MAD packets in an unbounded list, poses a risk of uncontrolled growth. As user-space applicati of extraction may not match the rate of incoming packets, leading to potential list overflow. To address this, we int considering typical scenarios, such as OpenSM processing, which can handle approximately 100k packets per seco packets, we set the list size limit to 200k. Packets received beyond this limit are dropped, assuming they are likely user-space. Notably, packets queued on the receive list due to reasons like timed-out sends are preserved even whe
CVE-2024-42146	In the Linux kernel, the following vulnerability has been resolved: drm/xe: Add outer runtime_pm protection to xe_ doing any memory access should get their own runtime_pm outer references since they don't use the standard drive from the same driver. Found by pre-merge CI on adding WARN calls for unprotected inner callers: <6> [318.6397 xe_test_dmabuf_import_same_driver <4> [318.639957] ------------[cut here]------------ <4> [318.639967] xe 0000 protection <4> [318.640049] WARNING: CPU: 117 PID: 3832 at drivers/gpu/drm/xe/xe_pm.c:533 xe_pm_runtim
CVE-2024-42147	In the Linux kernel, the following vulnerability has been resolved: crypto: hisilicon/debugfs - Fix debugfs uninit pr the debugfs failure does not stop the probe. When debugfs initialization fails, jumping to the error branch will also operation. As a result, it may be released repeatedly during the regs uninit process. Therefore, the null check needs

CVE-2024-42148	In the Linux kernel, the following vulnerability has been resolved: bnx2x: Fix multiple UBSAN array-index-out-of when using a system with 32 physical cpu cores or more, or when the user defines a number of Ethernet queues gre using the num_queues module parameter. Currently there is a read/write out of bounds that occurs on the array "str the "bnx2x_fw_stats_req" struct in "drivers/net/ethernet/broadcom/bnx2x/bnx2x.h". Looking at the definition of th struct stats_query_entry query[FP_SB_MAX_E1x+ BNX2X_FIRST_QUEUE_QUERY_IDX]; FP_SB_MAX_E1 fast path interrupts and has a value of 16, while BNX2X_FIRST_QUEUE_QUERY_IDX has a value of 3 meaning accesses to "struct stats_query_entry query" are offset-ted by BNX2X_FIRST_QUEUE_QUERY_IDX, that means should not exceed FP_SB_MAX_E1x (16). However one of these queues is reserved for FCOE and thus the numbe [FP_SB_MAX_E1x -1] (15) if FCOE is enabled or [FP_SB_MAX_E1x] (16) if it is not. This is also described in a ethernet/broadcom/bnx2x/bnx2x.h just above the Macro definition of FP_SB_MAX_E1x. Below is the part of this * The total number of L2 queues, MSIX vectors and HW contexts (CIDs) is * control by the number of fast-path st FW). Each fast-path status block (FP-SB) aka non-default * status block represents an independent interrupts conte queue. However special L2 queues such * as the FCoE queue do not require a FP-SB and other components like * number of possible L2 queues * * If the maximum number of FP-SB available is X then: * a. If CNIC is supported of * regular L2 queues is Y=X-1 * b. In MF mode the actual number of L2 queues is Y= (X-1/MF_factor) * c. If th number of L2 queues * is Y+1 * d. The number of irqs (MSIX vectors) is either Y+1 (one extra for * slow-path int additional * FP interrupt context for the CNIC). * e. The number of HW context (CID count) is always X or X+1 if the FCoE L2 queue is always X. */ However this driver also supports NICs that use the E2 controller which can ha represented by FP_SB_MAX_E2. Looking at the commits when the E2 support was added, it was originally using ("bnx2x: Add 57712 support"). Back then FP_SB_MAX_E2 was set to 16 the same as E1x. However the driver wa E2 instead of having it be limited to the capabilities of the E1x. But as far as we can tell, the array "stats_query_ent SB available to the E1x cards as part of an oversignt when the driver was updated to take full advantage of the E2, greater queue size supported by E2 NICs, it causes the UBSAN warnings seen in the stack traces below. This patch query" array by replacing FP_SB_MAX_E1x with FP_SB_MAX_E2 to be large enough to handle both types of NI of-bounds in drivers/net/ethernet/broadcom/bnx2x/bnx2x_stats.c:1529:11 index 20 is out of range for type 'stats_qu systemd-network Not tainted 6.9.0-060900rc7-generic #202405052133 Hardware name: HP ProLiant DL360 Gen9
CVE-2024-42151	In the Linux kernel, the following vulnerability has been resolved: bpf: mark bpf_dummy_struct_ops.test_1 parame dummy_init_ret_value passes NULL as the first parameter of the test_1() function. Mark this parameter as nullable Otherwise, NULL check in the test_1() code: SEC("struct_ops/test_1") int BPF_PROG(test_1, struct bpf_dummy_ access state ... } Might be removed by verifier, thus triggering NULL pointer dereference under certain conditions.
CVE-2024-42152	In the Linux kernel, the following vulnerability has been resolved: nvmet: fix a possible leak when destroy a ctrl du we capture sq->ctrl early and if it is non-NULL we know that a ctrl was allocated (in the admin connect request ha clear ctrl->sqs and sq->ctrl (for nvme-loop primarily), and drop the final reference on the ctrl. However, a small wi starts (as a result of the client giving up and disconnecting) concurrently with the nvme admin connect cmd (which kill_and_confirm of sq->ref (i.e. the admin connect managed to get an sq live reference). In this case, sq->ctrl was a local variable in nvmet_sq_destroy. This prevented the final reference drop on the ctrl. Solve this by re-capturing completed, where for sure sq->ctrl reference is final, and move forward based on that. This issue was observed in a multiple ctrls simoutanuosly, creating a delay in allocating a ctrl leading up to this race window.
CVE-2024-42153	In the Linux kernel, the following vulnerability has been resolved: i2c: pnx: Fix potential deadlock warning from d del_timer_sync() is called in an interrupt context it throws a warning because of potential deadlock. The timer is us after a timeout so replacing the call with wait_for_completion_timeout() allows to remove the problematic timer ar
CVE-2024-42154	In the Linux kernel, the following vulnerability has been resolved: tcp_metrics: validate source addr length I don't s TCP_METRICS_ATTR_SADDR_IPV4 is at least 4 bytes long, and the policy doesn't have an entry for this attribu manually validated).
CVE-2024-42155	In the Linux kernel, the following vulnerability has been resolved: s390/pkey: Wipe copies of protected- and secure protected- nor secure-keys is accessible, this key material should only be visible to the calling process. So wipe all stack, even in case of an error.
CVE-2024-42156	In the Linux kernel, the following vulnerability has been resolved: s390/pkey: Wipe copies of clear-key structures for all IOCTLs, which convert a clear-key into a protected- or secure-key.
CVE-2024-42157	In the Linux kernel, the following vulnerability has been resolved: s390/pkey: Wipe sensitive data on failure Wipe copy_to_user() fails.
CVE-2024-42158	In the Linux kernel, the following vulnerability has been resolved: s390/pkey: Use kfree_sensitive() to fix Coccinel and kfree() with kfree_sensitive() to fix warnings reported by Coccinelle: WARNING opportunity for kfree_sensiti opportunity for kfree_sensitive/kvfree_sensitive (line 1643) WARNING opportunity for kfree_sensitive/kvfree_sen
CVE-2024-42159	In the Linux kernel, the following vulnerability has been resolved: scsi: mpi3mr: Sanitise num_phys Information is larger then size of this field shouldn't be allowed.
CVE-2024-42160	In the Linux kernel, the following vulnerability has been resolved: f2fs: check validation of fault attrs in f2fs_build fault attrs in parse_options(), let's fix to add check condition in f2fs_build_fault_attr(). - Use f2fs_build_fault_attr()

CVE-2024-42161	In the Linux kernel, the following vulnerability has been resolved: bpf: Avoid uninitialized value in BPF_CORE_R Use a default branch in the switch statement to initialize `val'.] GCC warns that `val' may be used uninitialized in th defined in bpf_core_read.h as: [...] unsigned long long val; \ [...] \ switch (__CORE_RELO(s, field, BYTE_SIZE)) break; \ case 2: val = *(const unsigned short *)p; break; \ case 4: val = *(const unsigned int *)p; break; \ case 8: val \ [...] val; \ } \ This patch adds a default entry in the switch statement that sets `val' to zero in order to avoid the war __builtin_preserve_field_info returns unexpected values for BPF_FIELD_BYTE_SIZE. Tested in bpf-next master.
CVE-2024-42162	In the Linux kernel, the following vulnerability has been resolved: gve: Account for stopped queues when reading I NIC might send us stats for a subset of queues. Without this change, gve_get_ethtool_stats might make an invalid a
CVE-2024-42223	In the Linux kernel, the following vulnerability has been resolved: media: dvb-frontends: tda10048: Fix integer ove can overflow a 32 bit integer when multiplied by pll_mfactor. Create a new 64 bit variable to hold the calculations.
CVE-2024-42224	In the Linux kernel, the following vulnerability has been resolved: net: dsa: mv88e6xxx: Correct check for empty l mv88e6xxx: Support multiple MDIO busses") mv88e6xxx_default_mdio_bus() has checked that the return value o to be intended to guard against the list chip->mdios being empty. However, it is not the correct check as the implem return NULL for empty lists. Instead, use list_first_entry_or_null() which does return NULL if the list is empty. Fl
CVE-2024-42225	In the Linux kernel, the following vulnerability has been resolved: wifi: mt76: replace skb_put with skb_put_zero A
CVE-2024-42226	In the Linux kernel, the following vulnerability has been resolved: usb: xhci: prevent potential failure in handle_tx_ Some transfer events don't always point to a TRB, and consequently don't have a endpoint ring. In these cases, func because if 'ep->skip' is set, the pointer to the endpoint ring is used. To prevent a potential failure and make the code code for a Transfer event without TRBs.
CVE-2024-42227	In the Linux kernel, the following vulnerability has been resolved: drm/amd/display: Fix overlapping copy within d &mode_lib->mp.Watermark and &locals->Watermark are the same address. memcpy may lead to unexpected beha
CVE-2024-42228	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: Using uninitialized value *size wh the size before calling amdgpu_vce_cs_reloc, such as case 0x03000001. V2: To really improve the handling we wo 0xffffffff.(Christian)
CVE-2024-42229	In the Linux kernel, the following vulnerability has been resolved: crypto: aead,cipher - zeroize key buffer after use variables temporarily holding cryptographic information should be zeroized once they are no longer needed. Accor buffers that previously held the private key.
CVE-2024-42230	In the Linux kernel, the following vulnerability has been resolved: powerpc/pseries: Fix scv instruction crash with K (reloc_on_exc), required for scv instruction support, before other CPUs have been shut down. This means they can disabled, which causes an interrupt at an unexpected entry location that crashes the kernel. Change the kexec seque have been brought down. As a refresher, the real-mode scv interrupt vector is 0x17000, and the fixed-location head implementing such high addresses so it was just decided not to support that interrupt at all.
CVE-2024-4317	Missing authorization in PostgreSQL built-in views pg_stats_ext and pg_stats_ext_exprs allows an unprivileged da and other statistics from CREATE STATISTICS commands of other users. The most common values may reveal c not otherwise read or results of functions they cannot execute. Installing an unaffected version only fixes fresh Post are created with the initdb utility after installing that version. Current PostgreSQL installations will remain vulnera the release notes. Within major versions 14-16, minor versions before PostgreSQL 16.3, 15.7, and 14.12 are affecte unaffected.
DSA-5349-1	gnutls28 - security update
DSA-5402-1	linux - security update
DSA-5453-1	linux - security update
DSA-5461-1	linux - security update
DSA-5475-1	linux - security update
DSA-5480-1	linux - security update
DSA-5523-1	curl - security update
DSA-5523-1	curl - security update
DSA-5570-1	nghttp2 - security update
DSA-5587-1	curl - security update
DSA-5587-1	curl - security update
DSA-5594-1	linux - security update
DSA-5681-1	linux - security update
DSA-5703-1	linux - security update

DSA-5730-1	linux - security update
GHSA-9h6g-pr28-7cqp	### Summary A ReDoS vulnerability occurs when nodemailer tries to parse img files with the parameter `attachDa event loop. Another flaw was found when nodemailer tries to parse an attachments with a embedded file, causing t Regex: /^data:((?:[^;]*;)*(?:[^,]*)),(.*)$/ Path: compile -> getAttachments -> _processDataUrl Regex: /(<img\b[^> [^"'>\s]+)/ Path: _convertDataImages ### PoC https://gist.github.com/francoatmega/890dd5053375333e40c6fdbcc francoatmega/9aab042b0b24968d7b7039818e8b2698 ### Impact ReDoS causes the event loop to stuck a specially
RHSA-2022:4991	XZ Utils is an integrated collection of user-space file compression utilities based on the Lempel-Ziv-Markov chain data compression. The algorithm provides a high compression ratio while keeping the decompression time short.
RHSA-2022:5056	The Common UNIX Printing System (CUPS) provides a portable printing layer for Linux, UNIX, and similar oper
RHSA-2022:5311	The libgcrypt library provides general-purpose implementations of various cryptographic algorithms.
RHSA-2022:5313	The curl packages provide the libcurl library and the curl utility for downloading files from servers using various p
RHSA-2022:5314	Expat is a C library for parsing XML documents.
RHSA-2022:5317	The libxml2 library is a development toolbox providing the implementation of various XML standards.
RHSA-2022:5696	The java-1.8.0-openjdk packages provide the OpenJDK 8 Java Runtime Environment and the OpenJDK 8 Java Sof
RHSA-2022:5809	The pcre2 package contains a new generation of the Perl Compatible Regular Expression libraries for implementin the same syntax and semantics as Perl.
RHSA-2022:6159	The curl packages provide the libcurl library and the curl utility for downloading files from servers using various p
RHSA-2022:6180	The rsync utility enables the users to copy and synchronize files locally or across a network. Synchronization with differences in files over the network instead of sending whole files. The rsync utility is also used as a mirroring too
RHSA-2022:6206	The systemd packages contain systemd, a system and service manager for Linux, compatible with the SysV and LS parallelism capabilities, uses socket and D-Bus activation for starting services, offers on-demand starting of daemo Linux cgroups. In addition, it supports snapshotting and restoring of the system state, maintains mount and automo transactional dependency-based service control logic. It can also work as a drop-in replacement for sysvinit.
RHSA-2022:6457	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2022:6463	The GNU Privacy Guard (GnuPG or GPG) is a tool for encrypting data and creating digital signatures, compliant w
RHSA-2022:6878	Expat is a C library for parsing XML documents.
RHSA-2022:7006	The java-1.8.0-openjdk packages provide the OpenJDK 8 Java Runtime Environment and the OpenJDK 8 Java Sof
RHSA-2022:7089	KSBA (pronounced Kasbah) is a library to make X.509 certificates as well as the CMS easily accessible by other a blocks of S/MIME and TLS.
RHSA-2022:7105	The gnutls packages provide the GNU Transport Layer Security (GnuTLS) library, which implements cryptograph TLS, and DTLS.
RHSA-2022:7106	The zlib packages provide a general-purpose lossless data compression library that is used by many different progr
RHSA-2022:7108	SQLite is a C library that implements an SQL database engine. A large subset of SQL92 is supported. A complete API is designed for convenience and ease of use. Applications that link against SQLite can enjoy the power and fle administrative hassles of supporting a separate database server.
RHSA-2022:7704	WebKitGTK is the port of the portable web rendering engine WebKit to the GTK platform.
RHSA-2022:7715	The libxml2 library is a development toolbox providing the implementation of various XML standards.
RHSA-2022:7720	The e2fsprogs packages provide a number of utilities for creating, checking, modifying, and correcting the ext2, ex
RHSA-2022:7745	FreeType is a free, high-quality, portable font engine that can open and manage font files. FreeType loads, hints, ar
RHSA-2022:7793	The rsync utility enables the users to copy and synchronize files locally or across a network. Synchronization with differences in files over the network instead of sending whole files. The rsync utility is also used as a mirroring too
RHSA-2023:0110	SQLite is a C library that implements an SQL database engine. A large subset of SQL92 is supported. A complete API is designed for convenience and ease of use. Applications that link against SQLite can enjoy the power and fle administrative hassles of supporting a separate database server.
RHSA-2023:0200	The java-11-openjdk packages provide the OpenJDK 11 Java Runtime Environment and the OpenJDK 11 Java Sof
RHSA-2023:0208	The java-1.8.0-openjdk packages provide the OpenJDK 8 Java Runtime Environment and the OpenJDK 8 Java Sof
RHSA-2023:1095	The zlib packages provide a general-purpose lossless data compression library that is used by many different progr

RHSA-2023:1140	The curl packages provide the libcurl library and the curl utility for downloading files from servers using various p
RHSA-2023:1252	Network Security Services (NSS) is a set of libraries designed to support the cross-platform development of securit
RHSA-2023:1332	Network Security Services (NSS) is a set of libraries designed to support the cross-platform development of securit
RHSA-2023:1335	OpenSSL is a toolkit that implements the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protoco cryptography library.
RHSA-2023:1895	The java-11-openjdk packages provide the OpenJDK 11 Java Runtime Environment and the OpenJDK 11 Java Sof
RHSA-2023:1908	The java-1.8.0-openjdk packages provide the OpenJDK 8 Java Runtime Environment and the OpenJDK 8 Java Sof
RHSA-2023:2963	The curl packages provide the libcurl library and the curl utility for downloading files from servers using various p
RHSA-2023:2963	The curl packages provide the libcurl library and the curl utility for downloading files from servers using various p
RHSA-2023:3106	The curl packages provide the libcurl library and the curl utility for downloading files from servers using various p
RHSA-2023:3555	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2023:4175	The java-11-openjdk packages provide the OpenJDK 11 Java Runtime Environment and the OpenJDK 11 Java Sof
RHSA-2023:4176	The java-1.8.0-openjdk packages provide the OpenJDK 8 Java Runtime Environment and the OpenJDK 8 Java Sof
RHSA-2023:4864	The Common UNIX Printing System (CUPS) provides a portable printing layer for Linux, UNIX, and similar oper
RHSA-2023:5615	The libssh2 packages provide a library that implements the SSH2 protocol.
RHSA-2023:5731	The java-1.8.0-openjdk packages provide the OpenJDK 8 Java Runtime Environment and the OpenJDK 8 Java Sof
RHSA-2023:5742	The java-11-openjdk packages provide the OpenJDK 11 Java Runtime Environment and the OpenJDK 11 Java Sof
RHSA-2023:5998	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2023:6885	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2023:7034	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2023:7743	The curl packages provide the libcurl library and the curl utility for downloading files from servers using various p
RHSA-2023:7783	PostgreSQL is an advanced object-relational database management system (DBMS).
RHSA-2024:0266	The java-11-openjdk packages provide the OpenJDK 11 Java Runtime Environment and the OpenJDK 11 Java Sof
RHSA-2024:0533	The gnutls packages provide the GNU Transport Layer Security (GnuTLS) library, which implements cryptograph TLS, and DTLS.
RHSA-2024:0606	OpenSSH is an SSH protocol implementation supported by a number of Linux, UNIX, and similar operating syste both the OpenSSH client and server.
RHSA-2024:0606	OpenSSH is an SSH protocol implementation supported by a number of Linux, UNIX, and similar operating syste both the OpenSSH client and server.
RHSA-2024:0811	The sudo packages contain the sudo utility which allows system
RHSA-2024:0894	MySQL is a multi-user, multi-threaded SQL database server. It consists of the MySQL server daemon (mysqld) an
RHSA-2024:1129	The curl packages provide the libcurl library and the curl utility for downloading files from servers using various p
RHSA-2024:1431	Ruby is an extensible, interpreted, object-oriented, scripting language. It has features to process text files and to pe
RHSA-2024:1510	Node.js is a software development platform for building fast and scalable
RHSA-2024:1822	The java-11-openjdk packages provide the OpenJDK 11 Java Runtime Environment and the OpenJDK 11 Java Sof
RHSA-2024:1879	The gnutls package provide the GNU Transport Layer Security (GnuTLS) library, which implements cryptographi and DTLS.
RHSA-2024:2463	The systemd packages contain systemd, a system and service manager for Linux, compatible with the SysV and LS parallelism capabilities, uses socket and D-Bus activation for starting services, offers on-demand starting of daemo Linux cgroups. In addition, it supports snapshotting and restoring of the system state, maintains mount and automo transactional dependency-based service control logic. It can also work as a drop-in replacement for sysvinit.

RHSA-2024:2512	The file command is used to identify a particular file according to the type of data the file contains. It can identify r Executable and Linkable Format (ELF) binary files, system libraries, RPM packages, and different graphics format
RHSA-2024:2679	The libxml2 library is a development toolbox providing the implementation of various XML standards.
RHSA-2024:2780	Node.js is a software development platform for building fast and scalable network applications in the JavaScript pr
RHSA-2024:2987	Python is an interpreted, interactive, object-oriented programming language that supports modules, classes, excepti dynamic typing. The python27 packages provide a stable release of Python 2.7 with a number of additional utilities PostgreSQL.
RHSA-2024:2988	The container-tools module contains tools for working with containers, notably podman, buildah, skopeo, and runc
RHSA-2024:3254	The container-tools module contains tools for working with containers, notably podman, buildah, skopeo, and runc
RHSA-2024:3271	The Berkeley Internet Name Domain (BIND) is an implementation of the Domain Name System (DNS) protocols. resolver library (routines for applications to use when interfacing with DNS); and tools for verifying that the DNS s
RHSA-2024:3346	Git Large File Storage (LFS) replaces large files such as audio samples, videos, datasets, and graphics with text poi contents on a remote server.
RHSA-2024:3546	Ruby is an extensible, interpreted, object-oriented, scripting language. It has features to process text files and to per
RHSA-2024:3588	The glibc packages provide the standard C libraries (libc), POSIX thread
RHSA-2024:3834	The gdk-pixbuf2 packages provide an image loading library that can be extended
RHSA-2024:3968	The container-tools module contains tools for working with containers, notably podman, buildah, skopeo, and runc
SUSE-SU-2023:4659-1	Security update for curl
SUSE-SU-2023:4891-1	Security update for ncurses
SUSE-SU-2024:0070-1	Security update for tar
SUSE-SU-2024:0136-1	Security update for pam
SUSE-SU-2024:0140-1	Security update for libssh
SUSE-SU-2024:0305-1	Security update for cpio
SUSE-SU-2024:0549-1	Security update for openssl-1_1
SUSE-SU-2024:0555-1	Security update for libxml2
SUSE-SU-2024:0973-1	Security update for tiff
SUSE-SU-2024:0997-1	Security update for krb5
SUSE-SU-2024:1014-1	Security update for avahi
SUSE-SU-2024:1103-1	Security update for qemu
SUSE-SU-2024:1129-1	Security update for expat
SUSE-SU-2024:1133-1	Security update for ncurses
SUSE-SU-2024:1136-1	Security update for c-ares
SUSE-SU-2024:1151-1	Security update for curl
SUSE-SU-2024:1167-1	Security update for nghttp2
SUSE-SU-2024:1172-1	Security update for util-linux
SUSE-SU-2024:1271-1	Security update for gnutls
SUSE-SU-2024:1438-1	Security update for qemu
SUSE-SU-2024:1981-1	Security update for iperf
TEMP-0000000-F7A20F	Kernel: Unpriviledged user can freeze journald

Cloudera Data Services on premises 1.5.4-CHF3

The cumulative hotfixes for new features, known issues, and fixed issues for 1.5.4-CHF3.

Note: ECS Customers: Direct upgrade path is not available for customers currently on Cloudera Data Services on premises 1.5.2. Customers must upgrade to Cloudera Data Services on premises 1.5.4 prior to consuming any CHFs built on top of 1.5.4.

Note:

OCP Customers: Direct upgrade path is available. Customers can directly upgrade from Cloudera Data Services on premises 1.5.2 to any 1.5.4 CHFs.

Whats new in Cloudera Data Services on premises 1.5.4-CHF3

New features introduced in this cumulative hotfix release of Cloudera Data Services on premises 1.5.4-CHF3.

Note: Cloudera Manager 7.11.3 CHF9.1 (version: 7.11.3.24) supports Cloudera Data Services on premises 1.5.4 CHF3 release.

Note: Cloudera Manager 7.11.3 CHF8 does not support any Cloudera Data Services on premises release.

Note: Base 7.1.7 SP3, 7.1.9 CHF6, and 7.1.9 SP1 supports Cloudera Data Services on premises 1.5.4 CHF3 release.

Note: Cloudera Data Services on premises 1.5.4 CHF3 is certified with RHEL 8.10 and RHEL 9.4 (RHCK kernel only).

Kubernetes Ingress NGINX Controller vulnerabilities

Kubernetes Ingress NGINX Controller vulnerabilities.

Five vulnerabilities affecting the Ingress Nginx Controller for Kubernetes were publicly disclosed on March 24, 2025, and were given the nickname "IngressNightmare."

The ‘IngressNightmare’ vulnerabilities may allow Remote Code Execution (RCE) and potentially expose Kubernetes clusters to malicious configuration modifications. Exploitation requires specially crafted HTTP requests that bypass security measures, such as a Web Application Firewall (WAF). Successful exploitation may lead to complete cluster compromise, data exfiltration, and denial of service.

Details of the CVEs:

- CVE-2025-1974 (CVSS score: 9.8) – An unauthenticated attacker with access to the pod network can achieve arbitrary code execution in the context of the ingress-nginx controller under certain conditions
- CVE-2025-24514 (CVSS score: 8.8) – The auth-url Ingress annotation can be used to inject configuration into NGINX, resulting in arbitrary code execution in the context of the ingress-nginx controller and disclosure of secrets accessible to the controller
- CVE-2025-1097 (CVSS score: 8.8) – The auth-tls-match-cn Ingress annotation can be used to inject configuration into NGINX, resulting in arbitrary code execution in the context of the ingress-nginx controller and disclosure of secrets accessible to the controller
- CVE-2025-1098 (CVSS score: 8.8) – The mirror-target and mirror-host Ingress annotations can be used to inject arbitrary configuration into NGINX, resulting in arbitrary code execution in the context of the ingress-nginx controller and disclosure of secrets accessible to the controller
- CVE-2025-24513 (CVSS score: 4.8) – An improper input validation vulnerability that could result in directory traversal within the container, leading to denial-of-service (DoS) or limited disclosure of secret objects from the cluster when combined with other vulnerabilities

Note: Cloudera will not provide remediation options for unsupported versions, and has not tested mitigations on unsupported versions. Customers are advised to upgrade to a supported product version. For more information, refer to the Support Lifecycle Policy page.

Action required

Multiple mitigation scenarios are possible depending on the version of Cloudera Data Services deployed.

1. Check your Data Services version.

   To determine which version of Data Services is running, log into Cloudera Manager UI, click “Data Services” in the left sidebar, and find each Data Services cluster's version(s).
2. Check your Cloudera Manager version.

   To determine which version of Cloudera Manager is running, log into Cloudera Manager UI and click on the version number at the lower-left corner of the home page. The Cloudera Manager 7.11.3 CHF8 Build number shows 56304673.

Mitigation steps for Cloudera Data Services On Premises:

1. No action is needed in the following scenarios:
   - Running Data Services on OpenShift Container Platform (OCP) and using the default HAproxy ingress controller. If you have customized OCP to use NGINX as the ingress controller, please contact RedHat support for further information
   - Running Data Services 1.5.4 or later with Cloudera Manager 7.11.3 CHF8 or later. These versions already disable the affected admission webhook component.
2. Data Services before 1.5.4, or Cloudera Manager prior to 7.11.3 CHF8, perform one of the following:
   - (Recommended) Upgrade Cloudera Manager to version 7.13.1 CHF1 or higher and all Embedded Container Service clusters to 1.5.4 SP2 or higher. This will not upgrade the nginx version but will automatically disable the webhook via Cloudera Manager, which mitigates the CVEs.
   - If an upgrade is not possible, follow the below steps to disable the webhook to mitigate the CVEs manually.

     Note: These steps should be run on each ECS cluster. If you create a new ECS cluster while on an affected release, reapply these steps on the new cluster(s).

1. SSH into the ECS server host.
2. Create the following YAML file to patch the built-in Nginx Ingress Controller Helm Chart:

```
apiVersion: helm.cattle.io/v1
kind: HelmChartConfig
metadata:
 name: rke2-ingress-nginx
 namespace: kube-system
spec:
 valuesContent: |-
  controller:
   admissionWebhooks:
    enabled: false
```

3. Run the following commands to apply the patch:

```
alias kubectl=/var/lib/rancher/rke2/bin/kubectl
export KUBECONFIG=/etc/rancher/rke2/rke2.yaml
kubectl apply -f <PATCH_FILENAME>.yaml
```

Note: No cluster restart or ECS refresh is required.

Note: Cloudera will not provide remediation options for unsupported versions, and has not tested mitigations on unsupported versions. Customers are advised to upgrade to a supported product version. For more information, refer to the Support Lifecycle Policy page

Cloudera Data Services on premises version 1.5.4 SP2 CHF1 will contain fixes to upgrade the Ingress NGINX controller. Target release date: April 25. 2025.

For the latest update on this issue, see the corresponding Knowledge article: TSB 2025-839: Critical Kubernetes Ingress NGINX Controller Vulnerability Allows RCE Without Authentication

For the latest update on this issue, see the corresponding Data Services Knowledge article: https://my.cloudera.com/knowledge/TSB-2025-839-Mitigation-steps-for-Cloudera-Data-Services-on?id=405006

Known Issues in Cloudera Data Services on premises 1.5.4-CHF3

New known issues in the 1.5.4 cumulative hotfix CHF3 release of Cloudera Data Services on premises.

OPSX-5529 - Cloudera Data Services on premises ECS longhorn upgrade failure

The helm-install-longhorn pod enters a crash loop state during ECS upgrade.

Provide the Longhorn diagnostic bundle to facilitate issue identification.

Note: Starting from 1.5.4 CHF3, Longhorn upgrade in ECS has a failure policy set to "abort" to prevent unexpected uninstallation triggers during retries.

Resuming Longhorn Upgrade: After resolving the underlying longhorn upgrade failure issues, follow these steps to resume the upgrade:

```
# Get the history of longhorn helm chart so that we can identify
 the chart for which installation is failing. # helm history lon
ghorn -n longhorn-system
REVISION    UPDATED                     STATUS        CHART
       APP VERSION    DESCRIPTION
1           Thu Sep 26 21:31:05 2024    superseded    longhorn-1
.5.4    v1.5.4         Install complete
2           Fri Sep 27 05:17:44 2024    failed        longhorn-1
.6.2    v1.6.2         Upgrade "longhorn" failed: post-upgrade h
ooks failed: 1 error occurr...# Get the log of the failing helm-
install-longhorn job in the longhorn namespace

The job log should indicate that the due to failure policy being
 "abort", it is waiting for manaul intervention:
"Release status is 'failed' and failure policy is 'abort', not
'reinstall'; waiting for operator intervention"# We want to roll
back

# Find all jobs in longhorn-system and delete those. These jobs
will be re-triggered as part of the manual patch.
# kubectl get jobs -n longhorn-system
NAME                    COMPLETIONS   DURATION   AGE
helm-install-longhorn   0/1           9h         9h
longhorn-post-upgrade   1/1           11m        10h

# Delete all the longhorn jobs if any

# kubectl delete job helm-install-longhorn longhorn-post-upgrade
 -n longhorn-system# Rollback longhorn to the version prior to t
he upgrade
# In this case, revision 1 marks the step of a successful longhor
n 1.5.4 install
# helm rollback longhorn -n longhorn-system <revision number># R
esume longhorn upgrade from Cloudera Manager UI
```

```
Running commands > All recent commands > find the failed upgrade
 command and click on resume
```

OPSX-5573/OPSAPS-69892 - Intermittent kube-proxy failures - 1.5.4 CHF3

After rebooting/restarting an ECS agent node, the kube-proxy Linux process may not start due to a race condition in the kubelet. When this happens, ECS cluster networking and other services – such as Vault, DNS, authentication, Longhorn storage, etc. – are affected. At the Kubernetes pod level, errors such as "connection refused", "connection timed out" and "i/o timeout" may be observed. If you suspect possible networking issues in your ECS cluster, checking kube-proxy is a good first step.

To fix this issue, perform the following steps on all of the affected nodes:

1. To identify which agent needs to be restarted, check the status of each kube-proxy pod to make sure it is in the "ready" state by running the following command on each host in the cluster.

   ```
   kubectl describe pod [***POD-NAME***] -n kube-system
   ```

   Here, *[***POD-NAME***]* should have a format such as: kube-proxy-<hostname>.

   In the Conditions section of the describe pod output, confirm that the "ready" condition is "True".

   ```
   Conditions:
     Type              Status
     Initialized       True
     Ready             True
     ContainersReady   True
     PodScheduled      True
   ```

   Another option is to run the following command:

   ```
   kubectl get pods -n kube-system -l component=kube-proxy -o go-
   template='{{range .items}}
   {{.metadata.name}}{{"\n"}}{{"  "}}{{range .status.conditions}}
   {{ if eq .type "Ready" }}
   Ready:{{.status}}{{"\n\n"}}{{end}}{{end}}{{end}}'
   ```

   The sample output displays the status of all of the kube-proxy pods in the cluster:

   ```
   kube-proxy-host-1.cloudera.com
     Ready:True

   kube-proxy-host-2.cloudera.com
     Ready:True

   kube-proxy-host-3.cloudera.com
     Ready:True
   ```

2. If the "ready" state is False, kube-proxy is not functioning properly, regardless of whether the kube-proxy process is running on that host or not. On each of the affected nodes, run the following command to delete the kube-proxy pod manifest:

   ```
    rm /var/lib/rancher/rke2/agent/pod-manifests/kube-proxy.yaml
   ```

3. Start the agent role.

   After the agent role is started, you may not immediately see the kube-proxy process running, but a new kube-proxy process should start shortly. Check the pod status to make sure it is ready.

After all of the problem agents have been restarted, the cluster may complain that the vault is sealed – if so, unseal it. At this point, the Control Plane should be functioning properly.

If current version is 1.5.2 perform above steps. If it is 1.5.3 or above, perform step 1 from the above procedure to identify the problematic nodes. Perform stop and start of ECS roles on the hosts where the problem exists.

Additional details about this issue are available here: https://www.suse.com/support/kb/doc/?id=000021284

COMPX-18031 - [ECS] [154CHF1-CHF3] Any new pods are stuck in pending state post ecs upgrade

When pods are stuck in pending state post ECS upgrade, you will not be able to schedule new workload (warehouse, virtual cluster etc) . Pending pods are normal if resources or quotas are exhausted. After high cluster loads, pods will be left in a pending state even if enough resources are available. The pods are not evaluated as part of normal scheduling.

The issue will be fixed after restarting the YuniKorn scheduler pod in yunikorn namespace. YuniKorn scheduler will trigger a re-evaluation of all pods and schedule the pending pods.

OBS-4176 - Prometheus reports duplicate metric alert for Kubernetes state metrics

After installing Cloudera Data Services on premises 1.5.4, the EnvPrometheusDuplicateTimestamps warning message appears on the Control Plane Monitoring dashboard.

Perform the following steps:

1. On the Management Console home page, select Administration > Alerts .
2. On the Alerts page, search for the EnvPrometheusDuplicateTimestamps rule alert and from the drop-down menu select Disable Alert Rule to disable the alert.

Repository Locations for 1.5.4-CHF3

The URLs for Cloudera Data Services on premises 1.5.4-CHF3 are listed in the following table:

URL Type	Repository Location
Index	`https://username:password@archive.cloudera.com/p/cdp-pvc-ds/1.5.4-h4/`
Manifest	Repository: `https://username:password@archive.cloudera.com/p/cdp-pvc-ds/1.5.4-h4/manifest.json`
Parcels	Repository: `https://username:password@archive.cloudera.com/p/cdp-pvc-ds/1.5.4-h4/parcels/`

Fixed Common Vulnerabilities and Exposures in 1.5.4 CHF3

Review the Common vulnerabilities and Exposures (CVEs) that were fixed in 1.5.4 CHF3 release of Cloudera Data Services on premises.

Issue ID	Description
CVE-2013-0340	expat 2.1.0 and earlier does not properly handle entities expansion unless an application developer uses the XML_S remote attackers to cause a denial of service (resource consumption), send HTTP requests to intranet servers, or rea aka an XML External Entity (XXE) issue. NOTE: it could be argued that because expat already provides the ability responsibility for resolving this issue lies with application developers; according to this argument, this entry should would need its own CVE.

CVE-2014-9471	The parse_datetime function in GNU coreutils allows remote attackers to cause a denial of service (crash) or possib string, as demonstrated by the "--date=TZ="123"345" @1" string to the touch or date command.
CVE-2015-4041	The keycompare_mb function in sort.c in sort in GNU Coreutils through 8.23 on 64-bit platforms performs a size c bytes occupied by multibyte characters, which allows attackers to cause a denial of service (heap-based buffer over unspecified other impact via long UTF-8 strings.
CVE-2015-4042	Integer overflow in the keycompare_mb function in sort.c in sort in GNU Coreutils through 8.23 might allow attacl crash) or possibly have unspecified other impact via long strings.
CVE-2017-18018	In GNU Coreutils through 8.29, chown-core.c in chown and chgrp does not prevent replacement of a plain file with options, which allows local users to modify the ownership of arbitrary files by leveraging a race condition.
CVE-2018-13410	Info-ZIP Zip 3.0, when the -T and -TT command-line options are used, allows attackers to cause a denial of service possibly have unspecified other impact because of an off-by-one error. NOTE: it is unclear whether there are realis controls the -TT value, given that the entire purpose of -TT is execution of arbitrary commands
CVE-2019-5068	An exploitable shared memory permissions vulnerability exists in the functionality of X11 Mesa 3D Graphics Libr memory without any specific permissions to trigger this vulnerability.
CVE-2019-9704	Vixie Cron before the 3.0pl1-133 Debian package allows local users to cause a denial of service (daemon crash) vi value is not checked.
CVE-2019-9705	Vixie Cron before the 3.0pl1-133 Debian package allows local users to cause a denial of service (memory consump unlimited number of lines is accepted.
CVE-2020-25659	python-cryptography 3.2 is vulnerable to Bleichenbacher timing attacks in the RSA decryption API, via timed proc
CVE-2020-8201	Node.js < 12.18.4 and < 14.11 can be exploited to perform HTTP desync attacks and deliver malicious payloads to be crafted by an attacker to hijack user sessions, poison cookies, perform clickjacking, and a multitude of other atta underlying system. The attack was possible due to a bug in processing of carrier-return symbols in the HTTP heade
CVE-2021-20312	A flaw was found in ImageMagick in versions 7.0.11, where an integer overflow in WriteTHUMBNAILImage of c behavior via a crafted image file that is submitted by an attacker and processed by an application using ImageMagi is to system availability.
CVE-2021-20313	A flaw was found in ImageMagick in versions before 7.0.11. A potential cipher leak when the calculate signatures highest threat from this vulnerability is to data confidentiality.
CVE-2022-1125	Use after free in Portals in Google Chrome prior to 100.0.4896.60 allowed a remote attacker who convinced a user potentially exploit heap corruption via user interaction.
CVE-2022-1127	Use after free in QR Code Generator in Google Chrome prior to 100.0.4896.60 allowed a remote attacker who conv interaction to potentially exploit heap corruption via user interaction.
CVE-2022-1131	Use after free in Cast UI in Google Chrome prior to 100.0.4896.60 allowed a remote attacker to potentially exploit
CVE-2022-1133	Use after free in WebRTC Perf in Google Chrome prior to 100.0.4896.60 allowed a remote attacker to potentially e page.
CVE-2022-1134	Type confusion in V8 in Google Chrome prior to 100.0.4896.60 allowed a remote attacker to potentially exploit he
CVE-2022-1135	Use after free in Shopping Cart in Google Chrome prior to 100.0.4896.60 allowed a remote attacker to potentially e user interaction.
CVE-2022-1136	Use after free in Tab Strip in Google Chrome prior to 100.0.4896.60 allowed an attacker who convinced a user to i exploit heap corruption via specific set of user gestures.
CVE-2022-1137	Inappropriate implementation in Extensions in Google Chrome prior to 100.0.4896.60 allowed an attacker who cor to leak potentially sensitive information via a crafted HTML page.
CVE-2022-1138	Inappropriate implementation in Web Cursor in Google Chrome prior to 100.0.4896.60 allowed a remote attacker v obscure the contents of the Omnibox (URL bar) via a crafted HTML page.
CVE-2022-1139	Inappropriate implementation in Background Fetch API in Google Chrome prior to 100.0.4896.60 allowed a remot crafted HTML page.
CVE-2022-1141	Use after free in File Manager in Google Chrome prior to 100.0.4896.60 allowed a remote attacker who convinced potentially exploit heap corruption via specific user gesture.
CVE-2022-1142	Heap buffer overflow in WebUI in Google Chrome prior to 100.0.4896.60 allowed a remote attacker who convince to potentially exploit heap corruption via specific input into DevTools.
CVE-2022-1143	Heap buffer overflow in WebUI in Google Chrome prior to 100.0.4896.60 allowed a remote attacker who convince to potentially exploit heap corruption via specific input into DevTools.

CVE-2022-1144	Use after free in WebUI in Google Chrome prior to 100.0.4896.60 allowed a remote attacker who convinced a user potentially exploit heap corruption via specific input into DevTools.
CVE-2022-1145	Use after free in Extensions in Google Chrome prior to 100.0.4896.60 allowed an attacker who convinced a user to exploit heap corruption via specific user interaction and profile destruction.
CVE-2022-1146	Inappropriate implementation in Resource Timing in Google Chrome prior to 100.0.4896.60 allowed a remote attac HTML page.
CVE-2022-1232	Type confusion in V8 in Google Chrome prior to 100.0.4896.75 allowed a remote attacker to potentially exploit he
CVE-2022-1305	Use after free in storage in Google Chrome prior to 100.0.4896.88 allowed a remote attacker to potentially exploit
CVE-2022-1306	Inappropriate implementation in compositing in Google Chrome prior to 100.0.4896.88 allowed a remote attacker t bar) via a crafted HTML page.
CVE-2022-1308	Use after free in BFCache in Google Chrome prior to 100.0.4896.88 allowed a remote attacker to potentially explo
CVE-2022-1309	Insufficient policy enforcement in developer tools in Google Chrome prior to 100.0.4896.88 allowed a remote attac via a crafted HTML page.
CVE-2022-1310	Use after free in regular expressions in Google Chrome prior to 100.0.4896.88 allowed a remote attacker to potenti HTML page.
CVE-2022-1312	Use after free in storage in Google Chrome prior to 100.0.4896.88 allowed an attacker who convinced a user to ins perform a sandbox escape via a crafted Chrome Extension.
CVE-2022-1313	Use after free in tab groups in Google Chrome prior to 100.0.4896.88 allowed a remote attacker to potentially expl
CVE-2022-1314	Type confusion in V8 in Google Chrome prior to 100.0.4896.88 allowed a remote attacker to potentially exploit he
CVE-2022-1364	Type confusion in V8 Turbofan in Google Chrome prior to 100.0.4896.127 allowed a remote attacker to potentially page.
CVE-2022-1477	Use after free in Vulkan in Google Chrome prior to 101.0.4951.41 allowed a remote attacker to potentially exploit
CVE-2022-1478	Use after free in SwiftShader in Google Chrome prior to 101.0.4951.41 allowed a remote attacker to potentially ex page.
CVE-2022-1479	Use after free in ANGLE in Google Chrome prior to 101.0.4951.41 allowed a remote attacker to potentially exploit
CVE-2022-1482	Inappropriate implementation in WebGL in Google Chrome prior to 101.0.4951.41 allowed a remote attacker to po HTML page.
CVE-2022-1483	Heap buffer overflow in WebGPU in Google Chrome prior to 101.0.4951.41 allowed a remote attacker who had co potentially exploit heap corruption via a crafted HTML page.
CVE-2022-1484	Heap buffer overflow in Web UI Settings in Google Chrome prior to 101.0.4951.41 allowed a remote attacker to p HTML page.
CVE-2022-1485	Use after free in File System API in Google Chrome prior to 101.0.4951.41 allowed a remote attacker to potentiall page.
CVE-2022-1486	Type confusion in V8 in Google Chrome prior to 101.0.4951.41 allowed a remote attacker to obtain potentially sen crafted HTML page.
CVE-2022-1487	Use after free in Ozone in Google Chrome prior to 101.0.4951.41 allowed a remote attacker to potentially exploit h
CVE-2022-1488	Inappropriate implementation in Extensions API in Google Chrome prior to 101.0.4951.41 allowed an attacker wh extension to leak cross-origin data via a crafted Chrome Extension.
CVE-2022-1490	Use after free in Browser Switcher in Google Chrome prior to 101.0.4951.41 allowed a remote attacker who convi interaction to potentially exploit heap corruption via a crafted HTML page.
CVE-2022-1491	Use after free in Bookmarks in Google Chrome prior to 101.0.4951.41 allowed a remote attacker to potentially exp user interaction.
CVE-2022-1492	Insufficient data validation in Blink Editing in Google Chrome prior to 101.0.4951.41 allowed a remote attacker to HTML page.
CVE-2022-1493	Use after free in Dev Tools in Google Chrome prior to 101.0.4951.41 allowed a remote attacker to potentially expl interaction.
CVE-2022-1494	Insufficient data validation in Trusted Types in Google Chrome prior to 101.0.4951.41 allowed a remote attacker t HTML page.

CVE-2022-1496	Use after free in File Manager in Google Chrome prior to 101.0.4951.41 allowed a remote attacker to potentially ex user interaction.
CVE-2022-1497	Inappropriate implementation in Input in Google Chrome prior to 101.0.4951.41 allowed a remote attacker to spoo crafted HTML page.
CVE-2022-1498	Inappropriate implementation in HTML Parser in Google Chrome prior to 101.0.4951.41 allowed a remote attacke page.
CVE-2022-1499	Inappropriate implementation in WebAuthentication in Google Chrome prior to 101.0.4951.41 allowed a remote at crafted HTML page.
CVE-2022-1500	Insufficient data validation in Dev Tools in Google Chrome prior to 101.0.4951.41 allowed a remote attacker to by HTML page.
CVE-2022-1501	Inappropriate implementation in iframe in Google Chrome prior to 101.0.4951.41 allowed a remote attacker to leak
CVE-2022-1638	Heap buffer overflow in V8 Internationalization in Google Chrome prior to 101.0.4951.64 allowed a remote attack crafted HTML page.
CVE-2022-1639	Use after free in ANGLE in Google Chrome prior to 101.0.4951.64 allowed a remote attacker to potentially exploit
CVE-2022-1640	Use after free in Sharing in Google Chrome prior to 101.0.4951.64 allowed a remote attacker who convinced a use potentially exploit heap corruption via a crafted HTML page.
CVE-2022-1853	Use after free in Indexed DB in Google Chrome prior to 102.0.5005.61 allowed a remote attacker to potentially per page.
CVE-2022-1854	Use after free in ANGLE in Google Chrome prior to 102.0.5005.61 allowed a remote attacker to potentially exploit
CVE-2022-1855	Use after free in Messaging in Google Chrome prior to 102.0.5005.61 allowed a remote attacker to potentially expl
CVE-2022-1856	Use after free in User Education in Google Chrome prior to 102.0.5005.61 allowed an attacker who convinced a us potentially exploit heap corruption via a crafted Chrome Extension or specific user interaction.
CVE-2022-1857	Insufficient policy enforcement in File System API in Google Chrome prior to 102.0.5005.61 allowed a remote atta crafted HTML page.
CVE-2022-1858	Out of bounds read in DevTools in Google Chrome prior to 102.0.5005.61 allowed a remote attacker to perform an interaction.
CVE-2022-1859	Use after free in Performance Manager in Google Chrome prior to 102.0.5005.61 allowed a remote attacker who co interaction to potentially exploit heap corruption via a crafted HTML page.
CVE-2022-1860	Use after free in UI Foundations in Google Chrome on Chrome OS prior to 102.0.5005.61 allowed a remote attack user interaction to potentially exploit heap corruption via specific user interactions.
CVE-2022-1861	Use after free in Sharing in Google Chrome on Chrome OS prior to 102.0.5005.61 allowed a remote attacker who c interactions to potentially exploit heap corruption via specific user interaction.
CVE-2022-1862	Inappropriate implementation in Extensions in Google Chrome prior to 102.0.5005.61 allowed an attacker who con to bypass profile restrictions via a crafted HTML page.
CVE-2022-1863	Use after free in Tab Groups in Google Chrome prior to 102.0.5005.61 allowed an attacker who convinced a user t exploit heap corruption via a crafted Chrome Extension and specific user interaction.
CVE-2022-1864	Use after free in WebApp Installs in Google Chrome prior to 102.0.5005.61 allowed an attacker who convinced a u potentially exploit heap corruption via a crafted Chrome Extension and specific user interaction.
CVE-2022-1865	Use after free in Bookmarks in Google Chrome prior to 102.0.5005.61 allowed an attacker who convinced a user to exploit heap corruption via a crafted Chrome Extension and specific user interaction.
CVE-2022-1866	Use after free in Tablet Mode in Google Chrome on Chrome OS prior to 102.0.5005.61 allowed a remote attacker v user interactions to potentially exploit heap corruption via specific user interactions.
CVE-2022-1867	Insufficient validation of untrusted input in Data Transfer in Google Chrome prior to 102.0.5005.61 allowed a remo crafted clipboard content.
CVE-2022-1868	Inappropriate implementation in Extensions API in Google Chrome prior to 102.0.5005.61 allowed an attacker who extension to bypass navigation restrictions via a crafted HTML page.
CVE-2022-1869	Type Confusion in V8 in Google Chrome prior to 102.0.5005.61 allowed a remote attacker to potentially exploit he
CVE-2022-1870	Use after free in App Service in Google Chrome prior to 102.0.5005.61 allowed an attacker who convinced a user t exploit heap corruption via a crafted Chrome Extension.

CVE-2022-1871	Insufficient policy enforcement in File System API in Google Chrome prior to 102.0.5005.61 allowed an attacker v extension to bypass file system policy via a crafted HTML page.
CVE-2022-1872	Insufficient policy enforcement in Extensions API in Google Chrome prior to 102.0.5005.61 allowed an attacker w extension to bypass downloads policy via a crafted HTML page.
CVE-2022-1873	Insufficient policy enforcement in COOP in Google Chrome prior to 102.0.5005.61 allowed a remote attacker to le page.
CVE-2022-1875	Inappropriate implementation in PDF in Google Chrome prior to 102.0.5005.61 allowed a remote attacker to leak c
CVE-2022-1876	Heap buffer overflow in DevTools in Google Chrome prior to 102.0.5005.61 allowed an attacker who convinced a potentially exploit heap corruption via a crafted HTML page.
CVE-2022-1919	Use after free in Codecs in Google Chrome prior to 101.0.4951.41 allowed a remote attacker to potentially exploit
CVE-2022-2007	Use after free in WebGPU in Google Chrome prior to 102.0.5005.115 allowed a remote attacker to potentially expl
CVE-2022-2008	Double free in WebGL in Google Chrome prior to 102.0.5005.115 allowed a remote attacker to potentially exploit
CVE-2022-2010	Out of bounds read in compositing in Google Chrome prior to 102.0.5005.115 allowed a remote attacker who had c potentially perform a sandbox escape via a crafted HTML page.
CVE-2022-2011	Use after free in ANGLE in Google Chrome prior to 102.0.5005.115 allowed a remote attacker to potentially explo
CVE-2022-2156	Use after free in Core in Google Chrome prior to 103.0.5060.53 allowed a remote attacker to potentially exploit he
CVE-2022-2157	Use after free in Interest groups in Google Chrome prior to 103.0.5060.53 allowed a remote attacker who had comp exploit heap corruption via a crafted HTML page.
CVE-2022-2158	Type confusion in V8 in Google Chrome prior to 103.0.5060.53 allowed a remote attacker to potentially exploit he
CVE-2022-2161	Use after free in WebApp Provider in Google Chrome prior to 103.0.5060.53 allowed a remote attacker who convi interactions to potentially exploit heap corruption via specific UI interactions.
CVE-2022-2163	Use after free in Cast UI and Toolbar in Google Chrome prior to 103.0.5060.134 allowed an attacker who convince potentially exploit heap corruption via UI interaction.
CVE-2022-2164	Inappropriate implementation in Extensions API in Google Chrome prior to 103.0.5060.53 allowed an attacker wh extension to bypass discretionary access control via a crafted HTML page.
CVE-2022-2165	Insufficient data validation in URL formatting in Google Chrome prior to 103.0.5060.53 allowed a remote attacker homographs via a crafted domain name.
CVE-2022-2294	Heap buffer overflow in WebRTC in Google Chrome prior to 103.0.5060.114 allowed a remote attacker to potentia HTML page.
CVE-2022-2295	Type confusion in V8 in Google Chrome prior to 103.0.5060.114 allowed a remote attacker to potentially exploit h
CVE-2022-2399	Use after free in WebGPU in Google Chrome prior to 100.0.4896.88 allowed a remote attacker to potentially explo
CVE-2022-2415	Heap buffer overflow in WebGL in Google Chrome prior to 103.0.5060.53 allowed a remote attacker to potentially page.
CVE-2022-2477	Use after free in Guest View in Google Chrome prior to 103.0.5060.134 allowed an attacker who convinced a user exploit heap corruption via a crafted HTML page.
CVE-2022-2478	Use after free in PDF in Google Chrome prior to 103.0.5060.134 allowed a remote attacker to potentially exploit h
CVE-2022-2480	Use after free in Service Worker API in Google Chrome prior to 103.0.5060.134 allowed a remote attacker to pote HTML page.
CVE-2022-2481	Use after free in Views in Google Chrome prior to 103.0.5060.134 allowed a remote attacker who convinced a use potentially exploit heap corruption via UI interaction.
CVE-2022-2603	Use after free in Omnibox in Google Chrome prior to 104.0.5112.79 allowed a remote attacker to potentially explo
CVE-2022-2604	Use after free in Safe Browsing in Google Chrome prior to 104.0.5112.79 allowed a remote attacker to potentially page.
CVE-2022-2605	Out of bounds read in Dawn in Google Chrome prior to 104.0.5112.79 allowed a remote attacker to potentially exp page.
CVE-2022-2606	Use after free in Managed devices API in Google Chrome prior to 104.0.5112.79 allowed a remote attacker who co policy to potentially exploit heap corruption via a crafted HTML page.

CVE-2022-2610	Insufficient policy enforcement in Background Fetch in Google Chrome prior to 104.0.5112.79 allowed a remote a HTML page.
CVE-2022-2612	Side-channel information leakage in Keyboard input in Google Chrome prior to 104.0.5112.79 allowed a remote at process to obtain potentially sensitive information from process memory via a crafted HTML page.
CVE-2022-2614	Use after free in Sign-In Flow in Google Chrome prior to 104.0.5112.79 allowed a remote attacker to potentially ex page.
CVE-2022-2615	Insufficient policy enforcement in Cookies in Google Chrome prior to 104.0.5112.79 allowed a remote attacker to page.
CVE-2022-2616	Inappropriate implementation in Extensions API in Google Chrome prior to 104.0.5112.79 allowed an attacker wh extension to spoof the contents of the Omnibox (URL bar) via a crafted Chrome Extension.
CVE-2022-2617	Use after free in Extensions API in Google Chrome prior to 104.0.5112.79 allowed an attacker who convinced a us potentially exploit heap corruption via specific UI interactions.
CVE-2022-2618	Insufficient validation of untrusted input in Internals in Google Chrome prior to 104.0.5112.79 allowed a remote at malicious file .
CVE-2022-2619	Insufficient validation of untrusted input in Settings in Google Chrome prior to 104.0.5112.79 allowed an attacker extension to inject scripts or HTML into a privileged page via a crafted HTML page.
CVE-2022-2621	Use after free in Extensions in Google Chrome prior to 104.0.5112.79 allowed an attacker who convinced a user to exploit heap corruption via specific UI interactions.
CVE-2022-2624	Heap buffer overflow in PDF in Google Chrome prior to 104.0.5112.79 allowed a remote attacker who convinced a to potentially exploit heap corruption via a crafted PDF file.
CVE-2022-2743	Integer overflow in Window Manager in Google Chrome on Chrome OS and Lacros prior to 104.0.5112.79 allowe engage in specific UI interactions to perform an out of bounds memory write via crafted UI interactions. (Chrome s
CVE-2022-2852	Use after free in FedCM in Google Chrome prior to 104.0.5112.101 allowed a remote attacker to potentially exploi
CVE-2022-2854	Use after free in SwiftShader in Google Chrome prior to 104.0.5112.101 allowed a remote attacker to potentially ex page.
CVE-2022-2855	Use after free in ANGLE in Google Chrome prior to 104.0.5112.101 allowed a remote attacker to potentially explo
CVE-2022-2857	Use after free in Blink in Google Chrome prior to 104.0.5112.101 allowed a remote attacker to potentially exploit h
CVE-2022-2858	Use after free in Sign-In Flow in Google Chrome prior to 104.0.5112.101 allowed a remote attacker to potentially e interaction.
CVE-2022-2859	Use after free in Chrome OS Shell in Google Chrome prior to 104.0.5112.101 allowed a remote attacker who conv interactions to potentially exploit heap corruption via specific UI interactions.
CVE-2022-2860	Insufficient policy enforcement in Cookies in Google Chrome prior to 104.0.5112.101 allowed a remote attacker to crafted HTML page.
CVE-2022-2861	Inappropriate implementation in Extensions API in Google Chrome prior to 104.0.5112.101 allowed an attacker wh extension to inject arbitrary scripts into WebUI via a crafted HTML page.
CVE-2022-2998	Use after free in Browser Creation in Google Chrome prior to 104.0.5112.101 allowed a remote attacker who had c interaction to potentially exploit heap corruption via a crafted HTML page.
CVE-2022-3038	Use after free in Network Service in Google Chrome prior to 105.0.5195.52 allowed a remote attacker to potentiall page.
CVE-2022-3039	Use after free in WebSQL in Google Chrome prior to 105.0.5195.52 allowed a remote attacker to potentially explo
CVE-2022-3040	Use after free in Layout in Google Chrome prior to 105.0.5195.52 allowed a remote attacker to potentially exploit h
CVE-2022-3041	Use after free in WebSQL in Google Chrome prior to 105.0.5195.52 allowed a remote attacker to potentially explo
CVE-2022-3044	Inappropriate implementation in Site Isolation in Google Chrome prior to 105.0.5195.52 allowed a remote attacker to bypass site isolation via a crafted HTML page.
CVE-2022-3045	Insufficient validation of untrusted input in V8 in Google Chrome prior to 105.0.5195.52 allowed a remote attacker crafted HTML page.
CVE-2022-3046	Use after free in Browser Tag in Google Chrome prior to 105.0.5195.52 allowed an attacker who convinced a user exploit heap corruption via a crafted HTML page.
CVE-2022-3047	Insufficient policy enforcement in Extensions API in Google Chrome prior to 105.0.5195.52 allowed an attacker w extension to bypass downloads policy via a crafted HTML page.

CVE-2022-3054	Insufficient policy enforcement in DevTools in Google Chrome prior to 105.0.5195.52 allowed a remote attacker t crafted HTML page.
CVE-2022-3055	Use after free in Passwords in Google Chrome prior to 105.0.5195.52 allowed a remote attacker who convinced a u potentially exploit heap corruption via a crafted HTML page.
CVE-2022-3056	Insufficient policy enforcement in Content Security Policy in Google Chrome prior to 105.0.5195.52 allowed a rem via a crafted HTML page.
CVE-2022-3057	Inappropriate implementation in iframe Sandbox in Google Chrome prior to 105.0.5195.52 allowed a remote attack HTML page.
CVE-2022-3058	Use after free in Sign-In Flow in Google Chrome prior to 105.0.5195.52 allowed a remote attacker who convinced potentially exploit heap corruption via crafted UI interaction.
CVE-2022-3075	Insufficient data validation in Mojo in Google Chrome prior to 105.0.5195.102 allowed a remote attacker who had potentially perform a sandbox escape via a crafted HTML page.
CVE-2022-3195	Out of bounds write in Storage in Google Chrome prior to 105.0.5195.125 allowed a remote attacker to perform an HTML page. (Chromium security severity: High)
CVE-2022-3196	Use after free in PDF in Google Chrome prior to 105.0.5195.125 allowed a remote attacker to potentially exploit h (Chromium security severity: High)
CVE-2022-3197	Use after free in PDF in Google Chrome prior to 105.0.5195.125 allowed a remote attacker to potentially exploit h (Chromium security severity: High)
CVE-2022-3198	Use after free in PDF in Google Chrome prior to 105.0.5195.125 allowed a remote attacker to potentially exploit h (Chromium security severity: High)
CVE-2022-3199	Use after free in Frames in Google Chrome prior to 105.0.5195.125 allowed a remote attacker to potentially exploi (Chromium security severity: High)
CVE-2022-3200	Heap buffer overflow in Internals in Google Chrome prior to 105.0.5195.125 allowed a remote attacker to potentia HTML page. (Chromium security severity: High)
CVE-2022-3304	Use after free in CSS in Google Chrome prior to 106.0.5249.62 allowed a remote attacker to potentially exploit hea (Chromium security severity: High)
CVE-2022-3307	Use after free in media in Google Chrome prior to 106.0.5249.62 allowed a remote attacker to potentially exploit h (Chromium security severity: High)
CVE-2022-3308	Insufficient policy enforcement in developer tools in Google Chrome prior to 106.0.5249.62 allowed a remote attac via a crafted HTML page. (Chromium security severity: Medium)
CVE-2022-3311	Use after free in import in Google Chrome prior to 106.0.5249.62 allowed a remote attacker who had compromised sandbox escape via a crafted HTML page. (Chromium security severity: Medium)
CVE-2022-3312	Insufficient validation of untrusted input in VPN in Google Chrome on ChromeOS prior to 106.0.5249.62 allowed restrictions via physical access to the device. (Chromium security severity: Medium)
CVE-2022-3313	Incorrect security UI in full screen in Google Chrome prior to 106.0.5249.62 allowed a remote attacker to spoof se (Chromium security severity: Medium)
CVE-2022-3314	Use after free in logging in Google Chrome prior to 106.0.5249.62 allowed a remote attacker who had compromise sandbox escape via a crafted HTML page. (Chromium security severity: Medium)
CVE-2022-3315	Type confusion in Blink in Google Chrome prior to 106.0.5249.62 allowed a remote attacker to potentially exploit (Chromium security severity: Low)
CVE-2022-3316	Insufficient validation of untrusted input in Safe Browsing in Google Chrome prior to 106.0.5249.62 allowed a rem crafted HTML page. (Chromium security severity: Low)
CVE-2022-3370	Use after free in Custom Elements in Google Chrome prior to 106.0.5249.91 allowed a remote attacker to potential HTML page. (Chromium security severity: High)
CVE-2022-3373	Out of bounds write in V8 in Google Chrome prior to 106.0.5249.91 allowed a remote attacker to perform an out o page. (Chromium security severity: High)
CVE-2022-3443	Insufficient data validation in File System API in Google Chrome prior to 106.0.5249.62 allowed a remote attacker crafted HTML page. (Chromium security severity: Low)
CVE-2022-3444	Insufficient data validation in File System API in Google Chrome prior to 106.0.5249.62 allowed a remote attacker crafted HTML page and malicious file. (Chromium security severity: Low)

CVE-2022-3445	Use after free in Skia in Google Chrome prior to 106.0.5249.119 allowed a remote attacker to potentially exploit he (Chromium security severity: High)
CVE-2022-3446	Heap buffer overflow in WebSQL in Google Chrome prior to 106.0.5249.119 allowed a remote attacker to potentia HTML page. (Chromium security severity: High)
CVE-2022-3448	Use after free in Permissions API in Google Chrome prior to 106.0.5249.119 allowed a remote attacker who convi to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2022-3449	Use after free in Safe Browsing in Google Chrome prior to 106.0.5249.119 allowed an attacker who convinced a us potentially exploit heap corruption via a crafted Chrome Extension. (Chromium security severity: High)
CVE-2022-3450	Use after free in Peer Connection in Google Chrome prior to 106.0.5249.119 allowed a remote attacker to potential HTML page. (Chromium security severity: High)
CVE-2022-3652	Type confusion in V8 in Google Chrome prior to 107.0.5304.62 allowed a remote attacker to potentially exploit he (Chromium security severity: High)
CVE-2022-3653	Heap buffer overflow in Vulkan in Google Chrome prior to 107.0.5304.62 allowed a remote attacker to potentially page. (Chromium security severity: High)
CVE-2022-3654	Use after free in Layout in Google Chrome prior to 107.0.5304.62 allowed a remote attacker to potentially exploit (Chromium security severity: High)
CVE-2022-3655	Heap buffer overflow in Media Galleries in Google Chrome prior to 107.0.5304.62 allowed an attacker who convi potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: Medium)
CVE-2022-3656	Insufficient data validation in File System in Google Chrome prior to 107.0.5304.62 allowed a remote attacker to b HTML page. (Chromium security severity: Medium)
CVE-2022-3657	Use after free in Extensions in Google Chrome prior to 107.0.5304.62 allowed an attacker who convinced a user to exploit heap corruption via a crafted Chrome Extension. (Chromium security severity: Medium)
CVE-2022-3661	Insufficient data validation in Extensions in Google Chrome prior to 107.0.5304.62 allowed a remote attacker who leak cross-origin data via a crafted Chrome extension. (Chromium security severity: Low)
CVE-2022-3723	Type confusion in V8 in Google Chrome prior to 107.0.5304.87 allowed a remote attacker to potentially exploit he (Chromium security severity: High)
CVE-2022-3842	Use after free in Passwords in Google Chrome prior to 105.0.5195.125 allowed a remote attacker who had compro exploit heap corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2022-3863	Use after free in Browser History in Google Chrome prior to 100.0.4896.75 allowed a remote attacker to potentiall page. (Chrome security severity: High)
CVE-2022-3885	Use after free in V8 in Google Chrome prior to 107.0.5304.106 allowed a remote attacker to potentially exploit hea (Chromium security severity: High)
CVE-2022-3886	Use after free in Speech Recognition in Google Chrome prior to 107.0.5304.106 allowed a remote attacker to poten HTML page. (Chromium security severity: High)
CVE-2022-3887	Use after free in Web Workers in Google Chrome prior to 107.0.5304.106 allowed a remote attacker to potentially page. (Chromium security severity: High)
CVE-2022-3888	Use after free in WebCodecs in Google Chrome prior to 107.0.5304.106 allowed a remote attacker to potentially ex page. (Chromium security severity: High)
CVE-2022-3889	Type confusion in V8 in Google Chrome prior to 107.0.5304.106 allowed a remote attacker to potentially exploit h (Chromium security severity: High)
CVE-2022-3890	Heap buffer overflow in Crashpad in Google Chrome on Android prior to 107.0.5304.106 allowed a remote attacke to potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)
CVE-2022-4135	Heap buffer overflow in GPU in Google Chrome prior to 107.0.5304.121 allowed a remote attacker who had comp perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)
CVE-2022-41724	Large handshake records may cause panics in crypto/tls. Both clients and servers may send large TLS handshake re clients, respectively, to panic when attempting to construct responses. This affects all TLS 1.3 clients, TLS 1.2 clie resumption (by setting Config.ClientSessionCache to a non-nil value), and TLS 1.3 servers which request client ce RequestClientCert).

CVE-2022-41725	A denial of service is possible from excessive resource consumption in net/http and mime/multipart. Multipart form multipart.Reader.ReadForm can consume largely unlimited amounts of memory and disk files. This also affects for Request methods FormFile, FormValue, ParseMultipartForm, and PostFormValue. ReadForm takes a maxMemory to maxMemory bytes +10MB (reserved for non-file parts) in memory". File parts which cannot be stored in memor unconfigurable 10MB reserved for non-file parts is excessively large and can potentially open a denial of service v not properly account for all memory consumed by a parsed form, such as map entry overhead, part names, and MI form to consume well over 10MB. In addition, ReadForm contained no limit on the number of disk files created, p create a large number of disk temporary files. With fix, ReadForm now properly accounts for various forms of mer its documented limit of 10MB + maxMemory bytes of memory consumption. Users should still be aware that this l In addition, ReadForm now creates at most one on-disk temporary file, combining multiple form parts into a single interface type's documentation states, "If stored on disk, the File's underlying concrete type will be an *os.File.". Th more than one file part, due to this coalescing of parts into a single file. The previous behavior of using distinct file the environment variable GODEBUG=multipartfiles=distinct. Users should be aware that multipart.ReadForm and limit the amount of disk consumed by temporary files. Callers can limit the size of form data with http.MaxBytesR
CVE-2022-4174	Type confusion in V8 in Google Chrome prior to 108.0.5359.71 allowed a remote attacker to potentially exploit he (Chromium security severity: High)
CVE-2022-4175	Use after free in Camera Capture in Google Chrome prior to 108.0.5359.71 allowed a remote attacker to potentially page. (Chromium security severity: High)
CVE-2022-4177	Use after free in Extensions in Google Chrome prior to 108.0.5359.71 allowed an attacker who convinced a user to heap corruption via a crafted Chrome Extension and UI interaction. (Chromium security severity: High)
CVE-2022-4178	Use after free in Mojo in Google Chrome prior to 108.0.5359.71 allowed a remote attacker who had compromised heap corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2022-4179	Use after free in Audio in Google Chrome prior to 108.0.5359.71 allowed an attacker who convinced a user to insta heap corruption via a crafted Chrome Extension. (Chromium security severity: High)
CVE-2022-4180	Use after free in Mojo in Google Chrome prior to 108.0.5359.71 allowed an attacker who convinced a user to insta heap corruption via a crafted Chrome Extension. (Chromium security severity: High)
CVE-2022-4181	Use after free in Forms in Google Chrome prior to 108.0.5359.71 allowed a remote attacker to potentially exploit h (Chromium security severity: High)
CVE-2022-4182	Inappropriate implementation in Fenced Frames in Google Chrome prior to 108.0.5359.71 allowed a remote attack crafted HTML page. (Chromium security severity: Medium)
CVE-2022-4183	Insufficient policy enforcement in Popup Blocker in Google Chrome prior to 108.0.5359.71 allowed a remote attac crafted HTML page. (Chromium security severity: Medium)
CVE-2022-4184	Insufficient policy enforcement in Autofill in Google Chrome prior to 108.0.5359.71 allowed a remote attacker to b HTML page. (Chromium security severity: Medium)
CVE-2022-4186	Insufficient validation of untrusted input in Downloads in Google Chrome prior to 108.0.5359.71 allowed an attack extension to bypass Downloads restrictions via a crafted HTML page. (Chromium security severity: Medium)
CVE-2022-4189	Insufficient policy enforcement in DevTools in Google Chrome prior to 108.0.5359.71 allowed an attacker who co to bypass navigation restrictions via a crafted Chrome Extension. (Chromium security severity: Medium)
CVE-2022-4190	Insufficient data validation in Directory in Google Chrome prior to 108.0.5359.71 allowed a remote attacker to byp HTML page. (Chromium security severity: Medium)
CVE-2022-4191	Use after free in Sign-In in Google Chrome prior to 108.0.5359.71 allowed a remote attacker who convinced a user potentially exploit heap corruption via profile destruction. (Chromium security severity: Medium)
CVE-2022-4192	Use after free in Live Caption in Google Chrome prior to 108.0.5359.71 allowed a remote attacker who convinced potentially exploit heap corruption via UI interaction. (Chromium security severity: Medium)
CVE-2022-4193	Insufficient policy enforcement in File System API in Google Chrome prior to 108.0.5359.71 allowed a remote atta crafted HTML page. (Chromium security severity: Medium)
CVE-2022-4194	Use after free in Accessibility in Google Chrome prior to 108.0.5359.71 allowed a remote attacker to potentially ex page. (Chromium security severity: Medium)
CVE-2022-4195	Insufficient policy enforcement in Safe Browsing in Google Chrome prior to 108.0.5359.71 allowed a remote attac malicious file. (Chromium security severity: Medium)
CVE-2022-4262	Type confusion in V8 in Google Chrome prior to 108.0.5359.94 allowed a remote attacker to potentially exploit he (Chromium security severity: High)
CVE-2022-4436	Use after free in Blink Media in Google Chrome prior to 108.0.5359.124 allowed a remote attacker to potentially e page. (Chromium security severity: High)

CVE-2022-4437	Use after free in Mojo IPC in Google Chrome prior to 108.0.5359.124 allowed a remote attacker to potentially exp (Chromium security severity: High)
CVE-2022-4438	Use after free in Blink Frames in Google Chrome prior to 108.0.5359.124 allowed a remote attacker who convince interactions to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2022-4440	Use after free in Profiles in Google Chrome prior to 108.0.5359.124 allowed a remote attacker to potentially exploi (Chromium security severity: Medium)
CVE-2022-46751	Improper Restriction of XML External Entity Reference, XML Injection (aka Blind XPath Injection) vulnerability Ivy.This issue affects any version of Apache Ivy prior to 2.5.2. When Apache Ivy prior to 2.5.2 parses XML files - Apache Maven POMs - it will allow downloading external document type definitions and expand any entity referer used to exfiltrate data, access resources only the machine running Ivy has access to or disturb the execution of Ivy i processing is disabled by default except when parsing Maven POMs where the default is to allow DTD processing with Ivy that is needed to deal with existing Maven POMs that are not valid XML files but are nevertheless accepte lenient via newly introduced system properties where needed. Users of Ivy prior to version 2.5.2 can use Java syste DTDs, see the section about "JAXP Properties for External Access restrictions" inside Oracle's "Java API for XML
CVE-2022-4906	Inappropriate implementation in Blink in Google Chrome prior to 108.0.5359.71 allowed a remote attacker to perfo page. (Chromium security severity: High)
CVE-2022-4907	Uninitialized Use in FFmpeg in Google Chrome prior to 108.0.5359.71 allowed a remote attacker to execute arbitra page. (Chromium security severity: Medium)
CVE-2022-4908	Inappropriate implementation in iFrame Sandbox in Google Chrome prior to 107.0.5304.62 allowed a remote attac HTML page. (Chromium security severity: Medium)
CVE-2022-4909	Inappropriate implementation in XML in Google Chrome prior to 107.0.5304.62 allowed a remote attacker to pote HTML page. (Chromium security severity: Low)
CVE-2022-4910	Inappropriate implementation in Autofill in Google Chrome prior to 107.0.5304.62 allowed a remote attacker to by HTML page. (Chromium security severity: Medium)
CVE-2022-4911	Insufficient data validation in DevTools in Google Chrome prior to 106.0.5249.62 allowed a remote attacker to byp HTML page. (Chromium security severity: Low)
CVE-2022-4912	Type Confusion in MathML in Google Chrome prior to 105.0.5195.52 allowed a remote attacker to potentially exp page. (Chromium security severity: High)
CVE-2022-4913	Inappropriate implementation in Extensions in Google Chrome prior to 105.0.5195.52 allowed a remote attacker w spoof extension storage via a crafted HTML page. (Chromium security severity: High)
CVE-2022-4914	Heap buffer overflow in PrintPreview in Google Chrome prior to 104.0.5112.79 allowed an attacker who convince potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: Medium)
CVE-2022-4915	Inappropriate implementation in URL Formatting in Google Chrome prior to 103.0.5060.134 allowed a remote atta crafted HTML page. (Chromium security severity: Medium)
CVE-2022-4916	Use after free in Media in Google Chrome prior to 103.0.5060.53 allowed a remote attacker to perform arbitrary re (Chromium security severity: High)
CVE-2022-4918	Use after free in UI in Google Chrome prior to 102.0.5005.61 allowed a remote attacker to perform arbitrary read/v security severity: Medium)
CVE-2022-4919	Use after free in Base Internals in Google Chrome prior to 101.0.4951.41 allowed a remote attacker to perform arb (Chromium security severity: High)
CVE-2022-4920	Heap buffer overflow in Blink in Google Chrome prior to 101.0.4951.41 allowed a remote attacker who convinced potentially perform a sandbox escape via a crafted HTML page. (Chromium security severity: High)
CVE-2022-4955	Inappropriate implementation in DevTools in Google Chrome prior to 108.0.5359.71 allowed an attacker who conv to bypass file access restrictions via a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-0129	Heap buffer overflow in Network Service in Google Chrome prior to 109.0.5414.74 allowed an attacker who convi potentially exploit heap corruption via a crafted HTML page and specific interactions. (Chromium security severity
CVE-2023-0131	Inappropriate implementation in in iframe Sandbox in Google Chrome prior to 109.0.5414.74 allowed a remote att a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-0134	Use after free in Cart in Google Chrome prior to 109.0.5414.74 allowed an attacker who convinced a user to install heap corruption via database corruption and a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-0135	Use after free in Cart in Google Chrome prior to 109.0.5414.74 allowed an attacker who convinced a user to install heap corruption via database corruption and a crafted HTML page. (Chromium security severity: Medium)

CVE-2023-0138	Heap buffer overflow in libphonenumber in Google Chrome prior to 109.0.5414.74 allowed a remote attacker to po HTML page. (Chromium security severity: Low)
CVE-2023-0141	Insufficient policy enforcement in CORS in Google Chrome prior to 109.0.5414.74 allowed a remote attacker to le page. (Chromium security severity: Low)
CVE-2023-0471	Use after free in WebTransport in Google Chrome prior to 109.0.5414.119 allowed a remote attacker to potentially page. (Chromium security severity: High)
CVE-2023-0472	Use after free in WebRTC in Google Chrome prior to 109.0.5414.119 allowed a remote attacker to potentially expl (Chromium security severity: High)
CVE-2023-0473	Type Confusion in ServiceWorker API in Google Chrome prior to 109.0.5414.119 allowed a remote attacker to po HTML page. (Chromium security severity: Medium)
CVE-2023-0474	Use after free in GuestView in Google Chrome prior to 109.0.5414.119 allowed an attacker who convinced a user exploit heap corruption via a Chrome web app. (Chromium security severity: Medium)
CVE-2023-0696	Type confusion in V8 in Google Chrome prior to 110.0.5481.77 allowed a remote attacker to potentially exploit he (Chromium security severity: High)
CVE-2023-0698	Out of bounds read in WebRTC in Google Chrome prior to 110.0.5481.77 allowed a remote attacker to perform an HTML page. (Chromium security severity: High)
CVE-2023-0699	Use after free in GPU in Google Chrome prior to 110.0.5481.77 allowed a remote attacker to potentially exploit he browser shutdown. (Chromium security severity: Medium)
CVE-2023-0700	Inappropriate implementation in Download in Google Chrome prior to 110.0.5481.77 allowed a remote attacker to (URL bar) via a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-0701	Heap buffer overflow in WebUI in Google Chrome prior to 110.0.5481.77 allowed a remote attacker who convince to potentially exploit heap corruption via UI interaction . (Chromium security severity: Medium)
CVE-2023-0702	Type confusion in Data Transfer in Google Chrome prior to 110.0.5481.77 allowed a remote attacker who convinc to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-0703	Type confusion in DevTools in Google Chrome prior to 110.0.5481.77 allowed a remote attacker who convinced a potentially exploit heap corruption via UI interactions. (Chromium security severity: Medium)
CVE-2023-0704	Insufficient policy enforcement in DevTools in Google Chrome prior to 110.0.5481.77 allowed a remote attacker t settings via a crafted HTML page. (Chromium security severity: Low)
CVE-2023-0705	Integer overflow in Core in Google Chrome prior to 110.0.5481.77 allowed a remote attacker who had one a race c corruption via a crafted HTML page. (Chromium security severity: Low)
CVE-2023-0928	Use after free in SwiftShader in Google Chrome prior to 110.0.5481.177 allowed a remote attacker to potentially e page. (Chromium security severity: High)
CVE-2023-0929	Use after free in Vulkan in Google Chrome prior to 110.0.5481.177 allowed a remote attacker to potentially exploi (Chromium security severity: High)
CVE-2023-0930	Heap buffer overflow in Video in Google Chrome prior to 110.0.5481.177 allowed a remote attacker to potentially page. (Chromium security severity: High)
CVE-2023-0931	Use after free in Video in Google Chrome prior to 110.0.5481.177 allowed a remote attacker to potentially exploit (Chromium security severity: High)
CVE-2023-0933	Integer overflow in PDF in Google Chrome prior to 110.0.5481.177 allowed a remote attacker to potentially exploi (Chromium security severity: Medium)
CVE-2023-0941	Use after free in Prompts in Google Chrome prior to 110.0.5481.177 allowed a remote attacker to potentially explo (Chromium security severity: Critical)
CVE-2023-1213	Use after free in Swiftshader in Google Chrome prior to 111.0.5563.64 allowed a remote attacker to potentially exp page. (Chromium security severity: High)
CVE-2023-1214	Type confusion in V8 in Google Chrome prior to 111.0.5563.64 allowed a remote attacker to potentially exploit he (Chromium security severity: High)
CVE-2023-1215	Type confusion in CSS in Google Chrome prior to 111.0.5563.64 allowed a remote attacker to potentially exploit h (Chromium security severity: High)
CVE-2023-1216	Use after free in DevTools in Google Chrome prior to 111.0.5563.64 allowed a remote attacker who had convience to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)

CVE-2023-1218	Use after free in WebRTC in Google Chrome prior to 111.0.5563.64 allowed a remote attacker to potentially explo (Chromium security severity: High)
CVE-2023-1219	Heap buffer overflow in Metrics in Google Chrome prior to 111.0.5563.64 allowed a remote attacker who had com exploit heap corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2023-1220	Heap buffer overflow in UMA in Google Chrome prior to 111.0.5563.64 allowed a remote attacker who had compr exploit heap corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2023-1221	Insufficient policy enforcement in Extensions API in Google Chrome prior to 111.0.5563.64 allowed an attacker w extension to bypass navigation restrictions via a crafted Chrome Extension. (Chromium security severity: Medium)
CVE-2023-1222	Heap buffer overflow in Web Audio API in Google Chrome prior to 111.0.5563.64 allowed a remote attacker to po HTML page. (Chromium security severity: Medium)
CVE-2023-1224	Insufficient policy enforcement in Web Payments API in Google Chrome prior to 111.0.5563.64 allowed a remote a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-1226	Insufficient policy enforcement in Web Payments API in Google Chrome prior to 111.0.5563.64 allowed a remote a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-1229	Inappropriate implementation in Permission prompts in Google Chrome prior to 111.0.5563.64 allowed a remote a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-1232	Insufficient policy enforcement in Resource Timing in Google Chrome prior to 111.0.5563.64 allowed a remote att information from API via a crafted HTML page. (Chromium security severity: Low)
CVE-2023-1233	Insufficient policy enforcement in Resource Timing in Google Chrome prior to 111.0.5563.64 allowed an attacker extension to obtain potentially sensitive information from API via a crafted Chrome Extension. (Chromium securit
CVE-2023-1235	Type confusion in DevTools in Google Chrome prior to 111.0.5563.64 allowed a remote attacker who had compro exploit heap corruption via a crafted UI interaction. (Chromium security severity: Low)
CVE-2023-1236	Inappropriate implementation in Internals in Google Chrome prior to 111.0.5563.64 allowed a remote attacker to sp HTML page. (Chromium security severity: Low)
CVE-2023-1528	Use after free in Passwords in Google Chrome prior to 111.0.5563.110 allowed a remote attacker who had compro exploit heap corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2023-1529	Out of bounds memory access in WebHID in Google Chrome prior to 111.0.5563.110 allowed a remote attacker to malicious HID device. (Chromium security severity: High)
CVE-2023-1530	Use after free in PDF in Google Chrome prior to 111.0.5563.110 allowed a remote attacker to potentially exploit h (Chromium security severity: High)
CVE-2023-1531	Use after free in ANGLE in Google Chrome prior to 111.0.5563.110 allowed a remote attacker to potentially explo (Chromium security severity: High)
CVE-2023-1532	Out of bounds read in GPU Video in Google Chrome prior to 111.0.5563.110 allowed a remote attacker to potentia HTML page. (Chromium security severity: High)
CVE-2023-1533	Use after free in WebProtect in Google Chrome prior to 111.0.5563.110 allowed a remote attacker to potentially ex page. (Chromium security severity: High)
CVE-2023-1534	Out of bounds read in ANGLE in Google Chrome prior to 111.0.5563.110 allowed a remote attacker who had com exploit heap corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2023-1810	Heap buffer overflow in Visuals in Google Chrome prior to 112.0.5615.49 allowed a remote attacker who had com exploit heap corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2023-1811	Use after free in Frames in Google Chrome prior to 112.0.5615.49 allowed a remote attacker who convinced a user potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2023-1812	Out of bounds memory access in DOM Bindings in Google Chrome prior to 112.0.5615.49 allowed a remote attack via a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-1813	Inappropriate implementation in Extensions in Google Chrome prior to 112.0.5615.49 allowed an attacker who cor to bypass file access restrictions via a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-1814	Insufficient validation of untrusted input in Safe Browsing in Google Chrome prior to 112.0.5615.49 allowed a rem a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-1815	Use after free in Networking APIs in Google Chrome prior to 112.0.5615.49 allowed a remote attacker who convin to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: Medium)

CVE-2023-1816	Incorrect security UI in Picture In Picture in Google Chrome prior to 112.0.5615.49 allowed a remote attacker to p crafted HTML page. (Chromium security severity: Medium)
CVE-2023-1817	Insufficient policy enforcement in Intents in Google Chrome on Android prior to 112.0.5615.49 allowed a remote a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-1818	Use after free in Vulkan in Google Chrome prior to 112.0.5615.49 allowed a remote attacker to potentially exploit (Chromium security severity: Medium)
CVE-2023-1819	Out of bounds read in Accessibility in Google Chrome prior to 112.0.5615.49 allowed a remote attacker to perform HTML page. (Chromium security severity: Medium)
CVE-2023-1820	Heap buffer overflow in Browser History in Google Chrome prior to 112.0.5615.49 allowed a remote attacker who interaction to potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-1821	Inappropriate implementation in WebShare in Google Chrome prior to 112.0.5615.49 allowed a remote attacker to (URL bar) via a crafted HTML page. (Chromium security severity: Low)
CVE-2023-1822	Incorrect security UI in Navigation in Google Chrome prior to 112.0.5615.49 allowed a remote attacker to perform (Chromium security severity: Low)
CVE-2023-1823	Inappropriate implementation in FedCM in Google Chrome prior to 112.0.5615.49 allowed a remote attacker to by HTML page. (Chromium security severity: Low)
CVE-2023-2033	Type confusion in V8 in Google Chrome prior to 112.0.5615.121 allowed a remote attacker to potentially exploit h (Chromium security severity: High)
CVE-2023-20883	In Spring Boot versions 3.0.0 - 3.0.6, 2.7.0 - 2.7.11, 2.6.0 - 2.6.14, 2.5.0 - 2.5.14 and older unsupported versions, th attack if Spring MVC is used together with a reverse proxy cache.
CVE-2023-2133	Out of bounds memory access in Service Worker API in Google Chrome prior to 112.0.5615.137 allowed a remote corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2023-2134	Out of bounds memory access in Service Worker API in Google Chrome prior to 112.0.5615.137 allowed a remote corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2023-2135	Use after free in DevTools in Google Chrome prior to 112.0.5615.137 allowed a remote attacker who convinced a potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2023-2136	Integer overflow in Skia in Google Chrome prior to 112.0.5615.137 allowed a remote attacker who had compromis a sandbox escape via a crafted HTML page. (Chromium security severity: High)
CVE-2023-2137	Heap buffer overflow in sqlite in Google Chrome prior to 112.0.5615.137 allowed a remote attacker to potentially page. (Chromium security severity: Medium)
CVE-2023-2311	Insufficient policy enforcement in File System API in Google Chrome prior to 112.0.5615.49 allowed a remote att crafted HTML page. (Chromium security severity: Medium)
CVE-2023-2314	Insufficient data validation in DevTools in Google Chrome prior to 111.0.5563.64 allowed a remote attacker to byp HTML page. (Chromium security severity: Low)
CVE-2023-23931	cryptography is a package designed to expose cryptographic primitives and recipes to Python developers. In affecte accept Python objects which implement the buffer protocol, but provide only immutable buffers. This would allow mutated, thus violating fundamental rules of Python and resulting in corrupted output. This now correctly raises an `update_into` was originally introduced in cryptography 1.8.
CVE-2023-24534	HTTP and MIME header parsing can allocate large amounts of memory, even when parsing small inputs, potential unusual patterns of input data can cause the common function used to parse HTTP and MIME headers to allocate s hold the parsed headers. An attacker can exploit this behavior to cause an HTTP server to allocate large amounts of leading to memory exhaustion and a denial of service. With fix, header parsing now correctly allocates only the me
CVE-2023-24536	Multipart form parsing can consume large amounts of CPU and memory when processing form inputs containing v several causes: 1. mime/multipart.Reader.ReadForm limits the total memory a parsed multipart form can consume. memory consumed, leading it to accept larger inputs than intended. 2. Limiting total memory does not account for from large numbers of small allocations in forms with many parts. 3. ReadForm can allocate a large number of sho on the garbage collector. The combination of these factors can permit an attacker to cause an program that parses m of CPU and memory, potentially resulting in a denial of service. This affects programs that use mime/multipart.Rea the net/http package with the Request methods FormFile, FormValue, ParseMultipartForm, and PostFormValue. W estimating the memory consumption of parsed forms, and performs many fewer short-lived allocations. In addition the following limits on the size of parsed forms: 1. Forms parsed with ReadForm may contain no more than 1000 p environment variable GODEBUG=multipartmaxparts=. 2. Form parts parsed with NextPart and NextRawPart may addition, forms parsed with ReadForm may contain no more than 10,000 header fields across all parts. This limit m GODEBUG=multipartmaxheaders=.

CVE-2023-2459	Inappropriate implementation in Prompts in Google Chrome prior to 113.0.5672.63 allowed a remote attacker to by HTML page. (Chromium security severity: Medium)
CVE-2023-2460	Insufficient validation of untrusted input in Extensions in Google Chrome prior to 113.0.5672.63 allowed an attack extension to bypass file access checks via a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-2462	Inappropriate implementation in Prompts in Google Chrome prior to 113.0.5672.63 allowed a remote attacker to ob page. (Chromium security severity: Medium)
CVE-2023-2464	Inappropriate implementation in PictureInPicture in Google Chrome prior to 113.0.5672.63 allowed an attacker wh extension to perform an origin spoof in the security UI via a crafted HTML page. (Chromium security severity: Me
CVE-2023-2465	Inappropriate implementation in CORS in Google Chrome prior to 113.0.5672.63 allowed a remote attacker to leak (Chromium security severity: Medium)
CVE-2023-2466	Inappropriate implementation in Prompts in Google Chrome prior to 113.0.5672.63 allowed a remote attacker to sp crafted HTML page. (Chromium security severity: Low)
CVE-2023-2468	Inappropriate implementation in PictureInPicture in Google Chrome prior to 113.0.5672.63 allowed a remote attac process to obfuscate the security UI via a crafted HTML page. (Chromium security severity: Low)
CVE-2023-26112	All versions of the package configobj are vulnerable to Regular Expression Denial of Service (ReDoS) via the vali This is only exploitable in the case of a developer, putting the offending value in a server side configuration file.
CVE-2023-2721	Use after free in Navigation in Google Chrome prior to 113.0.5672.126 allowed a remote attacker to potentially ex page. (Chromium security severity: Critical)
CVE-2023-2723	Use after free in DevTools in Google Chrome prior to 113.0.5672.126 allowed a remote attacker who had compron exploit heap corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2023-2724	Type confusion in V8 in Google Chrome prior to 113.0.5672.126 allowed a remote attacker to potentially exploit h (Chromium security severity: High)
CVE-2023-2725	Use after free in Guest View in Google Chrome prior to 113.0.5672.126 allowed an attacker who convinced a user exploit heap corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2023-2726	Inappropriate implementation in WebApp Installs in Google Chrome prior to 113.0.5672.126 allowed an attacker v web app to bypass install dialog via a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-2929	Out of bounds write in Swiftshader in Google Chrome prior to 114.0.5735.90 allowed a remote attacker to potentia HTML page. (Chromium security severity: High)
CVE-2023-2930	Use after free in Extensions in Google Chrome prior to 114.0.5735.90 allowed an attacker who convinced a user to exploit heap corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2023-2931	Use after free in PDF in Google Chrome prior to 114.0.5735.90 allowed a remote attacker to potentially exploit hea (Chromium security severity: High)
CVE-2023-2932	Use after free in PDF in Google Chrome prior to 114.0.5735.90 allowed a remote attacker to potentially exploit hea (Chromium security severity: High)
CVE-2023-2933	Use after free in PDF in Google Chrome prior to 114.0.5735.90 allowed a remote attacker to potentially exploit hea (Chromium security severity: High)
CVE-2023-2934	Out of bounds memory access in Mojo in Google Chrome prior to 114.0.5735.90 allowed a remote attacker to pote HTML page. (Chromium security severity: High)
CVE-2023-2935	Type Confusion in V8 in Google Chrome prior to 114.0.5735.90 allowed a remote attacker to potentially exploit he (Chromium security severity: High)
CVE-2023-2936	Type Confusion in V8 in Google Chrome prior to 114.0.5735.90 allowed a remote attacker to potentially exploit he (Chromium security severity: High)
CVE-2023-2937	Inappropriate implementation in Picture In Picture in Google Chrome prior to 114.0.5735.90 allowed a remote atta process to spoof the contents of the Omnibox (URL bar) via a crafted HTML page. (Chromium security severity: M
CVE-2023-2938	Inappropriate implementation in Picture In Picture in Google Chrome prior to 114.0.5735.90 allowed a remote atta process to spoof the contents of the Omnibox (URL bar) via a crafted HTML page. (Chromium security severity: M
CVE-2023-2940	Inappropriate implementation in Downloads in Google Chrome prior to 114.0.5735.90 allowed an attacker who co to bypass file access restrictions via a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-29406	The HTTP/1 client does not fully validate the contents of the Host header. A maliciously crafted Host header can in With fix, the HTTP/1 client now refuses to send requests containing an invalid Request.Host or Request.URL.Host

CVE-2023-2941	Inappropriate implementation in Extensions API in Google Chrome prior to 114.0.5735.90 allowed an attacker wh extension to spoof the contents of the UI via a crafted Chrome Extension. (Chromium security severity: Low)
CVE-2023-30581	The use of __proto__ in process.mainModule.__proto__.require() can bypass the policy mechanism and require mo This vulnerability affects all users using the experimental policy mechanism in all active release lines: v16, v18 an was issued, the policy is an experimental feature of Node.js
CVE-2023-30588	When an invalid public key is used to create an x509 certificate using the crypto.X509Certificate() API a non-expe to DoS attacks when the attacker could force interruptions of application processing, as the process terminates whe certificates from user code. The current context of the users will be gone, and that will cause a DoS scenario. This versions v16, v18, and, v20.
CVE-2023-30589	The llhttp parser in the http module in Node v20.2.0 does not strictly use the CRLF sequence to delimit HTTP requ Smuggling (HRS). The CR character (without LF) is sufficient to delimit HTTP header fields in the llhttp parser. A CRLF sequence should delimit each header-field. This impacts all Node.js active versions: v16, v18, and, v20
CVE-2023-30590	The generateKeys() API function returned from crypto.createDiffieHellman() only generates missing (or outdated) if none has been set yet, but the function is also needed to compute the corresponding public key after calling setPr this API call: "Generates private and public Diffie-Hellman key values". The documented behavior is very differen could easily lead to security issues in applications that use these APIs as the DiffieHellman may be used as the basi are consequently broad.
CVE-2023-3079	Type confusion in V8 in Google Chrome prior to 114.0.5735.110 allowed a remote attacker to potentially exploit h (Chromium security severity: High)
CVE-2023-30861	Flask is a lightweight WSGI web application framework. When all of the following conditions are met, a response client may be cached and subsequently sent by the proxy to other clients. If the proxy also caches `Set-Cookie` hea `session` cookie to other clients. The severity depends on the application's use of the session and the proxy's behavi depends on all these conditions being met. 1. The application must be hosted behind a caching proxy that does not cookies. 2. The application sets `session.permanent = True` 3. The application does not access or modify the sessio `SESSION_REFRESH_EACH_REQUEST` enabled (the default). 5. The application does not set a `Cache-Contro should not be cached. This happens because vulnerable versions of Flask only set the `Vary: Cookie` header when it is refreshed (re-sent to update the expiration) without being accessed or modified. This issue has been fixed in ve
CVE-2023-3214	Use after free in Autofill payments in Google Chrome prior to 114.0.5735.133 allowed a remote attacker to potenti HTML page. (Chromium security severity: Critical)
CVE-2023-3215	Use after free in WebRTC in Google Chrome prior to 114.0.5735.133 allowed a remote attacker to potentially expl (Chromium security severity: High)
CVE-2023-3216	Type confusion in V8 in Google Chrome prior to 114.0.5735.133 allowed a remote attacker to potentially exploit h (Chromium security severity: High)
CVE-2023-3217	Use after free in WebXR in Google Chrome prior to 114.0.5735.133 allowed a remote attacker to potentially explo (Chromium security severity: High)
CVE-2023-33953	gRPC contains a vulnerability that allows hpack table accounting errors could lead to unwanted disconnects betwe ¬†Three vectors were found that allow the following DOS attacks: - Unbounded memory buffering in the HPACK the HPACK parser The unbounded CPU consumption is down to a copy that occurred per-input-block in the parse to the memory copy bug we end up with an O(n^2) parsing loop, with n selected by the client. The unbounded mem check was behind the string reading code, so we needed to first buffer up to a 4 gigabyte string before rejecting it a have an encoding quirk whereby an infinite number of 0,Äôs can be added at the start of an integer. gRPC,Äôs hpa concluding a parse. - gRPC,Äôs metadata overflow check was performed per frame, so that the following sequence HEADERS: containing a: 1 CONTINUATION: containing a: 2 CONTINUATION: containing a: 3 etc,Ä¶
CVE-2023-34054	In Reactor Netty HTTP Server, versions 1.1.x prior to 1.1.13 and versions 1.0.x prior to 1.0.39, it is possible for a u requests that may cause a denial-of-service (DoS) condition. Specifically, an application is vulnerable if Reactor N Micrometer is enabled.
CVE-2023-34055	In Spring Boot versions 2.7.0 - 2.7.17, 3.0.0-3.0.12 and 3.1.0-3.1.5, it is possible for a user to provide specially cra of-service (DoS) condition. Specifically, an application is vulnerable when all of the following are true: * the appli org.springframework.boot:spring-boot-actuator¬†is on the classpath
CVE-2023-34062	In Reactor Netty HTTP Server, versions 1.1.x prior to 1.1.13 and versions 1.0.x prior to 1.0.39, a malicious user ca URL that can lead to a directory traversal attack. Specifically, an application is vulnerable if Reactor Netty HTTP S
CVE-2023-34110	Flask-AppBuilder is an application development framework, built on top of Flask. Prior to version 4.3.2, an authen privileges, could by adding a special character on the add, edit User forms trigger a database error, this error is surf database engines this error can include the entire user row including the pbkdf2:sha256 hashed password. This vuln
CVE-2023-3420	Type Confusion in V8 in Google Chrome prior to 114.0.5735.198 allowed a remote attacker to potentially exploit h (Chromium security severity: High)

CVE-2023-3421	Use after free in Media in Google Chrome prior to 114.0.5735.198 allowed a remote attacker to potentially exploit (Chromium security severity: High)
CVE-2023-3422	Use after free in Guest View in Google Chrome prior to 114.0.5735.198 allowed an attacker who convinced a user exploit heap corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2023-3598	Out of bounds read and write in ANGLE in Google Chrome prior to 114.0.5735.90 allowed a remote attacker to po HTML page. (Chromium security severity: High)
CVE-2023-3727	Use after free in WebRTC in Google Chrome prior to 115.0.5790.98 allowed a remote attacker to potentially explo (Chromium security severity: High)
CVE-2023-37276	aiohttp is an asynchronous HTTP client/server framework for asyncio and Python. aiohttp v3.8.4 and earlier are bu code is used by aiohttp for its HTTP request parser when available which is the default case when installing from a users of aiohttp as an HTTP server (ie `aiohttp.Application`), you are not affected by this vulnerability if you are us (ie `aiohttp.ClientSession`). Sending a crafted HTTP request will cause the server to misinterpret one of the HTTP request smuggling. This issue has been addressed in version 3.8.5. Users are advised to upgrade. Users unable to up `AIOHTTP_NO_EXTENSIONS=1` as an environment variable to disable the llhttp HTTP request parser impleme vulnerable.
CVE-2023-3728	Use after free in WebRTC in Google Chrome prior to 115.0.5790.98 allowed a remote attacker to potentially explo (Chromium security severity: High)
CVE-2023-3730	Use after free in Tab Groups in Google Chrome prior to 115.0.5790.98 allowed a remote attacker who convinced a potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2023-3732	Out of bounds memory access in Mojo in Google Chrome prior to 115.0.5790.98 allowed a remote attacker who ha potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2023-3733	Inappropriate implementation in WebApp Installs in Google Chrome prior to 115.0.5790.98 allowed a remote attac Omnibox (URL bar) via a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-3734	Inappropriate implementation in Picture In Picture in Google Chrome prior to 115.0.5790.98 allowed a remote atta Omnibox (URL bar) via a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-3735	Inappropriate implementation in Web API Permission Prompts in Google Chrome prior to 115.0.5790.98 allowed a a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-3737	Inappropriate implementation in Notifications in Google Chrome prior to 115.0.5790.98 allowed a remote attacker via a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-3738	Inappropriate implementation in Autofill in Google Chrome prior to 115.0.5790.98 allowed a remote attacker to ob (Chromium security severity: Medium)
CVE-2023-3740	Insufficient validation of untrusted input in Themes in Google Chrome prior to 115.0.5790.98 allowed a remote att a user via a crafted background URL. (Chromium security severity: Low)
CVE-2023-37415	Improper Input Validation vulnerability in Apache Software Foundation Apache Airflow Apache Hive Provider. Pa Before¬†6.1.2¬†the proxy_user option can also inject semicolon. This issue affects Apache Airflow Apache Hive updating provider version to 6.1.2 in order to avoid this vulnerability.
CVE-2023-39321	Processing an incomplete post-handshake message for a QUIC connection can cause a panic.
CVE-2023-39322	QUIC connections do not set an upper bound on the amount of data buffered when reading post-handshake messag cause unbounded memory growth. With fix, connections now consistently reject messages larger than 65KiB in siz
CVE-2023-4068	Type Confusion in V8 in Google Chrome prior to 115.0.5790.170 allowed a remote attacker to perform arbitrary re (Chromium security severity: High)
CVE-2023-4069	Type Confusion in V8 in Google Chrome prior to 115.0.5790.170 allowed a remote attacker to potentially exploit h (Chromium security severity: High)
CVE-2023-4070	Type Confusion in V8 in Google Chrome prior to 115.0.5790.170 allowed a remote attacker to perform arbitrary re (Chromium security severity: High)
CVE-2023-4071	Heap buffer overflow in Visuals in Google Chrome prior to 115.0.5790.170 allowed a remote attacker to potentiall page. (Chromium security severity: High)
CVE-2023-4072	Out of bounds read and write in WebGL in Google Chrome prior to 115.0.5790.170 allowed a remote attacker to p HTML page. (Chromium security severity: High)
CVE-2023-4074	Use after free in Blink Task Scheduling in Google Chrome prior to 115.0.5790.170 allowed a remote attacker to po HTML page. (Chromium security severity: High)
CVE-2023-4075	Use after free in Cast in Google Chrome prior to 115.0.5790.170 allowed a remote attacker to potentially exploit he (Chromium security severity: High)

CVE-2023-4076	Use after free in WebRTC in Google Chrome prior to 115.0.5790.170 allowed a remote attacker to potentially expl session. (Chromium security severity: High)
CVE-2023-4077	Insufficient data validation in Extensions in Google Chrome prior to 115.0.5790.170 allowed an attacker who conv inject scripts or HTML into a privileged page via a crafted Chrome Extension. (Chromium security severity: Mediu
CVE-2023-4078	Inappropriate implementation in Extensions in Google Chrome prior to 115.0.5790.170 allowed an attacker who co extension to inject scripts or HTML into a privileged page via a crafted Chrome Extension. (Chromium security se
CVE-2023-41419	An issue in Gevent before version 23.9.0 allows a remote attacker to escalate privileges via a crafted script to the W
CVE-2023-4349	Use after free in Device Trust Connectors in Google Chrome prior to 116.0.5845.96 allowed a remote attacker to p crafted HTML page. (Chromium security severity: High)
CVE-2023-4351	Use after free in Network in Google Chrome prior to 116.0.5845.96 allowed a remote attacker who has elicited a br corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2023-4352	Type confusion in V8 in Google Chrome prior to 116.0.5845.96 allowed a remote attacker to potentially exploit he (Chromium security severity: High)
CVE-2023-4353	Heap buffer overflow in ANGLE in Google Chrome prior to 116.0.5845.96 allowed a remote attacker to potentially page. (Chromium security severity: High)
CVE-2023-4354	Heap buffer overflow in Skia in Google Chrome prior to 116.0.5845.96 allowed a remote attacker who had compro exploit heap corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2023-4355	Out of bounds memory access in V8 in Google Chrome prior to 116.0.5845.96 allowed a remote attacker to potenti HTML page. (Chromium security severity: High)
CVE-2023-4356	Use after free in Audio in Google Chrome prior to 116.0.5845.96 allowed a remote attacker who has convinced a u potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-4357	Insufficient validation of untrusted input in XML in Google Chrome prior to 116.0.5845.96 allowed a remote attac crafted HTML page. (Chromium security severity: Medium)
CVE-2023-4358	Use after free in DNS in Google Chrome prior to 116.0.5845.96 allowed a remote attacker to potentially exploit he (Chromium security severity: Medium)
CVE-2023-4360	Inappropriate implementation in Color in Google Chrome prior to 116.0.5845.96 allowed a remote attacker to obfu (Chromium security severity: Medium)
CVE-2023-4362	Heap buffer overflow in Mojom IDL in Google Chrome prior to 116.0.5845.96 allowed a remote attacker who had gained control of a WebUI process to potentially exploit heap corruption via a crafted HTML page. (Chromium sec
CVE-2023-4364	Inappropriate implementation in Permission Prompts in Google Chrome prior to 116.0.5845.96 allowed a remote a HTML page. (Chromium security severity: Medium)
CVE-2023-4365	Inappropriate implementation in Fullscreen in Google Chrome prior to 116.0.5845.96 allowed a remote attacker to page. (Chromium security severity: Medium)
CVE-2023-4366	Use after free in Extensions in Google Chrome prior to 116.0.5845.96 allowed an attacker who convinced a user to exploit heap corruption via a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-4367	Insufficient policy enforcement in Extensions API in Google Chrome prior to 116.0.5845.96 allowed an attacker w extension to bypass an enterprise policy via a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-4368	Insufficient policy enforcement in Extensions API in Google Chrome prior to 116.0.5845.96 allowed an attacker w extension to bypass an enterprise policy via a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-4427	Out of bounds memory access in V8 in Google Chrome prior to 116.0.5845.110 allowed a remote attacker to perfo crafted HTML page. (Chromium security severity: High)
CVE-2023-4428	Out of bounds memory access in CSS in Google Chrome prior to 116.0.5845.110 allowed a remote attacker to perf crafted HTML page. (Chromium security severity: High)
CVE-2023-4429	Use after free in Loader in Google Chrome prior to 116.0.5845.110 allowed a remote attacker to potentially exploit (Chromium security severity: High)
CVE-2023-4430	Use after free in Vulkan in Google Chrome prior to 116.0.5845.110 allowed a remote attacker to potentially exploi (Chromium security severity: High)
CVE-2023-4431	Out of bounds memory access in Fonts in Google Chrome prior to 116.0.5845.110 allowed a remote attacker to per crafted HTML page. (Chromium security severity: Medium)
CVE-2023-4572	Use after free in MediaStream in Google Chrome prior to 116.0.5845.140 allowed a remote attacker to potentially page. (Chromium security severity: High)

CVE-2023-47248	Deserialization of untrusted data in IPC and Parquet readers in PyArrow versions 0.14.0 to 14.0.0 allows arbitrary c if it reads Arrow IPC, Feather or Parquet data from untrusted sources (for example user-supplied input files). This v other Apache Arrow implementations or bindings. It is recommended that users of PyArrow upgrade to 14.0.1. Sim libraries upgrade their dependency requirements to PyArrow 14.0.1 or later. PyPI packages are already available, a be available soon. If it is not possible to upgrade, we provide a separate package `pyarrow-hotfix` that disables the https://pypi.org/project/pyarrow-hotfix/ for instructions.
CVE-2023-4761	Out of bounds memory access in FedCM in Google Chrome prior to 116.0.5845.179 allowed a remote attacker wh perform an out of bounds memory read via a crafted HTML page. (Chromium security severity: High)
CVE-2023-4762	Type Confusion in V8 in Google Chrome prior to 116.0.5845.179 allowed a remote attacker to execute arbitrary co security severity: High)
CVE-2023-47627	aiohttp is an asynchronous HTTP client/server framework for asyncio and Python. The HTTP parser in AIOHTTP which could lead to request smuggling. This parser is only used when AIOHTTP_NO_EXTENSIONS is enabled (have been addressed in commit `d5c12ba89` which has been included in release version 3.8.6. Users are advised to for these issues.
CVE-2023-4763	Use after free in Networks in Google Chrome prior to 116.0.5845.179 allowed a remote attacker to potentially expl (Chromium security severity: High)
CVE-2023-4764	Incorrect security UI in BFCache in Google Chrome prior to 116.0.5845.179 allowed a remote attacker to spoof the crafted HTML page. (Chromium security severity: High)
CVE-2023-4901	Inappropriate implementation in Prompts in Google Chrome prior to 117.0.5938.62 allowed a remote attacker to po HTML page. (Chromium security severity: Medium)
CVE-2023-4902	Inappropriate implementation in Input in Google Chrome prior to 117.0.5938.62 allowed a remote attacker to spoo (Chromium security severity: Medium)
CVE-2023-4904	Insufficient policy enforcement in Downloads in Google Chrome prior to 117.0.5938.62 allowed a remote attacker crafted download. (Chromium security severity: Medium)
CVE-2023-4905	Inappropriate implementation in Prompts in Google Chrome prior to 117.0.5938.62 allowed a remote attacker to sp (Chromium security severity: Medium)
CVE-2023-4906	Insufficient policy enforcement in Autofill in Google Chrome prior to 117.0.5938.62 allowed a remote attacker to b HTML page. (Chromium security severity: Low)
CVE-2023-4908	Inappropriate implementation in Picture in Picture in Google Chrome prior to 117.0.5938.62 allowed a remote atta page. (Chromium security severity: Low)
CVE-2023-49081	aiohttp is an asynchronous HTTP client/server framework for asyncio and Python. Improper validation made it pos request (e.g. to insert a new header) or create a new HTTP request if the attacker controls the HTTP version. The v control the HTTP version of the request. This issue has been patched in version 3.9.0.
CVE-2023-49082	aiohttp is an asynchronous HTTP client/server framework for asyncio and Python. Improper validation makes it po request (e.g. insert a new header) or even create a new HTTP request if the attacker controls the HTTP method. Th can control the HTTP method (GET, POST etc.) of the request. If the attacker can control the HTTP version of the (request smuggling). This issue has been patched in version 3.9.0.
CVE-2023-4909	Inappropriate implementation in Interstitials in Google Chrome prior to 117.0.5938.62 allowed a remote attacker to page. (Chromium security severity: Low)
CVE-2023-50658	The jose2go component before 1.6.0 for Go allows attackers to cause a denial of service (CPU consumption) via a
CVE-2023-51764	Postfix through 3.8.5 allows SMTP smuggling unless configured with smtpd_data_restrictions=reject_unauth_pipe smtpd_discard_ehlo_keywords=chunking (or certain other options that exist in recent versions). Remote attackers c to inject e-mail messages with a spoofed MAIL FROM address, allowing bypass of an SPF protection mechanism. <LF>.<CR><LF> but some other popular e-mail servers do not. To prevent attack variants (by always disallowing required, such as the smtpd_forbid_bare_newline=yes option with a Postfix minimum version of 3.5.23, 3.6.13, 3.7
CVE-2023-5186	Use after free in Passwords in Google Chrome prior to 117.0.5938.132 allowed a remote attacker who convinced a potentially exploit heap corruption via crafted UI interaction. (Chromium security severity: High)
CVE-2023-5187	Use after free in Extensions in Google Chrome prior to 117.0.5938.132 allowed an attacker who convinced a user t exploit heap corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2023-5218	Use after free in Site Isolation in Google Chrome prior to 118.0.5993.70 allowed a remote attacker to potentially ex page. (Chromium security severity: Critical)
CVE-2023-5346	Type confusion in V8 in Google Chrome prior to 117.0.5938.149 allowed a remote attacker to potentially exploit h (Chromium security severity: High)

CVE-2023-5472	Use after free in Profiles in Google Chrome prior to 118.0.5993.117 allowed a remote attacker to potentially exploi (Chromium security severity: High)
CVE-2023-5473	Use after free in Cast in Google Chrome prior to 118.0.5993.70 allowed a remote attacker who had compromised th corruption via a crafted HTML page. (Chromium security severity: Low)
CVE-2023-5474	Heap buffer overflow in PDF in Google Chrome prior to 118.0.5993.70 allowed a remote attacker who convinced a to potentially exploit heap corruption via a crafted PDF file. (Chromium security severity: Medium)
CVE-2023-5475	Inappropriate implementation in DevTools in Google Chrome prior to 118.0.5993.70 allowed an attacker who conv to bypass discretionary access control via a crafted Chrome Extension. (Chromium security severity: Medium)
CVE-2023-5476	Use after free in Blink History in Google Chrome prior to 118.0.5993.70 allowed a remote attacker to potentially e page. (Chromium security severity: Medium)
CVE-2023-5477	Inappropriate implementation in Installer in Google Chrome prior to 118.0.5993.70 allowed a local attacker to bypa command. (Chromium security severity: Low)
CVE-2023-5478	Inappropriate implementation in Autofill in Google Chrome prior to 118.0.5993.70 allowed a remote attacker to lea (Chromium security severity: Low)
CVE-2023-5479	Inappropriate implementation in Extensions API in Google Chrome prior to 118.0.5993.70 allowed an attacker who extension to bypass an enterprise policy via a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-5480	Inappropriate implementation in Payments in Google Chrome prior to 119.0.6045.105 allowed a remote attacker to (Chromium security severity: High)
CVE-2023-5481	Inappropriate implementation in Downloads in Google Chrome prior to 118.0.5993.70 allowed a remote attacker to (Chromium security severity: Medium)
CVE-2023-5482	Insufficient data validation in USB in Google Chrome prior to 119.0.6045.105 allowed a remote attacker to perform HTML page. (Chromium security severity: High)
CVE-2023-5483	Inappropriate implementation in Intents in Google Chrome prior to 118.0.5993.70 allowed a remote attacker to byp HTML page. (Chromium security severity: Medium)
CVE-2023-5484	Inappropriate implementation in Navigation in Google Chrome prior to 118.0.5993.70 allowed a remote attacker to (Chromium security severity: Medium)
CVE-2023-5485	Inappropriate implementation in Autofill in Google Chrome prior to 118.0.5993.70 allowed a remote attacker to by page. (Chromium security severity: Low)
CVE-2023-5486	Inappropriate implementation in Input in Google Chrome prior to 118.0.5993.70 allowed a remote attacker to spoo (Chromium security severity: Low)
CVE-2023-5487	Inappropriate implementation in Fullscreen in Google Chrome prior to 118.0.5993.70 allowed an attacker who con to bypass navigation restrictions via a crafted Chrome Extension. (Chromium security severity: Medium)
CVE-2023-5849	Integer overflow in USB in Google Chrome prior to 119.0.6045.105 allowed a remote attacker to potentially explo (Chromium security severity: High)
CVE-2023-5850	Incorrect security UI in Downloads in Google Chrome prior to 119.0.6045.105 allowed a remote attacker to perfor name. (Chromium security severity: Medium)
CVE-2023-5851	Inappropriate implementation in Downloads in Google Chrome prior to 119.0.6045.105 allowed a remote attacker page. (Chromium security severity: Medium)
CVE-2023-5852	Use after free in Printing in Google Chrome prior to 119.0.6045.105 allowed a remote attacker who convinced a us potentially exploit heap corruption via specific UI gestures. (Chromium security severity: Medium)
CVE-2023-5853	Incorrect security UI in Downloads in Google Chrome prior to 119.0.6045.105 allowed a remote attacker to obfusc (Chromium security severity: Medium)
CVE-2023-5854	Use after free in Profiles in Google Chrome prior to 119.0.6045.105 allowed a remote attacker who convinced a us potentially exploit heap corruption via specific UI gestures. (Chromium security severity: Medium)
CVE-2023-5855	Use after free in Reading Mode in Google Chrome prior to 119.0.6045.105 allowed a remote attacker who convinc potentially exploit heap corruption via specific UI gestures. (Chromium security severity: Medium)
CVE-2023-5856	Use after free in Side Panel in Google Chrome prior to 119.0.6045.105 allowed a remote attacker who convinced a potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: Medium)
CVE-2023-5857	Inappropriate implementation in Downloads in Google Chrome prior to 119.0.6045.105 allowed a remote attacker malicious file. (Chromium security severity: Medium)

CVE-2023-5858	Inappropriate implementation in WebApp Provider in Google Chrome prior to 119.0.6045.105 allowed a remote at HTML page. (Chromium security severity: Low)
CVE-2023-5859	Incorrect security UI in Picture In Picture in Google Chrome prior to 119.0.6045.105 allowed a remote attacker to HTML page. (Chromium security severity: Low)
CVE-2023-5996	Use after free in WebAudio in Google Chrome prior to 119.0.6045.123 allowed a remote attacker to potentially exp page. (Chromium security severity: High)
CVE-2023-5997	Use after free in Garbage Collection in Google Chrome prior to 119.0.6045.159 allowed a remote attacker to poten HTML page. (Chromium security severity: High)
CVE-2023-6112	Use after free in Navigation in Google Chrome prior to 119.0.6045.159 allowed a remote attacker to potentially exp page. (Chromium security severity: High)
CVE-2023-6345	Integer overflow in Skia in Google Chrome prior to 119.0.6045.199 allowed a remote attacker who had compromis a sandbox escape via a malicious file. (Chromium security severity: High)
CVE-2023-6346	Use after free in WebAudio in Google Chrome prior to 119.0.6045.199 allowed a remote attacker to potentially exp page. (Chromium security severity: High)
CVE-2023-6347	Use after free in Mojo in Google Chrome prior to 119.0.6045.199 allowed a remote attacker to potentially exploit h (Chromium security severity: High)
CVE-2023-6348	Type Confusion in Spellcheck in Google Chrome prior to 119.0.6045.199 allowed a remote attacker who had comp exploit heap corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2023-6350	Use after free in libavif in Google Chrome prior to 119.0.6045.199 allowed a remote attacker to potentially exploit (Chromium security severity: High)
CVE-2023-6351	Use after free in libavif in Google Chrome prior to 119.0.6045.199 allowed a remote attacker to potentially exploit (Chromium security severity: High)
CVE-2023-6508	Use after free in Media Stream in Google Chrome prior to 120.0.6099.62 allowed a remote attacker to potentially e page. (Chromium security severity: High)
CVE-2023-6509	Use after free in Side Panel Search in Google Chrome prior to 120.0.6099.62 allowed a remote attacker who convi interaction to potentially exploit heap corruption via specific UI interaction. (Chromium security severity: High)
CVE-2023-6510	Use after free in Media Capture in Google Chrome prior to 120.0.6099.62 allowed a remote attacker who convince potentially exploit heap corruption via specific UI interaction. (Chromium security severity: Medium)
CVE-2023-6511	Inappropriate implementation in Autofill in Google Chrome prior to 120.0.6099.62 allowed a remote attacker to by page. (Chromium security severity: Low)
CVE-2023-6512	Inappropriate implementation in Web Browser UI in Google Chrome prior to 120.0.6099.62 allowed a remote attac iframe dialog context menu via a crafted HTML page. (Chromium security severity: Low)
CVE-2023-6702	Type confusion in V8 in Google Chrome prior to 120.0.6099.109 allowed a remote attacker to potentially exploit h (Chromium security severity: High)
CVE-2023-6703	Use after free in Blink in Google Chrome prior to 120.0.6099.109 allowed a remote attacker to potentially exploit h (Chromium security severity: High)
CVE-2023-6704	Use after free in libavif in Google Chrome prior to 120.0.6099.109 allowed a remote attacker to potentially exploit (Chromium security severity: High)
CVE-2023-6705	Use after free in WebRTC in Google Chrome prior to 120.0.6099.109 allowed a remote attacker to potentially expl (Chromium security severity: High)
CVE-2023-6706	Use after free in FedCM in Google Chrome prior to 120.0.6099.109 allowed a remote attacker who convinced a us potentially exploit heap corruption via a crafted HTML page. (Chromium security severity: High)
CVE-2023-6707	Use after free in CSS in Google Chrome prior to 120.0.6099.109 allowed a remote attacker to potentially exploit he (Chromium security severity: Medium)
CVE-2024-0406	A flaw was discovered in the mholt/archiver package. This flaw allows an attacker to create a specially crafted tar f to restricted files or directories. This issue can allow the creation or overwriting of files with the user's or applicatio
CVE-2024-0853	curl inadvertently kept the SSL session ID for connections in its cache even when the verify status (*OCSP staplin same hostname could then succeed if the session ID cache was still fresh, which then skipped the verify status chec
CVE-2024-2004	When a protocol selection parameter option disables all protocols without adding any then the default set of protoc error in the logic for removing protocols. The below command would perform a request to curl.se with a plaintext p curl --proto -all,-http http://curl.se The flaw is only present if the set of selected protocols disables the entire set of no practical use and therefore unlikely to be encountered in real situations. The curl security team has thus assessed

CVE-2024-23334	aiohttp is an asynchronous HTTP client/server framework for asyncio and Python. When using aiohttp as a web ser necessary to specify the root path for static files. Additionally, the option 'follow_symlinks' can be used to determi outside the static root directory. When 'follow_symlinks' is set to True, there is no validation to check if reading a f lead to directory traversal vulnerabilities, resulting in unauthorized access to arbitrary files on the system, even wh follow_symlinks and using a reverse proxy are encouraged mitigations. Version 3.9.2 fixes this issue.
CVE-2024-23829	aiohttp is an asynchronous HTTP client/server framework for asyncio and Python. Security-sensitive parts of the P in allowable character sets, that must trigger error handling to robustly match frame boundaries of proxies in order requests. Additionally, validation could trigger exceptions that were not handled consistently with processing of ot internet standards require could, depending on deployment environment, assist in request smuggling. The unhandle consumption on the application server and/or its logging facilities. This vulnerability exists due to an incomplete fi this vulnerability.
CVE-2024-2466	libcurl did not check the server certificate of TLS connections done to a host specified as an IP address, when built avoid using the set hostname function when the specified hostname was given as an IP address, therefore complete all uses of TLS protocols (HTTPS, FTPS, IMAPS, POPS3, SMTPS, etc).
CVE-2024-24788	A malformed DNS message in response to a query can cause the Lookup functions to get stuck in an infinite loop.
CVE-2024-24790	The various Is methods (IsPrivate, IsLoopback, etc) did not work as expected for IPv4-mapped IPv6 addresses, retu true in their traditional IPv4 forms.
CVE-2024-24791	The net/http HTTP/1.1 client mishandled the case where a server responds to a request with an "Expect: 100-contin higher) status. This mishandling could leave a client connection in an invalid state, where the next request sent on t a request to a net/http/httputil.ReverseProxy proxy can exploit this mishandling to cause a denial of service by send elicit a non-informational response from the backend. Each such request leaves the proxy with an invalid connectio that connection to fail.
CVE-2024-25128	Flask-AppBuilder is an application development framework, built on top of Flask. When Flask-AppBuilder is set t an attacker to forge an HTTP request, that could deceive the backend into using any requested OpenID service. Thi unauthorised privilege access if a custom OpenID service is deployed by the attacker and accessible by the backend the application is using the OpenID 2.0 authorization protocol. Upgrade to Flask-AppBuilder 4.3.11 to fix the vuln
CVE-2024-27289	pgx is a PostgreSQL driver and toolkit for Go. Prior to version 4.18.2, SQL injection can occur when all of the foll simple protocol is used; a placeholder for a numeric value must be immediately preceded by a minus; there must be the first placeholder; both must be on the same line; and both parameter values must be user-controlled. The proble do not use the simple protocol or do not place a minus directly before a placeholder.
CVE-2024-27304	pgx is a PostgreSQL driver and toolkit for Go. SQL injection can occur if an attacker can cause a single query or bi integer overflow in the calculated message size can cause the one large message to be sent as multiple messages un resolved in v4.18.2 and v5.5.4. As a workaround, reject user input large enough to cause a single query or bind mes
CVE-2024-27306	aiohttp is an asynchronous HTTP client/server framework for asyncio and Python. A XSS vulnerability exists on in This vulnerability is fixed in 3.9.4. We have always recommended using a reverse proxy server (e.g. nginx) for serv recommendation are unaffected. Other users can disable `show_index` if unable to upgrade.
CVE-2024-28757	libexpat through 2.6.1 allows an XML Entity Expansion attack when there is isolated use of external parsers (creat
CVE-2024-30251	aiohttp is an asynchronous HTTP client/server framework for asyncio and Python. In affected versions an attacker form-data) request. When the aiohttp server processes it, the server will enter an infinite loop and be unable to proc the application from serving requests after sending a single request. This issue has been addressed in version 3.9.4. upgrade may manually apply a patch to their systems. Please see the linked GHSA for instructions.
CVE-2024-37568	lepture Authlib before 1.3.1 has algorithm confusion with asymmetric public keys. Unless an algorithm is specified allowed with any asymmetric public key. (This is similar to CVE-2022-29217 and CVE-2024-33663.)
CVE-2024-39705	NLTK through 3.8.1 allows remote code execution if untrusted packages have pickled Python code, and the integra used. This affects, for example, averaged_perceptron_tagger and punkt.
CVE-2024-43796	Express.js minimalist web framework for node. In express < 4.20.0, passing untrusted user input - even after sanitiz untrusted code. This issue is patched in express 4.20.0.
CVE-2024-6197	libcurl's ASN1 parser has this utf8asn1str() function used for parsing an ASN.1 UTF-8 string. Itcan detect an invali doing so it also invokes `free()` on a 4 byte localstack buffer. Most modern malloc implementations detect this erro accept the input pointer and add that memory to its list of available chunks. This leads to the overwriting of nearby is decided by the `free()` implementation; likely to be memory pointers and a set of flags. The most likely outcome cannot be ruled out that more serious results can be had in special circumstances.
CVE-2024-6874	libcurl's URL API function [curl_url_get()](https://curl.se/libcurl/c/curl_url_get.html) offers punycode conversions that is exactly 256 bytes, libcurl ends up reading outside of a stack based buffer when built to use the *macidn* IDI up the provided buffer exactly - but does not null terminate the string. This flaw can lead to stack contents accident string.

CVE-2024-8986	The grafana plugin SDK bundles build metadata into the binaries it compiles; this metadata includes the repository by running `git remote get-url origin`. If credentials are included in the repository URI (for instance, to allow for fe binary will contain the full URI, including said credentials.
CVE-2024-9355	A vulnerability was found in Golang FIPS OpenSSL. This flaw allows a malicious user to randomly cause an unini buffer to be returned in FIPS mode. It may also be possible to force a false positive match between non-equal hash sum to an untrusted input sum if an attacker can send a zeroed buffer in place of a pre-computed sum.¬† It is also p instead of an unpredictable value.¬† This may have follow-on implications for the Go TLS stack.
DLA-3867-1	git - security update
DLA-3878-1	libxml2 - security update
DLA-3881-1	aom - security update
DSA-5746-1	postgresql-13 - security update
GHSA-5cpq-8wj7-hf2v	pyca/cryptography's wheels include a statically linked copy of OpenSSL. The versions of OpenSSL included in cry security issue. More details about the vulnerability itself can be found in https://www.openssl.org/news/secadv/202 source ("sdist") then you are responsible for upgrading your copy of OpenSSL. Only users installing from wheels b distributed on PyPI) need to update their cryptography versions.
GHSA-7jwh-3vrq-q3m8	### Impact SQL injection can occur if an attacker can cause a single query or bind message to exceed 4 GB in size message size can cause the one large message to be sent as multiple messages under the attacker's control. ### Patc Workarounds Reject user input large enough to cause a single query or bind message to exceed 4 GB in size.
GHSA-9763-4f94-gfch	### Impact On some platforms, when an attacker can time decapsulation of Kyber on forged cipher texts, they coul Does not apply to ephemeral usage, such as when used in the regular way in TLS. ### Patches Patched in 1.3.7. ## kyberslash.cr.yp.to/)
GHSA-mh55-gqvf-xfwm	Middleware causes a prohibitive amount of heap allocations when processing malicious preflight requests that inclu (ACRH) header whose value contains many commas. This behavior can be abused by attackers to produce undue lo cause a denial of service.
GHSA-mhpq-9638-x6pw	An attacker controlled input of a PBES2 encrypted JWE blob can have a very large p2c value that, when decrypted
GHSA-pjjw-qhg8-p2p9	### Summary llhttp 8.1.1 is vulnerable to two request smuggling vulnerabilities. Details have not been disclosed yo The issue is resolved by using llhttp 9+ (which is included in aiohttp 3.8.6+).
RHBA-2024:5691	The ca-certificates package contains a set of Certificate Authority (CA) certificates chosen by the Mozilla Foundati Infrastructure (PKI).
RHSA-2023:0050	Node.js is a software development platform for building fast and scalable network applications in the JavaScript pr
RHSA-2023:0095	The libtiff packages contain a library of functions for manipulating Tagged Image File Format (TIFF) files.
RHSA-2023:0096	D-Bus is a system for sending messages between applications. It is used both for the system-wide message bus serv messaging facility.
RHSA-2023:0100	The systemd packages contain systemd, a system and service manager for Linux, compatible with the SysV and LS parallelism capabilities, uses socket and D-Bus activation for starting services, offers on-demand starting of daemo Linux cgroups. In addition, it supports snapshotting and restoring of the system state, maintains mount and automo transactional dependency-based service control logic. It can also work as a drop-in replacement for sysvinit.
RHSA-2023:0103	Expat is a C library for parsing XML documents.
RHSA-2023:0116	A library that provides Abstract Syntax Notation One (ASN.1, as specified by the X.680 ITU-T recommendation) p Distinguished Encoding Rules (DER, as per X.690) encoding and decoding functions.
RHSA-2023:0173	The libxml2 library is a development toolbox providing the implementation of various XML standards.
RHSA-2023:0379	X.Org X11 libXpm runtime library.
RHSA-2023:0610	Git is a distributed revision control system with a decentralized architecture. As opposed to centralized version con ensures that each working copy of a Git repository is an exact copy with complete revision history. This not only al projects without the need to have permission to push the changes to their official repositories, but also makes it pos connection.
RHSA-2023:0625	KSBA (pronounced Kasbah) is a library to make X.509 certificates as well as the CMS easily accessible by other a blocks of S/MIME and TLS.
RHSA-2023:0833	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2023:0835	The python-setuptools package provides a collection of enhancements to Python distribution utilities allowing conv packages.

RHSA-2023:0837	The systemd packages contain systemd, a system and service manager for Linux, compatible with the SysV and LS parallelism capabilities, uses socket and D-Bus activation for starting services, offers on-demand starting of daemo Linux cgroups. In addition, it supports snapshotting and restoring of the system state, maintains mount and automo transactional dependency-based service control logic. It can also work as a drop-in replacement for sysvinit.
RHSA-2023:0842	The GNU tar program can save multiple files in an archive and restore files from an archive.
RHSA-2023:0852	The httpd packages provide the Apache HTTP Server, a powerful, efficient, and extensible web server.
RHSA-2023:1405	OpenSSL is a toolkit that implements the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protoco cryptography library.
RHSA-2023:1569	The gnutls packages provide the GNU Transport Layer Security (GnuTLS) library, which implements cryptograph TLS, and DTLS.
RHSA-2023:1673	The httpd packages provide the Apache HTTP Server, a powerful, efficient, and extensible web server.
RHSA-2023:1743	Node.js is a software development platform for building fast and scalable network applications in the JavaScript pr
RHSA-2023:1930	GNU Emacs is a powerful, customizable, self-documenting text editor. It provides special code editing features, a s to read e-mail and news.
RHSA-2023:2076	The libwebp packages provide a library and tools for the WebP graphics format. WebP is an image format with a lo images. WebP consists of a codec based on the VP8 format, and a container based on the Resource Interchange Fil developers and browser developers can use WebP to compress, archive, and distribute digital images more efficien
RHSA-2023:2763	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2023:2859	Git is a distributed revision control system with a decentralized architecture. As opposed to centralized version con ensures that each working copy of a Git repository is an exact copy with complete revision history. This not only al projects without the need to have permission to push the changes to their official repositories, but also makes it pos connection.
RHSA-2023:2883	The libtiff packages contain a library of functions for manipulating Tagged Image File Format (TIFF) files.
RHSA-2023:2951	The kernel packages contain the Linux kernel, the core of any Linux operating system.
RHSA-2023:3018	The libarchive programming library can create and read several different streaming archive formats, including GNU Libarchive is used notably in the bsdtar utility, scripting language bindings such as python-libarchive, and several p
RHSA-2023:3042	GNU Emacs is a powerful, customizable, self-documenting text editor. It provides special code editing features, a s to read e-mail and news.
RHSA-2023:3104	GNU Emacs is a powerful, customizable, self-documenting text editor. It provides special code editing features, a s to read e-mail and news.
RHSA-2023:3109	The Apache Portable Runtime (APR) is a portability library used by the
RHSA-2023:3246	Git is a distributed revision control system with a decentralized architecture. As opposed to centralized version con ensures that each working copy of a Git repository is an exact copy with complete revision history. This not only al projects without the need to have permission to push the changes to their official repositories, but also makes it pos connection.
RHSA-2023:3781	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2023:3827	The libtiff packages contain a library of functions for manipulating Tagged Image File Format (TIFF) files.
RHSA-2023:5050	The httpd packages provide the Apache HTTP Server, a powerful, efficient, and extensible web server.
RHSA-2023:5244	The kernel packages contain the Linux kernel, the core of any Linux operating system.
RHSA-2023:5309	The libwebp packages provide a library and tools for the WebP graphics format. WebP is an image format with a lo images. WebP consists of a codec based on the VP8 format, and a container based on the Resource Interchange Fil developers and browser developers can use WebP to compress, archive, and distribute digital images more efficien
RHSA-2023:5353	The libtiff packages contain a library of functions for manipulating Tagged Image File Format (TIFF) files.
RHSA-2023:6236	The binutils packages provide a collection of binary utilities for the manipulation of object code in various object fi nm, objcopy, objdump, ranlib, readelf, size, strings, strip, and addr2line utilities.
RHSA-2023:7029	The libX11 packages contain the core X11 protocol client library.
RHSA-2023:7050	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing

RHSA-2023:7077	The kernel packages contain the Linux kernel, the core of any Linux operating system.
RHSA-2023:7165	The Common UNIX Printing System (CUPS) provides a portable printing layer for Linux, UNIX, and similar oper
RHSA-2023:7190	Avahi is an implementation of the DNS Service Discovery and Multicast DNS specifications for Zero Configuratio on a local network. Avahi and Avahi-aware applications allow you to plug your computer into a network and, with with, view printers to print with, and find shared files on other computers.
RHSA-2023:7549	The kernel packages contain the Linux kernel, the core of any Linux operating system.
RHSA-2023:7836	Avahi is an implementation of the DNS Service Discovery and Multicast DNS specifications for Zero Configuratio on a local network. Avahi and Avahi-aware applications allow you to plug your computer into a network and, with with, view printers to print with, and find shared files on other computers.
RHSA-2024:0105	Network Security Services (NSS) is a set of libraries designed to support the cross-platform development of securit
RHSA-2024:0113	The kernel packages contain the Linux kernel, the core of any Linux operating system.
RHSA-2024:0265	The java-1.8.0-openjdk packages provide the OpenJDK 8 Java Runtime Environment and the OpenJDK 8 Java Sof
RHSA-2024:0965	The unbound packages provide a validating, recursive, and caching DNS or DNSSEC resolver.
RHSA-2024:1751	The unbound packages provide a validating, recursive, and caching DNS or DNSSEC resolver.
RHSA-2024:1786	The httpd packages provide the Apache HTTP Server, a powerful, efficient, and extensible web server.
RHSA-2024:1818	The java-1.8.0-openjdk packages provide the OpenJDK 8 Java Runtime Environment and the OpenJDK 8 Java Sof
RHSA-2024:2973	The libX11 packages contain the core X11 protocol client library.
RHSA-2024:2974	X.Org X11 libXpm runtime library.
RHSA-2024:3059	The libtiff packages contain a library of functions for manipulating Tagged Image File Format (TIFF) files.
RHSA-2024:3121	The httpd packages provide the Apache HTTP Server, a powerful, efficient, and extensible web server.
RHSA-2024:3138	The kernel packages contain the Linux kernel, the core of any Linux operating system.
RHSA-2024:3618	The kernel packages contain the Linux kernel, the core of any Linux operating system.
RHSA-2024:5529	The curl packages provide the libcurl library and the curl utility for downloading files from servers using various p
RHSA-2024:6166	Kerberos is a network authentication system, which can improve the security of your network by eliminating the in network in unencrypted form. It allows clients and servers to authenticate to each other with the help of a trusted th (KDC).
RHSA-2024:6969	The container-tools module contains tools for working with containers, notably podman, buildah, skopeo, and runc
RHSA-2024:7135	Git Large File Storage (LFS) replaces large files such as audio samples, videos, datasets, and graphics with text poi contents on a remote server.

Service packs

The service pack issues that have been shipped for Cloudera Data Services on premises 1.5.4-SP1.

Cloudera Data Services on premises 1.5.4-SP1

The service packs for new features, known issues, and fixed issues for 1.5.4-SP1.



Note: ECS Customers: Direct upgrade path is not available for customers currently on Cloudera Data Services on premises 1.5.2. Customers must upgrade to Cloudera Data Services on premises 1.5.4 prior to consuming any Service Packs (SPs) built on top of 1.5.4.



Note:

OCP Customers: Direct upgrade path is available. Customers can directly upgrade from Cloudera Data Services on premises 1.5.2 to any 1.5.4 Cumulative Hotfixes (CHFs) or Service Packs (SPs).

Note: Cloudera Data Services on premises 1.5.4 is not supported on Cloudera Base on premises 7.3.1. You must not install or upgrade to Cloudera Base on premises 7.3.1, if you are using Cloudera Data Services on premises 1.5.4 on your cluster as it is incompatible.

Certifications in 1.5.4-SP1

The following are the certifications supported in 1.5.4-SP1:

- Cloudera Base on premises 7.1.9 SP1, 7.1.9 CHF7 (7.1.9.14), 7.1.7 SP3 (7.1.7.3000)
- Cloudera Manager 7.11.3 CHF 11 and later
- Iceberg v2 GA on Cloudera Data Warehouse, Cloudera Data Engineering, & Cloudera AI with Ozone
- OEL (RHCK Kernel Only) 8.7, 8.8, 8.9, 8.10, 9.1, 9.2, 9.3, 9.4
- RHEL 8.7, 8.8, 8.9. 9.1, 9.2, 9.3, 9.4
- K8s 1.29 and OCP 4.16
- Longhorn version 1.6.2

What's new in Cloudera Data Services on premises 1.5.4-SP1

New features introduced in this service pack release of Cloudera Data Services on premises 1.5.4-SP1.

SAML authentication support

SAML authentication features:

- Support for SAML-based Single Sign-On (SSO), enabling seamless and secure access to Cloudera on premises.
- Administrators can integrate their IdP's with CDP for authentication.
- Administrators can migrate their existing users in CDP from LDAP to SAML.

Minor, Major, and Mixed OS upgrade on Cloudera Embedded Container Service hosts

- After installing Cloudera Data Services on premises on a particular RHEL OS, you can now upgrade RHEL OS to a new major or minor version. For example, you can upgrade from RHEL 7.x to RHEL 9.x or you can upgrade from RHEL 8.6 to RHEL 8.8.
- Additionally, Cloudera now supports mixed OS versions in the same cluster. For example, in a 10 node cluster, you can upgrade any number of hosts from RHEL 7.x to RHEL 9.x major version and keep the other hosts, running RHEL 8.x.

Rotate Cloudera Embedded Container Service internal certificates

Certificate rotation for Cloudera Embedded Container Service internal certificates, namely, vault, embedded database, ECS webhook, and Ingress Controller (if using the default certificate), can now be performed using the new action under Cloudera Embedded Container Service Service in Cloudera Manager UI.

Configurable Docker registry port

The port of the docker embedded registry is now configurable at install time. The default value is 5000.

Note: The port configuration cannot be changed once installation is complete and must remain unaltered during upgrades.

Note: Cloudera Manager 7.11.3 CHF11 (version: 7.11.3.28) support Cloudera Data Services on premises 1.5.4 SP1 release.Cloudera Manager 7.11.3 CHF8 does not support any Cloudera Data Services on premises release.

Kubernetes Ingress NGINX Controller vulnerabilities

Kubernetes Ingress NGINX Controller vulnerabilities.

Five vulnerabilities affecting the Ingress Nginx Controller for Kubernetes were publicly disclosed on March 24, 2025, and were given the nickname "IngressNightmare."

The ‘IngressNightmare’ vulnerabilities may allow Remote Code Execution (RCE) and potentially expose Kubernetes clusters to malicious configuration modifications. Exploitation requires specially crafted HTTP requests that bypass security measures, such as a Web Application Firewall (WAF). Successful exploitation may lead to complete cluster compromise, data exfiltration, and denial of service.

Details of the CVEs:

- CVE-2025-1974 (CVSS score: 9.8) – An unauthenticated attacker with access to the pod network can achieve arbitrary code execution in the context of the ingress-nginx controller under certain conditions
- CVE-2025-24514 (CVSS score: 8.8) – The auth-url Ingress annotation can be used to inject configuration into NGINX, resulting in arbitrary code execution in the context of the ingress-nginx controller and disclosure of secrets accessible to the controller
- CVE-2025-1097 (CVSS score: 8.8) – The auth-tls-match-cn Ingress annotation can be used to inject configuration into NGINX, resulting in arbitrary code execution in the context of the ingress-nginx controller and disclosure of secrets accessible to the controller
- CVE-2025-1098 (CVSS score: 8.8) – The mirror-target and mirror-host Ingress annotations can be used to inject arbitrary configuration into NGINX, resulting in arbitrary code execution in the context of the ingress-nginx controller and disclosure of secrets accessible to the controller
- CVE-2025-24513 (CVSS score: 4.8) – An improper input validation vulnerability that could result in directory traversal within the container, leading to denial-of-service (DoS) or limited disclosure of secret objects from the cluster when combined with other vulnerabilities

Note: Cloudera will not provide remediation options for unsupported versions, and has not tested mitigations on unsupported versions. Customers are advised to upgrade to a supported product version. For more information, refer to the Support Lifecycle Policy page.

Action required

Multiple mitigation scenarios are possible depending on the version of Cloudera Data Services deployed.

1. Check your Data Services version.

   To determine which version of Data Services is running, log into Cloudera Manager UI, click “Data Services” in the left sidebar, and find each Data Services cluster's version(s).
2. Check your Cloudera Manager version.

   To determine which version of Cloudera Manager is running, log into Cloudera Manager UI and click on the version number at the lower-left corner of the home page. The Cloudera Manager 7.11.3 CHF8 Build number shows 56304673.

Mitigation steps for Cloudera Data Services On Premises:

1. No action is needed in the following scenarios:
   - Running Data Services on OpenShift Container Platform (OCP) and using the default HAproxy ingress controller. If you have customized OCP to use NGINX as the ingress controller, please contact RedHat support for further information
   - Running Data Services 1.5.4 or later with Cloudera Manager 7.11.3 CHF8 or later. These versions already disable the affected admission webhook component.
2. Data Services before 1.5.4, or Cloudera Manager prior to 7.11.3 CHF8, perform one of the following:
   - (Recommended) Upgrade Cloudera Manager to version 7.13.1 CHF1 or higher and all Embedded Container Service clusters to 1.5.4 SP2 or higher. This will not upgrade the nginx version but will automatically disable the webhook via Cloudera Manager, which mitigates the CVEs.
   - If an upgrade is not possible, follow the below steps to disable the webhook to mitigate the CVEs manually.

     Note: These steps should be run on each ECS cluster. If you create a new ECS cluster while on an affected release, reapply these steps on the new cluster(s).

1. SSH into the ECS server host.

2. Create the following YAML file to patch the built-in Nginx Ingress Controller Helm Chart:

```
apiVersion: helm.cattle.io/v1
kind: HelmChartConfig
metadata:
 name: rke2-ingress-nginx
 namespace: kube-system
spec:
 valuesContent: |-
  controller:
   admissionWebhooks:
    enabled: false
```

3. Run the following commands to apply the patch:

```
alias kubectl=/var/lib/rancher/rke2/bin/kubectl
export KUBECONFIG=/etc/rancher/rke2/rke2.yaml
kubectl apply -f <PATCH_FILENAME>.yaml
```

Note: No cluster restart or ECS refresh is required.

Note: Cloudera will not provide remediation options for unsupported versions, and has not tested mitigations on unsupported versions. Customers are advised to upgrade to a supported product version. For more information, refer to the Support Lifecycle Policy page

Cloudera Data Services on premises version 1.5.4 SP2 CHF1 will contain fixes to upgrade the Ingress NGINX controller. Target release date: April 25. 2025.

For the latest update on this issue, see the corresponding Knowledge article: TSB 2025-839: Critical Kubernetes Ingress NGINX Controller Vulnerability Allows RCE Without Authentication

For the latest update on this issue, see the corresponding Data Services Knowledge article: https://my.cloudera.com/knowledge/TSB-2025-839-Mitigation-steps-for-Cloudera-Data-Services-on?id=405006

Known Issues in Cloudera Data Services on premises 1.5.4-SP1

The following are the new known issues in the 1.5.4 service pack SP1 release of Cloudera Data Services on premises.

CDPVC-1739 - Rollback workload username generation changes to pre 1.5.4-SP1 state in Cloudera Data Services on premises

The current system for generating workload usernames in CDE, which shares code with PbC, automatically converts usernames to lowercase. Additionally, it removes special characters/spaces and adds prefixes if the IdP (Identity Provider) user ID starts with a number.

This automated transformation creates a significant issue: Ranger, a critical component for authorization, requires the workload username to exactly match the IdP user ID. If an IdP user ID from LDAP/SAML contains uppercase or mixed-case characters, special characters, spaces, or starts with a number, the CDE transformation will result in a mismatch. This mismatch prevents Ranger from properly recognizing and authorizing users, leading to functionality issues.

To resolve this discrepancy and ensure Ranger functions correctly, the affected workload usernames in the UMS (User Management Service) database must be manually updated.

There are two methods to achieve this:

1. Manually Update the Usernames in the UMS Database.

   Run the following SQL command to update the workload username:

```
update ums_actors set workload_username = '123456' where wor
kload_username =
```

```
              'u_123456'
```

2. For a more comprehensive update of all affected workload usernames in the ECS - UMS database, utilize the provided script: rename_workload_username_runbook.md.

OBS-6044 - Warning alert in the Cloudera Embedded Container Service Health Test status when a cluster is restarted for stability execution

The following warning alert is shown in the Cloudera Embedded Container Service Health Test status when a cluster is restarted in Cloudera Manager for stability execution. Prometheus has issue s compacting blocks

This issue occurs when WAL (Write Ahead Logs) are corrupted.

1. Run the following command to access the Prometheus container's shell:

```
kubectl exec -i -t -n <prometheus server namespace> <prometh
eus server pod name> -c
          <prometheus server container name> -- sh -c "(bash
 || ash || sh)"
```

2. Change the current working directory to the WAL directory of Prometheus.
   a. For Infrastructure Prometheus: The WAL directory location is /Prometheus/wal. For example:

      cd /prometheus/wal

   b. For control plane/environment: The Prometheus directory location is /data/wal. For example:

      cd /data/wal

3. Note the corrupted segment from Prometheus's pod logs. Example logs:

```
21T09:00:07.036Z caller=db.go:1074 level=error component=tsd
b msg="compaction failed" err="WAL truncation in Compact: cr
eate checkpoint: read segments: corruption in segment
          /prometheus/wal/00000026 at 10978: unexpected full r
ecord"
```

4. Skip the compression of the corrupted segment by moving the checkpoint. This requires renaming the checkpoint folder in the WAL directory. For example, if the corrupted segment is 00000026 and the current checkpoint folder name is checkpoint.00000020, then rename the checkpoint folder to checkpoint.00000027. For example:

   mv checkpoint.00000020 checkpoint.00000027

OPSX-5810 - Cloudera Control Plane on premises installation fails at the vault initialization phase due to longhorn-manager pods

At times, longhorn-manager pods will fail to come up with repeating error messages like:

```
level=error msg="Failed to save TLS secret for longhorn-system/
longhorn-webhook-tls: Operation cannot be fulfilled on secrets
 \"longhorn-webhook-tls\": the object has been modified; please
 apply your changes to the latest version and try again"
```

This causes the Longhorn nodes to remain in a `NotReady` state, stopping volumes from successfully being created/attached.

The following steps can be taken on an ECS Server node to fix the issue:

1. Stop the Longhorn Manager daemonset by executing following command:

```
kubectl -n longhorn-system patch daemonset longhorn-manager -p
 '{"spec": {"template": {"spec": {"nodeSelector": {"non-existi
ng": "true"}}}}}'
```

2. Delete the Longhorn Webhook TLS secret by executing the following command:

```
kubectl delete secret longhorn-webhook-tls -n longhorn-system
```

3. Start the Longhorn Manager daemonset by executing the following command:

```
kubectl -n longhorn-system patch daemonset longhorn-manager --
type json -p='[{"op": "remove", "path": "/spec/template/spec/n
odeSelector/non-existing"}]'
```

OPSX-5403 - Typecasting fails when truststore password is integer

The truststore_password in the SCM configuration should not be an integer for Private Cloud installation.

Update truststore_password in the SCM configuration to a non-integer value.

OPSX-4684 - Start Cloudera Embedded Container Service command shows finished successfully even though start docker server failed on one of the hosts

Docker service starts with one or more docker roles failed to start because the corresponding host is unhealthy.

Make sure the host is healthy. Start the docker role in the host.

OPSX-4391 - External docker cert not base64 encoded

When using Cloudera Data Services on premises on Cloudera Embedded Container Service, in some rare situations, the CA certificate for the Docker registry in the cdp namespace is incorrectly encoded, resulting in TLS errors when connecting to the Docker registry.

Compare and edit the contents of the "cdp-private-installer-docker-cert" secret in the cdp namespace so that it matches the contents of the "cdp-private-installer-docker-cert" secret in other namespaces. The secrets and their corresponding namespaces can be identified using the command "kubectl get secret -A | grep cdp-private-installer-docker-cert". Inspect each secret using the command "kubectl get secret -n cdp cdp-private-installer-docker-cert -o yaml", replacing "cdp" with the different namespace names. If necessary, modify the secret in the cdp namespace using the command "kubectl edit secret -n cdp cdp-private-installer-docker-cert"

OPSX-3323 - Custom Log redaction does not work for JSON files in diag bundles

The JSON files within the diag bundle will not be redacted.

No workaround available.

OPSX-2772 - For Account Administrator user, update roles functionality should be disabled

When a user with administrative privileges accesses the User Management > Update Roles page in the Cloudera Management Console, the user is presented with options to select various roles. Selecting or deselecting these roles does not change this user's privileges -- an administrative user, by default, has all privileges, and those privileges cannot be changed.

No workaround available.

Fixed Issues in Cloudera Data Services on premises 1.5.4-SP1

The fixes in this service pack release of Cloudera Data Services on premises 1.5.4-SP1.

OBS-4176 - Prometheus reports duplicate metric alert for Kubernetes state metrics

After installing Cloudera Data Services on premises 1.5.4, the EnvPrometheusDuplicateTimestamps warning message appears on the Control Plane Monitoring dashboard.

OPSX-5511 - Prevent Auto-Retry on Longhorn Helm Install/Upgrade Failures

When longhorn install or upgrade fails, auto retry is disabled to prevent an unexpected uninstalling. In such cases, manual intervention is required to:

- Roll back to a previous version.

- Reapply the Longhorn upgrade step.

OPSX-5533 - DWX homepage URL fails with "504 Gateway Time-out" error

The homepage intermittently fails to load the list of Virtual Warehouses with the time-out error. The issue has been fixed by optimizing the API calls within the system.

Repository Locations for 1.5.4-SP1

The URLs for Cloudera Data Services on premises 1.5.4-SP1 are listed in the following table:

URL Type	Repository Location
Index	`https://username:password@archive.cloudera.com/p/cdp-pvc-ds/1.5.4-h5/`
Manifest	Repository: `https://username:password@archive.cloudera.com/p/cdp-pvc-ds/1.5.4-h5/manifest.json`
Parcels	Repository: `https://username:password@archive.cloudera.com/p/cdp-pvc-ds/1.5.4-h5/parcels/`

Fixed Common Vulnerabilities and Exposures in 1.5.4 SP1

Review the Common vulnerabilities and Exposures (CVEs) that were fixed in 1.5.4 SP1 release of Cloudera Data Services on premises.

Issue ID	Description
CVE-2005-0406	A design flaw in image processing software that modifies JPEG images might not modify the original EXIF thumb of potentially sensitive visual information that had been removed from the main JPEG image.
CVE-2007-1420	MySQL 5.x before 5.0.36 allows local users to cause a denial of service (database crash) by performing informatio BY to sort a single-row result, which prevents certain structure elements from being initialized and triggers a NUL
CVE-2007-2243	OpenSSH 4.6 and earlier, when ChallengeResponseAuthentication is enabled, allows remote attackers to determine to authenticate via S/KEY, which displays a different response if the user account exists, a similar issue to CVE-20
CVE-2007-2691	MySQL before 4.1.23, 5.0.x before 5.0.42, and 5.1.x before 5.1.18 does not require the DROP privilege for RENA authenticated users to rename arbitrary tables.
CVE-2007-2768	OpenSSH, when using OPIE (One-Time Passwords in Everything) for PAM, allows remote attackers to determine displays a different response if the user account exists and is configured to use one-time passwords (OTP), a simila
CVE-2007-3476	Array index error in gd_gif_in.c in the GD Graphics Library (libgd) before 2.0.35 allows user-assisted remote attac heap corruption) via large color index values in crafted image data, which results in a segmentation fault.
CVE-2007-3477	The (a) imagearc and (b) imagefilledarc functions in GD Graphics Library (libgd) before 2.0.35 allow attackers to via a large (1) start or (2) end angle degree value.
CVE-2007-3996	Multiple integer overflows in libgd in PHP before 5.2.4 allow remote attackers to cause a denial of service (applica code via a large (1) srcW or (2) srcH value to the (a) gdImageCopyResized function, or a large (3) sy (height) or (4 or the (c) gdImageCreateTrueColor function.
CVE-2007-5686	initscripts in rPath Linux 1 sets insecure permissions for the /var/log/btmp file, which allows local users to obtain s attempts. NOTE: because sshd detects the insecure permissions and does not log certain events, this also prevents s attempts by remote attackers.
CVE-2007-5686	initscripts in rPath Linux 1 sets insecure permissions for the /var/log/btmp file, which allows local users to obtain s attempts. NOTE: because sshd detects the insecure permissions and does not log certain events, this also prevents s attempts by remote attackers.

CVE-2008-3134	Multiple unspecified vulnerabilities in GraphicsMagick before 1.2.4 allow remote attackers to cause a denial of ser consumption) via (a) unspecified vectors in the (1) AVI, (2) AVS, (3) DCM, (4) EPT, (5) FITS, (6) MTV, (7) PAL and (b) the GetImageCharacteristics function in magick/image.c, as reachable from a crafted (10) PNG, (11) JPEG,
CVE-2008-3234	sshd in OpenSSH 4 on Debian GNU/Linux, and the 20070303 OpenSSH snapshot, allows remote authenticated use by appending a :/ (colon slash) sequence, followed by the role name, to the username.
CVE-2008-4677	autoload/netrw.vim (aka the Netrw Plugin) 109, 131, and other versions before 133k for Vim 7.1.266, other 7.1 ver session, and sends those credentials when attempting to establish subsequent FTP sessions to servers on different h obtain sensitive information in opportunistic circumstances by logging usernames and passwords. NOTE: the upstr different ports on the same host, stating "I'm assuming that they're using the same id and password on that unchang
CVE-2009-0819	sql/item_xmlfunc.cc in MySQL 5.1 before 5.1.32 and 6.0 before 6.0.10 allows remote authenticated users to cause expression employing a scalar expression as a FilterExpr with ExtractValue() or UpdateXML()," which triggers an
CVE-2009-3546	The _gdGetColors function in gd_gd.c in PHP 5.2.11 and 5.3.x before 5.3.1, and the GD Graphics Library 2.x, doe structure member, which might allow remote attackers to conduct buffer overflow or buffer over-read attacks via a CVE-2009-3293. NOTE: some of these details are obtained from third party information.
CVE-2009-4028	The vio_verify_callback function in viosslfactories.c in MySQL 5.0.x before 5.0.88 and 5.1.x before 5.1.41, when the depth of X.509 certificates, which allows man-in-the-middle attackers to spoof arbitrary SSL-based MySQL se by a certificate presented by a server linked against the yaSSL library.
CVE-2010-1621	The mysql_uninstall_plugin function in sql/sql_plugin.cc in MySQL 5.1 before 5.1.46 does not check privileges be remote attackers to uninstall arbitrary plugins via the UNINSTALL PLUGIN command.
CVE-2010-1626	MySQL before 5.1.46 allows local users to delete the data and index files of another user's MyISAM table via a sy TABLE command, a different vulnerability than CVE-2008-4098 and CVE-2008-7247.
CVE-2010-3677	Oracle MySQL 5.1 before 5.1.49 and 5.0 before 5.0.92 allows remote authenticated users to cause a denial of servi that uses a table with a unique SET column.
CVE-2010-3682	Oracle MySQL 5.1 before 5.1.49 and 5.0 before 5.0.92 allows remote authenticated users to cause a denial of servi EXPLAIN with crafted "SELECT ... UNION ... ORDER BY (SELECT ... WHERE ...)" statements, which triggers Item_singlerow_subselect::store function.
CVE-2010-4651	Directory traversal vulnerability in util.c in GNU patch 2.6.1 and earlier allows user-assisted remote attackers to cr that is specified with a .. (dot dot) or full pathname, a related issue to CVE-2010-1679.
CVE-2011-3374	It was found that apt-key in apt, all versions, do not correctly validate gpg keys with the master keyring, leading to
CVE-2011-3389	The SSL protocol, as used in certain configurations in Microsoft Windows and Microsoft Internet Explorer, Mozill products, encrypts data by using CBC mode with chained initialization vectors, which allows man-in-the-middle at a blockwise chosen-boundary attack (BCBA) on an HTTPS session, in conjunction with JavaScript code that uses URLConnection API, or (3) the Silverlight WebClient API, aka a "BEAST" attack.
CVE-2011-4116	_is_safe in the File::Temp module for Perl does not properly handle symlinks.
CVE-2012-0039	GLib 2.31.8 and earlier, when the g_str_hash function is used, computes hash values without restricting the ability allows context-dependent attackers to cause a denial of service (CPU consumption) via crafted input to an applicati issue may be disputed by the vendor; the existence of the g_str_hash function is not a vulnerability in the library, b g_hash_table_new_full can specify an arbitrary hash function that is appropriate for the application.
CVE-2013-0340	expat 2.1.0 and earlier does not properly handle entities expansion unless an application developer uses the XML_S remote attackers to cause a denial of service (resource consumption), send HTTP requests to intranet servers, or rea aka an XML External Entity (XXE) issue. NOTE: it could be argued that because expat already provides the ability responsibility for resolving this issue lies with application developers; according to this argument, this entry should would need its own CVE.
CVE-2013-4392	systemd, when updating file permissions, allows local users to change the permissions and SELinux security contex unspecified files.
CVE-2014-8166	The browsing feature in the server in CUPS does not filter ANSI escape sequences from shared printer names, whi arbitrary code via a crafted printer name.
CVE-2015-2575	Unspecified vulnerability in the MySQL Connectors component in Oracle MySQL 5.1.34 and earlier allows remote and integrity via unknown vectors related to Connector/J.
CVE-2015-3276	The nss_parse_ciphers function in libraries/libldap/tls_m.c in OpenLDAP does not properly parse OpenSSL-style r might cause a weaker than intended cipher to be used and allow remote attackers to have unspecified impact via un
CVE-2016-10505	NULL pointer dereference vulnerabilities in the imagetopnm function in convert.c, sycc444_to_rgb function in col and sycc422_to_rgb function in color.c in OpenJPEG before 2.2.0 allow remote attackers to cause a denial of servi

CVE-2016-10506	Division-by-zero vulnerabilities in the functions opj_pi_next_cprl, opj_pi_next_pcrl, and opj_pi_next_rpcl in pi.c i attackers to cause a denial of service (application crash) via crafted j2k files.
CVE-2016-8678	The IsPixelMonochrome function in MagickCore/pixel-accessor.h in ImageMagick 7.0.3.0 allows remote attackers read and crash) via a crafted file. NOTE: the vendor says "This is a Q64 issue and we do not support Q64."
CVE-2016-9113	There is a NULL pointer dereference in function imagetobmp of convertbmp.c:980 of OpenJPEG 2.1.2. image->co initialization(NULL). Impact is Denial of Service.
CVE-2016-9114	There is a NULL Pointer Access in function imagetopnm of convert.c:1943(jp2) of OpenJPEG 2.1.2. image->comp initialization(NULL). Impact is Denial of Service.
CVE-2016-9115	Heap Buffer Over-read in function imagetotga of convert.c(jp2):942 in OpenJPEG 2.1.2. Impact is Denial of Servic
CVE-2016-9116	NULL Pointer Access in function imagetopnm of convert.c:2226(jp2) in OpenJPEG 2.1.2. Impact is Denial of Serv
CVE-2016-9117	NULL Pointer Access in function imagetopnm of convert.c(jp2):1289 in OpenJPEG 2.1.2. Impact is Denial of Serv
CVE-2016-9580	An integer overflow vulnerability was found in tiftoimage function in openjpeg 2.1.2, resulting in heap buffer over
CVE-2016-9581	An infinite loop vulnerability in tiftoimage that results in heap buffer overflow in convert_32s_C1P1 was found in
CVE-2017-1000382	VIM version 8.0.1187 (and other versions most likely) ignores umask when creating a swap file ("[ORIGINAL_FI be world readable or otherwise accessible in ways not intended by the user running the vi binary.
CVE-2017-1000383	GNU Emacs version 25.3.1 (and other versions most likely) ignores umask when creating a backup save file ("[OR that may be world readable or otherwise accessible in ways not intended by the user running the emacs binary.
CVE-2017-11754	The WritePICONImage function in coders/xpm.c in ImageMagick 7.0.6-4 allows remote attackers to cause a denia that is mishandled in an OpenPixelCache call.
CVE-2017-11755	The WritePICONImage function in coders/xpm.c in ImageMagick 7.0.6-4 allows remote attackers to cause a denia that is mishandled in an AcquireSemaphoreInfo call.
CVE-2017-14159	slapd in OpenLDAP 2.4.45 and earlier creates a PID file after dropping privileges to a non-root account, which mig by leveraging access to this non-root account for PID file modification before a root script executes a "kill `cat /pat openldap-initscript.
CVE-2017-14988	Header::readfrom in IlmImf/ImfHeader.cpp in OpenEXR 2.2.0 allows remote attackers to cause a denial of service file that is accessed with the ImfOpenInputFile function in IlmImf/ImfCRgbaFile.cpp. NOTE: The maintainer and vulnerability isn't valid
CVE-2017-15131	It was found that system umask policy is not being honored when creating XDG user directories, since Xsession so policy. This only affects xdg-user-dirs before 0.15.5 as shipped with Red Hat Enterprise Linux.
CVE-2017-15945	The installation scripts in the Gentoo dev-db/mysql, dev-db/mariadb, dev-db/percona-server, dev-db/mysql-cluster, 2017-09-29 have chown calls for user-writable directory trees, which allows local users to gain privileges by levera of a link.
CVE-2017-16231	In PCRE 8.41, after compiling, a pcretest load test PoC produces a crash overflow in the function match() in pcre_ third parties dispute the relevance of this report, noting that there are options that can be used to limit the amount o
CVE-2017-16232	LibTIFF 4.0.8 has multiple memory leak vulnerabilities, which allow attackers to cause a denial of service (memor tif_open.c, tif_lzw.c, and tif_aux.c. NOTE: Third parties were unable to reproduce the issue
CVE-2017-17479	In OpenJPEG 2.3.0, a stack-based buffer overflow was discovered in the pgxtoimage function in jpwl/convert.c. Th write, which may lead to remote denial of service or possibly remote code execution.
CVE-2017-17740	contrib/slapd-modules/nops/nops.c in OpenLDAP through 2.4.45, when both the nops module and the memberof o that was allocated on the stack, which allows remote attackers to cause a denial of service (slapd crash) via a memb
CVE-2017-17973	In LibTIFF 4.0.8, there is a heap-based use-after-free in the t2p_writeproc function in tiff2pdf.c. NOTE: there is a t issue
CVE-2017-18018	In GNU Coreutils through 8.29, chown-core.c in chown and chgrp does not prevent replacement of a plain file with options, which allows local users to modify the ownership of arbitrary files by leveraging a race condition.
CVE-2017-5563	LibTIFF version 4.0.7 is vulnerable to a heap-based buffer over-read in tif_lzw.c resulting in DoS or code executio
CVE-2017-7245	Stack-based buffer overflow in the pcre32_copy_substring function in pcre_get.c in libpcre1 in PCRE 8.40 allows (WRITE of size 4) or possibly have unspecified other impact via a crafted file.
CVE-2017-7246	Stack-based buffer overflow in the pcre32_copy_substring function in pcre_get.c in libpcre1 in PCRE 8.40 allows (WRITE of size 268) or possibly have unspecified other impact via a crafted file.
CVE-2017-7275	The ReadPCXImage function in coders/pcx.c in ImageMagick 7.0.4.9 allows remote attackers to cause a denial of and application crash) via a crafted file. NOTE: this vulnerability exists because of an incomplete fix for CVE-201

CVE-2017-9117	In LibTIFF 4.0.7, the program processes BMP images without verifying that biWidth and biHeight in the bitmap-in leading to a heap-based buffer over-read in bmp2tiff.
CVE-2017-9937	In LibTIFF 4.0.8, there is a memory malloc failure in tif_jbig.c. A crafted TIFF document can lead to an abort resu
CVE-2018-0734	The OpenSSL DSA signature algorithm has been shown to be vulnerable to a timing side channel attack. An attack algorithm to recover the private key. Fixed in OpenSSL 1.1.1a (Affected 1.1.1). Fixed in OpenSSL 1.1.0j (Affected (Affected 1.0.2-1.0.2p).
CVE-2018-10126	ijg-libjpeg before 9d, as used in tiff2pdf (from LibTIFF) and other products, does not check for a NULL pointer at jfdctint.c.
CVE-2018-13410	Info-ZIP Zip 3.0, when the -T and -TT command-line options are used, allows attackers to cause a denial of service possibly have unspecified other impact because of an off-by-one error. NOTE: it is unclear whether there are realis controls the -TT value, given that the entire purpose of -TT is execution of arbitrary commands
CVE-2018-13419	An issue has been found in libsndfile 1.0.28. There is a memory leak in psf_allocate in common.c, as demonstrated and third parties were unable to reproduce and closed the issue
CVE-2018-15607	In ImageMagick 7.0.8-11 Q16, a tiny input file 0x50 0x36 0x36 0x36 0x36 0x4c 0x36 0x38 0x36 0x36 0x36 0x36 a hang of several minutes during which CPU and memory resources are consumed until ultimately an attempted lar could leverage this vulnerability to cause a denial of service via a crafted file.
CVE-2018-16375	An issue was discovered in OpenJPEG 2.3.0. Missing checks for header_info.height and header_info.width in the f can lead to a heap-based buffer overflow.
CVE-2018-16376	An issue was discovered in OpenJPEG 2.3.0. A heap-based buffer overflow was discovered in the function t2_enco vulnerability causes an out-of-bounds write, which may lead to remote denial of service or possibly unspecified oth
CVE-2018-20846	Out-of-bounds accesses in the functions pi_next_lrcp, pi_next_rlcp, pi_next_rpcl, pi_next_pcrl, pi_next_rpcl, and p through 2.3.0 allow remote attackers to cause a denial of service (application crash).
CVE-2018-5407	Simultaneous Multi-threading (SMT) in processors can enable local users to exploit software vulnerable to timing a 'port contention'.
CVE-2018-5709	An issue was discovered in MIT Kerberos 5 (aka krb5) through 1.16. There is a variable "dbentry->n_key_data" in data but unknowingly the developer has assigned a "u4" variable to it, which is for 32-bit data. An attacker can use the database as we know that a Kerberos database dump file contains trusted data.
CVE-2018-6829	cipher/elgamal.c in Libgcrypt through 1.8.2, when used to encrypt messages directly, improperly encodes plaintext information by reading ciphertext data (i.e., it does not have semantic security in face of a ciphertext-only attack). T assumption does not hold for Libgcrypt's ElGamal implementation.
CVE-2018-6951	An issue was discovered in GNU patch through 2.7.6. There is a segmentation fault, associated with a NULL point in the intuit_diff_type function in pch.c, aka a "mangled rename" issue.
CVE-2019-12900	BZ2_decompress in decompress.c in bzip2 through 1.0.6 has an out-of-bounds write when there are many selectors
CVE-2019-1547	Normally in OpenSSL EC groups always have a co-factor present and this is used in side channel resistant code pa construct a group using explicit parameters (instead of using a named curve). In those cases it is possible that such This can occur even where all the parameters match a known named curve. If such a curve is used then OpenSSL f code paths which may result in full key recovery during an ECDSA signature operation. In order to be vulnerable a to time the creation of a large number of signatures where explicit parameters with no co-factor present are in use b avoidance of doubt libssl is not vulnerable because explicit parameters are never used. Fixed in OpenSSL 1.1.1d (A 1.1.0l (Affected 1.1.0-1.1.0k). Fixed in OpenSSL 1.0.2t (Affected 1.0.2-1.0.2s).
CVE-2019-1551	There is an overflow bug in the x64_64 Montgomery squaring procedure used in exponentiation with 512-bit modu suggests that attacks against 2-prime RSA1024, 3-prime RSA1536, and DSA1024 as a result of this defect would b believed likely. Attacks against DH512 are considered just feasible. However, for an attack the target would have t not recommended anyway. Also applications directly using the low level API BN_mod_exp may be affected if they OpenSSL 1.1.1e (Affected 1.1.1-1.1.1d). Fixed in OpenSSL 1.0.2u (Affected 1.0.2-1.0.2t).
CVE-2019-1559	If an application encounters a fatal protocol error and then calls SSL_shutdown() twice (once to send a close_notify can respond differently to the calling application if a 0 byte record is received with invalid padding compared to if MAC. If the application then behaves differently based on that in a way that is detectable to the remote peer, then t used to decrypt data. In order for this to be exploitable "non-stitched" ciphersuites must be in use. Stitched ciphersu commonly used ciphersuites. Also the application must call SSL_shutdown() twice even if a protocol error has occ some do anyway). Fixed in OpenSSL 1.0.2r (Affected 1.0.2-1.0.2q).
CVE-2019-1563	In situations where an attacker receives automated notification of the success or failure of a decryption attempt an a of messages to be decrypted, can recover a CMS/PKCS7 transported encryption key or decrypt any RSA encrypted public RSA key, using a Bleichenbacher padding oracle attack. Applications are not affected if they use a certificat CMS_decrypt or PKCS7_decrypt functions to select the correct recipient info to decrypt. Fixed in OpenSSL 1.1.1d 1.1.0l (Affected 1.1.0-1.1.0k). Fixed in OpenSSL 1.0.2t (Affected 1.0.2-1.0.2s).

CVE-2019-20838	libpcre in PCRE before 8.43 allows a subject buffer over-read in JIT when UTF is disabled, and \X or \R has more CVE-2019-20454.
CVE-2019-5068	An exploitable shared memory permissions vulnerability exists in the functionality of X11 Mesa 3D Graphics Libr memory without any specific permissions to trigger this vulnerability.
CVE-2019-6129	png_create_info_struct in png.c in libpng 1.6.36 has a memory leak, as demonstrated by pngcp. NOTE: a third part free this buffer.
CVE-2019-8457	SQLite3 from 3.6.0 to and including 3.27.2 is vulnerable to heap out-of-bound read in the rtreenode() function whe
CVE-2020-0478	In extend_frame_lowbd of restoration.c, there is a possible out of bounds write due to a missing bounds check. Thi with no additional execution privileges needed. User interaction is needed for exploitation.Product: AndroidVersio
CVE-2020-13529	An exploitable denial-of-service vulnerability exists in Systemd 245. A specially crafted DHCP FORCERENEW p client to be vulnerable to a DHCP ACK spoofing attack. An attacker can forge a pair of FORCERENEW and DCH
CVE-2020-14145	The client side in OpenSSH 5.7 through 8.4 has an Observable Discrepancy leading to an information leak in the al middle attackers to target initial connection attempts (where no host key for the server has been cached by the clien are also affected.
CVE-2020-15719	libldap in certain third-party OpenLDAP packages has a certificate-validation flaw when the third-party package is even when there is a non-matching subjectAltName (SAN). This is fixed in, for example, openldap-2.4.46-10.el8 in
CVE-2020-15719	libldap in certain third-party OpenLDAP packages has a certificate-validation flaw when the third-party package is even when there is a non-matching subjectAltName (SAN). This is fixed in, for example, openldap-2.4.46-10.el8 in
CVE-2020-15778	scp in OpenSSH through 8.3p1 allows command injection in the scp.c toremote function, as demonstrated by backt NOTE: the vendor reportedly has stated that they intentionally omit validation of "anomalous argument transfers" b breaking existing workflows."
CVE-2020-1968	The Raccoon attack exploits a flaw in the TLS specification which can lead to an attacker being able to compute th have used a Diffie-Hellman (DH) based ciphersuite. In such a case this would result in the attacker being able to ea sent over that TLS connection. The attack can only be exploited if an implementation re-uses a DH secret across m only impacts DH ciphersuites and not ECDH ciphersuites. This issue affects OpenSSL 1.0.2 which is out of suppor OpenSSL 1.1.1 is not vulnerable to this issue. Fixed in OpenSSL 1.0.2w (Affected 1.0.2-1.0.2v).
CVE-2020-1971	The X.509 GeneralName type is a generic type for representing different types of names. One of those name types provides a function GENERAL_NAME_cmp which compares different instances of a GENERAL_NAME to see i incorrectly when both GENERAL_NAMEs contain an EDIPARTYNAME. A NULL pointer dereference and a cra service attack. OpenSSL itself uses the GENERAL_NAME_cmp function for two purposes: 1) Comparing CRL di CRL and a CRL distribution point embedded in an X509 certificate 2) When verifying that a timestamp response to name (exposed via the API functions TS_RESP_verify_response and TS_RESP_verify_token) If an attacker can co attacker could trigger a crash. For example if the attacker can trick a client or server into checking a malicious certi occur. Note that some applications automatically download CRLs based on a URL embedded in a certificate. This the certificate and CRL being verified. OpenSSL's s_server, s_client and verify tools have support for the "-crl_dow CRL downloading and this attack has been demonstrated to work against those tools. Note that an unrelated bug m cannot parse or construct correct encodings of EDIPARTYNAME. However it is possible to construct a malformed will accept and hence trigger this attack. All OpenSSL 1.1.1 and 1.0.2 versions are affected by this issue. Other Op not been checked. Fixed in OpenSSL 1.1.1i (Affected 1.1.1-1.1.1h). Fixed in OpenSSL 1.0.2x (Affected 1.0.2-1.0.2
CVE-2020-26140	An issue was discovered in the ALFA Windows 10 driver 6.1316.1209 for AWUS036H. The WEP, WPA, WPA2, frames in a protected Wi-Fi network. An adversary can abuse this to inject arbitrary data frames independent of the
CVE-2020-26146	An issue was discovered on Samsung Galaxy S3 i9305 4.4.4 devices. The WPA, WPA2, and WPA3 implementatio consecutive packet numbers. An adversary can abuse this to exfiltrate selected fragments. This vulnerability is expl fragmented frames and the WEP, CCMP, or GCMP data-confidentiality protocol is used. Note that WEP is vulnera
CVE-2020-36386	An issue was discovered in the Linux kernel before 5.8.1. net/bluetooth/hci_event.c has a slab out-of-bounds read i CID-51c19bf3d5cf.
CVE-2020-36558	A race condition in the Linux kernel before 5.5.7 involving VT_RESIZEX could lead to a NULL pointer dereferen
CVE-2020-36781	In the Linux kernel, the following vulnerability has been resolved: i2c: imx: fix reference leak when pm_runtime_g i2c_imx_remove(), the pm reference count is not expected to be incremented on return. However, pm_runtime_get failed. Forgetting to putting operation will result in a reference leak here. Replace it with pm_runtime_resume_and
CVE-2020-36782	In the Linux kernel, the following vulnerability has been resolved: i2c: imx-lpi2c: fix reference leak when pm_runt is not expected to be incremented on return in lpi2c_imx_master_enable. However, pm_runtime_get_sync will incr Forgetting to putting operation will result in a reference leak here. Replace it with pm_runtime_resume_and_get to
CVE-2020-5408	Spring Security versions 5.3.x prior to 5.3.2, 5.2.x prior to 5.2.4, 5.1.x prior to 5.1.10, 5.0.x prior to 5.0.16 and 4.2. vector with CBC Mode in the implementation of the queryable text encryptor. A malicious user with access to the encryptor may be able to derive the unencrypted values using a dictionary attack.

CVE-2020-8554	Kubernetes API server in all versions allow an attacker who is able to create a ClusterIP service and set the spec.ex address. Additionally, an attacker who is able to patch the status (which is considered a privileged operation and sh LoadBalancer service can set the status.loadBalancer.ingress.ip to similar effect.
CVE-2020-8561	A security issue was discovered in Kubernetes where actors that control the responses of MutatingWebhookConfig requests are able to redirect kube-apiserver requests to private networks of the apiserver. If that user can view kube they can view the redirected responses and headers in the logs.
CVE-2020-8564	In Kubernetes clusters using a logging level of at least 4, processing a malformed docker config file will result in th leaked, which can include pull secrets or other registry credentials. This affects < v1.19.3, < v1.18.10, < v1.17.13.
CVE-2021-20197	There is an open race window when writing output in the following utilities in GNU binutils version 2.35 and earli utilities are run as a privileged user (presumably as part of a script updating binaries across different users), an unp getting ownership of arbitrary files through a symlink.
CVE-2021-20197	There is an open race window when writing output in the following utilities in GNU binutils version 2.35 and earli utilities are run as a privileged user (presumably as part of a script updating binaries across different users), an unp getting ownership of arbitrary files through a symlink.
CVE-2021-20311	A flaw was found in ImageMagick in versions before 7.0.11, where a division by zero in sRGBTransformImage() i undefined behavior via a crafted image file that is submitted by an attacker processed by an application using Imag vulnerability is to system availability.
CVE-2021-22922	When curl is instructed to download content using the metalink feature, thecontents is verified against a hash provi XML file points out to the client how to get the same contentfrom a set of different URLs, potentially hosted by dif download the file from one or several of them. In a serial orparallel manner.If one of the servers hosting the conten specific file on that server is replaced with a modified payload, curlshould detect this when the hash of the file misr remove the contents and instead try getting the contentsfrom another URL. This is not done, and instead such a has potentially malicious content is kept in the file ondisk.
CVE-2021-22923	When curl is instructed to get content using the metalink feature, and a user name and password are used to downlo credentials are then subsequently passed on to each of the servers from which curl will download or try to downloa user's expectations and intentions and without telling the user it happened.
CVE-2021-23336	The package python/cpython from 0 and before 3.6.13, from 3.7.0 and before 3.7.10, from 3.8.0 and before 3.8.8, f to Web Cache Poisoning via urllib.parse.parse_qsl and urllib.parse.parse_qs by using a vector called parameter cloa parameters using a semicolon (;), they can cause a difference in the interpretation of the request between the proxy server. This can result in malicious requests being cached as completely safe ones, as the proxy would usually not s would not include it in a cache key of an unkeyed parameter.
CVE-2021-23841	The OpenSSL public API function X509_issuer_and_serial_hash() attempts to create a unique hash value based on within an X509 certificate. However it fails to correctly handle any errors that may occur while parsing the issuer f is maliciously constructed). This may subsequently result in a NULL pointer deref and a crash leading to a potentia X509_issuer_and_serial_hash() is never directly called by OpenSSL itself so applications are only vulnerable if the it on certificates that may have been obtained from untrusted sources. OpenSSL versions 1.1.1i and below are affec should upgrade to OpenSSL 1.1.1j. OpenSSL versions 1.0.2x and below are affected by this issue. However OpenS receiving public updates. Premium support customers of OpenSSL 1.0.2 should upgrade to 1.0.2y. Other users shou 1.1.1j (Affected 1.1.1-1.1.1i). Fixed in OpenSSL 1.0.2y (Affected 1.0.2-1.0.2x).
CVE-2021-25736	Kube-proxy on Windows can unintentionally forward traffic to local processes listening on the same port (‚Äúspec Service when the LoadBalancer controller does not set the ‚Äústatus.loadBalancer.ingress[].ip‚Äù field. Clusters w ‚Äústatus.loadBalancer.ingress[].ip‚Äù field are unaffected.
CVE-2021-25740	A security issue was discovered with Kubernetes that could enable users to send network traffic to locations they w confused deputy attack.
CVE-2021-25743	kubectl does not neutralize escape, meta or control sequences contained in the raw data it outputs to a terminal. Thi unstructured string fields in objects such as Events.
CVE-2021-26945	An integer overflow leading to a heap-buffer overflow was found in OpenEXR in versions before 3.0.1. An attacke compiled with OpenEXR.
CVE-2021-31535	LookupCol.c in X.Org X through X11R7.7 and libX11 before 1.7.1 might allow remote attackers to execute arbitra (intended for server-side color lookup) contains a flaw allowing a client to send color-name requests with a name lo protocol (and also longer than the maximum packet size for normal-sized packets). The user-controlled data exceed the server as additional X protocol requests and executed, e.g., to disable X server authorization completely. For ex terminal control sequences for color codes, then the attacker may be able to take full control of the running graphic
CVE-2021-33560	Libgcrypt before 1.8.8 and 1.9.x before 1.9.3 mishandles ElGamal encryption because it lacks exponent blinding to mpi_powm, and the window size is not chosen appropriately. This, for example, affects use of ElGamal in OpenPG
CVE-2021-33656	When setting font with malicous data by ioctl cmd PIO_FONT,kernel will write memory out of bounds.

CVE-2021-34866	This vulnerability allows local attackers to escalate privileges on affected installations of Linux Kernel 5.14-rc3. A execute low-privileged code on the target system in order to exploit this vulnerability. The specific flaw exists with issue results from the lack of proper validation of user-supplied eBPF programs, which can result in a type confusio vulnerability to escalate privileges and execute arbitrary code in the context of the kernel. Was ZDI-CAN-14689.
CVE-2021-3618	ALPACA is an application layer protocol content confusion attack, exploiting TLS servers implementing different such as multi-domain or wildcard certificates. A MiTM attacker having access to victim's traffic at the TCP/IP laye another, resulting in a valid TLS session. This breaks the authentication of TLS and cross-protocol attacks may be service may compromise the other at the application layer.
CVE-2021-36368	An issue was discovered in OpenSSH before 8.9. If a client is using public-key authentication with agent forwardin an attacker has silently modified the server to support the None authentication option, then the user cannot determi confirm that the user wishes to connect to that server, or that the user wishes to allow that server to connect to a dif vendor's position is "this is not an authentication bypass, since nothing is being bypassed.
CVE-2021-3712	ASN.1 strings are represented internally within OpenSSL as an ASN1_STRING structure which contains a buffer l buffer length. This contrasts with normal C strings which are repesented as a buffer for the string data which is term strict requirement, ASN.1 strings that are parsed using OpenSSL's own "d2i" functions (and other similar parsing f has been set with the ASN1_STRING_set() function will additionally NUL terminate the byte array in the ASN1_S for applications to directly construct valid ASN1_STRING structures which do not NUL terminate the byte array b fields in the ASN1_STRING array. This can also happen by using the ASN1_STRING_set0() function. Numerous have been found to assume that the ASN1_STRING byte array will be NUL terminated, even though this is not gua constructed. Where an application requests an ASN.1 structure to be printed, and where that ASN.1 structure conta constructed by the application without NUL terminating the "data" field, then a read buffer overrun can occur. The constraints processing of certificates (for example if a certificate has been directly constructed by the application in functions, and the certificate contains non NUL terminated ASN1_STRING structures). It can also occur in the X5 and X509_get1_ocsp() functions. If a malicious actor can cause an application to directly construct an ASN1_STRI affected OpenSSL functions then this issue could be hit. This might result in a crash (causing a Denial of Service a of private memory contents (such as private keys, or sensitive plaintext). Fixed in OpenSSL 1.1.1l (Affected 1.1.1- 1.0.2-1.0.2y).
CVE-2021-37159	hso_free_net_device in drivers/net/usb/hso.c in the Linux kernel through 5.13.4 calls unregister_netdev without ch state, leading to a use-after-free and a double free.
CVE-2021-3782	An internal reference count is held on the buffer pool, incremented every time a new buffer is created from the poo int; on LP64 systems this can cause the reference count to overflow if the client creates a large number of wl_shm b create a large number of external references to the buffer storage. With the reference count overflowing, a use-after tracking structure, where values may be incremented or decremented; it may also be possible to construct a limited to the attacking client at a time.
CVE-2021-3968	vim is vulnerable to Heap-based Buffer Overflow
CVE-2021-3973	vim is vulnerable to Heap-based Buffer Overflow
CVE-2021-3984	vim is vulnerable to Heap-based Buffer Overflow
CVE-2021-4019	vim is vulnerable to Heap-based Buffer Overflow
CVE-2021-40211	An issue was discovered with ImageMagick 7.1.0-4 via Division by zero in function ReadEnhMetaFile of coders/e
CVE-2021-4069	vim is vulnerable to Use After Free
CVE-2021-4136	vim is vulnerable to Heap-based Buffer Overflow
CVE-2021-4173	vim is vulnerable to Use After Free
CVE-2021-4187	vim is vulnerable to Use After Free
CVE-2021-4192	vim is vulnerable to Use After Free
CVE-2021-4193	vim is vulnerable to Out-of-bounds Read
CVE-2021-4202	A use-after-free flaw was found in nci_request in net/nfc/nci/core.c in NFC Controller Interface (NCI) in the Linux with user privileges to cause a data race problem while the device is getting removed, leading to a privilege escalat
CVE-2021-4204	An out-of-bounds (OOB) memory access flaw was found in the Linux kernel's eBPF due to an Improper Input Vali a special privilege to crash the system or leak internal information.
CVE-2021-4214	A heap overflow flaw was found in libpngs' pngimage.c program. This flaw allows an attacker with local network a the pngimage utility, causing an application to crash, leading to a denial of service.
CVE-2021-42739	The firewire subsystem in the Linux kernel through 5.14.13 has a buffer overflow related to drivers/media/firewire/ firedtv-ci.c, because avc_ca_pmt mishandles bounds checking.

CVE-2021-43975	In the Linux kernel through 5.15.2, hw_atl_utils_fw_rpc_wait in drivers/net/ethernet/aquantia/atlantic/hw_atl/hw_a introduce a crafted device) to trigger an out-of-bounds write via a crafted length value.
CVE-2021-44879	In gc_data_segment in fs/f2fs/gc.c in the Linux kernel before 5.16.3, special files are not considered, leading to a m
CVE-2021-45346	A Memory Leak vulnerability exists in SQLite Project SQLite3 3.35.1 and 3.37.0 via maliciously crafted SQL Que is possible to query a record, and leak subsequent bytes of memory that extend beyond the record, which could let a NOTE: The developer disputes this as a vulnerability stating that If you give SQLite a corrupted database file and s read parts of the database that you did not intend or expect.
CVE-2021-45940	libbpf 0.6.0 and 0.6.1 has a heap-based buffer overflow (4 bytes) in __bpf_object__open (called from bpf_object__
CVE-2021-45941	libbpf 0.6.0 and 0.6.1 has a heap-based buffer overflow (8 bytes) in __bpf_object__open (called from bpf_object__
CVE-2021-46310	An issue was discovered IW44Image.cpp in djvulibre 3.5.28 in allows attackers to cause a denial of service via div
CVE-2021-46312	An issue was discovered IW44EncodeCodec.cpp in djvulibre 3.5.28 in allows attackers to cause a denial of service
CVE-2021-46822	The PPM reader in libjpeg-turbo through 2.0.90 mishandles use of tjLoadImage for loading a 16-bit binary PPM fil binary PGM file into an RGB buffer. This is related to a heap-based buffer overflow in the get_word_rgb_row func
CVE-2021-46909	In the Linux kernel, the following vulnerability has been resolved: ARM: footbridge: fix PCI interrupt mapping Sin to pci_assign_irq() in pci_device_probe()"), the PCI code will call the IRQ mapping function whenever a PCI drive causes an oops if a PCI driver is loaded or bound after the kernel has initialised.
CVE-2021-46923	In the Linux kernel, the following vulnerability has been resolved: fs/mount_setattr: always cleanup mount_kattr M after mount_kattr was succesfully built in both the success and failure case to prevent leaking any references we too lookup failed thereby risking to leak an additional reference we took when building mount_kattr when an idmapped
CVE-2021-46931	In the Linux kernel, the following vulnerability has been resolved: net/mlx5e: Wrap the tx reporter dump callback t mlx5e_tx_reporter_dump_sq() casts its void * argument to struct mlx5e_txqsq *, but in TX-timeout-recovery flow struct mlx5e_tx_timeout_ctx *. mlx5_core 0000:08:00.1 enp8s0f1: TX timeout detected mlx5_core 0000:08:00.1 e SQ: 0x11ec, CQ: 0x146d, SQ Cons: 0x0 SQ Prod: 0x1, usecs since last trans: 21565000 BUG: stack guard page wa 00000000b66ea0dc..000000004d932dae) kernel stack overflow (page fault): 0000 [#1] SMP NOPTI CPU: 5 PID: 9 OE 5.13.0_mlnx #1 Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS rel-1.13.0-0-gf21b5a4aeb02 mlx5e mlx5e_tx_timeout_work [mlx5_core] RIP: 0010:mlx5e_tx_reporter_dump_sq+0xd3/0x180 [mlx5_core] Ca +0x43/0x1c0 [mlx5_core] devlink_health_do_dump.part.91+0x71/0xd0 devlink_health_report+0x157/0x1b0 mlx5 [mlx5_core] ? mlx5e_tx_reporter_err_cqe_recover+0x1d0/0x1d0 [mlx5_core] ? mlx5e_health_queue_dump+0xd0 +0x19b/0x550 ? set_next_entity+0x72/0x80 ? pick_next_task_fair+0x227/0x340 ? finish_task_switch+0xa2/0x280 [mlx5_core] process_one_work+0x1de/0x3a0 worker_thread+0x2d/0x3c0 ? process_one_work+0x3a0/0x3a0 kthre +0x90/0x90 ret_from_fork+0x1f/0x30 --[end trace 51ccabea504edaff]--- RIP: 0010:mlx5e_tx_reporter_dump_sq Kernel panic - not syncing: Fatal exception Kernel Offset: disabled end Kernel panic - not syncing: Fatal exception mlx5e_tx_reporter_dump_sq() which extracts the sq from struct mlx5e_tx_timeout_ctx and set it as the TX-timeou
CVE-2021-47033	In the Linux kernel, the following vulnerability has been resolved: mt76: mt7615: fix tx skb dma unmap The first p well, otherwise it will leak DMA mapping entries
CVE-2021-47063	In the Linux kernel, the following vulnerability has been resolved: drm: bridge/panel: Cleanup connector on bridge drm_connector_cleanup() manually in panel_bridge_detach(), the connector will be cleaned up with the other DRM drm_mode_config_cleanup(). However, since our drm_connector is devm-allocated, by the time drm_mode_config be long gone. Therefore, the connector must be cleaned up when the bridge is detached to avoid use-after-free cond created v3: Add FIXME v4: (Use connector->dev) directly in if() block
CVE-2021-47073	In the Linux kernel, the following vulnerability has been resolved: platform/x86: dell-smbios-wmi: Fix oops on rm only registers the dell_smbios_wmi_driver on systems where the Dell WMI interface is supported. While exit_dell this leads to the following oops: [175.722921] ------------[cut here]------------ [175.722925] Unexpected driver un PID: 3630 at drivers/base/driver.c:194 driver_unregister+0x38/0x40 ... [175.723089] Call Trace: [175.723094] cle [175.723148] ---[end trace 064c34e1ad49509d]--- Make the unregister happen on the same condition the register
CVE-2021-47089	In the Linux kernel, the following vulnerability has been resolved: kfence: fix memory leak when cat kfence object unreferenced object 0xffff93d1d8cc02e8 (size 248): comm "cat", pid 23327, jiffies 4624670141 (age 495992.217s) 93 ff ff 00 10 00 00 00 00 00 00 .@..... 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ..... backtrace +0x167/0x1e0 do_dentry_open+0x1e1/0x3a0 path_openat+0x961/0xa20 do_filp_open+0xae/0x120 do_sys_opena do_syscall_64+0x33/0x40 entry_SYSCALL_64_after_hwframe+0x44/0xa9 unreferenced object 0xffff93d4198540 jiffies 4624670141 (age 495992.217s) hex dump (first 32 bytes): 6b 66 65 6e 63 65 2d 23 32 35 30 3a 20 30 78 30 30 37 35 34 62 64 61 31 32 2d 0000000754bda12- backtrace: seq_read_iter+0x313/0x440 seq_read+0x14b/0x1a0 +0xa5/0x1b0 ksys_read+0xa0/0xf0 do_syscall_64+0x33/0x40 entry_SYSCALL_64_after_hwframe+0x44/0xa9 I f with the following commands: cat /sys/kernel/debug/kfence/objects echo scan > /sys/kernel/debug/kmemleak cat /s memory is allocated in the stack below: do_syscall_64 do_sys_open do_dentry_open full_proxy_open seq_open -- seq_read seq_read_iter traverse ---> alloc seq_buf And it should have been released in the following process: do_sy exit_to_user_mode_prepare task_work_run ____fput __fput full_proxy_release ---> free here However, the release not implemented in kfence. As a result, a memory leak occurs. Therefore, the solution to this problem is to implem

CVE-2021-47097	In the Linux kernel, the following vulnerability has been resolved: Input: elantech - fix stack out of bound access in The array param[] in elantech_change_report_id() must be at least 3 bytes, because elantech_read_reg_params() is PSMOUSE_CMD_GETINFO, that is going to access 3 bytes from param[], but it's defined in the stack as an array potential stack out-of-bounds access here, also confirmed by KASAN: [6.512374] BUG: KASAN: stack-out-of-bo [6.512397] Read of size 1 at addr ffff8881024d77c2 by task kworker/2:1/118 [6.512416] CPU: 2 PID: 118 Comm generic #22+arighi20211110 [6.512428] Hardware name: LENOVO 20T8000QGE/20T8000QGE, BIOS R1AET3 Workqueue: events_long serio_handle_event [6.512453] Call Trace: [6.512462] show_stack+0x52/0x58 [6.5124 print_address_description.constprop.0+0x1d/0x140 [6.512502] ? __ps2_command+0x372/0x7e0 [6.512516] __ka _raw_write_lock_irq+0x20/0xd0 [6.512539] ? __ps2_command+0x372/0x7e0 [6.512552] kasan_report+0x3c/0x5 [6.512575] __ps2_command+0x372/0x7e0 [6.512589] ? ps2_drain+0x240/0x240 [6.512601] ? dev_printk_emit+ +0xc5/0xc5 [6.512621] ? __kasan_check_write+0x14/0x20 [6.512634] ? mutex_lock+0x8f/0xe0 [6.512643] ? __ [6.512655] ps2_command+0x52/0x90 [6.512670] elantech_ps2_command+0x4f/0xc0 [psmouse] [6.512734] elan [psmouse] [6.512799] ? elantech_report_trackpoint.constprop.0.cold+0xd/0xd [psmouse] [6.512863] ? ps2_comm elantech_query_info.cold+0x6bd/0x9ed [psmouse] [6.512943] ? elantech_setup_ps2+0x460/0x460 [psmouse] [6. [psmouse] [6.513064] ? psmouse_attr_set_helper+0x2a0/0x2a0 [psmouse] [6.513122] ? phys_pmd_init+0x30e/0x [psmouse] [6.513200] ? elantech_init_ps2+0xf0/0xf0 [psmouse] [6.513249] ? elantech_query_info+0x440/0x440 +0x60/0x60 [psmouse] [6.513342] ? elantech_query_info+0x440/0x440 [psmouse] [6.513388] ? psmouse_try_pr psmouse_extensions+0x65d/0x6e0 [psmouse] [6.513476] ? psmouse_try_protocol+0x170/0x170 [psmouse] [6.51 [6.513526] ? ps2_command+0x7f/0x90 [6.513536] ? psmouse_probe+0xa3/0xf0 [psmouse] [6.513580] psmouse [psmouse] [6.513624] psmouse_connect+0x272/0x530 [psmouse] [6.513669] serio_driver_probe+0x55/0x70 [6. [6.513689] driver_probe_device+0x160/0x1f0 [6.513697] device_driver_attach+0x119/0x130 [6.513705] ? devic __driver_attach+0xe7/0x1a0 [6.513720] ? device_driver_attach+0x130/0x130 [6.513728] bus_for_each_dev+0xf +0x10/0x10 [6.513748] ? _raw_write_unlock_bh+0x30/0x30 [6.513757] driver_attach+0x2d/0x40 [6.513764] se process_one_work+0x471/0x740 [6.513785] worker_thread+0x2d2/0x790 [6.513794] ? process_one_work+0x74 [6.513809] ? set_kthread_struct+0x80/0x80 [6.513816] ret_from_fork+0x22/0x30 [6.513832] The buggy address page:00000000bc35e189 refcount:0 mapcount:0 mapping:0000000000000000 index:0x0 pfn:0x1024d7 [6.513847 lastcpupid=0x1fffff) [6.513860] raw: 0 ---truncated---
CVE-2021-47131	In the Linux kernel, the following vulnerability has been resolved: net/tls: Fix use-after-free after the TLS device g active TLS offload goes down, tls_device_down is called to stop the offload and tear down the TLS context. Howe points to the TLS context, which is now deallocated. If a netdev goes up, while the connection is still active, and th retransmissions, it will lead to a use-after-free of the TLS context. This commit addresses this bug by keeping the c and implements the necessary fallbacks, so that the connection can resume in software (non-offloaded) kTLS mode encrypt all packets. The RX side already has all the necessary fallbacks, because receiving non-decrypted packets i is to block resync requests, which are normally produced after receiving non-decrypted packets. The necessary syn teardown: first the fallbacks are deployed, then the driver resources are released (it used to be possible to have a tls called TLS_RX_DEV_DEGRADED is added to indicate the fallback mode. It's used to skip the RX resync logic c objects may be released (for example, resync_async, which is allocated and freed by the driver).
CVE-2021-47188	In the Linux kernel, the following vulnerability has been resolved: scsi: ufs: core: Improve SCSI abort handling Th a test setup: WARNING: CPU: 4 PID: 250 at drivers/scsi/ufs/ufshcd.c:2737 ufshcd_queuecommand+0x468/0x65c +0x468/0x65c scsi_send_eh_cmnd+0x224/0x6a0 scsi_eh_test_devices+0x248/0x418 scsi_eh_ready_devs+0xc34/0 kthread+0x150/0x1b4 ret_from_fork+0x10/0x30 That warning is triggered by the following statement: WARN_ON lrbp->cmd from the abort handler.
CVE-2021-47202	In the Linux kernel, the following vulnerability has been resolved: thermal: Fix NULL pointer dereferences in of_t parses the thermal-zones node and registers a thermal_zone device for each subnode. However, if a thermal zone is that thermal sensor device hasn't probed yet, an attempt to set trip_point_*_temp for that thermal zone device can c Fix it. console:/sys/class/thermal/thermal_zone87 # echo 120000 > trip_point_0_temp ... Unable to handle kernel N address 0000000000000020 ... Call trace: of_thermal_set_trip_temp+0x40/0xc4 trip_point_temp_store+0xc0/0x1d sysfs_kf_write+0x64/0xc0 kernfs_fop_write_iter+0x108/0x1d0 vfs_write+0x2f4/0x368 ksys_write+0x7c/0xec __a el0_svc_common.llvm.7279915941325364641+0xbc/0x1bc do_el0_svc+0x28/0xa0 el0_svc+0x14/0x24 el0_sync_ While at it, fix the possible NULL pointer dereference in other functions as well: of_thermal_get_temp(), of_therm
CVE-2021-47214	In the Linux kernel, the following vulnerability has been resolved: hugetlb, userfaultfd: fix reservation restore on u case in hugetlb_mcopy_atomic_pte(), if we bail out using "goto out_release_unlock;" in the cases where idx >= siz that new_pagecache_page == false, and so call restore_reserve_on_error(). In this case I see restore_reserve_on_er call to remove_inode_hugepages() will increment h->resv_hugepages causing a 100% reproducible leak. We shoul adding a page into the pagecache and set new_pagecache_page to true, to indicate that there is no reservation to res restore_reserve_on_error(). Rename new_pagecache_page to page_in_pagecache to make that clear.
CVE-2021-47236	In the Linux kernel, the following vulnerability has been resolved: net: cdc_eem: fix tx fixup skb leak when usbnet eem_tx_fixup(), if skb_copy_expand() failed, it return NULL, usbnet_start_xmit() will have no chance to free origi eem_tx_fixup() first, then check skb clone status, if failed, return NULL to usbnet.

CVE-2021-47237	In the Linux kernel, the following vulnerability has been resolved: net: hamradio: fix memory leak in mkiss_close mkiss_open()[1]. The problem was in missing free_netdev() in mkiss_close(). In mkiss_open() netdevice is allocat netdevice was only unregistered, but not freed. Fail log: BUG: memory leak unreferenced object 0xffff8880281ba0 pid 11443, jiffies 4295046091 (age 17.660s) hex dump (first 32 bytes): 61 78 30 00 00 00 00 00 00 00 00 00 00 00 ff ff 00 00 00 00 00 00 00 00 .'.*..... backtrace: [<ffffffff81a27201>] kvmalloc_node+0x61/0xf0 [<ffffffff8706 [<ffffffff84e64192>] mkiss_open+0xb2/0x6f0 [1] [<ffffffff842355db>] tty_ldisc_open+0x9b/0x110 [<ffffffff8423 [<ffffffff8421f7f3>] tty_ioctl+0xda3/0x1440 [<ffffffff81c9f273>] __x64_sys_ioctl+0x193/0x200 [<ffffffff891126 [<ffffffff89200068>] entry_SYSCALL_64_after_hwframe+0x44/0xae BUG: memory leak unreferenced object 0x "syz-executor.1", pid 11443, jiffies 4295046091 (age 17.660s) hex dump (first 32 bytes): e8 a2 1b 28 80 88 ff ff e8 92 9c aa b0 40 02 00 00 00 00 00 00 00 00 .....@..... backtrace: [<ffffffff8709f68b>] __hw_addr_create_ex+(__hw_addr_add_ex+0x1f8/0x2b0 [<ffffffff870a0c7b>] dev_addr_init+0x10b/0x1f0 [<ffffffff8706e88b>] alloc_ne mkiss_open+0xb2/0x6f0 [1] [<ffffffff842355db>] tty_ldisc_open+0x9b/0x110 [<ffffffff84236488>] tty_set_ldisc+ tty_ioctl+0xda3/0x1440 [<ffffffff81c9f273>] __x64_sys_ioctl+0x193/0x200 [<ffffffff8911263a>] do_syscall_64+ entry_SYSCALL_64_after_hwframe+0x44/0xae BUG: memory leak unreferenced object 0xffff8880219bfc00 (siz jiffies 4295046091 (age 17.660s) hex dump (first 32 bytes): 00 a0 1b 28 80 88 ff ff 80 8f b1 8d ff ff ff ff ...(..... 00 00 ..... backtrace: [<ffffffff81a27201>] kvmalloc_node+0x61/0xf0 [<ffffffff8706eec7>] alloc_netdev_m mkiss_open+0xb2/0x6f0 [1] [<ffffffff842355db>] tty_ldisc_open+0x9b/0x110 [<ffffffff84236488>] tty_set_ldisc+ tty_ioctl+0xda3/0x1440 [<ffffffff81c9f273>] __x64_sys_ioctl+0x193/0x200 [<ffffffff8911263a>] do_syscall_64+ entry_SYSCALL_64_after_hwframe+0x44/0xae BUG: memory leak unreferenced object 0xffff888029b2b200 (siz jiffies 4295046091 (age 17.660s) hex dump (first 32 bytes): 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ..... 00 00 00 00 ..... backtrace: [<ffffffff81a27201>] kvmalloc_node+0x61/0xf0 [<ffffffff8706f062>] alloc_netd mkiss_open+0xb2/0x6f0 [1] [<ffffffff842355db>] tty_ldisc_open+0x9b/0x110 [<ffffffff84236488>] tty_set_ldisc+ tty_ioctl+0xda3/0x1440 [<ffffffff81c9f273>] __x64_sys_ioctl+0x193/0x200 [<ffffffff8911263a>] do_syscall_64+ entry_SYSCALL_64_after_hwframe+0x44/0xae
CVE-2021-47257	In the Linux kernel, the following vulnerability has been resolved: net: ieee802154: fix null deref in parse dev addr deref if the user sets the mode incorrectly for the given addr type.
CVE-2021-47272	In the Linux kernel, the following vulnerability has been resolved: usb: dwc3: gadget: Bail from dwc3_gadget_exit possible scenario in which dwc3_gadget_init() can fail: during during host -> peripheral mode switch in dwc3_set_ to bind. Then, if the DRD undergoes another mode switch from peripheral->host the resulting dwc3_gadget_exit() dangling dwc->gadget pointer as well as call dma_free_coherent() on unmapped DMA pointers. The exact scenario in peripheral mode - Configure ConfigFS gadget with FunctionFS instance (or use g_ffs) - Run FunctionFS userspa etc) - Bind gadget driver to DWC3's UDC - Switch DWC3 to host mode => dwc3_gadget_exit() is called. usb_del_ instance on the gadget_driver_pending_list - Stop FunctionFS application (closes the ep files) - Switch DWC3 to p as usb_add_gadget() calls check_pending_gadget_drivers() and attempts to rebind the UDC to the ConfigFS gadge FFS instance is not in FFS_ACTIVE state (userspace has not re-opened and written the descriptors yet, i.e. desc_re => dwc3_gadget_exit() is called again, but this time dwc->gadget is invalid. Although it can be argued that userspa the FunctionFS application be ready prior to allowing the composite driver bind to the UDC, failure to do so shoul Fix this by setting dwc->gadget to NULL in the failure path of dwc3_gadget_init() and add a check to dwc3_gadge valid.
CVE-2021-47287	In the Linux kernel, the following vulnerability has been resolved: driver core: auxiliary bus: Fix memory leak whe returns with error we need to free the memory allocated for auxdrv->driver.name before returning from __auxiliary
CVE-2021-47318	In the Linux kernel, the following vulnerability has been resolved: arch_topology: Avoid use-after-free for scale_fr topology_scale_freq_tick() (which gets called from scheduler_tick()) may end up using a pointer to "struct scale_fr topology_clear_scale_freq_source(), as there is no protection in place here. The users of topology_clear_scale_freq previously cleared scale_freq_data isn't used anymore, so they can free the related resources. Since topology_scale_ don't want to add locking in there. Use the RCU update mechanism instead (which is already used by the scheduler free updates here. synchronize_rcu() makes sure that all RCU critical sections that started before it is called, will fi topology_clear_scale_freq_source() don't need to worry about their callback getting called anymore.
CVE-2021-47346	In the Linux kernel, the following vulnerability has been resolved: coresight: tmc-etf: Fix global-out-of-bounds in t 6f755e85c332 ("coresight: Add helper for inserting synchronization packets") removed trailing '\0' from barrier_pk etb_update_buffer() to have proper checks for barrier_pkt size before read but missed updating tmc_update_etf_bu array size resulting in KASAN out-of-bounds bug. Fix this by adding a check for barrier_pkt size before accessing KASAN: global-out-of-bounds in tmc_update_etf_buffer+0x4b8/0x698 Read of size 4 at addr ffffffd05b7d1030 by +0x0/0x27c show_stack+0x20/0x2c dump_stack+0x11c/0x188 print_address_description+0x3c/0x4a4 __kasan_re __asan_report_load4_noabort+0x1c/0x24 tmc_update_etf_buffer+0x4b8/0x698 etm_event_stop+0x248/0x2d8 etm +0x214/0x6f0 group_sched_out+0xd0/0x270 ctx_sched_out+0x2ec/0x518 __perf_event_task_sched_out+0x4fc/0x preempt_schedule_irq+0x88/0x170 arm64_preempt_schedule_irq+0xf0/0x18c el1_irq+0xe8/0x180 perf_event_ex +0x204/0x400 load_elf_binary+0x72c/0x18c0 search_binary_handler+0x13c/0x420 load_script+0x500/0x6c4 sear exec_binprm+0x118/0x654 __do_execve_file+0x77c/0xba4 __arm64_compat_sys_execve+0x98/0xac el0_svc_co +0x84/0xb0 el0_svc_compat+0x10/0x50 The buggy address belongs to the variable: barrier_pkt+0x10/0x40 Memo ffffffd05b7d0f00: fa fa fa fa 04 fa fa fa fa fa fa fa 00 00 00 00 ffffffd05b7d0f80: 00 00 00 00 00 00 00 00 00 00 00 00 00 fa fa fa fa fa fa 00 00 00 03 ^ ffffffd05b7d1080: fa fa fa fa 00 02 fa fa fa fa fa fa 03 fa fa fa ffffffd05b7d11 fa ==

CVE-2021-47365	In the Linux kernel, the following vulnerability has been resolved: afs: Fix page leak There's a loop in afs_extend_ write we want to make to improve the efficiency of the writeback by making it larger. This loop stops, however, if immediately, but it doesn't get rid of the page ref we speculatively acquired. This was caused by the removal of the using find_get_pages_contig() to xarray scanning as the latter only gets a single page at a time, not a batch. Fix this break from the loop. Unfortunately, we can't just add that page to the pagevec we're employing as we'll go through was found by the generic/074 test. It leaks ~4GiB of RAM each time it is run - which can be observed with "top".
CVE-2021-47373	In the Linux kernel, the following vulnerability has been resolved: irqchip/gic-v3-its: Fix potential VPE leak on err its_vpe_init() returns an error, there is an off-by-one in the number of VPEs to be freed. Fix it by simply passing th index of the loop iterating over the VPEs. [maz: fixed commit message]
CVE-2021-47383	In the Linux kernel, the following vulnerability has been resolved: tty: Fix out-of-bound vmalloc access in imagebl program does an ioctl FBIOPUT_VSCREENINFO passing the fb_var_screeninfo struct containing only the fields If this struct is the same as the previous ioctl, the vc_resize() detects it and doesn't call the resize_screen(), leaving this leads to the updatescrollmode() calculates a wrong value to fbcon_display->vrows, which makes the real_y() r eventually, causes the imageblit to access an out-of-bound address value. To solve this issue I made the resize_scre need any resizing, so it will "fix and fill" the fb_var_screeninfo independently.
CVE-2021-47393	In the Linux kernel, the following vulnerability has been resolved: hwmon: (mlxreg-fan) Return non-zero value wh sysfs Fan speed minimum can be enforced from sysfs. For example, setting current fan speed to 20 is used to enfor 19 - to be not below 90% speed, etcetera. This feature provides ability to limit fan speed according to some system of some replaceable units or high system ambient temperature. Request for changing fan minimum speed is configu through 'sysfs' write procedure. In this situation value of argument 'state' is above nominal fan speed maximum. Re thermal_cooling_device_stats_update() call, because in this case statistics update violates thermal statistics table ra is configured with option CONFIG_THERMAL_STATISTICS. Here is the trace from KASAN: [159.506659] BU thermal_cooling_device_stats_update+0x7d/0xb0 [159.516016] Read of size 4 at addr ffff888116163840 by task h Trace: [159.548366] dump_stack+0x92/0xc1 [159.552084] ? thermal_cooling_device_stats_update+0x7d/0xb0 [+0x345/0x780 [159.688711] thermal_zone_device_set_mode+0x7d/0xc0 [159.694174] mlxsw_thermal_modules [159.700972] ? mlxsw_thermal_set_cur_state+0x5a0/0x5a0 [mlxsw_core] [159.731827] mlxsw_thermal_init+0x RIP: 0033:0x7fd995909970 [160.074239] Code: 73 01 c3 48 8b 0d 28 d5 2b 00 f7 d8 64 89 01 48 83 c8 ff c3 66 0 b8 01 00 00 00 0f 05 <48> 3d 01 f0 ff .. [160.095242] RSP: 002b:00007fff54f5d938 EFLAGS: 00000246 ORIG_ RAX: ffffffffffffffda RBX: 0000000000000013 RCX: 00007fd995909970 [160.111710] RDX: 000000000000001 0000000000000001 [160.119699] RBP: 0000000001906008 R08: 00007fd995bc9760 R09: 00007fd996210700 [R11: 0000000000000246 R12: 0000000000000013 [160.135673] R13: 0000000000000001 R14: 00007fd995bc86 [160.145338] Allocated by task 2924: [160.149242] kasan_save_stack+0x19/0x40 [160.153541] __kasan_kmallo +0x1a2/0x2b0 [160.161552] thermal_cooling_device_setup_sysfs+0xf9/0x1a0 [160.167687] __thermal_cooling_ devm_thermal_of_cooling_device_register+0x60/0xa0 [160.180356] mlxreg_fan_probe+0x474/0x5e0 [mlxreg_fa address belongs to the object at ffff888116163400 [160.249807] which belongs to the cache kmalloc-1k of size 10 located 64 bytes to the right of [160.263814] 1024-byte region [ffff888116163400, ffff888116163800) [160.2775 page: [160.282898] page:0000000012275840 refcount:1 mapcount:0 mapping:0000000000000000 index:0xffff88 head:0000000012275840 order:3 compound_mapcount:0 compound_pincount:0 [160.303251] flags: 0x2000000 [160.309694] raw: 0200000000010200 ffffea00046f7208 ffffea0004928208 ffff88810004dbc0 [160.318367] raw 00000001ffffffff 0000000000000000 [160.327033] page dumped because: kasan: bad access detected [160.33327 buggy address: [160.356469] >ffff888116163800: fc ..
CVE-2021-47422	In the Linux kernel, the following vulnerability has been resolved: drm/nouveau/kms/nv50-: fix file release memor single_release() should be called, otherwise the 'op' allocated in single_open() will be leaked.
CVE-2021-47430	In the Linux kernel, the following vulnerability has been resolved: x86/entry: Clear X86_FEATURE_SMAP when 3c73b81a9164 ("x86/entry, selftests: Further improve user entry sanity checks") added a warning if AC is set when ("x86/entry: Fix AC assertion") changed the warning to only fire if the CPU supports SMAP. However, the warnin SMAP but where it's disabled in the kernel config and when running the syscall_nt selftest, for example: ---------- 0 PID: 49 at irqentry_enter_from_user_mode CPU: 0 PID: 49 Comm: init Tainted: G T 5.15.0-rc4+ #98 e6202628 Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 1.10.2-1ubuntu1 04/01/2014 RIP: 0010:irqent irqentry_enter ? exc_general_protection ? asm_exc_general_protection ? asm_exc_general_protectio IS_ENABLE the warning condition, but even this would not be enough in case SMAP is disabled at boot time with the "nosmap" behaviour, clear X86_FEATURE_SMAP when !CONFIG_X86_SMAP. Found using entry-fuzz + satrandconfig. [

CVE-2021-47441	In the Linux kernel, the following vulnerability has been resolved: mlxsw: thermal: Fix out-of-bounds memory acc states to be set above the maximum cooling state supported by the driver: # cat /sys/class/thermal/thermal_zone2/c thermal/thermal_zone2/cdev0/max_state 10 # echo 18 > /sys/class/thermal/thermal_zone2/cdev0/cur_state # echo 9 accesses when thermal state transition statistics are enabled (CONFIG_THERMAL_STATISTICS=y), as the transi (state) [1]. According to the thermal maintainer, it is the responsibility of the driver to reject such operations [2]. Th be set exceeds the maximum cooling state supported by the driver. To avoid dead code, as suggested by the therma a421ce088ac8 ("mlxsw: core: Extend cooling device with cooling levels") that tried to interpret these invalid cooli way. The cooling levels array is not removed in order to prevent the fans going below 20% PWM, which would ca BUG: KASAN: slab-out-of-bounds in thermal_cooling_device_stats_update+0x271/0x290 Read of size 4 at addr f CPU: 0 PID: 5 Comm: kworker/0:0 Not tainted 5.15.0-rc3-custom-45935-gce1adf704b14 #122 Hardware name: M CB2FO"/"SA000874", BIOS 4.6.5 03/08/2016 Workqueue: events_freezable_power_ thermal_zone_device_check print_address_description.constprop.0+0x1f/0x140 kasan_report.cold+0x7f/0x11b thermal_cooling_device_stats_u +0x15e/0x4e0 thermal_cdev_update+0x9f/0xe0 step_wise_throttle+0x770/0xee0 thermal_zone_device_update+0x worker_thread+0x62f/0x13e0 kthread+0x3ee/0x4e0 ret_from_fork+0x1f/0x30 Allocated by task 1: kasan_save_sta +0x7c/0x90 thermal_cooling_device_setup_sysfs+0x153/0x2c0 __thermal_cooling_device_register.part.0+0x25b/ +0xb3/0x100 mlxsw_thermal_init+0x5c5/0x7e0 __mlxsw_core_bus_device_register+0xcb3/0x19c0 mlxsw_core_ mlxsw_pci_probe+0x54f/0x710 local_pci_probe+0xc6/0x170 pci_device_probe+0x2b2/0x4d0 really_probe+0x29 driver_probe_device+0x51/0x1e0 __driver_attach+0x21b/0x530 bus_for_each_dev+0x14c/0x1d0 bus_add_driver- mlxsw_sp_module_init+0xa2/0x174 do_one_initcall+0xee/0x5f0 kernel_init_freeable+0x45a/0x4de kernel_init+0 The buggy address belongs to the object at ffff8881052f7800 which belongs to the cache kmalloc-1k of size 1024 bytes inside of 1024-byte region [ffff8881052f7800, ffff8881052f7c00) The buggy address belongs to the page: pa mapcount:0 mapping:0000000000000000 index:0x0 pfn:0x1052f0 head:0000000052355272 order:3 compound_m 0x200000000010200(slab|head|node=0|zone=2) raw: 0200000000010200 ffffea0005034800 0000000300000003 f 0000000000100010 00000001ffffffff 0000000000000000 page dumped because: kasan: bad access detected Memo ffff8881052f7a80: 00 00 00 00 00 00 04 fc fc fc fc fc fc fc fc ffff8881052f7b00: fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc ^ ffff8881052f7c00: fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc fc ffff8881052f7c80: fc fc fc f lore.kernel.org/linux-pm/9aca37cb-1629-5c67- ---truncated---
CVE-2021-47444	In the Linux kernel, the following vulnerability has been resolved: drm/edid: In connector_bad_edid() cap num_of_ e11f5bd8228f ("drm: Add support for DP 1.4 Compliance edid corruption test") the function connector_bad_edid() EDID passed to it was big enough to hold `edid[0x7e] + 1` blocks of data (1 extra for the base block). It completely `num_blocks` which indicated how much memory had been allocated for the EDID. Let's fix this by adding a boun the case where there's an error in the first block of the EDID. In that case we will call connector_bad_edid() withou `edid[0x7e]`.
CVE-2021-47499	In the Linux kernel, the following vulnerability has been resolved: iio: accel: kxcjk-1013: Fix possible memory lea type is ACPI_SMO8500, the data->dready_trig will not be set, the memory allocated by iio_triggered_buffer_setup memory leak as follows: unreferenced object 0xffff888009551400 (size 512): comm "i2c-SMO8500-125", pid 911 dump (first 32 bytes): 02 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ..... 00 00 00 00 00 00 00 00 20 e2 e [<0000000041ce75ee>] kmem_cache_alloc_trace+0x16d/0x360 [<000000000aeb17b0>] iio_kfifo_allocate+0x41/ iio_triggered_buffer_setup_ext+0x2c/0x210 [industrialio_triggered_buffer] [<000000004375b15f>] kxcjk1013_pr remove data->dready_trig condition in probe and remove.
CVE-2021-47500	In the Linux kernel, the following vulnerability has been resolved: iio: mma8452: Fix trigger reference couting The trigger to the struct iio_dev. The IIO core when done using this trigger will call `iio_trigger_put()` to drop the refer `iio_trigger_get()` in the driver the reference count can reach 0 too early, the trigger gets freed while still in use and reference to the trigger before assigning it to the IIO device.
CVE-2021-47516	In the Linux kernel, the following vulnerability has been resolved: nfp: Fix memory leak in nfp_cpp_area_cache_a allocates and initializes a CPP area structure. But in line 807 (#2), when the cache is allocated failed, this CPP area memory leak. We can fix it by freeing the CPP area when the cache is allocated failed (#2). 792 int nfp_cpp_area_ 793 { 794 struct nfp_cpp_area_cache *cache; 795 struct nfp_cpp_area *area; 800 area = nfp_cpp_area_alloc(cpp, N 0), 801 0, size); // #1: allocates and initializes 802 if (!area) 803 return -ENOMEM; 805 cache = kzalloc(sizeof(*ca return -ENOMEM; // #2: missing free 817 return 0; 818 }
CVE-2021-47519	In the Linux kernel, the following vulnerability has been resolved: can: m_can: m_can_read_fifo: fix memory leak the second call to m_can_fifo_read() fails, the function jump to the out_fail label and returns without calling m_can previously allocated by alloc_can_skb() is not freed. In other terms, this is a memory leak. This patch adds a goto l Issue was found with GCC -fanalyzer, please follow the link below for details.
CVE-2021-47527	In the Linux kernel, the following vulnerability has been resolved: serial: core: fix transmit-buffer reset and memle convert uart_close to use tty_port_close") converted serial core to use tty_port_close() but failed to notice that the t final close. Not freeing the transmit buffer means that the buffer is no longer cleared on next open so that any ioctl(wait indefinitely (e.g. on termios changes) or that stale data can end up being transmitted in case tx is restarted. Fur been opened would leak on driver unbind. Note that the port lock is held when clearing the buffer pointer due to th a5ba1d95e46e ("uart: fix race between uart_put_char() and uart_shutdown()"). Also note that the tty-port shutdown is not strictly necessary to free the buffer page after releasing the lock (cf. d72402145ace ("tty/serial: do not free tra

CVE-2021-47537	In the Linux kernel, the following vulnerability has been resolved: octeontx2-af: Fix a memleak bug in rvu_mbox_ not freed or passed out under the switch-default region, which could lead to a memory leak. Fix this bug by changi was found by a static analyzer. The analysis employs differential checking to identify inconsistent security operatic paths and confirms that the inconsistent operations are not recovered in the current function or the callers, so they c static analysis, it can be a false positive or hard to trigger. Multiple researchers have cross-reviewed the bug. Builds no new warnings, and our static analyzer no longer warns about this code.
CVE-2021-47538	In the Linux kernel, the following vulnerability has been resolved: rxrpc: Fix rxrpc_local leak in rxrpc_lookup_pee candidate before kfree() as it holds a ref to rxrpc_local. [DH: v2: Changed to abstract the peer freeing code out into
CVE-2021-47539	In the Linux kernel, the following vulnerability has been resolved: rxrpc: Fix rxrpc_peer leak in rxrpc_look_up_bu bundle candidate before kfree() as it holds a ref to rxrpc_peer. [DH: v2: Changed to abstract out the bundle freeing
CVE-2021-47561	In the Linux kernel, the following vulnerability has been resolved: i2c: virtio: disable timeout handling If a timeout I2C bus and/or memory corruptions in the guest since the device can still be operating on the buffers it was given v example, the start of a slub_debug splat which was triggered on the next transfer after one transfer was forced to ti (rust-vmm/vhost-device): BUG kmalloc-1k (Not tainted): Poison overwritten First byte 0x1 instead of 0x6b Alloca cpu=0 pid=29 __kmalloc+0xc2/0x1c9 virtio_i2c_xfer+0x65/0x35c __i2c_transfer+0x429/0x57d i2c_transfer+0x1 i2cdev_ioctl+0x247/0x2ed vfs_ioctl+0x21/0x30 sys_ioctl+0xb18/0xb41 Freed in virtio_i2c_xfer+0x32e/0x35c age virtio_i2c_xfer+0x32e/0x35c __i2c_transfer+0x429/0x57d i2c_transfer+0x115/0x134 i2cdev_ioctl_rdwr+0x16a/0 +0x21/0x30 sys_ioctl+0xb18/0xb41 There is no simple fix for this (the driver would have to always create bounce eventually returns the buffers), so just disable the timeout support for now.
CVE-2021-47572	In the Linux kernel, the following vulnerability has been resolved: net: nexthop: fix null pointer dereference when add an IPv6 nexthop and IPv6 is not enabled (!CONFIG_IPV6) we'll hit a NULL pointer dereference[1] in the erro calling ipv6_stub->fib6_nh_release. The bug has been present since the beginning of IPv6 nexthop gateway suppor Add fib6_nh_init and release to stubs") tells us that only fib6_nh_init has a dummy stub because fib6_nh_release s returns an error, but the commit below added a call to ipv6_stub->fib6_nh_release in its error path. To fix it return error directly without calling ipv6_stub->fib6_nh_release in nh_create_ipv6()'s error path. [1] Output is a bit trunca BUG: kernel NULL pointer dereference, address: 0000000000000000 #PF: supervisor instruction fetch in kerne not-present pagege PGD 0 P4D 0 Oops: 0010 [#1] PREEMPT SMP NOPTI CPU: 4 PID: 638 Comm: ip Kdump: lo Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 1.14.0-4.fc34 04/01/2014 RIP: 0010:0x0 Code 0xffffffffffffffd6. RSP: 0018:ffff888109f5b8f0 EFLAGS: 00010286^Ac RAX: 0000000000000000 RBX: ffff8881 RDX: 0000000000000000 RSI: 0000000000000000 RDI: ffff8881008a2860 RBP: ffff888109f5b9d8 R08: 000000 R10: ffff888109f5b978 R11: ffff888109f5b948 R12: 00000000ffffff9f R13: ffff8881008a2a80 R14: ffff8881008a2 00007f98de70f100(0000) GS:ffff88822bf00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 C ffffffffffffffd6 CR3: 0000000100efc000 CR4: 00000000000006e0 Call Trace: <TASK> nh_create_ipv6+0xed/0x1 check_preemption_disabled+0x3d/0xf2 ? lock_is_held_type+0xbe/0xfd rtnetlink_rcv_msg+0x23f/0x26a ? check_p rtnl_calcit.isra.0+0x147/0x147 netlink_rcv_skb+0x61/0xb2 netlink_unicast+0x100/0x187 netlink_sendmsg+0x37f sock_sendmsg_nosec+0x67/0x9b ____sys_sendmsg+0x19d/0x1f9 ? copy_msghdr_from_user+0x4c/0x5e ? rcu_rea ___sys_sendmsg+0x6c/0x8c ? asm_sysvec_apic_timer_interrupt+0x12/0x20 ? lockdep_hardirqs_on+0xd9/0x102 ? __sys_sendmsg+0x50/0x6e do_syscall_64+0xcb/0xf2 entry_SYSCALL_64_after_hwframe+0x44/0xae RIP: 0033 48 c7 c0 ff ff ff ff eb b5 0f 1f 80 00 00 00 00 48 8d 05 e9 5d 0c 00 8b 00 85 c0 75 13 b8 2e 00 00 00 0f 05 <48> 3d d4 55 48 89 f5 53 RSP: 002b:00007fff859f5e68 EFLAGS: 00000246 ORIG_RAX: 000000000000002e2e RAX: ff RCX: 00007f98dea28914 RDX: 0000000000000000 RSI: 00007fff859f5ed0 RDI: 0000000000000003 RBP: 0000 R09: 0000000000000008 R10: ffffffffffffce6 R11: 0000000000000246 R12: 0000000000000001 R13: 000055c00 00007fff859f63a0 </TASK> Modules linked in: bridge stp llc bonding virtio_net
CVE-2021-47579	In the Linux kernel, the following vulnerability has been resolved: ovl: fix warning in ovl_create_real() Syzbot trig ovl_workdir_create() -> ovl_create_real(): if (!err && WARN_ON(!newdentry->d_inode)) { The reason is that the without instantiating the new dentry. Weird filesystems such as this will be rejected by overlayfs at a later stage du ovl_mkdir_real() directly from ovl_workdir_create() and reject this case early.
CVE-2021-47601	In the Linux kernel, the following vulnerability has been resolved: tee: amdtee: fix an IS_ERR() vs NULL bug The error pointers it returns NULL so fix this condition to avoid a NULL dereference.
CVE-2021-47609	In the Linux kernel, the following vulnerability has been resolved: firmware: arm_scpi: Fix string overflow in SCP scpi_pd->name, it could result in the buffer overflow when copying the SCPI device name from the corresponding maximum size of 30. Let us fix it by using devm_kasprintf so that the string buffer is allocated dynamically.
CVE-2021-47613	In the Linux kernel, the following vulnerability has been resolved: i2c: virtio: fix completion handling The driver c only received when the device is done with all the queued buffers. However, this is not true, since the notify callba buffers being completed (for example, with virtio-pci and shared interrupts) or with only some of the buffers being available to the device in multiple separate virtqueue_add_sgs() calls). This can lead to incorrect data on the I2C bu device operates on buffers which are have been freed by the driver. (The WARN_ON in the driver is also triggered overwritten First byte 0x0 instead of 0x6b Allocated in i2cdev_ioctl_rdwr+0x9d/0x1de age=243 cpu=0 pid=28 me +0x9d/0x1de i2cdev_ioctl+0x247/0x2ed vfs_ioctl+0x21/0x30 sys_ioctl+0xb18/0xb41 Freed in i2cdev_ioctl_rdwr+ +0x1bd/0x1cc i2cdev_ioctl_rdwr+0x1bb/0x1de i2cdev_ioctl+0x247/0x2ed vfs_ioctl+0x21/0x30 sys_ioctl+0xb18/0 from the notify handler like other virtio drivers and by actually waiting for all the buffers to be completed.
CVE-2022-0156	vim is vulnerable to Use After Free

CVE-2022-0158	vim is vulnerable to Heap-based Buffer Overflow
CVE-2022-0213	vim is vulnerable to Heap-based Buffer Overflow
CVE-2022-0261	Heap-based Buffer Overflow in GitHub repository vim/vim prior to 8.2.
CVE-2022-0319	Out-of-bounds Read in vim/vim prior to 8.2.
CVE-2022-0359	Heap-based Buffer Overflow in GitHub repository vim/vim prior to 8.2.
CVE-2022-0361	Heap-based Buffer Overflow in GitHub repository vim/vim prior to 8.2.
CVE-2022-0368	Out-of-bounds Read in GitHub repository vim/vim prior to 8.2.
CVE-2022-0407	Heap-based Buffer Overflow in GitHub repository vim/vim prior to 8.2.
CVE-2022-0408	Stack-based Buffer Overflow in GitHub repository vim/vim prior to 8.2.
CVE-2022-0413	Use After Free in GitHub repository vim/vim prior to 8.2.
CVE-2022-0417	Heap-based Buffer Overflow GitHub repository vim/vim prior to 8.2.
CVE-2022-0443	Use After Free in GitHub repository vim/vim prior to 8.2.
CVE-2022-0500	A flaw was found in unrestricted eBPF usage by the BPF_BTF_LOAD, leading to a possible out-of-bounds memor subsystem due to the way a user loads BTF. This flaw allows a local user to crash or escalate their privileges on the
CVE-2022-0554	Use of Out-of-range Pointer Offset in GitHub repository vim/vim prior to 8.2.
CVE-2022-0563	A flaw was found in the util-linux chfn and chsh utilities when compiled with Readline support. The Readline libra to get a path to the library config file. When the library cannot parse the specified file, it prints an error message co an unprivileged user to read root-owned files, potentially leading to privilege escalation. This flaw affects util-linux
CVE-2022-0572	Heap-based Buffer Overflow in GitHub repository vim/vim prior to 8.2.
CVE-2022-0629	Stack-based Buffer Overflow in GitHub repository vim/vim prior to 8.2.
CVE-2022-0685	Use of Out-of-range Pointer Offset in GitHub repository vim/vim prior to 8.2.4418.
CVE-2022-0696	NULL Pointer Dereference in GitHub repository vim/vim prior to 8.2.4428.
CVE-2022-0714	Heap-based Buffer Overflow in GitHub repository vim/vim prior to 8.2.4436.
CVE-2022-0729	Use of Out-of-range Pointer Offset in GitHub repository vim/vim prior to 8.2.4440.
CVE-2022-0943	Heap-based Buffer Overflow occurs in vim in GitHub repository vim/vim prior to 8.2.4563.
CVE-2022-1154	Use after free in utf_ptr2char in GitHub repository vim/vim prior to 8.2.4646.
CVE-2022-1210	A vulnerability classified as problematic was found in LibTIFF 4.3.0. Affected by this vulnerability is the TIFF Fil leads to a denial of service. The attack can be launched remotely but requires user interaction. The exploit has been
CVE-2022-1292	The c_rehash script does not properly sanitise shell metacharacters to prevent command injection. This script is dis manner where it is automatically executed. On such operating systems, an attacker could execute arbitrary commar c_rehash script is considered obsolete and should be replaced by the OpenSSL rehash command line tool. Fixed in Fixed in OpenSSL 1.1.1o (Affected 1.1.1-1.1.1n). Fixed in OpenSSL 1.0.2ze (Affected 1.0.2-1.0.2zd).
CVE-2022-1350	A vulnerability classified as problematic was found in GhostPCL 9.55.0. This vulnerability affects the function chu The manipulation with a malicious file leads to a memory corruption. The attack can be initiated remotely but requ disclosed to the public as a POC and may be used. It is recommended to apply the patches to fix this issue.
CVE-2022-1420	Use of Out-of-range Pointer Offset in GitHub repository vim/vim prior to 8.2.4774.
CVE-2022-1616	Use after free in append_command in GitHub repository vim/vim prior to 8.2.4895. This vulnerability is capable of Mechanism, Modify Memory, and possible remote execution
CVE-2022-1620	NULL Pointer Dereference in function vim_regexec_string at regexp.c:2729 in GitHub repository vim/vim prior to function vim_regexec_string at regexp.c:2729 allows attackers to cause a denial of service (application crash) via a
CVE-2022-1621	Heap buffer overflow in vim_strncpy find_word in GitHub repository vim/vim prior to 8.2.4919. This vulnerability Protection Mechanism, Modify Memory, and possible remote execution
CVE-2022-1629	Buffer Over-read in function find_next_quote in GitHub repository vim/vim prior to 8.2.4925. This vulnerabilities Memory, and possible remote execution
CVE-2022-1674	NULL Pointer Dereference in function vim_regexec_string at regexp.c:2733 in GitHub repository vim/vim prior to function vim_regexec_string at regexp.c:2733 allows attackers to cause a denial of service (application crash) via a

CVE-2022-1725	NULL Pointer Dereference in GitHub repository vim/vim prior to 8.2.4959.
CVE-2022-1733	Heap-based Buffer Overflow in GitHub repository vim/vim prior to 8.2.4968.
CVE-2022-1769	Buffer Over-read in GitHub repository vim/vim prior to 8.2.4974.
CVE-2022-1771	Uncontrolled Recursion in GitHub repository vim/vim prior to 8.2.4975.
CVE-2022-1785	Out-of-bounds Write in GitHub repository vim/vim prior to 8.2.4977.
CVE-2022-1796	Use After Free in GitHub repository vim/vim prior to 8.2.4979.
CVE-2022-1851	Out-of-bounds Read in GitHub repository vim/vim prior to 8.2.
CVE-2022-1886	Heap-based Buffer Overflow in GitHub repository vim/vim prior to 8.2.
CVE-2022-1897	Out-of-bounds Write in GitHub repository vim/vim prior to 8.2.
CVE-2022-1898	Use After Free in GitHub repository vim/vim prior to 8.2.
CVE-2022-1927	Buffer Over-read in GitHub repository vim/vim prior to 8.2.
CVE-2022-1942	Heap-based Buffer Overflow in GitHub repository vim/vim prior to 8.2.
CVE-2022-1968	Use After Free in GitHub repository vim/vim prior to 8.2.
CVE-2022-1974	A use-after-free flaw was found in the Linux kernel's NFC core functionality due to a race condition between kobje allows a local attacker with CAP_NET_ADMIN privilege to leak kernel information.
CVE-2022-2000	Out-of-bounds Write in GitHub repository vim/vim prior to 8.2.
CVE-2022-20153	In rcu_cblist_dequeue of rcu_segcblist.c, there is a possible use-after-free due to improper locking. This could lead kernel with System execution privileges needed. User interaction is not needed for exploitation.Product: AndroidV A-222091980References: Upstream kernel
CVE-2022-2042	Use After Free in GitHub repository vim/vim prior to 8.2.
CVE-2022-2068	In addition to the c_rehash shell command injection identified in CVE-2022-1292, further circumstances where the shell metacharacters to prevent command injection were found by code review. When the CVE-2022-1292 was fix places in the script where the file names of certificates being hashed were possibly passed to a command executed by some operating systems in a manner where it is automatically executed. On such operating systems, an attacker privileges of the script. Use of the c_rehash script is considered obsolete and should be replaced by the OpenSSL re 3.0.4 (Affected 3.0.0,3.0.1,3.0.2,3.0.3). Fixed in OpenSSL 1.1.1p (Affected 1.1.1-1.1.1o). Fixed in OpenSSL 1.0.2z
CVE-2022-2097	AES OCB mode for 32-bit x86 platforms using the AES-NI assembly optimised implementation will not encrypt th circumstances. This could reveal sixteen bytes of data that was preexisting in the memory that wasn't written. In the bytes of the plaintext would be revealed. Since OpenSSL does not support OCB based cipher suites for TLS and D OpenSSL 3.0.5 (Affected 3.0.0-3.0.4). Fixed in OpenSSL 1.1.1q (Affected 1.1.1-1.1.1p).
CVE-2022-2231	NULL Pointer Dereference in GitHub repository vim/vim prior to 8.2.
CVE-2022-2257	Out-of-bounds Read in GitHub repository vim/vim prior to 9.0.
CVE-2022-2264	Heap-based Buffer Overflow in GitHub repository vim/vim prior to 9.0.
CVE-2022-2289	Use After Free in GitHub repository vim/vim prior to 9.0.
CVE-2022-2304	Stack-based Buffer Overflow in GitHub repository vim/vim prior to 9.0.
CVE-2022-23222	kernel/bpf/verifier.c in the Linux kernel through 5.15.14 allows local users to gain privileges because of the availab *_OR_NULL pointer types.
CVE-2022-24959	An issue was discovered in the Linux kernel before 5.16.5. There is a memory leak in yam_siocdevprivate in drive
CVE-2022-24975	The --mirror documentation for Git through 2.35.1 does not mention the availability of deleted content, aka the "Gi risk if information-disclosure auditing processes rely on a clone operation without the --mirror option. Note: This h believe this is an intended feature of the git binary and does not pose a security risk.
CVE-2022-24975	The --mirror documentation for Git through 2.35.1 does not mention the availability of deleted content, aka the "Gi risk if information-disclosure auditing processes rely on a clone operation without the --mirror option. Note: This h believe this is an intended feature of the git binary and does not pose a security risk.
CVE-2022-25265	In the Linux kernel through 5.16.10, certain binary files may have the exec-all attribute if they were built in approx kernel 2.4.20). This can cause execution of bytes located in supposedly non-executable regions of a file.
CVE-2022-2571	Heap-based Buffer Overflow in GitHub repository vim/vim prior to 9.0.0101.
CVE-2022-2581	Out-of-bounds Read in GitHub repository vim/vim prior to 9.0.0104.

CVE-2022-2598	Out-of-bounds Write to API in GitHub repository vim/vim prior to 9.0.0100.
CVE-2022-26878	drivers/bluetooth/virtio_bt.c in the Linux kernel before 5.16.3 has a memory leak (socket buffers have memory allo
CVE-2022-26966	An issue was discovered in the Linux kernel before 5.16.12. drivers/net/usb/sr9700.c allows attackers to obtain sen crafted frame lengths from a device.
CVE-2022-27672	When SMT is enabled, certain AMD processors may speculatively execute instructions using a target from the sibl potentially resulting in information disclosure.
CVE-2022-2785	There exists an arbitrary memory read within the Linux Kernel BPF - Constants provided to fill pointers in structs and can point anywhere, including memory not owned by BPF. An attacker with CAP_BPF can arbitrarily read me recommend upgrading past commit 86f44fcec22c
CVE-2022-2816	Out-of-bounds Read in GitHub repository vim/vim prior to 9.0.0212.
CVE-2022-2817	Use After Free in GitHub repository vim/vim prior to 9.0.0213.
CVE-2022-2862	Use After Free in GitHub repository vim/vim prior to 9.0.0221.
CVE-2022-2874	NULL Pointer Dereference in GitHub repository vim/vim prior to 9.0.0224.
CVE-2022-2889	Use After Free in GitHub repository vim/vim prior to 9.0.0225.
CVE-2022-29458	ncurses 6.3 before patch 20220416 has an out-of-bounds read and segmentation violation in convert_strings in tinf
CVE-2022-2982	Use After Free in GitHub repository vim/vim prior to 9.0.0260.
CVE-2022-3134	Use After Free in GitHub repository vim/vim prior to 9.0.0389.
CVE-2022-31782	ftbench.c in FreeType Demo Programs through 2.12.1 has a heap-based buffer overflow.
CVE-2022-3278	NULL Pointer Dereference in GitHub repository vim/vim prior to 9.0.0552.
CVE-2022-3297	Use After Free in GitHub repository vim/vim prior to 9.0.0579.
CVE-2022-33068	An integer overflow in the component hb-ot-shape-fallback.cc of Harfbuzz v4.3.0 allows attackers to cause a Denia
CVE-2022-3324	Stack-based Buffer Overflow in GitHub repository vim/vim prior to 9.0.0598.
CVE-2022-3491	Heap-based Buffer Overflow in GitHub repository vim/vim prior to 9.0.0742.
CVE-2022-3520	Heap-based Buffer Overflow in GitHub repository vim/vim prior to 9.0.0765.
CVE-2022-35737	SQLite 1.0.12 through 3.39.x before 3.39.2 sometimes allows an array-bounds overflow if billions of bytes are use
CVE-2022-3591	Use After Free in GitHub repository vim/vim prior to 9.0.0789.
CVE-2022-36227	In libarchive before 3.6.2, the software does not check for an error after calling calloc function that can return with leads to a resultant NULL pointer dereference. NOTE: the discoverer cites this CWE-476 remark but third parties c circumstances, when NULL is equivalent to the 0x0 memory address and privileged code can access it, then writin lead to code execution."
CVE-2022-3629	A vulnerability was found in Linux Kernel. It has been declared as problematic. This vulnerability affects the funct af_vsock.c. The manipulation leads to memory leak. The complexity of an attack is rather high. The exploitation ap apply a patch to fix this issue. VDB-211930 is the identifier assigned to this vulnerability.
CVE-2022-3633	A vulnerability classified as problematic has been found in Linux Kernel. Affected is the function j1939_session_d The manipulation leads to memory leak. It is recommended to apply a patch to fix this issue. The identifier of this v
CVE-2022-37454	The Keccak XKCP SHA-3 reference implementation before fdc6fef has an integer overflow and resultant buffer ov arbitrary code or eliminate expected cryptographic properties. This occurs in the sponge function interface.
CVE-2022-3821	An off-by-one Error issue was discovered in Systemd in format_timespan() function of time-util.c. An attacker cou accuracy that leads to buffer overrun in format_timespan(), leading to a Denial of Service.
CVE-2022-39348	Twisted is an event-based framework for internet applications. Started with version 0.9.4, when the host header doe `twisted.web.vhost.NameVirtualHost` will return a `NoResource` resource which renders the Host header unescape script injection. In practice this should be very difficult to exploit as being able to modify the Host header of a norm in a privileged position. This issue was fixed in version 22.10.0rc1. There are no known workarounds.
CVE-2022-40896	A ReDoS issue was discovered in pygments/lexers/smithy.py in pygments through 2.15.0 via SmithyLexer.
CVE-2022-4141	Heap based buffer overflow in vim/vim 9.0.0946 and below by allowing an attacker to CTRL-W gf in the expressio command.

CVE-2022-41741	NGINX Open Source before versions 1.23.2 and 1.22.1, NGINX Open Source Subscription before versions R2 P1 versions R27 P1 and R26 P1 have a vulnerability in the module ngx_http_mp4_module that might allow a local att resulting in its termination or potential other impact using a specially crafted audio or video file. The issue affects c ngx_http_mp4_module, when the mp4 directive is used in the configuration file. Further, the attack is possible only specially crafted audio or video file with the module ngx_http_mp4_module.
CVE-2022-41742	NGINX Open Source before versions 1.23.2 and 1.22.1, NGINX Open Source Subscription before versions R2 P1 R27 P1 and R26 P1 have a vulnerability in the module ngx_http_mp4_module that might allow a local attacker to c in worker process memory disclosure by using a specially crafted audio or video file. The issue affects only NGINX ngx_http_mp4_module, when the mp4 directive is used in the configuration file. Further, the attack is possible only specially crafted audio or video file with the module ngx_http_mp4_module.
CVE-2022-42916	In curl before 7.86.0, the HSTS check could be bypassed to trick it into staying with HTTP. Using its HSTS suppor directly (instead of using an insecure cleartext HTTP step) even when HTTP is provided in the URL. This mechani in the given URL uses IDN characters that get replaced with ASCII counterparts as part of the IDN conversion, e.g (IDEOGRAPHIC FULL STOP) instead of the common ASCII full stop of U+002E (.). The earliest affected versio
CVE-2022-4304	A timing based side channel exists in the OpenSSL RSA Decryption implementation which could be sufficient to r Bleichenbacher style attack. To achieve a successful decryption an attacker would have to be able to send a very la The vulnerability affects all RSA padding modes: PKCS#1 v1.5, RSA-OEAP and RSASVE. For example, in a TLS client to send an encrypted pre-master secret to the server. An attacker that had observed a genuine connection betw to send trial messages to the server and record the time taken to process them. After a sufficiently large number of master secret used for the original connection and thus be able to decrypt the application data sent over that connec
CVE-2022-43551	A vulnerability exists in curl <7.87.0 HSTS check that could be bypassed to trick it to keep using HTTP. Using its HTTPS instead of using an insecure clear-text HTTP step even when HTTP is provided in the URL. However, the host name in the given URL first uses IDN characters that get replaced to ASCII counterparts as part of the IDN co +3002 (IDEOGRAPHIC FULL STOP) instead of the common ASCII full stop (U+002E) `.`. Then in a subsequent makes a clear text transfer. Because it would store the info IDN encoded but look for it IDN decoded.
CVE-2022-4415	A vulnerability was found in systemd. This security flaw can cause a local information leak due to systemd-coredu kernel setting.
CVE-2022-4450	The function PEM_read_bio_ex() reads a PEM file from a BIO and parses and decodes the "name" (e.g. "CERTIFI payload data. If the function succeeds then the "name_out", "header" and "data" arguments are populated with poin decoded data. The caller is responsible for freeing those buffers. It is possible to construct a PEM file that results in PEM_read_bio_ex() will return a failure code but will populate the header argument with a pointer to a buffer that also frees this buffer then a double free will occur. This will most likely lead to a crash. This could be exploited by supply malicious PEM files for parsing to achieve a denial of service attack. The functions PEM_read_bio() and PE PEM_read_bio_ex() and therefore these functions are also directly affected. These functions are also called indirec functions including PEM_X509_INFO_read_bio_ex() and SSL_CTX_use_serverinfo_file() which are also vulnera functions are not vulnerable because the caller does not free the header argument if PEM_read_bio_ex() returns a f PEM_read_bio_TYPE() functions as well as the decoders introduced in OpenSSL 3.0. The OpenSSL asn1parse co this issue.
CVE-2022-47024	A null pointer dereference issue was discovered in function gui_x11_create_blank_mouse in gui_x11.c in vim 8.1.2 denial of service or other unspecified impacts.
CVE-2022-47695	An issue was discovered Binutils objdump before 2.39.3 allows attackers to cause a denial of service or other unspe bfd_mach_o_get_synthetic_symtab in match-o.c.
CVE-2022-48063	GNU Binutils before 2.40 was discovered to contain an excessive memory consumption vulnerability via the functi The attacker could supply a crafted ELF file and cause a DNS attack.
CVE-2022-48065	GNU Binutils before 2.40 was discovered to contain a memory leak vulnerability var the function find_abstract_ins
CVE-2022-48303	GNU Tar through 1.34 has a one-byte out-of-bounds read that results in use of uninitialized memory for a condition control has not been demonstrated. The issue occurs in from_header in list.c via a V7 archive in which mtime has a
CVE-2022-48622	In GNOME GdkPixbuf (aka gdk-pixbuf) through 2.42.10, the ANI (Windows animated cursor) decoder encounters in io-ani.c) when parsing chunks in a crafted .ani file. A crafted file could allow an attacker to overwrite heap meta execution attack. This occurs in gdk_pixbuf_set_option() in gdk-pixbuf.c.
CVE-2022-48627	In the Linux kernel, the following vulnerability has been resolved: vt: fix memory overlapping when deleting chars occurs when deleting a long line. This memory overlapping copy can cause data corruption when scr_memcpyw is not ensure its behavior if the destination buffer overlaps with the source buffer. The line buffer is not always broke acceleration, whose result is not deterministic. Fix this problem by using replacing the scr_memcpyw with scr_mem
CVE-2022-48632	In the Linux kernel, the following vulnerability has been resolved: i2c: mlxbf: prevent stack overflow in mlxbf_i2c called in a loop while 'operation->length' upper bound is not checked and 'data_idx' also increments.

CVE-2022-48644	In the Linux kernel, the following vulnerability has been resolved: net/sched: taprio: avoid disabling offload when API design decision, qdisc->destroy() gets called even if qdisc->init() never succeeded, not exclusively since comm recovery at qdisc creation"), but apparently also earlier (in the case of qdisc_create_dflt()). The taprio qdisc does n full offload, because it starts off with q->flags = TAPRIO_FLAGS_INVALID in taprio_init(), then it replaces q->f parsed from netlink (in taprio_change(), tail called from taprio_init()). But in taprio_destroy(), we call taprio_disab do based on FULL_OFFLOAD_IS_ENABLED(q->flags). But looking at the implementation of FULL_OFFLOAD bit 1 in q->flags), it is invalid to call this macro on q->flags when it contains TAPRIO_FLAGS_INVALID, becaus FULL_OFFLOAD_IS_ENABLED() will return true on an invalid set of flags. As a result, it is possible to crash th setting q->flags = TAPRIO_FLAGS_INVALID, and the calling of taprio_enable_offload(). This is because drivers when it was never enabled. The error that we force here is to attach taprio as a non-root qdisc, but instead as child dev swp0 root handle 1: \ mqprio num_tc 8 map 0 1 2 3 4 5 6 7 \ queues 1@0 1@1 1@2 1@3 1@4 1@5 1@6 1@7 1:1 \ taprio num_tc 8 map 0 1 2 3 4 5 6 7 \ queues 1@0 1@1 1@2 1@3 1@4 1@5 1@6 1@7 base-time 0 \ sched-e 100000 \ flags 0x0 clockid CLOCK_TAI Unable to handle kernel paging request at virtual address ffffffffffffffff8 [p4d=0000000000000000 Internal error: Oops: 96000004 [#1] PREEMPT SMP Call trace: taprio_dump+0x27c/0x felix_port_setup_tc+0x24/0x3c dsa_slave_setup_tc+0x54/0x27c taprio_disable_offload.isra.0+0x58/0xe0 taprio_d +0x240/0x470 tc_modify_qdisc+0x1fc/0x6b0 rtnetlink_rcv_msg+0x12c/0x390 netlink_rcv_skb+0x5c/0x130 rtnetl track of the operations we made, and undo the offload only if we actually did it. I've added "bool offloaded" inside "atomic64_t picos_per_byte". Now the first cache line looks like below: $ pahole -C taprio_sched net/sched/sch_ta qdiscs; /* 0 8 */ struct Qdisc * root; /* 8 8 */ u32 flags; /* 16 4 */ enum tk_offsets tk_offset; /* 20 4 */ int clockid; 3 bytes hole, try to pack */ atomic64_t picos_per_byte; /* 32 0 */ /* XXX 8 bytes hole, try to pack */ spinlock_t cu hole, try to pack */ struct sched_entry * current_entry; /* 48 8 */ struct sched_gate_list * oper_sched; /* 56 8 */ /*
CVE-2022-48655	In the Linux kernel, the following vulnerability has been resolved: firmware: arm_scmi: Harden accesses to the res descriptors by the index upon the SCMI drivers requests through the SCMI reset operations interface can potentiall driver misbehave. Add an internal consistency check before any such domains descriptors accesses.
CVE-2022-48657	In the Linux kernel, the following vulnerability has been resolved: arm64: topology: fix possible overflow in amu_ returns max frequency in kHz as *unsigned int*, while freq_inv_set_max_ratio() gets passed this frequency in Hz a can potentially result in overflow -- multiplying by 1000ULL instead should avoid that... Found by Linux Verificat static analysis tool.
CVE-2022-48660	In the Linux kernel, the following vulnerability has been resolved: gpiolib: cdev: Set lineevent_state::irq after IRQ test on nxp-ls1028 platform with below command gpiomon --num-events=3 --rising-edge gpiochip1 25 There will free_irq+0x204/0x360 lineevent_free+0x64/0x70 gpio_ioctl+0x598/0x6a0 __arm64_sys_ioctl+0xb4/0x100 invoke +0x1a0/0x1a4 The reason of this issue is that calling request_threaded_irq() function failed, and then lineevent_fre Since the lineevent_state::irq was already set, so the subsequent invocation of free_irq() would trigger the above w lineevent_state::irq after the IRQ register successfully.
CVE-2022-48674	In the Linux kernel, the following vulnerability has been resolved: erofs: fix pcluster use-after-free on UP platform disabled, KASAN reports as below: ========== free in __mutex_lock+0xe5/0xc30 Read of size 8 at addr ffff8881094223f8 by task stress/7789 CPU: 0 PID: 7789 g0d53d2e882f9 #3 Hardware name: Red Hat KVM, BIOS 0.5.1 01/01/2011 Call Trace: <TASK> .. __mutex_lock +0x8ce/0x1560 .. z_erofs_readahead+0x31c/0x580 .. Freed by task 7787 kasan_save_stack+0x1e/0x40 kasan_set_ +0x20/0x40 __kasan_slab_free+0x10c/0x190 kmem_cache_free+0xed/0x380 rcu_core+0x3d5/0xc90 __do_softirq creation: kasan_save_stack+0x1e/0x40 __kasan_record_aux_stack+0x97/0xb0 call_rcu+0x3d/0x3f0 erofs_shrink_ +0xdc/0x170 shrink_slab.constprop.0+0x296/0x530 drop_slab+0x1c/0x70 drop_caches_sysctl_handler+0x70/0x8 vfs_write+0x555/0x6c0 ksys_write+0xbe/0x160 do_syscall_64+0x3b/0x90 The root cause is that erofs_workgroup it causes a race that the pcluster reuses unexpectedly before freeing. Since UP platforms are quite rare now, such pa specific-designed path directly instead.
CVE-2022-48675	In the Linux kernel, the following vulnerability has been resolved: IB/core: Fix a nested dead lock as part of ODP f ODP flow by using mmput_async(). From the below call trace [1] can see that calling mmput() once we have the u required by ib_umem_odp_map_dma_and_lock() might trigger in the same task the exit_mmap()->__mmu_notifie which may dead lock when trying to lock the same mutex. Moving to use mmput_async() will solve the problem as be called in other task and will be executed once the lock will be available. [1] [64843.077665] task:kworker/u133: 2 flags:0x00004000 [64843.077672] Workqueue: mlx5_ib_page_fault mlx5_ib_eqe_pf_action [mlx5_ib] [64843.0 <TASK> [64843.077724] __schedule+0x23d/0x590 [64843.077729] schedule+0x4e/0xb0 [64843.077735] schedul [64843.077740] __mutex_lock.constprop.0+0x263/0x490 [64843.077747] __mutex_lock_slowpath+0x13/0x20 [6 [64843.077758] mlx5_ib_invalidate_range+0x48/0x270 [mlx5_ib] [64843.077808] __mmu_notifier_release+0x1a +0x1bc/0x200 [64843.077822] ? walk_page_range+0x9c/0x120 [64843.077828] ? __cond_resched+0x1a/0x50 [6 +0x13/0x40 [64843.077839] ? uprobe_clear_state+0xac/0x120 [64843.077860] mmput+0x5f/0x140 [64843.07786 +0x21b/0x580 [ib_core] [64843.077931] pagefault_real_mr+0x9a/0x140 [mlx5_ib] [64843.077962] pagefault_mr pagefault_single_data_segment.constprop.0+0x2ac/0x560 [mlx5_ib] [64843.078022] mlx5_ib_eqe_pf_action+0x5 process_one_work+0x22b/0x3d0 [64843.078059] worker_thread+0x53/0x410 [64843.078065] ? process_one_wor +0x12a/0x150 [64843.078079] ? set_kthread_struct+0x50/0x50 [64843.078085] ret_from_fork+0x22/0x30 [64843

CVE-2022-48689	In the Linux kernel, the following vulnerability has been resolved: tcp: TX zerocopy should not sense pfmemalloc [1] showing a possible misuse of pfmemalloc page status in TCP zerocopy paths. Indeed, for pages coming from us page_is_pfmemalloc() is moot, and possibly could give false positives. There has been attempts to make page_is_p using it in the first place in this context is probably better, removing cpu cycles. Note to stable teams : You need to introduce __skb_fill_page_desc_noacc") as a prereq. Race is more probable after commit c07aea3ef4d4 ("mm: add page_is_pfmemalloc() is now using low order bit from page->lru.next, which can change more often than page->in be set for lru.next (when used as an anchor in LRU list), so KCSAN report is mostly a false positive. Backporting t necessary. [1] BUG: KCSAN: data-race in lru_add_fn / tcp_build_frag write to 0xffffea0004a1d2c8 of 8 bytes by t linux/list.h:73 [inline] list_add include/linux/list.h:88 [inline] lruvec_add_folio include/linux/mm_inline.h:105 [inli swap.c:228 folio_batch_move_lru+0x1e1/0x2a0 mm/swap.c:246 folio_batch_add_and_move mm/swap.c:263 [inli swap.c:490 filemap_add_folio+0xf8/0x150 mm/filemap.c:948 __filemap_get_folio+0x510/0x6d0 mm/filemap.c:19 mm/folio-compat.c:104 grab_cache_page_write_begin+0x2a/0x30 mm/folio-compat.c:116 ext4_da_write_begin+0 generic_perform_write+0x1d4/0x3f0 mm/filemap.c:3738 ext4_buffered_write_iter+0x235/0x3e0 fs/ext4/file.c:270 call_write_iter include/linux/fs.h:2187 [inline] new_sync_write fs/read_write.c:491 [inline] vfs_write+0x468/0x76 fs/read_write.c:631 __do_sys_write fs/read_write.c:643 [inline] __se_sys_write fs/read_write.c:640 [inline] __x64_ do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x2b/0x70 arch/x86/entry/common.c:80 entry read to 0xffffea0004a1d2c8 of 8 bytes by task 18611 on cpu 1: page_is_pfmemalloc include/linux/mm.h:1740 [inli linux/skbuff.h:2422 [inline] skb_fill_page_desc include/linux/skbuff.h:2443 [inline] tcp_build_frag+0x613/0xb20 +0x3e8/0xaf0 net/ipv4/tcp.c:1075 tcp_sendpage_locked net/ipv4/tcp.c:1140 [inline] tcp_sendpage+0x89/0xb0 net/ net/ipv4/af_inet.c:833 kernel_sendpage+0x184/0x300 net/socket.c:3561 sock_sendpage+0x5a/0x70 net/socket.c:1(fs/splice.c:361 splice_from_pipe_feed fs/splice.c:415 [inline] __splice_from_pipe+0x222/0x4d0 fs/splice.c:559 spl generic_splice_sendpage+0x89/0xc0 fs/splice.c:743 do_splice_from fs/splice.c:764 [inline] direct_splice_actor+0x +0x305/0x620 fs/splice.c:886 do_splice_direct+0xfb/0x180 fs/splice.c:974 do_sendfile+0x3bf/0x910 fs/read_write read_write.c:1317 [inline] __se_sys_sendfile64 fs/read_write.c:1303 [inline] __x64_sys_sendfile64+0x10c/0x150 x86/entry/common.c:50 [inline] do_syscall_64+0x2b/0x70 arch/x86/entry/common.c:80 entry_SYSCALL_64_afte 0x0000000000000000 -> 0xffffea0004a1d288 Reported by Kernel Concurrency Sanitizer on: CPU: 1 PID: 18611 C syzkaller-00248-ge022620b5d05-dirty #0 Hardware name: Google Google Compute Engine/Google Compute Eng
CVE-2022-48707	In the Linux kernel, the following vulnerability has been resolved: cxl/region: Fix null pointer dereference for reset callback. The CXL specification allows a host bridge with a single root port to have no explicit HDM decoders. Cu none. As such the CXL core creates a special pass through decoder instance without a commit/reset callback. Prior unconditionally when calling cxl_region_decode_reset. Thus a configuration with 1 Host Bridge, 1 Root Port, and multiple CXL type 3 devices attached to downstream ports of a switch can cause a null pointer dereference. Before destroy the region, and a pass through decoder is reset. The issue can be reproduced as below, 1) create a region wi single root port under which a memdev is attached directly. 2) destroy the region with cxl destroy-region regionX -
CVE-2022-48719	In the Linux kernel, the following vulnerability has been resolved: net, neigh: Do not trigger immediate probes on syzkaller was able to trigger a deadlock for NTF_MANAGED entries [0]: kworker/0:16/14617 is trying to acquire {++-.}-{2:2}, at: ___neigh_create+0x9e1/0x2990 net/core/neighbour.c:652 [...] but task is already holding lock: fff {2:2}, at: neigh_managed_work+0x35/0x250 net/core/neighbour.c:1572 The neighbor entry turned to NUD_FAIL triggered an immediate probe as per commit cd28ca0a3dd1 ("neigh: reduce arp latency") via neigh_probe() given t situation is to defer the neigh_probe() back to the neigh_timer_handler() similarly as pre cd28ca0a3dd1. For the ca is acceptable given this only happens on actual failure state and regular / expected state is NUD_VALID with the e parameter to __neigh_event_send() in order to communicate whether immediate probe is allowed or disallowed. Ex default as-is to immediate probe. However, the neigh_managed_work() disables it via use of neigh_event_send_pro dump_stack.c:88 [inline] dump_stack_lvl+0xcd/0x134 lib/dump_stack.c:106 print_deadlock_bug kernel/locking/lo kernel/locking/lockdep.c:2999 [inline] validate_chain kernel/locking/lockdep.c:3788 [inline] __lock_acquire.cold+ lock_acquire kernel/locking/lockdep.c:5639 [inline] lock_acquire+0x1ab/0x510 kernel/locking/lockdep.c:5604 __r rwlock_api_smp.h:202 [inline] _raw_write_lock_bh+0x2f/0x40 kernel/locking/spinlock.c:334 ___neigh_create+0x ip6_finish_output2+0x1070/0x14f0 net/ipv6/ip6_output.c:123 __ip6_finish_output net/ipv6/ip6_output.c:191 [inlin net/ipv6/ip6_output.c:170 ip6_finish_output+0x32/0x200 net/ipv6/ip6_output.c:201 NF_HOOK_COND include/li +0x1e4/0x530 net/ipv6/ip6_output.c:224 dst_output include/net/dst.h:451 [inline] NF_HOOK include/linux/netfilt +0xa99/0x17f0 net/ipv6/ndisc.c:508 ndisc_send_ns+0x3a9/0x840 net/ipv6/ndisc.c:650 ndisc_solicit+0x2cd/0x4f0 +0xc2/0x110 net/core/neighbour.c:1040 __neigh_event_send+0x37d/0x1570 net/core/neighbour.c:1201 neigh_eve neigh_managed_work+0x162/0x250 net/core/neighbour.c:1574 process_one_work+0x9ac/0x1650 kernel/workque kernel/workqueue.c:2454 kthread+0x2e9/0x3a0 kernel/kthread.c:377 ret_from_fork+0x1f/0x30 arch/x86/entry/entr
CVE-2022-48720	In the Linux kernel, the following vulnerability has been resolved: net: macsec: Fix offload support for NETDEV_ notify handler handles NETDEV_UNREGISTER event by releasing relevant SW resources only, this causes resou the underlay driver was not notified to clean it's macsec offload resources. Fix by calling the underlay driver to clea handling from macsec_dellink() to macsec_common_dellink() when handling NETDEV_UNREGISTER event.
CVE-2022-48727	In the Linux kernel, the following vulnerability has been resolved: KVM: arm64: Avoid consuming a stale esr valu exception other than an IRQ occurs, the CPU updates the ESR_EL2 register with the exception syndrome. An SErr synchronised by KVM. KVM notes the exception type, and whether an SError was synchronised in exit_code. Whe fixup_guest_exit() updates vcpu->arch.fault.esr_el2 from the hardware register. When an SError was synchronised exception was due to an HVC. If so, ELR_EL2 is moved back one instruction. This is so that KVM can process the guest survives the SError. But if an IRQ synchronises an SError, the vcpu's esr value is stale. If the previous non-IF ELR_EL2, causing an unrelated guest instruction to be executed twice. Check ARM_EXCEPTION_CODE() befor this register so don't need to check.

CVE-2022-48741	In the Linux kernel, the following vulnerability has been resolved: ovl: fix NULL pointer dereference in copy up w dereference to get a recently introduced warning message working.
CVE-2022-48754	In the Linux kernel, the following vulnerability has been resolved: phylib: fix potential use-after-free Commit bafb support") added call to phy_device_reset(phydev) after the put_device() call in phy_detach(). The comment before might go away with put_device(). Fix potential use-after-free by calling phy_device_reset() before put_device().
CVE-2022-48757	In the Linux kernel, the following vulnerability has been resolved: net: fix information leakage in /proc/net/ptype I socket without binding it to a device, users in other net namespaces can observe the new `packet_type` added by th file. This is minor information leakage as packet socket is namespace aware. Add a net pointer in `packet_type` to packet socket. In `ptype_seq_show`, this net pointer must be checked when it is not NULL.
CVE-2022-48760	In the Linux kernel, the following vulnerability has been resolved: USB: core: Fix hang in usb_kill_urb by adding identified a bug in which processes hang waiting for usb_kill_urb() to return. It turns out the issue is not unlinking the problem arises when the wakeup notification that the URB has completed is not received. The reason is memor In outline form, usb_kill_urb() and __usb_hcd_giveback_urb() operating concurrently on different CPUs perform t 1 ---------- ---------- usb_kill_urb(): __usb_hcd_giveback_urb(): ... ... atomic_inc(>use_count); ... ... wait_event(usb_kill_urb_queue, atomic_read(&urb->use_count) == 0); if (atomic_read(&urb-> Confining your attention to urb->reject and urb->use_count, you can see that the overall pattern of accesses on CPU >use_count; whereas the overall pattern of accesses on CPU 1 is: write urb->use_count, then read urb->reject. This circles as SB (for "Store Buffering"), and it is well known that without suitable enforcement of the desired order of it is entirely possible for one or both CPUs to execute their reads ahead of their writes. The end result will be that s value of urb->use_count while CPU 1 sees the old un-incremented value of urb->reject. Consequently CPU 0 ends up, leading to the observed hang in usb_kill_urb(). The same pattern of accesses occurs in usb_poison_urb() and th The problem is fixed by adding suitable memory barriers. To provide proper memory-access ordering in the SB pa CPUs. The atomic_inc() and atomic_dec() accesses themselves don't provide any memory ordering, but since they smp_mb__after_atomic() memory barrier in the various routines to obtain the desired effect. This patch adds the ne
CVE-2022-48768	In the Linux kernel, the following vulnerability has been resolved: tracing/histogram: Fix a potential memory leak path to free the memory allocated by kstrdup(): p = param = kstrdup(data->params[i], GFP_KERNEL); So it is bet
CVE-2022-48773	In the Linux kernel, the following vulnerability has been resolved: xprtrdma: fix pointer derefs in error cases of rpc must not leave the non-NULL pointers with the error value, otherwise `rpcrdma_ep_destroy` gets confused and trie
CVE-2022-48781	In the Linux kernel, the following vulnerability has been resolved: crypto: af_alg - get rid of alg_memory_allocate not seem to be really used. alg_proto does have a .memory_allocated field, but no corresponding .sysctl_mem. This true, but all sk_prot_mem_limits() users will trigger a NULL dereference [1]. THis was not a problem until SO_RE protection fault, probably for non-canonical address 0xdffffc0000000001: 0000 [#1] PREEMPT SMP KASAN KA [0x0000000000000008-0x000000000000000f] CPU: 1 PID: 3591 Comm: syz-executor153 Not tainted 5.17.0-rc3- Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google 01/01/2011 RIP: 0010:s [inline] RIP: 0010:sock_reserve_memory+0x1d7/0x330 net/core/sock.c:1000 Code: 08 00 74 08 48 89 ef e8 27 20 c5 08 48 89 e8 48 c1 e8 03 48 b9 00 00 00 00 00 fc ff df <80> 3c 08 00 74 08 48 89 ef e8 fb 1f bb f9 48 8b 6d 00 4 EFLAGS: 00010202 RAX: 0000000000000001 RBX: ffff88814aabc000 RCX: dffffc0000000000 RDX: 00000000 RDI: ffffffff90e18120 RBP: 0000000000000008 R08: dffffc0000000000 R09: fffffbfff21c3025 R10: fffffbfff21c3 ffffffff8d109840 R13: 0000000000001002 R14: 0000000000000001 R15: 0000000000000001 FS: 0000555556e08 knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 00007fc74416f130 CR3: 00000000003506e0 DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 DR3: 0000000 DR7: 0000000000000400 Call Trace: <TASK> sock_setsockopt+0x14a9/0x3a30 net/core/sock.c:1446 __sys_setsc __do_sys_setsockopt net/socket.c:2191 [inline] __se_sys_setsockopt net/socket.c:2188 [inline] __x64_sys_setsock do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x44/0xd0 arch/x86/entry/common.c:80 entry RIP: 0033:0x7fc7440fddc9 Code: 28 00 00 00 75 05 48 83 c4 28 c3 e8 51 15 00 00 90 48 89 f8 48 89 f7 48 89 d6 4 0f 05 <48> 3d 01 f0 ff ff 73 01 c3 48 c7 c1 c0 ff ff ff f7 d8 64 89 01 48 RSP: 002b:00007ffe98f07968 EFLAGS: 0 RAX: ffffffffffffffda RBX: 0000000000000003 RCX: 00007fc7440fddc9 RDX: 0000000000000049 RSI: 0000000 RBP: 0000000000000000 R08: 0000000000000004 R09: 00007ffe98f07990 R10: 0000000020000000 R11: 00000 R13: 00007ffe98f079a0 R14: 00007ffe98f079e0 R15: 0000000000000000 </TASK> Modules linked in: ---[end tr 0010:sk_prot_mem_limits include/net/sock.h:1523 [inline] RIP: 0010:sock_reserve_memory+0x1d7/0x330 net/co ef e8 27 20 bb f9 4c 03 7c 24 10 48 8b 6d 00 48 83 c5 08 48 89 e8 48 c1 e8 03 48 b9 00 00 00 00 00 fc ff df <80> 48 8b 6d 00 4c 89 ff 48 RSP: 0018:ffffc90001f1fb68 EFLAGS: 00010202 RAX: 0000000000000001 RBX: ffff888 RDX: 0000000000000001 RSI: 0000000000000008 RDI: ffffffff90e18120 RBP: 0000000000000008 R08: dffffc0 R10: fffffbfff21c3025 R11: 0000000000000000 R12: ffffffff8d109840 R13: 0000000000001002 R14: 0000000000 FS: 0000555556e08300(0000) GS:ffff8880b9b00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 00 00007fc74416f130 CR3: 0000000073d9e000 CR4: 00000000003506e0 DR0: 0000000000000000 DR1: 00000000

CVE-2022-48785	In the Linux kernel, the following vulnerability has been resolved: ipv6: mcast: use rcu-safe version of ipv6_get_ll ("ipv6,mcast: always hold idev->lock before mca_lock") switched ipv6_get_lladdr() to __ipv6_get_lladdr(), which because idev->lock was held for these codepaths. In 88e2ca308094 ("mld: convert ifmcaddr6 to RCU") these exter need to restore the original rcu-safe call. Otherwise, we occasionally get a machine crashed/stalled with the followi general protection fault, probably for non-canonical address 0xdead00000000008c: 0000 [#1] SMP NOPTI [3405. Comm: kworker/44:3 Tainted: G O 5.15.19-cloudflare-2022.2.1 #1 [3405.998061][T230589] Hardware name: SU Workqueue: mld mld_ifc_work [3406.017224][T230589] RIP: 0010:__ipv6_get_lladdr+0x34/0x60 [3406.02578 89 e5 48 39 d7 74 3e 48 8d 82 38 ff ff ff eb 13 48 8b 90 d0 00 00 00 48 8d 82 38 ff ff ff 48 39 d7 74 22 <66> 83 7 dd 48 8b 50 08 48 8b [3406.055748][T230589] RSP: 0018:ffff94e4b3fc3d10 EFLAGS: 00010202 [3406.065617] RBX: ffff94e4b3fc3d30 RCX: 0000000000000040 [3406.077477][T230589] RDX: dead000000000122 RSI: ffff9 [3406.089389][T230589] RBP: ffff94e4b3fc3d10 R08: 0000000000000000 R09: 0000000000000000 [3406.1014 000000000000000b R12: ffff8c2c37887100 [3406.113553][T230589] R13: ffff8c3a39537000 R14: 00000000000 [T230589] FS: 0000000000000000(0000) GS:ffff8c3b9fc80000(0000) knlGS:0000000000000000 [3406.138992]| CR0: 0000000080050033 [3406.149895][T230589] CR2: 00007f0dfea1db60 CR3: 000000387b5f2000 CR4: 000 Call Trace: [3406.170235][T230589] <TASK> [3406.177736][T230589] mld_newpack+0xfe/0x1a0 [3406.1866 [3406.195498][T230589] add_grec+0x485/0x4e0 [3406.204310][T230589] ? newidle_balance+0x126/0x3f0 [34 +0x15d/0x450 [3406.223279][T230589] process_one_work+0x1e6/0x380 [3406.232982][T230589] worker_thre rescuer_thread+0x360/0x360 [3406.252175][T230589] kthread+0x127/0x150 [3406.261197][T230589] ? set_kth [T230589] ret_from_fork+0x22/0x30 [3406.280812][T230589] </TASK> [3406.288937][T230589] Modules link [3406.476714][T230589] ---[end trace 3525a7655f2f3b9e]---
CVE-2022-48791	In the Linux kernel, the following vulnerability has been resolved: scsi: pm8001: Fix use-after-free for aborted TM occur if a TMF sas_task is aborted before we handle the IO completion in mpi_ssp_completion(). The abort occurs SAS_TASK_STATE_ABORTED flag is set and the sas_task is freed in pm8001_exec_internal_tmf_task(). Howe O completion still thinks that the sas_task is available. Fix this by clearing the ccb->task if the TMF times out - the pointer is cleared.
CVE-2022-48797	In the Linux kernel, the following vulnerability has been resolved: mm: don't try to NUMA-migrate COW pages th that enabling NUMA balancing causes corruption with his Gaudi accelerator test load: "All the details are in the bu patch causes corruption when the numa balancing feature is enabled AND we don't use process affinity AND we us DMA to/from system memory. Either disabling numa balancing, using process affinity to bind to specific numa-no to disappear" and Oded bisected the issue to commit 09854ba94c6a ("mm: do_wp_page() simplification"). Now, th changing the writability of a page, and as such shouldn't matter for COW. But it appears it does. Suspicious. Howe enabling NUMA faults in change_pte_range() is nonsensical. It uses "page_mapcount(page)" to decide if a COW p and that makes absolutely no sense. The number of mappings a page has is irrelevant: not only does GUP get a refe other mappings migth be paged out and the only reference to them would be in the page count. Since we should ne can't move anyway due to other references, just fix the code to use 'page_count()'. Oded confirms that that fixes his in NUMA balancing ends up changing page protections (other than the obvious one of making the page inaccessibl Otherwise the COW simplification wouldn't matter - since doing the GUP on the page would make sure it's writabl be good to figure out too, since it clearly results in spurious COW events - but fixing the nonsensical test that just h CorrectThing(tm) to do regardless.
CVE-2022-48812	In the Linux kernel, the following vulnerability has been resolved: net: dsa: lantiq_gswip: don't use devres for mdio ("net: dsa: realtek: register the MDIO bus under devres") 5135e96a3dd2 ("net: dsa: don't allocate the slave_mii_bu when called from devm_mdiobus_free() <- devres_release_all() <- __device_release_driver(), and that mdiobus wa switch is a platform device, so the initial set of constraints that I thought would cause this (I2C or SPI buses which But there is one more which applies here. If the DSA master itself is on a bus that calls ->remove from ->shutdown there is a device link between the switch and the DSA master, and device_links_unbind_consumers() will unbind t the same treatment must be applied to all DSA switch drivers, which is: either use devres for both the mdiobus allo all. The gswip driver has the code structure in place for orderly mdiobus removal, so just replace devm_mdiobus_a manual free where necessary, to ensure that we don't let devres free a still-registered bus.
CVE-2022-48813	In the Linux kernel, the following vulnerability has been resolved: net: dsa: felix: don't use devres for mdiobus As dsa: realtek: register the MDIO bus under devres") 5135e96a3dd2 ("net: dsa: don't allocate the slave_mii_bus using called from devm_mdiobus_free() <- devres_release_all() <- __device_release_driver(), and that mdiobus was not switch is a PCI device, so the initial set of constraints that I thought would cause this (I2C or SPI buses which call there is one more which applies here. If the DSA master itself is on a bus that calls ->remove from ->shutdown (lik there is a device link between the switch and the DSA master, and device_links_unbind_consumers() will unbind t same treatment must be applied to all DSA switch drivers, which is: either use devres for both the mdiobus allocati The felix driver has the code structure in place for orderly mdiobus removal, so just replace devm_mdiobus_alloc_ manual free where necessary, to ensure that we don't let devres free a still-registered bus.

CVE-2022-48814	In the Linux kernel, the following vulnerability has been resolved: net: dsa: seville: register the mdiobus under dev ("net: dsa: realtek: register the MDIO bus under devres") 5135e96a3dd2 ("net: dsa: don't allocate the slave_mii_bu when called from devm_mdiobus_free() <- devres_release_all() <- __device_release_driver(), and that mdiobus wa VSC9959 switch is a platform device, so the initial set of constraints that I thought would cause this (I2C or SPI bu do not apply. But there is one more which applies here. If the DSA master itself is on a bus that calls ->remove fro the fsl-mc bus), there is a device link between the switch and the DSA master, and device_links_unbind_consumer shutdown. So the same treatment must be applied to all DSA switch drivers, which is: either use devres for both th use devres at all. The seville driver has a code structure that could accommodate both the mdiobus_unregister and dependency upon mscc_miim_setup() from mdio-mscc-miim.c, which calls devm_mdiobus_alloc_size() on its beh exporting yet one more symbol mscc_miim_teardown(), let's work with devres and replace of_mdiobus_register w devres, we can ensure that devres doesn't free a still-registered bus (it either runs both callbacks, or none).
CVE-2022-48815	In the Linux kernel, the following vulnerability has been resolved: net: dsa: bcm_sf2: don't use devres for mdiobus ("net: dsa: realtek: register the MDIO bus under devres") 5135e96a3dd2 ("net: dsa: don't allocate the slave_mii_bu when called from devm_mdiobus_free() <- devres_release_all() <- __device_release_driver(), and that mdiobus wa 2 is a platform device, so the initial set of constraints that I thought would cause this (I2C or SPI buses which call - there is one more which applies here. If the DSA master itself is on a bus that calls ->remove from ->shutdown (lik there is a device link between the switch and the DSA master, and device_links_unbind_consumers() will unbind t the same treatment must be applied to all DSA switch drivers, which is: either use devres for both the mdiobus allo all. The bcm_sf2 driver has the code structure in place for orderly mdiobus removal, so just replace devm_mdiobus manual free where necessary, to ensure that we don't let devres free a still-registered bus.
CVE-2022-48816	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: lock against ->sock changing during asynchronously unless ->recv_mutex is held. So it is important to hold that mutex. Otherwise a sysfs read can trigg ("SUNRPC: Check if the xprt is connected before handling sysfs reads") appears to attempt to fix this problem, but
CVE-2022-48817	In the Linux kernel, the following vulnerability has been resolved: net: dsa: ar9331: register the mdiobus under dev ("net: dsa: realtek: register the MDIO bus under devres") 5135e96a3dd2 ("net: dsa: don't allocate the slave_mii_bu when called from devm_mdiobus_free() <- devres_release_all() <- __device_release_driver(), and that mdiobus wa an MDIO device, so the initial set of constraints that I thought would cause this (I2C or SPI buses which call ->rem is one more which applies here. If the DSA master itself is on a bus that calls ->remove from ->shutdown (like dpa a device link between the switch and the DSA master, and device_links_unbind_consumers() will unbind the ar933 treatment must be applied to all DSA switch drivers, which is: either use devres for both the mdiobus allocation an ar9331 driver doesn't have a complex code structure for mdiobus removal, so just replace of_mdiobus_register with and ensure that we don't free a still-registered bus.
CVE-2022-48819	In the Linux kernel, the following vulnerability has been resolved: tcp: take care of mixed splice()/sendmsg(MSG_ mixing sendpage() and sendmsg(MSG_ZEROCOPY) calls over the same TCP socket would again trigger the infa WARN_ON(sk_forward_alloc_get(sk)); While Talal took into account a mix of regular copied data and MSG_ZE path has been forgotten. We want the charging to happen for sendpage(), because pages could be coming from a pi pure zerocopy status to make sure sk_forward_alloc will stay synced. Add tcp_downgrade_zcopy_pure() helper so
CVE-2022-48830	In the Linux kernel, the following vulnerability has been resolved: can: isotp: fix potential CAN frame reception ra frame the current code logic does not consider concurrently receiving processes which do not show up in real worl syz problem is one of the scenarios. so->rx.len is changed by isotp_rcv_ff() during isotp_rcv_cf(), so->rx.len equal alloc_skb(). That will trigger skb_over_panic() in skb_put(). ========== Comm: ksoftirqd/1 Not tainted 5.16.0-rc8-syzkaller #0 RIP: 0010:skb_panic+0x16c/0x16e net/core/skbuff.c:113 C core/skbuff.c:118 [inline] skb_put.cold+0x24/0x24 net/core/skbuff.c:1990 isotp_rcv_cf net/can/isotp.c:570 [inline] deliver net/can/af_can.c:574 [inline] can_rcv_filter+0x445/0x8d0 net/can/af_can.c:635 can_receive+0x31d/0x580 net/can/af_can.c:696 __netif_receive_skb_one_core+0x114/0x180 net/core/dev.c:5465 __netif_receive_skb+0x24/ make sure the state changes and data structures stay consistent at CAN frame reception time by adding a spin_lock by syzkaller but does not affect real world operation.
CVE-2022-48831	In the Linux kernel, the following vulnerability has been resolved: ima: fix reference leak in asymmetric_verify() D algorithm is unknown.
CVE-2022-48832	In the Linux kernel, the following vulnerability has been resolved: audit: don't deref the syscall args when checking Jeff, dereferencing the openat2 syscall argument in audit_match_perm() to obtain the open_how::flags can result in using the open_how struct that we store in the audit_context with audit_openat2_how(). Independent of this patch, the audit mailing list roughly 40 minutes after this patch was posted.
CVE-2022-48836	In the Linux kernel, the following vulnerability has been resolved: Input: aiptek - properly check endpoint type Syz which is caused by wrong endpoint type. There was a check for the number of endpoints, but not for the type of en desc.bNumEndpoints check with usb_find_common_endpoints() helper for finding endpoints Fail log: usb 5-1: BC CPU: 2 PID: 48 at drivers/usb/core/urb.c:502 usb_submit_urb+0xed2/0x18a0 drivers/usb/core/urb.c:502 Modules l kworker/2:2 Not tainted 5.17.0-rc6-syzkaller-00226-g07ebd38a0da2 #0 Hardware name: QEMU Standard PC (Q35 Workqueue: usb_hub_wq hub_event ... Call Trace: <TASK> aiptek_open+0xd5/0x130 drivers/input/tablet/aiptek.c drivers/input/input.c:629 kbd_connect+0xfe/0x160 drivers/tty/vt/keyboard.c:1593

CVE-2022-48840	In the Linux kernel, the following vulnerability has been resolved: iavf: Fix hang during reboot/shutdown Recent c so the port is initialized in remove") adds a wait-loop at the beginning of iavf_remove() to ensure that port initializ net device. This causes a regression in reboot/shutdown scenario because in this case callback iavf_shutdown() is c device, makes it down if it is running and sets its state to __IAVF_REMOVE. Later shutdown callback of associate called. That callback calls among other things sriov_disable() that calls indirectly iavf_remove() (see stack trace be __IAVF_REMOVE then the mentioned loop is end-less and shutdown process hangs. The patch fixes this by chec of iavf_remove() and skips the rest of the function if the adapter is already in remove state (shutdown is in progress driven by ice or i40e driver 2. Ensure that the VF is bound to iavf driver 3. Reboot [52625.981294] sysrq: SysRq : task:reboot state:D stack: 0 pid:17359 ppid: 1 f2 [52625.996732] Call Trace: [52625.999187] __schedule+0x2d1/0 [52626.010545] schedule_hrtimeout_range_clock+0x83/0x100 [52626.020046] usleep_range+0x5b/0x80 [52626.0 [iavf] [52626.027645] pci_device_remove+0x3b/0xc0 [52626.031572] device_release_driver_internal+0x103/0x1 +0x72/0xa0 [52626.040904] pci_stop_and_remove_bus_device+0xe/0x20 [52626.045870] pci_iov_remove_virtfn +0x2f/0xe0 [52626.053813] ice_free_vfs+0x7c/0x340 [ice] [52626.057946] ice_remove+0x220/0x240 [ice] [5262 [52626.065987] pci_device_shutdown+0x34/0x60 [52626.070086] device_shutdown+0x165/0x1c5 [52626.07401 __do_sys_reboot+0x1d2/0x210 [52626.093815] do_syscall_64+0x5b/0x1a0 [52626.097483] entry_SYSCALL_64
CVE-2022-48866	In the Linux kernel, the following vulnerability has been resolved: HID: hid-thrustmaster: fix OOB read in thrustm out-of-bounds Read in thrustmaster_probe() bug. The root case is in missing validation check of actual number of e usb_host_interface::endpoint array, since it may contain less endpoints than code expects. Fix it by adding missing of endpoints do not match expected number
CVE-2022-48879	In the Linux kernel, the following vulnerability has been resolved: efi: fix NULL-deref in init error path In cases w have been disabled, the runtime services workqueue will never have been allocated. Do not try to destroy the work that EFI initialisation fails to avoid dereferencing a NULL pointer.
CVE-2022-48883	In the Linux kernel, the following vulnerability has been resolved: net/mlx5e: IPoIB, Block PKEY interfaces with to configure an arbitrary number of rx queues when creating an interface via netlink. This doesn't work for child PK uses the parent receive channels. Although the child shares the parent's receive channels, the number of rx queues i parent's rx channel index is used to access the child's channel_stats. So the array has to be at least as large as the pa correctly and to prevent out of bound accesses. This patch checks for the mentioned scenario and returns an error w propagated to the user.
CVE-2022-48884	In the Linux kernel, the following vulnerability has been resolved: net/mlx5: Fix command stats access after free C and can't accept FW commands till command interface is reinitialized. Such command failure is being logged to co access as command stats structure is being freed and reallocated during mlx5 devlink reload (see kernel log below) allocated on driver probe. Kernel log: [2394.808802] BUG: unable to handle kernel paging request at 0000000000 [2394.811811] Oops: 0002 [#1] SMP NOPTI ... [2394.815482] RIP: 0010:native_queued_spin_lock_slowpath+0x [2394.830667] _raw_spin_lock_irq+0x23/0x26 [2394.831858] cmd_status_err+0x55/0x110 [mlx5_core] [2394.8 [mlx5_core] [2394.834175] mlx5_query_port_ptys+0x78/0xa0 [mlx5_core] [2394.835337] mlx5e_ethtool_get_li [2394.836454] ? kmem_cache_alloc_trace+0x140/0x1c0 [2394.837562] __rh_call_get_link_ksettings+0x33/0x10 [2394.839755] __ethtool_get_link_ksettings+0x72/0x150 [2394.840862] duplex_show+0x6e/0xc0 [2394.841963 sysfs_kf_seq_show+0x9b/0x100 [2394.844123] seq_read+0x153/0x410 [2394.845187] vfs_read+0x91/0x140 [2 [2394.847234] do_syscall_64+0x5b/0x1a0 [2394.848228] entry_SYSCALL_64_after_hwframe+0x65/0xca
CVE-2022-48885	In the Linux kernel, the following vulnerability has been resolved: ice: Fix potential memory leak in ice_gnss_tty_ directly if the write_buf alloc failed, leaking the cmd_buf. Fix by free cmd_buf if write_buf alloc failed.
CVE-2022-48886	In the Linux kernel, the following vulnerability has been resolved: ice: Add check for kzalloc Add the check for the NULL pointer dereference. Moreover, use the goto-label to share the clean code.
CVE-2022-48921	In the Linux kernel, the following vulnerability has been resolved: sched/fair: Fix fault in reweight_entity Syzbot f been bisected to commit 4ef0c5c6b5ba ("kernel/sched: Fix sched_fork() access an invalid sched_task_group") The setpriority(PRIO_PGRP) within a thread group that causes a null-ptr-deref¬†in reweight_entity() in CFS. The scen of new threads, which then call setpriority(PRIO_PGRP, 0, -20), wait, and exit. For each of the new threads the co new task_struct and calls sched_post_fork() for it. In the above scenario there is a possibility that setpriority(PRIO for a thread in the group that is just being created by copy_process(), and for which the sched_post_fork() has not b pointer dereference in reweight_entity(),¬†as it will try to access the run queue pointer, which hasn't been set. Befo for the task has been set in sched_fork(), which is called much earlier in copy_process(), before the new task is add sched_post_fork(), which is called after that. To fix the issue the remove the update_load param from the update_lo only if the task flag doesn't have the TASK_NEW flag set.
CVE-2022-48934	In the Linux kernel, the following vulnerability has been resolved: nfp: flower: Fix a potential leak in nfp_tunnel_a an id between min (0) and max (NFP_MAX_MAC_INDEX) inclusive. So NFP_MAX_MAC_INDEX (0xff) is a v work correctly, the 'invalid' value for 'ida_idx' should not be in the 0..NFP_MAX_MAC_INDEX range, inclusive.
CVE-2022-48944	In the Linux kernel, the following vulnerability has been resolved: sched: Fix yet more sched_fork() races Where c sched_fork() access an invalid sched_task_group") fixed a fork race vs cgroup, it opened up a race vs syscalls by n it gets exposed through the pidhash. Commit 13765de8148f ("sched/fair: Fix fault in reweight_entity") is trying to whole class of issues, effectively reverting this commit.
CVE-2022-48963	In the Linux kernel, the following vulnerability has been resolved: net: wwan: iosm: fix memory leak in ipc_mux_i >ul_adb.pp_qlt in ipc_mux_init(), ipc_mux is not released.

CVE-2022-48968	In the Linux kernel, the following vulnerability has been resolved: octeontx2-pf: Fix potential memory leak in otx2 rhashtable_init() failed, it does not free tc->tc_entries_bitmap which is allocated in otx2_tc_alloc_ent_bitmap().
CVE-2022-48974	In the Linux kernel, the following vulnerability has been resolved: netfilter: conntrack: fix using __this_cpu_add in nf_conntrack_hash_check_insert(), when it fails in nf_ct_ext_valid_pre/post(), NF_CT_STAT_INC() will be called can be triggered: BUG: using __this_cpu_add() in preemptible [00000000] code: conntrack/1636 caller is nf_connt [nf_conntrack] Call Trace: <TASK> dump_stack_lvl+0x33/0x46 check_preemption_disabled+0xc3/0xf0 nf_connt [nf_conntrack] ctnetlink_create_conntrack+0x3cd/0x4e0 [nf_conntrack_netlink] ctnetlink_new_conntrack+0x1c0/ nfnetlink_rcv_msg+0x277/0x2f0 [nfnetlink] netlink_rcv_skb+0x50/0x100 nfnetlink_rcv+0x65/0x144 [nfnetlink] n +0x257/0x4f0 sock_sendmsg+0x5f/0x70 This patch is to fix it by changing to use NF_CT_STAT_INC_ATOMIC(check in nf_conntrack_hash_check_insert(), as well as nf_ct_ext_valid_post() in __nf_conntrack_confirm(). Note t __nf_conntrack_confirm() is safe to use NF_CT_STAT_INC(), as it's under local_bh_disable().
CVE-2022-48981	In the Linux kernel, the following vulnerability has been resolved: drm/shmem-helper: Remove errant put in error p this reference, resulting in the GEM object getting prematurely freed leading to a later use-after-free.
CVE-2022-48983	In the Linux kernel, the following vulnerability has been resolved: io_uring: Fix a null-ptr-deref in io_tctx_exit_cb(follows: BUG: KASAN: null-ptr-deref in io_tctx_exit_cb+0x53/0xd3 Read of size 4 at addr 0000000000000138 by file1 Not tainted 6.1.0-rc7-00103-gef4d3ea40565 #75 Hardware name: QEMU Standard PC (i440FX + PIIX, 1996 <TASK> dump_stack_lvl+0xcd/0x134 ? io_tctx_exit_cb+0x53/0xd3 kasan_report+0xbb/0x1f0 ? io_tctx_exit_cb+ io_tctx_exit_cb+0x53/0xd3 task_work_run+0x164/0x250 ? task_work_cancel+0x30/0x30 get_signal+0x1c3/0x244 lock_downgrade+0x6e0/0x6e0 ? exit_signals+0x8b0/0x8b0 ? do_raw_read_unlock+0x3b/0x70 ? do_raw_spin_unl +0x82/0x2470 ? kmem_cache_free+0x260/0x4b0 ? putname+0xfe/0x140 ? get_sigframe_size+0x10/0x10 ? do_ex lockdep_hardirqs_on+0x79/0x100 ? putname+0xfe/0x140 ? do_execveat_common.isra.0+0x238/0x710 exit_to_us syscall_exit_to_user_mode+0x19/0x50 do_syscall_64+0x42/0xb0 entry_SYSCALL_64_after_hwframe+0x63/0xc opcode bytes at 0xffffffffffffffd6. RSP: 002b:00000000fffb7790 EFLAGS: 00000200 ORIG_RAX: 0000000000000 0000000000000000 RCX: 0000000000000000 RDX: 0000000000000000 RSI: 0000000000000000 RDI: 0000000 0000000000000000 R09: 0000000000000000 R10: 0000000000000000 R11: 0000000000000000 R12: 0000000000 0000000000000000 R15: 0000000000000000 </TASK> Kernel panic - not syncing: panic_on_warn set ... This hap io_ring_exit_work() isn't synchronized with canceling all work items from eg exec. The execution of the two are or itself, but if io_tctx_exit_cb() is queued while we're canceling all work items off exec AND gets executed when the main loop in io_uring_cancel_generic(), then we can find current->io_uring == NULL and hit the above crash. It's the execution of the two paths are done by the task itself. [axboe: add code comment and also put an explanation in
CVE-2022-48984	In the Linux kernel, the following vulnerability has been resolved: can: slcan: fix freed work crash The LTP test pty kernel NULL pointer dereference, address: 0000000000000008 #PF: supervisor read access in kernel mode #PF: er PGD 0 P4D 0 Oops: 0000 [#1] PREEMPT SMP NOPTI CPU: 0 PID: 348 Comm: kworker/0:3 Not tainted 6.0.8-1- 9d20364b934f5aab0a9bdf84e8f45cfdfae39dab Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 04/01/2014 Workqueue: 0x0 (events) RIP: 0010:process_one_work (/home/rich/kernel/linux/kernel/workqueue.c:7 workqueue.c:2185) Code: 49 89 ff 41 56 41 55 41 54 55 53 48 89 f3 48 83 ec 10 48 8b 06 48 8b 6f 48 49 89 c4 45 e0 <49> 8b 44 24 08 44 8b a8 00 01 00 00 41 83 e5 20 f6 45 10 04 75 0e RSP: 0018:ffffaf7b40f47e98 EFLAGS: 0 RBX: ffff9d644e1b8b48 RCX: ffff9d649e439968 RDX: 00000000ffff8455 RSI: ffff9d644e1b8b48 RDI: ffff9d647 0000000000000c00 R09: ffff9d64764aa734 R10: 0000000000000007 R11: 0000000000000001 R12: 0000000000 ffff9d64490da780 R15: ffff9d64764aa6c0 FS: 0000000000000000(0000) GS:ffff9d649e400000(0000) knlGS:0000 0000 CR0: 0000000080050033 CR2: 0000000000000008 CR3: 0000000036424000 CR4: 00000000000006f0 Cal rich/kernel/linux/kernel/workqueue.c:2436) kthread (/home/rich/kernel/linux/kernel/kthread.c:376) ret_from_fork (entry_64.S:312) Apparently, the slcan's tx_work is freed while being scheduled. While slcan_netdev_close() (netde slcan_close() (tty side) does not. So when the netdev is never set UP, but the tty is stuffed with bytes and forced to never flushed. So add an additional flush_work() to slcan_close() to be sure the work is flushed under all circumsta flush_work() from slcan_close() to slcan_netdev_close(). What was the rationale behind it? Maybe we can drop the pattern in can327. So it perhaps needs the very same fix.
CVE-2022-48989	In the Linux kernel, the following vulnerability has been resolved: fscache: Fix oops due to race with cookie_lru ar the LRU and the LRU_DISCARD flag is set, but the state machine has not run yet, it's possible another thread can use it. When the cookie_worker finally runs, it will see the LRU_DISCARD flag set, transition the cookie->state to withdraw the cookie. Once the cookie is withdrawn the object is removed the below oops will occur because the ob Fix the oops by clearing the LRU_DISCARD bit if another thread uses the cookie before the cookie_worker runs. I address: 0000000000000008 ... CPU: 31 PID: 44773 Comm: kworker/u130:1 Tainted: G E 6.0.0-5.dneg.x86_64 #1 Google Compute Engine, BIOS Google 08/26/2022 Workqueue: events_unbound netfs_rreq_write_to_cache_work +0x28/0x90 [cachefiles] ... Call Trace: netfs_rreq_write_to_cache_work+0x11c/0x320 [netfs] process_one_work+ kthread+0xd6/0x100
CVE-2022-48999	In the Linux kernel, the following vulnerability has been resolved: ipv4: Handle attempt to delete multipath route w Gwangun Jung reported a slab-out-of-bounds access in fib_nh_match: fib_nh_match+0xf98/0x1130 linux-6.0-rc7/ +0x5f3/0xa40 linux-6.0-rc7/net/ipv4/fib_trie.c:1753 inet_rtm_delroute+0x2b3/0x380 linux-6.0-rc7/net/ipv4/fib_fro mutually exclusive with the legacy multipath spec. Fix fib_nh_match to return if the config for the to be deleted rou fib_info is using a nexthop object.
CVE-2022-49012	In the Linux kernel, the following vulnerability has been resolved: afs: Fix server->active leak in afs_put_server Th atomic_inc_return, which prevents the server from getting cleaned up and causes rmmod to hang with a warning: C

CVE-2022-49015	In the Linux kernel, the following vulnerability has been resolved: net: hsr: Fix potential use-after-free The skb is d calling this, dereferencing skb may trigger use-after-free.
CVE-2022-49016	In the Linux kernel, the following vulnerability has been resolved: net: mdiobus: fix unbalanced node reference cou device(mscc-miim) load test with CONFIG_OF_UNITTEST and CONFIG_OF_DYNAMIC enabled: OF: ERROR of 2, of_node_get()/of_node_put() unbalanced - destroy cset entry: attach overlay node /spi/soc@0/mdio@710700 an acpi node, the refcount is get in fwnode_mdiobus_phy_device_register(), but it has never been put when the dev fwnode_handle_put() in phy_device_release() to avoid leak. If it's an acpi node, it has never been get, but it's put in before phy_device_register() to keep get/put operation balanced.
CVE-2022-49024	In the Linux kernel, the following vulnerability has been resolved: can: m_can: pci: add missing m_can_class_free_ In m_can_pci_remove() and error handling path of m_can_pci_probe(), m_can_class_free_dev() should be called t m_can_class_allocate_dev(), otherwise there will be memleak.
CVE-2023-0051	Heap-based Buffer Overflow in GitHub repository vim/vim prior to 9.0.1144.
CVE-2023-0215	The public API function BIO_new_NDEF is a helper function used for streaming ASN.1 data via a BIO. It is prima support the SMIME, CMS and PKCS7 streaming capabilities, but may also be called directly by end user applicatio from the caller, prepends a new BIO_f_asn1 filter BIO onto the front of it to form a BIO chain, and then returns the caller. Under certain conditions, for example if a CMS recipient public key is invalid, the new filter BIO is freed ar indicating a failure. However, in this case, the BIO chain is not properly cleaned up and the BIO passed by the call previously freed filter BIO. If the caller then goes on to call BIO_pop() on the BIO then a use-after-free will occur. This scenario occurs directly in the internal function B64_write_ASN1() which may cause BIO_new_NDEF() to be BIO_pop() on the BIO. This internal function is in turn called by the public API functions PEM_write_bio_ASN1_ PEM_write_bio_PKCS7_stream, SMIME_write_ASN1, SMIME_write_CMS and SMIME_write_PKCS7. Other p by this include i2d_ASN1_bio_stream, BIO_new_CMS, BIO_new_PKCS7, i2d_CMS_bio_stream and i2d_PKCS7 command line applications are similarly affected.
CVE-2023-0286	There is a type confusion vulnerability relating to X.400 address processing inside an X.509 GeneralName. X.400 a but the public structure definition for GENERAL_NAME incorrectly specified the type of the x400Address field as interpreted by the OpenSSL function GENERAL_NAME_cmp as an ASN1_TYPE rather than an ASN1_STRING application sets the X509_V_FLAG_CRL_CHECK flag), this vulnerability may allow an attacker to pass arbitrary read memory contents or enact a denial of service. In most cases, the attack requires the attacker to provide both the need to have a valid signature. If the attacker only controls one of these inputs, the other input must already contain point, which is uncommon. As such, this vulnerability is most likely to only affect applications which have implem CRLs over a network.
CVE-2023-1118	A flaw use after free in the Linux kernel integrated infrared receiver/transceiver driver was found in the way user d flaw to crash the system or potentially escalate their privileges on the system.
CVE-2023-1355	NULL Pointer Dereference in GitHub repository vim/vim prior to 9.0.1402.
CVE-2023-1513	A flaw was found in KVM. When calling the KVM_GET_DEBUGREGS ioctl, on 32-bit systems, there might be s kvm_debugregs structure that could be copied to userspace, causing an information leak.
CVE-2023-1579	Heap based buffer overflow in binutils-gdb/bfd/libbfd.c in bfd_getl64.
CVE-2023-1838	A use-after-free flaw was found in vhost_net_set_backend in drivers/vhost/net.c in virtio network subcomponent in flaw could allow a local attacker to crash the system, and could even lead to a kernel information leak problem.
CVE-2023-2019	A flaw was found in the Linux kernel's netdevsim device driver, within the scheduling of events. This issue results reference count. This may allow an attacker to create a denial of service condition on the system.
CVE-2023-2156	A flaw was found in the networking subsystem of the Linux kernel within the handling of the RPL protocol. This is of user-supplied data, which can lead to an assertion failure. This may allow an unauthenticated remote attacker to system.
CVE-2023-2162	A use-after-free vulnerability was found in iscsi_sw_tcp_session_create in drivers/scsi/iscsi_tcp.c in SCSI sub-com attacker could leak kernel internal information.
CVE-2023-22742	libgit2 is a cross-platform, linkable library implementation of Git. When using an SSH remote with the optional lib certificate checking by default. Prior versions of libgit2 require the caller to set the `certificate_check` field of libgi a certificate check callback is not set, libgit2 does not perform any certificate checking. This means that by default callback, clients will not perform validation on the server SSH keys and may be subject to a man-in-the-middle atta v1.4.5 or v1.5.1. Users unable to upgrade should ensure that all relevant certificates are manually checked.
CVE-2023-22995	In the Linux kernel before 5.17, an error path in dwc3_qcom_acpi_register_core in drivers/usb/dwc3/dwc3-qcom.c calls.
CVE-2023-23000	In the Linux kernel before 5.17, drivers/phy/tegra/xusb.c mishandles the tegra_xusb_find_port_node return value. C error pointer is used.
CVE-2023-23004	In the Linux kernel before 5.19, drivers/gpu/drm/arm/malidp_planes.c misinterprets the get_sg_table return value (whereas it is actually an error pointer).

CVE-2023-23914	A cleartext transmission of sensitive information vulnerability exists in curl <v7.88.0 that could cause HSTS functi requested serially. Using its HSTS support, curl can be instructed to use HTTPS instead of usingan insecure clear-t in the URL. ThisHSTS mechanism would however surprisingly be ignored by subsequent transferswhen done on th would not be properlycarried on.
CVE-2023-23915	A cleartext transmission of sensitive information vulnerability exists in curl <v7.88.0 that could cause HSTS functi URLs are requested in parallel. Using its HSTS support, curl can be instructed to use HTTPS instead of using an in HTTP is provided in the URL. This HSTS mechanism would however surprisingly fail when multiple transfers are overwritten by the most recentlycompleted transfer. A later HTTP-only transfer to the earlier host name would ther
CVE-2023-24023	Bluetooth BR/EDR devices with Secure Simple Pairing and Secure Connections pairing in Bluetooth Core Specific the-middle attacks that force a short key length, and might lead to discovery of the encryption key and live injectio
CVE-2023-2426	Use of Out-of-range Pointer Offset in GitHub repository vim/vim prior to 9.0.1499.
CVE-2023-2431	A security issue was discovered in Kubelet that allows pods to bypass the seccomp profile enforcement. Pods that u specify an empty profile field, are affected by this issue. In this scenario, this vulnerability allows the pod to run in bug affects Kubelet.
CVE-2023-25696	Improper Input Validation vulnerability in the Apache Airflow Hive Provider. This issue affects Apache Airflow H
CVE-2023-28320	A denial of service vulnerability exists in curl <v8.1.0 in the way libcurl provides several different backends for res is built to use the synchronous resolver, it allows name resolves to time-out slow operations using `alarm()` and `sig global buffer that was not mutex protected and a multi-threaded application might therefore crash or otherwise mis
CVE-2023-28327	A NULL pointer dereference flaw was found in the UNIX protocol in net/unix/diag.c In unix_diag_get_exact in the not have sk, leading to a NULL pointer. This flaw allows a local user to crash or potentially cause a denial of servic
CVE-2023-2908	A null pointer dereference issue was found in Libtiff's tif_dir.c file. This issue may allow an attacker to pass a craft triggers a runtime error that causes undefined behavior. This will result in an application crash, eventually leading t
CVE-2023-29491	ncurses before 6.4 20230408, when used by a setuid application, allows local users to trigger security-relevant mem terminfo database file that is found in $HOME/.terminfo or reached via the TERMINFO or TERM environment va
CVE-2023-2953	A vulnerability was found in openldap. This security flaw causes a null pointer dereference in ber_memalloc_x() fu
CVE-2023-30571	Libarchive through 3.6.2 can cause directories to have world-writable permissions. The umask() call inside archive the whole process for a very short period of time; a race condition with another thread can lead to a permanent uma to implicit directory creation with permissions 0777 (without the sticky bit), which means that any low-privileged l those directories.
CVE-2023-31437	An issue was discovered in systemd 253. An attacker can modify a sealed log file such that, in some views, not all displayed. NOTE: the vendor reportedly sent "a reply denying that any of the finding was a security vulnerability."
CVE-2023-31438	An issue was discovered in systemd 253. An attacker can truncate a sealed log file and then resume log sealing suc despite modifications. NOTE: the vendor reportedly sent "a reply denying that any of the finding was a security vu
CVE-2023-31439	An issue was discovered in systemd 253. An attacker can modify the contents of past events in a sealed log file and integrity shows no error, despite modifications. NOTE: the vendor reportedly sent "a reply denying that any of the
CVE-2023-31486	HTTP::Tiny before 0.083, a Perl core module since 5.13.9 and available standalone on CPAN, has an insecure defa in to verify certificates.
CVE-2023-31486	HTTP::Tiny before 0.083, a Perl core module since 5.13.9 and available standalone on CPAN, has an insecure defa in to verify certificates.
CVE-2023-32570	VideoLAN dav1d before 1.2.0 has a thread_task.c race condition that can lead to an application crash, related to da
CVE-2023-34152	A vulnerability was found in ImageMagick. This security flaw cause a remote code execution vulnerability in Oper
CVE-2023-36054	lib/kadm5/kadm_rpc_xdr.c in MIT Kerberos 5 (aka krb5) before 1.20.2 and 1.21.x before 1.21.1 frees an uninitializ trigger a kadmind crash. This occurs because _xdr_kadm5_principal_ent_rec does not validate the relationship betw count.
CVE-2023-36479	Eclipse Jetty Canonical Repository is the canonical repository for the Jetty project. Users of the CgiServlet with a v wrong command executed. If a user sends a request to a org.eclipse.jetty.servlets.CGI Servlet for a binary with a sp command by wrapping it in quotation marks. This wrapped command, plus an optional command prefix, will then the original binary name provided by the user contains a quotation mark followed by a space, the resulting comman one. This issue was patched in version 9.4.52, 10.0.16, 11.0.16 and 12.0.0-beta2.
CVE-2023-38325	The cryptography package before 41.0.2 for Python mishandles SSH certificates that have critical options.

CVE-2023-38552	When the Node.js policy feature checks the integrity of a resource against a trusted manifest, the application can in checksum to the node's policy implementation, thus effectively disabling the integrity check. Impacts: This vulnera policy mechanism in all active release lines: 18.x and, 20.x. Please note that at the time this CVE was issued, the po Node.js.
CVE-2023-38560	An integer overflow flaw was found in pcl/pl/plfont.c:418 in pl_glyph_name in ghostscript. This issue may allow a transforming a crafted PCL file to PDF format.
CVE-2023-3896	Divide By Zero in vim/vim from¬†9.0.1367-1 to¬†9.0.1367-3
CVE-2023-39321	Processing an incomplete post-handshake message for a QUIC connection can cause a panic.
CVE-2023-39322	QUIC connections do not set an upper bound on the amount of data buffered when reading post-handshake messag cause unbounded memory growth. With fix, connections now consistently reject messages larger than 65KiB in siz
CVE-2023-39327	A flaw was found in OpenJPEG. Maliciously constructed pictures can cause the program to enter a large loop and c terminal.
CVE-2023-39615	Xmlsoft Libxml2 v2.11.0 was discovered to contain an out-of-bounds read via the xmlSAX2StartElement() functio allows attackers to cause a Denial of Service (DoS) via supplying a crafted XML file. NOTE: the vendor's position legacy SAX1 interface with custom callbacks; there is a crash even without crafted input.
CVE-2023-4244	A use-after-free vulnerability in the Linux kernel's netfilter: nf_tables component can be exploited to achieve local between nf_tables netlink control plane transaction and nft_set element garbage collection, it is possible to underflo free vulnerability. We recommend upgrading past commit 3e91b0ebd994635df2346353322ac51ce84ce6d8.
CVE-2023-43804	urllib3 is a user-friendly HTTP client library for Python. urllib3 doesn't treat the `Cookie` HTTP header special or over HTTP, that is the responsibility of the user. However, it is possible for a user to specify a `Cookie` header and redirects to a different origin if that user doesn't disable redirects explicitly. This issue has been patched in urllib3 v
CVE-2023-4408	The DNS message parsing code in `named` includes a section whose computational complexity is overly high. It d traffic, but crafted queries and responses may cause excessive CPU load on the affected `named` instance by explo authoritative servers and recursive resolvers. This issue affects BIND 9 versions 9.0.0 through 9.16.45, 9.18.0 throu through 9.11.37-S1, 9.16.8-S1 through 9.16.45-S1, and 9.18.11-S1 through 9.18.21-S1.
CVE-2023-4459	A NULL pointer dereference flaw was found in vmxnet3_rq_cleanup in drivers/net/vmxnet3/vmxnet3_drv.c in the Linux Kernel. This issue may allow a local attacker with normal user privilege to cause a denial of service due to a
CVE-2023-45139	fontTools is a library for manipulating fonts, written in Python. The subsetting module has a XML External Entity an attacker to resolve arbitrary entities when a candidate font (OT-SVG fonts), which contains a SVG table, is pars files from the filesystem fontTools is running on or make web requests from the host system. This vulnerability has
CVE-2023-45143	Undici is an HTTP/1.1 client written from scratch for Node.js. Prior to version 5.26.2, Undici already cleared Auth redirects, but did not clear `Cookie` headers. By design, `cookie` headers are forbidden request headers, disallowin browser environments. Since undici handles headers more liberally than the spec, there was a disconnect from the a implementation of fetch. As such this may lead to accidental leakage of cookie to a third-party site or a malicious a (ie. an open redirector) to leak the cookie to the third party site. This was patched in version 5.26.2. There are no kn
CVE-2023-45853	MiniZip in zlib through 1.3 has an integer overflow and resultant heap-based buffer overflow in zipOpenNewFileIn extra field. NOTE: MiniZip is not a supported part of the zlib product. NOTE: pyminizip through 0.2.6 is also vuln version, and exposes the applicable MiniZip code through its compress API.
CVE-2023-45862	An issue was discovered in drivers/usb/storage/ene_ub6250.c for the ENE UB6250 reader driver in the Linux kern extend beyond the end of an allocation.
CVE-2023-46137	Twisted is an event-based framework for internet applications. Prior to version 23.10.0rc1, when sending multiple twisted.web will process the requests asynchronously without guaranteeing the response order. If one of the endpoi can delay the response on purpose to manipulate the response of the second request when a victim launched two re 23.10.0rc1 contains a patch for this issue.
CVE-2023-46219	When saving HSTS data to an excessively long file name, curl could end up removing all contents, making subsequ HSTS status they should otherwise use.
CVE-2023-46361	Artifex Software jbig2dec v0.20 was discovered to contain a SEGV vulnerability via jbig2_error at /jbig2dec/jbig2
CVE-2023-46838	Transmit requests in Xen's virtual network protocol can consist of multiple parts. While not really useful, except fo zero length, i.e. carry no data at all. Besides a certain initial portion of the to be transferred data, these parts are dire fragments. Such converted request parts can, when for a particular SKB they are all of length zero, lead to a de-refe
CVE-2023-49083	cryptography is a package designed to expose cryptographic primitives and recipes to Python developers. Calling ` `load_der_pkcs7_certificates` could lead to a NULL-pointer dereference and segfault. Exploitation of this vulnerab (DoS) for any application attempting to deserialize a PKCS7 blob/certificate. The consequences extend to potential stability. This vulnerability has been patched in version 41.0.6.
CVE-2023-49463	libheif v1.17.5 was discovered to contain a segmentation violation via the function find_exif_tag at /libheif/exif.cc.

CVE-2023-50387	Certain DNSSEC aspects of the DNS protocol (in RFC 4033, 4034, 4035, 6840, and related RFCs) allow remote at consumption) via one or more DNSSEC responses, aka the "KeyTrap" issue. One of the concerns is that, when ther records, the protocol specification implies that an algorithm must evaluate all combinations of DNSKEY and RRSI
CVE-2023-50387	Certain DNSSEC aspects of the DNS protocol (in RFC 4033, 4034, 4035, 6840, and related RFCs) allow remote at consumption) via one or more DNSSEC responses, aka the "KeyTrap" issue. One of the concerns is that, when ther records, the protocol specification implies that an algorithm must evaluate all combinations of DNSKEY and RRSI
CVE-2023-50868	The Closest Encloser Proof aspect of the DNS protocol (in RFC 5155 when RFC 9276 guidance is skipped) allows (CPU consumption for SHA-1 computations) via DNSSEC responses in a random subdomain attack, aka the "NSE implies that an algorithm must perform thousands of iterations of a hash function in certain situations.
CVE-2023-5090	A flaw was found in KVM. An improper check in svm_set_x2apic_msr_interception() may allow direct access to h apic, potentially leading to a denial of service condition.
CVE-2023-51779	bt_sock_recvmsg in net/bluetooth/af_bluetooth.c in the Linux kernel through 6.6.8 has a use-after-free because of a
CVE-2023-51780	An issue was discovered in the Linux kernel before 6.6.8. do_vcc_ioctl in net/atm/ioctl.c has a use-after-free becau
CVE-2023-51781	An issue was discovered in the Linux kernel before 6.6.8. atalk_ioctl in net/appletalk/ddp.c has a use-after-free bec
CVE-2023-51782	An issue was discovered in the Linux kernel before 6.6.8. rose_ioctl in net/rose/af_rose.c has a use-after-free becau
CVE-2023-51792	Buffer Overflow vulnerability in libde265 v1.0.12 allows a local attacker to cause a denial of service via the allocat size of 0x10000000000.
CVE-2023-5197	A use-after-free vulnerability in the Linux kernel's netfilter: nf_tables component can be exploited to achieve local and removal of rules from chain bindings within the same transaction causes leads to use-after-free. We recommen f15f29fd4779be8a418b66e9d52979bb6d6c2325.
CVE-2023-52340	The IPv6 implementation in the Linux kernel before 6.3 has a net/ipv6/route.c max_size threshold that can be const service (network is unreachable errors) when IPv6 packets are sent in a loop via a raw socket.
CVE-2023-52429	dm_table_create in drivers/md/dm-table.c in the Linux kernel through 6.7.4 can attempt to (in alloc_targets) allocat because of a missing check for struct dm_ioctl.target_count.
CVE-2023-52434	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix potential OOBs in smb2_parse_ before dereferencing create contexts in smb2_parse_contexts(). This fixes following oops when accessing invalid c unable to handle page fault for address: ffff8881178d8cc3 #PF: supervisor read access in kernel mode #PF: error_c 4a01067 P4D 4a01067 PUD 0 Oops: 0000 [#1] PREEMPT SMP NOPTI CPU: 3 PID: 1736 Comm: mount.cifs No QEMU Standard PC (Q35 + ICH9, 2009), BIOS rel-1.16.2-3-gd478f380-rebuilt.opensuse.org 04/01/2014 RIP: 001 [cifs] Code: f8 10 75 13 48 b8 93 ad 25 50 9c b4 11 e7 49 39 06 0f 84 d2 00 00 00 8b 45 00 85 c0 74 61 41 29 c5 4 b7 7d 04 0f b7 45 06 4c 8d 74 3d 00 66 83 f8 04 75 bc ba 04 00 RSP: 0018:ffffc900007939e0 EFLAGS: 00010216 ffff8880180cc000 RCX: ffffc90000793c90 RDX: ffffc90000793cc0 RSI: ffff8880178d8cc0 RDI: ffff8880180cc00 ffffc90000793c22 R09: 0000000000000000 R10: ffff8880180cc000 R11: 0000000000000024 R12: 000000000000 0000000000000000 R15: ffffc90000793c22 FS: 00007f873753cbc0(0000) GS:ffff88806bc00000(0000) knlGS:000 0000 CR0: 0000000080050033 CR2: ffff8881178d8cc3 CR3: 00000000181ca000 CR4: 0000000000750ef0 PKRU +0x23/0x70 ? page_fault_oops+0x181/0x480 ? search_module_extables+0x19/0x60 ? srso_alias_return_thunk+0x asm_exc_page_fault+0x26/0x30 ? smb2_parse_contexts+0xa0/0x3a0 [cifs] SMB2_open+0x38d/0x5f0 [cifs] ? smb smb2_is_path_accessible+0x138/0x260 [cifs] cifs_is_path_remote+0x8d/0x230 [cifs] cifs_mount+0x7e/0x350 [cif [cifs] smb3_get_tree+0xd9/0x290 [cifs] vfs_get_tree+0x2c/0x100 ? capable+0x37/0x70 path_mount+0x2d7/0xb80 _raw_spin_unlock_irqrestore+0x44/0x60 __x64_sys_mount+0x11a/0x150 do_syscall_64+0x47/0xf0 entry_SYSC 0033:0x7f8737657b1e
CVE-2023-52436	In the Linux kernel, the following vulnerability has been resolved: f2fs: explicitly null-terminate the xattr list Wher xattr list. This eliminates the fragile assumption that the unused xattr space is always zeroed.
CVE-2023-52438	In the Linux kernel, the following vulnerability has been resolved: binder: fix use-after-free in shinker's callback Th the shrinker's callback, which means that using alloc->vma pointer isn't safe as it can race with munmap(). As of co zap pages with read mmap_sem in munmap") the mmap lock is downgraded after the vma has been isolated. I was manually adding some delays and triggering page reclaiming through the shrinker's debug sysfs. The following KA == BUG: KASAN: slab- +0x470/0x4b8 Read of size 8 at addr ffff356ed50e50f0 by task bash/478 CPU: 1 PID: 478 Comm: bash Not tainted #70 Hardware name: linux,dummy-virt (DT) Call trace: zap_page_range_single+0x470/0x4b8 binder_alloc_free_p +0x130/0x3b0 list_lru_walk_node+0xc4/0x22c binder_shrink_scan+0x108/0x1dc shrinker_debugfs_scan_write+0 vfs_write+0x1ac/0x758 ksys_write+0xf0/0x1dc __arm64_sys_write+0x6c/0x9c Allocated by task 492: kmem_cac +0x2c/0x190 mmap_region+0x258/0x18bc do_mmap+0x694/0xa60 vm_mmap_pgoff+0x170/0x29c ksys_mmap_ +0xcc/0x144 Freed by task 491: kmem_cache_free+0x17c/0x3c8 vm_area_free_rcu_cb+0x74/0x98 rcu_core+0xa3 __do_softirq+0x2fc/0xd24 Last potentially related work creation: __call_rcu_common.constprop.0+0x6c/0xba0 ca remove_vma+0xe4/0x118 do_vmi_align_munmap.isra.0+0x718/0xb5c do_vmi_munmap+0xdc/0x1fc __vm_munm +0x58/0x7c Fix this issue by performing instead a vma_lookup() which will fail to find the vma that was isolated b that this option has better performance than upgrading to a mmap write lock which would increase contention. Plus removed anyway.

CVE-2023-52439	In the Linux kernel, the following vulnerability has been resolved: uio: Fix use-after-free in uio_open core-1 core-2 ---------- uio_unregister_device uio_open idev = idr_find() device_unregister(>dev) uio_device_release get_device(&idev->dev) kfree(idev) uio_free_minor(minor) uio_release put_device(&id ---------- In the core-1 uio_unregister_device(), the device_unregister will kfr But after core-1 device_unregister, put_device and before doing kfree, the core-2 may get_device. Then: 1. After c free for idev. 2. When core-2 do uio_release and put_device, the idev will be double freed. To address this issue, w with minor_lock.
CVE-2023-52443	In the Linux kernel, the following vulnerability has been resolved: apparmor: avoid crash when parsed profile nam profile in unpack_profile() described like "profile :ns::samba-dcerpcd /usr/lib*/samba/{,samba/}samba-dcerpcd {... unpacked as a fully-qualified name and then passed to aa_splitn_fqname(). aa_splitn_fqname() treats ":samba-dcer Thus it returns NULL for tmpname, meanwhile tmpns is non-NULL. Later aa_alloc_profile() crashes as the new pr protection fault, probably for non-canonical address 0xdffffc0000000000: 0000 [#1] PREEMPT SMP KASAN NO [0x0000000000000000-0x0000000000000007] CPU: 6 PID: 1657 Comm: apparmor_parser Not tainted 6.7.0-rc2-d Standard PC (i440FX + PIIX, 1996), BIOS rel-1.16.2-3-gd478f380-rebuilt.opensuse.org 04/01/2014 RIP: 0010:strl strlen+0x1e/0xa0 aa_policy_init+0x1bb/0x230 aa_alloc_profile+0xb1/0x480 unpack_profile+0x3bc/0x4960 aa_un +0x213/0x33c0 policy_update+0x261/0x370 profile_replace+0x20e/0x2a0 vfs_write+0x2af/0xe00 ksys_write+0x1 entry_SYSCALL_64_after_hwframe+0x6e/0x76 </TASK> ---[end trace 0000000000000000]--- RIP: 0010:strlen aa_splitn_fqname() is expected and checked in other places where it is called (e.g. aa_remove_profiles). Well, there following profile is allowed" inside. AFAICS, nothing can prevent unpacked "name" to be in form like ":samba-dc whole profile set replacement in such case and inform user with EPROTO and an explaining message. Found by Li
CVE-2023-52444	In the Linux kernel, the following vulnerability has been resolved: f2fs: fix to avoid dirent corruption As Al reporte = new_dir && !whiteout) f2fs_set_link(old_inode, old_dir_entry, old_dir_page, new_dir); else f2fs_put_page(old_ in the ".." link. And cross-directory rename does move the source to new parent, even if you'd been asked to leave a lore.kernel.org/all/20231017055040.GN800259@ZenIV/ With below testcase, it may cause dirent corruption, due t ".." link to new directory. - mkdir -p dir/foo - renameat2 -w dir/foo bar [ASSERT] (__chk_dots_dentries:1421) --> ino is [0x3] [FSCK] other corrupted bugs [Fail]
CVE-2023-52445	In the Linux kernel, the following vulnerability has been resolved: media: pvrusb2: fix use after free on context dis created targeting the pvr2_context_thread_func function, which may call pvr2_context_destroy and thus call kfree(happen before the usb hub_event handler is able to notify the driver. This patch adds a sanity check before the inva disconnection call stack.
CVE-2023-52448	In the Linux kernel, the following vulnerability has been resolved: gfs2: Fix kernel NULL pointer dereference in gf NULL pointer dereference when accessing rgd->rd_rgl in gfs2_rgrp_dump(). This can happen when creating rgd-> NULL pointer check in gfs2_rgrp_dump() to prevent that.
CVE-2023-52449	In the Linux kernel, the following vulnerability has been resolved: mtd: Fix gluebi NULL pointer dereference caus and gluebi.ko are loaded, the notifier of ftl triggers NULL pointer dereference when trying to access ,Äògluebi->de ubi_gluebi_init ubi_register_volume_notifier ubi_enumerate_volumes ubi_notify_all gluebi_notify nb->notifier_ca mtd_device_parse_register add_mtd_device blktrans_notify_add not->add() ftl_add_mtd tr->add_mtd() scan_head gluebi_read mtd->read() gluebi->desc - NULL Detailed reproduction information available at the Link [1], In the n gluebi_get_device(), and access gluebi->desc in the gluebi_read(). However, gluebi_get_device() is not executed in leads to NULL pointer dereference. The solution for the gluebi module is to run jffs2 on the UBI volume without c Therefore, this problem can be avoided by preventing gluebi from creating the mtdblock device after creating mtd p
CVE-2023-52451	In the Linux kernel, the following vulnerability has been resolved: powerpc/pseries/memhp: Fix access beyond end dlpar_memory_remove_by_index() may access beyond the bounds of the drmem lmb array when the LMB lookup DRC index. When the search fails, the cursor is left pointing to &drmem_info->lmbs[drmem_info->n_lmbs], whic entry in the array. The debug message at the end of the function then dereferences this pointer: pr_debug("Failed to lmb->base_addr); This was found by inspection and confirmed with KASAN: pseries-hotplug-mem: Attempting to ========== BUG: KASAN: slab- +0x298/0x1658 Read of size 8 at addr c000000364e97fd0 by task bash/949 dump_stack_lvl+0xa4/0xfc (unreliable +0x140/0x2e0 __asan_load8+0xa8/0xe0 dlpar_memory+0x298/0x1658 handle_dlpar_errorlog+0x130/0x1d0 dlpar +0x68/0xa0 sysfs_kf_write+0xc4/0x110 kernfs_fop_write_iter+0x26c/0x390 vfs_write+0x2d4/0x4e0 ksys_write+0 +0x268/0x530 system_call_vectored_common+0x15c/0x2ec Allocated by task 1: kasan_save_stack+0x48/0x80 ka kasan_save_alloc_info+0x34/0x50 __kasan_kmalloc+0xd0/0x120 __kmalloc+0x8c/0x320 kmalloc_array.constpro do_one_initcall+0xe0/0x5c0 kernel_init_freeable+0x4ec/0x5a0 kernel_init+0x30/0x1e0 ret_from_kernel_user_thre the object at c000000364e80000 which belongs to the cache kmalloc-128k of size 131072 The buggy address is loc byte region [c000000364e80000, c000000364e97fd0) ========== hotplug-mem: Failed to hot-remove memory at 0 Log failed lookups with a separate message and dereference the c
CVE-2023-52457	In the Linux kernel, the following vulnerability has been resolved: serial: 8250: omap: Don't skip resource freeing i Returning an error code from .remove() makes the driver core emit the little helpful error message: remove callbac ignored. and then remove the device anyhow. So all resources that were not freed are leaked in this case. Skipping to keep enough of the UART around to trigger a use-after-free. So replace the error return (and with it the little help message and continue to cleanup.

CVE-2023-52464	In the Linux kernel, the following vulnerability has been resolved: EDAC/thunderx: Fix possible out-of-bounds str globally exposes a warning for a common bug in the usage of strncat(): drivers/edac/thunderx_edac.c: In function ' drivers/edac/thunderx_edac.c:1136:17: error: 'strncat' specified bound 1024 equals destination size [-Werror=string OCX_MESSAGE_SIZE); | ^~~~~~~ ... 1145 | strncat(msg, other, OCX_M other, OCX_MESSAGE_SIZE); ... Apparently the author of this driver expected strncat() to behave the way that st destination buffer as its third argument rather than the length of the source buffer. The result is that there is no chec it to strlcat(). [bp: Trim compiler output, fixup commit message.]
CVE-2023-52469	In the Linux kernel, the following vulnerability has been resolved: drivers/amd/pm: fix a use-after-free in kv_parse equals to NULL, kv_parse_power_table frees adev->pm.dpm.ps that allocated before. However, after the control fl kv_parse_power_table |-> kv_dpm_init |-> kv_dpm_sw_init |-> kv_dpm_fini The adev->pm.dpm.ps is used in the kv_parse_power_table and causes a use-after-free bug.
CVE-2023-52470	In the Linux kernel, the following vulnerability has been resolved: drm/radeon: check the alloc_workqueue return alloc_workqueue return value in radeon_crtc_init() to avoid null-ptr-deref.
CVE-2023-52472	In the Linux kernel, the following vulnerability has been resolved: crypto: rsa - add a check for allocation failure S allocation can fail so add a check to prevent a NULL dereference. Small allocations like this can't actually fail in cu simple and makes the static checkers happy.
CVE-2023-52478	In the Linux kernel, the following vulnerability has been resolved: HID: logitech-hidpp: Fix kernel crash on receiv has *four* time-of-check vs time-of-use (TOCTOU) races when it races with itself. hidpp_connect_event() primari runs on probe() and if a "device-connected" packet is received by the hw when the thread running hidpp_connect_e the hw, then a second thread running hidpp_connect_event() will be started from the workqueue. This opens the fo simplified): 1. Retrieving + printing the protocol (harmless race): if (!hidpp->protocol_major) { hidpp_root_get_pr response.rap.params[0]; } We can actually see this race hit in the dmesg in the abrt output attached to rhbz#222796 0003:046D:4071.0049: HID++ 4.5 device connected. [3064.658184] logitech-hidpp-device 0003:046D:4071.0049 with extra logging added has shown that after this the 2 threads take turn grabbing the hw access mutex (send_mut TOCTOU cases managing to hit all of them: 2. Updating the name to the HIDPP name (harmless race): if (hidpp-> new_name; } 3. Initializing the power_supply class for the battery (problematic!): hidpp_initialize_battery() { if (h Blocks, threads take turns executing this */ hidpp->battery.desc.properties = devm_kmemdup(dev, hidpp_battery_ = devm_power_supply_register(&hidpp->hid_dev->dev, &hidpp->battery.desc, cfg); } 4. Creating delayed input_ >delayed_input) return; hidpp->delayed_input = hidpp_allocate_input(hdev); The really big problem here is 3. Hitt hidpp->battery.desc.properties = devm_kmemdup(dev, hidpp_battery_props, cnt, GFP_KERNEL); hidpp->battery >hid_dev->dev, &hidpp->battery.desc, cfg); ... hidpp->battery.desc.properties = devm_kmemdup(dev, hidpp_batte >battery.ps = devm_power_supply_register(&hidpp->hid_dev->dev, &hidpp->battery.desc, cfg); So now we have battery, which looks a bit weird from userspace's pov but this is not even the really big problem. Notice how: 1. Th >battery.desc struct is shared between the 2 power supplies 3. hidpp->battery.desc.properties points to the result fr causes a use after free scenario on USB disconnect of the receiver: 1. The last registered power supply class device the last devm_kmemdup() call gets freed, hidpp->battery.desc.properties now points to freed memory 3. The first r unregistered, this involves sending a remove uevent to userspace which invokes power_supply_uevent() to fill the uses hidpp->battery.desc.properties which now points to freed memory leading to backtraces like this one: Sep 22 handle page fault for address: ffffb2140e017f08 ... Sep 22 20:01:35 eric kernel: Workqueue: usb_hub_wq hub_eve 0010:power_supply_uevent+0xee/0x1d0 ... Sep 22 20:01:35 eric kernel: ? asm_exc_page_fault+0x26/0x30 Sep 22 +0xee/0x1d0 Sep 22 20:01:35 eric kernel: ? power_supply_uevent+0x10d/0x1d0 Sep 22 20:01:35 eric kernel: dev_ kernel: kobject_uevent_env+0x291/0x680 Sep 22 20:01:35 eric kernel: ---truncated---
CVE-2023-52484	In the Linux kernel, the following vulnerability has been resolved: iommu/arm-smmu-v3: Fix soft lockup triggered When running an SVA case, the following soft lockup is triggered: ---------- - CPU#244 stuck for 26s! pstate: 83400009 (Nzcv daif +PAN -UAO +TCO +DIT -SSBS BTYPE=--) pc : arm_sm lr : arm_smmu_cmdq_issue_cmdlist+0x150/0xa50 sp : ffff8000d83ef290 x29: ffff8000d83ef290 x28: 000000003b x26: ffff8000d83ef3c0 x25: da86c0812194a0e8 x24: 0000000000000000 x23: 0000000000000040 x22: ffff8000d8 x20: 0000000000000001 x19: ffff0000c6398080 x18: 0000000000000000 x17: 0000000000000000 x16: 0000000 x14: ffff3000b4a30888 x13: ffff3000b4a3cf60 x12: 0000000000000000 x11: 0000000000000000 x10: 0000000 0000000000000000 x7 : 0000000000000000 x6 : 00000000048cfa x5 : 0000000000000000 x4 : 0000000000 0000000080000000 x1 : 0000000000000000 x0 : 0000000000000001 Call trace: arm_smmu_cmdq_issue_cmdlist +0x118/0x254 arm_smmu_tlb_inv_range_asid+0x6c/0x130 arm_smmu_mm_invalidate_range+0xa0/0xa4 __mmu +0x88/0x120 unmap_vmas+0x194/0x1e0 unmap_region+0xb4/0x144 do_mas_align_munmap+0x290/0x490 do_m +0xa8/0x19c __arm64_sys_munmap+0x28/0x50 invoke_syscall+0x78/0x11c el0_svc_common.constprop.0+0x58 +0x2c/0xd4 el0t_64_sync_handler+0x114/0x140 el0t_64_sync+0x1a4/0x1a8 ---------- rc1 the arm_smmu_mm_invalidate_range above is renamed to "arm_smmu_mm_arch_invalidate_secondary_tlbs", 06ff87bae8d3 ("arm64: mm: remove unused functions and variable protoypes") fixed a similar lockup on the CPU too, since arm_smmu_mm_arch_invalidate_secondary_tlbs() is called typically next to MMU tlb flush function, e.g { __flush_tlb_range { // check MAX_TLBI_OPS } } mmu_notifier_arch_invalidate_secondary_tlbs { arm_smmu_ not check MAX_TLBI_OPS } } } Clone a CMDQ_MAX_TLBI_OPS from the MAX_TLBI_OPS in tlbflush.h, si page table, so it makes sense to align with the tlbflush code. Then, replace per-page TLBI commands with a single hits this threshold.

CVE-2023-52486	In the Linux kernel, the following vulnerability has been resolved: drm: Don't unref the same fb many times by mis get a deadlock after the fb lookup in drm_mode_page_flip_ioctl() we proceed to unref the fb and then retry the whc to reset the fb pointer back to NULL, and so if we then get another error during the retry, before the fb lookup, we without having gotten another reference. The end result is that the fb will (eventually) end up being freed while it's we've unreffed it to avoid doing it again until we've done another fb lookup. This turned out to be pretty easy to hit (and CONFIG_DEBUG_WW_MUTEX_SLOWPATH=y). The first symptom I saw that drm_closefb() simply got framebuffer list. Fortunately I was able to convince it to oops instead, and from there it was easier to track down th
CVE-2023-52487	In the Linux kernel, the following vulnerability has been resolved: net/mlx5e: Fix peer flow lists handling The cite mlx5e_tc_del_fdb_peer_flow() to only clear DUP flag when list of peer flows has become empty. However, if any flow (for example, the neighbor update workqueue task is updating peer flow's parent encap entry concurrently), th peer list and, consecutively, DUP flag will remain set. Since mlx5e_tc_del_fdb_peers_flow() calls mlx5e_tc_del_f index the algorithm will try to remove the flow from eswitch instances that it has never peered with causing either remove the flow peer list head of peer_index that was never initialized or a warning if the list debug config is enabl peer flow from the list even when not releasing the last reference to it. [0]: [3102.985806] ------------[cut here]---- ffff888139110698->next is NULL [3102.986757] WARNING: CPU: 2 PID: 22109 at lib/list_debug.c:53 __list_d [3102.987561] Modules linked in: act_ct nf_flow_table bonding act_tunnel_key act_mirred act_skbedit vxlan cls_ cls_flower sch_ingress mlx5_vdpa vringh vhost_iotlb vdpa openvswitch nsh xt_MASQUERADE nf_conntrack_ne xt_conntrack nf_nat br_netfilter rpcsec_gss_krb5 auth_rpcg ss oid_registry overlay rpcrdma rdma_ucm ib_iser libi ib_ipoib iw_cm ib_cm mlx5_ib ib_uverbs ib_core mlx5_core [last unloaded: bonding] [3102.991113] CPU: 2 PID 6.6.0-rc6+ #3 [3102.991695] Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS rel-1.13.0-0-gf21b [3102.992605] RIP: 0010:__list_del_entry_valid_or_report+0x4f/0xc0 [3102.993122] Code: 39 c2 74 56 48 8b 3 73 b8 01 00 00 00 c3 48 89 fe 48 c7 c7 48 fd 0a 82 e8 41 0b ad ff <0f> 0b 31 c0 c3 48 89 fe 48 c7 c7 70 fd 0a 82 e 0018:ffff8881383e7710 EFLAGS: 00010286 [3102.995078] RAX: 0000000000000000 RBX: 000000000000002 RDX: 0000000000000001 RSI: ffff88885f89b640 RDI: ffff88885f89b640 [3102.997188] DEL flow 00000000be3 RBP: dead000000000122 R08: 0000000000000000 R09: c0000000ffffdfff [3102.999604] R10: 000000000000000 dead000000000100 [3103.000198] R13: 0000000000000002 R14: ffff888139110000 R15: ffff888101901240 [31 GS:ffff88885f880000(0000) knlGS:0000000000000000 [3103.001486] CS: 0010 DS: 0000 ES: 0000 CR0: 00000 00007fd42e8dcb70 CR3: 000000011e68a003 CR4: 0000000000370ea0 [3103.002596] DR0: 0000000000000000 0000000000000000 [3103.003190] DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 [<TASK> [3103.004297] ? __warn+0x7d/0x130 [3103.004623] ? __list_del_entry_valid_or_report+0x4f/0xc0 [3 [3103.005439] ? console_unlock+0x4a/0xd0 [3103.005806] ? handle_bug+0x3f/0x70 [3103.006149] ? exc_inval asm_exc_invalid_op+0x16/0x20 [3103.007430] ? __list_del_entry_valid_or_report+0x4f/0xc0 [3103.007910] ml [mlx5_core] [3103.008463] mlx5e_tc_del_flow+0x46/0x270 [mlx5_core] [3103.008944] mlx5e_flow_put+0x26/ mlx5e_delete_flower+0x25f/0x380 [mlx5_core] [3103.009901] tc_setup_cb_destroy+0xab/0x180 [3103.010292] [3103.010779] __fl_delete+0x2d4/0x2f0 [cls_flower] [3103.0 ---truncated---
CVE-2023-52504	In the Linux kernel, the following vulnerability has been resolved: x86/alternatives: Disable KASAN in apply_alte KASAN triggers during apply_alternatives() on a 5-level paging machine: BUG: KASAN: out-of-bounds in rcu_is ff110003ee6419a0 by task swapper/0/0 ... __asan_load4() rcu_is_watching() trace_hardirqs_on() text_poke_early() with 5-level paging, cpu_feature_enabled(X86_FEATURE_LA57) gets patched. It includes KASAN code, where I depends on __VIRTUAL_MASK_SHIFT, which is defined with cpu_feature_enabled(). KASAN gets confused wh KASAN_SHADOW_START users. A test patch that makes KASAN_SHADOW_START static, by replacing __V around the issue. Fix it for real by disabling KASAN while the kernel is patching alternatives. [mingo: updated the
CVE-2023-52518	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: hci_codec: Fix leaking content of loc leak can be observed when the controller supports codecs which are stored in local_codecs list but the elements are 0xffff88800221d840 (size 32): comm "kworker/u3:0", pid 36, jiffies 4294898739 (age 127.060s) hex dump (first 3 80 88 ff ff .....!..... 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ..... backtrace: [<ffffffffb324f557>] hci_codec_list_add.isra.0+0x2d/0x160 [<ffffffffb39ef643>] hci_read_codec_capabilities+0x183/0x270 [<ffffffffb3 +0x1bb/0x2d0 [<ffffffffb39f162e>] hci_read_local_codecs_sync+0x3e/0x60 [<ffffffffb39ff1b3>] hci_dev_open_s hci_power_on+0x10d/0x3f0 [<ffffffffb30c99b4>] process_one_work+0x404/0x800 [<ffffffffb30ca134>] worker_t kthread+0x188/0x1c0 [<ffffffffb304db6b>] ret_from_fork+0x2b/0x50 [<ffffffffb300206a>] ret_from_fork_asm+0
CVE-2023-52530	In the Linux kernel, the following vulnerability has been resolved: wifi: mac80211: fix potential key use-after-free ieee80211_gtk_rekey_add() but returns 0 due to KRACK protection (identical key reinstall), ieee80211_gtk_rekey key, in a potential use-after-free. This normally doesn't happen since it's only called by iwlwifi in case of WoWLA protection, but still better to fix, do that by returning an error code and converting that to success on the cfg80211 b callers of ieee80211_gtk_rekey_add().
CVE-2023-52582	In the Linux kernel, the following vulnerability has been resolved: netfs: Only call folio_start_fscache() one time fo using netfs implements a clamp_length() function, it can set subrequest lengths smaller than a page size. When we netfs_rreq_unlock_folios() to set any folios to be written back, we need to make sure we only call folio_start_fscac simple testcase: mount -o fsc,rsize=1024,wsize=1024 127.0.0.1:/export /mnt/nfs dd if=/dev/zero of=/mnt/nfs/file.bi records out 4096 bytes (4.1 kB, 4.0 KiB) copied, 0.0126359 s, 324 kB/s echo 3 > /proc/sys/vm/drop_caches cat /m oops similar to the following: page dumped because: VM_BUG_ON_FOLIO(folio_test_private_2(folio)) ---------- include/linux/netfs.h:44! ... CPU: 5 PID: 134 Comm: kworker/u16:5 Kdump: loaded Not tainted 6.4.0-rc5 ... RIP: 0 [netfs] ... Call Trace: netfs_rreq_assess+0x497/0x660 [netfs] netfs_subreq_terminated+0x32b/0x610 [netfs] nfs_ne [nfs] nfs_read_completion+0x2f9/0x330 [nfs] rpc_free_task+0x72/0xa0 [sunrpc] rpc_async_release+0x46/0x70 [su worker_thread+0x89/0x610 kthread+0x181/0x1c0 ret_from_fork+0x29/0x50

CVE-2023-52585	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: Fix possible NULL dereference in Return invalid error code -EINVAL for invalid block id. Fixes the below: drivers/gpu/drm/amd/amdgpu/amdgpu_r amdgpu_ras_query_error_status_helper() error: we previously assumed 'info' could be null (see line 1176)
CVE-2023-52594	In the Linux kernel, the following vulnerability has been resolved: wifi: ath9k: Fix potential array-index-out-of-bou array-index-out-of-bounds read in ath9k_htc_txstatus(). The bug occurs when txs->cnt, data from a URB provided of the array txs->txstatus, which is HTC_MAX_TX_STATUS. WARN_ON() already checks it, but there is no bug the function return if that is the case. Found by a modified version of syzkaller. UBSAN: array-index-out-of-bound range for type '__wmi_event_txstatus [12]' Call Trace: ath9k_htc_txstatus ath9k_wmi_event_tasklet tasklet_action sysvec_apic_timer_interrupt
CVE-2023-52595	In the Linux kernel, the following vulnerability has been resolved: wifi: rt2x00: restart beacon queue when hardwa all registers are reset, so all queues are forced to stop in hardware interface. However, mac80211 will not automatic the beacon queue, the queue will be deadlocked and unable to start again. This patch fixes the issue where Apple d ieee80211_restart_hw().
CVE-2023-52597	In the Linux kernel, the following vulnerability has been resolved: KVM: s390: fix setting of fpc register kvm_arch floating point control (fpc) register of a guest cpu. The new value is tested for validity by temporarily loading it int of the fpc register of the host process: if an interrupt happens while the value is temporarily loaded into the fpc regi point or vector registers are used, the current fp/vx registers are saved with save_fpu_regs() assuming they belong t registers when returning to user space. test_fp_ctl() restores the original user space / host process fpc register value to user space. In result the host process will incorrectly continue to run with the value that was supposed to be used the test. There is another test right before the SIE context is entered which will handles invalid values. This results now be accepted instead of that the ioctl fails with -EINVAL. This seems to be acceptable, given that this interface in addition the same behaviour implemented with the memory mapped interface (replace invalid values with zero)
CVE-2023-52598	In the Linux kernel, the following vulnerability has been resolved: s390/ptrace: handle setting of fpc register correc (fpc) register of a traced process is modified with the ptrace interface the new value is tested for validity by tempor may lead to corruption of the fpc register of the tracing process: if an interrupt happens while the value is temporar interrupt context floating point or vector registers are used, the current fp/vx registers are saved with save_fpu_regs will be loaded into fp/vx registers when returning to user space. test_fp_ctl() restores the original user space fpc reg when returning to user space. In result the tracer will incorrectly continue to run with the value that was supposed t saving fpu register contents with save_fpu_regs() before using test_fp_ctl().
CVE-2023-52599	In the Linux kernel, the following vulnerability has been resolved: jfs: fix array-index-out-of-bounds in diNewExt out-of-bounds in fs/jfs/jfs_imap.c:2360:2 index -878706688 is out of range for type 'struct iagctl[128]' CPU: 1 PID tainted 6.7.0-rc4-syzkaller-00009-gbee0e7762ad2 #0 Hardware name: Google Google Compute Engine/Google Co Call Trace: <TASK> __dump_stack lib/dump_stack.c:88 [inline] dump_stack_lvl+0x1e7/0x2d0 lib/dump_stack.c: [inline] __ubsan_handle_out_of_bounds+0x11c/0x150 lib/ubsan.c:348 diNewExt+0x3cf3/0x4000 fs/jfs/jfs_imap.c [inline] diAllocAG+0xbe8/0x1e50 fs/jfs/jfs_imap.c:1666 diAlloc+0x1d3/0x1760 fs/jfs/jfs_imap.c:1587 ialloc+0x8 +0x1c5/0xb90 fs/jfs/namei.c:225 vfs_mkdir+0x2f1/0x4b0 fs/namei.c:4106 do_mkdirat+0x264/0x3a0 fs/namei.c:41 [inline] __se_sys_mkdir fs/namei.c:4147 [inline] __x64_sys_mkdir+0x6e/0x80 fs/namei.c:4147 do_syscall_x64 ar do_syscall_64+0x45/0x110 arch/x86/entry/common.c:82 entry_SYSCALL_64_after_hwframe+0x63/0x6b RIP: 00 c0 c3 0f 1f 40 00 48 c7 c2 b8 ff ff ff f7 d8 64 89 02 b8 ff ff ff ff c3 66 0f 1f 44 00 00 b8 53 00 00 00 0f 05 <48> 3d f7 d8 64 89 01 48 RSP: 002b:00007ffd83023038 EFLAGS: 00000286 ORIG_RAX: 0000000000000053 RAX: fff 00007fcb7e6a0b57 RDX: 00000000000a1020 RSI: 00000000000001ff RDI: 0000000020000140 RBP: 000000002 0000000000000000 R10: 0000000000000000 R11: 0000000000000286 R12: 00007ffd830230d0 R13: 000000000 0000000000000000 [Analysis] When the agstart is too large, it can cause agno overflow. [Fix] After obtaining agn process. Modified the test from agno > MAXAG to agno >= MAXAG based on linux-next report by kernel test rob
CVE-2023-52600	In the Linux kernel, the following vulnerability has been resolved: jfs: fix uaf in jfs_evict_inode When the executio ipimap that has been released may be accessed in diFreeSpecial(). Asynchronous ipimap release occurs when rcu_ diMount(ipimap) fails, sbi->ipimap should not be initialized as ipimap.
CVE-2023-52601	In the Linux kernel, the following vulnerability has been resolved: jfs: fix array-index-out-of-bounds in dbAdjTree the dbAdjTree while accessing the dmt_stree. To add the required check added the bool is_ctl which is required to commit. https://lore.kernel.org/linux-kernel-mentees/f9475918-2186-49b8-b801-6f0f9e75f4fa@oracle.com/
CVE-2023-52602	In the Linux kernel, the following vulnerability has been resolved: jfs: fix slab-out-of-bounds Read in dtSearch Cur sorted entry table of the page there is a out of bound access. Added a bound check to fix the error. Dave: Set return

CVE-2023-52603	In the Linux kernel, the following vulnerability has been resolved: UBSAN: array-index-out-of-bounds in dtSplitR oop0: detected capacity change from 0 to 32768 UBSAN: array-index-out-of-bounds in fs/jfs/jfs_dtree.c:1971:9 in [128]' CPU: 0 PID: 3613 Comm: syz-executor270 Not tainted 6.0.0-syzkaller-09423-g493ffd6605b2 #0 Hardware Google Compute Engine, BIOS Google 09/22/2022 Call Trace: <TASK> __dump_stack lib/dump_stack.c:88 [inli dump_stack.c:106 ubsan_epilogue lib/ubsan.c:151 [inline] __ubsan_handle_out_of_bounds+0xdb/0x130 lib/ubsan jfs_dtree.c:1971 dtSplitUp fs/jfs/jfs_dtree.c:985 [inline] dtInsert+0x1189/0x6b80 fs/jfs/jfs_dtree.c:863 jfs_mkdir+0 +0x3b3/0x590 fs/namei.c:4013 do_mkdirat+0x279/0x550 fs/namei.c:4038 __do_sys_mkdirat fs/namei.c:4053 [inli [inline] __x64_sys_mkdirat+0x85/0x90 fs/namei.c:4051 do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_s common.c:80 entry_SYSCALL_64_after_hwframe+0x63/0xcd RIP: 0033:0x7fcdc0113fd9 Code: ff ff c3 66 2e 0f 89 f7 48 89 d6 48 89 ca 4d 89 c2 4d 89 c8 4c 8b 4c 24 08 0f 05 <48> 3d 01 f0 ff ff 73 01 c3 48 c7 c1 c0 ff ff ff f7 d EFLAGS: 00000246 ORIG_RAX: 0000000000000102 RAX: ffffffffffffffda RBX: 0000000000000000 RCX: 0000 RSI: 0000000020000340 RDI: 0000000000000003 RBP: 00007fcdc00d37a0 R08: 0000000000000000 R09: 0000 0000000000000246 R12: 00000000f8008000 R13: 0000000000000000 R14: 00083878000000f8 R15: 0000000000 the value of fsi becomes less than -1. The check to break the loop when fsi value becomes -1 is present but syzbot v cause the error. This patch simply add the change for the values less than 0. The patch is tested via syzbot.
CVE-2023-52604	In the Linux kernel, the following vulnerability has been resolved: FS:JFS:UBSAN:array-index-out-of-bounds in d issue: UBSAN: array-index-out-of-bounds in fs/jfs/jfs_dmap.c:2867:6 index 196694 is out of range for type 's8[130 PID: 109 Comm: jfsCommit Not tainted 6.6.0-rc3-syzkaller #0 Hardware name: Google Google Compute Engine/G 08/04/2023 Call Trace: <TASK> __dump_stack lib/dump_stack.c:88 [inline] dump_stack_lvl+0x1e7/0x2d0 lib/du ubsan.c:217 [inline] __ubsan_handle_out_of_bounds+0x11c/0x150 lib/ubsan.c:348 dbAdjTree+0x474/0x4f0 fs/jfs jfs/jfs_dmap.c:2834 dbFreeBits+0x4eb/0xda0 fs/jfs/jfs_dmap.c:2331 dbFreeDmap fs/jfs/jfs_dmap.c:2080 [inline] d txFreeMap+0x798/0xd50 fs/jfs/jfs_txnmgr.c:2534 txUpdateMap+0x342/0x9e0 txLazyCommit fs/jfs/jfs_txnmgr.c:2 fs/jfs/jfs_txnmgr.c:2732 kthread+0x2d3/0x370 kernel/kthread.c:388 ret_from_fork+0x48/0x80 arch/x86/kernel/pro arch/x86/entry/entry_64.S:304 </TASK> == Kernel panic - not syncing: UBSAN: panic_on_warn set ... CPU: 1 PID: 109 Comm: jfsCommit Not tainted 6.6.0-r Google Compute Engine/Google Compute Engine, BIOS Google 08/04/2023 Call Trace: <TASK> __dump_stack +0x1e7/0x2d0 lib/dump_stack.c:106 panic+0x30f/0x770 kernel/panic.c:340 check_panic_on_warn+0x82/0xa0 ker ubsan.c:223 [inline] __ubsan_handle_out_of_bounds+0x13c/0x150 lib/ubsan.c:348 dbAdjTree+0x474/0x4f0 fs/jfs jfs/jfs_dmap.c:2834 dbFreeBits+0x4eb/0xda0 fs/jfs/jfs_dmap.c:2331 dbFreeDmap fs/jfs/jfs_dmap.c:2080 [inline] d txFreeMap+0x798/0xd50 fs/jfs/jfs_txnmgr.c:2534 txUpdateMap+0x342/0x9e0 txLazyCommit fs/jfs/jfs_txnmgr.c:2 fs/jfs/jfs_txnmgr.c:2732 kthread+0x2d3/0x370 kernel/kthread.c:388 ret_from_fork+0x48/0x80 arch/x86/kernel/pro arch/x86/entry/entry_64.S:304 </TASK> Kernel Offset: disabled Rebooting in 86400 seconds.. The issue is caused CTLTREESIZE which is the max size of stree. Adding a simple check solves this issue. Dave: As the function retu a more intrusive code reorganization, so I modified Osama's patch at use WARN_ON_ONCE for lack of a cleaner
CVE-2023-52605	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.
CVE-2023-52606	In the Linux kernel, the following vulnerability has been resolved: powerpc/lib: Validate size for vector operations a certain maximum size for the instructions being emulated. The size of those operations however is determined se validate the assumption on the maximum size of the operations, so as to prevent any unintended kernel stack corru
CVE-2023-52607	In the Linux kernel, the following vulnerability has been resolved: powerpc/mm: Fix null-pointer dereference in pg to dynamically allocated memory which can be NULL upon failure. Ensure the allocation was successful by check
CVE-2023-52609	In the Linux kernel, the following vulnerability has been resolved: binder: fix race between mmput() and do_exit() to allocate and insert pages on a remote address space from Task B. For this, Task A pins the remote mm via mmg Task B do_exit() and the final mmput() refcount decrement will come from Task A. Task A | Task B ---------- do_exit() | exit_mm() | mmput() mmput() | exit_mmap() | remove_vma() | fput() | In this case, the work of ____fput TWA_RESUME. So in theory, Task A returns to userspace and the cleanup work gets executed. However, Task A B that never comes (it's dead). This means the binder_deferred_release() is blocked until an unrelated binder event the associated death notifications will also be delayed until then. In order to fix this use mmput_async() that will sc >async_put_work WQ instead of Task A.
CVE-2023-52612	In the Linux kernel, the following vulnerability has been resolved: crypto: scomp - fix req->dst buffer overflow Th before copying from the scomp_scratch->dst to avoid req->dst buffer overflow problem.
CVE-2023-52615	In the Linux kernel, the following vulnerability has been resolved: hwrng: core - Fix page fault dead lock on mmap device read path. This triggers when the user reads from /dev/hwrng into memory also mmap-ed from /dev/hwrng. read which then dead-locks. Fix this by using a stack buffer when calling copy_to_user.
CVE-2023-52617	In the Linux kernel, the following vulnerability has been resolved: PCI: switchtec: Fix stdev_release() crash after s removal may occur while stdev->cdev is held open. The call to stdev_release() then happens during close or exit, a Otherwise the last ref would vanish with the trailing put_device(), just before return. At that later point in time, the >mmio_mrpc mapping. Also, the stdev->pdev reference was not a counted one. Therefore, in DMA mode, the iow page fault, and the subsequent dma_free_coherent(), if reached, would pass a stale &stdev->pdev->dev pointer. Fix switchtec_pci_remove(), after stdev_kill(). Counting the stdev->pdev ref is now optional, but may prevent future ac lore.kernel.org/r/20231113212150.96410-1-dns@arista.com

CVE-2023-52619	In the Linux kernel, the following vulnerability has been resolved: pstore/ram: Fix crash when setting number of cp cpu cores is adjusted to 7 or other odd numbers, the zone size will become an odd number. The address of the zone zone1 = BASE + zone_size addr of zone2 = BASE + zone_size*2 ... The address of zone1/3/5/7 will be mapped to occur when accessing these va. So, use ALIGN_DOWN() to make sure the zone size is even to avoid this bug.
CVE-2023-52620	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_tables: disallow timeout for anonym these parameters.
CVE-2023-52622	In the Linux kernel, the following vulnerability has been resolved: ext4: avoid online resizing failures due to oversi resize an ext4 filesystem with a oversized flexbg_size, mkfs.ext4 -F -G 67108864 $dev -b 4096 100M mount $dev following WARN_ON is triggered: ========== CPU: 0 PID: 427 at mm/page_alloc.c:4402 __alloc_pages+0x411/0x550 Modules linked in: sg(E) CPU: 0 PID: 42 E 6.6.0-rc5+ #314 RIP: 0010:__alloc_pages+0x411/0x550 Call Trace: <TASK> __kmalloc_large_node+0xa2/0x2 ext4_resize_fs+0x481/0xd80 __ext4_ioctl+0x1616/0x1d90 ext4_ioctl+0x12/0x20 __x64_sys_ioctl+0xf0/0x150 do ========== This is because flexbg new_group_data array to be allocated exceeds MAX_ORDER. Currently, the minimum value of MAX_ORDER is 4096, the corresponding maximum number of groups that can be allocated is: (PAGE_SIZE << MAX_ORDER) / s And the value that is down-aligned to the power of 2 is 16384. Therefore, this value is defined as MAX_RESIZE_ time does not exceed this value during resizing, and is added multiple times to complete the online resizing. The di be more dispersed.
CVE-2023-52623	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: Fix a suspicious RCU usage warning while running cthon against an ontap server running pNFS: [57.202521] ========== suspicious RCU usage [57.202523] 6.7.0-rc3-g2cc14f52aeb7 #41492 Not tainted [57.202525] ---------- xprtmultipath.c:349 RCU-list traversed in non-reader section!! [57.202527] other info that might help us debug thi = 2, debug_locks = 1 [57.202529] no locks held by test5/3567. [57.202530] stack backtrace: [57.202532] CPU: 0 tainted 6.7.0-rc3-g2cc14f52aeb7 #41492 5b09971b4965c0aceba19f3eea324a4a806e227e [57.202534] Hardware n + ICH9, 2009), BIOS unknown 2/2/2022 [57.202536] Call Trace: [57.202537] <TASK> [57.202540] dump_stac lockdep_rcu_suspicious+0x154/0x1a0 [57.202556] rpc_xprt_switch_has_addr+0x17c/0x190 [sunrpc ebe02571b9a [57.202596] rpc_clnt_setup_test_and_add_xprt+0x50/0x180 [sunrpc ebe02571b9a8ceebf7d98e71675af20c19bdb1 rpc_clnt_add_xprt+0x254/0x300 [sunrpc ebe02571b9a8ceebf7d98e71675af20c19bdb1f6] [57.202646] rpc_clnt_a [sunrpc ebe02571b9a8ceebf7d98e71675af20c19bdb1f6] [57.202671] ? __pfx_rpc_clnt_setup_test_and_add_xprt+ ebe02571b9a8ceebf7d98e71675af20c19bdb1f6] [57.202696] nfs4_pnfs_ds_connect+0x345/0x760 [nfsv4 c716d88 [57.202728] ? __pfx_nfs4_test_session_trunk+0x10/0x10 [nfsv4 c716d88496ded0ea6d289bbea684fa996f9b57a9] +0x75/0xc0 [nfs_layout_nfsv41_files e3a4187f18ae8a27b630f9feae6831b584a9360a] [57.202760] filelayout_writ [nfs_layout_nfsv41_files e3a4187f18ae8a27b630f9feae6831b584a9360a] [57.202765] pnfs_generic_pg_writepag c716d88496ded0ea6d289bbea684fa996f9b57a9] [57.202788] __nfs_pageio_add_request+0x3fd/0x520 [nfs 6c976 [57.202813] nfs_pageio_add_request+0x18b/0x390 [nfs 6c976fa593a7c2976f5a0aeb4965514a828e6902] [57.202 6c976fa593a7c2976f5a0aeb4965514a828e6902] [57.202849] nfs_writepages_callback+0x13/0x30 [nfs 6c976fa59 [57.202866] write_cache_pages+0x265/0x450 [57.202870] ? __pfx_nfs_writepages_callback+0x10/0x10 [nfs 6c [57.202891] nfs_writepages+0x141/0x230 [nfs 6c976fa593a7c2976f5a0aeb4965514a828e6902] [57.202913] do_ filemap_fdatawrite_wbc+0x5c/0x80 [57.202921] filemap_fdatawrite_wbc+0x67/0x80 [57.202924] filemap_write [57.202930] nfs_wb_all+0x49/0x180 [nfs 6c976fa593a7c2976f5a0aeb4965514a828e6902] [57.202947] nfs4_file c716d88496ded0ea6d289bbea684fa996f9b57a9] [57.202969] __se_sys_close+0x46/0xd0 [57.202972] do_syscall do_syscall_64+0x77/0x100 [57.202976] ? do_syscall_64+0x77/0x100 [57.202979] entry_SYSCALL_64_after_h 0033:0x7fe2b12e4a94 [57.202985] Code: 00 f7 d8 64 89 01 48 83 c8 ff c3 66 2e 0f 1f 84 00 00 00 00 00 90 f3 0f 00 00 0f 05 <48> 3d 00 f0 ff ff 77 44 c3 0f 1f 00 48 83 ec 18 89 7c 24 0c e8 c3 [57.202987] RSP: 002b:00007ffe8 0000000000000003 [57.202989] RAX: ffffffffffffffda RBX: 00007ffe857dfd68 RCX: 00007fe2b12e4a94 [57.202 00007ffe857ddc40 RDI: 0000000000000003 [57.202992] RBP: 00007ffe857dfc50 R08: 7fffffffffffffff R09: 0000 truncated---

CVE-2023-52623	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: Fix a suspicious RCU usage warning while running cthon against an ontap server running pNFS: [57.202521] ============================= suspicious RCU usage [57.202523] 6.7.0-rc3-g2cc14f52aeb7 #41492 Not tainted [57.202525] ----------------------- xprtmultipath.c:349 RCU-list traversed in non-reader section!! [57.202527] other info that might help us debug thi = 2, debug_locks = 1 [57.202529] no locks held by test5/3567. [57.202530] stack backtrace: [57.202532] CPU: 0 tainted 6.7.0-rc3-g2cc14f52aeb7 #41492 5b09971b4965c0aceba19f3eea324a4a806e227e [57.202534] Hardware n + ICH9, 2009), BIOS unknown 2/2/2022 [57.202536] Call Trace: [57.202537] <TASK> [57.202540] dump_stac lockdep_rcu_suspicious+0x154/0x1a0 [57.202556] rpc_xprt_switch_has_addr+0x17c/0x190 [sunrpc ebe02571b9a [57.202596] rpc_clnt_setup_test_and_add_xprt+0x50/0x180 [sunrpc ebe02571b9a8ceebf7d98e71675af20c19bdb1 rpc_clnt_add_xprt+0x254/0x300 [sunrpc ebe02571b9a8ceebf7d98e71675af20c19bdb1f6] [57.202646] rpc_clnt_a [sunrpc ebe02571b9a8ceebf7d98e71675af20c19bdb1f6] [57.202671] ? __pfx_rpc_clnt_setup_test_and_add_xprt+ ebe02571b9a8ceebf7d98e71675af20c19bdb1f6] [57.202696] nfs4_pnfs_ds_connect+0x345/0x760 [nfsv4 c716d88 [57.202728] ? __pfx_nfs4_test_session_trunk+0x10/0x10 [nfsv4 c716d88496ded0ea6d289bbea684fa996f9b57a9] +0x75/0xc0 [nfs_layout_nfsv41_files e3a4187f18ae8a27b630f9feae6831b584a9360a] [57.202760] filelayout_writ [nfs_layout_nfsv41_files e3a4187f18ae8a27b630f9feae6831b584a9360a] [57.202765] pnfs_generic_pg_writepag c716d88496ded0ea6d289bbea684fa996f9b57a9] [57.202788] __nfs_pageio_add_request+0x3fd/0x520 [nfs 6c976 [57.202813] nfs_pageio_add_request+0x18b/0x390 [nfs 6c976fa593a7c2976f5a0aeb4965514a828e6902] [57.202 6c976fa593a7c2976f5a0aeb4965514a828e6902] [57.202849] nfs_writepages_callback+0x13/0x30 [nfs 6c976fa59 [57.202866] write_cache_pages+0x265/0x450 [57.202870] ? __pfx_nfs_writepages_callback+0x10/0x10 [nfs 6c [57.202891] nfs_writepages+0x141/0x230 [nfs 6c976fa593a7c2976f5a0aeb4965514a828e6902] [57.202913] do_ filemap_fdatawrite_wbc+0x5c/0x80 [57.202921] filemap_fdatawrite_wbc+0x67/0x80 [57.202924] filemap_write [57.202930] nfs_wb_all+0x49/0x180 [nfs 6c976fa593a7c2976f5a0aeb4965514a828e6902] [57.202947] nfs4_file c716d88496ded0ea6d289bbea684fa996f9b57a9] [57.202969] __se_sys_close+0x46/0xd0 [57.202972] do_syscall do_syscall_64+0x77/0x100 [57.202976] ? do_syscall_64+0x77/0x100 [57.202979] entry_SYSCALL_64_after_h 0033:0x7fe2b12e4a94 [57.202985] Code: 00 f7 d8 64 89 01 48 83 c8 ff c3 66 2e 0f 1f 84 00 00 00 00 00 90 f3 0f 00 00 0f 05 <48> 3d 00 f0 ff ff 77 44 c3 0f 1f 00 48 83 ec 18 89 7c 24 0c e8 c3 [57.202987] RSP: 002b:00007ffe8 0000000000000003 [57.202989] RAX: ffffffffffffffda RBX: 00007ffe857dfd68 RCX: 00007fe2b12e4a94 [57.202 00007ffe857ddc40 RDI: 0000000000000003 [57.202992] RBP: 00007ffe857dfc50 R08: 7fffffffffffffff R09: 0000 truncated---
CVE-2023-52628	In the Linux kernel, the following vulnerability has been resolved: netfilter: nftables: exthdr: fix 4-byte stack OOB dst[len / 4] can write past the destination array which leads to stack corruption. This construct is necessary to clean NOT a multiple of the register size, so make it conditional just like nft_payload.c does. The bug was added in 4.1 c and ip option support was added. Bug reported by Zero Day Initiative project (ZDI-CAN-21950, ZDI-CAN-21951,
CVE-2023-52629	In the Linux kernel, the following vulnerability has been resolved: sh: push-switch: Reorder cleanup operations to code puts flush_work() before timer_shutdown_sync() in switch_drv_remove(). Although we use flush_work() to s in switch_timer(). As a result, a use-after-free bug can occur. The details are shown below: (cpu 0) | (cpu 1) switch_ switch_timer // timer | schedule_work(&psw->work) timer_shutdown_sync() | ... | switch_work_handler // worker This patch puts timer_shutdown_sync() before flush_work() to mitigate the bugs. As a result, the worker and timer operations.
CVE-2023-52630	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.
CVE-2023-52637	In the Linux kernel, the following vulnerability has been resolved: can: j1939: Fix UAF in j1939_sk_match_filter c Lock jsk->sk to prevent UAF when setsockopt(..., SO_J1939_FILTER, ...) modifies jsk->filters while receiving pa affected system: ========== BUG j1939_sk_recv_match_one+0x1af/0x2d0 [can_j1939] Read of size 4 at addr ffff888012144014 by task j1939/350 W OE 6.5.0-rc5 #1 Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 1.13.0-1ubuntu1.1 04/01/2 kasan_complete_mode_report_info+0x7d/0x200 ? j1939_sk_recv_match_one+0x1af/0x2d0 [can_j1939] kasan_rep +0x1af/0x2d0 [can_j1939] __asan_load4+0x84/0xb0 j1939_sk_recv_match_one+0x1af/0x2d0 [can_j1939] j1939_ __kasan_check_write+0x18/0x20 ? __pfx_j1939_sk_recv+0x10/0x10 [can_j1939] ? j1939_simple_recv+0x69/0x2 +0x5e/0x310 [can_j1939] j1939_can_recv+0x43f/0x580 [can_j1939] ? __pfx_j1939_can_recv+0x10/0x10 [can_j1 __pfx_j1939_can_recv+0x10/0x10 [can_j1939] can_rcv_filter+0x11f/0x350 [can] can_receive+0x12f/0x190 [can] can_rcv+0xdd/0x130 [can] ? __pfx_can_rcv+0x10/0x10 [can] __netif_receive_skb_one_core+0x13d/0x150 ? __pf +0x10/0x10 ? __kasan_check_write+0x18/0x20 ? _raw_spin_lock_irq+0x8c/0xe0 __netif_receive_skb+0x23/0xb0 __napi_poll+0x69/0x310 net_rx_action+0x2a1/0x580 ? __pfx_net_rx_action+0x10/0x10 ? __pfx__raw_spin_lock __do_softirq+0xf3/0x3f8 do_softirq+0x53/0x80 </IRQ> <TASK> __local_bh_enable_ip+0x6e/0x70 netif_rx+0x1 [can] ? __pfx_can_send+0x10/0x10 [can] ? __check_object_size+0x299/0x410 raw_sendmsg+0x572/0x6d0 [can_ [can_raw] ? apparmor_socket_sendmsg+0x2f/0x40 ? __pfx_raw_sendmsg+0x10/0x10 [can_raw] sock_sendmsg+0 __pfx_sock_write_iter+0x10/0x10 ? __rtnl_unlock+0x47/0x80 ? security_file_permission+0x54/0x320 vfs_write+ __fget_light+0x1ca/0x1f0 ? __rcu_read_unlock+0x5b/0x280 ksys_write+0x143/0x170 ? __pfx_ksys_write+0x10/0 fpregs_assert_state_consistent+0x62/0x70 __x64_sys_write+0x47/0x60 do_syscall_64+0x60/0x90 ? do_syscall_64 exc_page_fault+0x79/0xf0 entry_SYSCALL_64_after_hwframe+0x6e/0xd8 Allocated by task 348: kasan_save_st kasan_save_alloc_info+0x1f/0x30 __kasan_kmalloc+0xb5/0xc0 __kmalloc_node_track_caller+0x67/0x160 j1939_ __sys_setsockopt+0x15c/0x2f0 __x64_sys_setsockopt+0x6b/0x80 do_syscall_64+0x60/0x90 entry_SYSCALL_64 task 349: kasan_save_stack+0x2a/0x50 kasan_set_track+0x29/0x40 kasan_save_free_info+0x2f/0x50 __kasan_sla +0x1b9/0x380 kfree+0x7a/0x120 j1939_sk_setsockopt+0x3b2/0x450 [can_j1939] __sys_setsockopt+0x15c/0x2f0 do_syscall_64+0x60/0x90 entry_SYSCALL_64_after_hwframe+0x6e/0xd8

CVE-2023-52644	In the Linux kernel, the following vulnerability has been resolved: wifi: b43: Stop/wake correct queue in DMA Tx disabled, the queue priority value will not map to the correct ieee80211 queue since there is only one queue. Stop/v trying to stop/wake a non-existent queue and failing to stop/wake the actual queue instantiated. Log of issue before [+5.112651] ------------[cut here]------------ [+0.000005] WARNING: CPU: 7 PID: 25513 at net/mac80211/util.c [mac80211] [+0.000067] Modules linked in: b43(O) snd_seq_dummy snd_hrtimer snd_seq snd_seq_device nft_ch xfrm_user xfrm_algo xt_addrtype overlay ccm af_packet amdgpu snd_hda_codec_cirrus snd_hda_codec_generic l xt_conntrack nf_conntrack nf_defrag_ipv6 nf_defrag_ipv4 ip6t_rpfilter ipt_rpfilter xt_pkttype xt_LOG nf_log_sys nfnetlink sch_fq_codel btusb uinput iTCO_wdt ctr btrtl intel_pmc_bxt i915 intel_rapl_msr mei_hdcp mei_pxp joyd radeon btbcm vivaldi_fmap btmtk intel_rapl_common snd_hda_codec_hdmi bluetooth uvcvideo nls_iso8859_1 ap snd_hda_intel intel_powerclamp vfat videobuf2_vmalloc coretemp fat snd_intel_dspcfg crc32_pclmul uvc polyval videobuf2_memops snd_hda_codec tun drm_suballoc_helper polyval_generic drm_ttm_helper drm_buddy tap ecd macvlan ttm ghash_clmulni_intel ecc tg3 [+0.000044] videodev bridge snd_hda_core rapl crc16 drm_display_help intel_cstate bcm5974 hid_appleir videobuf2_common stp mac_hid libphy snd_pcm drm_kms_helper acpi_als mei_ industrialio_triggered_buffer apple_mfi_fastcharge i2c_i801 mei snd lpc_ich agpgart ptp i2c_smbus thunderbolt ap industrialio soundcore pps_core wmi tiny_power_button sbs sbshc button ac cordic bcma mac80211 cfg80211 ssb fuse backlight firmware_class efi_pstore configfs efivarfs dmi_sysfs ip_tables x_tables autofs4 dm_crypt cbc encry tpm rng_core input_leds hid_apple led_class hid_generic usbhid hid sd_mod t10_pi crc64_rocksoft crc64 crc_t10d uhci_hcd ehci_pci ehci_hcd crct10dif_pclmul crct10dif_common sha512_ssse3 sha512_generic sha256_ssse3 sha1 crypto_simd cryptd scsi_common [+0.000055] usb_common rtc_cmos btrfs blake2b_generic libcrc32c crc32c_ge dm_bufio dm_mod dax [last unloaded: b43(O)] [+0.000009] CPU: 7 PID: 25513 Comm: irq/17-b43 Tainted: G W Hardware name: Apple Inc. MacBookPro8,3/Mac-942459F5819B171B, BIOS 87.0.0.0.0 06/13/2019 [+0.000001] +0xd5/0x180 [mac80211] [+0.000046] Code: 00 45 85 e4 0f 85 9b 00 00 00 48 8d bd 40 09 00 00 f0 48 0f ba ad 4 41 5e e9 cb 6d 3c d0 <0f> 0b 5b 5d 41 5c 41 5d 41 5e c3 cc cc cc cc 48 8d b4 16 94 00 00 [+0.000002] RSP: 0018 [+0.000001] RAX: 0000000000000001 RBX: 0000000000000002 RCX: 0000000000000000 [+0.000001] RDX: RDI: ffff88820b924900 [+0.000002] RBP: ffff88820b924900 R08: ffffc90003c77d90 R09: 000000000003bfd0 [R11: ffffc90003c77c68 R12: 0000000000000000 [+0.000001] R13: 0000000000000000 R14: ffffc90003c77d90 F 0000000000000000(0000) GS:ffff88846fb80000(0000) knlGS:0000000000000000 [+0.000001] CS: 0010 DS: 0 -
CVE-2023-52650	In the Linux kernel, the following vulnerability has been resolved: drm/tegra: dsi: Add missing check for of_find_d value of of_find_device_by_node() and return the error if it fails in order to avoid NULL pointer dereference.
CVE-2023-52654	In the Linux kernel, the following vulnerability has been resolved: io_uring/af_unix: disable sending io_uring over lots of problems for io_uring in the past, and it still doesn't work exactly right and races with unix_stream_read_ge disallow sending io_uring files via sockets via SCM_RIGHT, so there are no possible cycles invloving registered f the io_uring side unnecessary.
CVE-2023-52655	In the Linux kernel, the following vulnerability has been resolved: usb: aqc111: check packet for fixup for true lim 0 and sizeof(u64) the value passed to skb_trim() as length will wrap around ending up as some very large value. Th located at that position, which will either oops or process some random value. The fix is to check against sizeof(u6 does. The issue exists since the introduction of the driver.
CVE-2023-52656	In the Linux kernel, the following vulnerability has been resolved: io_uring: drop any code related to SCM_RIGHT for passing io_uring fds over SCM_RIGHTS, get rid of it.
CVE-2023-52670	In the Linux kernel, the following vulnerability has been resolved: rpmsg: virtio: Free driver_override when rpmsg when rpmsg_remove(), otherwise the following memory leak will occur: unreferenced object 0xffff0000d55d7080 pid 56, jiffies 4294893188 (age 214.272s) hex dump (first 32 bytes): 72 70 6d 73 67 5f 6e 73 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ..... backtrace: [<000000009c94c9c1>] __kmem_cache_alloc_node+0x1f8 __kmalloc_node_track_caller+0x44/0x70 [<00000000228a60c3>] kstrndup+0x4c/0x90 [<0000000077158695>] d [<000000003e9c4ea5>] rpmsg_register_device_override+0x98/0x170 [<000000001c0c89a8>] rpmsg_ns_register_ rpmsg_probe+0x2e0/0x3ec [<00000000e65a68df>] virtio_dev_probe+0x1c0/0x280 [<00000000443331cc>] really __driver_probe_device+0x78/0xe0 [<00000000a41c9a5b>] driver_probe_device+0xd8/0x160 [<000000009c3bd5 [<0000000043cd7614>] bus_for_each_drv+0x7c/0xd4 [<000000003b929a36>] __device_attach+0x9c/0x19c [<00 +0x14/0x20 [<000000003c999637>] bus_probe_device+0xa0/0xac
CVE-2023-52672	In the Linux kernel, the following vulnerability has been resolved: pipe: wakeup wr_wait after setting max_usage C notification queue support") a regression was introduced that would lock up resized pipes under certain conditions. resizing the pipe ring size was moved to a different function, doing that moved the wakeup for pipe->wr_wait befo pipe was full before the resize occured it would result in the wakeup never actually triggering pipe_write. Set @ma writers if this isn't a watch queue. [Christian Brauner <brauner@kernel.org>: rewrite to account for watch queues]
CVE-2023-52675	In the Linux kernel, the following vulnerability has been resolved: powerpc/imc-pmu: Add a null pointer check in u a pointer to dynamically allocated memory which can be NULL upon failure.
CVE-2023-52679	In the Linux kernel, the following vulnerability has been resolved: of: Fix double free in of_parse_phandle_with_a of_parse_phandle_with_args_map() the inner loop that iterates through the map entries calls of_node_put(new) to f iteration of the inner loop. This assumes that the value of "new" is NULL on the first iteration of the inner loop. Ma outer loop by setting "new" to NULL after its value is assigned to "cur". Extend the unittest to detect the double fre triggers this path.

CVE-2023-52683	In the Linux kernel, the following vulnerability has been resolved: ACPI: LPIT: Avoid u32 multiplication overflow a possibility of overflow in multiplication, if tsc_khz is large enough (> UINT_MAX/1000). Change multiplication Verification Center (linuxtesting.org) with SVACE.
CVE-2023-52686	In the Linux kernel, the following vulnerability has been resolved: powerpc/powernv: Add a null pointer check in c to dynamically allocated memory which can be NULL upon failure.
CVE-2023-52690	In the Linux kernel, the following vulnerability has been resolved: powerpc/powernv: Add a null pointer check to s pointer to dynamically allocated memory which can be NULL upon failure. Add a null pointer check, and release '
CVE-2023-52691	In the Linux kernel, the following vulnerability has been resolved: drm/amd/pm: fix a double-free in si_dpm_init V >pm.dpm.dyn_state.vddc_dependency_on_dispclk.entries fails, amdgpu_free_extended_power_table is called to fr control flow returns to si_dpm_sw_init, it goes to label dpm_failed and calls si_dpm_fini, which calls amdgpu_free fields again. Thus a double-free is triggered.
CVE-2023-52693	In the Linux kernel, the following vulnerability has been resolved: ACPI: video: check for error while searching fo called in acpi_video_dev_register_backlight() fails, for example, because acpi_ut_acquire_mutex() fails inside acpi (uninitialized) acpi_parent handle being passed to acpi_get_pci_dev() for detecting the parent pci device. Check ac only in case of success. Found by Linux Verification Center (linuxtesting.org) with SVACE.
CVE-2023-52696	In the Linux kernel, the following vulnerability has been resolved: powerpc/powernv: Add a null pointer check in c pointer to dynamically allocated memory which can be NULL upon failure.
CVE-2023-52699	In the Linux kernel, the following vulnerability has been resolved: sysv: don't call sb_bread() with pointers_lock he in SysV filesystem [1], for sb_bread() is called with rw_spinlock held. A "write_lock(&pointers_lock) => read_loc "sb_bread() with write_lock(&pointers_lock)" bug were introduced by "Replace BKL for chain locking with sysvfs "[PATCH] err1-40: sysvfs locking fix" in Linux 2.6.8 fixed the former bug by moving pointers_lock lock to the cal with read_lock(&pointers_lock)" bug (which made this problem easier to hit). Al Viro suggested that why not to do in Minix filesystem does. And doing like that is almost a revert of "[PATCH] err1-40: sysvfs locking fix" except th called without write_lock(&pointers_lock).
CVE-2023-52704	In the Linux kernel, the following vulnerability has been resolved: freezer,umh: Fix call_usermode_helper_exec() v f5d39b020809 ("freezer,sched: Rewrite core freezer logic") broke call_usermodehelper_exec() for the KILLABLE second, unconditional, wait_for_completion() was not optional and ensures the on-stack completion is unused befo
CVE-2023-52706	In the Linux kernel, the following vulnerability has been resolved: gpio: sim: fix a memory leak Fix an inverted log to GPIO hog structures never being freed.
CVE-2023-52752	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix use-after-free bug in cifs_debug that are being teared down (e.g. @ses->ses_status == SES_EXITING) in cifs_debug_data_proc_show() to avoid us following GPF when reading from /proc/fs/cifs/DebugData while mounting and umounting [816.251274] general p canonical address 0x6b6b6b6b6b6b6d81: 0000 [#1] PREEMPT SMP NOPTI ... [816.260138] Call Trace: [816.26 die_addr+0x36/0x90 [816.260762] ? exc_general_protection+0x1b3/0x410 [816.261126] ? asm_exc_general_pro cifs_debug_tcon+0xbd/0x240 [cifs] [816.261878] ? cifs_debug_tcon+0xab/0x240 [cifs] [816.262249] cifs_debug [816.262689] ? seq_read_iter+0x379/0x470 [816.262995] seq_read_iter+0x118/0x470 [816.263291] proc_reg_re srso_alias_return_thunk+0x5/0x7f [816.263945] vfs_read+0x201/0x350 [816.264211] ksys_read+0x75/0x100 [8 [816.264750] entry_SYSCALL_64_after_hwframe+0x6e/0xd8 [816.265135] RIP: 0033:0x7fd5e669d381
CVE-2023-52753	In the Linux kernel, the following vulnerability has been resolved: drm/amd/display: Avoid NULL dereference of t whether assigned timing generator is NULL or not before accessing its funcs to prevent NULL dereference.
CVE-2023-52764	In the Linux kernel, the following vulnerability has been resolved: media: gspca: cpia1: shift-out-of-bounds in set_ UBSAN: shift-out-of-bounds in drivers/media/usb/gspca/cpia1.c:1031:27 shift exponent 245 is too large for 32-bit >params.exposure.gain" exceeds the number of bits in an integer, a shift-out-of-bounds error is reported. It is trigge be left-shifted by more than the number of bits in an integer. In order to avoid invalid range during left-shift, the co
CVE-2023-52771	In the Linux kernel, the following vulnerability has been resolved: cxl/port: Fix delete_endpoint() vs parent unregis ->probe() time, establishes a lineage of ports (struct cxl_port objects) between an endpoint and the root of a CXL to port is attached to the cxl_port driver. Given that setup, it follows that when either any port in that lineage goes thro memdev goes through a cxl_mem ->remove() event. The hierarchy below the removed port, or the entire hierarchy down. The delete_endpoint() callback is careful to check whether it is being called to tear down the hierarchy, or if memdev because an ancestor port is going through ->remove(). That care needs to take the device_lock() of the enc be fixed: 1/ A reference on the parent is needed to prevent use-after-free scenarios like this signature: BUG: spinloc Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS edk2-20230524-3.fc38 05/24/2023 Workqueue: 0010:spin_bug+0x65/0xa0 Call Trace: do_raw_spin_lock+0x69/0xa0 __mutex_lock+0x695/0xb80 delete_endpoin +0xb8/0x110 device_unbind_cleanup+0xe/0x70 device_release_driver_internal+0x1d2/0x210 detach_memdev+0x +0x1e3/0x4c0 worker_thread+0x1dd/0x3d0 2/ In the case of RCH topologies, the parent device that needs to be loc cxl_mem_find_port(), use endpoint->dev.parent instead.

CVE-2023-52774	In the Linux kernel, the following vulnerability has been resolved: s390/dasd: protect device queue against concurr of requests on the device queue are counted. The access to the device queue is unprotected against concurrent acces with alias devices enabled, the device queue can change while dasd_profile_start() is accessing the queue. In the w incorrect pointer accesses. Fix this by taking the device lock before accessing the queue and counting the requests. pointer can be done earlier to avoid unnecessary locking in a hot path.
CVE-2023-52778	In the Linux kernel, the following vulnerability has been resolved: mptcp: deal with large GSO size After the blam the MPTCP subflows) can build egress packets larger than 64K. That exceeds the maximum DSS data size, the len stream being corrupted, as later observed on the receiver: WARNING: CPU: 0 PID: 9696 at net/mptcp/protocol.c:7 +0x2604/0x26e0 CPU: 0 PID: 9696 Comm: syz-executor.7 Not tainted 6.6.0-rc5-gcd8bdf563d46 #45 Hardware na 1996), BIOS 1.11.0-2.el7 04/01/2014 netlink: 8 bytes leftover after parsing attributes in process `syz-executor.4'. R +0x2604/0x26e0 net/mptcp/protocol.c:705 RSP: 0018:ffffc90000006e80 EFLAGS: 00010246 RAX: ffffffff83e9f6 ffff888102ad0000 netlink: 8 bytes leftover after parsing attributes in process `syz-executor.4'. RDX: 000000008000 RDI: 0000000000003908 RBP: ffffc90000007110 R08: ffffffff83e9e078 R09: 1ffff1100e548c8a R10: dfffc00000 0000000000013908 R13: dfffc0000000000 R14: 0000000000003908 R15: 000000000031cf29 FS: 00007f239c47 knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 00007f239c45cd78 CR3: 0000000000770ef0 DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000000000000 DR3: 0000000 0000000000000600 PKRU: 55555554 Call Trace: <IRQ> mptcp_data_ready+0x263/0xac0 net/mptcp/protocol.c:8 net/mptcp/subflow.c:1409 tcp_data_queue+0x21a1/0x7a60 net/ipv4/tcp_input.c:5151 tcp_rcv_established+0x950/0 tcp_v6_do_rcv+0x554/0x12f0 net/ipv6/tcp_ipv6.c:1483 tcp_v6_rcv+0x2e26/0x3810 net/ipv6/tcp_ipv6.c:1749 ip6_ ipv6/ip6_input.c:438 ip6_input+0x1c5/0x470 net/ipv6/ip6_input.c:483 ipv6_rcv+0xef/0x2c0 include/linux/netfilte net/core/dev.c:5532 process_backlog+0x353/0x660 net/core/dev.c:5974 __napi_poll+0xc6/0x5a0 net/core/dev.c:65 dev.c:6603 __do_softirq+0x184/0x524 kernel/softirq.c:553 do_softirq+0xdd/0x130 kernel/softirq.c:454 Address th GSO size to what MPTCP actually allows.
CVE-2023-52784	In the Linux kernel, the following vulnerability has been resolved: bonding: stop the device in bond_setup_by_slav only support ethernet devices") has been able to keep syzbot away from net/lapb, until today. In the following splat has been created on a bonding device without members. Then adding a non ARPHRD_ETHER member forced the is to make sure we call dev_close() in bond_setup_by_slave() so that the potential linked lapbether devices (or any physical device) are removed. A similar bug has been addressed in commit 40baec225765 ("bonding: fix panic on [1] skbuff: skb_under_panic: text:ffff800089508810 len:44 put:40 head:ffff0000c78e7c00 data:ffff0000c78e7bea t BUG at net/core/skbuff.c:192 ! Internal error: Oops - BUG: 00000000f2000800 [#1] PREEMPT SMP Modules lin executor383 Not tainted 6.6.0-rc3-syzkaller-gbf6547d8715b #0 Hardware name: Google Google Compute Engine/ 08/04/2023 pstate: 60400005 (nZCv daif +PAN -UAO -TCO -DIT -SSBS BTYPE=--) pc : skb_panic net/core/skbu +0x13c/0x140 net/core/skbuff.c:202 lr : skb_panic net/core/skbuff.c:188 [inline] lr : skb_under_panic+0x13c/0x14 x29: ffff800096a06ab0 x28: ffff800096a06ba0 x27: dfff800000000000 x26: ffff0000ce9b9b50 x25: 000000000000 ffff0000c78e7c00 x22: 000000000000002c x21: 0000000000000140 x20: 0000000000000028 x19: ffff800089508 0000000000000000 x16: ffff80008a629a3c x15: 0000000000000001 x14: 1fffe00036837a32 x13: 000000000000 0000000000000201 x10: 0000000000000000 x9 : cb50b496c519aa00 x8 : cb50b496c519aa00 x7 : 000000000000 ffff800096a063b8 x4 : ffff80008e280f80 x3 : ffff8000805ad11c x2 : 0000000000000001 x1 : 0000000100000201 net/core/skbuff.c:188 [inline] skb_under_panic+0x13c/0x140 net/core/skbuff.c:202 skb_push+0xf0/0x108 net/core net/ipv6/ip6_gre.c:1384 dev_hard_header include/linux/netdevice.h:3136 [inline] lapbeth_data_transmit+0x1c4/0x lapb_data_transmit+0x8c/0xb0 net/lapb/lapb_iface.c:447 lapb_transmit_buffer+0x178/0x204 net/lapb/lapb_out.c:1 net/lapb/lapb_subr.c:251 __lapb_disconnect_request+0x9c/0x17c net/lapb/lapb_iface.c:326 lapb_device_event+0x notifier_call_chain+0x1a4/0x510 kernel/notifier.c:93 raw_notifier_call_chain+0x3c/0x50 kernel/notifier.c:461 call dev.c:1970 [inline] call_netdevice_notifiers_extack net/core/dev.c:2008 [inline] call_netdevice_notifiers net/core/d +0x1b8/0x3c4 net/core/dev.c:1508 dev_close_many+0x1e0/0x470 net/core/dev.c:1559 dev_close+0x174/0x250 ne +0x2e4/0x958 drivers/net/wan/lapbether.c:466 notifier_call_chain+0x1a4/0x510 kernel/notifier.c:93 raw_notifier_ call_netdevice_notifiers_info net/core/dev.c:1970 [inline] call_netdevice_notifiers_extack net/core/dev.c:2008 [inli dev.c:2022 [inline] __dev_close_many+0x1b8/0x3c4 net/core/dev.c:1508 dev_close_many+0x1e0/0x470 net/core/ core/dev.c:1585 bond_enslave+0x2298/0x30cc drivers/net/bonding/bond_main.c:2332 bond_do_ioctl+0x268/0xc6 dev_ifsioc+0x754/0x9ac dev_ioctl+0x4d8/0xd34 net/core/dev_ioctl.c:786 sock_do_ioctl+0x1d4/0x2d0 net/socket. socket.c:1322 vfs_ioctl fs/ioctl.c:51 [inline] __do_ ---truncated---
CVE-2023-52786	In the Linux kernel, the following vulnerability has been resolved: ext4: fix racy may inline data check in dio write from ext4_iomap_begin() triggers as of the commit referenced below: if (WARN_ON_ONCE(ext4_has_inline_dat during a dio write, which is never expected to encounter an inode with inline data. To enforce this behavior, ext4_c state of the inode and clears the MAY_INLINE_DATA state flag to either fall back to buffered writes, or enforce t are not allowed to create inline data. The problem is that the check for existing inline data and the state flag can spa originally locked shared and subsequently upgraded to exclusive, another writer may have reacquired the lock and acquires the lock and proceeds. The commit referenced below loosens the lock requirements to allow some forms o lock, but AFAICT the inline data check was technically already racy for any dio write that would have involved a l bit to the same lock critical section that checks for preexisting inline data on the inode to close the race.
CVE-2023-52787	In the Linux kernel, the following vulnerability has been resolved: blk-mq: make sure active queue usage is held fo blk_integrity_unregister() can come if queue usage counter isn't held for one bio with integrity prepared, so this rec >complete_fn, then kernel panic. Another constraint is that bio_integrity_prep() needs to be called before bio merg with one queue usage counter grabbed reliably - call bio_integrity_prep() before bio merge
CVE-2023-52789	In the Linux kernel, the following vulnerability has been resolved: tty: vcc: Add check for kstrdup() in vcc_probe() and return the error, if it fails in order to avoid NULL pointer dereference.

CVE-2023-52791	In the Linux kernel, the following vulnerability has been resolved: i2c: core: Run atomic i2c xfer when !preemptibl non-atomic if preemption is disabled. However, non-atomic i2c transfers require preemption (e.g. in wait_for_com panic() calls preempt_disable_notrace() before calling emergency_restart(). Therefore, if an i2c device is used for t This avoids warnings like: [12.667612] WARNING: CPU: 1 PID: 1 at kernel/rcu/tree_plugin.h:318 rcu_note_cont Voluntary context switch within RCU read-side critical section! ... [12.742376] schedule_timeout from wait_for_c wait_for_completion_timeout from tegra_i2c_wait_completion+0x40/0x70 ... [12.994527] atomic_notifier_call_c [13.001050] machine_restart from panic+0x2a8/0x32c Use !preemptible() instead, which is basically the same che
CVE-2023-52792	In the Linux kernel, the following vulnerability has been resolved: cxl/region: Do not try to cleanup after cxl_regio ("cxl/region: decrement ->nr_targets on error in cxl_region_attach()") tried to avoid 'eiw' initialization errors when decrementing ->nr_targets when cxl_region_setup_targets() failed. Commit 86987c766276 ("cxl/region: Cleanup t cleanup to also clear cxled->pos and p->targets[pos]. The initialization error was incidentally fixed separately by: C port setup uninitialized variable warnings") which was merged a few days after 5e42bcbc3fef. But now the original fails prevents endpoint and switch decoder resources from being reused: 1) the cleanup does not set the decoder's re dpa_size_store() calls returning -EBUSY 2) the decoder is not properly freed, which results in future commit errors that the initialization errors were fixed separately, the proper cleanup for this case is to just return immediately. The get cleanup up as normal when the failed region is deleted. The ->nr_targets decrement in the error case also helpe add a new check to prevent against that overflow. Tested by trying to create an invalid region for a 2 switch * 2 end creating a valid region.
CVE-2023-52794	In the Linux kernel, the following vulnerability has been resolved: thermal: intel: powerclamp: fix mismatch in get [444.853098] BUG: KASAN: global-out-of-bounds in param_get_int+0x77/0x90 [444.853111] Read of size 4 at [444.853442] The buggy address belongs to the variable: [444.853443] max_idle+0x0/0xffffffffffffcde0 [intel_po param_get_int and the definition of max_idle. Replacing param_get_int with param_get_byte resolves this issue.
CVE-2023-52795	In the Linux kernel, the following vulnerability has been resolved: vhost-vdpa: fix use after free in vhost_vdpa_pro vhost_vdpa_release_dev() which calls ida_simple_remove() and frees "v". So this call to ida_simple_remove() is a
CVE-2023-52796	In the Linux kernel, the following vulnerability has been resolved: ipvlan: add ipvlan_route_v6_outbound() helper multiple ipvlan devices. Reduce stack size needed in ipvlan_process_v6_outbound() by moving the flowi6 struct us helper. ipvlan_route_v6_outbound() needs 120 bytes on the stack, immediately reclaimed. Also make sure ipvlan_p might also have to lower MAX_NEST_DEV, because only syzbot uses setups with more than four stacked devices ffffc9000e803ff8 (stack is ffffc9000e804000..ffffc9000e808000) stack guard page: 0000 [#1] SMP KASAN CPU: tainted 6.1.52-syzkaller #0 Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Goo +0x4/0x2a0 mm/kasan/generic.c:188 Code: 48 01 c6 48 89 c7 e8 db 4e c1 03 31 c0 5d c3 cc 0f 0b eb 02 0f 0b b8 e 48 89 e5 <41> 57 41 56 41 55 41 54 53 b0 01 48 85 f6 0f 84 a4 01 00 00 48 89 RSP: 0018:ffffc9000e804000 EFLA RBX: 0000000000000000 RCX: ffffffff817e5bf2 RDX: 0000000000000000 RSI: 0000000000000008 RDI: ffffffff 0000000000000000 R09: 0000000000000000 R10: 0000000000000000 R11: dffffc0000000001 R12: 1ffff92001d ffffffff87e6b100 R15: 0000000000000000 FS: 00007fd0c55826c0(0000) GS:ffff8881f6800000(0000) knlGS:0000 CR0: 0000000080050033 CR2: ffffc9000e803ff8 CR3: 0000000170ef7000 CR4: 00000000003506f0 DR0: 00000 0000000000000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 Call Trace: <#D __kasan_check_read+0x11/0x20 mm/kasan/shadow.c:31 [<ffffffff817e5bf2>] instrument_atomic_read include/linu [<ffffffff817e5bf2>] _test_bit include/asm-generic/bitops/instrumented-non-atomic.h:141 [inline] [<ffffffff817e5b cpumask.h:506 [inline] [<ffffffff817e5bf2>] cpu_online include/linux/cpumask.h:1092 [inline] [<ffffffff817e5bf2> events/lock.h:24 [inline] [<ffffffff817e5bf2>] lock_acquire+0xe2/0x590 kernel/locking/lockdep.c:5632 [<ffffffff85 include/linux/rcupdate.h:306 [<ffffffff8561464d>] rcu_read_lock include/linux/rcupdate.h:747 [inline] [<ffffffff85 net/ipv6/route.c:2221 [<ffffffff85618120>] ip6_pol_route_output+0x50/0x80 net/ipv6/route.c:2606 [<ffffffff856f6 ip6_fib.h:584 [inline] [<ffffffff856f65b5>] fib6_rule_lookup+0x265/0x620 net/ipv6/fib6_rules.c:116 [<ffffffff856 +0x2d9/0x3a0 net/ipv6/route.c:2638 [<ffffffff8561821a>] ip6_route_output_flags+0xca/0x340 net/ipv6/route.c:26 include/net/ip6_route.h:100 [inline] [<ffffffff838bd5a3>] ipvlan_process_v6_outbound drivers/net/ipvlan/ipvlan_c ipvlan_process_outbound drivers/net/ipvlan/ipvlan_core.c:529 [inline] [<ffffffff838bd5a3>] ipvlan_xmit_mode_l3 [<ffffffff838bd5a3>] ipvlan_queue_xmit+0xc33/0x1be0 drivers/net/ipvlan/ipvlan_core.c:677 [<ffffffff838c2909>] ipvlan/ipvlan_main.c:229 [<ffffffff84d03900>] netdev_start_xmit include/linux/netdevice.h:4966 [inline] [<ffffffff [inline] [<ffffffff84d03900>] dev_hard_start_xmit+0x320/0x980 net/core/dev.c:3660 [<ffffffff84d080e2>] __dev_ dev.c:4324 [<ffffffff855ce4cd>] dev_queue_xmit include/linux/netdevice.h:3067 [inline] [<ffffffff855ce4cd>] neig [inline] [<f ---truncated---
CVE-2023-52799	In the Linux kernel, the following vulnerability has been resolved: jfs: fix array-index-out-of-bounds in dbFindLea for sufficient free blocks there is an array out of bounds while getting element in tp->dm_stree. To add the required determine the type of dmtree. Thus added an extra parameter to dbFindLeaf so that the type of tree can be determin
CVE-2023-52801	In the Linux kernel, the following vulnerability has been resolved: iommufd: Fix missing update of domains_itree a if the original iopt_area has filled a domain and is linked to domains_itree, pages_nodes have to be properly reinser corrupted and we will UAF.

CVE-2023-52803	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: Fix RPC client cleaned up the freed p dentries cleanup is in separated rpc_remove_pipedir() workqueue,which takes care about pipefs superblock locking kernel frees the pipefs sb of the current client and immediately alloctes a new pipefs sb, rpc_remove_pipedir functi of pipefs sb which is not the one it used to hold. As a result, the rpc_remove_pipedir would clean the released freec rpc_remove_pipedir should check whether the current pipefs sb is consistent with the original pipefs sb. This error ========== [250.497700] BUG: KASAN: sla +0x195/0x200 [250.498315] Read of size 4 at addr ffff88800a2ab804 by task kworker/0:18/106503 [250.500549] [250.501001] Call Trace: [250.502880] kasan_report+0xb6/0xf0 [250.503209] ? dget_parent+0x195/0x200 [250 [250.503897] ? __pfx_rpc_clntdir_depopulate+0x10/0x10 [250.504384] rpc_rmdir_depopulate+0x1b/0x90 [250. [250.505195] rpc_free_client_work+0xe4/0x230 [250.505598] process_one_work+0x8ee/0x13b0 ... [22.039056] kasan_save_stack+0x22/0x50 [22.039758] kasan_set_track+0x25/0x30 [22.040109] __kasan_slab_alloc+0x59/0x +0xf0/0x240 [22.040889] __d_alloc+0x31/0x8e0 [22.041207] d_alloc+0x44/0x1f0 [22.041514] __rpc_lookup_c rpc_mkdir_populate.constprop.0+0x5f/0x110 [22.042459] rpc_create_client_dir+0x34/0x150 [22.042874] rpc_se rpc_client_register+0x136/0x4e0 [22.043689] rpc_new_client+0x911/0x1020 [22.044057] rpc_create_xprt+0xcb +0x36b/0x6c0 ... [22.049524] Freed by task 0: [22.049803] kasan_save_stack+0x22/0x50 [22.050165] kasan_set kasan_save_free_info+0x2b/0x50 [22.050921] __kasan_slab_free+0x10e/0x1a0 [22.051306] kmem_cache_free+ +0x62c/0x1930 [22.051995] __do_softirq+0x165/0x52a [22.052347] [22.052503] Last potentially related work c +0x22/0x50 [22.053313] __kasan_record_aux_stack+0x8e/0xa0 [22.053739] __call_rcu_common.constprop.0+0 +0xb2/0x140 [22.054540] __dentry_kill+0x3be/0x540 [22.054900] shrink_dentry_list+0x199/0x510 [22.055293 [22.055703] do_one_tree+0x11/0x40 [22.056028] shrink_dcache_for_umount+0x61/0x140 [22.056461] generic_ kill_anon_super+0x3a/0x60 [22.057234] rpc_kill_sb+0x121/0x200
CVE-2023-52804	In the Linux kernel, the following vulnerability has been resolved: fs/jfs: Add validity check for db_maxag and db_ used as the index of the db_agfree array, but there is currently no validity check for db_maxag and db_agpref, whic bug reported by Syzbot: UBSAN: array-index-out-of-bounds in fs/jfs/jfs_dmap.c:639:20 index 7936 is out of range the values of db_maxag and db_agpref are valid indexes for the db_agfree array.
CVE-2023-52805	In the Linux kernel, the following vulnerability has been resolved: jfs: fix array-index-out-of-bounds in diAlloc Cu the iag while allocating new inodes to avoid fragmentation problem. Added the check which is required.
CVE-2023-52810	In the Linux kernel, the following vulnerability has been resolved: fs/jfs: Add check for negative db_l2nbperpage l and the minimum legal value should be 0, not negative. In the case of l2nbperpage being negative, an error will occ Syzbot reported this bug: UBSAN: shift-out-of-bounds in fs/jfs/jfs_dmap.c:799:12 shift exponent -16777216 is neg
CVE-2023-52813	In the Linux kernel, the following vulnerability has been resolved: crypto: pcrypt - Fix hungtask for PADATA_RE test_aead_vec_cfg as follows: INFO: task cryptomgr_test:391009 blocked for more than 120 seconds. "echo 0 > /p disables this message. Call trace: __switch_to+0x98/0xe0 __schedule+0x6c4/0xf40 schedule+0xd8/0x1b4 schedul +0x368/0x4e0 wait_for_completion+0x20/0x30 wait_for_completion+0x20/0x30 test_aead_vec_cfg+0xab4/0xd50 +0xd8/0x1e0 alg_test+0x634/0x890 cryptomgr_test+0x40/0x70 kthread+0x1e0/0x220 ret_from_fork+0x10/0x18 H blocked tasks For padata_do_parallel, when the return err is 0 or -EBUSY, it will call wait_for_completion(&wait- In normal case, aead_request_complete() will be called in pcrypt_aead_serial and the return err is 0 for padata_do_ PADATA_RESET, the return err is -EBUSY for padata_do_parallel, and it won't call aead_request_complete(). Th wait_for_completion(&wait->completion), which will cause hungtask. The problem comes as following: (padata_c EINVAL; | (padata_replace) | pinst->flags |= PADATA_RESET; err = -EBUSY | if (pinst->flags & PADATA_RES order to resolve the problem, we replace the return err -EBUSY with -EAGAIN, which means parallel_data is chan remove retry and just change the return err. v2: introduce padata_try_do_parallel() in pcrypt_aead_encrypt and pcr

CVE-2023-52817	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: Fix a null pointer access when the types of chips, such as VEGA20, reading the amdgpu_regs_smc file could result in an abnormal null pointer access Below are the steps to reproduce this issue and the corresponding exception log: 1. Navigate to the directory: /sys/k cat amdgpu_regs_smc 3. Exception Log:: [4005007.702554] BUG: kernel NULL pointer dereference, address: 000 supervisor instruction fetch in kernel mode [4005007.702567] #PF: error_code(0x0010) - not-present page [40050 Oops: 0010 [#1] SMP NOPTI [4005007.702581] CPU: 4 PID: 62563 Comm: cat Tainted: G OE 5.15.0-43-generic 0010:0x0 [4005007.702598] Code: Unable to access opcode bytes at RIP 0xffffffffffffffd6. [4005007.702600] RSP 00010206 [4005007.702605] RAX: 0000000000000000 RBX: 0000000000000000 RCX: ffffa82b46d27e68 [4005 0000000000000000 RDI: ffff9940656e0000 [4005007.702612] RBP: ffffa82b46d27dd8 R08: 0000000000000000 R10: 0000000000020000 R11: 0000000000000000 R12: 00007f5e06753000 [4005007.702618] R13: ffff9940656e 00007f5e06753000 [4005007.702622] FS: 00007f5e0755b740(0000) GS:ffff99479d300000(0000) knlGS:0000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [4005007.702629] CR2: ffffffffffffffd6 CR3: 0000000325 [4005007.702633] Call Trace: [4005007.702636] <TASK> [4005007.702640] amdgpu_debugfs_regs_smc_read+0 full_proxy_read+0x5c/0x80 [4005007.703011] vfs_read+0x9f/0x1a0 [4005007.703019] ksys_read+0x67/0xe0 [40 [4005007.703028] do_syscall_64+0x5c/0xc0 [4005007.703034] ? do_user_addr_fault+0x1e3/0x670 [4005007.703 +0x37/0xb0 [4005007.703047] ? irqentry_exit_to_user_mode+0x9/0x20 [4005007.703052] ? irqentry_exit+0x19/0 +0x89/0x160 [4005007.703062] ? asm_exc_page_fault+0x8/0x30 [4005007.703068] entry_SYSCALL_64_after_h RIP: 0033:0x7f5e07672992 [4005007.703079] Code: c0 e9 b2 fe ff ff 50 48 8d 3d fa b2 0c 00 e8 c5 1d 02 00 0f 1f 00 00 85 c0 75 10 0f 05 <48> 3d 00 f0 ff ff 77 56 c3 0f 1f 44 00 00 48 83 e c 28 48 89 54 24 [4005007.703083] RS 00000246 ORIG_RAX: 0000000000000000 [4005007.703088] RAX: ffffffffffffffda RBX: 0000000000020000 RC RDX: 0000000000020000 RSI: 00007f5e06753000 RDI: 0000000000000003 [4005007.703094] RBP: 00007f5e0 R09: 00007f5e06752010 [4005007.703096] R10: 0000000000000022 R11: 0000000000000246 R12: 0000000000 0000000000000003 R14: 0000000000020000 R15: 0000000000020000 [4005007.703105] </TASK> [4005007.70 nfnetlink algif_hash af_alg binfmt_misc nls_ iso8859_1 ipmi_ssif ast intel_rapl_msr intel_rapl_common drm_vram edac_mce_amd kvm_amd ccp mac_hid k10temp kvm acpi_ipmi ipmi_si rapl sch_fq_codel ipmi_devintf ipm i_msg pstore_blk efi_pstore ramoops pstore_zone reed_solo mon ip_tables x_tables autofs4 ib_uverbs ib_core amdgpu(O iommu_v 2 amd_sched(OE) amdkcl(OE) drm_kms_helper syscopyarea sysfillrect sysimgblt fb_sys_fops cec rc_co i2c_algo_bit xhci_pci_renesas dca [4005007.703184] CR2: 0000000000000000 [4005007.703188] ---[en ---trunca
CVE-2023-52832	In the Linux kernel, the following vulnerability has been resolved: wifi: mac80211: don't return unset power in ieee warning if ieee80211_get_tx_power() returns the INT_MIN value mac80211 internally uses for "unset power level net/wireless/nl80211.c:3816:5 -2147483648 * 100 cannot be represented in type 'int' CPU: 0 PID: 20433 Comm: in dump_stack+0x74/0x92 ubsan_epilogue+0x9/0x50 handle_overflow+0x8d/0xd0 __ubsan_handle_mul_overflow+0 [cfg80211] [...] cfg80211_register_wdev+0x78/0xb0 [cfg80211] cfg80211_netdev_notifier_call+0x200/0x620 [cfg [mac80211] ieee80211_register_hw+0xda5/0x1170 [mac80211] In this case, simply return an error instead, to indi
CVE-2023-52835	In the Linux kernel, the following vulnerability has been resolved: perf/core: Bail out early if the request AUX area large AUX area, e.g 4GB, it fails with: #perf record -C 0 -m ,4G -e arm_spe_0// -- sleep 1 failed to mmap with 12 (WARNING with __alloc_pages(): ------------[cut here]------------ WARNING: CPU: 44 PID: 17573 at mm/page_a trace: __alloc_pages+0x1ec/0x248 __kmalloc_large_node+0xc0/0x1f8 __kmalloc_node+0x134/0x1e8 rb_alloc_au mmap_region+0x308/0x8a8 do_mmap+0x3c0/0x528 vm_mmap_pgoff+0xf4/0x1b8 ksys_mmap_pgoff+0x18c/0x2 invoke_syscall+0x50/0x128 el0_svc_common.constprop.0+0x58/0x188 do_el0_svc+0x34/0x50 el0_svc+0x34/0x1 el0t_64_sync+0x1a4/0x1a8 'rb->aux_pages' allocated by kcalloc() is a pointer array which is used to maintains AU this array is physically contiguous (and virtually contiguous) with an order of 0..MAX_ORDER. If the size of poin MAX_ORDER, it reveals a WARNING. So bail out early with -ENOMEM if the request AUX area is out of bound arm_spe_0// -- sleep 1 failed to mmap with 12 (Cannot allocate memory)
CVE-2023-52836	In the Linux kernel, the following vulnerability has been resolved: locking/ww_mutex/test: Fix potential workqueu test-ww_mutex code, I was seeing odd behavior where sometimes it seemed flush_workqueue was returning before this would cause strange crashes as the mutexes would be freed while they were being used. Looking at the code, th thread that spawns the work allocates the "struct stress" structures that are passed to the workqueue threads. Then w they free the stress struct that was passed to them. Unfortunately the workqueue work_struct node is in the stress st before the work thread returns and while flush_workqueue is waiting. It seems like a better idea to have the control structures, so that we can be sure we don't corrupt the workqueue by freeing the structure prematurely. So this patc change I no longer see the early flush_workqueue returns.
CVE-2023-52838	In the Linux kernel, the following vulnerability has been resolved: fbdev: imsttfb: fix a resource leak in probe I've n that if init_imstt() fails we need to call iounmap(par->cmap_regs).
CVE-2023-52840	In the Linux kernel, the following vulnerability has been resolved: Input: synaptics-rmi4 - fix use after free in rmi_ rmi_release_function() which frees "fn" so the dereference on the next line "fn->num_of_irqs" is a use after free. M

CVE-2023-52842	In the Linux kernel, the following vulnerability has been resolved: virtio/vsock: Fix uninit-value in virtio_transport the following uninit-value access issue: ========== virtio_transport_recv_pkt+0x1dfb/0x26a0 net/vmw_vsock/virtio_transport_common.c:1421 virtio_transport_recv_ virtio_transport_common.c:1421 vsock_loopback_work+0x3bb/0x5a0 net/vmw_vsock/vsock_loopback.c:120 proc [inline] process_scheduled_works+0xff6/0x1e60 kernel/workqueue.c:2703 worker_thread+0xeca/0x14d0 kernel/w kernel/kthread.c:388 ret_from_fork+0x66/0x80 arch/x86/kernel/process.c:147 ret_from_fork_asm+0x11/0x20 arch stored to memory at: virtio_transport_space_update net/vmw_vsock/virtio_transport_common.c:1274 [inline] virti net/vmw_vsock/virtio_transport_common.c:1415 vsock_loopback_work+0x3bb/0x5a0 net/vmw_vsock/vsock_loo workqueue.c:2630 [inline] process_scheduled_works+0xff6/0x1e60 kernel/workqueue.c:2703 worker_thread+0xec +0x3cc/0x520 kernel/kthread.c:388 ret_from_fork+0x66/0x80 arch/x86/kernel/process.c:147 ret_from_fork_asm+0 Uninit was created at: slab_post_alloc_hook+0x105/0xad0 mm/slab.h:767 slab_alloc_node mm/slub.c:3478 [inline mm/slub.c:3523 kmalloc_reserve+0x13c/0x4a0 net/core/skbuff.c:559 __alloc_skb+0x2fd/0x770 net/core/skbuff.c: [inline] virtio_vsock_alloc_skb include/linux/virtio_vsock.h:66 [inline] virtio_transport_alloc_skb+0x90/0x11e0 n virtio_transport_reset_no_sock net/vmw_vsock/virtio_transport_common.c:957 [inline] virtio_transport_recv_pkt+ virtio_transport_common.c:1387 vsock_loopback_work+0x3bb/0x5a0 net/vmw_vsock/vsock_loopback.c:120 proc [inline] process_scheduled_works+0xff6/0x1e60 kernel/workqueue.c:2703 worker_thread+0xeca/0x14d0 kernel/w kernel/kthread.c:388 ret_from_fork+0x66/0x80 arch/x86/kernel/process.c:147 ret_from_fork_asm+0x11/0x20 arch 10664 Comm: kworker/1:5 Not tainted 6.6.0-rc3-00146-g9f3ebbef746f #3 Hardware name: QEMU Standard PC (i 04/01/2014 Workqueue: vsock-loopback vsock_loopback_work ========== simple reproducer can cause the issue described above: int main(void) { int sock; struct sockaddr_vm addr = { .svm VMADDR_CID_ANY, .svm_port = 1234, }; sock = socket(AF_VSOCK, SOCK_STREAM, 0); connect(sock, (str return 0; } This issue occurs because the `buf_alloc` and `fwd_cnt` fields of the `struct virtio_vsock_hdr` are not in `virtio_transport_init_hdr()`. This patch resolves the issue by initializing these fields during allocation.
CVE-2023-52843	In the Linux kernel, the following vulnerability has been resolved: llc: verify mac len before reading mac header Ll without verifying that the skb has an Ethernet header. Syzbot was able to enter llc_rcv on a tun device. Tun can ins configurable skb->protocol (passing a tun_pi header when not configuring IFF_NO_PI). BUG: KMSAN: uninit-va llc_station.c:81 [inline] BUG: KMSAN: uninit-value in llc_station_rcv+0x6fb/0x1290 net/llc/llc_station.c:111 llc_ [inline] llc_station_rcv+0x6fb/0x1290 net/llc/llc_station.c:111 llc_rcv+0xc5d/0x14a0 net/llc/llc_input.c:218 __neti [inline] __netif_receive_skb+0x1a6/0x5a0 net/core/dev.c:5637 netif_receive_skb_internal net/core/dev.c:5723 [inl core/dev.c:5782 tun_rx_batched+0x3ee/0x980 drivers/net/tun.c:1555 tun_get_user+0x54c5/0x69c0 drivers/net/tun eth_hdr(skb) calls under net/llc. There are further uses in include/net/llc_pdu.h. All these are protected by a test skb not protect against this tun scenario. But the mac_len test added in this patch in llc_fixup_skb will indirectly protec any other LLC code. It is tempting to just add a blanket mac_len check in llc_rcv, but not sure whether that could b Ethernet header. 802.2 LLC may be used on top of non-802.3 protocols in principle. The below referenced commit least one of the three eth_hdr uses goes back to before the start of git history. But the one that syzbot exercises is in enough (2008), that effectively all stable kernels should receive this.
CVE-2023-52845	In the Linux kernel, the following vulnerability has been resolved: tipc: Change nla_policy for bearer-related name the following uninit-value access issue [1]: ========== lib/string.c:418 [inline] BUG: KMSAN: uninit-value in strstr+0xb8/0x2f0 lib/string.c:756 strlen lib/string.c:418 [in tipc_nl_node_reset_link_stats+0x3ea/0xb50 net/tipc/node.c:2595 genl_family_rcv_msg_doit net/netlink/genetlink. netlink/genetlink.c:1051 [inline] genl_rcv_msg+0x11ec/0x1290 net/netlink/genetlink.c:1066 netlink_rcv_skb+0x3 genl_rcv+0x40/0x60 net/netlink/genetlink.c:1075 netlink_unicast_kernel net/netlink/af_netlink.c:1342 [inline] netl af_netlink.c:1368 netlink_sendmsg+0x1238/0x13d0 net/netlink/af_netlink.c:1910 sock_sendmsg_nosec net/socket. socket.c:753 [inline] ____sys_sendmsg+0x9c2/0xd60 net/socket.c:2541 ___sys_sendmsg+0x28d/0x3c0 net/socket [inline] __do_sys_sendmsg net/socket.c:2633 [inline] __se_sys_sendmsg net/socket.c:2631 [inline] __x64_sys_sen do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_syscall_64+0x41/0xc0 arch/x86/entry/common.c:80 entry Uninit was created at: slab_post_alloc_hook+0x12f/0xb70 mm/slab.h:767 slab_alloc_node mm/slub.c:3478 [inline mm/slub.c:3523 kmalloc_reserve+0x13d/0x4a0 net/core/skbuff.c:559 __alloc_skb+0x318/0x740 net/core/skbuff.c: [inline] netlink_alloc_large_skb net/netlink/af_netlink.c:1214 [inline] netlink_sendmsg+0xb34/0x13d0 net/netlink/ net/socket.c:730 [inline] sock_sendmsg net/socket.c:753 [inline] ____sys_sendmsg+0x9c2/0xd60 net/socket.c:254 socket.c:2595 __sys_sendmsg net/socket.c:2624 [inline] __do_sys_sendmsg net/socket.c:2633 [inline] __se_sys_se __x64_sys_sendmsg+0x307/0x490 net/socket.c:2631 do_syscall_x64 arch/x86/entry/common.c:50 [inline] do_sys common.c:80 entry_SYSCALL_64_after_hwframe+0x63/0xcd TIPC bearer-related names including link names m which is not null-terminated is passed through netlink, strstr() and similar functions can cause buffer overrun. This the nla_policy for bearer-related names from NLA_STRING to NLA_NUL_STRING. This resolves the issue by en accepted as bearer-related names. syzbot reported similar uninit-value issue related to bearer names [2]. The root ca bearer name was passed. This patch also resolved this issue.
CVE-2023-52846	In the Linux kernel, the following vulnerability has been resolved: hsr: Prevent use after free in prp_create_tagged_ In that situation, it frees the skb and returns NULL. Meanwhile on the success path, it returns the original skb. So it returned value.
CVE-2023-52847	In the Linux kernel, the following vulnerability has been resolved: media: bttv: fix use after free error due to btv->t condition between timer function bttv_irq_timeout and bttv_remove. The timer is setup in probe and there is no tim When it hit kfree btv, the function might still be invoked, which will cause use after free bug. This bug is found by it by adding del_timer_sync invoking to the remove function. cpu0 cpu1 bttv_probe ->timer_setup ->bttv_set_dma >bttv_irq_timeout ->USE btv

CVE-2023-52849	In the Linux kernel, the following vulnerability has been resolved: cxl/mem: Fix shutdown order Ira reports that rer the following trace: BUG: kernel NULL pointer dereference, address: 0000000000000044 [..] RIP: 0010:cxl_regio [..] Call Trace: <TASK> cxl_region_detach+0xe8/0x210 [cxl_core] cxl_decoder_kill_region+0x27/0x40 [cxl_core devres_release_all+0xb8/0x110 device_unbind_cleanup+0xe/0x70 device_release_driver_internal+0x1d2/0x210 b +0x155/0x3e0 device_unregister+0x13/0x60 devm_release_action+0x4d/0x90 ? __pfx_unregister_port+0x10/0x10 [cxl_core] devres_release_all+0xb8/0x110 device_unbind_cleanup+0xe/0x70 device_release_driver_internal+0x1c device_del+0x155/0x3e0 ? lock_release+0x142/0x290 cdev_device_del+0x15/0x50 cxl_memdev_unregister+0x54 clearing out the cxl_memdev's driver context (@cxlds) before the subsystem is done with it. This is ultimately due member, being torn down and expecting to be able to de-reference @cxlds, like here: static int cxl_region_decode_ (cxlds->rcd) goto endpoint_reset; ... Fix it by keeping the driver context valid until memdev-device unregistration, dependencies, unwinds.
CVE-2023-52853	In the Linux kernel, the following vulnerability has been resolved: hid: cp2112: Fix duplicate workqueue initializat INIT_DELAYED_WORK within cp2112_gpio_irq_startup, resulting in duplicate initilizations of the workqueue o initial request. This resulted in a warning in set_work_data in workqueue.c, as well as a rare NULL dereference wi Initialize the workqueue within _probe instead.
CVE-2023-52855	In the Linux kernel, the following vulnerability has been resolved: usb: dwc2: fix possible NULL pointer dereferen In _dwc2_hcd_urb_enqueue(), "urb->hcpriv = NULL" is executed without holding the lock "hsotg->lock". In _dwc spin_lock_irqsave(&hsotg->lock, flags); ... if (!urb->hcpriv) { dev_dbg(hsotg->dev, "## urb->hcpriv is NULL ##\r dwc2_hcd_urb_dequeue(hsotg, urb->hcpriv); // Use urb->hcpriv ... out: spin_unlock_irqrestore(&hsotg->lock, flag _dwc2_hcd_urb_dequeue() are concurrently executed, the NULL check of "urb->hcpriv" can be executed before "u is NULL, it can be used in the function call to dwc2_hcd_urb_dequeue(), which can cause a NULL pointer derefer experimental static analysis tool developed by myself. This tool analyzes the locking APIs to extract function pairs analyzes the instructions in the paired functions to identify possible concurrency bugs including data races and ator is reported, when my tool analyzes the source code of Linux 6.5. To fix this possible bug, "urb->hcpriv = NULL" s "hsotg->lock". After using this patch, my tool never reports the possible bug, with the kernelconfiguration allyesco hardware, I cannot test the patch in runtime testing, and just verify it according to the code logic.
CVE-2023-52858	In the Linux kernel, the following vulnerability has been resolved: clk: mediatek: clk-mt7629: Add check for mtk_ value of mtk_alloc_clk_data() in order to avoid NULL pointer dereference.
CVE-2023-52864	In the Linux kernel, the following vulnerability has been resolved: platform/x86: wmi: Fix opening of char device S pass miscdevice pointer via file private data"), the miscdevice stores a pointer to itself inside filp->private_data, wh NULL when wmi_char_open() is called. This might cause memory corruption should wmi_char_open() be unable happen when the associated WMI device is deleted in wmi_free_devices(). Fix the problem by using the miscdevic associated with a char device using container_of(). This also avoids wmi_char_open() picking a wrong WMI devic original driver.
CVE-2023-52865	In the Linux kernel, the following vulnerability has been resolved: clk: mediatek: clk-mt6797: Add check for mtk_ value of mtk_alloc_clk_data() in order to avoid NULL pointer dereference.
CVE-2023-52866	In the Linux kernel, the following vulnerability has been resolved: HID: uclogic: Fix user-memory-access bug in uc When CONFIG_HID_UCLOGIC=y and CONFIG_KUNIT_ALL_TESTS=y, launch kernel and then the below use In hid_test_uclogic_params_cleanup_event_hooks(),it call uclogic_params_ugee_v2_init_event_hooks() with the f uclogic_params_ugee_v2_has_battery(), the hid_get_drvdata() will access hdev->dev with hdev=NULL, which wi So add a fake_device with quirks member and call hid_set_drvdata() to assign hdev->dev->driver_data which avoi >quirks in uclogic_params_ugee_v2_has_battery(). After applying this patch, the below user-memory-access bug r fault, probably for non-canonical address 0xdffffc0000000329: 0000 [#1] PREEMPT SMP KASAN KASAN: prob [0x0000000000001948-0x000000000000194f] CPU: 5 PID: 2189 Comm: kunit_try_catch Tainted: G B W N 6.6.0 Standard PC (i440FX + PIIX, 1996), BIOS 1.15.0-1 04/01/2014 RIP: 0010:uclogic_params_ugee_v2_init_event_h 48 8b 14 25 28 00 00 00 48 89 54 24 60 31 d2 48 89 fa c7 44 24 30 00 00 00 00 48 c7 44 24 28 02 f8 02 01 48 c1 e 00 48 8b 9d 48 19 00 00 48 b8 00 00 00 00 RSP: 0000:ffff88810679fc88 EFLAGS: 00010202 RAX: dfffc000000000 0000000000000000 RDX: 0000000000000329 RSI: ffff88810679fd88 RDI: 0000000000001948 RBP: 000000000 R09: ffffed1020f639f0 R10: ffff888107b1cf87 R11: 0000000000000400 R12: 1ffff11020cf3f92 R13: ffff888106790 R15: ffff8881030bb080 FS: 0000000000000000(0000) GS:ffff888119e80000(0000) knlGS:0000000000000000 CS 0000000080050033 CR2: 0000000000000000 CR3: 0000000005286001 CR4: 0000000000770ee0 DR0: ffffffff8f ffffffff8fdd6cf6 DR3: ffffffff8fdd6cf7 DR6: 00000000fffe0ff0 DR7: 0000000000000600 PKRU: 55555554 Call Tr exc_general_protection+0x144/0x220 ? asm_exc_general_protection+0x22/0x30 ? uclogic_params_ugee_v2_init_ +0x69/0x550 ? uclogic_parse_ugee_v2_desc_gen_params+0x70/0x70 ? load_balance+0x2950/0x2950 ? rcu_trc_c hid_test_uclogic_params_cleanup_event_hooks+0x9e/0x1a0 ? uclogic_params_ugee_v2_init_event_hooks+0x600 migrate_enable+0x260/0x260 ? __kthread_parkme+0x83/0x150 ? kunit_try_run_case_cleanup+0xe0/0xe0 kunit_g kunit_try_catch_throw+0x80/0x80 kthread+0x2b5/0x380 ? kthread_complete_and_exit+0x20/0x20 ret_from_fork +0x20/0x20 ret_from_fork_asm+0x11/0x20 </TASK> Modules linked in: Dumping ftrace buffer: (ftrace buffer en RIP: 0010:uclogic_params_ugee_v2_init_event_hooks+0x87/0x600 Code: f3 f3 65 48 8b 14 25 28 00 00 00 48 89 00 00 00 00 48 c7 44 24 28 02 f8 02 01 48 c1 ea 03 <80> 3c 02 00 0f 85 2c 04 00 00 48 8b 9d 48 19 00 00 48 b8 0 EFLAGS: 00010202 RAX: dfffc0000000000 RBX: 0000000000000004 RCX: 0000000000000000 RDX: 000000 RDI: 0000000000001948 RBP: 0000000000000000 R08: 0000000000000000 R09: ffffed1020f639f0 R10: ffff888 R12: 1ffff11020cf3f92 R13: ffff88810679fd88 R14: ffff888100b97b08 R15: ffff8881030bb080 FS: 000000000000 knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 0000000000000000 CR3 0000000000770ee0 DR0: ffffffff8fdd6cf4 DR1: ---truncated---

CVE-2023-52868	In the Linux kernel, the following vulnerability has been resolved: thermal: core: prevent potential string overflow it's a number between zero and INT_MAX. If it's too high then these sprintf()s will overflow.
CVE-2023-52871	In the Linux kernel, the following vulnerability has been resolved: soc: qcom: llcc: Handle a second device without device. But if there were a second, even a failed probe call would modify the global drv_data pointer. So check if d
CVE-2023-52873	In the Linux kernel, the following vulnerability has been resolved: clk: mediatek: clk-mt6779: Add check for mtk_ value of mtk_alloc_clk_data() in order to avoid NULL pointer dereference.
CVE-2023-52875	In the Linux kernel, the following vulnerability has been resolved: clk: mediatek: clk-mt2701: Add check for mtk_ value of mtk_alloc_clk_data() in order to avoid NULL pointer dereference.
CVE-2023-52876	In the Linux kernel, the following vulnerability has been resolved: clk: mediatek: clk-mt7629-eth: Add check for m return value of mtk_alloc_clk_data() in order to avoid NULL pointer dereference.
CVE-2023-52881	In the Linux kernel, the following vulnerability has been resolved: tcp: do not accept ACK of bytes we never sent T ideas from Yepeng Pan and Christian Rossow. ACK seq validation is currently following RFC 5961 5.2 guidelines only if it is in the range of ((SND.UNA - MAX.SND.WND) <= SEG.ACK <= SND.NXT). All incoming segments above condition MUST be discarded and an ACK sent back. It needs to be noted that RFC 793 on page 72 (fifth ch (SEG.ACK < SND.UNA), it can be ignored. If the ACK acknowledges something not yet sent (SEG.ACK > SND. and return". The "ignored" above implies that the processing of the incoming data segment continues, which means This mitigation makes the ACK check more stringent since any ACK < SND.UNA wouldn't be accepted, instead o - MAX.SND.WND) <= SEG.ACK <= SND.NXT) get through. This can be refined for new (and possibly spoofed) were never sent. This greatly improves TCP security at a little cost. I added a Fixes: tag to make sure this patch wil was adhering to the RFC. tp->bytes_acked was added in linux-4.2 Following packetdrill test (courtesy of Yepeng P SOCK_STREAM, IPPROTO_TCP) = 3 +0 setsockopt(3, SOL_SOCKET, SO_REUSEADDR, [1], 4) = 0 +0 bind(---------- Handshake ---------- // // when window scale is set to 14 the window size can be extended to / would accept an ACK packet // with ack number in (Server_ISN+1-1073725440. Server_ISN+1) // ,though this ack sent by the server. +0 < S 0:0(0) win 65535 <mss 1400,nop,wscale 14> +0 > S. 0:0(0) ack 1 <...> +0 < . 1:1(0) ack the established connection, we send an ACK packet, // the ack packet uses ack number 1 - 1073725300 + 2^32, // w we used 1073725300 instead of 1073725440 to avoid possible // edge cases. // 1 - 1073725300 + 2^32 = 32212419 packet. +0 < . 1:1001(1000) ack 3221241997 win 65535 // After the kernel fix the following will be replaced by a c would be dropped. +0 > . 1:1(0) ack 1001
CVE-2023-52887	In the Linux kernel, the following vulnerability has been resolved: net: can: j1939: enhanced error handling for tigh xtp_rx_rts_session_new This patch enhances error handling in scenarios with RTS (Request to Send) messages arri informative WARN_ON_ONCE backtraces with a new error handling method. This provides clearer error message of problematic sessions. Previously, sessions were only released at the end of j1939_xtp_rx_rts(). Potentially this c like: testj1939 -r vcan0:0x80 & while true; do # send first RTS cansend vcan0 18EC8090#1014000303002301; # s 18EC8090#1014000303002301; # send abort cansend vcan0 18EC8090#ff00000000002301; done
CVE-2023-52905	In the Linux kernel, the following vulnerability has been resolved: octeontx2-pf: Fix resource leakage in VF driver to support the Ntuple feature and hash tables for the tc feature are not getting freed in driver unbind. This patch fixe
CVE-2023-52909	In the Linux kernel, the following vulnerability has been resolved: nfsd: fix handling of cached open files in nfsd4_ ("NFSD: Instantiate a struct file when creating a regular NFSv4 file") added the ability to cache an open fd over a c with the way this currently works: It's racy, as a newly-created nfsd_file can end up with its PENDING bit cleared is still zeroed out. Other tasks can find it in this state and they expect to see a valid nf_file, and can oops if nf_file i end up creating a new nfsd_file if one is already in the hash. If an extant entry is in the hash with a valid nf_file, nf with the value of op_file and the old nf_file will leak. Fix both issues by making a new nfsd_file_acquirei_opened one is present when this is called, we'll take a new reference to it instead of trying to open the file. If the nfsd_file a optional file and pass the nfsd_file back as-is. Also rework the tracepoints a bit to allow for an "opened" variant an the case where we already have a cached open file.
CVE-2023-5517	A flaw in query-handling code can cause `named` to exit prematurely with an assertion failure when: - `nxdomain- the resolver receives a PTR query for an RFC 1918 address that would normally result in an authoritative NXDOM versions 9.12.0 through 9.16.45, 9.18.0 through 9.18.21, 9.19.0 through 9.19.19, 9.16.8-S1 through 9.16.45-S1, and
CVE-2023-6121	An out-of-bounds read vulnerability was found in the NVMe-oF/TCP subsystem in the Linux kernel. This issue ma TCP packet, triggering a heap-based buffer overflow that results in kmalloc data being printed and potentially leak
CVE-2023-6228	An issue was found in the tiffcp utility distributed by the libtiff package where a crafted TIFF file on processing ma to an application crash.
CVE-2023-6270	A flaw was found in the ATA over Ethernet (AoE) driver in the Linux kernel. The aoecmd_cfg_pkts() function imp net_device`, and a use-after-free can be triggered by racing between the free on the struct and the access through th denial of service condition or potential code execution.
CVE-2023-6356	A flaw was found in the Linux kernel's NVMe driver. This issue may allow an unauthenticated malicious actor to s using NVMe over TCP, leading the NVMe driver to a NULL pointer dereference in the NVMe driver and causing

CVE-2023-6516	To keep its cache database efficient, `named` running as a recursive resolver occasionally attempts to clean up the c some that are asynchronous: a small chunk of memory pointing to the cache element that can be cleaned up is first processing. It was discovered that if the resolver is continuously processing query patterns triggering this type of ca be able to handle the cleanup events in a timely manner. This in turn enables the list of queued cleanup events to gr configured `max-cache-size` limit to be significantly exceeded. This issue affects BIND 9 versions 9.16.0 through 9
CVE-2023-6535	A flaw was found in the Linux kernel's NVMe driver. This issue may allow an unauthenticated malicious actor to s using NVMe over TCP, leading the NVMe driver to a NULL pointer dereference in the NVMe driver, causing kern
CVE-2023-6536	A flaw was found in the Linux kernel's NVMe driver. This issue may allow an unauthenticated malicious actor to s using NVMe over TCP, leading the NVMe driver to a NULL pointer dereference in the NVMe driver, causing kern
CVE-2023-6597	An issue was found in the CPython `tempfile.TemporaryDirectory` class affecting versions 3.12.1, 3.11.7, 3.10.13, tempfile.TemporaryDirectory class would dereference symlinks during cleanup of permissions-related errors. This programs are potentially able to modify permissions of files referenced by symlinks in some circumstances.
CVE-2023-6879	Increasing the resolution of video frames, while performing a multi-threaded encode, can result in a heap overflow
CVE-2023-6915	A Null pointer dereference problem was found in ida_free in lib/idr.c in the Linux Kernel. This issue may allow an service problem due to a missing check at a function return.
CVE-2023-7192	A memory leak problem was found in ctnetlink_create_conntrack in net/netfilter/nf_conntrack_netlink.c in the Linu attacker with CAP_NET_ADMIN privileges to cause a denial of service (DoS) attack due to a refcount overflow.
CVE-2023-7207	Debian's cpio contains a path traversal vulnerability. This issue was introduced by reverting CVE-2015-1197 patch absolute-filenames. Upstream has since provided a proper fix to --no-absolute-filenames.
CVE-2024-0340	A vulnerability was found in vhost_new_msg in drivers/vhost/vhost.c in the Linux kernel, which does not properly between virtual guests and the host operating system in the vhost/vhost.c:vhost_new_msg() function. This issue ca kernel memory contents when reading from the /dev/vhost-net device file.
CVE-2024-0443	A flaw was found in the blkgs destruction path in block/blk-cgroup.c in the Linux kernel, leading to a cgroup blkio is being destroyed, cgroup_rstat_flush() is only called at css_release_work_fn(), which is called when the blkcg ref dependency will prevent blkcg and some blkgs from being freed after they are made offline. This issue may allow a instability, such as an out of memory error.
CVE-2024-0444	GStreamer AV1 Video Parsing Stack-based Buffer Overflow Remote Code Execution Vulnerability. This vulnerab arbitrary code on affected installations of GStreamer. Interaction with this library is required to exploit this vulnera on the implementation. The specific flaw exists within the parsing of tile list data within AV1-encoded video files. validation of the length of user-supplied data prior to copying it to a fixed-length stack-based buffer. An attacker ca in the context of the current process. Was ZDI-CAN-22873.
CVE-2024-0450	An issue was found in the CPython `zipfile` module affecting versions 3.12.1, 3.11.7, 3.10.13, 3.9.18, and 3.8.18 a ,Äúquoted-overlap,Äù zip-bombs which exploit the zip format to create a zip-bomb with a high compression ratio. zipfile module reject zip archives which overlap entries in the archive.
CVE-2024-0565	An out-of-bounds memory read flaw was found in receive_encrypted_standard in fs/smb/client/smb2ops.c in the SM Kernel. This issue occurs due to integer underflow on the memcpy length, leading to a denial of service.
CVE-2024-0567	A vulnerability was found in GnuTLS, where a cockpit (which uses gnuTLS) rejects a certificate chain with distrib certificate chain with cockpit-certificate-ensure. This flaw allows an unauthenticated, remote client or attacker to in
CVE-2024-0646	An out-of-bounds memory write flaw was found in the Linux kernel,Äôs Transport Layer Security functionality in socket as the destination. This flaw allows a local user to crash or potentially escalate their privileges on the system
CVE-2024-0775	A use-after-free flaw was found in the __ext4_remount in fs/ext4/super.c in ext4 in the Linux kernel. This flaw allo problem while freeing the old quota file names before a potential failure, leading to a use-after-free.
CVE-2024-0841	A null pointer dereference flaw was found in the hugetlbfs_fill_super function in the Linux kernel hugetlbfs (Huge allow a local user to crash the system or potentially escalate their privileges on the system.
CVE-2024-0985	Late privilege drop in REFRESH MATERIALIZED VIEW CONCURRENTLY in PostgreSQL allows an object cr the command issuer. The command intends to run SQL functions as the owner of the materialized view, enabling s The victim is a superuser or member of one of the attacker's roles. The attack requires luring the victim into runnin CONCURRENTLY on the attacker's materialized view. Versions before PostgreSQL 16.2, 15.6, 14.11, 13.14, and
CVE-2024-1013	An out-of-bounds stack write flaw was found in unixODBC on 64-bit architectures where the caller has 4 bytes and unnoticed on little-endian architectures, while big-endian architectures can be broken.
CVE-2024-1086	A use-after-free vulnerability in the Linux kernel's netfilter: nf_tables component can be exploited to achieve local function allows positive values as drop error within the hook verdict, and hence the nf_hook_slow() function can ca NF_DROP is issued with a drop error which resembles NF_ACCEPT. We recommend upgrading past commit f342

CVE-2024-11053	When asked to both use a `.netrc` file for credentials and to follow HTTP redirects, curl could leak the password us under certain circumstances. This flaw only manifests itself if the netrc file has an entry that matches the redirect ta the password or omits both login and password.
CVE-2024-2201	A flaw was found in some Intel CPUs where mitigations for the Spectre V2/BHI vulnerability were incomplete. Th arbitrary memory, compromising system integrity and exposing sensitive information.
CVE-2024-22190	GitPython is a python library used to interact with Git repositories. There is an incomplete fix for CVE-2023-4059(search path if it uses a shell to run `git`, as well as when it runs `bash.exe` to interpret hooks. If either of those featu `git.exe` or `bash.exe` may be run from an untrusted repository. This issue has been patched in version 3.1.41.
CVE-2024-22195	Jinja is an extensible templating engine. Special placeholders in the template allow writing code similar to Python s attributes into the rendered HTML template, potentially leading to Cross-Site Scripting (XSS). The Jinja `xmlattr` f attribute keys and values, bypassing the auto escaping mechanism and potentially leading to XSS. It may also be po they are blacklist-based.
CVE-2024-23254	The issue was addressed with improved UI handling. This issue is fixed in tvOS 17.4, macOS Sonoma 14.4, vision 10.4, Safari 17.4. A malicious website may exfiltrate audio data cross-origin.
CVE-2024-23263	A logic issue was addressed with improved validation. This issue is fixed in tvOS 17.4, macOS Sonoma 14.4, visio watchOS 10.4, iOS 16.7.6 and iPadOS 16.7.6, Safari 17.4. Processing maliciously crafted web content may preven enforced.
CVE-2024-23280	An injection issue was addressed with improved validation. This issue is fixed in Safari 17.4, macOS Sonoma 14.4 tvOS 17.4. A maliciously crafted webpage may be able to fingerprint the user.
CVE-2024-23284	A logic issue was addressed with improved state management. This issue is fixed in tvOS 17.4, macOS Sonoma 14 17.4, watchOS 10.4, iOS 16.7.6 and iPadOS 16.7.6, Safari 17.4. Processing maliciously crafted web content may p enforced.
CVE-2024-23307	Integer Overflow or Wraparound vulnerability in Linux Linux kernel kernel on Linux, x86, ARM (md, raid, raid5 r
CVE-2024-2379	libcurl skips the certificate verification for a QUIC connection under certain conditions, when built to use wolfSSL curve, the error path accidentally skips the verification and returns OK, thus ignoring any certificate problems.
CVE-2024-23849	In rds_recv_track_latency in net/rds/af_rds.c in the Linux kernel through 6.7.1, there is an off-by-one error for an R comparison, resulting in out-of-bounds access.
CVE-2024-23851	copy_params in drivers/md/dm-ioctl.c in the Linux kernel through 6.7.1 can attempt to allocate more than INT_MA param_kernel->data_size check. This is related to ctl_ioctl.
CVE-2024-2398	When an application tells libcurl it wants to allow HTTP/2 server push, and the amount of received headers for the (1000), libcurl aborts the server push. When aborting, libcurl inadvertently does not free all the previously allocate Further, this error condition fails silently and is therefore not easily detected by an application.
CVE-2024-2398	When an application tells libcurl it wants to allow HTTP/2 server push, and the amount of received headers for the (1000), libcurl aborts the server push. When aborting, libcurl inadvertently does not free all the previously allocate Further, this error condition fails silently and is therefore not easily detected by an application.
CVE-2024-2398	When an application tells libcurl it wants to allow HTTP/2 server push, and the amount of received headers for the (1000), libcurl aborts the server push. When aborting, libcurl inadvertently does not free all the previously allocate Further, this error condition fails silently and is therefore not easily detected by an application.
CVE-2024-2398	When an application tells libcurl it wants to allow HTTP/2 server push, and the amount of received headers for the (1000), libcurl aborts the server push. When aborting, libcurl inadvertently does not free all the previously allocate Further, this error condition fails silently and is therefore not easily detected by an application.
CVE-2024-24577	libgit2 is a portable C implementation of the Git core methods provided as a linkable library with a solid API, allow application. Using well-crafted inputs to `git_index_add` can cause heap corruption that could be leveraged for arb `has_dir_name` function in `src/libgit2/index.c`, which frees an entry that should not be freed. The freed entry is la actor-controlled data leading to controlled heap corruption. Depending on the application that uses libgit2, this cou has been patched in version 1.6.5 and 1.7.2.
CVE-2024-24762	`python-multipart` is a streaming multipart parser for Python. When using form data, `python-multipart` uses a Reg Type` header, including options. An attacker could send a custom-made `Content-Type` option that is very difficul resources and stalling indefinitely (minutes or more) while holding the main event loop. This means that process ca regular expression denial of service. This vulnerability has been patched in version 0.0.7.

CVE-2024-24806	libuv is a multi-platform support library with a focus on asynchronous I/O. The `uv_getaddrinfo` function in `src/u counterpart `src/win/getaddrinfo.c`), truncates hostnames to 256 characters before calling `getaddrinfo`. This behav `0x00007f000001`, which are considered valid by `getaddrinfo` and could allow an attacker to craft payloads that r developer checks. The vulnerability arises due to how the `hostname_ascii` variable (with a length of 256 bytes) is in `uv__idna_toascii`. When the hostname exceeds 256 characters, it gets truncated without a terminating null byte internal APIs or for websites (similar to MySpace) that allows users to have `username.example.com` pages. Intern pages can be exposed to SSRF attacks if a malicious user chooses a long vulnerable username. This issue has been advised to upgrade. There are no known workarounds for this vulnerability.
CVE-2024-24857	A race condition was found in the Linux kernel's net/bluetooth device driver in conn_info_{min,max}_age_set() fu issue, possibly leading to bluetooth connection abnormality or denial of service.
CVE-2024-24860	A race condition was found in the Linux kernel's bluetooth device driver in {min,max}_key_size_set() function. Th issue, possibly leading to a kernel panic or denial of service issue.
CVE-2024-24861	A race condition was found in the Linux kernel's media/xc4000 device driver in xc4000 xc4000_get_frequency() fu overflow issue, possibly leading to malfunction or denial of service issue.
CVE-2024-25062	An issue was discovered in libxml2 before 2.11.7 and 2.12.x before 2.12.5. When using the XML Reader interface enabled, processing crafted XML documents can lead to an xmlValidatePopElement use-after-free.
CVE-2024-25062	An issue was discovered in libxml2 before 2.11.7 and 2.12.x before 2.12.5. When using the XML Reader interface enabled, processing crafted XML documents can lead to an xmlValidatePopElement use-after-free.
CVE-2024-25269	libheif <= 1.17.6 contains a memory leak in the function JpegEncoder::Encode. This flaw allows an attacker to cau
CVE-2024-25739	create_empty_lvol in drivers/mtd/ubi/vtbl.c in the Linux kernel through 6.7.4 can attempt to allocate zero bytes, an >leb_size.
CVE-2024-26581	In the Linux kernel, the following vulnerability has been resolved: netfilter: nft_set_rbtree: skip end interval eleme collect an end interval element that has been just added in this transactions, skip end interval elements that are not
CVE-2024-26583	In the Linux kernel, the following vulnerability has been resolved: tls: fix race between async notify and socket clo recvmsg/sendmsg) may exit as soon as the async crypto handler calls complete() so any code past that point risks t locking and extra flags altogether. Have the main thread hold an extra reference, this way we can depend solely on Don't futz with reiniting the completion, either, we are now tightly controlling when completion fires.
CVE-2024-26584	In the Linux kernel, the following vulnerability has been resolved: net: tls: handle backlogging of crypto requests S CRYPTO_TFM_REQ_MAY_BACKLOG flag on our requests to the crypto API, crypto_aead_{encrypt,decrypt} instead of -EINPROGRESS in valid situations. For example, when the cryptd queue for AESNI is full (easy to trig cryptd.cryptd_max_cpu_qlen), requests will be enqueued to the backlog but still processed. In that case, the async first with err == -EINPROGRESS, which it seems we can just ignore, then with err == 0. Compared to Sabrina's or tls_*crypt_async_wait() helpers and converts the EBUSY to EINPROGRESS to avoid having to modify all the err
CVE-2024-26585	In the Linux kernel, the following vulnerability has been resolved: tls: fix race between tx work scheduling and soc submitting thread (recvmsg/sendmsg) may exit as soon as the async crypto handler calls complete(). Reorder sched This seems more logical in the first place, as it's the inverse order of what the submitting thread will do.
CVE-2024-26586	In the Linux kernel, the following vulnerability has been resolved: mlxsw: spectrum_acl_tcam: Fix stack corruptio device, the corresponding local port gets bound to an ACL group in the device. The group contains a list of ACLs. TCAM region where the filters are stored. During forwarding, the ACLs are sequentially evaluated until a match is regions is when they are added with decreasing priorities and in an alternating order so that two consecutive filters their key usage. In Spectrum-2 and newer ASICs the firmware started to report that the maximum number of ACLs of the register that configures ACL groups (PAGT) was not updated to account for that. It is therefore possible to h more than 16 ACLs in a group are required. Fix by limiting the maximum ACL group size to the minimum betwee ACLs that fit in the PAGT register. Add a test case to make sure the machine does not crash when this condition is protector: Kernel stack is corrupted in: mlxsw_sp_acl_tcam_group_update+0x116/0x120 [...] dump_stack_lvl+0x3 +0x15/0x20 mlxsw_sp_acl_tcam_group_update+0x116/0x120 mlxsw_sp_acl_tcam_group_region_attach+0x69/0x +0x492/0xa20 mlxsw_sp_acl_tcam_ventry_add+0x25/0xe0 mlxsw_sp_acl_rule_add+0x47/0x240 mlxsw_sp_flow +0xdc/0x1c0 fl_hw_replace_filter+0x146/0x1f0 fl_change+0xc17/0x1360 tc_new_tfilter+0x472/0xb90 rtnetlink_r +0x58/0x100 netlink_unicast+0x244/0x390 netlink_sendmsg+0x1e4/0x440 ____sys_sendmsg+0x164/0x260 ___s +0x7a/0xc0 do_syscall_64+0x40/0xe0 entry_SYSCALL_64_after_hwframe+0x63/0x6b
CVE-2024-26593	In the Linux kernel, the following vulnerability has been resolved: i2c: i801: Fix block process call transactions Ac reset the block buffer index twice for block process call transactions: once before writing the outgoing data to the b incoming data from the buffer. The driver is currently missing the second reset, causing the wrong portion of the bl

CVE-2024-26597	In the Linux kernel, the following vulnerability has been resolved: net: qualcomm: rmnet: fix global oob in rmnet_ rmnet_link_ops assign a *bigger* maxtype which leads to a global out-of-bounds read when parsing the netlink att ========== BUG: KASAN: globa nlattr.c:386 [inline] BUG: KASAN: global-out-of-bounds in __nla_validate_parse+0x24af/0x2750 lib/nlattr.c:600 task syz-executor.6/84207 CPU: 0 PID: 84207 Comm: syz-executor.6 Tainted: G N 6.1.0 #3 Hardware name: QEM BIOS 1.13.0-1ubuntu1.1 04/01/2014 Call Trace: <TASK> __dump_stack lib/dump_stack.c:88 [inline] dump_stack print_address_description mm/kasan/report.c:284 [inline] print_report+0x172/0x475 mm/kasan/report.c:395 kasan_ validate_nla lib/nlattr.c:386 [inline] __nla_validate_parse+0x24af/0x2750 lib/nlattr.c:600 __nla_parse+0x3e/0x50 include/net/netlink.h:1248 [inline] __rtnl_newlink+0x50a/0x1880 net/core/rtnetlink.c:3485 rtnl_newlink+0x64/0xa +0x43c/0xd70 net/core/rtnetlink.c:6091 netlink_rcv_skb+0x14f/0x410 net/netlink/af_netlink.c:2540 netlink_unicas [inline] netlink_unicast+0x54e/0x800 net/netlink/af_netlink.c:1345 netlink_sendmsg+0x930/0xe50 net/netlink/af_ net/socket.c:714 [inline] sock_sendmsg+0x154/0x190 net/socket.c:734 ____sys_sendmsg+0x6df/0x840 net/socket net/socket.c:2536 __sys_sendmsg+0xf3/0x1c0 net/socket.c:2565 do_syscall_x64 arch/x86/entry/common.c:50 [inl x86/entry/common.c:80 entry_SYSCALL_64_after_hwframe+0x63/0xcd RIP: 0033:0x7fdcf2072359 Code: 28 00 00 90 48 89 f8 48 89 f7 48 89 d6 48 89 ca 4d 89 c2 4d 89 c8 4c 8b 4c 24 08 0f 05 <48> 3d 01 f0 ff ff 73 01 c3 48 c 002b:00007fdcf13e3168 EFLAGS: 00000246 ORIG_RAX: 000000000000002e RAX: ffffffffffffffda RBX: 00007 RDX: 0000000000000000 RSI: 0000000020000200 RDI: 0000000000000003 RBP: 00007fdcf20bd493 R08: 0000 R10: 0000000000000000 R11: 0000000000000246 R12: 0000000000000000 R13: 00007fffbb8d7bdf R14: 00007f </TASK> The buggy address belongs to the variable: rmnet_policy+0x30/0xe0 The buggy address belongs to the p refcount:1 mapcount:0 mapping:0000000000000000 index:0x0 pfn:0x155243 flags: 0x200000000001000(reserved ffffea00055490c8 ffffea00055490c8 0000000000000000 raw: 0000000000000000 0000000000000000 00000001f because: kasan: bad access detected Memory state around the buggy address: ffffffff92c43780: f9 f9 f9 f9 00 00 00 f9 f9 f9 f9 00 00 00 05 f9 f9 f9 f9 06 f9 f9 f9 >ffffffff92c43880: f9 f9 f9 f9 00 00 00 00 00 00 f9 f9 f9 f9 f9 f9 ^ fff f9 f9 f9 f9 f9 f9 f9 ffffffff92c43980: 00 00 00 07 f9 f9 f9 f9 00 00 00 05 f9 f9 f9 f9 According to the comment of `r should be len(destination array) - 1. Hence use `IFLA_RMNET_MAX` here.
CVE-2024-26600	In the Linux kernel, the following vulnerability has been resolved: phy: ti: phy-omap-usb2: Fix NULL pointer dere together with phy-omap-usb2 does not implement send_srp(), we may still attempt to call it. This can happen on an example: configfs-gadget.g1 gadget.0: ECM Suspend configfs-gadget.g1 gadget.0: Port suspended. Triggering wak dereference at virtual address 00000000 when execute ... PC is at 0x0 LR is at musb_gadget_wakeup+0x1d4/0x254 [musb_hdrc] from usb_gadget_wakeup+0x1c/0x3c [udc_core] usb_gadget_wakeup [udc_core] from eth_start_xmi [u_ether] from dev_hard_start_xmit+0x94/0x24c dev_hard_start_xmit from sch_direct_xmit+0x104/0x2e4 sch_dir +0x334/0xd88 __dev_queue_xmit from arp_solicit+0xf0/0x268 arp_solicit from neigh_probe+0x54/0x7c neigh_pr __neigh_event_send from neigh_resolve_output+0x14c/0x1c0 neigh_resolve_output from ip_finish_output2+0x1c +0x40/0xd8 ip_send_skb from udp_send_skb+0x124/0x340 udp_send_skb from udp_sendmsg+0x780/0x984 udp_ __sys_sendto from ret_fast_syscall+0x0/0x58 Let's fix the issue by checking for send_srp() and set_vbus() before c these both could be NULL.
CVE-2024-26602	In the Linux kernel, the following vulnerability has been resolved: sched/membarrier: reduce the ability to hammer sys_membarrier can be very expensive, causing overall slowdowns for everything. So put a lock on the path in orde ability for this to be called at too high of a frequency and saturate the machine.
CVE-2024-26606	In the Linux kernel, the following vulnerability has been resolved: binder: signal epoll threads of self-work In (e)po to determine when data is ready for consumption. Within binder, a thread may initiate a command via BINDER_W make use of epoll_wait() or similar to consume any responses afterwards. It is then crucial that epoll threads are sig work. Otherwise, they risk waiting indefinitely for an event leaving their work unhandled. What is worse, subseque as the thread has pending work.

CVE-2024-26614	In the Linux kernel, the following vulnerability has been resolved: tcp: make sure init the accept_queue's spinlocks C program locally, it causes the following issue: pvqspinlock: lock 0xffff9d181cd5c660 has corrupted value 0x0! W at __pv_queued_spin_unlock_slowpath (kernel/locking/qspinlock_paravirt.h:508) Hardware name: Red Hat KVM, 0010:__pv_queued_spin_unlock_slowpath (kernel/locking/qspinlock_paravirt.h:508) Code: 73 56 3a ff 90 c3 cc cc cc cc 8b 17 48 89 fe 48 c7 c7 30 20 ce 8f e8 ad 56 42 ff <0f> 0b c3 cc cc cc cc 0f 0b 0f 1f 40 00 90 90 90 90 90 EFLAGS: 00010282 RAX: 0000000000000000 RBX: 0000000000000000 RCX: ffff9d1ef60e0908 RDX: 0000000 RDI: ffff9d1ef60e0900 RBP: ffff9d181cd5c280 R08: 0000000000000000 R09: 00000000ffff7fff R10: ffffa8d2006 0000000000000000 R13: ffff9d181cd5c660 R14: ffff9d1813a3f330 R15: 0000000000001000 FS: 00007fa1101840 knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 0000000020000000 CR3 00000000000006f0 Call Trace: <IRQ> _raw_spin_unlock (kernel/locking/spinlock.c:186) inet_csk_reqsk_queue_a inet_csk_complete_hashdance (net/ipv4/inet_connection_sock.c:1358) tcp_check_req (net/ipv4/tcp_minisocks.c:80 ip_protocol_deliver_rcu (net/ipv4/ip_input.c:205) ip_local_deliver_finish (net/ipv4/ip_input.c:234) __netif_receive process_backlog (./include/linux/rcupdate.h:779) __napi_poll (net/core/dev.c:6533) net_rx_action (net/core/dev.c:6 jump_label.h:27) do_softirq (kernel/softirq.c:454 kernel/softirq.c:441) </IRQ> <TASK> __local_bh_enable_ip (ke core/dev.c:4374) ip_finish_output2 (./include/net/neighbour.h:540 net/ipv4/ip_output.c:235) __ip_queue_xmit (net (net/ipv4/tcp_output.c:1462) tcp_rcv_synsent_state_process (net/ipv4/tcp_input.c:6469) tcp_rcv_state_process (ne (net/ipv4/tcp_ipv4.c:1929) __release_sock (./include/net/sock.h:1121 net/core/sock.c:2968) release_sock (net/core/ ipv4/af_inet.c:609) __inet_stream_connect (net/ipv4/af_inet.c:702) inet_stream_connect (net/ipv4/af_inet.c:748) _ net/socket.c:2064) __x64_sys_connect (net/socket.c:2073 net/socket.c:2070 net/socket.c:2070) do_syscall_64 (arch common.c:82) entry_SYSCALL_64_after_hwframe (arch/x86/entry/entry_64.S:129) RIP: 0033:0x7fa10ff05a3d C 00 f3 0f 1e fa 48 89 f8 48 89 f7 48 89 d6 48 89 ca 4d 89 c2 4d 89 c8 4c 8b 4c 24 08 0f 05 <48> 3d 01 f0 ff ff 73 01 RSP: 002b:00007fa110183de8 EFLAGS: 00000202 ORIG_RAX: 000000000000002a RAX: ffffffffffffffda RBX: RDX: 000000000000001c RSI: 0000000020000040 RDI: 0000000000000003 RBP: 00007fa110183e20 R08: 0000 R10: 0000000000000000 R11: 0000000000000202 R12: 00007fa110184640 R13: 0000000000000000 R14: 00007 </TASK> The issue triggering process is analyzed as follows: Thread A Thread B tcp_v4_rcv //receive ack TCP pa tcp_disconnect //disconnect sock ... tcp_set_state(sk, TCP_CLOSE) inet_csk_complete_hashdance ... inet_csk_reqs
CVE-2024-26615	In the Linux kernel, the following vulnerability has been resolved: net/smc: fix illegal rmb_desc access in SMC-D (dumping SMC-D connections. It can be reproduced by following steps: - run nginx/wrk test: smc_run nginx smc_r 'Connection: Close' <URL> - continuously dump SMC-D connections in parallel: watch -n 1 'smcss -D' BUG: kern 0000000000000030 CPU: 2 PID: 7204 Comm: smcss Kdump: loaded Tainted: G E 6.7.0+ #55 RIP: 0010:__smc_d [smc_diag] Call Trace: <TASK> ? __die+0x24/0x70 ? page_fault_oops+0x66/0x150 ? exc_page_fault+0x69/0x14 __smc_diag_dump.constprop.0+0x5e5/0x620 [smc_diag] ? __kmalloc_node_track_caller+0x35d/0x430 ? __alloc_ +0xd0/0xf0 [smc_diag] smc_diag_dump+0x26/0x60 [smc_diag] netlink_dump+0x19f/0x320 __netlink_dump_star +0x6a/0x80 [smc_diag] ? __pfx_smc_diag_dump+0x10/0x10 [smc_diag] sock_diag_rcv_msg+0x121/0x140 ? __p netlink_rcv_skb+0x5a/0x110 sock_diag_rcv+0x28/0x40 netlink_unicast+0x22a/0x330 netlink_sendmsg+0x1f8/0x ____sys_sendmsg+0x24e/0x300 ? copy_msghdr_from_user+0x62/0x80 ___sys_sendmsg+0x7c/0xd0 ? __do_fault do_fault+0xb0/0x110 ? __handle_mm_fault+0x2b0/0x6c0 __sys_sendmsg+0x4d/0x80 do_syscall_64+0x69/0x180 +0x6e/0x76 It is possible that the connection is in process of being established when we dump it. Assumed that the group by smc_conn_create() but the rmb_desc has not yet been initialized by smc_buf_create(), thus causing the ill checking before dump.
CVE-2024-26622	In the Linux kernel, the following vulnerability has been resolved: tomoyo: fix UAF write bug in tomoyo_write_co head->write_buf when write() of long lines is requested, we need to fetch head->write_buf after head->io_sem is h can cause use-after-free-write and double-free problems.

CVE-2024-26625	In the Linux kernel, the following vulnerability has been resolved: llc: call sock_orphan() at release time syzbot rep by a stale sk->sk_wq pointer in a closed llc socket. In commit ff7b11aa481f ("net: socket: set sock->sk to NULL af Biggers hinted that some protocols are missing a sock_orphan(), we need to perform a full audit. In net-next, I plan and amend Eric patch to add a warning. [1] BUG: KASAN: slab-use-after-free in list_empty include/linux/list.h:37 free in waitqueue_active include/linux/wait.h:127 [inline] BUG: KASAN: slab-use-after-free in sock_def_write_sp BUG: KASAN: slab-use-after-free in sock_wfree+0x9a8/0x9d0 net/core/sock.c:2468 Read of size 8 at addr ffff888 1 PID: 27 Comm: ksoftirqd/1 Not tainted 6.8.0-rc1-syzkaller-00049-g6098d87eaf31 #0 Hardware name: QEMU St 1.16.2-debian-1.16.2-1 04/01/2014 Call Trace: <TASK> __dump_stack lib/dump_stack.c:88 [inline] dump_stack_ print_address_description mm/kasan/report.c:377 [inline] print_report+0xc4/0x620 mm/kasan/report.c:488 kasan_ list_empty include/linux/list.h:373 [inline] waitqueue_active include/linux/wait.h:127 [inline] sock_def_write_spac sock_wfree+0x9a8/0x9d0 net/core/sock.c:2468 skb_release_head_state+0xa3/0x2b0 net/core/skbuff.c:1080 skb_re [inline] napi_consume_skb+0x119/0x2b0 net/core/skbuff.c:1404 e1000_unmap_and_free_tx_resource+0x144/0x2 e1000_main.c:1970 e1000_clean_tx_irq drivers/net/ethernet/intel/e1000/e1000_main.c:3860 [inline] e1000_clean+ e1000/e1000_main.c:3801 __napi_poll.constprop.0+0xb4/0x540 net/core/dev.c:6576 napi_poll net/core/dev.c:664 core/dev.c:6778 __do_softirq+0x21a/0x8de kernel/softirq.c:553 run_ksoftirqd kernel/softirq.c:921 [inline] run_kso smpboot_thread_fn+0x660/0xa10 kernel/smpboot.c:164 kthread+0x2c6/0x3a0 kernel/kthread.c:388 ret_from_fork ret_from_fork_asm+0x11/0x20 arch/x86/entry/entry_64.S:242 </TASK> Allocated by task 5167: kasan_save_stac kasan_save_track+0x14/0x30 mm/kasan/common.c:68 unpoison_slab_object mm/kasan/common.c:314 [inline] __ common.c:340 kasan_slab_alloc include/linux/kasan.h:201 [inline] slab_post_alloc_hook mm/slub.c:3813 [inline] kmem_cache_alloc_lru+0x142/0x6f0 mm/slub.c:3879 alloc_inode_sb include/linux/fs.h:3019 [inline] sock_alloc_i alloc_inode+0x5d/0x220 fs/inode.c:260 new_inode_pseudo+0x16/0x80 fs/inode.c:1005 sock_alloc+0x40/0x270 ne net/socket.c:1535 sock_create net/socket.c:1622 [inline] __sys_socket_create net/socket.c:1659 [inline] __sys_sock __do_sys_socket net/socket.c:1720 [inline] __se_sys_socket net/socket.c:1718 [inline] __x64_sys_socket+0x72/0x x86/entry/common.c:52 [inline] do_syscall_64+0xd3/0x250 arch/x86/entry/common.c:83 entry_SYSCALL_64_af kasan_save_stack+0x33/0x50 mm/kasan/common.c:47 kasan_save_track+0x14/0x30 mm/kasan/common.c:68 kas generic.c:640 poison_slab_object mm/kasan/common.c:241 [inline] __kasan_slab_free+0x121/0x1b0 mm/kasan/c kasan.h:184 [inline] slab_free_hook mm/slub.c:2121 [inlin ---truncated---
CVE-2024-26630	In the Linux kernel, the following vulnerability has been resolved: mm: cachestat: fix folio read-after-free in cache from the page cache's xarray to compute its page offset, and check for its dirty and writeback flags. However, we d performing these actions, which means the folio can concurrently be released and reused as another folio/page/slab xarray's existing machinery for the folio page offsets and dirty/writeback states. This changes behavior for tmpfs fi and writeback counters. This is okay as tmpfs doesn't follow conventional writeback cache behavior: its pages get " no longer resident etc.
CVE-2024-26635	In the Linux kernel, the following vulnerability has been resolved: llc: Drop support for ETH_P_TR_802_2. syzbo [0] llc supports ETH_P_802_2 (0x0004) and used to support ETH_P_TR_802_2 (0x0011), and syzbot abused the l &(0x7f0000000040)={@val={0x0, 0x11}, @val, @mpls={[], @llc={@snap={0xaa, 0x1, ')', "90e5dd"}}}}, 0x16) variables {saddr,daddr}.mac based on skb in llc_pdu_decode_sa()/llc_pdu_decode_da() and passes them to __llc_l done only when skb->protocol is htons(ETH_P_802_2), otherwise, __llc_lookup_established() and __llc_lookup_l initialisation existed prior to commit 211ed865108e ("net: delete all instances of special processing for token ring") token ring stuff but forgot to close the door allowing ETH_P_TR_802_2 packets to sneak into llc_rcv(). Let's remo the deprecation. [0]: BUG: KMSAN: uninit-value in __llc_lookup_established+0xe9d/0xf90 __llc_lookup_establis llc_conn.c:611 [inline] llc_conn_handler+0x4bd/0x1360 net/llc/llc_conn.c:791 llc_rcv+0xfbb/0x14a0 net/llc/llc_in net/core/dev.c:5527 [inline] __netif_receive_skb+0x1a6/0x5a0 net/core/dev.c:5641 netif_receive_skb_internal net/ +0x58/0x660 net/core/dev.c:5786 tun_rx_batched+0x3ee/0x980 drivers/net/tun.c:1555 tun_get_user+0x53af/0x66 +0x3af/0x5d0 drivers/net/tun.c:2048 call_write_iter include/linux/fs.h:2020 [inline] new_sync_write fs/read_write. fs/read_write.c:584 ksys_write+0x20f/0x4c0 fs/read_write.c:637 __do_sys_write fs/read_write.c:649 [inline] __se_ __x64_sys_write+0x93/0xd0 fs/read_write.c:646 do_syscall_x64 arch/x86/entry/common.c:51 [inline] do_syscall_ entry_SYSCALL_64_after_hwframe+0x63/0x6b Local variable daddr created at: llc_conn_handler+0x53/0x1360 net/llc/llc_input.c:206 CPU: 1 PID: 5004 Comm: syz-executor994 Not tainted 6.6.0-syzkaller-14500-g1c41041124 Compute Engine/Google Compute Engine, BIOS Google 10/09/2023

CVE-2024-26636	In the Linux kernel, the following vulnerability has been resolved: llc: make llc_ui_sendmsg() more robust against llc_ui_sendmsg(), allocating an skb with no headroom, but subsequently trying to push 14 bytes of Ethernet header releases the socket lock before calling sock_alloc_send_skb(). Then it acquires it again, but does not redo all the sa fix: - Uses LL_RESERVED_SPACE() to reserve space. - Check all conditions again after socket lock is held again mtu limitation. [1] skbuff: skb_under_panic: text:ffff800088baa334 len:1514 put:14 head:ffff0000c9c37000 data:ff dev:bond0 kernel BUG at net/core/skbuff.c:193 ! Internal error: Oops - BUG: 00000000f2000800 [#1] PREEMPT 6875 Comm: syz-executor.0 Not tainted 6.7.0-rc8-syzkaller-00101-g0802e17d9aca-dirty #0 Hardware name: Goog Compute Engine, BIOS Google 11/17/2023 pstate: 60400005 (nZCv daif +PAN -UAO -TCO -DIT -SSBS BTYPE [inline] pc : skb_under_panic+0x13c/0x140 net/core/skbuff.c:203 lr : skb_panic net/core/skbuff.c:189 [inline] lr : s skbuff.c:203 sp : ffff800096f97000 x29: ffff800096f97010 x28: ffff80008cc8d668 x27: dfff800000000000 x26: fff x24: ffff0000c9c36ff2 x23: ffff0000c9c37000 x22: 00000000000005ea x21: 00000000000006c0 x20: 0000000000 x18: 1fffe000368261ce x17: ffff80008e4ed000 x16: ffff80008a8310f8 x15: 0000000000000001 x14: 1ffff00012df 0000000000000000 x11: 0000000000000001 x10: 0000000000ff0100 x9 : e28a51f1087e8400 x8 : e28a51f1087e8 0000000000000000 x5 : 0000000000000001 x4 : 0000000000000001 x3 : ffff800082b78714 x2 : 0000000000000 0000000000000089 Call trace: skb_panic net/core/skbuff.c:189 [inline] skb_under_panic+0x13c/0x140 net/core/sk skbuff.c:2451 eth_header+0x44/0x1f8 net/ethernet/eth.c:83 dev_hard_header include/linux/netdevice.h:3188 [inlin llc_output.c:33 llc_sap_action_send_xid_c+0x170/0x344 net/llc/llc_s_ac.c:85 llc_exec_sap_trans_actions net/llc/ll net/llc/llc_sap.c:182 [inline] llc_sap_state_process+0x1ec/0x774 net/llc/llc_sap.c:209 llc_build_and_send_xid_pkt llc_ui_sendmsg+0x7bc/0xb1c net/llc/af_llc.c:997 sock_sendmsg_nosec net/socket.c:730 [inline] __sock_sendmsg +0x194/0x274 net/socket.c:767 splice_to_socket+0x7cc/0xd58 fs/splice.c:881 do_splice_from fs/splice.c:933 [inli fs/splice.c:1142 splice_direct_to_actor+0x2a0/0x7e4 fs/splice.c:1088 do_splice_direct+0x20c/0x348 fs/splice.c:11 read_write.c:1254 __do_sys_sendfile64 fs/read_write.c:1322 [inline] __se_sys_sendfile64 fs/read_write.c:1308 [in fs/read_write.c:1308 __invoke_syscall arch/arm64/kernel/syscall.c:37 [inline] invoke_syscall+0x98/0x2b8 arch/arm +0x130/0x23c arch/arm64/kernel/syscall.c:136 do_el0_svc+0x48/0x58 arch/arm64/kernel/syscall.c:155 el0_svc+0 common.c:678 el0t_64_sync_handler+0x84/0xfc arch/arm64/kernel/entry-common.c:696 el0t_64_sync+0x190/0x aa1803e6 aa1903e7 a90023f5 94792f6a (d4210000)
CVE-2024-26637	In the Linux kernel, the following vulnerability has been resolved: wifi: ath11k: rely on mac80211 debugfs handlin entries in certain cases, causing a ath11k to crash when it tried to delete the entries later. Fix this by relying on mac and adding them from the vif_add_debugfs handler.
CVE-2024-26642	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_tables: disallow anonymous set with with timeout from userspace, reject this. Exception to this rule is NFT_SET_EVAL to ensure legacy meters still wo
CVE-2024-26643	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_tables: mark set as dead when unbi the rhashtable set gc runs asynchronously, a race allows it to collect elements from anonymous sets with timeouts v path. Mingi Cho originally reported this issue in a different path in 6.1.x with a pipapo set with low timeouts which ("netfilter: nf_tables: use timestamp to check for set element timeout"). Fix this by setting on the dead flag for anon According to 08e4c8c5919f ("netfilter: nf_tables: mark newset as dead on transaction abort"), Florian plans to acce workqueue, therefore, this sets on the dead flag for abort path too.
CVE-2024-26645	In the Linux kernel, the following vulnerability has been resolved: tracing: Ensure visibility when inserting an elem two commands in parallel on a multi-processor AArch64 machine can sporadically produce an unexpected warning while true; do echo hist:key=id.syscall:val=hitcount > \ /sys/kernel/debug/tracing/events/raw_syscalls/sys_enter/tri raw_syscalls/sys_enter/hist sleep 0.001 done $ stress-ng --sysbadaddr $(nproc) The warning looks as follows: [291 [2911.173111] Duplicates detected: 1 [2911.173574] WARNING: CPU: 2 PID: 12247 at kernel/trace/tracing_ma +0x3e0/0x408 [2911.174702] Modules linked in: iscsi_ibft(E) iscsi_boot_sysfs(E) rfkill(E) af_packet(E) nls_iso88 ena(E) tiny_power_button(E) qemu_fw_cfg(E) button(E) fuse(E) efi_pstore(E) ip_tables(E) x_tables(E) xfs(E) libc crct10dif_ce(E) polyval_ce(E) polyval_generic(E) ghash_ce(E) gf128mul(E) sm4_ce_gcm(E) sm4_ce_ccm(E) sm4 sm3_ce(E) sm3(E) sha3_ce(E) sha512_ce(E) sha512_arm64(E) sha2_ce(E) sha256_arm64(E) nvme(E) sha1_ce(E) sg(E) scsi_mod(E) scsi_common(E) efivarfs(E) [2911.174738] Unloaded tainted modules: cppc_cpufreq(E):1 [29 Comm: cat Kdump: loaded Tainted: G E 6.7.0-default #2 1b58bbb22c97e4399dc09f92d309344f69c44a01 [2911.1 EC2 c7g.8xlarge/, BIOS 1.0 11/1/2018 [2911.183208] pstate: 61400005 (nZCv daif +PAN -UAO -TCO +DIT -SS tracing_map_sort_entries+0x3e0/0x408 [2911.184667] lr : tracing_map_sort_entries+0x3e0/0x408 [2911.185310 ffff8000a1513900 x28: ffff0003f272fe80 x27: 0000000000000001 [2911.186600] x26: ffff0003f272fe80 x25: 000 [2911.187458] x23: ffff0003c5788000 x22: ffff0003c16710c8 x21: ffff80008017f180 [2911.188310] x20: ffff800 ffffffffffffffff [2911.189160] x17: 0000000000000000 x16: 0000000000000000 x15: ffff8000a15134b8 [2911.19 205d373432323154 x12: 5b5d313131333731 [2911.190844] x11: 00000000fffeffff x10: 00000000fffeffff x9 : fff 000000000017ffe8 x7 : c0000000fffeffff x6 : 000000000057ffa8 [2911.192554] x5 : ffff0012f6c24ec0 x4 : 00000 [2911.193404] x2 : 0000000000000000 x1 : 0000000000000000 x0 : ffff0003ff254480 [2911.194259] Call trace: +0x3e0/0x408 [2911.195220] hist_show+0x124/0x800 [2911.195692] seq_read_iter+0x1d4/0x4e8 [2911.19619 vfs_read+0xc8/0x300 [2911.197078] ksys_read+0x70/0x108 [2911.197534] __arm64_sys_read+0x24/0x38 [291 [2911.198553] el0_svc_common.constprop.0+0xd0/0xf8 [2911.199157] do_el0_svc+0x28/0x40 [2911.199613] e el0t_64_sync_handler+0x13c/0x158 [2911.200621] el0t_64_sync+0x1a8/0x1b0 [2911.201115] ---[end trace 000 be caused by CPU reordering of writes issued from __tracing_map_insert(). The check for the presence of an eleme READ_ONCE(entry->val); if (val && keys_match(key, val->key, map->key_size)) ... The write of a new entry is: key, map->key_size); entry->val = elt; The "memcpy(elt->key, key, map->key_size);" and "entry->val = elt;" store another CPU. This second CPU might then incorrectly determine that a new key doesn't match an already present v

CVE-2024-26649	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: Fix the null pointer when load rlc because of wrong header size, the pointer to the rlc firmware is released in function amdgpu_ucode_request. There So skip validation to fix it.
CVE-2024-26651	In the Linux kernel, the following vulnerability has been resolved: sr9800: Add check for usbnet_get_endpoints Ac return the error if it fails in order to transfer the error.
CVE-2024-26654	In the Linux kernel, the following vulnerability has been resolved: ALSA: sh: aica: reorder cleanup operations to a could schedule the spu_dma_work and the spu_dma_work could also arm the dreamcastcard->timer. When the snd will be deallocated. But it could still be dereferenced in the worker thread. The reason is that del_timer() will return handler is running or not and the worker could be rescheduled in the timer handler. As a result, the UAF bug will h (Thread 1) | (Thread 2) snd_aicapcm_pcm_close() | ... | run_spu_dma() //worker | mod_timer() flush_work() | del_t kfree(dreamcastcard->channel) | schedule_work() | run_spu_dma() //worker ... | dreamcastcard->channel-> //USE I corner cases, call mod_timer() conditionally in run_spu_dma(), then implement PCM sync_stop op to cancel both t be called from PCM core appropriately when needed.
CVE-2024-26664	In the Linux kernel, the following vulnerability has been resolved: hwmon: (coretemp) Fix out-of-bounds memory before out-of-bounds check. The problem might be triggered on systems with more than 128 cores per package.
CVE-2024-26666	In the Linux kernel, the following vulnerability has been resolved: wifi: mac80211: fix RCU use in TDLS fast-xmi protection, but isn't guaranteed to actually have protection. Fix that.
CVE-2024-26671	In the Linux kernel, the following vulnerability has been resolved: blk-mq: fix IO hang from sbitmap wakeup race __add_wait_queue() may be re-ordered with the following blk_mq_get_driver_tag() in case of getting driver tag fa waitqueue_active() may not observe the added waiter in blk_mq_mark_tag_wait() and wake up nothing, meantime tag successfully. This issue can be reproduced by running the following test in loop, and fio hang can be observed i in laptop. modprobe -r scsi_debug modprobe scsi_debug delay=0 dev_size_mb=4096 max_queue=1 host_max_que pseudo/drivers/scsi_debug/adapter*/host*/target*/*/block/* | head -1 | xargs basename` fio --filename=/dev/"$dev" \ --runtime=100 --numjobs=40 --time_based --name=test \ --ioengine=libaio Fix the issue by adding one explicit ba just fine in case of running out of tag.
CVE-2024-26673	In the Linux kernel, the following vulnerability has been resolved: netfilter: nft_ct: sanitize layer 3 and 4 protocol r families other than NFPROTO_{IPV4,IPV6,INET}. - Disallow layer 4 protocol with no ports, since destination po
CVE-2024-26675	In the Linux kernel, the following vulnerability has been resolved: ppp_async: limit MRU to 64K syzbot triggered WARN_ON_ONCE_GFP(order > MAX_PAGE_ORDER, gfp) Willem fixed a similar issue in commit c0a2a1b0d the same sanity check for ppp_async_ioctl(PPPIOCSMRU) [1]: WARNING: CPU: 1 PID: 11 at mm/page_alloc.c: mm/page_alloc.c:4543 Modules linked in: CPU: 1 PID: 11 Comm: kworker/u4:0 Not tainted 6.8.0-rc2-syzkaller-g4 Google Google Compute Engine/Google Compute Engine, BIOS Google 11/17/2023 Workqueue: events_unbound (nzCv daIF +PAN -UAO -TCO -DIT -SSBS BTYPE=--) pc : __alloc_pages+0x308/0x698 mm/page_alloc.c:4543 mm/page_alloc.c:4537 sp : ffff800093967580 x29: ffff800093967660 x28: ffff8000939675a0 x27: dfff800000000 0000000000000000 x24: ffff8000939675c0 x23: 0000000000000000 x22: 0000000000060820 x21: 1ffff0001272c 0000000000000010 x18: ffff800093967120 x17: ffff800083bded5c x16: ffff80008ac97500 x15: 0000000000000 0000000000000000 x12: 0000000000000000 x11: ffff70001272cec1 x10: 1ffff0001272cec0 x9 : 0000000000000 0000000000000000 x6 : 000000000000003f x5 : 00000000ffffffff x4 : 0000000000000000 x3 : 0000000000000 0000000000000000 x0 : ffff8000939675e0 Call trace: __alloc_pages+0x308/0x698 mm/page_alloc.c:4543 __alloc [inline] alloc_pages_node include/linux/gfp.h:261 [inline] __kmalloc_large_node+0xbc/0x1fc mm/slub.c:3926 __c __kmalloc_node_track_caller+0x418/0x620 mm/slub.c:4001 kmalloc_reserve+0x17c/0x23c net/core/skbuff.c:590 skbuff.c:651 __netdev_alloc_skb+0xb8/0x3e8 net/core/skbuff.c:715 netdev_alloc_skb include/linux/skbuff.h:3235 skbuff.h:3248 [inline] ppp_async_input drivers/net/ppp/ppp_async.c:863 [inline] ppp_asynctty_receive+0x588/0x1 tty_ldisc_receive_buf+0x12c/0x15c drivers/tty/tty_buffer.c:390 tty_port_default_receive_buf+0x74/0xac drivers/tt tty/tty_buffer.c:444 [inline] flush_to_ldisc+0x284/0x6e4 drivers/tty/tty_buffer.c:494 process_one_work+0x694/0x process_scheduled_works kernel/workqueue.c:2706 [inline] worker_thread+0x938/0xef4 kernel/workqueue.c:2787 ret_from_fork+0x10/0x20 arch/arm64/kernel/entry.S:860
CVE-2024-26677	In the Linux kernel, the following vulnerability has been resolved: rxrpc: Fix delayed ACKs to not set the reference delayed ACKs to not set the reference serial number as they can't be used as an RTT reference.
CVE-2024-26679	In the Linux kernel, the following vulnerability has been resolved: inet: read sk->sk_family once in inet_recv_error the socket lock. IPv6 socket could mutate to IPv4 with IPV6_ADDRFORM socket option and trigger a KCSAN wa
CVE-2024-26682	In the Linux kernel, the following vulnerability has been resolved: wifi: mac80211: improve CSA/ECSA connectio commit, we pretty quickly found that some APs have ECSA elements stuck in their probe response, so using that to happening we never connect to such an AP. Improve this situation by checking more carefully and ignoring the EC ECSA element being stuck in the probe response. Additionally, allow connecting to an AP that's switching to a cha mode. In this case, we may just have to adjust bandwidth later. If it's actually switching channels, it's better not to t

CVE-2024-26683	In the Linux kernel, the following vulnerability has been resolved: wifi: cfg80211: detect stuck ECSA element in p that we don't try to connect to an AP that is currently in a channel switch process, since that might want the channe connect in time to hear the switching in a beacon. This was in commit c09c4f31998b ("wifi: mac80211: don't conn However, we promptly got a report that this caused new connection failures, and it turns out that the AP that we no advertising an extended channel switch announcement, even with quiet. The AP in question was an Asus RT-AC53 g21a5898. As a first step, attempt to detect that we're dealing with such a situation, so mac80211 can use this later.
CVE-2024-26684	In the Linux kernel, the following vulnerability has been resolved: net: stmmac: xgmac: fix handling of DPP safety 56e58d6c8a56 ("net: stmmac: Implement Safety Features in XGMAC core") checks and reports safety errors, but l channel in DMA unhandled at all, lead to a storm of interrupt. Fix it by checking and clearing the DMA_DPP_Inte
CVE-2024-26685	In the Linux kernel, the following vulnerability has been resolved: nilfs2: fix potential bug in end_buffer_async_wr end_buffer_async_write(), which handles the completion of block device writes, may detect abnormal condition of a BUG_ON failure when using nilfs2. Nilfs2 itself does not use end_buffer_async_write(). But, the async_write fla 7f42ec394156 ("nilfs2: fix issue with race condition of competition between segments for dirty blocks") as a means blocks in nilfs_lookup_dirty_data_buffers() and nilfs_lookup_node_buffers() and the resulting crash. This modifica and b-tree node blocks where the page caches are independent. However, it was irrelevant and redundant to also int and super root blocks that share buffers with the backing device. This led to the possibility that the BUG_ON check described above, if independent writebacks of the backing device occurred in parallel. The use of async_write for s removed in a previous change. Fix this issue by removing the manipulation of the async_write flag for the remainin
CVE-2024-26696	In the Linux kernel, the following vulnerability has been resolved: nilfs2: fix hang in nilfs_lookup_dirty_data_buff migrate_pages_batch() called by mbind() and nilfs_lookup_dirty_data_buffers() called in the log writer of nilfs2. W folio and waits for the writeback to complete, the log writer thread that should bring the writeback to completion pi nilfs_lookup_dirty_data_buffers() that it calls for subsequent log creation and was trying to lock the folio. Thus cau unexpected that folios/pages in the middle of writeback will be updated and become dirty. Nilfs2 adds a checksum and uses it for recovery at mount, so data changes during writeback are suppressed. Since this is broken, an unclear to fail. Investigation revealed that the root cause is that the wait for writeback completion in nilfs_page_mkwrite() not require stable writes, data may be modified without waiting. Fix these issues by making nilfs_page_mkwrite() w stable write requirement of the backing device.
CVE-2024-26697	In the Linux kernel, the following vulnerability has been resolved: nilfs2: fix data corruption in dsync block recove nilfs_recovery_copy_block() of nilfs_recovery_dsync_blocks(), which recovers data from logs created by data sync shutdown, incorrectly calculates the on-page offset when copying repair data to the file's page cache. In environme page size, this flaw can cause data corruption and leak uninitialized memory bytes during the recovery process. Fix calculation on the page.
CVE-2024-26698	In the Linux kernel, the following vulnerability has been resolved: hv_netvsc: Fix race condition between netvsc_p ac5047671758 ("hv_netvsc: Disable NAPI before closing the VMBus channel"), napi_disable was getting called fo without confirming if they are enabled or not. This caused hv_netvsc getting hung at napi_disable, when netvsc_pr >subchan_work has not started yet. netvsc_subchan_work() -> rndis_set_subchannel() has not created the sub-chan is not running. netvsc_remove() calls cancel_work_sync(&nvdev->subchan_work), for which netvsc_subchan_wor bit NAPI_STATE_SCHED because it ensures NAPI cannot be scheduled. Then netvsc_sc_open() -> napi_enable w bit, so it can be scheduled. napi_disable() does the opposite. Now during netvsc_device_remove(), when napi_disal napi_disable gets stuck on infinite msleep. This fix addresses this problem by ensuring that napi_disable() is not ge But netif_napi_del() is still necessary for these non-enabled NAPI struct for cleanup purpose. Call trace: [654.559 2321 ppid: 1091 flags:0x00004002 [654.568030] Call Trace: [654.571221] <TASK> [654.573790] __schedule+0 +0x69/0xf0 [654.581214] schedule_timeout+0x87/0x140 [654.585463] ? __bpf_trace_tick_stop+0x20/0x20 [654 napi_disable+0x2b/0x80 [654.597437] netvsc_device_remove+0x8a/0x1f0 [hv_netvsc] [654.603935] rndis_filter [654.611101] ? do_wait_intr+0xb0/0xb0 [654.615753] netvsc_remove+0x7c/0x120 [hv_netvsc] [654.621675] vm
CVE-2024-26702	In the Linux kernel, the following vulnerability has been resolved: iio: magnetometer: rm3100: add boundary check RM3100_REG_TMRC Recently, we encounter kernel crash in function rm3100_common_probe caused by out of (because of underlying hardware failures). Add boundary check to prevent out of bound access.
CVE-2024-26704	In the Linux kernel, the following vulnerability has been resolved: ext4: fix double-free of blocks due to wrong ext moved_len is only updated when all moves are successfully executed, and only discards orig_inode and donor_inod When the loop fails to exit after successfully moving some extents, moved_len is not updated and remains at 0, so i moved extents overlap with the preallocated extents, the overlapped extents are freed twice in ext4_mb_release_inc described in commit 94d7c16cbbbd ("ext4: Fix double-free of blocks with EXT4_IOC_MOVE_EXT")), and bb_fr executed, a zero-division bug is triggered in mb_update_avg_fragment_size() because bb_free is not zero and bb_f after each extent move to avoid the issue.
CVE-2024-26712	In the Linux kernel, the following vulnerability has been resolved: powerpc/kasan: Fix addr error caused by page a is not page aligned, at the begin of for loop, k_cur = k_start & PAGE_MASK is less than k_start, and then `va = bl addr va is invalid, because the memory address space from va to block is not alloced by memblock_alloc, which w it will be used by other places. As a result, memory overwriting occurs. for example: int __init __weak kasan_init_ block(dcd97000) k_start(feef7400) k_end(feeff3fe) */ block = memblock_alloc(k_end - k_start, PAGE_SIZE); [...] < k_end; k_cur += PAGE_SIZE) { /* at the begin of for loop * block(dcd97000) va(dcd96c00) k_cur(feef7000) k_ block(dcd97000), va is invalid */ void *va = block + k_cur - k_start; [...] } [...] } Therefore, page alignment is perfo ensure the validity of the VA address.

CVE-2024-26714	In the Linux kernel, the following vulnerability has been resolved: interconnect: qcom: sc8180x: Mark CO0 BCM all times, otherwise some hardware (like the UFS controller) loses its connection to the rest of the SoC, resulting in spectacular logspam. Mark it as keepalive to prevent such cases.
CVE-2024-26722	In the Linux kernel, the following vulnerability has been resolved: ASoC: rt5645: Fix deadlock in rt5645_jack_det rt5645_jack_detect_work(), where rt5645->jd_mutex is left locked forever. That may lead to deadlock when rt564 time. Found by Linux Verification Center (linuxtesting.org) with SVACE.
CVE-2024-26733	In the Linux kernel, the following vulnerability has been resolved: arp: Prevent overflow in arp_req_get(). syzkalle arp_req_get(). [0] When ioctl(SIOCGARP) is issued, arp_req_get() looks up an neighbour entry and copies neigh-> arp_ha here is struct sockaddr, not struct sockaddr_storage, so the sa_data buffer is just 14 bytes. In the splat below arp_flags. We initialise the field just after the memcpy(), so it's not a problem. However, when dev->addr_len is gr arp_netmask is overwritten, which could be set as htonl(0xFFFFFFFFUL) in arp_ioctl() before calling arp_req_get length of memcpy(). Note that commit b5f0de6df6dc ("net: dev: Convert sa_data to flexible array in struct sockadd detected field-spanning write (size 16) of single field "r->arp_ha.sa_data" at net/ipv4/arp.c:1128 (size 14) WARNI arp.c:1128 arp_req_get+0x411/0x4a0 net/ipv4/arp.c:1128 Modules linked in: CPU: 0 PID: 144638 Comm: syz-exe name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 1.16.0-debian-1.16.0-5 04/01/2014 RIP: 0010:arp_req_g ff ff e8 41 42 de fb b9 0e 00 00 00 4c 89 fe 48 c7 c2 20 6d ab 87 48 c7 c7 80 6d ab 87 c6 05 25 af 72 04 01 e8 5f 8 be 03 00 00 00 4c 89 e7 e8 a6 RSP: 0018:ffffc900050b7998 EFLAGS: 00010286 RAX: 0000000000000000 RBX: RDX: 0000000000000000 RSI: ffffffff8641a44a RDI: 0000000000000001 RBP: ffffc900050b7a98 R08: 0000000 R10: 0000000000000000 R11: 203a7970636d656d R12: ffff888039c54000 R13: 1ffff92000a16f37 R14: ffff88803 FS: 00007f172bf306c0(0000) GS:ffff88805aa00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 000 00007f172b3569f0 CR3: 0000000057f12005 CR4: 00000000000770ef0 DR0: 0000000000000000 DR1: 0000000 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 PKRU: 55555554 Call Trace: <TASK> arp inet_ioctl+0x314/0x3a0 net/ipv4/af_inet.c:981 sock_do_ioctl+0xdf/0x260 net/socket.c:1204 sock_ioctl+0x3ef/0x6 [inline] __do_sys_ioctl fs/ioctl.c:870 [inline] __se_sys_ioctl fs/ioctl.c:856 [inline] __x64_sys_ioctl+0x18e/0x220 f common.c:51 [inline] do_syscall_64+0x37/0x90 arch/x86/entry/common.c:81 entry_SYSCALL_64_after_hwfram Code: 66 2e 0f 1f 84 00 00 00 00 00 0f 1f 00 f3 0f 1e fa 48 89 f8 48 89 f7 48 89 d6 48 89 ca 4d 89 c2 4d 89 c8 4c 8 c3 48 c7 c1 b8 ff ff ff f7 d8 64 89 01 48 RSP: 002b:00007f172bf300b8 EFLAGS: 00000246 ORIG_RAX: 0000000 00007f172b3abf80 RCX: 00007f172b262b8d RDX: 0000000020000000 RSI: 0000000000008954 RDI: 0000000 0000000000000000 R09: 0000000000000000 R10: 0000000000000000 R11: 0000000000000246 R12: 0000000 00007f172b3abf80 R15: 00007f172bf10000 </TASK>
CVE-2024-26734	In the Linux kernel, the following vulnerability has been resolved: devlink: fix possible use-after-free and memory structure for the subsystem must be registered before registering the generic netlink family. Make an unregister in c
CVE-2024-26735	In the Linux kernel, the following vulnerability has been resolved: ipv6: sr: fix possible use-after-free and null-ptr- subsystem must be registered before registering the generic netlink family.
CVE-2024-26736	In the Linux kernel, the following vulnerability has been resolved: afs: Increase buffer size in afs_update_volume_ is 20 characters. So increase idbuf[] size up to 24 to avoid overflow. Found by Linux Verification Center (linuxtest + NUL, so increase it to 24 and use snprintf()]
CVE-2024-26748	In the Linux kernel, the following vulnerability has been resolved: usb: cdns3: fix memory double free when handl { 830 spin_unlock(&priv_dev->lock); 831 usb_gadget_giveback_request(&priv_ep->endpoint, 832 request); 833 s 836 if (request->buf == priv_dev->zlp_buf) 837 cdns3_gadget_ep_free_request(&priv_ep->endpoint, request); Dri request when queue a packet, which length mod max packet size is 0. When transfer complete, run to line 831, usb requestion. 836 condition is true, so cdns3_gadget_ep_free_request() free this request again. Log: [1920.140696][in cdns3_gadget_giveback+0x134/0x2c0 [cdns3] [1920.140696][T150] [1920.151837][T150] Use-after-free rea [1920.159082][T150] cdns3_gadget_giveback+0x134/0x2c0 [cdns3] [1920.164988][T150] cdns3_transfer_com at line 829, skip call usb_gadget_giveback_request() if it is additional zero length packet request. Needn't call usb_ allocated in this driver.
CVE-2024-26749	In the Linux kernel, the following vulnerability has been resolved: usb: cdns3: fixed memory use after free at cdns3 cdns3_gadget_ep_free_request(&priv_ep->endpoint, &priv_req->request); list_del_init(&priv_req->list); ... 'priv_r cdns3_gadget_ep_free_request(). But list_del_init() use priv_req->list after it. [1542.642868][T534] BUG: KFEN __list_del_entry_valid+0x10/0xd4 [1542.642868][T534] [1542.653162][T534] Use-after-free read at 0x0000000 [T534] __list_del_entry_valid+0x10/0xd4 [1542.665375][T534] cdns3_gadget_ep_disable+0x1f8/0x388 [cdns3] +0x44/0xe4 [1542.675948][T534] ffs_func_eps_disable+0x64/0xc8 [1542.680839][T534] ffs_func_set_alt+0x7 ffs_func_disable+0x18/0x28 Move list_del_init() before cdns3_gadget_ep_free_request() to resolve this problem.
CVE-2024-26752	In the Linux kernel, the following vulnerability has been resolved: l2tp: pass correct message length to ip6_append accounting for the transport header twice when splicing more data into an already partially-occupied skbuff. To ma contains data using skb_queue_empty when deciding how much data to append using ip6_append_data. However, incorrect: ulen = len + skb_queue_empty(&sk->sk_write_queue) ? transhdrlen : 0; ...due to C operator precedence, messages with a non-zero length, which results in corrupted packets on the wire. Add parentheses to correct the cal

CVE-2024-26754	In the Linux kernel, the following vulnerability has been resolved: gtp: fix use-after-free and null-ptr-deref in gtp_g pernet operations structure for the subsystem must be registered before registering the generic netlink family. Syzk gtp_genl_dump_pdp' bug: general protection fault, probably for non-canonical address 0xdffffc0000000002: 0000 KASAN: null-ptr-deref in range [0x0000000000000010-0x0000000000000017] CPU: 1 PID: 5826 Comm: gtp No name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS 1.16.0-alt1 04/01/2014 RIP: 0010:gtp_genl_dump_pdp+0x 86 df 58 45 85 f6 0f 85 4e 04 00 00 e8 c5 ee 86 df 48 8b 54 24 18 48 b8 00 00 00 00 00 fc ff df 48 c1 ea 03 <80> 3 18 4c 8b 30 4c 39 f0 74 RSP: 0018:ffff888014107220 EFLAGS: 00010202 RAX: dffffc0000000000 RBX: 00000 RDX: 0000000000000002 RSI: 0000000000000000 RDI: 0000000000000000 RBP: 0000000000000000 R08: 000 R10: 0000000000000000 R11: 0000000000000000 R12: 0000000000000000 R13: ffff88800fcda588 R14: 000000 FS: 00007f1be4eb05c0(0000) GS:ffff88806ce80000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 000 00007f1be4e766cf CR3: 000000000c33e000 CR4: 0000000000750ef0 PKRU: 55555554 Call Trace: <TASK> ? sl +0x50/0xd0 ? exc_general_protection+0x148/0x220 ? asm_exc_general_protection+0x22/0x30 ? gtp_genl_dump_ +0x1dd/0x350 ? __pfx___alloc_skb+0x10/0x10 genl_dumpit+0x11d/0x230 netlink_dump+0x5b9/0xce0 ? lockdep __pfx_netlink_dump+0x10/0x10 ? kasan_save_track+0x10/0x40 ? __kasan_kmalloc+0x9b/0xa0 ? genl_start+0x67 genl_family_rcv_msg_dumpit+0x1bb/0x2d0 ? __pfx_genl_family_rcv_msg_dumpit+0x10/0x10 ? genl_op_from_s +0x1d0/0x240 ? __pfx_genl_start+0x10/0x10 ? __pfx_genl_dumpit+0x10/0x10 ? __pfx_genl_done+0x10/0x10 ? s
CVE-2024-26760	In the Linux kernel, the following vulnerability has been resolved: scsi: target: pscsi: Fix bio_put() for error case A bio_kmalloc into a simple kmalloc wrapper"), a bio allocated by bio_kmalloc() must be freed by bio_uninit() and k case, hitting WARN and NULL pointer dereference in bio_free().
CVE-2024-26764	In the Linux kernel, the following vulnerability has been resolved: fs/aio: Restrict kiocb_set_cancel_fn() to I/O sub is called for I/O submitted via io_uring, the following kernel warning appears: WARNING: CPU: 3 PID: 368 at fs/ Call trace: kiocb_set_cancel_fn+0x9c/0xa8 ffs_epfile_read_iter+0x144/0x1d0 io_read+0x19c/0x498 io_issue_sqe+ __arm64_sys_io_uring_enter+0x104/0xab0 invoke_syscall+0x58/0x11c el0_svc_common+0xb4/0xf4 do_el0_svc+ el0t_64_sync_handler+0x68/0xb4 el0t_64_sync+0x1a4/0x1a8 Fix this by setting the IOCB_AIO_RW flag for read
CVE-2024-26766	In the Linux kernel, the following vulnerability has been resolved: IB/hfi1: Fix sdma.h tx->num_descs off-by-one e `fd8958efe877` introduced another error causing the `descs` array to overflow. This reults in further crashes easily call. [1080.836473] general protection fault, probably for non-canonical address 0x400300015528b00a: 0000 [#1] RIP: 0010:hfi1_ipoib_build_ib_tx_headers.constprop.0+0xe1/0x2b0 [hfi1] -- [1080.974535] Call Trace: [1080.97 hfi1_ipoib_send_dma_common+0x7a/0x2e0 [hfi1] [1081.027364] hfi1_ipoib_send_dma_list+0x62/0x270 [hfi1] [[hfi1] [1081.042001] ipoib_start_xmit+0x2a9/0x2d0 [ib_ipoib] [1081.046978] dev_hard_start_xmit+0xc4/0x210 +0x59/0xa0 crash> ipoib_txreq 0xffff9cfeba229f00 struct ipoib_txreq { txreq = { list = { next = 0xffff9cfeba229f0 = 0xffff9cfeba229f40, coalesce_buf = 0x0, wait = 0xffff9cfea4e69a48, complete = 0xffffffffc0fe0760 <hfi1_ipoib_ tlen = 0x0, num_desc = 0x0, desc_limit = 0x6, next_descq_idx = 0x45c, coalesce_idx = 0x0, flags = 0x0, descs = { SDMA_DESC0_FIRST_DESC_FLAG (bit 63) }, { qw = { 0x3800014231b108, 0x4} }, { qw = { 0x310000e4ee0 0x8} }, { qw = { 0x59000dfb9d0000, 0x8} }, { qw = { 0x78000e02e40000, 0x8} }} }, sdma_hdr = 0x4003000155 structure sdma_status = 0x0, SDMA_DESC0_LAST_DESC_FLAG (bit 62) complete = 0x0, priv = 0x0, txq = 0xf If an SDMA send consists of exactly 6 descriptors and requires dword padding (in the 7th descriptor), the sdma_tx and the packet will overflow into the container structure. This results in a panic when the send completion runs. Th elements of the container structure get corrupted. The fix is to use the correct expression in _pad_sdma_tx_descs() array. With this patch the crashes are no longer reproducible and the machine is stable.
CVE-2024-26772	In the Linux kernel, the following vulnerability has been resolved: ext4: avoid allocating blocks from corrupted gro logic for checking if the group's block bitmap is corrupt under the protection of the group lock to avoid allocating b bitmap.
CVE-2024-26773	In the Linux kernel, the following vulnerability has been resolved: ext4: avoid allocating blocks from corrupted gro if the group block bitmap is corrupted before using ac_b_ex in ext4_mb_try_best_found() to avoid allocating block bitmap in the following concurrency and making the situation worse. ext4_mb_regular_allocator ext4_lock_group(check if the group bbitmap is corrupted ext4_mb_complex_scan_group // Scan group gets ac_b_ex but doesn't use ext4_mark_group_bitmap_corrupted(group) // The block bitmap was corrupted during // the group unlock gap. ext4 >ac_sb, group) ext4_mb_use_best_found mb_mark_used // Allocating blocks in block bitmap corrupted group
CVE-2024-26779	In the Linux kernel, the following vulnerability has been resolved: wifi: mac80211: fix race condition on enabling after the sta has been uploaded to the driver, otherwise it could end up passing the not-yet-uploaded sta via drv_tx because of uninitialized drv_priv data. Add a missing sta->uploaded check and re-check fast xmit after inserting a s
CVE-2024-26779	In the Linux kernel, the following vulnerability has been resolved: wifi: mac80211: fix race condition on enabling after the sta has been uploaded to the driver, otherwise it could end up passing the not-yet-uploaded sta via drv_tx because of uninitialized drv_priv data. Add a missing sta->uploaded check and re-check fast xmit after inserting a s

CVE-2024-26787	In the Linux kernel, the following vulnerability has been resolved: mmc: mmci: stm32: fix DMA API overlapping CONFIG_DMA_API_DEBUG_SG results in the following warning: DMA-API: mmci-pl18x 48220000.mmc: cac mappings aren't supported WARNING: CPU: 1 PID: 51 at kernel/dma/debug.c:568 add_dma_entry+0x234/0x2f4 Comm: kworker/1:2 Not tainted 6.1.28 #1 Hardware name: STMicroelectronics STM32MP257F-EV1 Evaluation mmc_rescan Call trace: add_dma_entry+0x234/0x2f4 debug_dma_map_sg+0x198/0x350 __dma_map_sg_attrs+0 sdmmc_idma_prep_data+0x80/0xc0 mmci_prep_data+0x38/0x84 mmci_start_data+0x108/0x2dc mmci_request+0 +0x68/0x140 mmc_start_request+0x94/0xc0 mmc_wait_for_req+0x70/0x100 mmc_send_tuning+0x108/0x1ac sdr mmc_execute_tuning+0x48/0xec mmc_sd_init_uhs_card.part.0+0x208/0x464 mmc_sd_init_card+0x318/0x89c mm +0x244/0x320 DMA API debug brings to light leaking dma-mappings as dma_map_sg and dma_unmap_sg are no mmci_cmd_irq function, only mmci_dma_error function is called and as this API is not managed on stm32 variant path.
CVE-2024-26788	In the Linux kernel, the following vulnerability has been resolved: dmaengine: fsl-qdma: init irq after reg initializa registers are configured so that interrupts that may have been pending from a primary kernel don't get processed by panic with the following trace: Call trace: fsl_qdma_queue_handler+0xf8/0x3e8 __handle_irq_event_percpu+0x78 handle_irq_event+0x44/0x78 handle_fasteoi_irq+0xc8/0x178 generic_handle_irq+0x24/0x38 __handle_domain_i el1_irq+0xb8/0x180 _raw_spin_unlock_irqrestore+0x14/0x40 __setup_irq+0x4bc/0x798 request_threaded_irq+0x +0x74/0xe8 fsl_qdma_probe+0x4d4/0xca8 platform_drv_probe+0x50/0xa0 really_probe+0xe0/0x3f8 driver_probe +0x6c/0x78 __driver_attach+0xbc/0x158 bus_for_each_dev+0x5c/0x98 driver_attach+0x20/0x28 bus_add_driver __platform_driver_register+0x44/0x50 fsl_qdma_driver_init+0x18/0x20 do_one_initcall+0x48/0x258 kernel_init_ ret_from_fork+0x10/0x18
CVE-2024-26790	In the Linux kernel, the following vulnerability has been resolved: dmaengine: fsl-qdma: fix SoC may hang on 16 b errata: The SoC may hang on 16 byte unaligned read transactions by QDMA. Unaligned read transactions initiated On-Chip), causing a deadlock condition. Stalled transactions will trigger completion timeouts in PCIe controller. W source descriptor prefetchable bit (SD[PF] = 1). Implement this workaround.
CVE-2024-26791	In the Linux kernel, the following vulnerability has been resolved: btrfs: dev-replace: properly validate device nam buffers passed to device replace are not properly checked for string termination which could lead to a read out of b validates both source and target device name buffers. For devid as the source initialize the buffer to empty string in originally analyzed and fixed in a different way by Edward Adam Davis (see links).
CVE-2024-26793	In the Linux kernel, the following vulnerability has been resolved: gtp: fix use-after-free and null-ptr-deref in gtp_n structure for the subsystem must be registered after registering the gtp_net_ops pernet operations structure. Syzkall gtp_genl_dump_pdp' bug: [1010.702740] gtp: GTP module unloaded [1010.715877] general protection fault, prob 0xdffffc0000000001: 0000 [#1] SMP KASAN NOPTI [1010.715888] KASAN: null-ptr-deref in range [0x000000 [1010.715895] CPU: 1 PID: 128616 Comm: a.out Not tainted 6.8.0-rc6-std-def-alt1 #1 [1010.715899] Hardware ICH9, 2009), BIOS 1.16.0-alt1 04/01/2014 [1010.715908] RIP: 0010:gtp_newlink+0x4d7/0x9c0 [gtp] [1010.7159 00 48 8b bb d8 05 00 00 e8 ed f6 ff ff 48 89 c2 48 89 c5 48 b8 00 00 00 00 00 fc ff df 48 c1 ea 03 <80> 3c 02 00 0 48 b8 00 00 00 [1010.715920] RSP: 0018:ffff888020fbf180 EFLAGS: 00010203 [1010.715929] RAX: dffffc000 0000000000000000 [1010.715933] RDX: 0000000000000001 RSI: ffffffff84805280 RDI: 0000000000000282 [1 R08: 0000000000000001 R09: 0000000000000000 [1010.715942] R10: 0000000000000001 R11: 0000000000000 R13: 0000000000000000 R14: 0000000000000000 R15: 0000000000000400 [1010.715953] FS: 00007fd1509ab5 knlGS:0000000000000000 [1010.715958] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [1010.715962 000000001c07a000 CR4: 0000000000750ee0 [1010.715968] PKRU: 55555554 [1010.715972] Call Trace: [1010 [1010.715995] ? die_addr+0x43/0x70 [1010.716002] ? exc_general_protection+0x199/0x2f0 [1010.716016] ? as [1010.716026] ? gtp_newlink+0x4d7/0x9c0 [gtp] [1010.716034] ? gtp_net_exit+0x150/0x150 [gtp] [1010.71604 [1010.716051] ? rtnl_setlink+0x3c0/0x3c0 [1010.716063] ? is_bpf_text_address+0xc0/0x1f0 [1010.716070] ? ke [1010.716076] ? __kernel_text_address+0x56/0xa0 [1010.716084] ? unwind_get_return_address+0x5a/0xa0 [10 +0x30/0x30 [1010.716098] ? arch_stack_walk+0x9e/0xf0 [1010.716106] ? stack_trace_save+0x91/0xd0 [1010.7 +0x170/0x170 [1010.716121] ? __lock_acquire+0x15c5/0x5380 [1010.716139] ? mark_held_locks+0x9e/0xe0 [+0x35f/0x3c0 [1010.716155] ? __rtnl_newlink+0x1700/0x1700 [1010.716160] rtnl_newlink+0x69/0xa0 [1010.7 [1010.716172] ? rtnl_fdb_dump+0x9f0/0x9f0 [1010.716179] ? lock_acquire+0x1fe/0x560 [1010.716188] ? netli netlink_rcv_skb+0x14d/0x440 [1010.716202] ? rtnl_fdb_dump+0x9f0/0x9f0 [1010.716208] ? netlink_ack+0xab0 +0x202/0xd50 [1010.716220] ? netlink_deliver_tap+0x218/0xd50 [1010.716226] ? __virt_addr_valid+0x30b/0x5 +0x54b/0x800 [1010.716240] ? netlink_attachskb+0x870/0x870 [1010.716248] ? __check_object_size+0x2de/0x +0x938/0xe40 [1010.716261] ? netlink_unicast+0x800/0x800 [1010.716269] ? __import_iovec+0x292/0x510 [1 [1010.716284] __sock_sendmsg+0x159/0x190 [1010.716290] ____sys_sendmsg+0x712/0x880 [1010.716297] ? [1010.716304] ? __ia32_sys_recvmmsg+0x270/0x270 [1010.716309] ? lock_acquire+0x1fe/0x560 [1010.71631 [1010.716324] ___sys_sendmsg+0xf8/0x170 [1010.716331] ? sendmsg_copy_msghdr+0x170/0x170 [1010.7163
CVE-2024-26801	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: Avoid potential use-after-free in hci_ HCI_EV_HARDWARE_ERROR event, if the underlying BT controller is not responding, the GPIO reset mechani use-after-free in hci_error_reset. Here's the call trace observed on a ChromeOS device with Intel AX201: queue_w +0x2ee/0x4c0 [bluetooth <HASH:3b4a6>] ? init_wait_entry+0x31/0x31 __hci_cmd_sync+0x16/0x20 [bluetooth < [bluetooth <HASH:3b4a 6>] process_one_work+0x1d8/0x33f worker_thread+0x21b/0x373 kthread+0x13a/0x152 +0x31/0x31 ret_from_fork+0x1f/0x30 This patch holds the reference count on the hci_dev while processing a HCI potential crash.

CVE-2024-26804	In the Linux kernel, the following vulnerability has been resolved: net: ip_tunnel: prevent perpetual headroom grov splat: BUG: KASAN: use-after-free in __skb_flow_dissect+0x19d1/0x7a50 net/core/flow_dissector.c:1170 Read o task syz-executor183/5191 [..] kasan_report+0xda/0x110 mm/kasan/report.c:588 __skb_flow_dissect+0x19d1/0x7 skb_flow_dissect_flow_keys include/linux/skbuff.h:1514 [inline] ___skb_get_hash net/core/flow_dissector.c:1791 core/flow_dissector.c:1856 skb_get_hash include/linux/skbuff.h:1556 [inline] ip_tunnel_xmit+0x1855/0x33c0 net/ +0x3cc/0x4e0 net/ipv4/ipip.c:308 __netdev_start_xmit include/linux/netdevice.h:4940 [inline] netdev_start_xmit i xmit_one net/core/dev.c:3548 [inline] dev_hard_start_xmit+0x13d/0x6d0 net/core/dev.c:3564 __dev_queue_xmit+ dev_queue_xmit include/linux/netdevice.h:3134 [inline] neigh_connected_output+0x42c/0x5d0 net/core/neighbour net/ipv4/ip_output.c:235 ip_finish_output+0x31/0x310 net/ipv4/ip_output.c:323 .. iptunnel_xmit+0x5b4/0x9b0 net +0x1dbc/0x33c0 net/ipv4/ip_tunnel.c:831 ipgre_xmit+0x4a1/0x980 net/ipv4/ip_gre.c:665 __netdev_start_xmit inc netdev_start_xmit include/linux/netdevice.h:4954 [inline] xmit_one net/core/dev.c:3548 [inline] dev_hard_start_xm The splat occurs because skb->data points past skb->head allocated area. This is because neigh layer does: __skb_p skb_network_offset() returns a negative offset and __skb_pull() arg is unsigned. IOW, we skb->data gets "adjusted is returned because skb->head and skb->data distance is more than 64k and skb->network_header (u16) has wrappe infrastructure, which can cause dev->needed_headroom to increment ad infinitum. The syzkaller reproducer consis and route of gre encapsulated packets pointing at another (ipip) tunnel. The ipip encapsulation finds gre0 as next ou pattern: 1). First packet is to be sent out via gre0. Route lookup found an output device, ipip0. 2). ip_tunnel_xmit fo based on the future output device, rt.dev->needed_headroom (ipip0). 3). ip output / start_xmit moves skb on to ipip (xmit recursion). 4). Routing step for the post-gre0-encap packet finds gre0 as output device to use for ipip0 encaps then incremented based on the (already bumped) gre0 device headroom. This repeats for every future packet: gre0- previous packets' ipip0 step incremented rt->dev (gre0) headroom, and ipip0 incremented because gre0 needed_hea packet, gre/ipip0->needed_headroom grows until post-expand-head reallocations result in a skb->head/data distanc skb->network_header (u16) wraps around when pskb_expand_head tries to make sure that skb_network_offset() is reallocation. After this skb_network_offset(skb) returns a different (and negative) result post headroom expansion. that would __skb_pull the network header) makes skb->data point to a memory location outside skb->head area. v2 arbitarily chosen upperlimit to prevent perpetual increase instead of dropping the headroom increment completely.
CVE-2024-26805	In the Linux kernel, the following vulnerability has been resolved: netlink: Fix kernel-infoleak-after-free in __skb_ following uninit-value access issue [1]: netlink_to_full_skb() creates a new `skb` and puts the `skb->data` passed a new `skb`. The data size is specified as `len` and passed to skb_put_data(). This `len` is based on `skb->end` that is `skb->end` contains data and tailroom. Since the tailroom is not initialized when the new `skb` created, KMSAN de copying the data. This patch resolved this issue by correct the len from `skb->end` to `skb->len`, which is the actua infoleak-after-free in instrument_copy_to_user include/linux/instrumented.h:114 [inline] BUG: KMSAN: kernel-in iov_iter.c:24 [inline] BUG: KMSAN: kernel-infoleak-after-free in iterate_ubuf include/linux/iov_iter.h:29 [inline] in iterate_and_advance2 include/linux/iov_iter.h:245 [inline] BUG: KMSAN: kernel-infoleak-after-free in iterate_ [inline] BUG: KMSAN: kernel-infoleak-after-free in _copy_to_iter+0x364/0x2520 lib/iov_iter.c:186 instrument_c [inline] copy_to_user_iter lib/iov_iter.c:24 [inline] iterate_ubuf include/linux/iov_iter.h:29 [inline] iterate_and_adv iterate_and_advance include/linux/iov_iter.h:271 [inline] _copy_to_iter+0x364/0x2520 lib/iov_iter.c:186 copy_to_ simple_copy_to_iter+0x68/0xa0 net/core/datagram.c:532 __skb_datagram_iter+0x123/0xdc0 net/core/datagram.c: net/core/datagram.c:546 skb_copy_datagram_msg include/linux/skbuff.h:3960 [inline] packet_recvmsg+0xd9c/0x2 sock_recvmsg_nosec net/socket.c:1044 [inline] sock_recvmsg net/socket.c:1066 [inline] sock_read_iter+0x467/0x linux/fs.h:2014 [inline] new_sync_read fs/read_write.c:389 [inline] vfs_read+0x8f6/0xe00 fs/read_write.c:470 ksys __do_sys_read fs/read_write.c:623 [inline] __se_sys_read fs/read_write.c:621 [inline] __x64_sys_read+0x93/0xd0 x86/entry/common.c:52 [inline] do_syscall_64+0x44/0x110 arch/x86/entry/common.c:83 entry_SYSCALL_64_af to memory at: skb_put_data include/linux/skbuff.h:2622 [inline] netlink_to_full_skb net/netlink/af_netlink.c:181 [i netlink/af_netlink.c:298 [inline] __netlink_deliver_tap+0x5be/0xc90 net/netlink/af_netlink.c:325 netlink_deliver_t netlink_deliver_tap_kernel net/netlink/af_netlink.c:347 [inline] netlink_unicast_kernel net/netlink/af_netlink.c:134 netlink/af_netlink.c:1368 netlink_sendmsg+0x1238/0x13d0 net/netlink/af_netlink.c:1910 sock_sendmsg_nosec net socket.c:745 [inline] ____sys_sendmsg+0x9c2/0xd60 net/socket.c:2584 ___sys_sendmsg+0x28d/0x3c0 net/socket [inline] __do_sys_sendmsg net/socket.c:2676 [inline] __se_sys_sendmsg net/socket.c:2674 [inline] __x64_sys_sen do_syscall_x64 arch/x86/entry/common.c:52 [inline] do_syscall_64+0x44/0x110 arch/x86/entry/common.c:83 entr Uninit was created at: free_pages_prepare mm/page_alloc.c:1087 [inline] free_unref_page_prepare+0xb0/0xa40 m +0xeb/0x1100 mm/page_alloc.c:2533 release_pages+0x23d3/0x2410 mm/swap.c:1042 free_pages_and_swap_cach tlb_batch_pages ---truncated---
CVE-2024-26810	In the Linux kernel, the following vulnerability has been resolved: vfio/pci: Lock external INTx masking ops Mask DisINTx may race INTx configuration changes via ioctl. Create wrappers that add locking for paths outside of the is updated holding igate, therefore testing is_intx() requires holding igate. For example clearing DisINTx from con interrupt configuration. This aligns interfaces which may trigger the INTx eventfd into two camps, one side serializ INTx is configured. A subsequent patch introduces synchronization for the latter flows.
CVE-2024-26813	In the Linux kernel, the following vulnerability has been resolved: vfio/platform: Create persistent IRQ handlers Th allows loopback triggering of an interrupt before a signaling eventfd has been configured by the user, which thereb than register the IRQ relative to a valid trigger, register all IRQs in a disabled state in the device open path. This all within the overall enable state governed by a valid eventfd signal. This decouples @masked, protected by the @loc the @igate mutex. In doing so, it's guaranteed that changes to @trigger cannot race the IRQ handlers because the I modifying the trigger, and loopback triggering of the IRQ via ioctl is safe due to serialization with trigger changes failures are maintained to be local to the SET_IRQS ioctl rather than a fatal error in the open device path. This allo polling mode support to continue to work regardless of moving the request_irq() call site. This necessarily blocks a

CVE-2024-26816	In the Linux kernel, the following vulnerability has been resolved: x86, relocs: Ignore relocations in .notes section CONFIG_XEN_PV=y, .text symbols are emitted into the .notes section so that Xen can find the "startup_xen" entr to booting the kernel, so relocations are not useful. In fact, performing relocations against the .notes section means sys/kernel/notes is world-readable. To avoid leaking the KASLR base without breaking unprivileged tools that are performing relocations in the .notes section. The values readable in .notes are then identical to those found in Syste
CVE-2024-26817	In the Linux kernel, the following vulnerability has been resolved: amdkfd: use calloc instead of kzalloc to avoid in doing the multiplication which might overflow.
CVE-2024-26820	In the Linux kernel, the following vulnerability has been resolved: hv_netvsc: Register VF in netvsc_probe if NET driver is unloaded and reloaded, the NET_DEVICE_REGISTER handler cannot perform VF register successfully a netvsc_probe is finished. This is because we register register_netdevice_notifier() very early(even before vmbus_c register each such matching VF(if it is visible as a netdevice) at the end of netvsc_probe.
CVE-2024-26820	In the Linux kernel, the following vulnerability has been resolved: hv_netvsc: Register VF in netvsc_probe if NET driver is unloaded and reloaded, the NET_DEVICE_REGISTER handler cannot perform VF register successfully a netvsc_probe is finished. This is because we register register_netdevice_notifier() very early(even before vmbus_c register each such matching VF(if it is visible as a netdevice) at the end of netvsc_probe.
CVE-2024-26825	In the Linux kernel, the following vulnerability has been resolved: nfc: nci: free rx_data_reassembly skb on NCI de is stored during NCI data exchange for processing fragmented packets. It is dropped only when the last fragment is NCI_OP_RF_DEACTIVATE_NTF opcode is received. However, the NCI device may be deallocated before that v rx_data_reassembly skb is bound to the NCI device and nothing prevents the device to be freed before the skb is pr the NCI device cleanup. Found by Linux Verification Center (linuxtesting.org) with Syzkaller.
CVE-2024-26826	In the Linux kernel, the following vulnerability has been resolved: mptcp: fix data re-injection from stale subflow V is stale, all the packet scheduler must re-inject all the mptcp-level unacked data. To avoid acquiring unneeded locks is present at all in the RTX queue, but such check is currently broken, as it uses TCP-specific helper on an MPTCP checkers are happy, as the accessed memory still belongs to the mptcp_sock struct, and even from a functional pers as the short-cut test always failed. A recent unrelated TCP change - commit d5fed5addb2b ("tcp: reorganize tcp_so issue, as the tcp field reorganization makes the mptcp code always skip the re-inection. Fix the issue dropping the b optimization proved once again to be evil.
CVE-2024-26831	In the Linux kernel, the following vulnerability has been resolved: net/handshake: Fix handshake_req_destroy_test started failing: Expected handshake_req_destroy_test == req, but handshake_req_destroy_test == 00000000000000 req_destroy works This is because "sock_release(sock)" was replaced with "fput(filp)" to address a memory leak. N fput() usually delays the final close and clean-up. The delay is not consequential in the other cases that were change that handshake_req_cancel() followed by closing the file actually does call the ->hp_destroy method. Thus the PTR final close is complete before it checks the pointer. We cannot use a completion here because if ->hp_destroy is ne test will hang. Reported by: Guenter Roeck <linux@roeck-us.net>
CVE-2024-26835	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_tables: set dormant flag on hook reg again if we fail to register the hooks. During memory pressure hook registration can fail and we end up with a table table/base chain deletion, nf_tables will attempt to unregister the hook again which yields a warn splat from the nft
CVE-2024-26839	In the Linux kernel, the following vulnerability has been resolved: IB/hfi1: Fix a memleak in init_credit_return Wh dd->cr_base[i].va, init_credit_return should deallocate dd->cr_base and dd->cr_base[i] that allocated before. Or the memleak is triggered.
CVE-2024-26840	In the Linux kernel, the following vulnerability has been resolved: cachefiles: fix memory leak in cachefiles_add_c was reported after unbinding /dev/cachefiles: == object 0xffff9b674176e3c0 (size 192): comm "cachefilesd2", pid 680, jiffies 4294881224 hex dump (first 32 bytes) 00 00 00 00 ..... 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ..... backtrace (crc ea38a44b): [<fff +0x2d5/0x370 [<ffffffff8e917f86>] prepare_creds+0x26/0x2e0 [<ffffffffc002eeef>] cachefiles_determine_cache_s cachefiles_add_cache+0x13c/0x3a0 [<ffffffffc0025216>] cachefiles_daemon_write+0x146/0x1c0 [<ffffffff8ebc4a [<ffffffff8ebc5069>] ksys_write+0x69/0xf0 [<ffffffff8f6d4662>] do_syscall_64+0x72/0x140 [<ffffffff8f8000aa>] +0x6e/0x76 ========== Put the r cachefiles_daemon_unbind() to fix the problem. And also put cache_cred in cachefiles_add_cache() error branch to
CVE-2024-26848	In the Linux kernel, the following vulnerability has been resolved: afs: Fix endless loop in directory parsing If a dir files in it (from uncompleted silly-rename), these .__afsXXXX files are skipped but without advancing the file posi afs_dir_iterate() repeating the block again and again. Fix this by making the code that skips the .__afsXXXX file al The symptoms are a soft lookup: watchdog: BUG: soft lockup - CPU#3 stuck for 52s! [check:5737] ... RIP: 0010:a watchdog_timer_fn+0x1a6/0x213 ... ? asm_sysvec_apic_timer_interrupt+0x16/0x20 ? afs_dir_iterate_block+0x39 afs_readdir+0x30/0x4a iterate_dir+0x93/0xd3 __do_sys_getdents64+0x6b/0xd4 This is almost certainly the actual show_bug.cgi?id=218496

CVE-2024-26851	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_conntrack_h323: Add protection fo load reports an exception of BRK#5515 SHIFT_ISSUE:Bitwise shifts that are out of bounds for their data type. vm <net/netfilter/nf_conntrack_h323_asn1.c:0> vmlinux decode_seq(bs=0xFFFFFFD008037000, f=0xFFFFFFD0080 1956 <net/netfilter/nf_conntrack_h323_asn1.c:592> vmlinux decode_choice(base=0xFFFFFFD0080370F0, level= nf_conntrack_h323_asn1.c:814> vmlinux decode_seq(f=0xFFFFFFD0080371A8, level=134443500) + 812 <net/n vmlinux decode_choice(base=0xFFFFFFD008037280, level=0) + 1216 <net/netfilter/nf_conntrack_h323_asn1.c:8 <net/netfilter/nf_conntrack_h323_asn1.c:833> vmlinux ras_help() + 684 <net/netfilter/nf_conntrack_h323_main.c: netfilter/nf_conntrack_proto.c:137> Due to abnormal data in skb->data, the extension bitmap length exceeds 32 wh length to make a shift operation. It will change into negative after several loop. UBSAN load could detect a negativ exception. So we add the protection to avoid the length exceeding 32. Or else it will return out of range error and st
CVE-2024-26852	In the Linux kernel, the following vulnerability has been resolved: net/ipv6: avoid possible UAF in ip6_route_mpa use-after-free in ip6_route_mpath_notify() [1] Commit f7225172f25a ("net/ipv6: prevent use after free in ip6_route the root cause. We need to defer the fib6_info_release() calls after ip6_route_mpath_notify(), in the cleanup phase. free in rt6_fill_node+0x1460/0x1ac0 Read of size 4 at addr ffff88809a07fc64 by task syz-executor.2/23037 CPU: Not tainted 6.8.0-rc4-syzkaller-01035-gea7f3cfaa588 #0 Hardware name: Google Google Compute Engine/Google 01/25/2024 Call Trace: <TASK> __dump_stack lib/dump_stack.c:88 [inline] dump_stack_lvl+0x1e7/0x2e0 lib/du mm/kasan/report.c:377 [inline] print_report+0x167/0x540 mm/kasan/report.c:488 kasan_report+0x142/0x180 mm +0x1460/0x1ac0 inet6_rt_notify+0x13b/0x290 net/ipv6/route.c:6184 ip6_route_mpath_notify net/ipv6/route.c:519 net/ipv6/route.c:5404 [inline] inet6_rtm_newroute+0x1d0f/0x2300 net/ipv6/route.c:5517 rtnetlink_rcv_msg+0x885 netlink_rcv_skb+0x1e3/0x430 net/netlink/af_netlink.c:2543 netlink_unicast_kernel net/netlink/af_netlink.c:1341 [netlink/af_netlink.c:1367 netlink_sendmsg+0xa3b/0xd70 net/netlink/af_netlink.c:1908 sock_sendmsg_nosec net/so +0x221/0x270 net/socket.c:745 ____sys_sendmsg+0x525/0x7d0 net/socket.c:2584 ___sys_sendmsg net/socket.c:2 net/socket.c:2667 do_syscall_64+0xf9/0x240 entry_SYSCALL_64_after_hwframe+0x6f/0x77 RIP: 0033:0x7f73d 83 c4 28 c3 e8 e1 20 00 00 90 48 89 f8 48 89 f7 48 89 d6 48 89 ca 4d 89 c2 4d 89 c8 4c 8b 4c 24 08 0f 05 <48> 3d f7 d8 64 89 01 48 RSP: 002b:00007f73de6550c8 EFLAGS: 00000246 ORIG_RAX: 000000000000002e RAX: fff RCX: 00007f73dd87dda9 RDX: 0000000000000000 RSI: 0000000020000140 RDI: 0000000000000005 RBP: 000 R09: 0000000000000000 R10: 0000000000000000 R11: 0000000000000246 R12: 0000000000000000 R13: 0000 R15: 00007ffdbdeb7858 </TASK> Allocated by task 23037: kasan_save_stack mm/kasan/common.c:47 [inline] ka common.c:68 poison_kmalloc_redzone mm/kasan/common.c:372 [inline] __kasan_kmalloc+0x98/0xb0 mm/kasan linux/kasan.h:211 [inline] __do_kmalloc_node mm/slub.c:3981 [inline] __kmalloc+0x22e/0x490 mm/slub.c:3994 kzalloc include/linux/slab.h:711 [inline] fib6_info_alloc+0x2e/0xf0 net/ipv6/ip6_fib.c:155 ip6_route_info_create+ ip6_route_multipath_add net/ipv6/route.c:5298 [inline] inet6_rtm_newroute+0x744/0x2300 net/ipv6/route.c:5517 rtnetlink.c:6597 netlink_rcv_skb+0x1e3/0x430 net/netlink/af_netlink.c:2543 netlink_unicast_kernel net/netlink/af_ +0x7ea/0x980 net/netlink/af_netlink.c:1367 netlink_sendmsg+0xa3b/0xd70 net/netlink/af_netlink.c:1908 sock_sen __sock_sendmsg+0x221/0x270 net/socket.c:745 ____sys_sendmsg+0x525/0x7d0 net/socket.c:2584 ___sys_sendm +0x2b0/0x3a0 net/socket.c:2667 do_syscall_64+0xf9/0x240 entry_SYSCALL_64_after_hwframe+0x6f/0x77 Free common.c:47 [inline] kasan_save_track+0x3f/0x80 mm/kasan/common.c:68 kasan_save_free_info+0x4e/0x60 mn +0xa6/0xe0 m ---truncated---
CVE-2024-26854	In the Linux kernel, the following vulnerability has been resolved: ice: fix uninitialized dplls mutex usage The pf-> after its first use. Move it to the top of ice_dpll_init. Note that the "err_exit" error path destroys the mutex. And the ice_dpll_deinit. This fixes the following warning with CONFIG_DEBUG_MUTEXES: ice 0000:10:00.0: The DDI ICE OS Default Package version 1.3.36.0 ice 0000:10:00.0: 252.048 Gb/s available PCIe bandwidth (16.0 GT/s PC init successful ------------[cut here]------------ DEBUG_LOCKS_WARN_ON(lock->magic != lock) WARNING: mutex.c:587 __mutex_lock+0x773/0xd40 Modules linked in: crct10dif_pclmul crc32_pclmul crc32c_intel polyval nvme_c> CPU: 0 PID: 410 Comm: kworker/0:4 Not tainted 6.8.0-rc5+ #3 Hardware name: HPE ProLiant DL110 Plus, BIOS U56 10/19/2023 Workqueue: events work_for_cpu_fn RIP: 0010:__mutex_lock+0x773/0xd40 Code: c 01 45 85 f6 0f 85 0d f9 ff ff 48 c7 c6 12 a2 a9 85 48 c7 c7 12 f1 a> RSP: 0018:ff7eb1a3417a7ae0 EFLAGS: 00010 0000000000000002 RCX: 0000000000000000 RDX: 0000000000000002 RSI: ffffffff85ac2bff RDI: 00000000fff 0000000000000000 R09: 00000000ffffbfff R10: ff7eb1a3417a7978 R11: ff32b80f7fd2e568 R12: 0000000000000 0000000000000000 R15: ff32b7f02c50e0d8 FS: 0000000000000000(0000) GS:ff32b80efe800000(0000) knlGS:00 0000 ES: 0000 CR0: 0000000080050033 CR2: 000055b5852cc000 CR3: 000000003c43a004 CR4: 000000000077 0000000000000000 DR2: 0000000000000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 000000000 <TASK> ? __warn+0x84/0x170 ? __mutex_lock+0x773/0xd40 ? report_bug+0x1c7/0x1d0 ? prb_read_valid+0x1b exc_invalid_op+0x18/0x70 ? asm_exc_invalid_op+0x1a/0x20 ? __mutex_lock+0x773/0xd40 ? rcu_is_watching+0 +0x346/0x490 ? ice_dpll_lock_status_get+0x28/0x50 [ice] ? __pfx_ice_dpll_lock_status_get+0x10/0x10 [ice] ? ic ice_dpll_lock_status_get+0x28/0x50 [ice] dpll_device_get_one+0x14f/0x2e0 dpll_device_event_send+0x7d/0x150 ice_dpll_init_dpll+0x7b/0xd0 [ice] ice_dpll_init+0x224/0xa40 [ice] ? _dev_info+0x70/0x90 ice_load+0x468/0x69 _raw_spin_unlock_irqrestore+0x4f/0x80 ? process_one_work+0x1a3/0x500 local_pci_probe+0x47/0xa0 work_for +0x20d/0x500 worker_thread+0x1df/0x3e0 ? __pfx_worker_thread+0x10/0x10 kthread+0x103/0x140 ? __pfx_kth __pfx_kthread+0x10/0x10 ret_from_fork_asm+0x1b/0x30 </TASK> irq event stamp: 125197 hardirqs last enable finish_task_switch.isra.0+0x12d/0x3d0 hardirqs last disabled at (125196): [<ffffffff85134044>] __schedule+0xea4 [<ffffffff84e1e65a>] napi_get_frags_check+0x1a/0x60 softirqs last disabled at (105332): [<ffffffff84e1e65a>] nap 0000000000000000]---
CVE-2024-26855	In the Linux kernel, the following vulnerability has been resolved: net: ice: Fix potential NULL pointer dereference ice_bridge_setlink() may encounter a NULL pointer dereference if nlmsg_find_attr() returns NULL and br_spec is nla_for_each_nested(). To address this issue, add a check to ensure that br_spec is not NULL before proceeding wi

CVE-2024-26859	In the Linux kernel, the following vulnerability has been resolved: net/bnx2x: Prevent access to a freed page in pag crash during EEH error handling During EEH error recovery, the bnx2x driver's transmit timeout logic could cause The bnx2x_tx_timeout() schedules reset tasks via bnx2x_sp_rtnl_task(), which ultimately leads to bnx2x_nic_unlo freed using bnx2x_free_rx_sge_range(). However, this could overlap with the EEH driver's attempt to reset the dev also tries to free SGEs. This race condition can result in system crashes due to accessing freed memory locations in void bnx2x_free_rx_sge(struct bnx2x *bp, 800 struct bnx2x_fastpath *fp, u16 index) 801 { 802 struct sw_rx_page 803 struct page *page = sw_buf->page; .... where sw_buf was set to NULL after the call to dma_unmap_page() by 'slot_reset' PCI 0011:01:00.0#10000: EEH: Invoking bnx2x->slot_reset() bnx2x: [bnx2x_io_slot_reset:14228(eth1) 0011:01:00.0: enabling device (0140 -> 0142) bnx2x: [bnx2x_io_slot_reset:14244(eth1)]IO slot reset --> driver un page (0) - exploit attempt? (uid: 0) BUG: Kernel NULL pointer dereference on read at 0x00000000 Faulting instru Oops: Kernel access of bad area, sig: 11 [#1] ..... Call Trace: [c000000003c67a20] [c00800000250658c] bnx2x_io_ (unreliable) [c000000003c67af0] [c0000000000518a8] eeh_report_reset+0xb8/0xf0 [c000000003c67b60] [c00000 [c000000003c67c70] [c00000000005318c] eeh_handle_normal_event+0x84c/0xa60 [c000000003c67d50] [c00000 [c000000003c67da0] [c000000000194c58] kthread+0x1c8/0x1d0 [c000000003c67e10] [c00000000000cf64] ret_f issue, we need to verify page pool allocations before freeing.
CVE-2024-26859	In the Linux kernel, the following vulnerability has been resolved: net/bnx2x: Prevent access to a freed page in pag crash during EEH error handling During EEH error recovery, the bnx2x driver's transmit timeout logic could cause The bnx2x_tx_timeout() schedules reset tasks via bnx2x_sp_rtnl_task(), which ultimately leads to bnx2x_nic_unlo freed using bnx2x_free_rx_sge_range(). However, this could overlap with the EEH driver's attempt to reset the dev also tries to free SGEs. This race condition can result in system crashes due to accessing freed memory locations in void bnx2x_free_rx_sge(struct bnx2x *bp, 800 struct bnx2x_fastpath *fp, u16 index) 801 { 802 struct sw_rx_page 803 struct page *page = sw_buf->page; .... where sw_buf was set to NULL after the call to dma_unmap_page() by 'slot_reset' PCI 0011:01:00.0#10000: EEH: Invoking bnx2x->slot_reset() bnx2x: [bnx2x_io_slot_reset:14228(eth1) 0011:01:00.0: enabling device (0140 -> 0142) bnx2x: [bnx2x_io_slot_reset:14244(eth1)]IO slot reset --> driver un page (0) - exploit attempt? (uid: 0) BUG: Kernel NULL pointer dereference on read at 0x00000000 Faulting instru Oops: Kernel access of bad area, sig: 11 [#1] ..... Call Trace: [c000000003c67a20] [c00800000250658c] bnx2x_io_ (unreliable) [c000000003c67af0] [c0000000000518a8] eeh_report_reset+0xb8/0xf0 [c000000003c67b60] [c00000 [c000000003c67c70] [c00000000005318c] eeh_handle_normal_event+0x84c/0xa60 [c000000003c67d50] [c00000 [c000000003c67da0] [c000000000194c58] kthread+0x1c8/0x1d0 [c000000003c67e10] [c00000000000cf64] ret_f issue, we need to verify page pool allocations before freeing.
CVE-2024-26862	In the Linux kernel, the following vulnerability has been resolved: packet: annotate data-races around ignore_outgo from dev_queue_xmit_nit() and packet_getsockopt() Add appropriate READ_ONCE()/WRITE_ONCE() annotatio data-race in dev_queue_xmit_nit / packet_setsockopt write to 0xffff888107804542 of 1 bytes by task 22618 on cpu net/packet/af_packet.c:4003 do_sock_setsockopt net/socket.c:2311 [inline] __sys_setsockopt+0x1d8/0x250 net/soc socket.c:2343 [inline] __se_sys_setsockopt net/socket.c:2340 [inline] __x64_sys_setsockopt+0x66/0x80 net/socke entry_SYSCALL_64_after_hwframe+0x6d/0x75 read to 0xffff888107804542 of 1 bytes by task 27 on cpu 1: dev_ dev.c:2248 xmit_one net/core/dev.c:3527 [inline] dev_hard_start_xmit+0xcc/0x3f0 net/core/dev.c:3547 __dev_que dev_queue_xmit include/linux/netdevice.h:3091 [inline] batadv_send_skb_packet+0x264/0x300 net/batman-adv/se +0x24/0x30 net/batman-adv/send.c:127 batadv_iv_ogm_send_to_if net/batman-adv/bat_iv_ogm.c:392 [inline] bata bat_iv_ogm.c:420 [inline] batadv_iv_send_outstanding_bat_ogm_packet+0x3f0/0x4b0 net/batman-adv/bat_iv_ogm workqueue.c:3254 [inline] process_scheduled_works+0x465/0x990 kernel/workqueue.c:3335 worker_thread+0x52 kthread+0x1d1/0x210 kernel/kthread.c:388 ret_from_fork+0x4b/0x60 arch/x86/kernel/process.c:147 ret_from_fork entry_64.S:243 value changed: 0x00 -> 0x01 Reported by Kernel Concurrency Sanitizer on: CPU: 1 PID: 27 Comm syzkaller-08073-g480e035fc4c7 #0 Hardware name: Google Google Compute Engine/Google Compute Engine, B bat_events batadv_iv_send_outstanding_bat_ogm_packet
CVE-2024-26863	In the Linux kernel, the following vulnerability has been resolved: hsr: Fix uninit-value access in hsr_get_node() K uninit-value access issue [1]: == BUG: KM +0xa2e/0xa40 net/hsr/hsr_framereg.c:246 hsr_get_node+0xa2e/0xa40 net/hsr/hsr_framereg.c:246 fill_frame_info r hsr_forward_skb+0xe12/0x30e0 net/hsr/hsr_forward.c:615 hsr_dev_xmit+0x1a1/0x270 net/hsr/hsr_device.c:223 _ netdevice.h:4940 [inline] netdev_start_xmit include/linux/netdevice.h:4954 [inline] xmit_one net/core/dev.c:3548 [net/core/dev.c:3564 __dev_queue_xmit+0x33b8/0x5130 net/core/dev.c:4349 dev_queue_xmit include/linux/netdev +0x9c/0x6b0 net/packet/af_packet.c:276 packet_snd net/packet/af_packet.c:3087 [inline] packet_sendmsg+0x8b1c sock_sendmsg_nosec net/socket.c:730 [inline] __sock_sendmsg net/socket.c:745 [inline] __sys_sendto+0x735/0xa net/socket.c:2203 [inline] __se_sys_sendto net/socket.c:2199 [inline] __x64_sys_sendto+0x125/0x1c0 net/socket.c common.c:52 [inline] do_syscall_64+0x6d/0x140 arch/x86/entry/common.c:83 entry_SYSCALL_64_after_hwfram slab_post_alloc_hook+0x129/0xa70 mm/slab.h:768 slab_alloc_node mm/slub.c:3478 [inline] kmem_cache_alloc_ kmalloc_reserve+0x13d/0x4a0 net/core/skbuff.c:560 __alloc_skb+0x318/0x740 net/core/skbuff.c:651 alloc_skb in alloc_skb_with_frags+0xc8/0xbd0 net/core/skbuff.c:6334 sock_alloc_send_pskb+0xa80/0xbf0 net/core/sock.c:278 af_packet.c:2936 [inline] packet_snd net/packet/af_packet.c:3030 [inline] packet_sendmsg+0x70e8/0x9f30 net/pac net/socket.c:730 [inline] __sock_sendmsg net/socket.c:745 [inline] __sys_sendto+0x735/0xa10 net/socket.c:2191 [inline] __se_sys_sendto net/socket.c:2199 [inline] __x64_sys_sendto+0x125/0x1c0 net/socket.c:2199 do_syscall_ do_syscall_64+0x6d/0x140 arch/x86/entry/common.c:83 entry_SYSCALL_64_after_hwframe+0x63/0x6b CPU: 1 tainted 6.7.0-syzkaller-00562-g9f8413c4a66f #0 Hardware name: Google Google Compute Engine/Google Compu == If the packet type ID field in the Ethern ETH_P_HSR, but it is not followed by an HSR tag, hsr_get_skb_sequence_nr() reads an invalid value as a sequenc patch fixes the issue by returning NULL if the Ethernet header is not followed by an HSR tag.

CVE-2024-26874	In the Linux kernel, the following vulnerability has been resolved: drm/mediatek: Fix a null pointer crash in mtk_d mtk_crtc->event is NULL in mtk_drm_crtc_finish_page_flip(). pending_needs_vblank value is set by mtk_crtc->e it's is not guarded by the same lock in mtk_drm_finish_page_flip(), thus a race condition happens. Consider the fol mtk_drm_crtc_atomic_begin() mtk_crtc->event is not null, step 1: mtk_drm_crtc_atomic_flush: mtk_drm_crtc_up mtk_crtc_ddp_irq -> mtk_drm_finish_page_flip: lock mtk_crtc->event set to null, pending_needs_vblank set to fal true, step 2: mtk_crtc_ddp_irq -> mtk_drm_finish_page_flip called again, pending_needs_vblank is still true //null mtk_drm_crtc_atomic_flush(), it's more efficient to just check if mtk_crtc->event is null before use.
CVE-2024-26875	In the Linux kernel, the following vulnerability has been resolved: media: pvrusb2: fix uaf in pvr2_context_set_not slab-use-after-free in pvr2_context_set_notify+0x2c4/0x310 drivers/media/usb/pvrusb2/pvrusb2-context.c:35 Read by task kworker/1:1/26 CPU: 1 PID: 26 Comm: kworker/1:1 Not tainted 6.8.0-rc1-syzkaller-00046-gf1a27f081c1f Compute Engine/Google Compute Engine, BIOS Google 01/25/2024 Workqueue: usb_hub_wq hub_event Call Tra dump_stack.c:88 [inline] dump_stack_lvl+0xd9/0x1b0 lib/dump_stack.c:106 print_address_description mm/kasan/ mm/kasan/report.c:488 kasan_report+0xda/0x110 mm/kasan/report.c:601 pvr2_context_set_notify+0x2c4/0x310 d context.c:35 pvr2_context_notify drivers/media/usb/pvrusb2/pvrusb2-context.c:95 [inline] pvr2_context_disconnec pvrusb2-context.c:272 Freed by task 906: kasan_save_stack+0x33/0x50 mm/kasan/common.c:47 kasan_save_track kasan_save_free_info+0x3f/0x60 mm/kasan/generic.c:640 poison_slab_object mm/kasan/common.c:241 [inline] _ kasan/common.c:257 kasan_slab_free include/linux/kasan.h:184 [inline] slab_free_hook mm/slub.c:2121 [inline] s +0x105/0x340 mm/slub.c:4409 pvr2_context_check drivers/media/usb/pvrusb2/pvrusb2-context.c:137 [inline] pvr media/usb/pvrusb2/pvrusb2-context.c:158 [Analyze] Task A set disconnect_flag = !0, which resulted in Task B's c to this issue. [Fix] Place the disconnect_flag assignment operation after all code in pvr2_context_disconnect() to av
CVE-2024-26878	In the Linux kernel, the following vulnerability has been resolved: quota: Fix potential NULL pointer dereference I dereference P1 P2 dquot_free_inode quota_off drop_dquot_ref remove_dquot_ref dquots = i_dquot(inode) dquots dquots[cnt]) != NULL (1) dquots[type] = NULL (2) spin_lock(&dquots[cnt]->dq_dqb_lock) (3) .... If dquot_free_i pointers (1) before quota_off sets it to NULL(2) and use it (3) after that, NULL pointer dereference will be triggere to avoid this issue.
CVE-2024-26880	In the Linux kernel, the following vulnerability has been resolved: dm: call the resume method on internal suspend experimenting with the lvm2 testsuite. The list corruption is caused by the fact that the postsuspend and resume me were two consecutive calls to the origin_postsuspend function. The second call attempts to remove the "hash_list" removed by the first call. Fix __dm_internal_resume so that it calls the preresume and resume methods of the table some target fails, we are in a tricky situation. We can't return an error because dm_internal_resume isn't supposed t because then the "resume" and "postsuspend" methods would not be paired correctly. So, we set the DMF_SUSPE - it may confuse userspace tools, but it won't cause a kernel crash. ----------[cut here]---------- kernel BUG at lib 0000 [#1] PREEMPT SMP CPU: 1 PID: 8343 Comm: dmsetup Not tainted 6.8.0-rc6 #4 Hardware name: QEMU S BIOS 1.14.0-2 04/01/2014 RIP: 0010:__list_del_entry_valid_or_report+0x77/0xc0 <snip> RSP: 0018:ffff8881b83 000000000000004e RBX: ffff888143b6eb80 RCX: 0000000000000000 RDX: 0000000000000001 RSI: ffffffff81 ffff8881b83a3400 R08: 00000000fffeffff R09: 0000000000000058 R10: 0000000000000000 R11: ffffffff81a2408 ffff88814538e000 R14: ffff888143bc6dc0 R15: fffffffa02e4bb0 FS: 00000000f7c0f780(0000) GS:ffff8893f0a400 0010 DS: 002b ES: 002b CR0: 0000000080050033 CR2: 0000000057fb5000 CR3: 0000000143474000 CR4: 0000 +0x2d/0x80 ? do_trap+0xeb/0xf0 ? __list_del_entry_valid_or_report+0x77/0xc0 ? do_error_trap+0x60/0x80 ? __l exc_invalid_op+0x49/0x60 ? __list_del_entry_valid_or_report+0x77/0xc0 ? asm_exc_invalid_op+0x16/0x20 ? tab __list_del_entry_valid_or_report+0x77/0xc0 origin_postsuspend+0x1a/0x50 [dm_snapshot] dm_table_postsuspend +0xd8/0xf0 [dm_mod] dev_suspend+0x1f2/0x2f0 [dm_mod] ? table_deps+0x1b0/0x1b0 [dm_mod] ctl_ioctl+0x30 +0x7/0x10 [dm_mod] __x64_compat_sys_ioctl+0x104/0x170 do_syscall_64+0x184/0x1b0 entry_SYSCALL_64_ 0033:0xf7e6aead <snip> ---[end trace 0000000000000000]---
CVE-2024-26884	In the Linux kernel, the following vulnerability has been resolved: bpf: Fix hashtab overflow check on 32-bit arche roundup_pow_of_two() to compute the number of hash buckets, and contains an overflow check by checking if the arches, the roundup code itself can overflow by doing a 32-bit left-shift of an unsigned long value, which is undefi truncate neatly. This was triggered by syzbot on the DEVMAP_HASH type, which contains the same check, copie fix to hashtab, by moving the overflow check to before the roundup.
CVE-2024-26889	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: hci_core: Fix possible buffer overflo name[8] field so in the event that hdev->name is bigger than that strcpy would attempt to write past its size, so this strscpy.
CVE-2024-26894	In the Linux kernel, the following vulnerability has been resolved: ACPI: processor_idle: Fix memory leak in acpi_ the CPU idle device, the memory associated with it is not freed, leading to a memory leak: unreferenced object 0xf "swapper/0", pid 1, jiffies 4294893170 hex dump (first 32 bytes): 00 00 00 00 0b 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ..... backtrace (crc 8836a742): [<ffffffff993495ed>] kmalloc_trace+0x29d/0x340 [<ffffffff +0xf3/0x1c0 [<ffffffff9972d263>] __acpi_processor_start+0xd3/0xf0 [<ffffffff9972d2bc>] acpi_processor_start+0 +0xe2/0x480 [<ffffffff99805c98>] __driver_probe_device+0x78/0x160 [<ffffffff99805daf>] driver_probe_device+ __driver_attach+0xce/0x1c0 [<ffffffff99803170>] bus_for_each_dev+0x70/0xc0 [<ffffffff99804822>] bus_add_dr driver_register+0x55/0x100 [<ffffffff9aee4acb>] acpi_processor_driver_init+0x3b/0xc0 [<ffffffff990012d1>] do_ kernel_init_freeable+0x320/0x470 [<ffffffff99b231f6>] kernel_init+0x16/0x1b0 [<ffffffff99042e6d>] ret_from_fo idle device after unregistering it.

CVE-2024-26898	In the Linux kernel, the following vulnerability has been resolved: aoe: fix the potential use-after-free problem in a CVE-2023-6270. The description of cve is: A flaw was found in the ATA over Ethernet (AoE) driver in the Linux improperly updates the refcnt on `struct net_device`, and a use-after-free can be triggered by racing between the fre `skbtxq` global queue. This could lead to a denial of service condition or potential code execution. In aoecmd_cfg_ skb initial code is finished. But the net_device ifp will still be used in later tx()->dev_queue_xmit() in kthread. Whi be called in the success path of skb initial code in aoecmd_cfg_pkts(). Otherwise tx() may run into use-after-free be removed the dev_put(ifp) in the success path in aoecmd_cfg_pkts(), and added dev_put() after skb xmit in tx().
CVE-2024-26901	In the Linux kernel, the following vulnerability has been resolved: do_sys_name_to_handle(): use kzalloc() to fix k a kernel information leak vulnerability in do_sys_name_to_handle() and issued the following report [1]. [1] "BUG: instrument_copy_to_user include/linux/instrumented.h:114 [inline] BUG: KMSAN: kernel-infoleak in _copy_to_u instrument_copy_to_user include/linux/instrumented.h:114 [inline] _copy_to_user+0xbc/0x100 lib/usercopy.c:40 c [inline] do_sys_name_to_handle fs/fhandle.c:73 [inline] __do_sys_name_to_handle_at fs/fhandle.c:112 [inline] __ fs/fhandle.c:94 __x64_sys_name_to_handle_at+0xe4/0x140 fs/fhandle.c:94 ... Uninit was created at: slab_post_allo slab_alloc_node mm/slub.c:3478 [inline] __kmem_cache_alloc_node+0x5c9/0x970 mm/slub.c:3517 __do_kmalloc [inline] __kmalloc+0x121/0x3c0 mm/slab_common.c:1020 kmalloc include/linux/slab.h:604 [inline] do_sys_name __do_sys_name_to_handle_at fs/fhandle.c:112 [inline] __se_sys_name_to_handle_at+0x441/0xb10 fs/fhandle.c:94 fs/fhandle.c:94 ... Bytes 18-19 of 20 are uninitialized Memory access of size 20 starts at ffff888128a46380 Data co Chuck Lever's suggestion, use kzalloc() instead of kmalloc() to solve the problem.
CVE-2024-26907	In the Linux kernel, the following vulnerability has been resolved: RDMA/mlx5: Fix fortify source warning while a [cut here]------------ memcpy: detected field-spanning write (size 56) of single field "eseg->inline_hdr.start" at /va build/drivers/infiniband/hw/mlx5/wr.c:131 (size 2) WARNING: CPU: 0 PID: 293779 at /var/lib/dkms/mlnx-ofed-k hw/mlx5/wr.c:131 mlx5_ib_post_send+0x191b/0x1a60 [mlx5_ib] Modules linked in: 8021q garp mrp stp llc rdma_ ib_ipoib(OE) ib_cm(OE) ib_umad(OE) mlx5_ib(OE) ib_uverbs(OE) ib_core(OE) mlx5_core(OE) pci_hyperv_intf mlxfw(OE) psample nft_fib_inet nft_fib_ipv4 nft_fib_ipv6 nft_fib nft_reject_inet nf_reject_ipv4 nf_reject_ipv6 nf nf_conntrack nf_defrag_ipv6 nf_defrag_ipv4 ip_set nf_tables libcrc32c nfnetlink mst_pciconf(OE) knem(OE) vfio vfio iommufd irqbypass cuse nfsv3 nfs fscache netfs xfrm_user xfrm_algo ipmi_devintf ipmi_msghandler binfmt_ polyval_clmulni polyval_generic ghash_clmulni_intel sha512_ssse3 snd_pcsp aesni_intel crypto_simd cryptd snd_ input_leds serio_raw evbug nfsd auth_rpcgss nfs_acl lockd grace sch_fq_codel sunrpc drm efi_pstore ip_tables x_t net_failover failover floppy [last unloaded: mlx_compat(OE)] CPU: 0 PID: 293779 Comm: ssh Tainted: G OE 6.2. Hardware name: Red Hat KVM, BIOS 0.5.1 01/01/2011 RIP: 0010:mlx5_ib_post_send+0x191b/0x1a60 [mlx5_ib] a0 b9 02 00 00 00 48 c7 c2 10 5b fd c0 48 c7 c7 80 5b fd c0 c6 05 57 0c 03 00 01 e8 95 4d 93 da <0f> 0b 44 8b 4d ff ff 41 0f b7 RSP: 0018:ffffb5b48478b570 EFLAGS: 00010046 RAX: 0000000000000000 RBX: 0000000000000 0000000000000000 RSI: 0000000000000000 RDI: 0000000000000000 RBP: ffffb5b48478b628 R08: 0000000000 R10: 0000000000000000 R11: 0000000000000000 R12: ffffb5b48478b5e8 R13: ffff963a3c609b5e R14: ffff9639c FS: 00007fc03b444c80(0000) GS:ffff963a3dc00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 000 0000556f46bdf000 CR3: 0000000006ac6003 CR4: 0000000003706f0 DR0: 0000000000000000 DR1: 0000000 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 Call Trace: <TASK> ? show_regs+0x72/0x [mlx5_ib] ? __warn+0x8d/0x160 ? mlx5_ib_post_send+0x191b/0x1a60 [mlx5_ib] ? report_bug+0x1bb/0x1d0 ? ha +0x19/0x80 ? asm_exc_invalid_op+0x1b/0x20 ? mlx5_ib_post_send+0x191b/0x1a60 [mlx5_ib] mlx5_ib_post_se +0x2ec/0x770 [ib_ipoib] ipoib_start_xmit+0x5a0/0x770 [ib_ipoib] dev_hard_start_xmit+0x8e/0x1e0 ? validate_x +0x116/0x3a0 __dev_xmit_skb+0x1fd/0x580 __dev_queue_xmit+0x284/0x6b0 ? _raw_spin_unlock_irq+0xe/0x5 push_pseudo_header+0x17/0x40 [ib_ipoib] neigh_connected_output+0xcd/0x110 ip_finish_output2+0x179/0x480 __ip_finish_output+0xc3/0x190 ip_finish_output+0x2e/0xf0 ip_output+0x78/0x110 ? __pfx_ip_finish_output+0x1 __ip_queue_xmit+0x18a/0x460 ip_queue_xmit+0x15/0x30 __tcp_transmit_skb+0x914/0x9c0 tcp_write_xmit+0x3 tcp_sendmsg_locked+0x2e1/0xac0 tcp_sendmsg+0x2d/0x50 inet_sendmsg+0x43/0x90 sock_sendmsg+0x68/0x80 +0x326/0x3c0 ksys_write+0xbd/0xf0 ? do_syscall_64+0x69/0x90 __x64_sys_write+0x19/0x30 do_syscall_ ---tru
CVE-2024-26910	In the Linux kernel, the following vulnerability has been resolved: netfilter: ipset: fix performance regression in sw fix race condition between swap/destroy and kernel side add/del/test", commit 28628fa9 fixes a race condition. But function unnecessarily slows it down: it can safely be moved to destroy and use call_rcu() instead. Eric Dumazet p functions as rcu callback does not work: sets with timeout use garbage collectors which need cancelling at destroy functions are split into two: cancelling garbage collectors safely at executing the command received by netlink and callback.
CVE-2024-26917	In the Linux kernel, the following vulnerability has been resolved: scsi: Revert "scsi: fcoe: Fix potential deadlock c 1a1975551943f681772720f639ff42fbaa746212. This commit causes interrupts to be lost for FCoE devices, since it Instead, a work queue should be used, and will be addressed in a separate commit.
CVE-2024-26918	In the Linux kernel, the following vulnerability has been resolved: PCI: Fix active state requirement in PME pollin requirement that runtime PM managed devices need to be in the RPM_ACTIVE state for PME polling. In fact, onl polled. However there's still a requirement that the device config space must be accessible, which has implications and the parent bridge, when present. It's not sufficient to assume the bridge remains in D0 and cases have been obs the PM state indicates RPM_SUSPENDING and config space of the polled device becomes inaccessible during pc is already effectively required to be in the RPM_ACTIVE state, formalize this in the code and elevate the PM usag subordinate device. This resolves a regression reported in the bugzilla below where a Thunderbolt/USB4 hierarchy downstream of a bridge in a D3hot power state.
CVE-2024-26920	In the Linux kernel, the following vulnerability has been resolved: tracing/trigger: Fix to return error if failed to all return error code if it failed to allocate a snapshot instead of 0 (success). Unless that, it will register snapshot trigge

CVE-2024-26921	In the Linux kernel, the following vulnerability has been resolved: inet: inet_defrag: prevent sk release while still i pass skb->sk as function argument. If the skb is a fragment and reassembly happens before such function call retur skb fragments reassembled via netfilter or similar modules, e.g. openvswitch or ct_act.c, when run as part of tx pip of this bug. Quoting Eric: Calling ip_defrag() in output path is also implying skb_orphan(), which is buggy because A relevant old patch about the issue was : 8282f27449bf ("inet: frag: Always orphan skbs inside ip_defrag()") [..] being set, and probably to an inet socket, not an arbitrary one. If we orphan the packet in ipvlan, then downstream t properly. We need to change ip_defrag() to only use skb_orphan() when really needed, ie whenever frag_list is goi fragment queue and made an initial patch. However there is a problem with this: If skb is refragmented again right to the new fragments, and sets up destructor to sock_wfree. IOW, we have no choice but to fix up sk_wmem accou wmem will underflow. This change moves the orphan down into the core, to last possible moment. As ip_defrag_o we must move the offset into the FRAG_CB, else skb->sk gets clobbered. This allows to delay the orphaning long or if the skb is completing the reasm queue. In the former case, things work as before, skb is orphaned. This is safe continue past reasm engine. In the latter case, we will steal the skb->sk reference, reattach it to the head skb, and fi truesize.
CVE-2024-26923	In the Linux kernel, the following vulnerability has been resolved: af_unix: Fix garbage collector racing against co account the risk of embryo getting enqueued during the garbage collection. If such embryo has a peer that carries S scan_children() may see a different set of children. Leading to an incorrectly elevated inflight count, and then a dan sockets are AF_UNIX/SOCK_STREAM S is an unconnected socket L is a listening in-flight socket bound to addr, via sendmsg(), gets inflight count bumped connect(S, addr) sendmsg(S, [V]); close(V) __unix_gc() ---------- -- unix_create1() skb1 = sock_wmalloc(NS) L = unix_find_other(addr) unix_state_lock(L) unix_peer(S) = NS // V co = sock_alloc() skb_queue_tail(NS, skb2[V]) // V became in-flight // V count=2 inflight=1 close(V) // V count=1 in u in gc_inflight_list: if (total_refs == inflight_refs) add u to gc_candidates // gc_candidates={L, V} for u in gc_can embryo (skb1) was not // reachable from L yet, so V's // inflight remains unchanged __skb_queue_tail(L, skb1) uni (u.inflight) scan_children(u, inc_inflight_move_tail) // V count=1 inflight=2 (!) If there is a GC-candidate listening wait until the end of any ongoing connect() to that socket. After flipping the lock, a possibly SCM-laden embryo is embryo coming, it can not possibly carry SCM_RIGHTS. At this point, unix_inflight() can not happen because uni remains unaffected.
CVE-2024-26923	In the Linux kernel, the following vulnerability has been resolved: af_unix: Fix garbage collector racing against co account the risk of embryo getting enqueued during the garbage collection. If such embryo has a peer that carries S scan_children() may see a different set of children. Leading to an incorrectly elevated inflight count, and then a dan sockets are AF_UNIX/SOCK_STREAM S is an unconnected socket L is a listening in-flight socket bound to addr, via sendmsg(), gets inflight count bumped connect(S, addr) sendmsg(S, [V]); close(V) __unix_gc() ---------- -- unix_create1() skb1 = sock_wmalloc(NS) L = unix_find_other(addr) unix_state_lock(L) unix_peer(S) = NS // V co = sock_alloc() skb_queue_tail(NS, skb2[V]) // V became in-flight // V count=2 inflight=1 close(V) // V count=1 in u in gc_inflight_list: if (total_refs == inflight_refs) add u to gc_candidates // gc_candidates={L, V} for u in gc_can embryo (skb1) was not // reachable from L yet, so V's // inflight remains unchanged __skb_queue_tail(L, skb1) uni (u.inflight) scan_children(u, inc_inflight_move_tail) // V count=1 inflight=2 (!) If there is a GC-candidate listening wait until the end of any ongoing connect() to that socket. After flipping the lock, a possibly SCM-laden embryo is embryo coming, it can not possibly carry SCM_RIGHTS. At this point, unix_inflight() can not happen because uni remains unaffected.
CVE-2024-26925	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_tables: release mutex after nft_gc_s should not be released during the critical section between nft_gc_seq_begin() and nft_gc_seq_end(), otherwise, asy and get the released commit lock within the same GC sequence. nf_tables_module_autoload() temporarily releases it goes back to replay the transaction again. Move it at the end of the abort phase after nft_gc_seq_end() is called.
CVE-2024-26934	In the Linux kernel, the following vulnerability has been resolved: USB: core: Fix deadlock in usb_deauthorize_int routines in drivers/usb/core/sysfs.c, the interface_authorized_store() function is the only one which acquires a devi usb_deauthorize_interface(), which locks the interface's parent USB device. The will lead to deadlock if another pr remove the interface, whether through a configuration change or because the device has been disconnected. As part waits for all ongoing sysfs attribute callbacks to complete. But usb_deauthorize_interface() can't complete until the won't be released until the removal has finished. The mechanism provided by sysfs to prevent this kind of deadlock function, which tells sysfs not to wait for the attribute callback. Reported-and-tested by: Yue Sun <samsun100621 <xrivendell7@gmail.com>

CVE-2024-26937	In the Linux kernel, the following vulnerability has been resolved: drm/i915/gt: Reset queue_priority_hint on parki execution, we could complete execution only when the queue was empty. Preempt-to-busy allows replacement of a before the preemption is processed by HW. If that happens, the request is retired from the queue, but the queue_pri submission until after the next CS interrupt is processed. This preempt-to-busy race can be triggered by the heartbe management barrier and upon completion allow us to idle the HW. We may process the completion of the heartbea CS event that restores the queue_priority_hint, causing us to fail the assertion that it is MIN. <3>[166.210729] __e >sched_engine->queue_priority_hint != (-((int)(~0U >> 1)) - 1)) <0>[166.210781] Dumping ftrace buffer: <0>[1 <0>[167.302811] drm_fdin-1097 2..s1. 165741070us : trace_ports: 0000:00:02.0 rcs0: promote { ccid:20 1217:2 p 2d.s2. 165741072us : execlists_submission_tasklet: 0000:00:02.0 rcs0: preempting last=1217:2, prio=0, hint=2147 2d.s2. 165741072us : __i915_request_unsubmit: 0000:00:02.0 rcs0: fence 1217:2, current 0 <0>[167.302992] drm __i915_request_submit: 0000:00:02.0 rcs0: fence 3:4660, current 4659 <0>[167.303044] drm_fdin-1097 2d.s1. 16 0000:00:02.0 rcs0: context:3 schedule-in, ccid:40 <0>[167.303095] drm_fdin-1097 2d.s1. 165741077us : trace_po 3:4660* prio 2147483646 } <0>[167.303159] kworker/-89 11..... 165741139us : i915_request_retire.part.0: 0000:0 <0>[167.303208] kworker/-89 11..... 165741148us : __intel_context_do_unpin: 0000:00:02.0 rcs0: context:c90 un 11..... 165741159us : i915_request_retire.part.0: 0000:00:02.0 rcs0: fence 1217:2, current 2 <0>[167.303321] kwo __intel_context_do_unpin: 0000:00:02.0 rcs0: context:1217 unpin <0>[167.303384] kworker/-89 11..... 16574117 rcs0: fence 3:4660, current 4660 <0>[167.303434] kworker/-89 11d..1. 165741172us : __intel_context_retire: 000 runtime: { total:56028ns, avg:56028ns } <0>[167.303484] kworker/-89 11..... 165741198us : __engine_park: 000 <idle>-0 5d.H3. 165741207us : execlists_irq_handler: 0000:00:02.0 rcs0: semaphore yield: 00000040 <0>[167.30 __intel_context_retire: 0000:00:02.0 rcs0: context:1217 retire runtime: { total:325575ns, avg:0ns } <0>[167.3037 __intel_context_retire: 0000:00:02.0 rcs0: context:c90 retire runtime: { total:0ns, avg:0ns } <0>[167.303806] kwo __engine_park:283 GEM_BUG_ON(engine->sched_engine->queue_priority_hint != (-((int)(~0U >> 1)) - 1)) <0>[<4>[167.304722] ------------[cut here]------------ <2>[167.304725] kernel BUG at drivers/gpu/drm/i915/gt/intel_ invalid opcode: 0000 [#1] PREEMPT SMP NOPTI <4>[167.304734] CPU: 11 PID: 89 Comm: kworker/11:1 Tain gc655e0fd2804+ #1 <4>[167.304736] Hardware name: Intel Corporation Rocket Lake Client Platform/RocketLak RKLSFWI1.R00.3173.A03.2204210138 04/21/2022 <4>[167.304738] Workqueue: i915-unordered retire_work_h
CVE-2024-26955	In the Linux kernel, the following vulnerability has been resolved: nilfs2: prevent kernel bug at submit_bh_wbc() F successful status when searching and inserting the specified block both fail inconsistently. If this inconsistent behav an unexpected race is occurring, so return a temporary error -EAGAIN instead. This prevents callers such as __bloc into a buffer that is not mapped, which would cause the BUG_ON check for the BH_Mapped flag in submit_bh_wb
CVE-2024-26956	In the Linux kernel, the following vulnerability has been resolved: nilfs2: fix failure to detect DAT corruption in bt fix kernel bug at submit_bh_wbc()". This resolves a kernel BUG reported by syzbot. Since there are two flaws invo The first patch alone resolves the syzbot-reported bug, but I think both fixes should be sent to stable, so I've tagged has reported a kernel bug in submit_bh_wbc() when writing file data to a nilfs2 file system whose metadata is corru issue. The first flaw is that when nilfs_get_block() locates a data block using btree or direct mapping, if the disk ad fails with internal code -ENOENT due to DAT metadata corruption, it can be passed back to nilfs_get_block(). This existing block as non-existent, causing both data block lookup and insertion to fail inconsistently. The second flaw status in this inconsistent state. This causes the caller __block_write_begin_int() or others to request a read even th in a BUG_ON check for the BH_Mapped flag in submit_bh_wbc() failing. This fixes the first issue by changing th a conversion using DAT fails with code -ENOENT, avoiding the conflicting condition that leads to the kernel bug (indicates that metadata corruption was detected during the block lookup, which will be properly handled as a file sy passing through the nilfs2 bmap layer.

CVE-2024-26957	In the Linux kernel, the following vulnerability has been resolved: s390/zcrypt: fix reference counting on zcrypt ca cards on KVM guests with debug kernel build revealed an use after free for the load field of the struct zcrypt_card. handling of the zcrypt card object which could lead to a free of the zcrypt card object while it was still in use. This kernel: 0x00000000885a7512-0x00000000885a7513 @offset=1298. First byte 0x68 instead of 0x6b kernel: Alloca [zcrypt] age=18046 cpu=3 pid=43 kernel: kmalloc_trace+0x3f2/0x470 kernel: zcrypt_card_alloc+0x36/0x70 [zcry +0x26/0x380 [zcrypt_cex4] kernel: ap_device_probe+0x15c/0x290 kernel: really_probe+0xd2/0x468 kernel: drive __device_attach_driver+0xc0/0x140 kernel: bus_for_each_drv+0x8c/0xd0 kernel: __device_attach+0x114/0x198 device_add+0x4d2/0x6e0 kernel: ap_scan_adapter+0x3d0/0x7c0 kernel: ap_scan_bus+0x5a/0x3b0 kernel: ap_scan process_one_work+0x26e/0x620 kernel: worker_thread+0x21c/0x440 kernel: Freed in zcrypt_card_put+0x54/0x8 kfree+0x37e/0x418 kernel: zcrypt_card_put+0x54/0x80 [zcrypt] kernel: ap_device_remove+0x4c/0xe0 kernel: dev kernel: bus_remove_device+0x100/0x188 kernel: device_del+0x164/0x3c0 kernel: device_unregister+0x30/0x90 kernel: ap_scan_bus+0x5a/0x3b0 kernel: ap_scan_bus_wq_callback+0x40/0x60 kernel: process_one_work+0x26e kernel: kthread+0x150/0x168 kernel: __ret_from_fork+0x3c/0x58 kernel: ret_from_fork+0xa/0x30 kernel: Slab 0x fp=0x00000000885a7c88 flags=0x3ffff00000000a00(workingset|slab|node=0|zone=1|lastcpupid=0x1ffff) kernel: C fp=0x00000000885a7c88 kernel: Redzone 00000000885a74b0: bb bb bb bb bb bb bb bb ..... kernel: Object 0000 6b 6b 6b 6b 6b 6b 6b kkkkkkkkkkkkkkk kernel: Object 00000000885a74c8: 6b 6b 6b 6b 6b 6b 6b 6b 6b 6b 6b Object 00000000885a74d8: 6b 6b 6b 6b 6b 6b 6b 6b 6b 6b 6b 6b 6b 6b 6b kkkkkkkkkkkkkkk kernel: Object 0 6b 6b 6b 6b 6b 6b 6b 6b kkkkkkkkkkkkkkk kernel: Object 00000000885a74f8: 6b 6b 6b 6b 6b 6b 6b 6b 6b kernel: Object 00000000885a7508: 6b 6b 6b 6b 6b 6b 6b 6b 6b 68 4b 6b 6b 6b a5 kkkkkkkkkkhKkkk. kernel: R bb bb bb bb bb ..... kernel: Padding 00000000885a756c: 5a 5a 5a 5a 5a 5a 5a 5a 5a 5a 5a ZZZZZZZZZZZ k udevd Not tainted 6.8.0-HF #2 kernel: Hardware name: IBM 3931 A01 704 (KVM/Linux) kernel: Call Trace: kern +0x90/0x120 kernel: [<00000000c99d78bc>] check_bytes_and_report+0x114/0x140 kernel: [<00000000c99d53c kernel: [<00000000c99d820c>] alloc_debug_processing+0xc4/0x1f8 kernel: [<00000000c99d852e>] get_partial_ [<00000000c99d94ec>] ___slab_alloc+0xaf4/0x13c8 kernel: [<00000000c99d9e38>] __slab_alloc.constprop.0+0x __kmalloc+0x434/0x590 kernel: [<00000000c9b4c0ce>] ext4_htree_store_dirent+0x4e/0x1c0 kernel: [<0000000 +0x17a/0x3f0 kernel: ---truncated---
CVE-2024-26960	In the Linux kernel, the following vulnerability has been resolved: mm: swap: fix race between free_swap_and_cac previously a theoretical window where swapoff() could run and teardown a swap_info_struct while a call to free_sw another thread. This could cause, amongst other bad possibilities, swap_page_trans_huge_swapped() (called by fre freed memory for swap_map. This is a theoretical problem and I haven't been able to provoke it from a test case. B on code review that this is possible (see link below). Fix it by using get_swap_device()/put_swap_device(), which extra check in _swap_info_get() to confirm that the swap entry was not free. This isn't present in get_swap_device(general due to the race between getting the reference and swapoff. So I've added an equivalent check directly in fre to provoke one possible issue (thanks to David Hildenbrand for deriving this): --8<----- __swap_entry_free() might SWAP_HAS_CACHE". swapoff->try_to_unuse() will stop as soon as soon as si->inuse_pages==0. So the questio turn si->inuse_pages==0, before we completed swap_page_trans_huge_swapped(). Imagine the following: 2 MiB f still references by swap entries. Process 1 still references subpage 0 via swap entry. Process 2 still references subpa free_swap_and_cache(). -> count == SWAP_HAS_CACHE [then, preempted in the hypervisor etc.] Process 2 quit == SWAP_HAS_CACHE Process 2 goes ahead, passes swap_page_trans_huge_swapped(), and calls __try_to_rec >folio_free_swap()->delete_from_swap_cache()-> put_swap_folio()->free_swap_slot()->swapcache_free_entries(> ... WRITE_ONCE(si->inuse_pages, si->inuse_pages - nr_entries); What stops swapoff to succeed after process 2 process1 finished its call to swap_page_trans_huge_swapped()? --8<-----
CVE-2024-26965	In the Linux kernel, the following vulnerability has been resolved: clk: qcom: mmcc-msm8974: fix terminating of f arrays are supposed to be terminated with an empty element. Add such entry to the end of the arrays where it is mis access when the table is traversed by functions like qcom_find_freq() or qcom_find_freq_floor(). Only compile tes
CVE-2024-26966	In the Linux kernel, the following vulnerability has been resolved: clk: qcom: mmcc-apq8084: fix terminating of fr arrays are supposed to be terminated with an empty element. Add such entry to the end of the arrays where it is mis access when the table is traversed by functions like qcom_find_freq() or qcom_find_freq_floor(). Only compile tes
CVE-2024-26969	In the Linux kernel, the following vulnerability has been resolved: clk: qcom: gcc-ipq8074: fix terminating of frequ are supposed to be terminated with an empty element. Add such entry to the end of the arrays where it is missing in when the table is traversed by functions like qcom_find_freq() or qcom_find_freq_floor(). Only compile tested.
CVE-2024-26972	In the Linux kernel, the following vulnerability has been resolved: ubifs: ubifs_symlink: Fix memleak of inode->i_ path in ubifs_symlink(), inode will be marked as bad first, then iput() is invoked. If inode->i_link is initialized by f scenario, inode->i_link won't be freed by callchain ubifs_free_inode -> fscrypt_free_inode in error handling path, b 'inode->i_mode' as 'S_IFREG'. Following kmemleak is easy to be reproduced by injecting error in ubifs_jnl_update scenario: unreferenced object 0xffff888103da3d98 (size 8): comm "ln", pid 1692, jiffies 4294914701 (age 12.045s) __fscrypt_encrypt_symlink+0xed/0x1c0 ubifs_symlink+0x210/0x300 [ubifs] vfs_symlink+0x216/0x360 do_symli There are two ways fixing it: 1. Remove make_bad_inode() in error handling path. We can do that because ubifs_e good symlink inode and bad symlink inode, for inode->i_nlink checking is before is_bad_inode(). 2. Free inode->i picked, it has less influence, personally, I think.
CVE-2024-26973	In the Linux kernel, the following vulnerability has been resolved: fat: fix uninitialized field in nostale filehandles handle without a parent it stores only first 10 bytes of the file handle. However the length of the file handle must be 12 bytes long and the last two bytes remain uninitialized. This is not great at we potentially leak uninitialized infor initialize the full handle length.

CVE-2024-26974	In the Linux kernel, the following vulnerability has been resolved: crypto: qat - resolve race condition during AER error recovery process, the kernel driver may encounter a race condition with freeing the reset_data structure's mem than 10 seconds the function scheduling that restart will exit due to a timeout, and the reset_data structure will be f completion notification after the restart is completed, which leads to a UAF bug. This results in a KFENCE bug no adf_device_reset_worker+0x38/0xa0 [intel_qat] Use-after-free read at 0x00000000bc56fddf (in kfence-#142): adf_ process_one_work+0x173/0x340 To resolve this race condition, the memory associated to the container of the wor expired, otherwise on the function that schedules the worker. The timeout detection can be done by checking if the using completion_done() function.
CVE-2024-26976	In the Linux kernel, the following vulnerability has been resolved: KVM: Always flush async #PF workqueue whe flush the per-vCPU async #PF workqueue when a vCPU is clearing its completion queue, e.g. when a VM and all i must ensure that none of its workqueue callbacks is running when the last reference to the KVM _module_ is put. C VM prevents the workqueue callback from dereferencing freed vCPU/VM memory, but does not prevent the KVM the callback completes. Drop the misguided VM refcount gifting, as calling kvm_put_kvm() from async_pf_execut async #PF workqueue will result in deadlock. async_pf_execute() can't return until kvm_put_kvm() finishes, and k async_pf_execute() finishes: WARNING: CPU: 8 PID: 251 at virt/kvm/kvm_main.c:1435 kvm_put_kvm+0x2d/0x vhost vhost_iotlb tap kvm_intel kvm irqbypass CPU: 8 PID: 251 Comm: kworker/8:1 Tainted: G W 6.6.0-rc1-e7af Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS 0.0.0 02/06/2015 Workqueue: events async_pf_e +0x2d/0x320 [kvm] Call Trace: <TASK> async_pf_execute+0x198/0x260 [kvm] process_one_work+0x145/0x2d0 +0xba/0xe0 ret_from_fork+0x2d/0x50 ret_from_fork_asm+0x11/0x20 </TASK> ---[end trace 0000000000000000 blocked for more than 120 seconds. Tainted: G W 6.6.0-rc1-e7af8d17224a-x86/gmem-vm #119 "echo 0 > /proc/sy disables this message. task:kworker/8:1 state:D stack:0 pid:251 ppid:2 flags:0x00004000 Workqueue: events async __schedule+0x33f/0xa40 schedule+0x53/0xc0 schedule_timeout+0x12a/0x140 __wait_for_common+0x8d/0x1d0 kvm_clear_async_pf_completion_queue+0x129/0x190 [kvm] kvm_arch_destroy_vm+0x78/0x1b0 [kvm] kvm_put +0x198/0x260 [kvm] process_one_work+0x145/0x2d0 worker_thread+0x27e/0x3a0 kthread+0xba/0xe0 ret_from_ +0x11/0x20 </TASK> If kvm_clear_async_pf_completion_queue() actually flushes the workqueue, then there's no because all invocations of async_pf_execute() will be forced to complete before the vCPU and its VM are destroye unloading bug as __fput() won't do module_put() on the last vCPU reference until the vCPU has been freed, e.g. if reference to the KVM module. Note that kvm_check_async_pf_completion() may also take the work item off the c the work queue, as the work will not be seen by kvm_clear_async_pf_completion_queue(). Waiting on the workqu waiting for the work to complete, but that's a very, very small chance, and likely a very small delay. kvm_arch_asy makes a new request, i.e. will effectively delay entering the guest, so the remaining work is really just: trace_kvm_ __kvm_vcpu_wake_up(vcpu); mmput(mm); and mmput() can't drop the last reference to the page tables if the vCP tearing down page tables. Add a helper to do the flushing, specifically to deal with "wakeup all" work items, as the placed in a workqueue. Trying to flush a bogus workqueue entry rightly makes __flush_work() complain (kudos to commit 5f6de5cbebee ("KVM: Prevent module exit until al ---truncated---
CVE-2024-26979	Rejected reason: This CVE ID has been rejected or withdrawn by its CVE Numbering Authority.
CVE-2024-26981	In the Linux kernel, the following vulnerability has been resolved: nilfs2: fix OOB in nilfs_set_de_type The size of nilfs2/dir.c file is defined as "S_IFMT >> S_SHIFT", but the nilfs_set_de_type() function, which uses this array, sp same way as "(mode & S_IFMT) >> S_SHIFT". static void nilfs_set_de_type(struct nilfs_dir_entry *de, struct inod de->file_type = nilfs_type_by_mode[(mode & S_IFMT)>>S_SHIFT]; // oob } However, when the index is determ occurs by referring to an index that is 1 larger than the array size when the condition "mode & S_IFMT == S_IFMT nilfs_type_by_mode array should be applied to prevent OOB errors.

CVE-2024-26991	In the Linux kernel, the following vulnerability has been resolved: KVM: x86/mmu: x86: Don't overflow lpage_inf KVM_SET_MEMORY_ATTRIBUTES to not overflow lpage_info array and trigger KASAN splat, as seen in the When memory attributes are set on a GFN range, that range will have specific properties applied to the TDP. A hug are inconsistent, so they are disabled for those the specific huge pages. For internal KVM reasons, huge pages are a regardless of whether the backing memory could be mapped as huge. What GFNs support which huge page sizes is on the memslot, of ,Äòkvm_lpage_info,Äô structs. Each index of lpage_info contains a vmalloc allocated array of The kvm_lpage_info denotes whether a specific huge page (GFN and page size) on the memslot is supported. Thes head and tail huge pages. Preventing huge pages from spanning adjacent memslot is covered by incrementing the c when the memslot is allocated, but disallowing huge pages for memory that has mixed attributes has to be done in a KVM_SET_MEMORY_ATTRIBUTES ioctl KVM updates lpage_info for each memslot in the range that has misr memslot at a time, and marks a special bit, KVM_LPAGE_MIXED_FLAG, in the kvm_lpage_info for any huge pa elevated count. So huge pages will not be mapped for the GFN at that page size if the count is elevated in either cas to the memslot or if KVM_LPAGE_MIXED_FLAG is set because it has mixed attributes. To determine whether a the KVM_SET_MEMORY_ATTRIBUTES operation checks an xarray to make sure it consistently has the incomi are aligned to level huge pages, it employs an optimization. As long as the level - 1 huge pages are checked first, it if each level - 1 huge page contained within the level sized huge page is not mixed, then the level size huge page is in the helper hugepage_has_attrs(). Unfortunately, although the kvm_lpage_info array representing page size 'level tail page of size level, the array for level - 1 will not contain an entry for each GFN at page size level. The level - 1 unaligned region covered by level - 1 huge page size, which can be a smaller region. So this causes the optimizatio and perform a vmalloc out of bounds read. In some cases of head and tail pages where an overflow could happen, c as KVM_LPAGE_MIXED_FLAG is not required to prevent huge pages as discussed earlier. But for memslots tha does call hugepage_has_attrs(). In this case the huge page is both the head and tail page. The issue can be observed CONFIG_KASAN_VMALLOC and running the selftest ,Äúprivate_mem_conversions_test,Äù, which produces th KASAN: vmalloc-out-of-bounds in hugepage_has_attrs+0x7e/0x110 Read of size 4 at addr ffffc900000a3008 by ta dump_stack_lvl print_report ? __virt_addr_valid ? hugepage_has_attrs ? hugepage_has_attrs kasan_report ? hugep kvm_arch_post_set_memory_attributes kvm_vm_ioctl It is a little ambiguous whether the unaligned head page (in expected to have KVM_LPAGE_MIXED_FLAG set. It is not functionally required, as the unal ---truncated---
CVE-2024-26993	In the Linux kernel, the following vulnerability has been resolved: fs: sysfs: Fix reference leak in sysfs_break_activ sysfs_break_active_protection() routine has an obvious reference leak in its error path. If the call to kernfs_find_an companion sysfs_unbreak_active_protection() routine won't get called (and would only cause an access violation b called). As a result, the reference to kobj acquired at the start of the function will never be released. Fix the leak by is NULL.
CVE-2024-26994	In the Linux kernel, the following vulnerability has been resolved: speakup: Avoid crash on very long word In case really long word (> 256 characters), we have to stop before the length of the word buffer.
CVE-2024-27000	In the Linux kernel, the following vulnerability has been resolved: serial: mxs-auart: add spinlock around changing function in serial_core expects the caller to hold uport->lock. For example, I have seen the below kernel splat, whe board. [85.119255] ------------[cut here]------------ [85.124413] WARNING: CPU: 0 PID: 27 at /drivers/tty/serial +0xb4/0xec [85.134694] Modules linked in: hci_uart bluetooth ecdh_generic ecc wlcore_sdio configfs [85.14331 tainted 6.6.3-00021-gd62a2f068f92 #1 [85.151396] Hardware name: Freescale MXS (Device Tree) [85.156679] W (...) [85.191765] uart_handle_cts_change from mxs_auart_irq_handle+0x380/0x3f4 [85.198787] mxs_auart_irq_h +0x88/0x210 (...)
CVE-2024-27001	In the Linux kernel, the following vulnerability has been resolved: comedi: vmk80xx: fix incomplete endpoint chec checking implemented, some things can fall through the cracks. Depending on the hardware model, URBs can have version of vmk80xx_find_usb_endpoints() function does not take that fully into account. While this warning does n it will crash systems with 'panic_on_warn' set on them. Fix the issue found by Syzkaller [1] by somewhat simplifyi usb_find_common_endpoints() and ensuring that only expected endpoint types are present. This patch has not been report: usb 1-1: BOGUS urb xfer, pipe 1 != type 3 WARNING: CPU: 0 PID: 781 at drivers/usb/core/urb.c:504 usb_ core/urb.c:503 ... Call Trace: <TASK> usb_start_wait_urb+0x113/0x520 drivers/usb/core/message.c:59 vmk80xx_ vmk80xx.c:227 [inline] vmk80xx_auto_attach+0xa1c/0x1a40 drivers/comedi/drivers/vmk80xx.c:818 comedi_auto drivers.c:1067 usb_probe_interface+0x5cd/0xb00 drivers/usb/core/driver.c:399 ... Similar issue also found by Syzk

CVE-2024-27012	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_tables: restore set elements when de nft_mapelem_activate() needs to restore refcounters to the original state. Currently, it uses the set->ops->walk() to existing set iterator skips inactive elements in the next generation, this does not work from the abort path to restore active elements instead (not inactive ones). This patch moves the check for inactive elements to the set iterator call the .activate case which needs to skip active elements. Toggle next generation bit for elements when delete set com from .activate (abort) path to restore the next generation bit. The splat below shows an object in mappings memleal here]---------- [43929.457532] WARNING: CPU: 0 PID: 1139 at include/net/netfilter/nf_tables.h:1237 nft_setele [...] [43929.458014] RIP: 0010:nft_setelem_data_deactivate+0xe4/0xf0 [nf_tables] [43929.458076] Code: 83 f8 01 de 49 8b 6c 24 08 48 8d 7d 50 e8 e9 5c d0 de 8b 45 50 8d 50 ff 89 55 50 85 c0 75 86 <0f> 0b eb 82 0f 0b eb b3 0f 90 90 [43929.458081] RSP: 0018:ffff888140f9f4b0 EFLAGS: 00010246 [43929.458086] RAX: 00000000000000 dffffc0000000000 [43929.458090] RDX: 00000000ffffffff RSI: ffffffffa26d28a7 RDI: ffff88810ecc9550 [43929.45 0000000000000001 R09: ffffed10281f3e8f [43929.458096] R10: 0000000000000003 R11: ffff0000ffff0000 R12: R13: ffff888140f9f5f4 R14: ffff888151c7a800 R15: 0000000000000002 [43929.458103] FS: 00007f0c687c4740(0 knlGS:0000000000000000 [43929.458107] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [43929.45811 0000000123602005 CR4: 00000000001706f0 [43929.458114] Call Trace: [43929.458118] <TASK> [43929.4581 nft_setelem_data_deactivate+0xe4/0xf0 [nf_tables] [43929.458188] ? report_bug+0x1b1/0x1e0 [43929.458196] ? exc_invalid_op+0x17/0x40 [43929.458211] ? nft_setelem_data_deactivate+0xd7/0xf0 [nf_tables] [43929.458271] [nf_tables] [43929.458332] nft_mapelem_deactivate+0x24/0x30 [nf_tables] [43929.458392] nft_rhash_walk+0xdc __pfx_nft_rhash_walk+0x10/0x10 [nf_tables] [43929.458512] ? rb_insert_color+0x2e/0x280 [43929.458520] nft_ [nf_tables] [43929.458582] ? __pfx_nft_map_deactivate+0x10/0x10 [nf_tables] [43929.458642] ? __pfx_nft_mape [43929.458701] ? __rcu_read_unlock+0x46/0x70 [43929.458709] nft_delset+0xff/0x110 [nf_tables] [43929.45876 [43929.458830] nf_tables_deltable+0x501/0x580 [nf_tables]
CVE-2024-27013	In the Linux kernel, the following vulnerability has been resolved: tun: limit printing rate when illegal packet recei tun call backs to receive packets. If too many illegal packets arrives, tun_do_read will keep dumping packet conten much more cpu time to dump packet and soft lockup will be detected. net_ratelimit mechanism can be used to limi ffff949da6f20000 CPU: 23 COMMAND: "vhost-32980" #0 [fffffe00003fce50] crash_nmi_callback at ffffffff8924 at ffffffff89225fa3 #2 [fffffe00003fceb0] default_do_nmi at ffffffff8922642e #3 [fffffe00003fced0] do_nmi at fffff end_repeat_nmi at ffffffff89c01663 [exception RIP: io_serial_in+20] RIP: ffffffff89792594 RSP: ffffa655314979e ffffffff89792500 RBX: ffffffff8af428a0 RCX: 0000000000000000 RDX: 00000000000003fd RSI: 0000000000000 0000000000002710 R8: 0000000000000004 R9: 000000000000000f R10: 0000000000000000 R11: ffffffff8acbf6 ffffffff8acbf698 R14: 0000000000000058 R15: 0000000000000000 ORIG_RAX: ffffffffffffffff CS: 0010 SS: 0018 at ffffffff89792594 #6 [ffffa655314979e8] wait_for_xmitr at ffffffff89793470 #7 [ffffa65531497a08] serial8250_c [ffffa65531497a20] uart_console_write at ffffffff8978b605 #9 [ffffa65531497a48] serial8250_console_write at ffff console_unlock at ffffffff89316124 #11 [ffffa65531497b10] vprintk_emit at ffffffff89317c07 #12 [ffffa65531497b [ffffa65531497bc8] print_hex_dump at ffffffff89650765 #14 [ffffa65531497ca8] tun_do_read at ffffffffc0b06c27 [ffffffffc0b06e34 [tun] #16 [ffffa65531497d68] handle_rx at ffffffffc0c5d682 [vhost_net] #17 [ffffa65531497ed0] v [ffffa65531497f10] kthread at ffffffff892d2e72 #19 [ffffa65531497f50] ret_from_fork at ffffffff89c0022f
CVE-2024-27019	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_tables: Fix potential data-race in __ can concurrent with __nft_obj_type_get(), and there is not any protection when iterate over nf_tables_objects list i is potential data-race of nf_tables_objects list entry. Use list_for_each_entry_rcu() to iterate over nf_tables_objects rcu_read_lock() in the caller nft_obj_type_get() to protect the entire type query process.
CVE-2024-27020	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_tables: Fix potential data-race in __ concurrent with __nft_expr_type_get(), and there is not any protection when iterate over nf_tables_expressions list potential data-race of nf_tables_expressions list entry. Use list_for_each_entry_rcu() to iterate over nf_tables_expr rcu_read_lock() in the caller nft_expr_type_get() to protect the entire type query process.
CVE-2024-27024	In the Linux kernel, the following vulnerability has been resolved: net/rds: fix WARNING in rds_conn_connect_if get_mr() will fail, trigger connection after get_mr().
CVE-2024-27026	In the Linux kernel, the following vulnerability has been resolved: vmxnet3: Fix missing reserved tailroom Use rbi dataring packet. Found issue: XDP_WARN: xdp_update_frame_from_buff(line:278): Driver BUG: missing reserv 0 PID: 0 at net/core/xdp.c:586 xdp_warn+0xf/0x20 CPU: 0 PID: 0 Comm: swapper/0 Tainted: G W O 6.5.1 #1 RIP xdp_warn+0xf/0x20 xdp_do_redirect+0x15f/0x1c0 vmxnet3_run_xdp+0x17a/0x400 [vmxnet3] vmxnet3_process_ vmxnet3_tq_tx_complete.isra.0+0x21e/0x2c0 [vmxnet3] vmxnet3_rq_rx_complete+0x7ad/0x1120 [vmxnet3] vmx __napi_poll+0x20/0x180 net_rx_action+0x177/0x390
CVE-2024-27027	In the Linux kernel, the following vulnerability has been resolved: dpll: fix dpll_xa_ref_*_del() for multiple registr registrations of the same pin on the same dpll device, following warnings are observed: WARNING: CPU: 5 PID: dpll_xa_ref_pin_del.isra.0+0x21e/0x230 WARNING: CPU: 5 PID: 2212 at drivers/dpll/dpll_core.c:223 __dpll_pi that in both dpll_xa_ref_dpll_del() and dpll_xa_ref_pin_del() registration is only removed from list in case the refe the registration has to be removed always. To fix this, remove the registration from the list and free it unconditiona reference counter reaches zero.
CVE-2024-27028	In the Linux kernel, the following vulnerability has been resolved: spi: spi-mt65xx: Fix NULL pointer access in int can be a NULL pointer, so the interrupt handler may end up writing to the invalid memory and cause crashes. Add
CVE-2024-27030	In the Linux kernel, the following vulnerability has been resolved: octeontx2-af: Use separate handlers for interrup vector same interrupt handler is registered which is causing race condition. When two interrupts are raised to two C event corrupting the data.

CVE-2024-27031	In the Linux kernel, the following vulnerability has been resolved: NFS: Fix nfs_netfs_issue_read() xarray locking nfs_netfs_issue_read() currently does not disable interrupts while iterating through pages in the xarray to submit fo since after taking xa_lock, another page in the mapping could be processed for writeback inside an interrupt, and d and clean if we use xa_for_each_range(), which handles the iteration with RCU while reducing code complexity. T with the following test: mount -o vers=3,fsc 127.0.0.1:/export /mnt/nfs dd if=/dev/zero of=/mnt/nfs/file1.bin bs=40 drop_caches dd if=/mnt/nfs/file1.bin of=/dev/null umount /mnt/nfs On the console with a lockdep-enabled kernel a be seen: ================================ WARNING: inconsistent lock state 6.7.0-lockdbg+ #10 Not inconsistent {IN-SOFTIRQ-W} -> {SOFTIRQ-ON-W} usage. test5/1708 [HC0[0]:SC0[0]:HE1:SE1] takes: ffff88 {3:3}, at: nfs_netfs_issue_read+0x1b2/0x4b0 [nfs] {IN-SOFTIRQ-W} state was registered at: lock_acquire+0x144 +0x4e/0xa0 __folio_end_writeback+0x17e/0x5c0 folio_end_writeback+0x93/0x1b0 iomap_finish_ioend+0xeb/0x blk_mq_end_request+0x30/0x1c0 blk_complete_reqs+0x7e/0xa0 __do_softirq+0x113/0x544 __irq_exit_rcu+0xfe sysvec_call_function_single+0x6f/0x90 asm_sysvec_call_function_single+0x1a/0x20 pv_native_safe_halt+0xf/0x +0x67/0xa0 do_idle+0x2b5/0x300 cpu_startup_entry+0x34/0x40 start_secondary+0x19d/0x1c0 secondary_startup event stamp: 176891 hardirqs last enabled at (176891): [<ffffffffa67a0be4>] _raw_spin_unlock_irqrestore+0x44/0 [<ffffffffa67a0899>] _raw_spin_lock_irqsave+0x79/0xa0 softirqs last enabled at (176646): [<ffffffffa515d91e>] _ last disabled at (176633): [<ffffffffa515d91e>] __irq_exit_rcu+0xfe/0x120 other info that might help us debug this CPU0 ---- lock(&xa->xa_lock#4); <Interrupt> lock(&xa->xa_lock#4); *** DEADLOCK *** 2 locks held by test5 >s_type->i_mutex_key#22){++++}-{4:4}, at: nfs_start_io_read+0x28/0x90 [nfs] #1: ffff888127baa650 (mapping. page_cache_ra_unbounded+0xa4/0x280 stack backtrace: CPU: 6 PID: 1708 Comm: test5 Kdump: loaded Not taint QEMU Standard PC (Q35 + ICH9, 2009), BIOS 1.16.3-1.fc39 04/01/2014 Call Trace: dump_stack_lvl+0x5b/0x90 +0x77b/0x3360 _raw_spin_lock+0x34/0x80 nfs_netfs_issue_read+0x1b2/0x4b0 [nfs] netfs_begin_read+0x77f/0x [nfs] nfs_readahead+0x323/0x5a0 [nfs] read_pages+0xf3/0x5c0 page_cache_ra_unbounded+0x1c8/0x280 filemap +0x206/0x5e0 nfs_file_read+0xb7/0x140 [nfs] vfs_read+0x2a9/0x460 ksys_read+0xb7/0x140
CVE-2024-27038	In the Linux kernel, the following vulnerability has been resolved: clk: Fix clk_core_get NULL dereference It is po NULL in the following sequence: clk_core_get() of_clk_get_hw_from_clkspec() __of_clk_get_hw_from_provider NULL which is dereferenced by clk_core_get() at hw->core. Prior to commit dde4eff47c82 ("clk: Look for parents IS_ERR_OR_NULL() was performed which would have caught the NULL. Reading the description of this functio cannot be so at the moment. Update the function to check for hw before dereferencing it and return NULL if hw is
CVE-2024-27039	In the Linux kernel, the following vulnerability has been resolved: clk: hisilicon: hi3559a: Fix an erroneous devm_ the for loop for all clk that need to be registered. It is incremented at each loop iteration. If a clk_register() call fails from what should be freed. The best we can do, is to avoid this wrong release of memory.
CVE-2024-27043	In the Linux kernel, the following vulnerability has been resolved: media: edia: dvbdev: fix a use-after-free In dvb_ equal to dvbdev, which is freed in several error-handling paths. However, *pdvbdev is not set to NULL after dvbde frees in many places, for example, in the following call chain: budget_register |-> dvb_dmxdev_init |-> dvb_registe dvb_unregister_device |-> dvb_remove_device |-> dvb_device_put |-> kref_put When calling dvb_unregister_devi dvb_register_device) could point to memory that had been freed in dvb_register_device. Thereafter, this pointer is after-free.
CVE-2024-27044	In the Linux kernel, the following vulnerability has been resolved: drm/amd/display: Fix potential NULL pointer d 'dcn10_set_output_transfer_func()' The 'stream' pointer is used in dcn10_set_output_transfer_func() before the che drivers/gpu/drm/amd/amdgpu/../display/dc/hwss/dcn10/dcn10_hwseq.c:1892 dcn10_set_output_transfer_func() wa 'stream' (see line 1875)
CVE-2024-27046	In the Linux kernel, the following vulnerability has been resolved: nfp: flower: handle acti_netdevs allocation failu nfp_fl_lag_do_work() will return null, if the physical memory has run out. As a result, if we dereference the acti_n happen. This patch adds a check to judge whether allocation failure occurs. If it happens, the delayed work will be
CVE-2024-27053	In the Linux kernel, the following vulnerability has been resolved: wifi: wilc1000: fix RCU usage in connect path function from cfg802.11 layer lead to the following warning: ============================ WARNING #333 Not tainted ----------------------------- drivers/net/wireless/microchip/wilc1000/hif.c:386 suspicious rcu_derefe CPU: 0 PID: 100 Comm: wpa_supplicant Not tainted 6.7.0-rc1-wt+ #333 Hardware name: Atmel SAMA5 unwind show_stack from dump_stack_lvl+0x34/0x48 dump_stack_lvl from wilc_parse_join_bss_param+0x7dc/0x7f4 wilc +0x2c4/0x648 connect from cfg80211_connect+0x30c/0xb74 cfg80211_connect from nl80211_connect+0x860/0x +0x3fc/0x59c genl_rcv_msg from netlink_rcv_skb+0xd0/0x1f8 netlink_rcv_skb from genl_rcv+0x2c/0x3c genl_rc netlink_unicast from netlink_sendmsg+0x368/0x688 netlink_sendmsg from ____sys_sendmsg+0x190/0x430 ____ +0x110/0x158 ___sys_sendmsg from sys_sendmsg+0xe8/0x150 sys_sendmsg from ret_fast_syscall+0x0/0x1c Thi path, when trying to parse target BSS parameters, we dereference a RCU pointer whithout being in RCU critical se it to a RCU read critical section. To avoid wrapping the whole wilc_parse_join_bss_param under the critical sectio
CVE-2024-27059	In the Linux kernel, the following vulnerability has been resolved: USB: usb-storage: Prevent divide-by-0 error in i in usb-storage uses the HEADS and SECTORS values in the ATA ID information to calculate cylinder and head va WRITE commands. The calculation involves division and modulus operations, which will cause a crash if either of with a genuine device, it could happen with a flawed or subversive emulation, as reported by the syzbot fuzzer. Pro bind to the device if either the ATA_ID_HEADS or ATA_ID_SECTORS value in the device's ID information is 0. a negative error code when initialization fails; currently it always returns 0 (even when there is an error).

CVE-2024-27059	In the Linux kernel, the following vulnerability has been resolved: USB: usb-storage: Prevent divide-by-0 error in i in usb-storage uses the HEADS and SECTORS values in the ATA ID information to calculate cylinder and head va WRITE commands. The calculation involves division and modulus operations, which will cause a crash if either of with a genuine device, it could happen with a flawed or subversive emulation, as reported by the syzbot fuzzer. Pro bind to the device if either the ATA_ID_HEADS or ATA_ID_SECTORS value in the device's ID information is 0. a negative error code when initialization fails; currently it always returns 0 (even when there is an error).
CVE-2024-27065	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_tables: do not compare internal tabl transaction if table update does not modify flags.
CVE-2024-27065	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_tables: do not compare internal tabl transaction if table update does not modify flags.
CVE-2024-27074	In the Linux kernel, the following vulnerability has been resolved: media: go7007: fix a memleak in go7007_load_ bounce(i.e. go->boot_fw), is allocated without a deallocation thereafter. After the following call chain: saa7134_go go7007_load_encoder |-> kfree(go) go is freed and thus bounce is leaked.
CVE-2024-27076	In the Linux kernel, the following vulnerability has been resolved: media: imx: csc/scaler: fix v4l2_ctrl_handler me v4l2_ctrl_handler_init on release.
CVE-2024-27077	In the Linux kernel, the following vulnerability has been resolved: media: v4l2-mem2mem: fix a memleak in v4l2_ name) is allocated in v4l2_m2m_register_entity but isn't freed in its following error-handling paths. This patch add entity->name.
CVE-2024-27077	In the Linux kernel, the following vulnerability has been resolved: media: v4l2-mem2mem: fix a memleak in v4l2_ name) is allocated in v4l2_m2m_register_entity but isn't freed in its following error-handling paths. This patch add entity->name.
CVE-2024-27078	In the Linux kernel, the following vulnerability has been resolved: media: v4l2-tpg: fix some memleaks in tpg_allo deallocated in each and every error-handling paths, since they are allocated in for statements. Otherwise there woul only when tpg_alloc return 0.
CVE-2024-27388	In the Linux kernel, the following vulnerability has been resolved: SUNRPC: fix some memleaks in gssx_dec_opti freed in the error-handling paths after their allocation. So this patch add these deallocations in the corresponding pa
CVE-2024-27390	In the Linux kernel, the following vulnerability has been resolved: ipv6: mcast: remove one synchronize_net() barr past (commit 2d3916f31891 ("ipv6: fix skb drops in igmp6_event_query() and igmp6_event_report()")) I think the not needed. Under load, synchronize_net() can last between 200 usec and 5 ms. KASAN seems to agree as well.
CVE-2024-27393	In the Linux kernel, the following vulnerability has been resolved: xen-netfront: Add missing skb_mark_for_recycl introduced later than fixes tag in commit 6a5bcd84e886 ("page_pool: Allow drivers to hint on SKB recycling"). It to page_pool_release_page() between v5.9 to v5.14, after which is should have used skb_mark_for_recycle(). Since were removed (in commit 535b9c61bdef ("net: page_pool: hide page_pool_release_page()") and remaining callers branch 'net-page_pool-remove-page_pool_release_page'")). This leak became visible in v6.8 via commit dba1b8a7 memory leaks").
CVE-2024-27395	In the Linux kernel, the following vulnerability has been resolved: net: openvswitch: Fix Use-After-Free in ovs_ct_ hlist_for_each_entry_rcu traversal of ovs_ct_limit_exit, is not part of the RCU read critical section, it is possible th traversal and the key will be free. To prevent this, it should be changed to hlist_for_each_entry_safe.
CVE-2024-27396	In the Linux kernel, the following vulnerability has been resolved: net: gtp: Fix Use-After-Free in gtp_dellink Since hlist_for_each_entry_rcu traversal of gtp_dellink, is not part of the RCU read critical section, it is possible that the traversal and the key will be free. To prevent this, it should be changed to hlist_for_each_entry_safe.
CVE-2024-27397	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_tables: use timestamp to check for s at the beginning of the transaction, store it in the nftables per-netns area. Update set backend .insert, .deactivate and avoids that an element expires while control plane transaction is still unfinished. .lookup and .update, which are use to check if the element has expired. And .get path and dump also since this runs lockless under rcu read size lock. T check the current time since it runs asynchronously from a workqueue.

CVE-2024-27399	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: l2cap: fix null-ptr-deref in l2cap_cha between l2cap_chan_timeout() and l2cap_chan_del(). When we use l2cap_chan_del() to delete the channel, the cha conn could be dereferenced again in the mutex_lock() of l2cap_chan_timeout(). As a result the null pointer derefer report triggered by POC is shown below: [472.074580] ========== [472.075284] BUG: KASAN: null-ptr-deref in mutex_lock+0x68/0xc0 [472.075308] Write of size 8 at addr 0000 [472.075308] [472.075308] CPU: 0 PID: 7 Comm: kworker/0:0 Not tainted 6.9.0-rc5-00356-g78c0094a146b #36 QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.14.0-0-g155821a1990b-prebuilt.qemu4 [472.075308] Wo [472.075308] Call Trace: [472.075308] <TASK> [472.075308] dump_stack_lvl+0x137/0x1a0 [472.075308] pri __virt_addr_valid+0x77/0x160 [472.075308] ? mutex_lock+0x68/0xc0 [472.075308] kasan_report+0x139/0x170 +0x68/0xc0 [472.075308] kasan_check_range+0x2c3/0x2e0 [472.075308] mutex_lock+0x68/0xc0 [472.075308] [472.075308] process_one_work+0x5d2/0xe00 [472.075308] worker_thread+0xe1d/0x1660 [472.075308] ? pr_c kthread+0x2b7/0x350 [472.075308] ? pr_cont_work+0x5e0/0x5e0 [472.075308] ? kthread_blkcg+0xd0/0xd0 [47 +0x4d/0x80 [472.075308] ? kthread_blkcg+0xd0/0xd0 [472.075308] ret_from_fork_asm+0x11/0x20 [472.07530 ========== [472.094860] Disabl taint [472.096136] BUG: kernel NULL pointer dereference, address: 0000000000000158 [472.096136] #PF: supe [472.096136] #PF: error_code(0x0002) - not-present page [472.096136] PGD 0 P4D 0 [472.096136] Oops: 0002 [472.096136] CPU: 0 PID: 7 Comm: kworker/0:0 Tainted: G B 6.9.0-rc5-00356-g78c0094a146b #36 [472.09613 PC (i440FX + PIIX, 1996), BIOS rel-1.14.0-0-g155821a1990b-prebuilt.qemu4 [472.096136] Workqueue: events l RIP: 0010:mutex_lock+0x88/0xc0 [472.096136] Code: be 08 00 00 00 e8 f8 23 1f fd 4c 89 f7 be 08 00 00 00 e8 e [472.096136] RSP: 0018:ffff88800744fc78 EFLAGS: 00000246 [472.096136] RAX: 0000000000000000 RBX: [472.096136] RDX: 0000000000000001 RSI: 0000000000000008 RDI: ffff88800744fc78 [472.096136] RBP: 00 R09: 1ffff11000e89f8f [472.096136] R10: dfffc0000000000 R11: ffffed1000e89f90 R12: dfffc0000000000 [47 ffff88800744fc78 R15: ffff888007405a00 [472.096136] FS: 0000000000000000(0000) GS:ffff88806d200000(000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [472.096136] CR2: 0000000000000158 CR3: 000000000 [472.096136] Call Trace: [472.096136] <TASK> [472.096136] ? __die_body+0x8d/0xe0 [472.096136] ? page_ kernelmode_fixup_or_oops+0x20c/0x2a0 [472.096136] ? do_user_addr_fault+0x1027/0x1340 [472.096136] ? _p mutex_lock+0x68/0xc0 [472.096136] ? add_taint+0x42/0xd0 [472.096136] ? exc_page_fault+0x6a/0x1b0 [472.0 [472.096136] ? mutex_lock+0x75/0xc0 [472.096136] ? mutex_lock+0x88/0xc0 [472.096136] ? mutex_lock+0x7 truncated---
CVE-2024-9486	A security issue was discovered in the Kubernetes Image Builder versions <= v0.1.37 where default credentials are Virtual machine images built using the Proxmox provider do not disable these default credentials, and nodes using these default credentials. The credentials can be used to gain root access. Kubernetes clusters are only affected if th Image Builder project with its Proxmox provider. **Note:** This CVE-2024-9486 does not impact Cloudera Data Services on premises.
CVE-2024-27401	In the Linux kernel, the following vulnerability has been resolved: firewire: nosy: ensure user_length is taken into a that packet_buffer_get respects the user_length provided. If the length of the head packet exceeds the user_length, to the user that no data were read and a larger buffer size is required. Helps prevent user space overflows.
CVE-2024-27405	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: ncm: Avoid dropping datagrams of sometimes when tethering is used over NCM with Windows 11 as host, at some instances, the gadget_giveback has NTB. When the NTB is parsed, unwrap call looks for any leftover bytes in SKB provided by u_ether and if there ar a separate NTB and parses it. But in case the second NTB (as per unwrap call) is faulty/corrupt, all the datagrams t NTB and saved in rx_list are dropped. Adding a few custom traces showed the following: [002] d..1 7828.532866: 000000003868811a length 1025/16384 zsI ==> 0 [002] d..1 7828.532867: ncm_unwrap_ntb: K: ncm_unwrap_ntb ncm_unwrap_ntb: K: ncm_unwrap_ntb nth: 1751999342 [002] d..1 7828.532868: ncm_unwrap_ntb: K: ncm_unwr ncm_unwrap_ntb: K: ncm_unwrap_ntb blk_len: 0x400 [002] d..1 7828.532868: ncm_unwrap_ntb: K: ncm_unwrap ncm_unwrap_ntb: K: Parsed NTB with 1 frames In this case, the giveback is of 1025 bytes and block length is 102 be parsed resulting in drop of all datagrams in rx_list. Same is case with packets of size 2048: [002] d..1 7828.5579 0000000011dfd96e length 2049/16384 zsI ==> 0 [002] d..1 7828.557949: ncm_unwrap_ntb: K: ncm_unwrap_ntb ncm_unwrap_ntb: K: ncm_unwrap_ntb blk_len: 0x800 Lecroy shows one byte coming in extra confirming that the Bytes Transferred(1025) Timestamp((18.524 843 590) - Transaction 8391 - Data(1025 bytes) Timestamp(18.524 8 Duration(2.117us) Idle(14.700ns) Timestamp(18.524 843 590) --- Packet 4063863 Data(1 byte) Duration(66.160ns 722) According to Windows driver, no ZLP is needed if wBlockLength is non-zero, because the non-zero wBlockL size of transfer to be expected. However, there are in-market NCM devices that rely on ZLP as long as the wBlockL deal with such devices, it pads an extra 0 at end so the transfer is no longer multiple of wMaxPacketSize.
CVE-2024-27410	In the Linux kernel, the following vulnerability has been resolved: wifi: nl80211: reject iftype change with mesh ID mesh ID when the interface isn't yet in mesh mode, at the same time as changing it into mesh mode. This leads to a for the interface type it currently has, causing cfg80211_change_iface() to do wrong things when switching. We co mesh while setting the mesh ID at the same time by doing a different order of operations here, but realistically there disallow changes in iftype when setting mesh ID.

CVE-2024-27412	In the Linux kernel, the following vulnerability has been resolved: power: supply: bq27xxx-i2c: Do not free non ex have an IRQ, in which case client->irq will be 0. bq27xxx_battery_i2c_probe() already has an if (client->irq) check bq27xxx_battery_i2c_remove() unconditionally calls free_irq(client->irq) leading to: [190.310742] ----------[cut already-free IRQ 0 [190.310861] WARNING: CPU: 2 PID: 1304 at kernel/irq/manage.c:1893 free_irq+0x1b8/0x3 the driver. Add an if (client->irq) to bq27xxx_battery_i2c_remove() mirroring probe() to fix this.
CVE-2024-27414	In the Linux kernel, the following vulnerability has been resolved: rtnetlink: fix error logic of IFLA_BRIDGE_FLA d73ef2d69c0d ("rtnetlink: let rtnl_bridge_setlink checks IFLA_BRIDGE_MODE length"), an adjustment was mad `rtnl_bridge_setlink` to enable the loop to also check the length of the IFLA_BRIDGE_MODE attribute. However, statement and led to an error logic of the flags writing back at the end of this function. if (have_flags) memcpy(nla_ should point to IFLA_BRIDGE_FLAGS NLA !!! Before the mentioned commit, the `attr` is granted to be IFLA_B necessarily true fow now as the updated loop will let the attr point to the last NLA, even an invalid NLA which cou introduces a new variable `br_flag` to save the NLA pointer that points to IFLA_BRIDGE_FLAGS and uses it to r
CVE-2024-27416	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: hci_event: Fix handling of HCI_EV_ HCI_EV_IO_CAPA_REQUEST while HCI_OP_READ_REMOTE_EXT_FEATURES is yet to be responded assu otherwise this event shouldn't be generated.
CVE-2024-27419	In the Linux kernel, the following vulnerability has been resolved: netrom: Fix data-races around sysctl_net_busy_ sysctl value because the value can be changed concurrently.
CVE-2024-28182	nghttp2 is an implementation of the Hypertext Transfer Protocol version 2 in C. The nghttp2 library prior to versio number of HTTP/2 CONTINUATION frames even after a stream is reset to keep HPACK context in sync. This ca stream. nghttp2 v1.61.0 mitigates this vulnerability by limiting the number of CONTINUATION frames it accepts vulnerability.
CVE-2024-28757	libexpat through 2.6.1 allows an XML Entity Expansion attack when there is isolated use of external parsers (creat
CVE-2024-29040	This repository hosts source code implementing the Trusted Computing Group's (TCG) TPM2 Software Stack (TS Fapi_Quote has to be deserialized by Fapi_VerifyQuote to the TPM Structure `TPMS_ATTEST`. For the field `TP any number can be used in the JSON structure. The verifier can receive a state which does not represent the actual, test. The malicious device might get access to data it shouldn't, or can use services it shouldn't be able to. This issu
CVE-2024-29511	Artifex Ghostscript before 10.03.1, when Tesseract is used for OCR, has a directory traversal issue that allows arbi messages to arbitrary files) via OCRLanguage. For example, exploitation can use debug_file /tmp/out and user_pat
CVE-2024-31047	An issue in Academy Software Foundation openexr v.3.2.3 and before allows a local attacker to cause a denial of s exrmultipart.cpp.
CVE-2024-31076	In the Linux kernel, the following vulnerability has been resolved: genirq/cpuhotplug, x86/vector: Prevent vector le IRQD_MOVE_PCNTXT prevents immediate effectiveness of interrupt affinity reconfiguration via procfs. Instead, instance of the interrupt being triggered on the original CPU. When the interrupt next triggers on the original CPU, __irq_move_irq(). A vector is allocated from the new CPU, but the old vector on the original CPU remains and is r >move_in_progress is flagged, and the reclaiming process is delayed until the next trigger of the interrupt on the ne of the interrupt on the new CPU, irq_complete_move() adds a task to the old CPU's vector_cleanup list if it remain old CPU iterates over its vector_cleanup list, reclaiming old vectors. However, a rare scenario arises if the old CPU again on the new CPU. In that case irq_force_complete_move() is not invoked on the outgoing CPU to reclaim the interrupt isn't currently affine to the outgoing CPU, and irq_needs_fixup() returns false. Even though __vector_sch CPU, it doesn't reclaim apicd->prev_vector; instead, it simply resets both apicd->move_in_progress and apicd->pr remains unreclaimed in vector_matrix, leading to a CPU vector leak. To address this issue, move the invocation of irq_needs_fixup() call to reclaim apicd->prev_vector, if the interrupt is currently or used to be affine to the outgoin __vector_schedule_cleanup() as well, following a warning message, although theoretically it should never see apic pointing to an offline CPU.
CVE-2024-32002	Git is a revision control system. Prior to versions 2.45.1, 2.44.1, 2.43.4, 2.42.2, 2.41.1, 2.40.2, and 2.39.4, repositor way that exploits a bug in Git whereby it can be fooled into writing files not into the submodule's worktree but into hook that will be executed while the clone operation is still running, giving the user no opportunity to inspect the c has been patched in versions 2.45.1, 2.44.1, 2.43.4, 2.42.2, 2.41.1, 2.40.2, and 2.39.4. If symbolic link support is di core.symlinks false`), the described attack won't work. As always, it is best to avoid cloning repositories from untru
CVE-2024-32004	Git is a revision control system. Prior to versions 2.45.1, 2.44.1, 2.43.4, 2.42.2, 2.41.1, 2.40.2, and 2.39.4, an attack way that, when cloned, will execute arbitrary code during the operation. The problem has been patched in versions and 2.39.4. As a workaround, avoid cloning repositories from untrusted sources.
CVE-2024-32020	Git is a revision control system. Prior to versions 2.45.1, 2.44.1, 2.43.4, 2.42.2, 2.41.1, 2.40.2, and 2.39.4, local clo target repository's object database when source and target repository reside on the same disk. If the source repositor hardlinked files may be rewritten at any point in time by the untrusted user. Cloning local repositories will cause G repository into the target repository. This significantly speeds up such local clones compared to doing a "proper" cl time. When cloning a repository located on the same disk that is owned by a different user than the current user we files will continue to be owned and controlled by the potentially-untrusted user and can be rewritten by them at wil in versions 2.45.1, 2.44.1, 2.43.4, 2.42.2, 2.41.1, 2.40.2, and 2.39.4.

CVE-2024-32021	Git is a revision control system. Prior to versions 2.45.1, 2.44.1, 2.43.4, 2.42.2, 2.41.1, 2.40.2, and 2.39.4, when clo symlinks via the filesystem, Git may create hardlinks to arbitrary user-readable files on the same filesystem as the t Cloning a local repository over the filesystem may creating hardlinks to arbitrary user-owned files on the same files ` directory. When cloning a repository over the filesystem (without explicitly specifying the `file://` protocol or `--n will be used, which include attempting to hard link the object files instead of copying them. While the code include repository, which were added during the fix for CVE-2022-39253, these checks can still be raced because the hard If the object on the filesystem appears as a file during the check, and then a symlink during the operation, this will create hardlinks in the destination objects directory to arbitrary, user-readable files. The problem has been patched 2.41.1, 2.40.2, and 2.39.4.
CVE-2024-32465	Git is a revision control system. The Git project recommends to avoid working in untrusted repositories, and instea to obtain a clean copy. Git has specific protections to make that a safe operation even with an untrusted source repo protections to be bypassed. In the context of cloning local repositories owned by other users, this vulnerability has are circumstances where the fixes for CVE-2024-32004 are not enough: For example, when obtaining a `.zip` file c it should not be trusted by default to be safe, as e.g. hooks could be configured to run within the context of that rep versions 2.45.1, 2.44.1, 2.43.4, 2.42.2, 2.41.1, 2.40.2, and 2.39.4. As a workaround, avoid using Git in repositories untrusted sources.
CVE-2024-33621	In the Linux kernel, the following vulnerability has been resolved: ipvlan: Dont Use skb->sk in ipvlan_process_v{ PF_PACKET socket ontop of an IPv6-backed ipvlan device will hit WARN_ON_ONCE() in sk_mc_loop() throug CPU: 2 PID: 0 at net/core/sock.c:775 sk_mc_loop+0x2d/0x70 Modules linked in: sch_netem ipvlan rfkill cirrus dr CPU: 2 PID: 0 Comm: swapper/2 Kdump: loaded Not tainted 6.9.0+ #279 Hardware name: QEMU Standard PC (i 1.15.0-1 04/01/2014 RIP: 0010:sk_mc_loop+0x2d/0x70 Code: fa 0f 1f 44 00 00 65 0f b7 15 f7 96 a3 4f 31 c0 66 8 0018:ffffa9584015cd78 EFLAGS: 00010212 RAX: 0000000000000011 RBX: ffff91e585793e00 RCX: 000000000 RSI: 0000000000000040 RDI: ffff91e589c0f000 RBP: ffff91e5855bd100 R08: 0000000000000000 R09: 3d00545 R11: 00000060dd8616f4 R12: ffff91e58132d000 R13: ffff91e584fdcc68 R14: ffff91e5869ce800 R15: ffff91e589c GS:ffff91e898100000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 C 0000000008e1a000 CR4: 00000000000006f0 DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000 00000000fffe0ff0 DR7: 0000000000000400 Call Trace: <IRQ> ? __warn (kernel/panic.c:693) ? sk_mc_loop (net/c lib/bug.c:219) ? handle_bug (arch/x86/kernel/traps.c:239) ? exc_invalid_op (arch/x86/kernel/traps.c:260 (discrimin include/asm/idtentry.h:621) ? sk_mc_loop (net/core/sock.c:760) ip6_finish_output2 (net/ipv6/ip6_output.c:83 (disc core.c:626) ip6_finish_output (net/ipv6/ip6_output.c:222) ? __pfx_ip6_finish_output (net/ipv6/ip6_output.c:215) ip ipvlan_core.c:602) ipvlan ipvlan_start_xmit (drivers/net/ipvlan/ipvlan_main.c:226) ipvlan dev_hard_start_xmit (ne sched/sch_generic.c:343) __qdisc_run (net/sched/sch_generic.c:416) net_tx_action (net/core/dev.c:5286) handle_s (kernel/softirq.c:589) sysvec_apic_timer_interrupt (arch/x86/kernel/apic/apic.c:1043) The warning triggers as this: >sk is packet sk __dev_queue_xmit __dev_xmit_skb //q->enqueue is not NULL __qdisc_run sch_direct_xmit dev_ ipvlan_xmit_mode_l3 //l3 mode ipvlan_process_outbound //vepa flag ipvlan_process_v6_outbound ip6_local_out multicast packet sk_mc_loop //sk->sk_family is AF_PACKET Call ip{6}_local_out() with NULL sk in ipvlan as o
CVE-2024-35789	In the Linux kernel, the following vulnerability has been resolved: wifi: mac80211: check/clear fast rx for non-4ad out of a VLAN and deleting the VLAN afterwards, the fast_rx entry still holds a pointer to the VLAN's netdev, wh immediately calling ieee80211_check_fast_rx after the VLAN change.
CVE-2024-35804	In the Linux kernel, the following vulnerability has been resolved: KVM: x86: Mark target gfn of emulated atomic an atomic access on behalf of the guest, mark the target gfn dirty if the CMPXCHG by KVM is attempted and does effectively corrupts guest memory during live migration by writing to guest memory without informing userspace t got unintentionally dropped when KVM's emulated CMPXCHG was converted to do a user access. Before that, KV kernel memory, and marked the page dirty during the unmap phase. Mark the page dirty even if the CMPXCHG fa i.e. the page is still written. The value written is guaranteed to be the same because the operation is atomic, but KV regardless of the value written. And more importantly, that's what KVM did before the buggy commit. Huge kudos who did all the actual work of triaging and debugging. base-commit: 6769ea8da8a93ed4630f1ce64df6aafcaabfce64
CVE-2024-35806	In the Linux kernel, the following vulnerability has been resolved: soc: fsl: qbman: Always disable interrupts when disables IRQs when executing the callback. To prevent deadlocks, we must disable IRQs when taking cgr_lock els qman_update_cgr and qman_delete_cgr; fix the other lockers.
CVE-2024-35807	In the Linux kernel, the following vulnerability has been resolved: ext4: fix corruption during on-line resize We ob of a file system that is larger than 16 TiB with 4k block size. With having more then 2^32 blocks resize_inode is tu can be reproduced on a smaller file system for convenience by explicitly turning off resize_inode. An on-line resiz meta block group in this setup) then leads to a corruption: dev=/dev/<some_dev> # should be >= 16 GiB mkdir -p ^resize_inode $dev $((2 * 2**21 - 2**15)) mount -t ext4 $dev /corruption dd if=/dev/zero bs=4096 of=/corruption/ corruption/test # 79d2658b39dcfd77274e435b0934028adafaab11 /corruption/test /sbin/resize2fs $dev $((2*2**21) from disk echo 1 > /proc/sys/vm/drop_caches sha1sum /corruption/test # 3c2abc63cbf1a94c9e6977e0fbd72cd832c equals 8 GiB whereof 2^15 is the number of blocks per block group and 2^6 are the number of block groups that m might be different depending on how the file is laid out across the physical blocks. The actual corruption occurs at would be the location of the backup of the meta block group's block descriptor. During the on-line resize the file sy s_first_meta_bg which is 2 in the example - meaning all block groups after 16 GiB. However, in ext4_flex_group_ part of the first meta block group yet. In the reproducer we achieved this by substracting the size of a whole block g group would start. This must be considered when updating the backup block group descriptors to follow the non-m the group to add is already part of the meta block group or not.

CVE-2024-35809	In the Linux kernel, the following vulnerability has been resolved: PCI/PM: Drain runtime-idle callbacks before dr the .runtime_idle() callback and the .remove() callback in the rtsx_pcr PCI driver leads to a kernel crash due to an u is that rtsx_pci_runtime_idle() is not expected to be running after pm_runtime_get_sync() has been called, but the l only guarantees that the suspend and resume callbacks will not be running when it returns. However, if a .runtime_ pm_runtime_get_sync() is called, the latter will notice that the runtime PM status of the device is RPM_ACTIVE a for the former to complete. In fact, it cannot wait for .runtime_idle() to complete because it may be called from tha sense to do that, but it is not strictly prohibited). Thus in general, whoever is providing a .runtime_idle() callback n with whatever code runs after pm_runtime_get_sync(). [Note that .runtime_idle() will not start after pm_runtime_g running then if it has started earlier.] One way to address that race condition is to call pm_runtime_barrier() after pr a nonzero value of the runtime PM usage counter is necessary to prevent runtime PM callbacks from being invoked complete should it be running at that point. A suitable place for doing that is in pci_device_remove() which calls pr driver, so it may as well call pm_runtime_barrier() subsequently, which will prevent the race in question from occu any PCI drivers providing .runtime_idle() callbacks.
CVE-2024-35811	In the Linux kernel, the following vulnerability has been resolved: wifi: brcmfmac: Fix use-after-free bug in brcmf patch of CVE-2023-47233 : https://nvd.nist.gov/vuln/detail/CVE-2023-47233 In brcm80211 driver,it starts with th timeout worker: ->brcmf_usb_probe ->brcmf_usb_probe_cb ->brcmf_attach ->brcmf_bus_started ->brcmf_cfg802 ->INIT_WORK(&cfg->escan_timeout_work, brcmf_cfg80211_escan_timeout_worker); If we disconnect the USB to make cleanup. The invoking chain is : brcmf_usb_disconnect ->brcmf_usb_disconnect_cb ->brcmf_detach ->br the timeout woker may still be running. This will cause a use-after-free bug on cfg in brcmf_cfg80211_escan_time canceling the worker in brcmf_cfg80211_detach. [arend.vanspriel@broadcom.com: keep timer delete as is and can
CVE-2024-35813	In the Linux kernel, the following vulnerability has been resolved: mmc: core: Avoid negative index with array acc mrq.sbc in close-ended ffu") assigns prev_idata = idatas[i - 1], but doesn't check that the iterator i is greater than ze
CVE-2024-35815	In the Linux kernel, the following vulnerability has been resolved: fs/aio: Check IOCB_AIO_RW before the struct kiocb_set_cancel_fn() argument may point at a struct kiocb that is not embedded inside struct aio_kiocb. With the req->ki_ctx read happens either before the IOCB_AIO_RW test or after that test. Move the req->ki_ctx read such t test happens first.
CVE-2024-35819	In the Linux kernel, the following vulnerability has been resolved: soc: fsl: qbman: Use raw spinlock for cgr_lock s hard IRQ context, even on PREEMPT_RT, where spinlocks can sleep. So we need to use a raw spinlock for cgr_lo task. Although this bug has existed for a while, it was not apparent until commit ef2a8d5478b9 ("net: dpaa: Adjust smp_call_function_single via qman_update_cgr_safe every time a link goes up or down.
CVE-2024-35821	In the Linux kernel, the following vulnerability has been resolved: ubifs: Set page uptodate in the correct place Pag freshly allocated page uptodate before we've overwritten it with the data it's supposed to have in it will allow a simu to SetPageUptodate into ubifs_write_end(), which is after we copied the new data into the page.
CVE-2024-35822	In the Linux kernel, the following vulnerability has been resolved: usb: udc: remove warning when queue disabled message from mass storage function, WARNING: CPU: 6 PID: 3839 at drivers/usb/gadget/udc/core.c:294 usb_ep_ +0x7c/0x104 lr : fsg_main_thread+0x494/0x1b3c Root cause is mass storage function try to queue request from ma disable ep when function disable. As there is no function failure in the driver, in order to avoid effort to fix warning usb_ep_queue() to pr_debug().
CVE-2024-35823	In the Linux kernel, the following vulnerability has been resolved: vt: fix unicode buffer corruption when deleting c fixed for the VGA text buffer in commit 39cdb68c64d8 ("vt: fix memory overlapping when deleting chars in the bu memcpy() with memmove() due to the overlaping buffers.
CVE-2024-35825	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: ncm: Fix handling of zero block len with CDC_NCM_NTB_DEF_SIZE_TX set to 65536, it has been observed that we receive short packets, which co and have block length zero but still contain 1-2 valid datagrams present. According to the NCM spec: "If wBlockL short packet. In this case, the USB transfer must still be shorter than dwNtbInMaxSize or dwNtbOutMaxSize. If ex bytes are sent, and the size is a multiple of wMaxPacketSize for the given pipe, then no ZLP shall be sent. wBlockL care, because of the possibility that the host and device may get out of sync, and because of test issues. wBlockLen latency by starting to send a very large NTB, and then shortening it when the sender discovers that there,Äôs not su However, there is a potential issue with the current implementation, as it checks for the occurrence of multiple NTB leftover bytes to be processed is zero or not. If the block length reads zero, we would process the same NTB infinte and it leads to a crash. Fix this by bailing out if block length reads zero.
CVE-2024-35828	In the Linux kernel, the following vulnerability has been resolved: wifi: libertas: fix some memleaks in lbs_allocate lbs_allocate_cmd_buffer(), if the allocation of cmdarray[i].cmdbuf fails, both cmdarray and cmdarray[i].cmdbuf ne memleaks in lbs_allocate_cmd_buffer().
CVE-2024-35830	In the Linux kernel, the following vulnerability has been resolved: media: tc358743: register v4l2 async device onl been setup correctly before registering the v4l2 async device, thus allowing userspace to access.
CVE-2024-35833	In the Linux kernel, the following vulnerability has been resolved: dmaengine: fsl-qdma: Fix a memory leak relate dma_alloc_coherent() is undone neither in the remove function, nor in the error handling path of fsl_qdma_probe() issues.
CVE-2024-35834	In the Linux kernel, the following vulnerability has been resolved: xsk: recycle buffer in case Rx queue was full Ac __xsk_rcv_zc() failed to produce descriptor to XSK Rx queue.

CVE-2024-35835	In the Linux kernel, the following vulnerability has been resolved: net/mlx5e: fix a double-free in arfs_create_grou arfs_create_groups will free ft->g and return an error. However, arfs_create_table, the only caller of arfs_create_gr mlx5e_destroy_flow_table, in which the ft->g will be freed again.
CVE-2024-35836	In the Linux kernel, the following vulnerability has been resolved: dpll: fix pin dump crash for rebound module Wh resources were not entirely freed (other kernel module instance of the same PCI device have had kept the reference bound, the pin properties would not be updated (the properties are only assigned when memory for the pin is alloca module memory of the kernel module which was deallocated on the unbind. If the pin dump is invoked in this state by storing persistent pin properties in dpll subsystem, copy the content from the kernel module when pin is allocate module.
CVE-2024-35847	In the Linux kernel, the following vulnerability has been resolved: irqchip/gic-v3-its: Prevent double free on error T its_vpe_irq_domain_alloc() causes a double free when its_vpe_init() fails after successfully allocating at least one i its_vpe_irq_domain_free() frees the interrupts along with the area bitmap and the vprop_page and its_vpe_irq_dom bitmap and the vprop_page again. Fix this by unconditionally invoking its_vpe_irq_domain_free() which handles a bitmap/vprop_page freeing from its_vpe_irq_domain_alloc(). [tglx: Massaged change log]
CVE-2024-35849	In the Linux kernel, the following vulnerability has been resolved: btrfs: fix information leak in btrfs_ioctl_logical_ information leak for in btrfs_ioctl_logical_to_ino(): BUG: KMSAN: kernel-infoleak in instrument_copy_to_user in BUG: KMSAN: kernel-infoleak in _copy_to_user+0xbc/0x110 lib/usercopy.c:40 instrument_copy_to_user include _copy_to_user+0xbc/0x110 lib/usercopy.c:40 copy_to_user include/linux/uaccess.h:191 [inline] btrfs_ioctl_logical btrfs_ioctl+0x714/0x1260 vfs_ioctl fs/ioctl.c:51 [inline] __do_sys_ioctl fs/ioctl.c:904 [inline] __se_sys_ioctl+0x26 +0x96/0xe0 fs/ioctl.c:890 x64_sys_call+0x1883/0x3b50 arch/x86/include/generated/asm/syscalls_64.h:17 do_sysc [inline] do_syscall_64+0xcf/0x1e0 arch/x86/entry/common.c:83 entry_SYSCALL_64_after_hwframe+0x77/0x7f +0x231/0x370 mm/slub.c:3921 __do_kmalloc_node mm/slub.c:3954 [inline] __kmalloc_node+0xb07/0x1060 mm linux/slab.h:648 [inline] kvmalloc_node+0xc0/0x2d0 mm/util.c:634 kvmalloc include/linux/slab.h:766 [inline] init backref.c:2779 btrfs_ioctl_logical_to_ino+0x17c/0x750 fs/btrfs/ioctl.c:3480 btrfs_ioctl+0x714/0x1260 vfs_ioctl fs fs/ioctl.c:904 [inline] __se_sys_ioctl+0x261/0x450 fs/ioctl.c:890 __x64_sys_ioctl+0x96/0xe0 fs/ioctl.c:890 x64_sy include/generated/asm/syscalls_64.h:17 do_syscall_x64 arch/x86/entry/common.c:52 [inline] do_syscall_64+0xcf/ entry_SYSCALL_64_after_hwframe+0x77/0x7f Bytes 40-65535 of 65536 are uninitialized Memory access of size happens, because we're copying a 'struct btrfs_data_container' back to user-space. This btrfs_data_container is alloc which does not zero-fill the memory. Fix this by using kvzalloc() which zeroes out the memory on allocation.
CVE-2024-35852	In the Linux kernel, the following vulnerability has been resolved: mlxsw: spectrum_acl_tcam: Fix memory leak w delayed work is rescheduled with a delay if the number of credits at end of the work is not negative as supposedly i it is rescheduled immediately. After "mlxsw: spectrum_acl_tcam: Fix possible use-after-free during rehash" the abo number of credits is no longer indicative of the migration being done. It can also happen if the work encountered ar resume the next time the work is scheduled. The significance of the above is that it is possible for the work to be pe allocated when the migration started. This leads to the hints being leaked [1] when the work is canceled while pend Fix by freeing the hints if hints are associated with a work that was canceled while pending. Blame the original con pending work associated with hints is fragile. [1] unreferenced object 0xffff88810e7c3000 (size 256): comm "kwor dump (first 32 bytes): 00 30 95 11 81 88 ff ff 61 00 00 00 00 00 00 80 .0......a..... 00 00 61 00 40 00 00 00 00 00 0 (crc 2544ddb9): [<00000000cf8cfab3>] kmalloc_trace+0x23f/0x2a0 [<000000004d9a1ad9>] objagg_hints_get+0x mlxsw_sp_acl_erp_rehash_hints_get+0xca/0x400 [<0000000059bdb60a>] mlxsw_sp_acl_tcam_vregion_rehash_w process_one_work+0x59c/0xf20 [<00000000ceee9e81>] worker_thread+0x799/0x12c0 [<00000000bda6fe39>] kt ret_from_fork+0x34/0x70 [<00000000dea2b93e>] ret_from_fork_asm+0x1a/0x30
CVE-2024-35853	In the Linux kernel, the following vulnerability has been resolved: mlxsw: spectrum_acl_tcam: Fix memory leak du migrates filters from one region to another. This is done by iterating over all chunks (all the filters with the same pr iterating over all the filters. If the migration fails, the code tries to migrate the filters back to the old region. Howev case another migration will be erroneously performed. Besides the fact that this ping pong is not a very good idea, i references two chunks: The currently used one ('vchunk->chunk') and a backup ('vchunk->chunk2'). During migrat migrate filters to and the second holds the chunk we are migrating filters from. The code currently assumes - but do not exist (NULL) if the currently used chunk does not reference the target region. This assumption breaks when we in the backup chunk being overwritten and leaked [1]. Fix by not rolling back a failed rollback and add a warning t CPU: 5 PID: 1063 at lib/parman.c:291 parman_destroy+0x17/0x20 Modules linked in: CPU: 5 PID: 1063 Comm: k custom-00784-gc6a05c468a0b #14 Hardware name: Mellanox Technologies Ltd. MSN3700/VMOD0005, BIOS 5. mlxsw_sp_acl_tcam_vregion_rehash_work RIP: 0010:parman_destroy+0x17/0x20 [...] Call Trace: <TASK> mlxsw mlxsw_sp_acl_tcam_region_destroy+0x49/0xf0 mlxsw_sp_acl_tcam_vregion_rehash_work+0x1f1/0x470 process +0x2cb/0x3e0 kthread+0xd0/0x100 ret_from_fork+0x34/0x50 ret_from_fork_asm+0x1a/0x30 </TASK>

CVE-2024-35854	In the Linux kernel, the following vulnerability has been resolved: mlxsw: spectrum_acl_tcam: Fix possible use-aft work migrates filters from one region to another according to the number of available credits. The migrated from re if the number of credits is non-negative as the assumption is that this is indicative of migration being complete. Thi negative number of credits can also be the result of a failed migration. The destruction of a region that still has filte after-free [1]. Fix by not destroying the region if migration failed. [1] BUG: KASAN: slab-use-after-free in mlxsw_ +0x21d/0x230 Read of size 8 at addr ffff8881735319e8 by task kworker/0:31/3858 CPU: 0 PID: 3858 Comm: kwo custom-00782-gf2275c2157d8 #5 Hardware name: Mellanox Technologies Ltd. MSN3700/VMOD0005, BIOS 5.1 mlxsw_sp_acl_tcam_vregion_rehash_work Call Trace: <TASK> dump_stack_lvl+0xc6/0x120 print_report+0xce/0 mlxsw_sp_acl_ctcam_region_entry_remove+0x21d/0x230 mlxsw_sp_acl_ctcam_entry_del+0x2e/0x70 mlxsw_sp_ mlxsw_sp_acl_tcam_vchunk_migrate_all+0x3cd/0xb50 mlxsw_sp_acl_tcam_vregion_rehash_work+0x157/0x130 worker_thread+0x6c9/0xf70 kthread+0x2c9/0x3b0 ret_from_fork+0x4d/0x80 ret_from_fork_asm+0x1a/0x30 </TA kasan_save_stack+0x33/0x60 kasan_save_track+0x14/0x30 __kasan_kmalloc+0x8f/0xa0 __kmalloc+0x19c/0x360 +0xdf/0x9c0 mlxsw_sp_acl_tcam_vregion_rehash_work+0x954/0x1300 process_one_work+0x8eb/0x19b0 worker ret_from_fork+0x4d/0x80 ret_from_fork_asm+0x1a/0x30 Freed by task 7: kasan_save_stack+0x33/0x60 kasan_sa +0x3b/0x60 poison_slab_object+0x102/0x170 __kasan_slab_free+0x14/0x30 kfree+0xc1/0x290 mlxsw_sp_acl_tc mlxsw_sp_acl_tcam_vregion_rehash_work+0x731/0x1300 process_one_work+0x8eb/0x19b0 worker_thread+0x6 +0x4d/0x80 ret_from_fork_asm+0x1a/0x30
CVE-2024-35855	In the Linux kernel, the following vulnerability has been resolved: mlxsw: spectrum_acl_tcam: Fix possible use-aft activity update delayed work periodically traverses the list of configured rules and queries their activity from the de entry pointed by 'ventry->entry', but this entry can be changed concurrently by the rehash delayed work, leading to race and perform the activity query under the 'vregion->lock' mutex. [1] BUG: KASAN: slab-use-after-free in mlxs +0x121/0x140 Read of size 8 at addr ffff8881054ed808 by task kworker/0:18/181 CPU: 0 PID: 181 Comm: kwork custom-00781-gd5ab772d32f7 #2 Hardware name: Mellanox Technologies Ltd. MSN3700/VMOD0005, BIOS 5.1 mlxsw_sp_acl_rule_activity_update_work Call Trace: <TASK> dump_stack_lvl+0xc6/0x120 print_report+0xce/0 mlxsw_sp_acl_tcam_flower_rule_activity_get+0x121/0x140 mlxsw_sp_acl_rule_activity_update_work+0x219/0x worker_thread+0x6c9/0xf70 kthread+0x2c9/0x3b0 ret_from_fork+0x4d/0x80 ret_from_fork_asm+0x1a/0x30 </TA kasan_save_stack+0x33/0x60 kasan_save_track+0x14/0x30 __kasan_kmalloc+0x8f/0xa0 __kmalloc+0x19c/0x360 +0x7b/0x1f0 mlxsw_sp_acl_tcam_vchunk_migrate_all+0x30d/0xb50 mlxsw_sp_acl_tcam_vregion_rehash_work+ +0x8eb/0x19b0 worker_thread+0x6c9/0xf70 kthread+0x2c9/0x3b0 ret_from_fork+0x4d/0x80 ret_from_fork_asm+ kasan_save_stack+0x33/0x60 kasan_save_track+0x14/0x30 kasan_save_free_info+0x3b/0x60 poison_slab_object+ kfree+0xc1/0x290 mlxsw_sp_acl_tcam_vchunk_migrate_all+0x3d7/0xb50 mlxsw_sp_acl_tcam_vregion_rehash_v +0x8eb/0x19b0 worker_thread+0x6c9/0xf70 kthread+0x2c9/0x3b0 ret_from_fork+0x4d/0x80 ret_from_fork_asm+
CVE-2024-35877	In the Linux kernel, the following vulnerability has been resolved: x86/mm/pat: fix VM_PAT handling in COW ma thing in COW mappings: the first PTE (or, in fact, all PTEs) can be replaced during write faults to point at anon fol cachemode using follow_phys() from PTEs will not work in COW mappings. Using follow_phys(), we might just g (which is very wrong), or fail on swap/nonswap entries, failing follow_phys() and triggering a WARN_ON_ONCE properly calling free_pfn_range(). In free_pfn_range(), we either wouldn't call memtype_free() or would call it wit To fix that, let's update follow_phys() to refuse returning anon folios, and fallback to using the stored PFN inside v into that. We will now properly handle untrack_pfn() with COW mappings, where we don't need the cachemode. W if the first page was replaced by an anon folio, though: we'd have to store the cachemode in the VMA to make this now, lets keep it simple and let track_pfn_copy() just fail in that case: it would have failed in the past with swap/no done the wrong thing with anon folios. Simple reproducer to trigger the WARN_ON_ONCE() in untrack_pfn(): <- #include <sys/mman.h> #include <unistd.h> #include <liburing.h> int main(void) { struct io_uring_params p = {} = io_uring_setup(1, &p); if (ring_fd < 0) { perror("io_uring_setup"); return 1; } size = p.sq_off.array + p.sq_entries queue ring MAP_PRIVATE */ map = mmap(0, size, PROT_READ | PROT_WRITE, MAP_PRIVATE, ring_fd, IO MAP_FAILED) { perror("mmap"); return 1; } /* We have at least one page. Let's COW it. */ *map = 0; pause(); re system with 16 GiB RAM and swap configured: # ./iouring & # memhog 16G # killall iouring [301.552930] ------ WARNING: CPU: 7 PID: 1402 at arch/x86/mm/pat/memtype.c:1060 untrack_pfn+0xf4/0x100 [301.553989] Mod nft_fib_ipv4 nft_fib_ipv6 nft_fib nft_reject_g [301.558232] CPU: 7 PID: 1402 Comm: iouring Not tainted 6.7.5-1 name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS rel-1.16.3-0-ga6ed6b701f0a-prebu4 [301.559569] RIP: 00 Code: 75 c4 eb cf 48 8b 43 10 8b a8 e8 00 00 00 3b 6b 28 74 b8 48 8b 7b 30 e8 ea 1a f7 000 [301.561189] RSP: 0 [301.561590] RAX: 00000000ffffffea RBX: ffff9208c8ce9cc0 RCX: 000000010455e047 [301.562105] RDX: 07 RDI: ffff9208c391d200 [301.562628] RBP: 0000000000000000 R08: ffffba2c0377fab8 R09: 0000000000000000 R11: 0000000000000002 R12: 00007fea890e0000 [301.563669] R13: 0000000000000000 R14: ffffba2c0377fc08 FS: 0000000000000000(0000) GS:ffff920c2fbc0000(0000) knlGS:0000000000000000 [301.564773] CS: 0010 DS [301.565197] CR2: 00007fea88ee8a20 CR3: 00000001033a8000 CR4: 0000000000750ef0 [301.565725] PKRU: [301.566148] <TASK> [301.566325] ? untrack_pfn+0xf4/0x100 [301.566618] ? __warn+0x81/0x130 [301.5668 truncated---
CVE-2024-35879	In the Linux kernel, the following vulnerability has been resolved: of: dynamic: Synchronize of_changeset_destroy sequence: 1) of_platform_depopulate() 2) of_overlay_remove() During the step 1, devices are destroyed and devlin are destroyed but __of_changeset_entry_destroy() can raise warnings related to missing of_node_put(): ERROR: m 2 ... Indeed, during the devlink removals performed at step 1, the removal itself releasing the device (and the attach workqueue and so, it is done asynchronously with respect to function calls. When the warning is present, of_node_ the workqueue job. In order to be sure that any ongoing devlink removals are done before the of_node destruction, the devlink removals.

CVE-2024-35888	In the Linux kernel, the following vulnerability has been resolved: erspan: make sure erspan_base_hdr is present in in ip6erspan_rcv() [1] Issue is that ip6erspan_rcv() (and erspan_rcv()) no longer make sure erspan_base_hdr is pres before getting @ver field from it. Add the missing pskb_may_pull() calls. v2: Reload iph pointer in erspan_rcv() af >head might have changed. [1] BUG: KMSAN: uninit-value in pskb_may_pull_reason include/linux/skbuff.h:2742 in pskb_may_pull include/linux/skbuff.h:2756 [inline] BUG: KMSAN: uninit-value in ip6erspan_rcv net/ipv6/ip6_ uninit-value in gre_rcv+0x11f8/0x1930 net/ipv6/ip6_gre.c:610 pskb_may_pull_reason include/linux/skbuff.h:2742 skbuff.h:2756 [inline] ip6erspan_rcv net/ipv6/ip6_gre.c:541 [inline] gre_rcv+0x11f8/0x1930 net/ipv6/ip6_gre.c:61 net/ipv6/ip6_input.c:438 ip6_input_finish net/ipv6/ip6_input.c:483 [inline] NF_HOOK include/linux/netfilter.h:31 net/ipv6/ip6_input.c:492 ip6_mc_input+0xa7e/0xc80 net/ipv6/ip6_input.c:586 dst_input include/net/dst.h:460 [inli ipv6/ip6_input.c:79 NF_HOOK include/linux/netfilter.h:314 [inline] ipv6_rcv+0xde/0x390 net/ipv6/ip6_input.c:31 core/dev.c:5538 [inline] __netif_receive_skb+0x1da/0xa00 net/core/dev.c:5652 netif_receive_skb_internal net/core +0x58/0x660 net/core/dev.c:5798 tun_rx_batched+0x3ee/0x980 drivers/net/tun.c:1549 tun_get_user+0x5566/0x69 +0x3af/0x5d0 drivers/net/tun.c:2048 call_write_iter include/linux/fs.h:2108 [inline] new_sync_write fs/read_write. fs/read_write.c:590 ksys_write+0x20f/0x4c0 fs/read_write.c:643 __do_sys_write fs/read_write.c:655 [inline] __se_ __x64_sys_write+0x93/0xe0 fs/read_write.c:652 do_syscall_64+0xd5/0x1f0 entry_SYSCALL_64_after_hwframe at: slab_post_alloc_hook mm/slub.c:3804 [inline] slab_alloc_node mm/slub.c:3845 [inline] kmem_cache_alloc_no kmalloc_reserve+0x13d/0x4a0 net/core/skbuff.c:577 __alloc_skb+0x35b/0x7a0 net/core/skbuff.c:668 alloc_skb in alloc_skb_with_frags+0xc8/0xbf0 net/core/skbuff.c:6504 sock_alloc_send_pskb+0xa81/0xbf0 net/core/sock.c:279 [inline] tun_get_user+0x209a/0x69e0 drivers/net/tun.c:1846 tun_chr_write_iter+0x3af/0x5d0 drivers/net/tun.c:204 [inline] new_sync_write fs/read_write.c:497 [inline] vfs_write+0xb63/0x1520 fs/read_write.c:590 ksys_write+0x2 fs/read_write.c:655 [inline] __se_sys_write fs/read_write.c:652 [inline] __x64_sys_write+0x93/0xe0 fs/read_write entry_SYSCALL_64_after_hwframe+0x6d/0x75 CPU: 1 PID: 5045 Comm: syz-executor114 Not tainted 6.9.0-rc1
CVE-2024-35891	In the Linux kernel, the following vulnerability has been resolved: net: phy: micrel: Fix potential null pointer deref lan8814_get_sig_tx() ptp_parse_header() may return NULL as ptp_header due to abnormal packet type or corrupte check. Found by Linux Verification Center (linuxtesting.org) with SVACE.
CVE-2024-35893	In the Linux kernel, the following vulnerability has been resolved: net/sched: act_skbmod: prevent kernel-infoleak copying four bytes from kernel stack to user space [1]. The issue here is that 'struct tc_skbmod' has a four bytes hol filling fields. [1] BUG: KMSAN: kernel-infoleak in instrument_copy_to_user include/linux/instrumented.h:114 [in in copy_to_user_iter lib/iov_iter.c:24 [inline] BUG: KMSAN: kernel-infoleak in iterate_ubuf include/linux/iov_ite infoleak in iterate_and_advance2 include/linux/iov_iter.h:245 [inline] BUG: KMSAN: kernel-infoleak in iterate_ar [inline] BUG: KMSAN: kernel-infoleak in _copy_to_iter+0x366/0x2520 lib/iov_iter.c:185 instrument_copy_to_us [inline] copy_to_user_iter lib/iov_iter.c:24 [inline] iterate_ubuf include/linux/iov_iter.h:29 [inline] iterate_and_adv [inline] iterate_and_advance include/linux/iov_iter.h:271 [inline] _copy_to_iter+0x366/0x2520 lib/iov_iter.c:185 c [inline] simple_copy_to_iter net/core/datagram.c:532 [inline] __skb_datagram_iter+0x185/0x1000 net/core/datagra +0x5c/0x200 net/core/datagram.c:546 skb_copy_datagram_msg include/linux/skbuff.h:4050 [inline] netlink_recvm af_netlink.c:1962 sock_recvmsg_nosec net/socket.c:1046 [inline] sock_recvmsg+0x2c4/0x340 net/socket.c:1068 _ socket.c:2242 __do_sys_recvfrom net/socket.c:2260 [inline] __se_sys_recvfrom net/socket.c:2256 [inline] __x64_ do_syscall_64+0xd5/0x1f0 entry_SYSCALL_64_after_hwframe+0x6d/0x75 Uninit was stored to memory at: pskb skbuff.c:2253 netlink_trim+0x2c2/0x330 net/netlink/af_netlink.c:1317 netlink_unicast+0x9f/0x1260 net/netlink/af net/netlink.h:1144 [inline] nlmsg_notify+0x21d/0x2f0 net/netlink/af_netlink.c:2610 rtnetlink_send+0x73/0x90 net/ include/linux/rtnetlink.h:17 [inline] tcf_add_notify net/sched/act_api.c:2048 [inline] tcf_action_add net/sched/act_a +0x146e/0x19d0 net/sched/act_api.c:2119 rtnetlink_rcv_msg+0x1737/0x1900 net/core/rtnetlink.c:6595 netlink_rcv af_netlink.c:2559 rtnetlink_rcv+0x34/0x40 net/core/rtnetlink.c:6613 netlink_unicast_kernel net/netlink/af_netlink.c net/netlink/af_netlink.c:1361 netlink_sendmsg+0x10df/0x11f0 net/netlink/af_netlink.c:1905 sock_sendmsg_nosec +0x30f/0x380 net/socket.c:745 ____sys_sendmsg+0x877/0xb60 net/socket.c:2584 ___sys_sendmsg+0x28d/0x3c0 socket.c:2667 [inline] __do_sys_sendmsg net/socket.c:2676 [inline] __se_sys_sendmsg net/socket.c:2674 [inline] _ socket.c:2674 do_syscall_64+0xd5/0x1f0 entry_SYSCALL_64_after_hwframe+0x6d/0x75 Uninit was stored to m nla_put+0x1c6/0x230 lib/nlattr.c:1099 tcf_skbmod_dump+0x23f/0xc20 net/sched/act_skbmod.c:256 tcf_action_du tcf_action_dump_1+0x85e/0x970 net/sched/act_api.c:1227 tcf_action_dump+0x1fd/0x460 net/sched/act_api.c:125 act_api.c:1628 tcf_add_notify_msg net/sched/act_api.c:2023 [inline] tcf_add_notify net/sched/act_api.c:2042 [inli [inline] tc_ctl_action+0x1365/0x19d0 net/sched/act_api.c:2119 rtnetlink_rcv_msg+0x1737/0x1900 net/core/rtnetli netlink/af_netli ---truncated---
CVE-2024-35897	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_tables: discard table flag update wit unregistration is deferred to the commit phase, same occurs with hook updates triggered by the table dormant flag. results in deleting a basechain while leaving its hook still registered in the core.
CVE-2024-35898	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_tables: Fix potential data-race in __ nft_unregister_flowtable_type() within nf_flow_inet_module_exit() can concurrent with __nft_flowtable_type_get And thhere is not any protection when iterate over nf_tables_flowtables list in __nft_flowtable_type_get(). Therefo nf_tables_flowtables list entry. Use list_for_each_entry_rcu() to iterate over nf_tables_flowtables list in __nft_flow the caller nft_flowtable_type_get() to protect the entire type query process.

CVE-2024-35899	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_tables: flush pending destroy work 2c9f0293280e ("netfilter: nf_tables: flush pending destroy work before netlink notifier") to address a race between The trace below shows an element to be released via destroy workqueue while exit_net path (triggered via module set that is used in such transaction. [1360.547789] BUG: KASAN: slab-use-after-free in nf_tables_trans_destroy_ [1360.547861] Read of size 8 at addr ffff888140500cc0 by task kworker/4:1/152465 [1360.547870] CPU: 4 PID: Not tainted 6.8.0+ #359 [1360.547882] Workqueue: events nf_tables_trans_destroy_work [nf_tables] [1360.5479 <TASK> [1360.547998] dump_stack_lvl+0x53/0x70 [1360.548014] print_report+0xc4/0x610 [1360.548026] ? [1360.548040] ? __pfx__raw_spin_lock_irqsave+0x10/0x10 [1360.548054] ? nf_tables_trans_destroy_work+0x3 kasan_report+0xae/0xe0 [1360.548189] ? nf_tables_trans_destroy_work+0x3f5/0x590 [nf_tables] [1360.548312] +0x3f5/0x590 [nf_tables] [1360.548447] ? __pfx_nf_tables_trans_destroy_work+0x10/0x10 [nf_tables] [1360.54 +0x18/0x30 [1360.548591] process_one_work+0x2f1/0x670 [1360.548610] worker_thread+0x4d3/0x760 [1360. +0x10/0x10 [1360.548640] kthread+0x16b/0x1b0 [1360.548653] ? __pfx_kthread+0x10/0x10 [1360.548665] ret __pfx_kthread+0x10/0x10 [1360.548690] ret_from_fork_asm+0x1a/0x30 [1360.548707] </TASK> [1360.54871 [1360.548726] kasan_save_stack+0x20/0x40 [1360.548739] kasan_save_track+0x14/0x30 [1360.548750] __kas __kmalloc_node+0x1f1/0x450 [1360.548771] nf_tables_newset+0x10c7/0x1b50 [nf_tables] [1360.548883] nfnet [nfnetlink] [1360.548909] nfnetlink_rcv+0x1a8/0x1e0 [nfnetlink] [1360.548927] netlink_unicast+0x367/0x4f0 [+0x34b/0x610 [1360.548944] ____sys_sendmsg+0x4d4/0x510 [1360.548953] ___sys_sendmsg+0xc9/0x120 [13 [1360.548971] do_syscall_64+0x55/0x120 [1360.548982] entry_SYSCALL_64_after_hwframe+0x55/0x5d [136 [1360.548999] kasan_save_stack+0x20/0x40 [1360.549009] kasan_save_track+0x14/0x30 [1360.549019] kasan_ poison_slab_object+0x100/0x180 [1360.549036] __kasan_slab_free+0x14/0x30 [1360.549042] kfree+0xb6/0x26 +0x473/0x6a0 [nf_tables] [1360.549131] nf_tables_exit_net+0x170/0x240 [nf_tables] [1360.549221] ops_exit_li +0x101/0x140 [1360.549236] unregister_pernet_operations+0x107/0x160 [1360.549245] unregister_pernet_subs nf_tables_module_exit+0x43/0x80 [nf_tables] [1360.549345] __do_sys_delete_module+0x253/0x370 [1360.549 [1360.549360] entry_SYSCALL_64_after_hwframe+0x55/0x5d (gdb) list *__nft_release_table+0x473 0x1e033 is nf_tables_api.c:11354). 11349 list_for_each_entry_safe(flowtable, nf, &table->flowtables, list) { 11350 list_del(& >use); 11352 nf_tables_flowtable_destroy(flowtable); 11353 } 11354 list_for_each_entry_safe(set, ns, &table->set nft_use_dec(&table->use); 11357 if (set->flags & (NFT_SET_MAP | NFT_SET_OBJECT)) 11358 nft_map_deact
CVE-2024-35900	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_tables: reject new basechain after ta is toggled, hooks are disabled in the commit phase by iterating over current chains in table (existing and new). The an inconsistent state: add table x add chain x y { type filter hook input priority 0; } add table x { flags dormant; } a priority 1; } which triggers the following warning when trying to unregister chain w which is already unregistered. PID: 1211 at net/netfilter/core.c:50 1 __nf_unregister_net_hook+0x21a/0x260 [...] [127.322519] Call Trace: [127 __warn+0x9f/0x1a0 [127.322531] ? __nf_unregister_net_hook+0x21a/0x260 [127.322537] ? report_bug+0x1b1/ +0x3c/0x70 [127.322552] ? exc_invalid_op+0x17/0x40 [127.322556] ? asm_exc_invalid_op+0x1a/0x20 [127.32 [127.322570] ? __nf_unregister_net_hook+0x6a/0x260 [127.322577] ? __nf_unregister_net_hook+0x21a/0x260 +0x6a/0x260 [127.322590] ? __nf_tables_unregister_hook+0x8a/0xe0 [nf_tables] [127.322655] nft_table_disable nf_tables_commit+0x2571/0x2620 [nf_tables]
CVE-2024-35903	In the Linux kernel, the following vulnerability has been resolved: x86/bpf: Fix IP after emitting call depth account calculates the correct offset for the CALL instruction if `x86_call_depth_emit_accounting` emits code. Otherwise v crash.
CVE-2024-35910	In the Linux kernel, the following vulnerability has been resolved: tcp: properly terminate timers for kernel sockets timers firing after the corresponding netns has been dismantled. Fortunately Josef Bacik could trigger the issue mo years ago. When TCP sockets are closed, we call inet_csk_clear_xmit_timers() to 'stop' the timers. inet_csk_clear_ including when socket lock is held. This is the reason it uses sk_stop_timer(), aka del_timer(). This means that ong sockets, this is fine because each running timer holds a reference on the socket, and the user socket holds a referenc that the netns is freed before timer can complete, because kernel sockets do not hold reference on the netns. This pa function that using sk_stop_timer_sync() to make sure all timers are terminated before the kernel socket is released their netns exit() handler. Also add sock_not_owned_by_me() helper to get LOCKDEP support : inet_csk_clear_x socket lock is held. It is very possible we can revert in the future commit 3a58f13a881e ("net: rds: acquire refcount the issue in rds only. (net/smc/af_smc.c and net/mptcp/subflow.c have similar code) We probably can remove the c and __tcp_close() in the future.
CVE-2024-35915	In the Linux kernel, the following vulnerability has been resolved: nfc: nci: Fix uninit-value in nci_dev_up and nci uninit-value access issue [1][2]: nci_rx_work() parses and processes received packet. When the payload length is z uninitialized payload and KMSAN detects this issue. The receipt of a packet with a zero-size payload is considered should be silently discarded. This patch resolved this issue by checking payload size before calling each message ty
CVE-2024-35917	In the Linux kernel, the following vulnerability has been resolved: s390/bpf: Fix bpf_plt pointer arithmetic Kui-Fe triggered by the dummy_st_ops/dummy_init_ptr_arg test [1]: [<0000000000000002>] 0x2 [<00000000009d5cde> [<000000000033145a>] __sys_bpf+0xa1a/0xd00 [<00000000003319dc>] __s390x_sys_bpf+0x44/0x50 [<000000 [<0000000000c59a40>] system_call+0x70/0x98 This is caused by GCC moving memcpy() after assignments in bp written instead of the return and the target addresses. Looking at the GCC internals, the reordering is allowed becau destination and the assignments' left-hand-sides are based on different objects: new_plt and bpf_plt_ret/bpf_plt_tar alias. This is in turn due to a violation of the C standard: When two pointers are subtracted, both shall point to elem last element of the array object ... From the C's perspective, bpf_plt_ret and bpf_plt are distinct objects and cannot confuses the GCC's alias analysis. The code was written this way in order to let the C side know a few offsets defi means necessary. Fix the noncompliance by hardcoding these offsets. [1] https://lore.kernel.org/bpf/c9923c1d-971

CVE-2024-35925	In the Linux kernel, the following vulnerability has been resolved: block: prevent division by zero in blk_rq_stat_s >nr_samples may have zero value on overflow. It is necessary to add a check to avoid division by zero. Found by L with Svace.
CVE-2024-35930	In the Linux kernel, the following vulnerability has been resolved: scsi: lpfc: Fix possible memory leak in lpfc_rcv in lpfc_rcv_padisc() may return an unsuccessful status. In such cases, the elsiocb is not issued, the completion is no leaked. Check return value after calling lpfc_sli4_resume_rpi() and conditionally release the elsiocb resource.
CVE-2024-35935	In the Linux kernel, the following vulnerability has been resolved: btrfs: send: handle path ref underflow in header proper error handling if building the path buffer fails. The pointers are not printed so we don't accidentally leak ker
CVE-2024-35936	In the Linux kernel, the following vulnerability has been resolved: btrfs: handle chunk tree lookup error in btrfs_rel btrfs_relocate_sys_chunks() loop is a corruption, as it could be caused only by two impossible conditions: - at first tree item, with offset -1, this is an inexact search and the key->offset will contain the correct offset upon a successf an offset -1 - after first successful search, the found_key corresponds to a chunk item, the offset is decremented by a chunk item there due to alignment and size constraints
CVE-2024-35944	In the Linux kernel, the following vulnerability has been resolved: VMCI: Fix memcpy() run-time warning in dg_d in dg_dispatch_as_host' bug. memcpy: detected field-spanning write (size 56) of single field "&dg_info->msg" at c (size 24) WARNING: CPU: 0 PID: 1555 at drivers/misc/vmw_vmci/vmci_datagram.c:237 dg_dispatch_as_host+0 vmci_datagram.c:237 Some code commentry, based on my understanding: 544 #define VMCI_DG_SIZE(_dg) (VM (_dg)->payload_size) /// This is 24 + payload_size memcpy(&dg_info->msg, dg, dg_size); Destination = dg_info-> vmci_datagram) Source = dg --> this is a 24 byte structure (struct vmci_datagram) Size = dg_size = 24 + payload_s managed to set payload_size to 32. 35 struct delayed_datagram_info { 36 struct datagram_entry *entry; 37 struct w 39 /* msg and msg_payload must be together. */ 40 struct vmci_datagram msg; 41 u8 msg_payload[]; 42 }; So tho msg_payload[], a run time warning is seen while fuzzing with Syzkaller. One possible way to fix the warning is to direct assignment of msg and second taking care of payload. Gustavo quoted: "Under FORTIFY_SOURCE we sho a structure."
CVE-2024-35947	In the Linux kernel, the following vulnerability has been resolved: dyndbg: fix old BUG_ON in >control parser Fix "unreachable" (I didn't really look), lets make sure by removing it, doing pr_err and return -EINVAL instead.
CVE-2024-35955	In the Linux kernel, the following vulnerability has been resolved: kprobes: Fix possible use-after-free issue on kpr a module, its state is changing MODULE_STATE_LIVE -> MODULE_STATE_GOING -> MODULE_STATE_U take a time. `is_module_text_address()` and `__module_text_address()` works with MODULE_STATE_LIVE and If we use `is_module_text_address()` and `__module_text_address()` separately, there is a chance that the first one failed because module->state becomes MODULE_STATE_UNFORMED between those operations. In `check_kpr `__module_text_address()` is failed, that is ignored because it expected a kernel_text address. But it may have faile been changed to MODULE_STATE_UNFORMED. In this case, arm_kprobe() will try to modify non-exist module problem, we should not use separated `is_module_text_address()` and `__module_text_address()`, but use only `__ `try_module_get(module)` which is only available with MODULE_STATE_LIVE.
CVE-2024-35960	In the Linux kernel, the following vulnerability has been resolved: net/mlx5: Properly link new fs rules into the tree add newly created rules from the handle into the tree when they had a refcount of 1. On the other hand, create_flow already existing identical rules instead of creating new ones. These two behaviors can result in a situation where cr references it, then 2) in a subsequent step during the same handle creation references it again, resulting in a rule wit tree, will have a NULL parent and root and will result in a crash when the flow group is deleted because del_sw_hw node->parent is != NULL. This happened in the wild, due to another bug related to incorrect handling of duplicate create_flow_handle incorrectly referencing a just-added rule in the same flow handle, resulting in the problem desc changes add_rule_fg to add new rules without parents into the tree, properly initializing them and avoiding the cras rules are added to an FTE in create_flow_handle.
CVE-2024-35964	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: ISO: Fix not validating setsockopt us copying data.
CVE-2024-35972	In the Linux kernel, the following vulnerability has been resolved: bnxt_en: Fix possible memory leak in bnxt_rdm the allocated edev will leak because it is not properly assigned and the cleanup path will not be able to free it. Fix it allocation.

CVE-2024-35973	In the Linux kernel, the following vulnerability has been resolved: geneve: fix header validation in geneve[6]_xmit value in geneve_xmit() [1] Problem : While most ip tunnel helpers (like ip_tunnel_get_dsfield()) uses skb_protocol is only using skb->protocol. If anything else than ETH_P_IPV6 or ETH_P_IP is found in skb->protocol, pskb_inet vlan tag was provided by the caller (af_packet in the syzbot case), the network header might not point to the correc be smaller than expected. Add skb_vlan_inet_prepare() to perform a complete mac validation. Use this in geneve fo adopt this more broadly. v4 - Jakub reported v3 broke l2_tos_ttl_inherit.sh selftest - Only call __vlan_get_protocol Sabrina comments on v1 and v2 [1] BUG: KMSAN: uninit-value in geneve_xmit_skb drivers/net/geneve.c:910 [in geneve_xmit+0x302d/0x5420 drivers/net/geneve.c:1030 geneve_xmit_skb drivers/net/geneve.c:910 [inline] geneve geneve.c:1030 __netdev_start_xmit include/linux/netdevice.h:4903 [inline] netdev_start_xmit include/linux/netdev dev.c:3531 [inline] dev_hard_start_xmit+0x247/0xa20 net/core/dev.c:3547 __dev_queue_xmit+0x348d/0x52c0 ne linux/netdevice.h:3091 [inline] packet_xmit+0x9c/0x6c0 net/packet/af_packet.c:276 packet_snd net/packet/af_pack +0x8bb0/0x9ef0 net/packet/af_packet.c:3113 sock_sendmsg_nosec net/socket.c:730 [inline] __sock_sendmsg+0x3 +0x685/0x830 net/socket.c:2191 __do_sys_sendto net/socket.c:2203 [inline] __se_sys_sendto net/socket.c:2199 [in net/socket.c:2199 do_syscall_64+0xd5/0x1f0 entry_SYSCALL_64_after_hwframe+0x6d/0x75 Uninit was created [inline] slab_alloc_node mm/slub.c:3845 [inline] kmem_cache_alloc_node+0x613/0xc50 mm/slub.c:3888 kmalloc __alloc_skb+0x35b/0x7a0 net/core/skbuff.c:668 alloc_skb include/linux/skbuff.h:1318 [inline] alloc_skb_with_fra sock_alloc_send_pskb+0xa81/0xbf0 net/core/sock.c:2795 packet_alloc_skb net/packet/af_packet.c:2930 [inline] pa packet_sendmsg+0x722d/0x9ef0 net/packet/af_packet.c:3113 sock_sendmsg_nosec net/socket.c:730 [inline] __soc __sys_sendto+0x685/0x830 net/socket.c:2191 __do_sys_sendto net/socket.c:2203 [inline] __se_sys_sendto net/soc +0x125/0x1d0 net/socket.c:2199 do_syscall_64+0xd5/0x1f0 entry_SYSCALL_64_after_hwframe+0x6d/0x75 CPU tainted 6.9.0-rc1-syzkaller-00005-g928a87efa423 #0 Hardware name: Google Google Compute Engine/Google Co
CVE-2024-35974	In the Linux kernel, the following vulnerability has been resolved: block: fix q->blkg_list corruption during disk re allocated/added for single request queue in case of disk rebind. blkg may still stay in q->blkg_list when calling blk becomes corrupted. Fix the list corruption issue by: - add blkg_init_queue() to initialize q->blkg_list & q->blkcg_n into blk_alloc_queue() The list corruption should be started since commit f1c006f1c685 ("blk-cgroup: synchronize blkcg_deactivate_policy()") which delays removing blkg from q->blkg_list into blkg_free_workfn().
CVE-2024-35980	In the Linux kernel, the following vulnerability has been resolved: arm64: tlb: Fix TLBI RANGE operand KVM/ar flush TLBs when the dirty pages are collected by VMM and the page table entries become write protected during li passed to the TLBI RANGE instruction isn't correctly sorted out due to the commit 117940aa6e5f ("KVM: arm64: It leads to crash on the destination VM after live migration because TLBs aren't flushed completely and some of th have a VM where 8GB memory is assigned, starting from 0x40000000 (1GB). Note that the host has 4KB as the ba kvm_tlb_flush_vmid_range() is executed to flush TLBs. It passes MAX_TLBI_RANGE_PAGES as the argument __flush_s2_tlb_range_op(). SCALE#3 and NUM#31, corresponding to MAX_TLBI_RANGE_PAGES, isn't suppo specific case, -1 has been returned from __TLBI_RANGE_NUM() for SCALE#3/2/1/0 and rejected by the loop in @scale underflows and becomes -9, 0xffff708000040000 is set as the operand. The operand is wrong since it's sort according to invalid @scale and @num. Fix it by extending __TLBI_RANGE_NUM() to support the combination changes, [-1 31] instead of [-1 30] can be returned from the macro, meaning the TLBs for 0x200000 pages in the al with SCALE#3 and NUM#31. The macro TLBI_RANGE_MASK is dropped since no one uses it any more. The co
CVE-2024-35981	In the Linux kernel, the following vulnerability has been resolved: virtio_net: Do not send RSS key if it is not supp options in virtio_net that can break the whole machine, getting the kernel into an infinite loop. Running the followi with virtionet will reproduce this problem: # ethtool -X eth0 hfunc toeplitz This is how the problem happens: 1) eth virtnet_set_rxfh() calls virtnet_commit_rss_command() 3) virtnet_commit_rss_command() populates 4 entries for above does not have a key, then the last scatter-gatter entry will be zeroed, since rss_key_size == 0. sg_buf_size = to qemu, but qemu is not happy with a buffer with zero length, and do the following in virtqueue_map_desc() (QEM "virtio: zero sized buffers are not allowed"); 6) virtio_error() (also QEMU function) set the device as broken vdev- not repond this crazy kernel. 8) The kernel is waiting for the response to come back (function virtnet_send_comma following : while (!virtqueue_get_buf(vi->cvq, &tmp) && !virtqueue_is_broken(vi->cvq)) cpu_relax(); 10) None the kernel loops here forever. Keeping in mind that virtqueue_is_broken() does not look at the qemu `vdev->broke at QEMU side. Fix it by not sending RSS commands if the feature is not available in the device.
CVE-2024-35982	In the Linux kernel, the following vulnerability has been resolved: batman-adv: Avoid infinite loop trying to resize interface becomes too small to transmit the local translation table then it must be resized to fit inside all fragments MTU becomes too low to transmit even the header + the VLAN specific part then the resizing of the local TT will when the usable space is 110 bytes and 11 VLANs are on top of batman-adv. In this case, at least 116 byte would b of batman_adv: batadv0: Forced to purge local tt entries to fit new maximum fragment MTU (110) in the log but th is that the timeout will be halved all the time and will then stagnate at 0 and therefore never be able to reduce the ta possible with a similar result. The number of BATADV_TT_CLIENT_NOPURGE entries in the local TT can for e Such a scenario can therefore happen also with only a single VLAN + 7 non-purgable addresses - requiring at least proactively when: * interface with too low MTU is added * VLAN is added * non-purgeable local mac is added * I fragmentation setting gets disabled (which most likely requires dropping attached interfaces) not all of these scenar only consuming events without the the possibility to prevent these actions (non-purgable MAC address added, MT therefore necessary to also make sure that the code is able to handle also the situations when there were already inc
CVE-2024-35984	In the Linux kernel, the following vulnerability has been resolved: i2c: smbus: fix NULL function pointer derefere designware controller as target only. Target-only modes break the assumption of one transfer function always being pointer in __i2c_transfer. [wsa: dropped the simplification in core-smbus to avoid theoretical regressions]

CVE-2024-35997	In the Linux kernel, the following vulnerability has been resolved: HID: i2c-hid: remove I2C_HID_READ_PENDI I2C_HID_READ_PENDING is used to serialize I2C operations. However, this is not necessary, because I2C core importantly, this flag can cause a lock-up: if the flag is set in i2c_hid_xfer() and an interrupt happens, the interrupt return immediately without doing anything, then the interrupt handler will be invoked again in an infinite loop. Sin over the CPU and the flag-clearing task never gets scheduled, thus we have a lock-up. Delete this unnecessary flag.
CVE-2024-36004	In the Linux kernel, the following vulnerability has been resolved: i40e: Do not use WQ_MEM_RECLAIM flag fo during SRIOV testing, call trace: When both i40e and the i40iw driver are loaded, a warning in check_flush_depen to be because of the i40e driver workqueue is allocated with the WQ_MEM_RECLAIM flag, and the i40iw one is on ice too and it was fixed by removing the flag. Do the same for i40e too. [Feb 9 09:08] ----------[cut here]------ WQ_MEM_RECLAIM i40e:i40e_service_task [i40e] is flushing !WQ_MEM_RECLAIM infiniband:0x0 [+0.000 at kernel/workqueue.c:2966 check_flush_dependency+0x10b/0x120 [+0.000007] Modules linked in: snd_seq_dun snd_seq_device snd soundcore nls_utf8 cifs cifs_arc4 nls_ucs2_utils rdma_cm iw_cm ib_cm cifs_md4 dns_resolve intel_rapl_msr intel_rapl_common irdma intel_uncore_frequency intel_uncore_frequency_common ice ipmi_ssif i x86_pkg_temp_thermal intel_powerclamp gnss coretemp ib_uverbs rapl intel_cstate ib_core iTCO_wdt iTCO_ven intel_uncore ioatdma i2c_i801 joydev pcspkr mei ipmi_devintf lpc_ich intel_pch_thermal i2c_smbus ipmi_msghan libcrc32c ast sd_mod drm_shmem_helper t10_pi drm_kms_helper sg ixgbe drm i40e ahci crct10dif_pclmul libahci ghash_clmulni_intel i2c_algo_bit mdio dca wmi dm_mirror dm_region_hash dm_log dm_mod fuse [+0.000050] C Kdump: loaded Not tainted 6.8.0-rc2-Feb-net_dev-Qiueue-00279-gbd43c5687e05 #1 [+0.000003] Hardware name S2600BPB, BIOS SE5C620.86B.02.01.0013.121520200651 12/15/2020 [+0.000001] Workqueue: i40e i40e_servi 0010:check_flush_dependency+0x10b/0x120 [+0.000003] Code: ff 49 8b 54 24 18 48 8d 8b b0 00 00 00 49 89 e8 97 fa 9f c6 05 8a cc 1f 02 01 e8 35 b3 fd ff <0f> 0b e9 10 ff ff ff 80 3d 78 cc 1f 02 00 75 94 e9 46 ff ff ff 90 [+0.0 EFLAGS: 00010282 [+0.000002] RAX: 0000000000000000 RBX: ffff94d4c483c000 RCX: 0000000000000027 [RSI: 0000000000000001 RDI: ffff94d47f620bc0 [+0.000001] RBP: 0000000000000000 R08: 0000000000000000 R10: ffffbd294976bb98 R11: ffffffffa0be65e8 R12: ffff94c5451ea180 [+0.000001] R13: ffff94c5ab5e8000 R14: f [+0.000001] FS: 0000000000000000(0000) GS:ffff94d47f600000(0000) knlGS:0000000000000000 [+0.000002] 0000000080050033 [+0.000001] CR2: 00007f9e6f1fca70 CR3: 0000000038e20004 CR4: 00000000007706f0 [+0 0000000000000000 DR2: 0000000000000000 [+0.000001] DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR 55555554 [+0.000001] Call Trace: [+0.000001] <TASK> [+0.000002] ? __warn+0x80/0x130 [+0.000003] ? che [+0.000002] ? report_bug+0x195/0x1a0 [+0.000005] ? handle_bug+0x3c/0x70 [+0.000003] ? exc_invalid_op+0 +0x16/0x20 [+0.000006] ? check_flush_dependency+0x10b/0x120 [+0.000002] ? check_flush_dependency+0x1 +0x126/0x3f0 [+0.000015] ib_cache_cleanup_one+0x1c/0xe0 [ib_core] [+0.000056] __ib_unregister_device+0x ib_unregister_device_and_put+0x34/0x50 [ib_core] [+0.000020] i40iw_close+0x4b/0x90 [irdma] [+0.000022] i4 [i40e] [+0.000035] i40e_service_task+0x126/0x190 [i40e] [+0.000024] process_one_work+0x174/0x340 [+0.00
CVE-2024-36006	In the Linux kernel, the following vulnerability has been resolved: mlxsw: spectrum_acl_tcam: Fix incorrect list A the chunks within a region and the function that migrates all the entries within a chunk call list_first_entry() on the are not empty. This is incorrect usage of the API, which leads to the following warning [1]. Fix by returning if the in this case. [1] WARNING: CPU: 0 PID: 6437 at drivers/net/ethernet/mellanox/mlxsw/spectrum_acl_tcam.c:1266 +0x1f1/0> Modules linked in: CPU: 0 PID: 6437 Comm: kworker/0:37 Not tainted 6.9.0-rc3-custom-00883-g94a6 Mellanox Technologies Ltd. MSN3700/VMOD0005, BIOS 5.11 01/06/2019 Workqueue: mlxsw_core mlxsw_sp_ RIP: 0010:mlxsw_sp_acl_tcam_vchunk_migrate_all+0x1f1/0x2c0 [...] Call Trace: <TASK> mlxsw_sp_acl_tcam_ process_one_work+0x151/0x370 worker_thread+0x2cb/0x3e0 kthread+0xd0/0x100 ret_from_fork+0x34/0x50 ret_
CVE-2024-36007	In the Linux kernel, the following vulnerability has been resolved: mlxsw: spectrum_acl_tcam: Fix warning during delayed work migrates filters from one region to another. This is done by iterating over all chunks (all the filters wi chunk iterating over all the filters. When the work runs out of credits it stores the current chunk and entry as marke know where to resume the migration from the next time the work is scheduled. Upon error, the chunk marker is res markers despite being relative to it. This can result in migration being resumed from an entry that does not belong t eventually lead to a chunk being iterated over as if it is an entry. Because of how the two structures happen to be de but to warnings such as [1]. Fix by creating a helper that resets all the markers and call it from all the places the cur measures also call it when starting a completely new rehash. Add a warning to avoid future cases. [1] WARNING: mellanox/mlxsw/core_acl_flex_keys.c:407 mlxsw_afk_encode+0x242/0x2f0 Modules linked in: CPU: 7 PID: 1076 rc3-custom-00880-g29e61d91b77b #29 Hardware name: Mellanox Technologies Ltd. MSN3700/VMOD0005, BIO mlxsw_sp_acl_tcam_vregion_rehash_work RIP: 0010:mlxsw_afk_encode+0x242/0x2f0 [...] Call Trace: <TASK> mlxsw_sp_acl_tcam_entry_create+0x5e/0xa0 mlxsw_sp_acl_tcam_vchunk_migrate_all+0x109/0x290 mlxsw_sp_ process_one_work+0x151/0x370 worker_thread+0x2cb/0x3e0 kthread+0xd0/0x100 ret_from_fork+0x34/0x50 </T

CVE-2024-36008	In the Linux kernel, the following vulnerability has been resolved: ipv4: check for NULL idev in ip_route_use_hin deref in fib_validate_source() in an old tree [1]. It appears the bug exists in latest trees. All calls to __in_dev_get_r result. [1] general protection fault, probably for non-canonical address 0xdffffc0000000000: 0000 [#1] SMP KASA [0x0000000000000000-0x0000000000000007] CPU: 2 PID: 3257 Comm: syz-executor.3 Not tainted 5.10.0-syzka (Q35 + ICH9, 2009), BIOS 1.16.3-debian-1.16.3-2~bpo12+1 04/01/2014 RIP: 0010:fib_validate_source+0xbf/0x1 f2 f2 f2 42 c7 44 20 23 f3 f3 f3 f3 48 89 44 24 78 42 c6 44 20 27 f3 e8 5d 88 48 fc 4c 89 e8 48 c1 e8 03 48 89 44 2 15 98 fc 48 89 5c 24 10 41 bf RSP: 0018:ffffc900015fee40 EFLAGS: 00010246 RAX: 0000000000000000 RBX: RDX: 0000000000000000 RSI: 0000000004001eac RDI: ffff8880160c64c0 RBP: ffffc900015ff060 R08: 0000000 R10: 0000000000000002 R11: ffff88800f4f90c0 R12: dffffc0000000000 R13: 0000000000000000 R14: 0000000 FS: 00007f938acfe6c0(0000) GS:ffff888058c00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 000 00007f938acddd58 CR3: 000000001248e000 CR4: 0000000000352ef0 DR0: 0000000000000000 DR1: 0000000 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 Call Trace: ip_route_use_hint+0x410/0x9b +0x2c4/0x1a30 net/ipv4/ip_input.c:327 ip_list_rcv_finish net/ipv4/ip_input.c:612 [inline] ip_sublist_rcv+0x3ed/0x +0x422/0x470 net/ipv4/ip_input.c:673 __netif_receive_skb_list_ptype net/core/dev.c:5572 [inline] __netif_receive dev.c:5620 __netif_receive_skb_list net/core/dev.c:5672 [inline] netif_receive_skb_list_internal+0x9f9/0xdc0 net/ +0x55/0x3e0 net/core/dev.c:5816 xdp_recv_frames net/bpf/test_run.c:257 [inline] xdp_test_run_batch net/bpf/test_ +0x1818/0x1d00 net/bpf/test_run.c:363 bpf_prog_test_run_xdp+0x81f/0x1170 net/bpf/test_run.c:1376 bpf_prog_t syscall.c:3736 __sys_bpf+0x45c/0x710 kernel/bpf/syscall.c:5115 __do_sys_bpf kernel/bpf/syscall.c:5201 [inline] [inline] __x64_sys_bpf+0x7c/0x90 kernel/bpf/syscall.c:5199
CVE-2024-36011	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: HCI: Fix potential null-ptr-deref Fix hci_le_big_sync_established_evt().
CVE-2024-36014	In the Linux kernel, the following vulnerability has been resolved: drm/arm/malidp: fix a possible null pointer dere new memory is allocated with kzalloc, but no check is performed. In order to prevent null pointer dereferencing, er __drm_atomic_helper_connector_reset.
CVE-2024-36016	In the Linux kernel, the following vulnerability has been resolved: tty: n_gsm: fix possible out-of-bounds in gsm0_ A configures the n_gsm in basic option mode - side B sends the header of a basic option mode frame with data leng mode - side B sends 2 data bytes which exceeds gsm->len Reason: gsm->len is not used in advanced option mode. side B keeps sending until gsm0_receive() writes past gsm->buf Reason: Neither gsm->state nor gsm->len have be changing gsm->count to gsm->len comparison from equal to less than. Also add upper limit checks against the con gsm1_receive() to harden against memory corruption of gsm->len and gsm->mru. All other checks remain as we st configuration and actual payload size.
CVE-2024-36017	In the Linux kernel, the following vulnerability has been resolved: rtnetlink: Correct nested IFLA_VF_VLAN_LIS a nested IFLA_VF_VLAN_LIST is assumed to be a struct ifla_vf_vlan_info so the size of such attribute needs to b which is 14 bytes. The current size validation in do_setvfinfo is against NLA_HDRLEN (4 bytes) which is less tha validation is not enough and a too small attribute might be cast to a struct ifla_vf_vlan_info, this might result in an saved (casted) entry in ivvl.
CVE-2024-36020	In the Linux kernel, the following vulnerability has been resolved: i40e: fix vf may be used uninitialized in this fun introduced by commit 52424f974bc5, which causes servers hang in very hard to reproduce conditions with resets ra the root cause. In this function before the fix bumping v didn't mean bumping vf pointer. But the code used this var to different/not intended vf. Remove redundant "v" variable and iterate via single VF pointer across whole function
CVE-2024-36039	PyMySQL through 1.1.0 allows SQL injection if used with untrusted JSON input because keys are not escaped by
CVE-2024-36270	In the Linux kernel, the following vulnerability has been resolved: netfilter: tproxy: bail out if IP has been disabled general protection fault, probably for non-canonical address 0xdffffc0000000003: 0000 [#1] PREEMPT SMP KAS range [0x0000000000000018-0x000000000000001f] [..] RIP: 0010:nf_tproxy_laddr4+0xb7/0x340 net/ipv4/netfilt nft_tproxy_eval_v4 net/netfilter/nft_tproxy.c:56 [inline] nft_tproxy_eval+0xa9a/0x1a00 net/netfilter/nft_tproxy.c:1 check for this.
CVE-2024-36286	In the Linux kernel, the following vulnerability has been resolved: netfilter: nfnetlink_queue: acquire rcu_read_loc reported that nf_reinject() could be called without rcu_read_lock() : WARNING: suspicious RCU usage 6.9.0-rc7- tainted net/netfilter/nfnetlink_queue.c:263 suspicious rcu_dereference_check() usage! other info that might help us debug_locks = 1 2 locks held by syz-executor.4/13427: #0: ffffffff8e334f60 (rcu_callback){....}-{0:0}, at: rcu_lock [inline] #0: ffffffff8e334f60 (rcu_callback){....}-{0:0}, at: rcu_do_batch kernel/rcu/tree.c:2190 [inline] #0: ffffffff8 rcu_core+0xa86/0x1830 kernel/rcu/tree.c:2471 #1: ffff88801ca92958 (&inst->lock){+.-.}-{2:2}, at: spin_lock_bh ffff88801ca92958 (&inst->lock){+.-.}-{2:2}, at: nfqnl_flush net/netfilter/nfnetlink_queue.c:405 [inline] #1: ffff888 at: instance_destroy_rcu+0x30/0x220 net/netfilter/nfnetlink_queue.c:172 stack backtrace: CPU: 0 PID: 13427 Com syzkaller-02060-g5c1672705a1a #0 Hardware name: Google Google Compute Engine/Google Compute Engine, B __dump_stack lib/dump_stack.c:88 [inline] dump_stack_lvl+0x241/0x360 lib/dump_stack.c:114 lockdep_rcu_susp lockdep.c:6712 nf_reinject net/netfilter/nfnetlink_queue.c:323 [inline] nfqnl_reinject+0x6ec/0x1120 net/netfilter/nf nfnetlink_queue.c:410 [inline] instance_destroy_rcu+0x1ae/0x220 net/netfilter/nfnetlink_queue.c:172 rcu_do_batc +0xafd/0x1830 kernel/rcu/tree.c:2471 handle_softirqs+0x2d6/0x990 kernel/softirq.c:554 __do_softirq kernel/softir softirq.c:428 [inline] __irq_exit_rcu+0xf4/0x1c0 kernel/softirq.c:637 irq_exit_rcu+0x9/0x30 kernel/softirq.c:649 i kernel/apic/apic.c:1043 [inline] sysvec_apic_timer_interrupt+0xa6/0xc0 arch/x86/kernel/apic/apic.c:1043 </IRQ>

CVE-2024-36489	In the Linux kernel, the following vulnerability has been resolved: tls: fix missing memory barrier in tls_init In tls_ and store-store reordering may cause NULL dereference in tls_{setsockopt,getsockopt}. CPU0 CPU1 ----- ----- // I kzalloc() ctx->sk_proto = READ_ONCE(sk->sk_prot) -(1) // In update_sk_prot() WRITE_ONCE(sk->sk_prot, tls_ READ_ONCE(sk->sk_prot)->setsockopt() // In tls_{setsockopt,getsockopt}() ctx->sk_proto->setsockopt() -(3) In reordered, (3) can observe the NULL value of ctx->sk_proto, causing NULL dereference. To fix it, we rely on rcu_ barrier semantic. By moving rcu_assign_pointer() after ctx->sk_proto is initialized, we can ensure that ctx->sk_pro
CVE-2024-3651	A vulnerability was identified in the kjd/idna library, specifically within the `idna.encode()` function, affecting vers handling of crafted input strings, which can lead to quadratic complexity and consequently, a denial of service cond crafted input that causes the `idna.encode()` function to process the input with considerable computational load, sig quadratic manner relative to the input size.
CVE-2024-36600	Buffer Overflow Vulnerability in libcdio v2.1.0 allows an attacker to execute arbitrary code via a crafted ISO 9660
CVE-2024-36883	In the Linux kernel, the following vulnerability has been resolved: net: fix out-of-bounds access in ops_init net_all is called without any locking. It reads max_gen_ptrs, which is changed under pernet_ops_rwsem. It is read twice, f s.len, which is later used to limit the bounds of the array access. It is possible that the array is allocated and another increments max_gen_ptrs, which is then used to set s.len with a larger than allocated length for the variable array. I net_alloc_generic. If max_gen_ptrs is later incremented, it will be caught in net_assign_generic.
CVE-2024-36886	In the Linux kernel, the following vulnerability has been resolved: tipc: fix UAF in error path Sam Page (sam4k) w Initiative reported a UAF in the tipc_buf_append() error path: BUG: KASAN: slab-use-after-free in kfree_skb_list skbuff.c:1183 Read of size 8 at addr ffff88804d2a7c80 by task poc/8034 CPU: 1 PID: 8034 Comm: poc Not tainted PC (i440FX + PIIX, 1996), BIOS 1.16.0-debian-1.16.0-5 04/01/2014 Call Trace: <IRQ> __dump_stack linux/lib/d linux/lib/dump_stack.c:106 print_address_description linux/mm/kasan/report.c:377 print_report+0xc4/0x620 linux +0xda/0x110 linux/mm/kasan/report.c:601 kfree_skb_list_reason+0x47e/0x4c0 linux/net/core/skbuff.c:1183 skb_r core/skbuff.c:1026 skb_release_all linux/net/core/skbuff.c:1094 __kfree_skb linux/net/core/skbuff.c:1108 kfree_sk skbuff.c:1144 kfree_skb linux/./include/linux/skbuff.h:1244 tipc_buf_append+0x425/0xb50 linux/net/tipc/msg.c:18 net/tipc/link.c:1324 tipc_link_rcv+0x76e/0x2d70 linux/net/tipc/link.c:1824 tipc_rcv+0x45f/0x10f0 linux/net/tipc/n linux/net/tipc/udp_media.c:390 udp_queue_rcv_one_skb+0xad2/0x1850 linux/net/ipv4/udp.c:2108 udp_queue_rcv udp.c:2186 udp_unicast_rcv_skb+0x165/0x3b0 linux/net/ipv4/udp.c:2346 __udp4_lib_rcv+0x2594/0x3400 linux/n +0x30c/0x4e0 linux/net/ipv4/ip_input.c:205 ip_local_deliver_finish+0x2e4/0x520 linux/net/ipv4/ip_input.c:233 N NF_HOOK linux/./include/linux/netfilter.h:308 ip_local_deliver+0x18e/0x1f0 linux/net/ipv4/ip_input.c:254 dst_in linux/net/ipv4/ip_input.c:449 NF_HOOK linux/./include/linux/netfilter.h:314 NF_HOOK linux/./include/linux/netf net/ipv4/ip_input.c:569 __netif_receive_skb_one_core+0x199/0x1e0 linux/net/core/dev.c:5534 __netif_receive_sk process_backlog+0x101/0x6b0 linux/net/core/dev.c:5976 __napi_poll.constprop.0+0xba/0x550 linux/net/core/dev. dev.c:6645 net_rx_action+0x95a/0xe90 linux/net/core/dev.c:6781 __do_softirq+0x21f/0x8e7 linux/kernel/softirq.c do_softirq+0xb2/0xf0 linux/kernel/softirq.c:441 </IRQ> <TASK> __local_bh_enable_ip+0x100/0x120 linux/kern include/linux/bottom_half.h:33 rcu_read_unlock_bh linux/./include/linux/rcupdate.h:851 __dev_queue_xmit+0x87 dev_queue_xmit linux/./include/linux/netdevice.h:3169 neigh_hh_output linux/./include/net/neighbour.h:526 neigh ip_finish_output2+0x169f/0x2550 linux/net/ipv4/ip_output.c:235 __ip_finish_output linux/net/ipv4/ip_output.c:31 net/ipv4/ip_output.c:295 ip_finish_output+0x31/0x310 linux/net/ipv4/ip_output.c:323 NF_HOOK_COND linux/./i +0x13b/0x2a0 linux/net/ipv4/ip_output.c:433 dst_output linux/./include/net/dst.h:451 ip_local_out linux/net/ipv4/i linux/net/ipv4/ip_output.c:1492 udp_send_skb+0x73f/0x1530 linux/net/ipv4/udp.c:963 udp_sendmsg+0x1a36/0x2 +0x105/0x140 linux/net/ipv4/af_inet.c:850 sock_sendmsg_nosec linux/net/socket.c:730 __sock_sendmsg linux/net linux/net/socket.c:2191 __do_sys_sendto linux/net/socket.c:2203 __se_sys_sendto linux/net/socket.c:2199 __x64_ socket.c:2199 do_syscall_x64 linux/arch/x86/entry/common.c:52 do_syscall_ ---truncated---
CVE-2024-36893	In the Linux kernel, the following vulnerability has been resolved: usb: typec: tcpm: Check for port partner validity does not guarantee partner registration to always succeed. In the event of failure, port->partner is set to the error va validity is not checked, this results in the following crash: Unable to handle kernel NULL pointer dereference at vi +0x1bc8/0x1c08 lr : run_state_machine+0x1b90/0x1c08 .. Call trace: run_state_machine+0x1bc8/0x1c08 tcpm_sta kthread_worker_fn+0x118/0x328 kthread+0x1d0/0x23c ret_from_fork+0x10/0x20 To prevent the crash, check for all the call sites.
CVE-2024-36894	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: f_fs: Fix race between aio_cancel() applications can utilize the aio_cancel() callback to dequeue pending USB requests submitted to the UDC. There is issues an AIO cancel call, while the UDC is handling a soft disconnect. For a DWC3 based implementation, the cal DWC3 Gadget FFS Application dwc3_gadget_soft_disconnect() ... --> dwc3_stop_active_transfers() --> dwc3_gad ffs_epfile_async_io_complete() ffs_aio_cancel() --> usb_ep_free_request() --> usb_ep_dequeue() There is currentl the AIO completion handler and AIO cancel, so the issue occurs if the completion routine is running in parallel to a the FFS application. As the completion call frees the USB request (io_data->req) the FFS application is also referen call. This can lead to accessing a stale/hanging pointer. commit b566d38857fc ("usb: gadget: f_fs: use io_data->sta usb_ep_free_request() into ffs_epfile_async_io_complete(). However, in order to properly implement locking to m to ffs_epfile_async_io_complete(), as usb_ep_dequeue() (if successfully dequeuing a USB request) will call the fun context. Hence, leading into a deadlock. Fix this issue by moving the usb_ep_free_request() back to ffs_user_copy_ sets io_data->req to NULL after freeing it within the ffs->eps_lock. This resolves the race condition above, as the f attempting to dequeue a request that has already been freed, or the ffs_user_copy_work() not freeing the USB requ This fix depends on commit b566d38857fc ("usb: gadget: f_fs: use io_data->status consistently")

CVE-2024-36898	In the Linux kernel, the following vulnerability has been resolved: gpiolib: cdev: fix uninitialised kfifo If a line is r debouncing in software, and the line is subsequently reconfigured to enable edge detection then the allocation of th This results in events being written to and read from an uninitialised kfifo. Read events are returned to userspace. I software debounce is already active.
CVE-2024-36899	In the Linux kernel, the following vulnerability has been resolved: gpiolib: cdev: Fix use after free in lineinfo_char as follows: when the GPIO chip device file is being closed by invoking gpio_chrdev_release(), watched_lines is fre of lineinfo_changed_nb notifier chain failed due to waiting write rwsem. Additionally, one of the GPIO chip's lines notifier chain's read rwsem. Consequently, a race condition leads to the use-after-free of watched_lines. Here is the gpio_chrdev_release() --> bitmap_free(cdev->watched_lines) <-- freed --> blocking_notifier_chain_unregister() -- rwsem --> __down_write_common() --> rwsem_down_write_slowpath() --> schedule_preempt_disabled() --> sche gpio_free() --> gpiod_free() --> gpiod_free_commit() --> gpiod_line_state_notify() --> blocking_notifier_call_cha rwsem --> notifier_call_chain() --> lineinfo_changed_notify() --> test_bit(xxxx, cdev->watched_lines) <-- use afte issue is that a GPIO line event is being generated for userspace where it shouldn't. However, since the chrdev is bei read that event anyway. To fix the issue, call the bitmap_free() function after the unregistration of lineinfo_changed
CVE-2024-36901	In the Linux kernel, the following vulnerability has been resolved: ipv6: prevent NULL dereference in ip6_output() chance that ip6_dst_idev() returns NULL in ip6_output(). Most places in IPv6 stack deal with a NULL idev just fin general protection fault, probably for non-canonical address 0xdffffc00000000bc: 0000 [#1] PREEMPT SMP KAS range [0x00000000000005e0-0x00000000000005e7] CPU: 0 PID: 9775 Comm: syz-executor.4 Not tainted 6.9.0-r #0 Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google 03/27/2024 RIP: 001 ip6_output.c:237 Code: 3c 1e 00 49 89 df 74 08 4c 89 ef e8 19 58 db f7 48 8b 44 24 20 49 89 45 00 49 89 c5 48 8d <42> 0f b6 04 38 84 c0 4c 8b 74 24 28 0f 85 61 01 00 00 8b 1b 31 ff RSP: 0018:ffffc9000927f0d8 EFLAGS: 0001 00000000000005e0 RCX: 0000000000040000 RDX: ffffc900131f9000 RSI: 0000000000004f47 RDI: 000000000 ffffffff8a1f0b9a R09: 1ffffffff1f51fad R10: dfffc0000000000 R11: fffffbfff1f51fae R12: ffff8880293ec8c0 R13: f R15: dfffc0000000000 FS: 00007f135c6856c0(0000) GS:ffff8880b9400000(0000) knlGS:0000000000000000 CS 0000000080050033 CR2: 0000000020000080 CR3: 0000000064096000 CR4: 00000000003506f0 DR0: 0000000 DR2: 0000000000000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 Call Trace netfilter.h:314 [inline] ip6_xmit+0xefe/0x17f0 net/ipv6/ip6_output.c:358 sctp_v6_xmit+0x9f2/0x13f0 net/sctp/ipv net/sctp/output.c:653 sctp_packet_singleton+0x22c/0x320 net/sctp/outqueue.c:783 sctp_outq_flush_ctrl net/sctp/ou +0x6d5/0x3e20 net/sctp/outqueue.c:1212 sctp_side_effects net/sctp/sm_sideeffect.c:1198 [inline] sctp_do_sm+0x5 sctp_primitive_ASSOCIATE+0x95/0xc0 net/sctp/primitive.c:73 __sctp_connect+0x9cd/0xe30 net/sctp/socket.c:12 [inline] sctp_inet_connect+0x149/0x1f0 net/sctp/socket.c:4834 __sys_connect_file net/socket.c:2048 [inline] __sys __do_sys_connect net/socket.c:2075 [inline] __se_sys_connect net/socket.c:2072 [inline] __x64_sys_connect+0x7 arch/x86/entry/common.c:52 [inline] do_syscall_64+0xf5/0x240 arch/x86/entry/common.c:83 entry_SYSCALL_6
CVE-2024-36902	In the Linux kernel, the following vulnerability has been resolved: ipv6: fib6_rules: avoid possible NULL derefere to trigger the following crash [1], caused by unsafe ip6_dst_idev() use. Indeed ip6_dst_idev() can return NULL, an general protection fault, probably for non-canonical address 0xdffffc0000000000: 0000 [#1] PREEMPT SMP KAS [0x0000000000000000-0x0000000000000007] CPU: 0 PID: 31648 Comm: syz-executor.0 Not tainted 6.9.0-rc4-n name: Google Google Compute Engine/Google Compute Engine, BIOS Google 03/27/2024 RIP: 0010:__fib6_rule [inline] RIP: 0010:fib6_rule_action+0x241/0x7b0 net/ipv6/fib6_rules.c:267 Code: 02 00 00 49 8d 9f d8 00 00 00 4 08 48 89 df e8 f9 32 bf f7 48 8b 1b 48 89 d8 48 c1 e8 03 <42> 80 3c 20 00 74 08 48 89 df e8 e0 32 bf f7 4c 8b 03 4 EFLAGS: 00010246 RAX: 0000000000000000 RBX: 0000000000000000 RCX: 1a772f98c8186700 RDX: 00000 ffffffff8c1f9760 RBP: ffff8880673fb980 R08: ffffffff8fac15ef R09: 1ffffffff1f582bd R10: dfffc0000000000 R11: 0000000000000080 R14: ffff888076509000 R15: ffff88807a029a00 FS: 00007f55e82ca6c0(0000) GS:ffff8880b94 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 0000001b31d23000 CR3: 0000000022b66000 CR4: 000 DR1: 0000000000000000 DR2: 0000000000000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000 fib_rules_lookup+0x62c/0xdb0 net/core/fib_rules.c:317 fib6_rule_lookup+0x1fd/0x790 net/ipv6/fib6_rules.c:108 i route.c:2637 [inline] ip6_route_output_flags+0x38e/0x610 net/ipv6/route.c:2649 ip6_route_output include/net/ip6_ +0x189/0x11a0 net/ipv6/ip6_output.c:1120 ip6_dst_lookup_flow+0xb9/0x180 net/ipv6/ip6_output.c:1250 sctp_v6 sctp_transport_route+0x12c/0x2e0 net/sctp/transport.c:455 sctp_assoc_add_peer+0x614/0x15c0 net/sctp/associola net/sctp/socket.c:1099 __sctp_connect+0x66d/0xe30 net/sctp/socket.c:1197 sctp_connect net/sctp/socket.c:4819 [i net/sctp/socket.c:4834 __sys_connect_file net/socket.c:2048 [inline] __sys_connect+0x2df/0x310 net/socket.c:206 [inline] __se_sys_connect net/socket.c:2072 [inline] __x64_sys_connect+0x7a/0x90 net/socket.c:2072 do_syscall_ do_syscall_64+0xf5/0x240 arch/x86/entry/common.c:83 entry_SYSCALL_64_after_hwframe+0x77/0x7f

CVE-2024-36904	In the Linux kernel, the following vulnerability has been resolved: tcp: Use refcount_inc_not_zero() in tcp_twsk_u use-after-free splat in tcp_twsk_unique() with nice analysis. Since commit ec94c2696f0b ("tcp/dccp: avoid one ato inet_twsk_hashdance() sets TIME-WAIT socket's sk_refcnt after putting it into ehash and releasing the bucket lock where other threads could try to reuse the port during connect() and call sock_hold() in tcp_twsk_unique() for the T If that happens, the refcnt taken by tcp_twsk_unique() is overwritten and sock_put() will cause underflow, triggerin else. To avoid the use-after-free, we need to use refcount_inc_not_zero() in tcp_twsk_unique() and give up on reus refcount_t: addition on 0; use-after-free. WARNING: CPU: 0 PID: 1039313 at lib/refcount.c:25 refcount_warn_sat 1039313 Comm: trigger Not tainted 6.8.6-200.fc39.x86_64 #1 Hardware name: VMware, Inc. VMware20,1/440BX VMW201.00V.21805430.B64.2305221830 05/22/2023 RIP: 0010:refcount_warn_saturate+0xe5/0x110 Code: 42 8 01 00 0f 85 5e ff ff ff 48 c7 c7 f8 8e b7 82 c6 05 96 13 ea 01 01 e8 7b 42 8e ff <0f> 0b c3 cc cc cc cc 48 c7 c7 50 8 0018:ffffc90006b43b60 EFLAGS: 00010282 RAX: 0000000000000000 RBX: ffff888009bb3ef0 RCX: 000000000 RSI: 0000000000000001 RDI: ffff88807be218c0 RBP: 0000000000069d70 R08: 0000000000000000 R09: ffffc90 R11: 0000000000000003 R12: ffff8880029ede84 R13: 0000000000004e20 R14: ffffffff84356dc0 R15: ffff888009 GS:ffff88807be00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 C 000000004628c005 CR4: 0000000000f70ef0 PKRU: 55555554 Call Trace: <TASK> ? refcount_warn_saturate+0x refcount_warn_saturate+0xe5/0x110 ? report_bug+0x171/0x1a0 ? refcount_warn_saturate+0xe5/0x110 ? handle_b +0x17/0x70 ? asm_exc_invalid_op+0x1a/0x20 ? refcount_warn_saturate+0xe5/0x110 tcp_twsk_unique+0x186/0x __inet_hash_connect+0x74/0x7d0 ? __pfx___inet_check_established+0x10/0x10 tcp_v4_connect+0x278/0x530 __ inet_stream_connect+0x3a/0x60 __sys_connect+0xa8/0xd0 __x64_sys_connect+0x18/0x20 do_syscall_64+0x83/0 +0x78/0x80 RIP: 0033:0x7f62c11a885d Code: ff c3 66 2e 0f 1f 84 00 00 00 00 00 90 f3 0f 1e fa 48 89 f8 48 89 f7 8b 4c 24 08 0f 05 <48> 3d 01 f0 ff ff 73 01 c3 48 8b 0d a3 45 0c 00 f7 d8 64 89 01 48 RSP: 002b:00007f62c1091e 000000000000002a RAX: ffffffffffffffda RBX: 0000000020ccb004 RCX: 00007f62c11a885d RDX: 00000000000 0000000000000003 RBP: 00007f62c1091e90 R08: 0000000000000000 R09: 0000000000000000 R10: 000000000 00007f62c10926c0 R13: ffffffffffffff88 R14: 0000000000000000 R15: 00007ffe237885b0 </TASK>
CVE-2024-36905	In the Linux kernel, the following vulnerability has been resolved: tcp: defer shutdown(SEND_SHUTDOWN) for ' TCP_SYN_RECV state is really special, it is only used by cross-syn connections, mostly used by fuzzers. In the fo to trigger a divide by zero in tcp_rcv_space_adjust() A socket makes the following state transitions, without ever ca tcp_init_buffer_space() is also not called. TCP_CLOSE connect() TCP_SYN_SENT TCP_SYN_RECV shutdown(TCP_FIN_WAIT1 To fix this issue, change tcp_shutdown() to not perform a TCP_SYN_RECV -> TCP_FIN_WA anyway. When tcp_rcv_state_process() later changes socket state from TCP_SYN_RECV to TCP_ESTABLISH, th finally enter TCP_FIN_WAIT1 state, and send a FIN packet from a sane socket state. This means tcp_send_fin() ca and must use GFP_ATOMIC allocations. [1] divide error: 0000 [#1] PREEMPT SMP KASAN NOPTI CPU: 1 PID Not tainted 6.9.0-rc6-syzkaller-00022-g98369dccd2f8 #0 Hardware name: Google Google Compute Engine/Googl 03/27/2024 RIP: 0010:tcp_rcv_space_adjust+0x2df/0x890 net/ipv4/tcp_input.c:767 Code: e3 04 4c 01 eb 48 8b 44 d5 0f 85 a5 03 00 00 41 8b 8e c8 09 00 00 89 e8 29 c8 48 0f af c3 31 d2 <48> f7 f1 48 8d 1c 43 49 8d 96 76 08 00 0018:ffffc900031ef3f0 EFLAGS: 00010246 RAX: 0c677a10441f8f42 RBX: 000000004fb95e7e RCX: 000000000 RSI: 0000000000000000 RDI: 0000000000000000 RBP: 0000000027d4b11f R08: ffffffff89e535a4 R09: 1ffffffff2 R11: ffffffff8135e920 R12: ffff88802a9f8d30 R13: dffffc0000000000 R14: ffff88802a9f8d00 R15: 1ffff1100553f2 GS:ffff8880b9500000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 C 000000002b9f2000 CR4: 0000000000350ef0 Call Trace: <TASK> tcp_recvmsg_locked+0x106d/0x25a0 net/ipv4/ net/ipv4/tcp.c:2578 inet6_recvmsg+0x16a/0x730 net/ipv6/af_inet6.c:680 sock_recvmsg_nosec net/socket.c:1046 [i net/socket.c:1068 ____sys_recvmsg+0x1db/0x470 net/socket.c:2803 ___sys_recvmsg net/socket.c:2845 [inline] do socket.c:2939 __sys_recvmmsg net/socket.c:3018 [inline] __do_sys_recvmmsg net/socket.c:3041 [inline] __se_sys __x64_sys_recvmmsg+0x199/0x250 net/socket.c:3034 do_syscall_x64 arch/x86/entry/common.c:52 [inline] do_sy common.c:83 entry_SYSCALL_64_after_hwframe+0x77/0x7f RIP: 0033:0x7faeb6363db9 Code: 28 00 00 00 75 0 90 48 89 f8 48 89 f7 48 89 d6 48 89 ca 4d 89 c2 4d 89 c8 4c 8b 4c 24 08 0f 05 <48> 3d 01 f0 ff ff 73 01 c3 48 c7 c 002b:00007ffcc1997168 EFLAGS: 00000246 ORIG_RAX: 000000000000012b RAX: ffffffffffffffda RBX: 00000 RDX: 0000000000000001 RSI: 0000000020000bc0 RDI: 0000000000000005 RBP: 0000000000000000 R08: 000 R10: 0000000000000122 R11: 0000000000000246 R12: 0000000000000000 R13: 0000000000000000 R14: 0000

CVE-2024-36906	In the Linux kernel, the following vulnerability has been resolved: ARM: 9381/1: kasan: clear stale stack poison W ========== [33.453513] BUG: K refresh_cpu_vm_stats.constprop.0+0xcc/0x2ec [33.454660] Write of size 164 at addr c1d03d30 by task swapper/0 Comm: swapper/0 Tainted: G O 6.1.25-mainline #1 [33.456880] Hardware name: Generic DT based system [33.4 +0x18/0x1c [33.458326] show_stack from dump_stack_lvl+0x40/0x4c [33.459072] dump_stack_lvl from print_r from kasan_report+0x9c/0x148 [33.460616] kasan_report from kasan_check_range+0x94/0x1a0 [33.461424] kas [33.462157] memset from refresh_cpu_vm_stats.constprop.0+0xcc/0x2ec [33.463064] refresh_cpu_vm_stats.con +0x180/0x53c [33.464181] tick_nohz_idle_stop_tick from do_idle+0x264/0x354 [33.465029] do_idle from cpu_ cpu_startup_entry from rest_init+0xf0/0xf4 [33.466528] rest_init from arch_post_acpi_subsys_init+0x0/0x18 [33 belongs to stack of task swapper/0/0 [33.468493] and is located at offset 112 in frame: [33.469172] refresh_cpu_v [33.470165] This frame has 2 objects: [33.470696] [32, 76) 'global_zone_diff' [33.470729] [112, 276) 'global_no buggy address belongs to the physical page: [33.472862] page:3cd72da8 refcount:1 mapcount:0 mapping:0000000 0x1000(reserved|zone=0) [33.474565] raw: 00001000 ed741470 ed741470 00000000 00000000 ffffffff [33.476050] page dumped because: kasan: bad access detected [33.476816] [33.477061] Memory state around th 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 [33.478630] c1d03c80: 00 00 00 00 00 00 00 00 f1 f1 f1 f1 00 0 f2 f2 f2 f2 00 00 00 00 00 00 f1 f1 f1 f1 [33.480415] ^ [33.481195] c1d03d80: 00 00 00 00 00 00 00 00 00 00 04 f3 f3 f3 00 00 00 00 00 00 00 00 00 00 00 00 [33.482978] ========== find the root cause of this OOB is that arm does not clear stale stack poison in the case of cpuidle. This patch refer issue. From cited commit [1] that explain the problem Functions which the compiler has instrumented for KASAN and remove this poison prior to returning. In the case of cpuidle, CPUs exit the kernel a number of levels deep in C critical path will leave portions of the stack shadow poisoned. If CPUs lose context and return to the kernel via a co __cpu_suspend_enter are forgotten, and we never remove the poison they placed in the stack shadow area by funct the kernel. Thus, (depending on stackframe layout) subsequent calls to instrumented functions may hit this stale po to the console. To avoid this, clear any stale poison from the idle thread for a CPU prior to bringing a CPU online. CONFIG_KASAN_STACK [1] commit 0d97e6d8024c ("arm64: kasan: clear stale stack poison") [2] commit d56a only with CONFIG_KASAN_STACK")
CVE-2024-36933	In the Linux kernel, the following vulnerability has been resolved: nsh: Restore skb->{protocol,data,mac_header} f triggered various splats (see [0] and links) by a crafted GSO packet of VIRTIO_NET_HDR_GSO_UDP layering th + ETH_P_NSH + ETH_P_IPV6 + IPPROTO_UDP NSH can encapsulate IPv4, IPv6, Ethernet, NSH, and MPLS. A NSH GSO handler, nsh_gso_segment(), calls skb_mac_gso_segment() to invoke inner protocol GSO handlers. nsh for the original skb before calling skb_mac_gso_segment() 1. reset skb->network_header 2. save the original skb-> variable 3. pull the NSH header 4. resets skb->mac_header 5. set up skb->mac_len and skb->protocol for the inner the segmented skb 6. set ntohs(ETH_P_NSH) to skb->protocol 7. push the NSH header 8. restore skb->mac_heade skb->network_header 10. restore skb->mac_len There are two problems in 6-7 and 8-9. (a) After 6 & 7, skb->data header (ETH_P_8021AD in this case) is stripped when skb is sent out of netdev. Also, if NSH is encapsulated by N skb_pull() in the first nsh_gso_segment() will make skb->data point to the middle of the outer NSH or Ethernet hea not pulled by the second nsh_gso_segment(). (b) While restoring skb->{mac_header,network_header} in 8 & 9, ns the data in the linear buffer is shifted. However, udp6_ufo_fragment() could shift the data and change skb->mac_h syzbot. If this happens, even the restored skb->mac_header points to the middle of the outer header. It seems nsh_g outer headers so far. At the end of nsh_gso_segment(), the outer header must be restored for the segmented skb, ins let's calculate the outer header position relatively from the inner header and set skb->{data,mac_header,protocol} p value in ipvlan_process_outbound drivers/net/ipvlan/ipvlan_core.c:524 [inline] BUG: KMSAN: uninit-value in ipv ipvlan_core.c:602 [inline] BUG: KMSAN: uninit-value in ipvlan_queue_xmit+0xf44/0x16b0 drivers/net/ipvlan/ipv drivers/net/ipvlan/ipvlan_core.c:524 [inline] ipvlan_xmit_mode_l3 drivers/net/ipvlan/ipvlan_core.c:602 [inline] ipv net/ipvlan/ipvlan_core.c:668 ipvlan_start_xmit+0x5c/0x1a0 drivers/net/ipvlan/ipvlan_main.c:222 __netdev_start_x netdev_start_xmit include/linux/netdevice.h:5003 [inline] xmit_one net/core/dev.c:3547 [inline] dev_hard_start_xr __dev_queue_xmit+0x33ed/0x51c0 net/core/dev.c:4351 dev_queue_xmit include/linux/netdevice.h:3171 [inline] p af_packet.c:276 packet_snd net/packet/af_packet.c:3081 [inline] packet_sendmsg+0x8aef/0x9f10 net/packet/af_pa socket.c:730 [inline] __sock_sendmsg net/socket.c:745 [inline] __sys_sendto+0x735/0xa10 net/socket.c:2191 __do __se_sys_sendto net/socket.c:2199 [inline] __x64_sys_sendto+0x125/0x1c0 net/socket.c:2199 do_syscall_x64 arch do_syscall_64+0xcf/0x1e0 arch/x86/entry/common.c:83 entry_SYSCALL_64_after_hwframe+0x63/0x6b Uninit v slub.c:3819 [inline] slab_alloc_node mm/slub.c:3860 [inline] __do_kmalloc_node mm/slub.c:3980 [inline] __kmal slub.c:4001 kmalloc_reserve+0x249/0x4a0 net/core/skbuff.c:582 __ ---truncated---
CVE-2024-36934	In the Linux kernel, the following vulnerability has been resolved: bna: ensure the copied buf is NUL terminated C buffer and copy nbytes from userspace to that buffer. Later, we use sscanf on this buffer but we don't ensure that th can lead to OOB read when using sscanf. Fix this issue by using memdup_user_nul instead of memdup_user.

CVE-2024-36939	In the Linux kernel, the following vulnerability has been resolved: nfs: Handle error of rpc_proc_register() in nfs_r [0] triggered while destroying immature netns. rpc_proc_register() was called in init_nfs_fs(), but its error has beer 1da177e4c3f4 ("Linux-2.6.12-rc2"). Recently, commit d47151b79e32 ("nfs: expose /proc/net/sunrpc/nfs in net nam netns and made the problem more visible. Even when rpc_proc_register() fails, nfs_net_init() could succeed, and th destroying the netns. Then, remove_proc_entry() will be called for non-existing proc directory and trigger the warn rpc_proc_register() properly in nfs_net_init(). [0]: name 'nfs' WARNING: CPU: 1 PID: 1710 at fs/proc/generic.c:7 fs/proc/generic.c:711 Modules linked in: CPU: 1 PID: 1710 Comm: syz-executor.2 Not tainted 6.8.0-12822-gcd51c Standard PC (i440FX + PIIX, 1996), BIOS rel-1.16.0-0-gd239552ce722-prebuilt.qemu.org 04/01/2014 RIP: 0010:r proc/generic.c:711 Code: 41 5d 41 5e c3 e8 85 09 b5 ff 48 c7 c7 88 58 64 86 e8 09 0e 71 02 e8 74 09 b5 ff 4c 89 e0 ff <0f> 0b eb b1 e8 5c 09 b5 ff 48 c7 c7 88 58 64 86 e8 e0 0d 71 02 eb RSP: 0018:ffffc9000c6d7ce0 EFLAGS: 000 RBX: ffff8880422b8b00 RCX: ffffffff8110503c RDX: ffff888030652f00 RSI: ffffffff81105045 RDI: 0000000000 R08: 0000000000000001 R09: 0000000000000000 R10: 0000000000000001 R11: ffffffff81bb62cb R12: ffffffff84 ffffffff84807ffc R15: ffffffff85741ff8 FS: 00007f30cfba8640(0000) GS:ffff88807dd00000(0000) knlGS:0000000000 0000000080050033 CR2: 00007ff51afe8000 CR3: 000000005a60a005 CR4: 0000000000770ef0 DR0: 0000000000 0000000000000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 PKRU: 5555555 +0x64/0x70 net/sunrpc/stats.c:310 nfs_net_exit+0x1c/0x30 fs/nfs/inode.c:2438 ops_exit_list+0x62/0xb0 net/core/n net/core/net_namespace.c:372 copy_net_ns+0x244/0x590 net/core/net_namespace.c:505 create_new_namespaces+ unshare_nsproxy_namespaces+0xae/0x160 kernel/nsproxy.c:228 ksys_unshare+0x342/0x760 kernel/fork.c:3322 _ [inline] __se_sys_unshare kernel/fork.c:3391 [inline] __x64_sys_unshare+0x1f/0x30 kernel/fork.c:3391 do_syscal do_syscall_64+0x4f/0x110 arch/x86/entry/common.c:83 entry_SYSCALL_64_after_hwframe+0x46/0x4e RIP: 00 84 00 00 00 00 90 f3 0f 1e fa 48 89 f8 48 89 f7 48 89 d6 48 89 ca 4d 89 c2 4d 89 c8 4c 8b 4c 24 08 0f 05 <48> f7 d8 64 89 01 48 RSP: 002b:00007f30cfba7cc8 EFLAGS: 00000246 ORIG_RAX: 0000000000000110 RAX: ffff 00007f30d0febe5d RDX: 0000000000000000 RSI: 0000000000000000 RDI: 000000006c020600 RBP: 0000000 0000000000000000 R10: 0000000000000000 R11: 0000000000000246 R12: 0000000000000002 R13: 0000000 0000000000000000 </TASK>
CVE-2024-36940	In the Linux kernel, the following vulnerability has been resolved: pinctrl: core: delete incorrect free in pinctrl_ena devm_pinctrl_register_and_init(). It's a devm_ managed pointer that is freed by devm_pinctrl_dev_release(), so fre double free. The devm_pinctrl_dev_release() function frees the pindescs and destroys the mutex as well.
CVE-2024-36941	In the Linux kernel, the following vulnerability has been resolved: wifi: nl80211: don't free NULL coalescing rule NULL pointer here.
CVE-2024-36946	In the Linux kernel, the following vulnerability has been resolved: phonet: fix rtm_phonet_notify() skb allocation f skb: - struct rtmsg - RTA_DST (u8) - RTA_OIF (u32) Therefore, rtm_phonet_notify() should use NLMSG_ALIGN nla_total_size(4)
CVE-2024-36950	In the Linux kernel, the following vulnerability has been resolved: firewire: ohci: mask bus reset interrupts betwee interrupt handler, if a bus reset interrupt has occurred, mask bus reset interrupts until bus_reset_work has serviced a always leave bus reset interrupts masked. We infer the bus reset from the self-ID interrupt that happens shortly ther reset interrupts was introduced in 2008 in a007bb857e0b26f5d8b73c2ff90782d9c0972620: If OHCI_PARAM_DE parameter bitmask, we will unmask bus reset interrupts so we can log them. irq_handler logs the bus reset interrupt flag in irq_handler, because we won't service the event until later. irq_handler exits with the event flag still set. If th the first bus reset will usually freeze the system due to irq_handler being called again each time it exits. This freeze with "modprobe firewire_ohci debug=-1" (to enable all debugging output). Apparently there are also some cases w enough to clear the event, and operation will continue normally. This freeze was first reported a few months after a never fixed. The debug level could safely be set to -1 through sysfs after the module was loaded, but this would be since they were only unmasked during initialization. irq_handler will now leave the event flag set but mask bus res again and there will be no freeze. If OHCI_PARAM_DEBUG_BUSRESETS is enabled, bus_reset_work will unma future interrupts will be caught as desired. As a side effect to this change, OHCI_PARAM_DEBUG_BUSRESETS to during initial module loading. However, when enabled through sysfs, logging of bus reset interrupts will be effe after bus_reset_work has executed.
CVE-2024-36950	In the Linux kernel, the following vulnerability has been resolved: firewire: ohci: mask bus reset interrupts betwee interrupt handler, if a bus reset interrupt has occurred, mask bus reset interrupts until bus_reset_work has serviced a always leave bus reset interrupts masked. We infer the bus reset from the self-ID interrupt that happens shortly ther reset interrupts was introduced in 2008 in a007bb857e0b26f5d8b73c2ff90782d9c0972620: If OHCI_PARAM_DE parameter bitmask, we will unmask bus reset interrupts so we can log them. irq_handler logs the bus reset interrupt flag in irq_handler, because we won't service the event until later. irq_handler exits with the event flag still set. If th the first bus reset will usually freeze the system due to irq_handler being called again each time it exits. This freeze with "modprobe firewire_ohci debug=-1" (to enable all debugging output). Apparently there are also some cases w enough to clear the event, and operation will continue normally. This freeze was first reported a few months after a never fixed. The debug level could safely be set to -1 through sysfs after the module was loaded, but this would be since they were only unmasked during initialization. irq_handler will now leave the event flag set but mask bus res again and there will be no freeze. If OHCI_PARAM_DEBUG_BUSRESETS is enabled, bus_reset_work will unma future interrupts will be caught as desired. As a side effect to this change, OHCI_PARAM_DEBUG_BUSRESETS to during initial module loading. However, when enabled through sysfs, logging of bus reset interrupts will be effe after bus_reset_work has executed.
CVE-2024-36954	In the Linux kernel, the following vulnerability has been resolved: tipc: fix a possible memleak in tipc_buf_append it fails, so move '*buf = NULL' after __skb_linearize(), so that the skb can be freed on the err path.

CVE-2024-36954	In the Linux kernel, the following vulnerability has been resolved: tipc: fix a possible memleak in tipc_buf_append it fails, so move '*buf = NULL' after __skb_linearize(), so that the skb can be freed on the err path.
CVE-2024-36959	In the Linux kernel, the following vulnerability has been resolved: pinctrl: devicetree: fix refcount leak in pinctrl_d buffer, we need to drop the reference count we just took. Because the pinctrl_dt_free_maps() includes the droping
CVE-2024-36960	In the Linux kernel, the following vulnerability has been resolved: drm/vmwgfx: Fix invalid reads in fence signalec drm_event to the size of the structure that's actually used. The length of the drm_event was set to the parent structu which is supposed to be read. drm_read uses the length parameter to copy the event to the user space thus resuling
CVE-2024-36964	In the Linux kernel, the following vulnerability has been resolved: fs/9p: only translate RWX permissions for plain bits is allowed through, which causes it to be able to set (among others) the suid bit. This was presumably not the ir explicitly and conditionally on .u.
CVE-2024-36971	In the Linux kernel, the following vulnerability has been resolved: net: fix __dst_negative_advice() race __dst_neg rules when sk->dst_cache must be cleared, leading to possible UAF. RCU rules are that we must first clear sk->sk_ Note that sk_dst_reset(sk) is implementing this protocol correctly, while __dst_negative_advice() uses the wrong o special logic against RTF_CACHE, this means each of the three ->negative_advice() existing methods must perfor check against NULL dst is centralized in __dst_negative_advice(), there is no need to duplicate it in various callbac tracking this issue. This old bug became visible after the blamed commit, using UDP sockets.
CVE-2024-36974	In the Linux kernel, the following vulnerability has been resolved: net/sched: taprio: always validate TCA_TAPRIO TCA_TAPRIO_ATTR_PRIOMAP attribute has been provided, taprio_parse_mqprio_opt() must validate it, or user the second time taprio_change() is called. First call (with valid attributes) sets dev->num_tc to a non zero value. Se returns early from taprio_parse_mqprio_opt() and bad things can happen.
CVE-2024-36975	In the Linux kernel, the following vulnerability has been resolved: KEYS: trusted: Do not use WARN when encode WARN is not the correct solution. 1. asn1_encode_sequence() is not an internal function (located in lib/asn1_encod the stack trace useless. 3. Results a crash if panic_on_warn is set. It is also noteworthy that the use of WARN is unc there is a carefully considered rationale to use it. Replace WARN with pr_err, and print the return value instead, wh
CVE-2024-36978	In the Linux kernel, the following vulnerability has been resolved: net: sched: sch_multiq: fix possible OOB write i qopt->bands to execute subsequent code logic after kmalloc. So the old q->bands should not be used in kmalloc. O
CVE-2024-37078	In the Linux kernel, the following vulnerability has been resolved: nilfs2: fix potential kernel bug due to lack of wr block device on which nilfs2 is mounted can cause a kernel bug in the folio/page writeback start routine or writeba log below): kernel BUG at mm/page-writeback.c:3070! Oops: invalid opcode: 0000 [#1] PREEMPT SMP KASAN +0xbaa/0x10e0 Code: 25 ff 0f 00 00 0f 84 18 01 00 00 e8 40 ca c6 ff e9 17 f6 ff ff e8 36 ca c6 ff 4c 89 f7 48 c7 c6 1f ca c6 ff 4c 89 f7 48 c7 c6 a0 c6 12 84 e8 d0 b3 0f 00 ... Call Trace: <TASK> nilfs_segctor_do_construct+0x465 +0x181/0x6b0 [nilfs2] nilfs_segctor_thread+0x548/0x11c0 [nilfs2] kthread+0x2f0/0x390 ret_from_fork+0x4b/0x8 This is because when the log writer starts a writeback for segment summary blocks or a super root block that use th wait for the ongoing folio/page writeback, resulting in an inconsistent writeback state. Fix this issue by waiting for pages on the backing device into writeback state.
CVE-2024-37356	In the Linux kernel, the following vulnerability has been resolved: tcp: Fix shift-out-of-bounds in dctcp_update_alp a module parameter dctcp_shift_g as follows: alpha -= min_not_zero(alpha, alpha >> dctcp_shift_g); ... delivered_ syzkaller started fuzzing module parameters and triggered shift-out-of-bounds [0] by setting 100 to dctcp_shift_g: r module/tcp_dctcp/parameters/dctcp_shift_g\000", 47); res = syscall(__NR_openat, /*fd=*/0xffffffffffffff9cul, /*fil *mode=*/0ul); memcpy((void*)0x20000000, "100\000", 4); syscall(__NR_write, /*fd=*/r[0], /*val=*/0x2000000 value of dctcp_shift_g by param_set_uint_minmax(). With this patch: # echo 10 > /sys/module/tcp_dctcp/paramete tcp_dctcp/parameters/dctcp_shift_g 10 # echo 11 > /sys/module/tcp_dctcp/parameters/dctcp_shift_g -bash: echo: w shift-out-of-bounds in net/ipv4/tcp_dctcp.c:143:12 shift exponent 100 is too large for 32-bit type 'u32' (aka 'unsigne syz-executor345 Not tainted 6.9.0-05151-g1b294a1f3561 #2 Hardware name: QEMU Standard PC (i440FX + PIIX 04/01/2014 Call Trace: <TASK> __dump_stack lib/dump_stack.c:88 [inline] dump_stack_lvl+0x201/0x300 lib/du ubsan.c:231 [inline] __ubsan_handle_shift_out_of_bounds+0x346/0x3a0 lib/ubsan.c:468 dctcp_update_alpha+0x5 tcp_in_ack_event net/ipv4/tcp_input.c:3802 [inline] tcp_ack+0x17b1/0x3bc0 net/ipv4/tcp_input.c:3948 tcp_rcv_st tcp_input.c:6711 tcp_v4_do_rcv+0x764/0xc40 net/ipv4/tcp_ipv4.c:1937 sk_backlog_rcv include/net/sock.h:1106 [sock.c:2983 release_sock+0x61/0x1f0 net/core/sock.c:3549 mptcp_subflow_shutdown+0x3d0/0x620 net/mptcp/pr +0x225/0x410 net/mptcp/protocol.c:2976 __mptcp_close+0x238/0xad0 net/mptcp/protocol.c:3072 mptcp_close+0 inet_release+0x190/0x1f0 net/ipv4/af_inet.c:437 __sock_release net/socket.c:659 [inline] sock_close+0xc0/0x240 file_table.c:422 task_work_run+0x23b/0x300 kernel/task_work.c:180 exit_task_work include/linux/task_work.h:3 exit.c:878 do_group_exit+0x201/0x2b0 kernel/exit.c:1027 __do_sys_exit_group kernel/exit.c:1038 [inline] __se_s __x64_sys_exit_group+0x3f/0x40 kernel/exit.c:1036 do_syscall_x64 arch/x86/entry/common.c:52 [inline] do_sysc common.c:83 entry_SYSCALL_64_after_hwframe+0x67/0x6f RIP: 0033:0x7f6c2b5005b6 Code: Unable to acces 002b:00007ffe883eb948 EFLAGS: 00000246 ORIG_RAX: 00000000000000e7 RAX: ffffffffffffffda RBX: 00007 RDX: 0000000000000001 RSI: 000000000000003c RDI: 0000000000000001 RBP: 0000000000000001 R08: 000 0000000000000006 R11: 0000000000000246 R12: 00007f6c2b5862f0 R13: 0000000000000001 R14: 000000000
CVE-2024-3772	Regular expression denial of service in Pydanic < 2.4.0, < 1.10.13 allows remote attackers to cause denial of servic

CVE-2024-38381	In the Linux kernel, the following vulnerability has been resolved: nfc: nci: Fix uninit-value in nci_rx_work syzbot issue [1] nci_rx_work() parses received packet from ndev->rx_q. It should be validated header size, payload size ar packet. If an invalid packet is detected, it should be silently discarded.
CVE-2024-38549	In the Linux kernel, the following vulnerability has been resolved: drm/mediatek: Add 0 size check to mtk_drm_ge if we attempt to allocate a GEM object of 0 bytes. Currently, no such check exists and the kernel will panic if a use GBM buffer. Tested by attempting to allocate a 0x0 GBM buffer on an MT8188 and verifying that we now return I
CVE-2024-38552	In the Linux kernel, the following vulnerability has been resolved: drm/amd/display: Fix potential index out of bou Fixes index out of bounds issue in the color transformation function. The issue could occur when the index 'i' excee (TRANSFER_FUNC_POINTS). The fix adds a check to ensure 'i' is within bounds before accessing the transfer fu an error message is logged and the function returns false to indicate an error. Reported by smatch: drivers/gpu/drm/ dcn10_cm_common.c:405 cm_helper_translate_curve_to_hw_format() error: buffer overflow 'output_tf->tf_pts.re amdgpu/../display/dc/dcn10/dcn10_cm_common.c:406 cm_helper_translate_curve_to_hw_format() error: buffer ov <= s32max drivers/gpu/drm/amd/amdgpu/../display/dc/dcn10/dcn10_cm_common.c:407 cm_helper_translate_curv 'output_tf->tf_pts.blue' 1025 <= s32max
CVE-2024-38558	In the Linux kernel, the following vulnerability has been resolved: net: openvswitch: fix overwriting ct original tup OVS_PACKET_CMD_EXECUTE has 3 main attributes: - OVS_PACKET_ATTR_KEY - Packet metadata in a ne OVS_PACKET_ATTR_PACKET - Binary packet content. - OVS_PACKET_ATTR_ACTIONS - Actions to exec OVS_PACKET_ATTR_KEY is parsed first to populate sw_flow_key structure with the metadata like conntrack st the packet itself gets parsed to populate the rest of the keys from the packet headers. Whenever the packet parsing c first zeroes out fields in the key corresponding to Neighbor Discovery information even if it is not an ND packet. It is a union that shares the space between 'nd' and 'ct_orig' that holds the original tuple conntrack metadata parsed fro ND packets should not normally have conntrack state, so it's fine to share the space, but normal ICMPv6 Echo pack can have the state attached and it should not be overwritten. The issue results in all but the last 4 bytes of the destin original conntrack tuple leading to incorrect packet matching and potentially executing wrong actions in case this p goes back to userspace. ND fields should not be accessed in non-ND packets, so not clearing them should be fine. I packets to avoid the issue. Initializing the whole thing before parsing is needed because ND packet may not contain the OVS_PACKET_CMD_EXECUTE path and doesn't affect packets entering OVS datapath from network interfa populated from skb after the packet is already parsed.
CVE-2024-38559	In the Linux kernel, the following vulnerability has been resolved: scsi: qedf: Ensure the copied buf is NUL termin kernel buffer and copy count from userspace to that buffer. Later, we use kstrtouint on this buffer but we don't ensu buffer, this can lead to OOB read when using kstrtouint. Fix this issue by using memdup_user_nul instead of memc
CVE-2024-38560	In the Linux kernel, the following vulnerability has been resolved: scsi: bfa: Ensure the copied buf is NUL terminat kernel buffer and copy nbytes from userspace to that buffer. Later, we use sscanf on this buffer but we don't ensure this can lead to OOB read when using sscanf. Fix this issue by using memdup_user_nul instead of memdup_user.
CVE-2024-38564	In the Linux kernel, the following vulnerability has been resolved: bpf: Add BPF_PROG_TYPE_CGROUP_SKB a BPF_LINK_CREATE bpf_prog_attach uses attach_type_to_prog_type to enforce proper attach type for BPF_PRO bpf_prog_get and relies on bpf_prog_attach_check_attach_type to properly verify prog_type <> attach_type associ for the link_create case. Otherwise, it's currently possible to attach cgroup_skb prog types to other cgroup hooks.
CVE-2024-38565	In the Linux kernel, the following vulnerability has been resolved: wifi: ar5523: enable proper endpoint verificatio an endpoint in use not having an expected type to it. Fix the issue by checking for the existence of all proper endpo this patch has not been tested on real hardware. [1] Syzkaller report: ------------[cut here]------------ usb 1-1: BOGU CPU: 0 PID: 3643 at drivers/usb/core/urb.c:504 usb_submit_urb+0xed6/0x1880 drivers/usb/core/urb.c:504 ... Call drivers/net/wireless/ath/ar5523/ar5523.c:275 ar5523_cmd_read drivers/net/wireless/ath/ar5523/ar5523.c:302 [inlin ath/ar5523/ar5523.c:1376 [inline] ar5523_probe+0x14b0/0x1d10 drivers/net/wireless/ath/ar5523/ar5523.c:1655 us usb/core/driver.c:396 call_driver_probe drivers/base/dd.c:560 [inline] really_probe+0x249/0xb90 drivers/base/dd.c drivers/base/dd.c:778 driver_probe_device+0x4c/0x1a0 drivers/base/dd.c:808 __device_attach_driver+0x1d4/0x2e +0x163/0x1e0 drivers/base/bus.c:427 __device_attach+0x1e4/0x530 drivers/base/dd.c:1008 bus_probe_device+0x +0xbd9/0x1e90 drivers/base/core.c:3517 usb_set_configuration+0x101d/0x1900 drivers/usb/core/message.c:2170 drivers/usb/core/generic.c:238 usb_probe_device+0xd8/0x2c0 drivers/usb/core/driver.c:293 call_driver_probe driv +0x249/0xb90 drivers/base/dd.c:639 __driver_probe_device+0x1df/0x4d0 drivers/base/dd.c:778 driver_probe_dev __device_attach_driver+0x1d4/0x2e0 drivers/base/dd.c:936 bus_for_each_drv+0x163/0x1e0 drivers/base/bus.c:42 dd.c:1008 bus_probe_device+0x1e8/0x2a0 drivers/base/bus.c:487 device_add+0xbd9/0x1e90 drivers/base/core.c:3 drivers/usb/core/hub.c:2573 hub_port_connect drivers/usb/core/hub.c:5353 [inline] hub_port_connect_change driv drivers/usb/core/hub.c:5653 [inline] hub_event+0x26cb/0x45d0 drivers/usb/core/hub.c:5735 process_one_work+0x worker_thread+0x669/0x1090 kernel/workqueue.c:2436 kthread+0x2e8/0x3a0 kernel/kthread.c:376 ret_from_fork TASK>
CVE-2024-38566	In the Linux kernel, the following vulnerability has been resolved: bpf: Fix verifier assumptions about socket->sk T socket' is valid and non-NULL when 'socket' pointer itself is trusted and non-NULL. That may not be the case whe socket_accept hook. Fix this verifier assumption and adjust tests.

CVE-2024-38567	In the Linux kernel, the following vulnerability has been resolved: wifi: carl9170: add a proper sanity check for end warning which is caused by presence of a wrong endpoint type at the URB sumbitting stage. While there was a che it can switch types between bulk and interrupt, other endpoints are trusted implicitly. Similar warning is triggered i Fix the issue by doing a comprehensive check of all endpoints taking into account difference between high- and ful report: ... WARNING: CPU: 0 PID: 4721 at drivers/usb/core/urb.c:504 usb_submit_urb+0xed6/0x1880 drivers/usb carl9170_usb_send_rx_irq_urb+0x273/0x340 drivers/net/wireless/ath/carl9170/usb.c:504 carl9170_usb_init_devic [inline] carl9170_usb_firmware_finish drivers/net/wireless/ath/carl9170/usb.c:999 [inline] carl9170_usb_firmware ath/carl9170/usb.c:1028 request_firmware_work_func+0x130/0x240 drivers/base/firmware_loader/main.c:1107 pr workqueue.c:2289 worker_thread+0x669/0x1090 kernel/workqueue.c:2436 kthread+0x2e8/0x3a0 kernel/kthread.c entry/entry_64.S:308 </TASK> [2] Related syzkaller crashes:
CVE-2024-38570	In the Linux kernel, the following vulnerability has been resolved: gfs2: Fix potential glock use-after-free on unmo and there ares still locks in that lockspace, DLM will unlock those locks automatically. Commit fb6791d100d1b sta filesystem unmount: gfs2 would simply free glocks it didn't want to unlock and then release the lockspace. This dic lock contention notifications into account, which remain active until until a lock is unlocked or its lockspace is rele accessing deallocated objects, put the glocks that should not be unlocked on the sd_dead_glocks list, release the loc an additional measure, ignore unexpected ast and bast callbacks if the receiving glock is dead.
CVE-2024-38573	In the Linux kernel, the following vulnerability has been resolved: cppc_cpufreq: Fix possible null pointer derefere hisi_cppc_cpufreq_get_rate() can be called from different places with various parameters. So cpufreq_cpu_get() ca circumstances. Fix this bug by adding null return check. Found by Linux Verification Center (linuxtesting.org) wit
CVE-2024-38578	In the Linux kernel, the following vulnerability has been resolved: ecryptfs: Fix buffer size for tag 66 packet The 'T description is missing the cipher code and checksum fields that are packed into the message packet. As a result, the is 3 bytes too small and write_tag_66_packet() will write up to 3 bytes past the end of the buffer. Fix this by increa so the whole packet will always fit in the buffer. This fixes the below kasan slab-out-of-bounds bug: BUG: KASAN ecryptfs_generate_key_packet_set+0x7d6/0xde0 Write of size 1 at addr ffff88800afbb2a5 by task touch/181 CPU: gnu #1 4c9534092be820851bb687b82d1f92a426598dc6 Hardware name: QEMU Standard PC (i440FX + PIIX, 19 Trace: <TASK> dump_stack_lvl+0x4c/0x70 print_report+0xc5/0x610 ? ecryptfs_generate_key_packet_set+0x7d6 +0x44/0x210 ? ecryptfs_generate_key_packet_set+0x7d6/0xde0 kasan_report+0xc2/0x110 ? ecryptfs_generate_ke __asan_store1+0x62/0x80 ecryptfs_generate_key_packet_set+0x7d6/0xde0 ? __pfx_ecryptfs_generate_key_packe +0x2e2/0x540 ? __pfx_ovl_open+0x10/0x10 [overlay 30837f11141636a8e1793533a02e6e2e885dad1d] ? dentry_o +0x30a/0x550 ? __pfx_ecryptfs_write_metadata+0x10/0x10 ? ecryptfs_get_lower_file+0x6b/0x190 ecryptfs_initia +0x1c2/0x2f0 path_openat+0x17cf/0x1ba0 ? __pfx_path_openat+0x10/0x10 do_filp_open+0x15e/0x290 ? __pfx_ __kasan_check_write+0x18/0x30 ? _raw_spin_lock+0x86/0xf0 ? __pfx__raw_spin_lock+0x10/0x10 ? __kasan_ch +0xf4/0x330 do_sys_openat2+0x122/0x160 ? __pfx_do_sys_openat2+0x10/0x10 __x64_sys_openat+0xef/0x170 do_syscall_64+0x60/0xd0 entry_SYSCALL_64_after_hwframe+0x6e/0xd8 RIP: 0033:0x7f00a703fd67 Code: 25 0 04 25 18 00 00 00 85 c0 75 5b 44 89 e2 48 89 ee bf 9c ff ff ff b8 01 01 00 00 0f 05 <48> 3d 00 f0 ff ff 0f 87 85 00 002b:00007ffc088e30b0 EFLAGS: 00000246 ORIG_RAX: 0000000000000101 RAX: ffffffffffffffda RBX: 00007 RDX: 0000000000000941 RSI: 00007ffc088e48d7 RDI: 00000000ffffff9c RBP: 00007ffc088e48d7 R08: 0000000 R10: 00000000000001b6 R11: 0000000000000246 R12: 0000000000000941 R13: 0000000000000000 R14: 0000 TASK> Allocated by task 181: kasan_save_stack+0x2f/0x60 kasan_set_track+0x29/0x40 kasan_save_alloc_info+0 __kmalloc+0x66/0x160 ecryptfs_generate_key_packet_set+0x6d2/0xde0 ecryptfs_write_metadata+0x30a/0x550 e ecryptfs_create+0x1c2/0x2f0 path_openat+0x17cf/0x1ba0 do_filp_open+0x15e/0x290 do_sys_openat2+0x122/0x do_syscall_64+0x60/0xd0 entry_SYSCALL_64_after_hwframe+0x6e/0xd8
CVE-2024-38582	In the Linux kernel, the following vulnerability has been resolved: nilfs2: fix potential hang in nilfs_detach_log_wr in nilfs_detach_log_writer() called during nilfs2 unmount. Analysis revealed that this is because nilfs_segctor_sync thread, can be called after nilfs_segctor_destroy() terminates that thread, as shown in the call trace below: nilfs_det nilfs_segctor_kill_thread --> Shut down log writer thread flush_work nilfs_iput_work_func nilfs_dispose_list iput nilfs_construct_segment (if inode needs sync) nilfs_segctor_sync --> Attempt to synchronize with log writer thread changing nilfs_segctor_sync() so that the log writer thread returns normally without synchronizing after it terminat waiting to complete once after the thread terminates. The skipped inode metadata flushout will then be processed to nilfs_segctor_destroy().
CVE-2024-38583	In the Linux kernel, the following vulnerability has been resolved: nilfs2: fix use-after-free of timer for log writer t related issues". This bug fix series covers three nilfs2 log writer-related issues, including a timer use-after-free issu and a potential freeze issue in event synchronization found during their analysis. Details are described in each com issue has been reported regarding the timer sc_timer on the nilfs_sc_info structure. The problem is that even though thread, sc_timer is not shut down until the nilfs_sc_info structure is about to be freed, and is used regardless of the use of sc_timer only while the log writer thread is alive.
CVE-2024-38587	In the Linux kernel, the following vulnerability has been resolved: speakup: Fix sizeof() vs ARRAY_SIZE() bug T This code should be using ARRAY_SIZE() (which is 256) instead of sizeof() (which is 512), otherwise it can the s

CVE-2024-38589	In the Linux kernel, the following vulnerability has been resolved: netrom: fix possible dead-lock in nr_rt_ioctl() sy deadlock in nr_rt_ioctl [1] Make sure we always acquire nr_node_list_lock before nr_node_lock(nr_node) [1] WA detected 6.9.0-rc7-syzkaller-02147-g654de42f3fc6 #0 Not tainted ---------- syz ffff8880186e2070 (&nr_node->node_lock){+...}-{2:2}, at: spin_lock_bh include/linux/spinlock.h:356 [inline] ffff {2:2}, at: nr_node_lock include/net/netrom.h:152 [inline] ffff8880186e2070 (&nr_node->node_lock){+...}-{2:2}, [inline] ffff8880186e2070 (&nr_node->node_lock){+...}-{2:2}, at: nr_rt_ioctl+0x1bb/0x1090 net/netrom/nr_route ffffffff8f7053b8 (nr_node_list_lock){+...}-{2:2}, at: spin_lock_bh include/linux/spinlock.h:356 [inline] ffffffff8f7 at: nr_dec_obs net/netrom/nr_route.c:462 [inline] ffffffff8f7053b8 (nr_node_list_lock){+...}-{2:2}, at: nr_rt_ioctl+ which lock already depends on the new lock. the existing dependency chain (in reverse order) is: -> #1 (nr_node_li +0x1ed/0x550 kernel/locking/lockdep.c:5754 __raw_spin_lock_bh include/linux/spinlock_api_smp.h:126 [inline] locking/spinlock.c:178 spin_lock_bh include/linux/spinlock.h:356 [inline] nr_remove_node net/netrom/nr_route.c: net/netrom/nr_route.c:355 nr_rt_ioctl+0xa95/0x1090 net/netrom/nr_route.c:683 sock_do_ioctl+0x158/0x460 net/s socket.c:1341 vfs_ioctl fs/ioctl.c:51 [inline] __do_sys_ioctl fs/ioctl.c:904 [inline] __se_sys_ioctl+0xfc/0x170 fs/ioc common.c:52 [inline] do_syscall_64+0xf5/0x240 arch/x86/entry/common.c:83 entry_SYSCALL_64_after_hwfran {+...}-{2:2}: check_prev_add kernel/locking/lockdep.c:3134 [inline] check_prevs_add kernel/locking/lockdep.c:32 kernel/locking/lockdep.c:3869 __lock_acquire+0x1346/0x1fd0 kernel/locking/lockdep.c:5137 lock_acquire+0x1ed __raw_spin_lock_bh include/linux/spinlock_api_smp.h:126 [inline] _raw_spin_lock_bh+0x35/0x50 kernel/locking linux/spinlock.h:356 [inline] nr_node_lock include/net/netrom.h:152 [inline] nr_dec_obs net/netrom/nr_route.c:464 netrom/nr_route.c:697 sock_do_ioctl+0x158/0x460 net/socket.c:1222 sock_ioctl+0x629/0x8e0 net/socket.c:1341 v fs/ioctl.c:904 [inline] __se_sys_ioctl+0xfc/0x170 fs/ioctl.c:890 do_syscall_x64 arch/x86/entry/common.c:52 [inlin entry/common.c:83 entry_SYSCALL_64_after_hwframe+0x77/0x7f other info that might help us debug this: Poss ---- ---- lock(nr_node_list_lock); lock(&nr_node->node_lock); lock(nr_node_list_lock); lock(&nr_node->node_loc by syz-executor350/5129: #0: ffffffff8f7053b8 (nr_node_list_lock){+...}-{2:2}, at: spin_lock_bh include/linux/spin (nr_node_list_lock){+...}-{2:2}, at: nr_dec_obs net/netrom/nr_route.c:462 [inline] #0: ffffffff8f70 ---truncated---
CVE-2024-38596	In the Linux kernel, the following vulnerability has been resolved: af_unix: Fix data races in unix_release_sock/uni has been identified in af_unix. In one data path, the write function unix_release_sock() atomically writes to sk->sk_ on the reader side, unix_stream_sendmsg() does not read it atomically. Consequently, this issue is causing the follo KCSAN: data-race in unix_release_sock / unix_stream_sendmsg write (marked) to 0xffff88867256ddbb of 1 bytes (net/unix/af_unix.c:640) unix_release (net/unix/af_unix.c:1050) sock_close (net/socket.c:659 net/socket.c:1421) __ (fs/file_table.c:508) __se_sys_close (fs/open.c:1559 fs/open.c:1541) __x64_sys_close (fs/open.c:1541) x64_sys_ca do_syscall_64 (arch/x86/entry/common.c:?) entry_SYSCALL_64_after_hwframe (arch/x86/entry/entry_64.S:130) task 989 on cpu 14: unix_stream_sendmsg (net/unix/af_unix.c:2273) __sock_sendmsg (net/socket.c:730 net/socket __sys_sendmmsg (net/socket.c:2638 net/socket.c:2724) __x64_sys_sendmmsg (net/socket.c:2753 net/socket.c:275 entry/syscall_64.c:33) do_syscall_64 (arch/x86/entry/common.c:?) entry_SYSCALL_64_after_hwframe (arch/x86 -> 0x03 The line numbers are related to commit dd5a440a31fa ("Linux 6.9-rc7"). Commit e1d09c2c2f57 ("af_unix addressed a comparable issue in the past regarding sk->sk_shutdown. However, it overlooked resolving this particu unix_stream_sendmsg() function, since the other reads seem to be protected by unix_state_lock() as discussed in
CVE-2024-38598	In the Linux kernel, the following vulnerability has been resolved: md: fix resync softlockup when bitmap size is le raid10, lvextend + lvchange --syncaction will trigger following softlockup: kernel:watchdog: BUG: soft lockup - C CPU: 7 PID: 3588 Comm: mdX_resync Kdump: loaded Not tainted 6.9.0-rc4-next-20240419 #1 RIP: 0010:_raw_s <TASK> md_bitmap_start_sync+0x6b/0xf0 raid10_sync_request+0x25c/0x1b40 [raid10] md_do_sync+0x64b/0x1 +0xcf/0x100 ret_from_fork+0x30/0x50 ret_from_fork_asm+0x1a/0x30 And the detailed process is as follows: md_ < max_sectors) sectors = raid10_sync_request(mddev, j, &skipped) if (!md_bitmap_start_sync(..., &sync_blocks)) to 0 return sync_blocks + sectors_skippe; // sectors = 0; j += sectors; // j never change Root cause is that commit 3 of-bounds in md_bitmap_get_counter") return early from md_bitmap_get_counter(), without setting returned block blocks from md_bitmap_get_counter"(), as it used to be. Noted that this patch just fix the softlockup problem in ke array size still need to be fixed.
CVE-2024-38599	In the Linux kernel, the following vulnerability has been resolved: jffs2: prevent xattr node from overflowing the e the requested xattr node size is no larger than the eraseblock minus the cleanmarker. Unlike the usual inode nodes, spread across multiple eraseblocks, which means that a xattr node must not occupy more than one eraseblock. If th xattr node can spill onto the next eraseblock, overwriting the nodes and causing errors such as: jffs2: argh. node ad nextblock 0x0000a000, expected at 0000b00c jffs2: error: (823) do_verify_xattr_datum: node CRC failed at 0x01e jffs2: notice: (823) jffs2_get_inode_nodes: Node header CRC failed at 0x01e00c. {848f,2fc4,0fef511f,59a3d171} j 0x00001044 would run over the end of the erase block jffs2: Perhaps the file system was created with the wrong er Magic bitmask 0x1985 not found at 0x00000010: 0x1044 instead This breaks the filesystem and can lead to KASA slab-out-of-bounds in jffs2_sum_add_kvec+0x125e/0x15d0 Read of size 4 at addr ffff88802c31e914 by task repro/ Not tainted 6.9.0-rc3+ #1 Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS Arch Linux 1.16.3-1 dump_stack_lvl+0xc6/0x120 print_report+0xc4/0x620 ? __virt_addr_valid+0x308/0x5b0 kasan_report+0xc1/0xf0 jffs2_sum_add_kvec+0x125e/0x15d0 jffs2_sum_add_kvec+0x125e/0x15d0 jffs2_flash_direct_writev+0xa8/0xd0 j __x64_sys_setxattr+0xc4/0x160 ? do_syscall_64+0x69/0x140 ? entry_SYSCALL_64_after_hwframe+0x76/0x7e (linuxtesting.org) with Syzkaller.
CVE-2024-38606	In the Linux kernel, the following vulnerability has been resolved: crypto: qat - validate slices count returned by FW enables the telemetry (TL) feature on a QAT device by sending the ICP_QAT_FW_TL_START message to the fir TL data to a DMA buffer in memory and returns an array containing the number of accelerators of each type (slices array is stored in the adf_tl_hw_data data structure called slice_cnt. The array slice_cnt is then used in the function statistics about the supported accelerators. An incorrect value of the elements in slice_cnt might lead to an out of b isn't an implementation of FW that returns a wrong value, but for robustness validate the slice count array returned

CVE-2024-38612	In the Linux kernel, the following vulnerability has been resolved: ipv6: sr: fix invalid unregister error path The err CONFIG_IPV6_SEG6_LWTUNNEL is not defined. In that case if seg6_hmac_init() fails, the genl_unregister_fan commit 46738b1317e1 ("ipv6: sr: add option to control lwtunnel support"), and commit 5559cea2d5aa ("ipv6: sr: f replaced unregister_pernet_subsys() with genl_unregister_family() in this error path.
CVE-2024-38613	In the Linux kernel, the following vulnerability has been resolved: m68k: Fix spinlock race in kernel thread creatio retain the correct lock owner across the switch from 'prev' to 'next' tasks. This does rely on interrupts remaining dis This condition is guaranteed for normal process creation and context switching between already running processes, have interrupts disabled in their saved copies of the status register. The situation is different for newly created kern PS_S in copy_thread(), which does leave the IPL at 0. Upon restoring the 'next' thread's status register in switch_to enabled prematurely. resume() then returns via ret_from_kernel_thread() and schedule_tail() where run queue lock finish_lock_switch()). A timer interrupt calling scheduler_tick() before the lock is released in finish_task_switch() current task as lock owner. This causes a spinlock recursion warning as reported by Guenter Roeck. As far as I can commit 533e6903bea0 ("m68k: split ret_from_fork(), simplify kernel_thread()") but I haven't done a detailed study commit. Interrupts cannot be disabled in the saved status register copy for kernel threads (init will complain about i space). Disable interrupts temporarily when switching the tasks' register sets in resume(). Note that a simple oriw 0 - this leaves enough of a race for the 'spinlock recursion' warning to still be observed. Tested on ARAnyM and qem
CVE-2024-38615	In the Linux kernel, the following vulnerability has been resolved: cpufreq: exit() callback is optional The exit() ca without checking a valid pointer first. Also, we must clear freq_table pointer even if the exit() callback isn't present
CVE-2024-38619	In the Linux kernel, the following vulnerability has been resolved: usb-storage: alauda: Check whether the media is struct alauda_info will remain 0 if alauda_init_media() fails, potentially causing divide errors in alauda_read_data("media_initialized" to struct alauda_info. - Change a condition in alauda_check_media() to ensure the first initializ value of alauda_init_media().
CVE-2024-38621	In the Linux kernel, the following vulnerability has been resolved: media: stk1160: fix bounds checking in stk1160 is reversed. The ->length is the length of the buffer. The ->bytesused is how many bytes we have copied thus far. V the result of the subtraction is always negative but since it's unsigned then the result is a very high positive value. T Additionally, the ->bytesused doesn't actually work for this purpose because we're not writing to "buf->mem + buf the destination where we are writing is a bit involved. You calculate the number of full lines already written, multip start on an odd numbered line, and add the offset into the line. To fix this buffer overflow, just take the actual desti already out of bounds print an error and return. Otherwise, write up to buf->length bytes.
CVE-2024-38627	In the Linux kernel, the following vulnerability has been resolved: stm class: Fix a double free in stm_register_dev trigger stm_device_release() which frees "stm" so the vfree(stm) on the next line is a double free.
CVE-2024-38630	In the Linux kernel, the following vulnerability has been resolved: watchdog: cpu5wdt.c: Fix use-after-free bug cau module is removing, the origin code uses del_timer() to de-activate the timer. If the timer handler is running, del_ti directly. If the port region is released by release_region() and then the timer handler cpu5wdt_trigger() calls outb() use-after-free bug will happen. Change del_timer() to timer_shutdown_sync() in order that the timer handler could
CVE-2024-38633	In the Linux kernel, the following vulnerability has been resolved: serial: max3100: Update uart_driver_registered MAX3100 device triggers the removal of the driver. However, code doesn't update the respective global variable a the kernel oopses: max3100 spi-PRP0001:01: max3100_probe: adding port 0 BUG: kernel NULL pointer dereferer 0010:serial_core_register_port+0xa0/0x840 ... max3100_probe+0x1b6/0x280 [max3100] spi_probe+0x8d/0xb0 Up driver will be registered again. Hugo also noticed, that the error path in the probe also affected by having the variab move the assignment after the successfull uart_register_driver() call.
CVE-2024-38634	In the Linux kernel, the following vulnerability has been resolved: serial: max3100: Lock port->lock when calling uart_handle_cts_change() has to be called with port lock taken, Since we run it in a separate work, the lock may no that it's taken by explicitly doing that. Without it we got a splat: WARNING: CPU: 0 PID: 10 at drivers/tty/serial/s +0xa6/0xb0 ... Workqueue: max3100-0 max3100_work [max3100] RIP: 0010:uart_handle_cts_change+0xa6/0xb0 [max3100] max3100_work+0x12a/0x340 [max3100]
CVE-2024-38637	In the Linux kernel, the following vulnerability has been resolved: greybus: lights: check return of get_channel_fro is not found we return null from get_channel_from_mode. Make sure we validate the return pointer before using it originally reported in [0]: Found by Linux Verification Center (linuxtesting.org) with SVACE. [0] https://lore.kern m.lobanov@rosalinux.ru
CVE-2024-38659	In the Linux kernel, the following vulnerability has been resolved: enic: Validate length of nl attributes in enic_set attribute IFLA_PORT_PROFILE is of length PORT_PROFILE_MAX and that the nl attributes IFLA_PORT_INS are of length PORT_UUID_MAX. These attributes are validated (in the function do_setlink in rtnetlink.c) using th defines IFLA_PORT_PROFILE as NLA_STRING, IFLA_PORT_INSTANCE_UUID as NLA_BINARY and IFL That means that the length validation using the policy is for the max size of the attributes and not on exact size so t the sizes that enic_set_vf_port expects. This might cause an out of bands read access in the memcpys of the data of

CVE-2024-38661	In the Linux kernel, the following vulnerability has been resolved: s390/ap: Fix crash in AP internal function modi address: 200000cb7df6f000 TEID: 200000cb7df6f403 Fault in home space mode while using kernel ASCE. AS:00 S:000000011a446000 P:000000015660c13d Oops: 0038 ilc:3 [#1] PREEMPT SMP Modules linked in: mlx5_ib ... 6.9.0-rc7 #8 Hardware name: IBM 3931 A01 704 (LPAR) Krnl PSW : 0704e00180000000 0000014b75e7b606 (ap T:1 IO:1 EX:1 Key:0 M:1 W:0 P:0 AS:3 CC:2 PM:0 RI:0 EA:3 Krnl GPRS: 0000000000000001 fffffffffffffc0 00 000000cb00000100 ffffffffffffffff ffffffffffffffff 000000cb7df6fce0 000000cb7df6fce0 00000000ffffffff 000000000 000003ff9b2dbc80 200000cb7df6fcd8 0000014bffffffc0 000000cb7df6fbc8 Krnl Code: 0000014b75e7b5fc: a7840 0000014b75e7b600: 18b2 lr %r11,%r2 #0000014b75e7b602: a7f4000a brc 15,0000014b75e7b616 >0000014b75e 0000014b75e7b60c: a7680001 lhi %r6,1 0000014b75e7b610: 187b lr %r7,%r11 0000014b75e7b612: 84960021 br 0000014b75e7b616: 18e9 lr %r14,%r9 Call Trace: [<0000014b75e7b606>] ap_parse_bitmap_str+0x10e/0x1f8 ([< +0xe4/0x1f8) [<0000014b75e7b758>] apmask_store+0x68/0x140 [<0000014b75679196>] kernfs_fop_write_iter+ vfs_write+0x1b4/0x448 [<0000014b7559894c>] ksys_write+0x74/0x100 [<0000014b7618a440>] __do_syscall+0 system_call+0x70/0x98 INFO: lockdep is turned off. Last Breaking-Event-Address: [<0000014b75e7b636>] ap_pa not syncing: Fatal exception: panic_on_oops occured when /sys/bus/ap/a[pq]mask was updated with a relative mas of the numeric values exceeding INT_MAX. The fix is simple: use unsigned long values for the internal variables. function but a simple int for the internal variables was used with the possibility to overflow.
CVE-2024-38780	In the Linux kernel, the following vulnerability has been resolved: dma-buf/sw-sync: don't enable IRQ from sync_ ("dma-buf/sw-sync: Reduce irqsave/irqrestore from known context") by error replaced spin_unlock_irqrestore() wi sync_debugfs_show() and sync_print_obj() despite sync_print_obj() is called from sync_debugfs_show(), lockdep Use plain spin_{lock,unlock}() for sync_print_obj(), for sync_debugfs_show() is already using spin_{lock,unlock}
CVE-2024-38828	Spring MVC controller methods with an @RequestBody byte[]¬†method parameter are vulnerable to a DoS attack
CVE-2024-38949	Heap Buffer Overflow vulnerability in Libde265 v1.0.15 allows attackers to crash the application via crafted paylo
CVE-2024-38950	Heap Buffer Overflow vulnerability in Libde265 v1.0.15 allows attackers to crash the application via crafted paylo
CVE-2024-39276	In the Linux kernel, the following vulnerability has been resolved: ext4: fix mb_cache_entry's e_refcnt leak in ext4 a warning as follows: ========== WARNING: CPU: 0 PID: 507 +0x224/0x290 Modules linked in: CPU: 0 PID: 5075 Comm: syz-executor199 Not tainted 6.9.0-rc6-gb947cc5bf6d +0x224/0x290 fs/mbcache.c:419 Call Trace: <TASK> ext4_put_super+0x6d4/0xcd0 fs/ext4/super.c:1375 generic_ kill_block_super+0x44/0x90 fs/super.c:1675 ext4_kill_sb+0x68/0xa0 fs/ext4/super.c:7327 [...] ========== This is because when finding an entry in ext4_xattr_block_cache_find(), if ext4_sb_bread() returns -ENOMEM, th grown in the __entry_find(), won't be put away, and eventually trigger the above issue in mb_cache_destroy() due mb_cache_entry_put() on the -ENOMEM error branch as a quick fix.
CVE-2024-39292	In the Linux kernel, the following vulnerability has been resolved: um: Add winch to winch_handlers before regist racy, an interrupt may occur before the winch is added to the winch_handlers list. If that happens, register_winch_i to be (or has already been) freed, causing a panic later in winch_cleanup(). Avoid the race by adding the winch to t IRQ, and rolling back if um_request_irq() fails.
CVE-2024-39301	In the Linux kernel, the following vulnerability has been resolved: net/9p: fix uninit-value in p9_client_rpc() Syzbo following error: BUG: KMSAN: uninit-value in trace_9p_client_res include/trace/events/9p.h:146 [inline] BUG: K +0x1314/0x1340 net/9p/client.c:754 trace_9p_client_res include/trace/events/9p.h:146 [inline] p9_client_rpc+0x13 p9_client_create+0x1551/0x1ff0 net/9p/client.c:1031 v9fs_session_init+0x1b9/0x28e0 fs/9p/v9fs.c:410 v9fs_mour legacy_get_tree+0x114/0x290 fs/fs_context.c:662 vfs_get_tree+0xa7/0x570 fs/super.c:1797 do_new_mount+0x71 +0x742/0x1f20 fs/namespace.c:3679 do_mount fs/namespace.c:3692 [inline] __do_sys_mount fs/namespace.c:389 fs/namespace.c:3875 __x64_sys_mount+0xe4/0x150 fs/namespace.c:3875 do_syscall_64+0xd5/0x1f0 entry_SYSC Uninit was created at: __alloc_pages+0x9d6/0xe70 mm/page_alloc.c:4598 __alloc_pages_node include/linux/gfp.h include/linux/gfp.h:261 [inline] alloc_slab_page mm/slub.c:2175 [inline] allocate_slab mm/slub.c:2338 [inline] ne ___slab_alloc+0x1184/0x33d0 mm/slub.c:3525 __slab_alloc mm/slub.c:3610 [inline] __slab_alloc_node mm/slub. slub.c:3835 [inline] kmem_cache_alloc+0x6d3/0xbe0 mm/slub.c:3852 p9_tag_alloc net/9p/client.c:278 [inline] p9_ net/9p/client.c:641 p9_client_rpc+0x27e/0x1340 net/9p/client.c:688 p9_client_create+0x1551/0x1ff0 net/9p/client. fs/9p/v9fs.c:410 v9fs_mount+0xe2/0x12b0 fs/9p/vfs_super.c:122 legacy_get_tree+0x114/0x290 fs/fs_context.c:66 super.c:1797 do_new_mount+0x71f/0x15e0 fs/namespace.c:3352 path_mount+0x742/0x1f20 fs/namespace.c:3679 __do_sys_mount fs/namespace.c:3898 [inline] __se_sys_mount+0x725/0x810 fs/namespace.c:3875 __x64_sys_mo do_syscall_64+0xd5/0x1f0 entry_SYSCALL_64_after_hwframe+0x6d/0x75 If p9_check_errors() fails early in p9_ initialized. However, trace_9p_client_res() ends up trying to print it out anyway before p9_client_rpc() finishes. Fi p9_fcall fields such as 'tag' and (just in case KMSAN unearths something new) 'id' during the tag allocation stage.

CVE-2024-39467	In the Linux kernel, the following vulnerability has been resolved: f2fs: fix to do sanity check on i_xattr_nid in sanity_check_inode() syzbot reports a kernel bug as below: F2FS-fs (loop0): Mounted with checkpoint version = 4 ========== BUG: KASAN: slab- f2fs/f2fs.h:2933 [inline] BUG: KASAN: slab-out-of-bounds in current_nat_addr fs/f2fs/node.h:213 [inline] BUG: f2fs_get_node_info+0xece/0x1200 fs/f2fs/node.c:600 Read of size 1 at addr ffff88807a58c76c by task syz-executo syz-executor280 Not tainted 6.9.0-rc5-syzkaller #0 Hardware name: Google Google Compute Engine/Google Com Call Trace: <TASK> __dump_stack lib/dump_stack.c:88 [inline] dump_stack_lvl+0x241/0x360 lib/dump_stack.c: report.c:377 [inline] print_report+0x169/0x550 mm/kasan/report.c:488 kasan_report+0x143/0x180 mm/kasan/repo [inline] current_nat_addr fs/f2fs/node.h:213 [inline] f2fs_get_node_info+0xece/0x1200 fs/f2fs/node.c:600 f2fs_xat f2fs_fiemap+0x55d/0x1ee0 fs/f2fs/data.c:1925 ioctl_fiemap fs/ioctl.c:220 [inline] do_vfs_ioctl+0x1c07/0x2e50 fs/ [inline] __se_sys_ioctl+0x81/0x170 fs/ioctl.c:890 do_syscall_x64 arch/x86/entry/common.c:52 [inline] do_syscall_ entry_SYSCALL_64_after_hwframe+0x77/0x7f The root cause is we missed to do sanity check on i_xattr_nid dur current_nat_addr() will access nat_bitmap w/ offset from invalid i_xattr_nid, result in triggering kasan bug report,
CVE-2024-39468	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix deadlock in smb2_find_smb_tco cifs_put_smb_ses() to avoid such deadlock.
CVE-2024-39469	In the Linux kernel, the following vulnerability has been resolved: nilfs2: fix nilfs_empty_dir() misjudgment and lo handling in nilfs_empty_dir() when a directory folio/page read fails is incorrect, as in the old ext2 implementation, nilfs_check_folio() fails, it will falsely determine the directory as empty and corrupt the file system. In addition, sin return on a failed folio/page read, but continues to loop, this can cause a long loop with I/O if i_size of the directory log writer thread to wait and hang, as reported by syzbot. Fix these issues by making nilfs_empty_dir() immediatel directory folio/page.
CVE-2024-39471	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: add error handle to avoid out-of-b -EINVAL, the process should be stop to avoid out-of-bounds read, so directly return -EINVAL.
CVE-2024-39472	In the Linux kernel, the following vulnerability has been resolved: xfs: fix log recovery buffer allocation for the leg detect and handle invalid iclog size set by mkfs") added a fixup for incorrect h_size values used for the initial umou commit 0c771b99d6c9 ("xfs: clean up calculation of LR header blocks") cleaned up the log reover buffer calculatio to size the log recovery buffer, which can lead to an out of bounds access when the incorrect h_size does not come by open coding xlog_logrec_hblks and taking the fixed h_size into account for this calculation.
CVE-2024-39475	In the Linux kernel, the following vulnerability has been resolved: fbdev: savage: Handle err return when savagefb 04e5eac8f3ab("fbdev: savage: Error out if pixclock equals zero") checks the value of pixclock to avoid divide-by-z savagefb_probe doesn't handle the error return of savagefb_check_var. When pixclock is 0, it will cause divide-by-
CVE-2024-39476	In the Linux kernel, the following vulnerability has been resolved: md/raid5: fix deadlock that raid5d() wait for itse Xiao reported that lvm2 test lvconvert-raid-takeover.sh can hang with small possibility, the root cause is exactly th "md/raid5: Wait for MD_SB_CHANGE_PENDING in raid5d"") However, Dan reported another hang after that, a found out that this is caused by plugged bio can't issue from raid5d(). Current implementation in raid5d() has a wei from raid5d() must hold 'reconfig_mutex' to clear MD_SB_CHANGE_PENDING; 2) raid5d() handles IO in a deac from raid5d() must wait for MD_SB_CHANGE_PENDING to be cleared; This behaviour is introduce before v2.6, hold 'reconfig_mutex', and md_check_recovery() can't update super_block, then raid5d() will waste one cpu 100% is released. Refer to the implementation from raid1 and raid10, fix this problem by skipping issue IO if MD_SB_C md_check_recovery(), daemon thread will be woken up when 'reconfig_mutex' is released. Meanwhile, the hang pr
CVE-2024-39484	In the Linux kernel, the following vulnerability has been resolved: mmc: davinci: Don't strip remove function when function results in the remove callback being discarded with CONFIG_MMC_DAVINCI=y. When such a device g the driver is just removed without the cleanup being performed. This results in resource leaks. Fix it by compiling i This also fixes a W=1 modpost warning: WARNING: modpost: drivers/mmc/host/davinci_mmc: section mismatch (section: .data) -> davinci_mmcsd_remove (section: .exit.text)
CVE-2024-39487	In the Linux kernel, the following vulnerability has been resolved: bonding: Fix out-of-bounds read in bond_option bond_option_arp_ip_targets_set(), if newval->string is an empty string, newval->string+1 will point to the byte aft read. BUG: KASAN: slab-out-of-bounds in strlen+0x7d/0xa0 lib/string.c:418 Read of size 1 at addr ffff8881119c4 1 PID: 8107 Comm: syz-executor665 Not tainted 6.7.0-rc7 #1 Hardware name: QEMU Standard PC (i440FX + PI Call Trace: <TASK> __dump_stack lib/dump_stack.c:88 [inline] dump_stack_lvl+0xd9/0x150 lib/dump_stack.c:1 report.c:364 [inline] print_report+0xc1/0x5e0 mm/kasan/report.c:475 kasan_report+0xbe/0xf0 mm/kasan/report.c:5 __fortify_strlen include/linux/fortify-string.h:210 [inline] in4_pton+0xa3/0x3f0 net/core/utils.c:130 bond_option_a bonding/bond_options.c:1201 __bond_opt_set+0x2a4/0x1030 drivers/net/bonding/bond_options.c:767 __bond_opt bond_options.c:792 bond_opt_tryset_rtnl+0xda/0x160 drivers/net/bonding/bond_options.c:817 bonding_sysfs_stor bond_sysfs.c:156 dev_attr_store+0x54/0x80 drivers/base/core.c:2366 sysfs_kf_write+0x114/0x170 fs/sysfs/file.c:1 fs/kernfs/file.c:334 call_write_iter include/linux/fs.h:2020 [inline] new_sync_write fs/read_write.c:491 [inline] vfs_ ksys_write+0x122/0x250 fs/read_write.c:637 do_syscall_x64 arch/x86/entry/common.c:52 [inline] do_syscall_64+ entry_SYSCALL_64_after_hwframe+0x63/0x6b ---[end trace]--- Fix it by adding a check of string length before
CVE-2024-39489	In the Linux kernel, the following vulnerability has been resolved: ipv6: sr: fix memleak in seg6_hmac_init_algo s up the previous allocations if one fails, so it's going to leak all that memory and the crypto tfms. Update seg6_hmac allocated, so we can reuse the code directly.

CVE-2024-39493	In the Linux kernel, the following vulnerability has been resolved: crypto: qat - Fix ADF_DEV_RESET_SYNC me to determine whether the caller has gone away only works after a complete call. Furthermore it's still possible that t wait_for_completion, resulting in another potential UAF. Fix this by making the caller use cancel_work_sync and t
CVE-2024-39495	In the Linux kernel, the following vulnerability has been resolved: greybus: Fix use-after-free bug in gb_interface_ condition. In gb_interface_create, &intf->mode_switch_completion is bound with gb_interface_mode_switch_wor by gb_interface_request_mode_switch. Here is the relevant code. if (!queue_work(system_long_wq, &intf->mode_ call gb_interface_release to make cleanup, there may be an unfinished work. This function will call kfree to free th gb_interface_mode_switch_work is scheduled to run after kfree, it may cause use-after-free error as gb_interface_r "intf". The possible execution flow that may lead to the issue is as follows: CPU0 CPU1 | gb_interface_create | gb_ gb_interface_release | kfree(intf) (free) | | gb_interface_mode_switch_work | mutex_lock(&intf->mutex) (use) Fix i
CVE-2024-39499	In the Linux kernel, the following vulnerability has been resolved: vmci: prevent speculation leaks by sanitizing ev that event_msg is controlled by user-space, event_msg->event_data.event is passed to event_deliver() and used as a ensures that the event index is sanitized to mitigate any possibility of speculative information leaks. This bug was d Analysis Security Testing (SAST) by Synopsys, Inc. Only compile tested, no access to HW.
CVE-2024-39501	In the Linux kernel, the following vulnerability has been resolved: drivers: core: synchronize really_probe() and de usage in really_probe() and dev_uevent(). These can run in different threads, what can result in the following race c Thread #1: ========== really_probe() { ... probe_failed: ... device_unbind_cleanup(dev) { ... dev->driver = NUL to NULL ... } ... } Thread #2: ========== dev_uevent() { ... if (dev->driver) // If dev->driver is NULLed from re check, the system crashes add_uevent_var(env, "DRIVER=%s", dev->driver->name); ... } really_probe() holds the done there. dev_uevent() is called with lock held, often, too. But not always. What implies that we can't add any loc race by adding the lock to the non-protected path. This is the path where above race is observed: dev_uevent+0x23 Add lock here dev_attr_show+0x3a/0xa0 sysfs_kf_seq_show+0x17c/0x250 kernfs_seq_show+0x7c/0x90 seq_read +0xc6/0x310 vfs_read+0x5bc/0x6b0 ksys_read+0xeb/0x1b0 __x64_sys_read+0x42/0x50 x64_sys_call+0x27ad/0x entry_SYSCALL_64_after_hwframe+0x77/0x7f Similar cases are reported by syzkaller in https://syzkaller.appspo these are regarding the *initialization* of dev->driver dev->driver = drv; As this switches dev->driver to non-NUL positives (which should be "fixed" by this commit, as well, though). The same issue was reported and tried to be fi lkml/1421259054-2574-1-git-send-email-a.sangwan@samsung.com/ already.
CVE-2024-39502	In the Linux kernel, the following vulnerability has been resolved: ionic: fix use after netif_napi_del() When queue napi_enable() are called. If there are 4 queues and only 3 queues are used for the current configuration, only 3 queu The ionic_qcq_enable() checks whether the .poll pointer is not NULL for enabling only the using queue' napi. Unu by netif_napi_add(), so the .poll pointer indicates NULL. But it couldn't distinguish whether the napi was unregiste doesn't reset the .poll pointer to NULL. So, ionic_qcq_enable() calls napi_enable() for the queue, which was unregi ethtool -L <interface name> rx 1 tx 1 combined 0 ethtool -L <interface name> rx 0 tx 0 combined 1 ethtool -L <int Splat looks like: kernel BUG at net/core/dev.c:6666! Oops: invalid opcode: 0000 [#1] PREEMPT SMP NOPTI CP Not tainted 6.10.0-rc2+ #16 Workqueue: events ionic_lif_deferred_work [ionic] RIP: 0010:napi_enable+0x3b/0x40 61 09 00 00 00 74 0d 48 83 bf 60 01 00 00 00 74 03 80 ce 01 f0 4f RSP: 0018:ffffb6ed83227d48 EFLAGS: 000102 ffff97560cda0828 RCX: 0000000000000029 RDX: 0000000000000001 RSI: 0000000000000000 RDI: ffff97560c 0000000000000400 R09: 0000000000000001 R10: 0000000000000001 R11: 0000000000000001 R12: 000000000 R14: 0000000000000000 R15: ffff975613ba0a20 FS: 0000000000000000(0000) GS:ffff975d5f780000(0000) knlG DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 00007f8f734ee200 CR3: 0000000103e50000 CR4: 000000000 Trace: <TASK> ? die+0x33/0x90 ? do_trap+0xd9/0x100 ? napi_enable+0x3b/0x40 ? do_error_trap+0x83/0xb0 ? n +0x3b/0x40 ? exc_invalid_op+0x4e/0x70 ? napi_enable+0x3b/0x40 ? asm_exc_invalid_op+0x16/0x20 ? napi_ena +0xb7/0x180 [ionic 59bdfc8a035436e1c4224ff7d10789e3f14643f8] ionic_start_queues+0xc4/0x290 [ionic 59bdfc ionic_link_status_check+0x11c/0x170 [ionic 59bdfc8a035436e1c4224ff7d10789e3f14643f8] ionic_lif_deferred_w 59bdfc8a035436e1c4224ff7d10789e3f14643f8] process_one_work+0x145/0x360 worker_thread+0x2bb/0x3d0 ? _ +0xcc/0x100 ? __pfx_kthread+0x10/0x10 ret_from_fork+0x2d/0x50 ? __pfx_kthread+0x10/0x10 ret_from_fork_a
CVE-2024-39503	In the Linux kernel, the following vulnerability has been resolved: netfilter: ipset: Fix race between namespace clea Ackermann reported that there is a race condition between namespace cleanup in ipset and the garbage collection o destroy the list:set type of sets while the gc of the set type is waiting to run in rcu cleanup. The latter uses data from free. The patch contains the following parts: - When destroying all sets, first remove the garbage collectors, then w the badly ordered "wait then remove gc" for the destroy a single set case. - Fix the missing rcu locking in the list:se RCU list handlings in the list:set type. The patch depends on c1193d9bbbd3 (netfilter: ipset: Add list flush to cance
CVE-2024-39504	In the Linux kernel, the following vulnerability has been resolved: netfilter: nft_inner: validate mandatory meta and attributes in payload and meta expression when used embedded from the inner expression, otherwise NULL pointe
CVE-2024-39505	In the Linux kernel, the following vulnerability has been resolved: drm/komeda: check for error-valued pointer kor error-valued pointer, thus check the pointer for negative or null value before dereferencing.
CVE-2024-39506	In the Linux kernel, the following vulnerability has been resolved: liquidio: Adjust a NULL pointer handling path i lio_vf_rep_copy_packet() pg_info->page is compared to a NULL value, but then it is unconditionally passed to skb could lead to null pointer dereference. lio_vf_rep_copy_packet() call trace looks like: octeon_droq_process_packet octeon_droq_dispatch_pkt octeon_create_recv_info ...search in the dispatch_list... ->disp_fn(rdisp->rinfo, ...) lio_v *recv_info, ...) In this path there is no code which sets pg_info->page to NULL. So this check looks unneeded and the author had reason to add a check and I have no such card and can't do real test. In addition, the code in the func lio_core.c does exactly the same. Based on this, I consider the most acceptable compromise solution to adjust this i conditional scope. Found by Linux Verification Center (linuxtesting.org) with SVACE.

CVE-2024-4032	The ,Äúipaddress‚Äù module contained incorrect information about whether certain IPv4 and IPv6 addresses were or ,Äúprivate‚Äù. This affected the is_private and is_global properties of the ipaddress.IPv4Address, ipaddress.IPv ipaddress.IPv6Network classes, where values wouldn‚Äôt be returned in accordance with the latest information fro Registries. CPython 3.12.4 and 3.13.0a6 contain updated information from these registries and thus have the intend
CVE-2024-4032	The ,Äúipaddress‚Äù module contained incorrect information about whether certain IPv4 and IPv6 addresses were or ,Äúprivate‚Äù. This affected the is_private and is_global properties of the ipaddress.IPv4Address, ipaddress.IPv ipaddress.IPv6Network classes, where values wouldn‚Äôt be returned in accordance with the latest information fro Registries. CPython 3.12.4 and 3.13.0a6 contain updated information from these registries and thus have the intend
CVE-2024-4068	The NPM package `braces`, versions prior to 3.0.3, fails to limit the number of characters it can handle, which coul parse.js,` if a malicious user sends "imbalanced braces" as input, the parsing will enter a loop, which will cause the without freeing it at any moment of the loop. Eventually, the JavaScript heap limit is reached, and the program will
CVE-2024-40901	In the Linux kernel, the following vulnerability has been resolved: scsi: mpt3sas: Avoid test/set_bit() operating in n potential out-of-bounds access when using test_bit() on a single word. The test_bit() and set_bit() functions operate or setting a single word, they can exceed the word boundary. KASAN detects this issue and produces a dump: BUG in _scsih_add_device.constprop.0 (./arch/x86/include/asm/bitops.h:60 ./include/asm-generic/bitops/instrumented-at mpt3sas_scsih.c:7331) mpt3sas Write of size 8 at addr ffff8881d26e3c60 by task kworker/u1536:2/2965 For full lo least the size of sizeof(unsigned long) so that set_bit() and test_bit() have sufficient room for read/write operations Link: https://lore.kernel.org/all/ZkNcALr3W3KGYYJG@gmail.com/
CVE-2024-40902	In the Linux kernel, the following vulnerability has been resolved: jfs: xattr: fix buffer overflow for invalid xattr W printed out to the kernel log in hex format as a form of debugging. But when that xattr size is bigger than the expec off the end of the buffer. Fix this all up by properly restricting the size of the debug hex dump in the kernel log.
CVE-2024-40904	In the Linux kernel, the following vulnerability has been resolved: USB: class: cdc-wdm: Fix CPU lockup caused b found that the interrupt-URB completion callback in the cdc-wdm driver was taking too long, and the driver's imm -EPROTO status combined with the dummy-hcd emulation to cause a CPU lockup: cdc_wdm 1-1:1.0: nonzero urb wdm_int_callback - 0 bytes watchdog: BUG: soft lockup - CPU#0 stuck for 26s! [syz-executor782:6625] CPU#0 U system, 0% softirq, 3% hardirq, 0% idle #2: 98% system, 0% softirq, 3% hardirq, 0% idle #3: 98% system, 0% soft 0% softirq, 3% hardirq, 0% idle #5: 98% system, 1% softirq, 3% hardirq, 0% idle Modules linked in: irq event stan [<ffff80008037bc00>] console_emit_next_record kernel/printk/printk.c:2935 [inline] hardirqs last enabled at (7309 +0x650/0xb74 kernel/printk/printk.c:2994 hardirqs last disabled at (73096): [<ffff80008af10b00>] __el1_irq arch/a hardirqs last disabled at (73096): [<ffff80008af10b00>] el1_interrupt+0x24/0x68 arch/arm64/kernel/entry-common [<ffff8000801ea530>] softirq_handle_end kernel/softirq.c:400 [inline] softirqs last enabled at (73048): [<ffff80008 kernel/softirq.c:582 softirqs last disabled at (73043): [<ffff800080020de8>] __do_softirq+0x14/0x20 kernel/softirq executor782 Tainted: G W 6.10.0-rc2-syzkaller-g8867bbd4a056 #0 Hardware name: Google Google Compute Eng 04/02/2024 Testing showed that the problem did not occur if the two error messages -- the first two lines above -- w to the kernel log takes a surprisingly large amount of time. In any case, the best approach for preventing these lock thousands of error messages per second is to ratelimit the two dev_err() calls. Therefore we replace them with dev_
CVE-2024-40908	In the Linux kernel, the following vulnerability has been resolved: bpf: Set run context for rawtp test_run callback executed through the test_run interface calls bpf_get_attach_cookie helper or any other helper that touches task->b >bpf_ctx pointer) for test_run callback.
CVE-2024-40912	In the Linux kernel, the following vulnerability has been resolved: wifi: mac80211: Fix deadlock in ieee80211_sta ieee80211_sta_ps_deliver_wakeup() function takes sta->ps_lock to synchronizes with ieee80211_tx_h_unicast_ps context. However using only spin_lock() to get sta->ps_lock in ieee80211_sta_ps_deliver_wakeup() does not preve CPU, to run ieee80211_tx_h_unicast_ps_buf() and try to take this same lock ending in deadlock. Below is an exam situation. rcu: INFO: rcu_sched self-detected stall on CPU rcu: 2-....: (42413413 ticks this GP) idle=b154/1/0x4000 fqs=21206996 rcu: (t=42586894 jiffies g=2057 q=362405 ncpus=4) CPU: 2 PID: 719 Comm: wpa_supplicant Tain #742 Hardware name: RPT (r1) (DT) pstate: 00000005 (nzcv daif -PAN -UAO -TCO -DIT -SSBS BTYPE=--) pc : +0x58/0x2d0 lr : invoke_tx_handlers_early+0x5b4/0x5c0 sp : ffff00001ef64660 x29: ffff00001ef64660 x28: ffff00 x26: ffff000009bc0900 x25: ffff00001ef647a8 x24: 0000000000000000 x23: ffff000009bc0900 x22: ffff000009bc x20: ffff00000a279e00 x19: ffff00001ef646e8 x18: 0000000000000000 x17: ffff800016468000 x16: ffff00001ef6 x14: 0010395c9faa3946 x13: 0000000000000000 x12: 00000000fa83b2da x11: 000000012edeceea x10: ffff00000 x8 : 000000000010533c x7 : ffff00000ad8b740 x6 : ffff00000c350880 x5 : 0000000000000007 x4 : 0000000000000 x2 : 0000000000000000 x1 : 0000000000000001 x0 : ffff00000ac0e0e8 Call trace: queued_spin_lock_slowpath+0 +0x80/0x12c ieee80211_tx_pending+0x110/0x278 tasklet_action_common.constprop.0+0x10c/0x144 tasklet_actio ____do_softirq+0xc/0x14 call_on_irq_stack+0x24/0x34 do_softirq_own_stack+0x18/0x20 do_softirq+0x74/0x7c _ieee80211_wake_txqs+0x3b0/0x4b8 __ieee80211_wake_queue+0x12c/0x168 ieee80211_add_pending_skbs+0xe +0x2a4/0x480 ieee80211_mps_sta_status_update.part.0+0xd8/0x11c ieee80211_mps_sta_status_update+0x18/0x2 ieee80211_change_station+0x1b8/0x2dc nl80211_set_station+0x444/0x49c genl_family_rcv_msg_doit.isra.0+0xa netlink_rcv_skb+0x38/0x10c genl_rcv+0x34/0x48 netlink_unicast+0x254/0x2bc netlink_sendmsg+0x190/0x3b4 _ ___sys_sendmsg+0x68/0x8c __sys_sendmsg+0x44/0x84 __arm64_sys_sendmsg+0x20/0x28 do_el0_svc+0x6c/0x +0xb0/0xb4 el0t_64_sync+0x14c/0x150 Using spin_lock_bh()/spin_unlock_bh() instead prevents softirq to raise o

CVE-2024-40916	In the Linux kernel, the following vulnerability has been resolved: drm/exynos: hdmi: report safe 640x480 mode as When reading EDID fails and driver reports no modes available, the DRM core adds an artificial 1024x786 mode t variants of the Exynos HDMI (like the one in Exynos4 SoCs) are not able to drive such mode, so report a safe 640x case of the EDID reading failure. This fixes the following issue observed on Trats2 board since commit 13d5b0403 negative values from .get_modes()"): [drm] Exynos DRM: using 11c00000.fimd device for DMA mapping operati 11c00000.fimd (ops fimd_component_ops) exynos-drm exynos-drm: bound 12c10000.mixer (ops mixer_compone [drm:samsung_dsim_host_attach] Attached s6e8aa0 device (lanes:4 bpp:24 mode-flags:0x10b) exynos-drm exynos exynos_dsi_component_ops) exynos-drm exynos-drm: bound 12d00000.hdmi (ops hdmi_component_ops) [drm] I for exynos-drm on minor 1 exynos-hdmi 12d00000.hdmi: [drm:hdmiphy_enable.part.0] *ERROR* PLL could not s6e8aa0 11c80000.dsi.0: ID: 0xa2, 0x20, 0x8c exynos-mixer 12c10000.mixer: timeout waiting for VSYNC -------- CPU: 1 PID: 11 at drivers/gpu/drm/drm_atomic_helper.c:1682 drm_atomic_helper_wait_for_vblanks.part.0+0x2b0 timed out Modules linked in: CPU: 1 PID: 11 Comm: kworker/u16:0 Not tainted 6.9.0-rc5-next-20240424 #14913 (Flattened Device Tree) Workqueue: events_unbound deferred_probe_work_func Call trace: unwind_backtrace fro from dump_stack_lvl+0x68/0x88 dump_stack_lvl from __warn+0x7c/0x1c4 __warn from warn_slowpath_fmt+0x drm_atomic_helper_wait_for_vblanks.part.0+0x2b0/0x2b8 drm_atomic_helper_wait_for_vblanks.part.0 from drm drm_atomic_helper_commit_tail_rpm from commit_tail+0x9c/0x184 commit_tail from drm_atomic_helper_comm from drm_atomic_commit+0xb4/0xe0 drm_atomic_commit from drm_client_modeset_commit_atomic+0x23c/0x2 from drm_client_modeset_commit_locked+0x60/0x1cc drm_client_modeset_commit_locked from drm_client_mo drm_client_modeset_commit from __drm_fb_helper_restore_fbdev_mode_unlocked+0x9c/0xc4 __drm_fb_helper drm_fb_helper_set_par+0x2c/0x3c drm_fb_helper_set_par from fbcon_init+0x3d8/0x550 fbcon_init from visual_i do_bind_con_driver+0x1b8/0x3a4 do_bind_con_driver from do_take_over_console+0x140/0x1ec do_take_over_c +0x70/0xd0 do_fbcon_takeover from fbcon_fb_registered+0x19c/0x1ac fbcon_fb_registered from register_frameb from __drm_fb_helper_initial_config_and_unlock+0x350/0x574 __drm_fb_helper_initial_config_and_unlock fro +0x6c/0xb0 exynos_drm_fbdev_client_hotplug from drm_client_register+0x58/0x94 drm_client_register from exy exynos_drm_bind from try_to_bring_up_aggregate_device+0x200/0x2d8 try_to_bring_up_aggregate_device from __component_add from mixer_probe+0x74/0xcc mixer_probe from platform_probe+0x5c/0xb8 platform_probe fr from __driver_probe_device+0x9c/0x1e4 __driver_probe_device from driver_probe_device+0x30/0xc0 driver_pro +0xa8/0x120 __device_attach_driver from bus_for_each_drv+0x80/0xcc bus_for_each_drv from __device_attach- bus_probe_device+0x8c/0x90 bus_probe_device from deferred_probe_work_func+0 ---truncated---
CVE-2024-40931	In the Linux kernel, the following vulnerability has been resolved: mptcp: ensure snd_una is properly initialized on fb7a0d334894 ("mptcp: ensure snd_nxt is properly initialized on connect"). It turns out that syzkaller can trigger th processing any other incoming packet - so that snd_una is still left uninitialized. Address the issue explicitly initiali write_seq.
CVE-2024-40932	In the Linux kernel, the following vulnerability has been resolved: drm/exynos/vidi: fix memory leak in .get_mode
CVE-2024-40934	In the Linux kernel, the following vulnerability has been resolved: HID: logitech-dj: Fix memory leak in logi_dj_re on logi_dj_recv_send_report() error path.
CVE-2024-40935	In the Linux kernel, the following vulnerability has been resolved: cachefiles: flush all requests after setting CACH the daemon is processing an open request, if the kernel flags the cache as CACHEFILES_DEAD, the cachefiles_da daemon can't pass the copen to the kernel. Then the kernel process that is waiting for the copen triggers a hung_tas only be exited by closing /dev/cachefiles. Therefore, after calling cachefiles_io_error() to mark the cache as CACH all requests to avoid the above hungtask. We may still be able to read some of the cached data before closing the fd patch that adds reference counting to the req, otherwise it may UAF.
CVE-2024-40936	In the Linux kernel, the following vulnerability has been resolved: cxl/region: Fix memregion leaks in devm_cxl_a __create_region() before allocating the memregion to avoid the memregion leaks.
CVE-2024-40939	In the Linux kernel, the following vulnerability has been resolved: net: wwan: iosm: Fix tainted pointer delete is ca creation fail in ipc_devlink_create_region(), previously created regions delete process starts from tainted pointer w bug by decreasing region index before delete. Found by Linux Verification Center (linuxtesting.org) with SVACE.
CVE-2024-40941	In the Linux kernel, the following vulnerability has been resolved: wifi: iwlwifi: mvm: don't read past the mfuart n notification that claims it has more data than it has, we will read past that was allocated for the notification. Remov default. If needed, we can see the content with tracing. This was reported by KFENCE.

CVE-2024-40942	In the Linux kernel, the following vulnerability has been resolved: wifi: mac80211: mesh: Fix leak of mesh_preq_ of type mesh_preq_queue, added to a list in ieee80211_if_mesh, to keep track of mpath we need to resolve. If the n removed, the entries in that list will never get cleaned. Fix this by flushing all corresponding items of the preq_que should take care of KASAN reports like this: unreferenced object 0xffff00000668d800 (size 128): comm "kworker 1836.444s) hex dump (first 32 bytes): 00 1f 05 09 00 00 ff ff 00 d5 68 06 00 00 ff ff .....h..... 8e 97 ea eb 3e b8 0 backtrace: [<000000007302a0b6>] __kmem_cache_alloc_node+0x1e0/0x35c [<00000000049bd418>] kmalloc_tra mesh_queue_preq+0x44/0x2a8 [<00000000c99c3696>] mesh_nexthop_resolve+0x198/0x19c [<00000000926bf59 [<00000000fc8c2284>] __ieee80211_subif_start_xmit+0x30c/0x764 [<000000005926ee38>] ieee80211_subif_sta dev_hard_start_xmit+0x174/0x440 [<0000000023495647>] __dev_queue_xmit+0xe24/0x111c [<00000000cfe9ca [<000000007bacc5d5>] batadv_v_elp_periodic_work+0x2f4/0x508 [<00000000adc3cd94>] process_one_work+0 worker_thread+0x9c/0x634 [<0000000005852dd5>] kthread+0x1bc/0x1c4 [<000000005fccd770>] ret_from_fork+ 0xffff000009051f00 (size 128): comm "kworker/u8:4", pid 67, jiffies 4295419553 (age 1836.440s) hex dump (first 68 06 00 00 ff ff .....h..... 36 27 92 e4 02 e0 01 00 00 58 79 06 00 00 ff ff 6'.....Xy..... backtrace: [<0000000073 +0x1e0/0x35c [<00000000049bd418>] kmalloc_trace+0x34/0x80 [<0000000000d792bb>] mesh_queue_preq+0x4 mesh_nexthop_resolve+0x198/0x19c [<00000000926bf598>] ieee80211_xmit+0x1d0/0x1f4 [<00000000fc8c2284 +0x30c/0x764 [<000000005926ee38>] ieee80211_subif_start_xmit+0x9c/0x7a4 [<000000004c86e916>] dev_hard [<0000000023495647>] __dev_queue_xmit+0xe24/0x111c [<00000000cfe9ca78>] batadv_send_skb_packet+0x1 batadv_v_elp_periodic_work+0x2f4/0x508 [<00000000adc3cd94>] process_one_work+0x4b8/0xa1c [<0000000 [<0000000005852dd5>] kthread+0x1bc/0x1c4 [<000000005fccd770>] ret_from_fork+0x10/0x20
CVE-2024-40943	In the Linux kernel, the following vulnerability has been resolved: ocfs2: fix races between hole punching and AIO commit "ocfs2: return real error code in ocfs2_dio_wr_get_block", fstests/generic/300 become from always failed t == [473.293420 [475.296983] JBD2: Ignoring recovery information on journal [475.302473] ocfs2: Mounting device (253,1) on data mode. [494.290998] OCFS2: ERROR (device dm-1): ocfs2_change_extent_flag: Owner 5668 has an extent a longer be found [494.291609] On-disk corruption discovered. Please run fsck.ocfs2 once the filesystem is unmoun File system is now read-only. [494.292224] (kworker/19:11,2628,19):ocfs2_mark_extent_written:5272 ERROR: (kworker/19:11,2628,19):ocfs2_dio_end_io_write:2374 ERROR: status = -3 fio: io_u error on file /mnt/scratch/rac write offset=460849152, buflen=131072 == In __blockdev_direct_IO, ocfs2_dio_wr_get_block is called to add unwritten extents to a list. extents are also inser ocfs2_write_begin_nolock. Then another thread call fallocate to puch a hole at one of the unwritten extent. The ext ocfs2_remove_extent(). At end io worker thread, ocfs2_search_extent_list found there is no such extent at the cpos inode unlock ocfs2_fallocate __ocfs2_change_file_space inode lock lock ip_alloc_sem ocfs2_remove_inode_range ocfs2_remove_extent ^---remove the extent at cpos 78723 ... unlock ip_alloc_sem inode unlock ocfs2_dio_end_io ocfs2_mark_extent_written ocfs2_change_extent_flag ocfs2_search_extent_list ^---failed to find extent ... unlock i not compatible with racing with AIO+DIO, so fix it by adding to wait for all dio before fallocate/punch_hole like e
CVE-2024-40958	In the Linux kernel, the following vulnerability has been resolved: netns: Make get_net_ns() handle zero refcount n addition on 0; use-after-free. WARNING: CPU: 3 PID: 7890 at lib/refcount.c:25 refcount_warn_saturate+0xdf/0x1 7890 Comm: tun Not tainted 6.10.0-rc3-00100-gcaa4f9578aba-dirty #310 Hardware name: QEMU Standard PC (i4 04/01/2014 RIP: 0010:refcount_warn_saturate+0xdf/0x1d0 Code: 41 49 04 31 ff 89 de e8 9f 1e cd fe 84 db 75 9c e 90 48 c7 c7 b8 8e 25 86 e8 d2 05 b5 fe 90 <0f> 0b 90 90 e9 79 ff ff ff e8 53 26 cd fe 0f b6 1 RSP: 0018:ffff888106 0000000000000000 RBX: 0000000000000000 RCX: ffffffff811c72ac RDX: ffff8881026a2140 RSI: ffffffff811c72 ffff8881067b7db0 R08: 0000000000000000 R09: 205b5d3730353139 R10: 0000000000000000 R11: 205d303938 ffff8881086500c4 R14: ffff8881086500b0 R15: ffff888108650040 FS: 00007f5b2961a4c0(0000) GS:ffff88823bd0 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 000055d7ed36fd18 CR3: 00000001482f6000 CR4: 0000 DR1: 0000000000000000 DR2: 0000000000000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000 show_regs+0xa3/0xc0 ? __warn+0xa5/0x1c0 ? refcount_warn_saturate+0xdf/0x1d0 ? report_bug+0x1fc/0x2d0 ? r handle_bug+0xa1/0x110 ? exc_invalid_op+0x3c/0xb0 ? asm_exc_invalid_op+0x1f/0x30 ? __warn_printk+0xcc/0x refcount_warn_saturate+0xdf/0x1d0 get_net_ns+0xa4/0xc0 ? __pfx_get_net_ns+0x10/0x10 open_related_ns+0x5a __sanitizer_cov_trace_switch+0x58/0xa0 ? __sanitizer_cov_trace_const_cmp2+0x1c/0x30 ? __pfx_tun_chr_ioctl+ __x64_sys_ioctl+0x11b/0x160 x64_sys_call+0x1211/0x20d0 do_syscall_64+0x9e/0x1d0 entry_SYSCALL_64_af 0033:0x7f5b28f165d7 Code: b3 66 90 48 8b 05 b1 48 2d 00 64 c7 00 26 00 00 00 48 c7 c0 ff ff ff ff c3 66 2e 0f 1f <48> 3d 01 f0 ff ff 73 01 c3 48 8b 0d 81 48 2d 00 8 RSP: 002b:00007ffc2b59c5e8 EFLAGS: 00000246 ORIG_RA RBX: 0000000000000000 RCX: 00007f5b28f165d7 RDX: 0000000000000000 RSI: 00000000000054e3 RDI: 000 R08: 00007f5b291ed8c0 R09: 00007f5b2961a4c0 R10: 0000000029690010 R11: 0000000000000246 R12: 00000 R14: 0000000000000000 R15: 0000000000000000 </TASK> Kernel panic - not syncing: kernel: panic_on_warn s tun_set_iff() //dev is tun0 tun->dev = dev //ip link set tun0 netns ns1 put_net() //ref is 0 __tun_chr_ioctl() //TUNGE open_related_ns(&net->ns, get_net_ns); //ns1 get_net_ns() get_net() //addition on 0 Use maybe_get_net() in get_n

CVE-2024-40959	In the Linux kernel, the following vulnerability has been resolved: xfrm6: check ip6_dst_idev() return value in xfr NULL, xfrm6_get_saddr() must act accordingly. syzbot reported: Oops: general protection fault, probably for non- 0000 [#1] PREEMPT SMP KASAN PTI KASAN: null-ptr-deref in range [0x0000000000000000-0x000000000000 u8:1 Not tainted 6.10.0-rc2-syzkaller-00383-gb8481381d4e2 #0 Hardware name: Google Google Compute Engine 04/02/2024 Workqueue: wg-kex-wg1 wg_packet_handshake_send_worker RIP: 0010:xfrm6_get_saddr+0x93/0x1 89 fa 48 c1 ea 03 80 3c 02 00 0f 85 97 00 00 00 4c 8b ab d8 00 00 00 48 b8 00 00 00 00 00 fc ff df 4c 89 ea 48 c1 8b 6d 00 e8 ca 13 47 01 48 b8 00 RSP: 0018:ffffc90000117378 EFLAGS: 00010246 RAX: dfffffc0000000000 RBX RDX: 0000000000000000 RSI: ffffffff89a0d6e9 RDI: ffff88807b079e98 RBP: ffff88807ad73248 R08: 000000000 ffff88807b079dc0 R11: 0000000000000007 R12: ffffc90000117480 R13: 0000000000000000 R14: 000000000000 FS: 0000000000000000(0000) GS:ffff8880b9300000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 00 00007f4586d00440 CR3: 0000000079042000 CR4: 00000000003506f0 DR0: 0000000000000000 DR1: 00000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 Call Trace: <TASK> xfrm_get_saddr xfrm_tmpl_resolve_one net/xfrm/xfrm_policy.c:2481 [inline] xfrm_tmpl_resolve+0xa26/0xf10 net/xfrm/xfrm_pol +0x140/0x2570 net/xfrm/xfrm_policy.c:2835 xfrm_bundle_lookup net/xfrm/xfrm_policy.c:3070 [inline] xfrm_loo xfrm_policy.c:3201 xfrm_lookup net/xfrm/xfrm_policy.c:3298 [inline] xfrm_lookup_route+0x3b/0x200 net/xfrm/x +0x15c/0x1d0 net/ipv6/ip6_output.c:1256 send6+0x611/0xd20 drivers/net/wireguard/socket.c:139 wg_socket_send wireguard/socket.c:178 wg_socket_send_buffer_to_peer+0x12b/0x190 drivers/net/wireguard/socket.c:200 wg_pac drivers/net/wireguard/send.c:40 wg_packet_handshake_send_worker+0x1c/0x30 drivers/net/wireguard/send.c:51 p workqueue.c:3231 process_scheduled_works kernel/workqueue.c:3312 [inline] worker_thread+0x6c8/0xf70 kerne kernel/kthread.c:389 ret_from_fork+0x45/0x80 arch/x86/kernel/process.c:147 ret_from_fork_asm+0x1a/0x30 arch
CVE-2024-40960	In the Linux kernel, the following vulnerability has been resolved: ipv6: prevent possible NULL dereference in rt6_ in rt6_probe() [1] Bail out if __in6_dev_get() returns NULL. [1] Oops: general protection fault, probably for non-c 0000 [#1] PREEMPT SMP KASAN PTI KASAN: null-ptr-deref in range [0x0000000000000658-0x000000000000 executor.0 Not tainted 6.10.0-rc2-syzkaller-00383-gb8481381d4e2 #0 Hardware name: Google Google Compute E Google 04/02/2024 RIP: 0010:rt6_probe net/ipv6/route.c:656 [inline] RIP: 0010:find_match+0x8c4/0xf50 net/ipv6 38 ff ff ff 48 c7 45 b0 00 00 00 00 48 8d b8 5c 06 00 00 48 b8 00 00 00 00 00 fc ff df 48 89 fa 48 c1 ea 03 <0f> b6 d0 7c 08 84 d2 0f 85 19 RSP: 0018:ffffc900034af070 EFLAGS: 00010203 RAX: dfffffc0000000000 RBX: 000000 RDX: 00000000000000cb RSI: ffffffff8990d0cd RDI: 000000000000065c RBP: ffffc900034af150 R08: 00000000 R10: 0000000000000001 R11: 0000000000000002 R12: 000000000000000a R13: 1ffff92000695e18 R14: ffff888 FS: 00007f4844a5a6c0(0000) GS:ffff8880b9300000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 000 0000001b31b27000 CR3: 000000002d42c000 CR4: 00000000003506f0 DR0: 0000000000000000 DR1: 00000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 Call Trace: <TASK> rt6_nh_find_ma nexthop_for_each_fib6_nh+0x26d/0x4a0 net/ipv4/nexthop.c:1496 __find_rr_leaf+0x6e7/0xe00 net/ipv6/route.c:82 [inline] rt6_select net/ipv6/route.c:897 [inline] fib6_table_lookup+0x57e/0xa30 net/ipv6/route.c:2195 ip6_pol_rou pol_lookup_func include/net/ip6_fib.h:616 [inline] fib6_rule_lookup+0x386/0x720 net/ipv6/fib6_rules.c:121 ip6_r route.c:2639 [inline] ip6_route_output_flags+0x1d0/0x640 net/ipv6/route.c:2651 ip6_dst_lookup_tail.constprop.0+ ip6_dst_lookup_flow+0x99/0x1d0 net/ipv6/ip6_output.c:1250 rawv6_sendmsg+0xdab/0x4340 net/ipv6/raw.c:898 af_inet.c:853 sock_sendmsg_nosec net/socket.c:730 [inline] __sock_sendmsg net/socket.c:745 [inline] sock_write new_sync_write fs/read_write.c:497 [inline] vfs_write+0x6b6/0x1140 fs/read_write.c:590 ksys_write+0x1f8/0x26 x86/entry/common.c:52 [inline] do_syscall_64+0xcd/0x250 arch/x86/entry/common.c:83 entry_SYSCALL_64_af
CVE-2024-40961	In the Linux kernel, the following vulnerability has been resolved: ipv6: prevent possible NULL deref in fib6_nh_i can return NULL. fib6_nh_init() ip6_validate_gw(&idev) ip6_route_check_nh(idev) *idev = in6_dev_get(dev); protection fault, probably for non-canonical address 0xdffffc00000000bc: 0000 [#1] PREEMPT SMP KASAN PTI [0x00000000000005e0-0x00000000000005e7] CPU: 0 PID: 11237 Comm: syz-executor.3 Not tainted 6.10.0-rc2-s name: Google Google Compute Engine/Google Compute Engine, BIOS Google 06/07/2024 RIP: 0010:fib6_nh_ini Code: 00 00 fc ff df 4c 8b 64 24 58 48 8b 44 24 28 4c 8b 74 24 30 48 89 c1 48 89 44 24 28 48 8d 98 e0 05 00 00 4 c0 0f 85 b3 17 00 00 8b 1b 31 ff 89 de e8 b8 8b RSP: 0018:ffffc900032775a0 EFLAGS: 00010202 RAX: 0000000 RCX: 0000000000000000 RDX: 0000000000000010 RSI: ffffc90003277a54 RDI: ffff88802b3a08d8 RBP: ffffc90 R09: 0000000000000000 R10: 00000000000002fc R11: 0000000000000000 R12: ffff88802b3a08b8 R13: 1ffff920 R15: dfffffc0000000000 FS: 00007f940feb06c0(0000) GS:ffff8880b9400000(0000) knlGS:0000000000000000 CS 0000000080050033 CR2: 0000000000000000 CR3: 00000000245e8000 CR4: 00000000003506f0 DR0: 0000000 0000000000000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 Call Trace: <TA net/ipv6/route.c:3809 ip6_route_add+0x28/0x160 net/ipv6/route.c:3853 ipv6_route_ioctl+0x588/0x870 net/ipv6/ro ipv6/af_inet6.c:579 sock_do_ioctl+0x158/0x460 net/socket.c:1222 sock_ioctl+0x629/0x8e0 net/socket.c:1341 vfs_ ioctl.c:907 [inline] __se_sys_ioctl+0xfc/0x170 fs/ioctl.c:893 do_syscall_x64 arch/x86/entry/common.c:52 [inline] common.c:83 entry_SYSCALL_64_after_hwframe+0x77/0x7f RIP: 0033:0x7f940f07cea9
CVE-2024-40963	In the Linux kernel, the following vulnerability has been resolved: mips: bmips: BCM6358: make sure CBR is corr have CBR address set to 0 causing kernel panic when arch_sync_dma_for_cpu_all is called. This was notice in situ and BMIPS_GET_CBR() returns 0 instead of a valid address and !!(read_c0_brcm_cmt_local() & (1 << 31)); not f should be disabled or not are not enough hence lets check if CBR is a valid address or not.
CVE-2024-40968	In the Linux kernel, the following vulnerability has been resolved: MIPS: Octeon: Add PCIe link status check The interface is used to access the configuration space of the peripheral PCIe devices of the mips processor after the PC panic caused by "Data bus error". So it is necessary to add PCIe link status check for system protection. When the value of 0 to the configuration address can prevent read-write behavior to the configuration space of peripheral PC

CVE-2024-40974	In the Linux kernel, the following vulnerability has been resolved: powerpc/pseries: Enforce hcall result buffer vali and related functions expect callers to provide valid result buffers of certain minimum size. Currently this is comm code and the compiler has no idea. For example, if I write a bug like this: long retbuf[PLPAR_HCALL_BUFSIZE] plpar_hcall9(H_ALLOCATE_VAS_WINDOW, retbuf, ...); This compiles with no diagnostics emitted, but likely r plpar_hcall9() stores results past the end of the array. (To be clear this is a contrived example and I have not found error less likely, we can use explicitly-sized array parameters instead of pointers in the declarations for the hcall AI the code above now provokes a diagnostic like this: error: array argument is too small; is of size 32, callee requires | plpar_hcall9(H_ALLOCATE_VAS_WINDOW, retbuf, | ^ ~~~~~~ [1] Enabled for LLVM builds but not GCC fo disable '-Warray-bounds' for gcc-13 too") and related changes.
CVE-2024-40978	In the Linux kernel, the following vulnerability has been resolved: scsi: qedi: Fix crash while reading debugfs attrib qedi_dbg_do_not_recover_cmd_read() function invokes sprintf() directly on a __user pointer, which results into th small local stack buffer for sprintf() and then call simple_read_from_buffer(), which in turns make the copy_to_use page fault for address: 00007f4801111000 PGD 8000000864df6067 P4D 8000000864df6067 PUD 864df7067 PM [#1] PREEMPT SMP PTI Hardware name: HPE ProLiant DL380 Gen10/ProLiant DL380 Gen10, BIOS U30 06/15 +0xcd/0x130 RSP: 0018:ffffb7a18c3ffc40 EFLAGS: 00010202 RAX: 00007f4801111000 RBX: 00007f480111100 RDX: 000000000000000f RSI: ffffffffc0bfd7a0 RDI: 00007f4801111000 RBP: ffffffffc0bfd7a0 R08: 725f746f6e5 R10: ffffb7a18c3ffd08 R11: 0000000000000000 R12: 00007f4881110fff R13: 000000007fffffff R14: ffffb7a18c3f 00007f480118a740(0000) GS:ffff98e38af00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 C 00007f4801111000 CR3: 0000000864b8e001 CR4: 00000000007706e0 DR0: 0000000000000000 DR1: 00000000 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 PKRU: 55555554 Call Trace: <TASK> ? _ +0x183/0x510 ? exc_page_fault+0x69/0x150 ? asm_exc_page_fault+0x22/0x30 ? memcpy_orig+0xcd/0x130 vsnp qedi_dbg_do_not_recover_cmd_read+0x2f/0x50 [qedi 6bcfdeeecdea037da47069eca2ba717c84a77324] full_proxy folio_add_new_anon_rmap+0x44/0xa0 ? set_pte_at+0x15/0x30 ? do_pte_missing+0x426/0x7f0 ksys_read+0xa5/0 __count_memcg_events+0x46/0x90 ? count_memcg_event_mm+0x3d/0x60 ? handle_mm_fault+0x196/0x2f0 ? d exc_page_fault+0x69/0x150 entry_SYSCALL_64_after_hwframe+0x72/0xdc RIP: 0033:0x7f4800f20b4d
CVE-2024-40981	In the Linux kernel, the following vulnerability has been resolved: batman-adv: bypass empty buckets in batadv_pu pointing to soft lockups in batadv_purge_orig_ref() [1] Root cause is unknown, but we can avoid spending too muc reports. [1] watchdog: BUG: soft lockup - CPU#0 stuck for 27s! [kworker/u4:6:621] Modules linked in: irq event s at (6182793): [<ffff8000801dae10>] __local_bh_enable_ip+0x224/0x44c kernel/softirq.c:386 hardirqs last disable __el1_irq arch/arm64/kernel/entry-common.c:533 [inline] hardirqs last disabled at (6182794): [<ffff80008ad66a78 kernel/entry-common.c:551 softirqs last enabled at (6182792): [<ffff80008aab71c4>] spin_unlock_bh include/linu enabled at (6182792): [<ffff80008aab71c4>] batadv_purge_orig_ref+0x114c/0x1228 net/batman-adv/originator.c: [<ffff80008aab61dc>] spin_lock_bh include/linux/spinlock.h:356 [inline] softirqs last disabled at (6182790): [<fff +0x164/0x1228 net/batman-adv/originator.c:1271 CPU: 0 PID: 621 Comm: kworker/u4:6 Not tainted 6.8.0-rc7-syz Google Google Compute Engine/Google Compute Engine, BIOS Google 02/29/2024 Workqueue: bat_events batad +PAN -UAO -TCO -DIT -SSBS BTYPE=--) pc : should_resched arch/arm64/include/asm/preempt.h:79 [inline] pc kernel/softirq.c:388 lr : __local_bh_enable_ip+0x224/0x44c kernel/softirq.c:386 sp : ffff800099007970 x29: ffff80 x27: dfff800000000000 x26: ffff0000d2620008 x25: ffff0000c7e70de8 x24: 0000000000000001 x23: 1fffe00018e x21: ffff80008aab71c4 x20: ffff0001b40136c0 x19: ffff0000c72bbc08 x18: 1fffe0001a817bb0 x17: ffff800125414 0000000000000001 x14: 1fffe0001ee9d610 x13: 0000000000000000 x12: 0000000000000003 x11: 00000000000 0000000000000000 x8 : 00000000005e5789 x7 : ffff80008aab61dc x6 : 0000000000000000 x5 : 00000000000000 0000000000000000 x2 : 0000000000000006 x1 : 0000000000000080 x0 : ffff800125414000 Call trace: __daif_lo asm/irqflags.h:27 [inline] arch_local_irq_enable arch/arm64/include/asm/irqflags.h:49 [inline] __local_bh_enable_ __raw_spin_unlock_bh include/linux/spinlock_api_smp.h:167 [inline] _raw_spin_unlock_bh+0x3c/0x4c kernel/lo include/linux/spinlock.h:396 [inline] batadv_purge_orig_ref+0x114c/0x1228 net/batman-adv/originator.c:1287 bat adv/originator.c:1300 process_one_work+0x694/0x1204 kernel/workqueue.c:2633 process_scheduled_works kern +0x938/0xef4 kernel/workqueue.c:2787 kthread+0x288/0x310 kernel/kthread.c:388 ret_from_fork+0x10/0x20 arc from CPU 0 to CPUs 1: NMI backtrace for cpu 1 CPU: 1 PID: 0 Comm: swapper/1 Not tainted 6.8.0-rc7-syzkaller- Google Google Compute Engine/Google Compute Engine, BIOS Google 02/29/2024 pstate: 80400005 (Nzcv daif pc : arch_local_irq_enable+0x8/0xc arch/arm64/include/asm/irqflags.h:51 lr : default_idle_call+0xf8/0x128 kernel x29: ffff800093a17d30 x28: dfff800000000000 x27: 1ffff00012742fb4 x26: ffff80008ec9d000 x25: 000000000000 x23: 1ffff00011d93a74 x22: ffff80008ec9d3a0 x21: 0000000000000000 x20: ffff0000c19dbc00 x19: ffff8000802d ffff80008ec9d000 x16: ffff8000802d089c x15: 0000000000000001 ---truncated---
CVE-2024-40984	In the Linux kernel, the following vulnerability has been resolved: ACPICA: Revert "ACPICA: avoid Info: mappin Undo the modifications made in commit d410ee5109a1 ("ACPICA: avoid "Info: mapping multiple BARs. Your ke commit was to stop memory mappings for operation regions from overlapping page boundaries, as it can trigger wa However, it was found that when this situation arises, mapping continues until the boundary's end, but there is still the map, leading to a NULL pointer deference. For example, if a four-byte mapping request is made but only one b boundary's end, a four-byte read/write attempt is still made, resulting in a NULL pointer deference. Instead, map th does not mandate that it must be within the same page boundary. It is permissible for it to be mapped across differe
CVE-2024-40987	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: fix UBSAN warning in kv_dpm.c sumo_vid_mapping_entry.
CVE-2024-40988	In the Linux kernel, the following vulnerability has been resolved: drm/radeon: fix UBSAN warning in kv_dpm.c A sumo_vid_mapping_entry.

CVE-2024-40992	In the Linux kernel, the following vulnerability has been resolved: RDMA/rxe: Fix responder length checking for U specification: If a UD request packet is detected with an invalid length, the request shall be an invalid request and i The responder then waits for a new request packet. commit 689c5421bfe0 ("RDMA/rxe: Fix incorrect responder le check for UD QPs in function `copy_data`. But it introduces a regression issue for UD QPs. When the packet size i `copy_data` will return error code -EINVAL. Then `send_data_in` will return RESPST_ERR_MALFORMED_WC
CVE-2024-40994	In the Linux kernel, the following vulnerability has been resolved: ptp: fix integer overflow in max_vclocks_store (overflow. Use kcalloc() to do the allocation to prevent this.
CVE-2024-40995	In the Linux kernel, the following vulnerability has been resolved: net/sched: act_api: fix possible infinite loop in t tasks waiting on rtnl_lock [1] A reproducer is available in the syzbot bug. When a request to add multiple actions w will block forever on the first request. This holds rtnl_lock, and causes tasks to hang. Return -EAGAIN to prevent i behavior. [1] INFO: task kworker/1:0:5088 blocked for more than 143 seconds. Not tainted 6.9.0-rc4-syzkaller-001 sys/kernel/hung_task_timeout_secs" disables this message. task:kworker/1:0 state:D stack:23744 pid:5088 tgid:508 events_power_efficient reg_check_chans_work Call Trace: <TASK> context_switch kernel/sched/core.c:5409 [inl core.c:6746 __schedule_loop kernel/sched/core.c:6823 [inline] schedule+0xe7/0x350 kernel/sched/core.c:6838 sch sched/core.c:6895 __mutex_lock_common kernel/locking/mutex.c:684 [inline] __mutex_lock+0x5b8/0x9c0 kerne net/cfg80211.h:5953 [inline] reg_leave_invalid_chans net/wireless/reg.c:2466 [inline] reg_check_chans_work+0x1
CVE-2024-41003	In the Linux kernel, the following vulnerability has been resolved: bpf: Fix reg_set_min_max corruption of fake_re some changes to buzzer [0] and implementing a new fuzzing strategy guided by coverage, they noticed the followi (79) r6 = *(u64 *)(r0 +0) ; R0=map_value(ks=4,vs=8) R6_w=scalar() 14: (b7) r0 = 0 ; R0_w=0 15: (b4) w0 = -1 ; >>= 1 ; R0_w=0x7fffffff 17: (5c) w6 &= w0 ; R0_w=0x7fffffff R6_w=scalar(smin=smin32=0,smax=umax=umax 0x7fffffff)) 18: (44) w6 |= 2 ; R6_w=scalar(smin=umin=smin32=umin32=2,smax=umax=umax32=0x7fffffff,var_o if w6 != 0x7ffffffd goto pc+1 REG INVARIANTS VIOLATION (true_reg2): range bounds violation u64=[0x7ffff 0x7ffffffd] u32=[0x7fffffff, 0x7ffffffd] s32=[0x7fffffff, 0x7ffffffd] var_off=(0x7fffffff, 0x0) REG INVARIANTS range bounds violation u64=[0x7fffffff, 0x7ffffffd] s64=[0x7fffffff, 0x7ffffffd] u32=[0x7fffffff, 0x7ffffffd] s32=[0 var_off=(0x7fffffff, 0x0) REG INVARIANTS VIOLATION (false_reg2): const tnum out of sync with range bound s64=[0x8000000000000000, 0x7fffffffffffffff] u32=[0x0, 0xffffffff] s32=[0x80000000, 0x7fffffff] var_off=(0x7fff 20: (95) exit from 19 to 21: R0=0x7fffffff R6=scalar(smin=umin=smin32=umin32=2,smax=umax=smax32=umax3 0x7ffffffd)) R7=map_ptr(ks=4,vs=8) R9=ctx() R10=fp0 fp-24=map_ptr(ks=4,vs=8) fp-40=mmmmmmmm 21: R0= R6=scalar(smin=umin=smin32=umin32=2,smax=umax=smax32=umax32=0x7ffffffe,var_off=(0x2; 0x7ffffffd)) R7=map_ptr(ks=4,vs=8) R9=ctx() R10=fp0 fp-24=map_ptr(ks=4,vs=8) fp-40=mmmmmmmm 21: (14) w6 -= 2147 R6_w=scalar(smin=umin=umin32=2,smax=umax=0xffffffff,smin32=0x80000012,smax32=14,var_off=(0x2; 0xfff +1 ; R6_w=scalar(smin=umin=umin32=2,smax=umax=0xffffffff,smin32=0x80000012,smax32=13,var_off=(0x2; (24: R0=0x7fffffff R6_w=14 R7=map_ptr(ks=4,vs=8) R9=ctx() R10=fp0 fp-24=map_ptr(ks=4,vs=8) fp-40=mmmm R7=map_ptr(ks=4,vs=8) R9=ctx() R10=fp0 fp-24=map_ptr(ks=4,vs=8) fp-40=mmmmmmmm 24: (14) w6 -= 14 ; register invariant violation on line 19. After the binary-or in line 18, the verifier knows that bit 2 is set but knows n loaded from a map value, meaning, range is [2,0x7fffffff] with var_off=(0x2; 0x7ffffffd). When in line 19 the verif states in reg_set_min_max() into the registers of the true branch (true_reg1, true_reg2) and the registers of the false test is w6 != 0x7ffffffd, the src_reg is a known constant. Internally, the verifier creates a "fake" register initialized a passes it onto reg_set_min_max(). Now, for line 19, it is mathematically impossible to take the false branch of this impossible because the second bit of r6 will be set due to the prior or operation and the constant in the condition ha 1101). When the verifier first analyzes the false / fall-through branch, it will compute an intersection between the v because the verifier creates a "fake" register initialized to the value of the constant. The intersection result later refi [...] t = tnum_intersect(tnum_subreg(reg1->var_off), tnum_subreg(reg2->var_off)); reg1->var_o ---truncated---

CVE-2024-41004	In the Linux kernel, the following vulnerability has been resolved: tracing: Build event generation tests only as mo test modules add events and lock (get a reference) those event file reference in module init function, and unlock an is because those are designed for playing as modules. If we make those modules as built-in, those events are left lo This causes kprobe event self-test failure as below. [97.349708] ----------[cut here]---------- [97.353453] WA trace_kprobe.c:2133 kprobe_trace_self_tests_init+0x3f1/0x480 [97.357106] Modules linked in: [97.358488] CP 6.9.0-g699646734ab5-dirty #14 [97.361556] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 0010:kprobe_trace_self_tests_init+0x3f1/0x480 [97.365538] Code: a8 24 08 82 e9 ae fd ff ff 90 0f 0b 90 48 c7 c7 48 c7 c7 2d 61 06 82 e9 8e fd ff ff 90 <0f> 0b 90 48 c7 c7 33 0b 0c 82 89 c6 e8 6e 03 1f ff 41 ff c7 e9 90 [97.370 EFLAGS: 00010286 [97.371852] RAX: 0000000fffffff0 RBX: ffff888005919c00 RCX: 0000000000000000 [97 ffffffff8236a598 RDI: ffff888003f40a68 [97.375715] RBP: 0000000000000000 R08: 0000000000000001 R09: 00 ffffffff811c9ae5 R11: ffffffff8120c4e0 R12: 0000000000000000 [97.379591] R13: 0000000000000001 R14: 0000 [97.381536] FS: 0000000000000000(0000) GS:ffff88807dcc0000(0000) knlGS:0000000000000000 [97.383813] 0000000080050033 [97.385449] CR2: 0000000000000000 CR3: 0000000002244000 CR4: 00000000000006b0 [DR1: 0000000000000000 DR2: 0000000000000000 [97.389277] DR3: 0000000000000000 DR6: 00000000fffe0f Call Trace: [97.391967] <TASK> [97.392647] ? __warn+0xcc/0x180 [97.393640] ? kprobe_trace_self_tests_init +0xbd/0x150 [97.396234] ? handle_bug+0x3e/0x60 [97.397311] ? exc_invalid_op+0x1a/0x50 [97.398434] ? asm trace_kprobe_is_busy+0x20/0x20 [97.400904] ? tracing_reset_all_online_cpus+0x15/0x90 [97.402304] ? kprobe [97.403773] ? init_kprobe_trace+0x50/0x50 [97.404972] do_one_initcall+0x112/0x240 [97.406113] do_initcall +0x1a/0x1a0 [97.408401] do_initcalls+0x3f/0x70 [97.409452] kernel_init_freeable+0x16f/0x1e0 [97.410662] ? kernel_init+0x1a/0x1a0 [97.412788] ret_from_fork+0x39/0x50 [97.413817] ? rest_init+0x1f0/0x1f0 [97.414844 [97.416285] </TASK> [97.417134] irq event stamp: 13437323 [97.418376] hardirqs last enabled at (13437337): +0x11c/0x150 [97.421285] hardirqs last disabled at (13437370): [<ffffffff8110bbf1>] console_unlock+0x101/0x1 at (13437366): [<ffffffff8108e17f>] handle_softirqs+0x23f/0x2a0 [97.426450] softirqs last disabled at (13437393 +0x66/0xd0 [97.428850] ---[end trace 0000000000000000]--- And also, since we can not cleanup dynamic_even issues, build these tests only as modules.
CVE-2024-41006	In the Linux kernel, the following vulnerability has been resolved: netrom: Fix a memory leak in nr_heartbeat_exp in nr_create() [0]. Commit 409db27e3a2e ("netrom: Fix use-after-free of a listening socket.") added sock_hold() to where a) a socket has a SOCK_DESTROY flag or b) a listening socket has a SOCK_DEAD flag. But in the case "a is set, the file descriptor has already been closed and the nr_release() function has been called. So it makes no sens no one will call another nr_destroy_socket() and put it as in the case "b." nr_connect nr_establish_data_link nr_sta >state) case NR_STATE_3 nr->state = NR_STATE_2 sock_set_flag(sk, SOCK_DESTROY); nr_rx_frame nr_pro NR_STATE_2 nr_state2_machine() nr_disconnect() nr_sk(sk)->state = NR_STATE_0 sock_set_flag(sk, SOCK_D >state) case NR_STATE_0 if (sock_flag(sk, SOCK_DESTROY) || (sk->sk_state == TCP_LISTEN && sock_flag(nr_destroy_socket() To fix the memory leak, let's call sock_hold() only for a listening socket. Found by InfoTeCS (linuxtesting.org) with Syzkaller. [0]: https://syzkaller.appspot.com/bug?extid=d327a1f3b12e1e206c16
CVE-2024-41007	In the Linux kernel, the following vulnerability has been resolved: tcp: avoid too many retransmit packets If a TCP and the other peer retracted its window to zero, tcp_retransmit_timer() can retransmit a packet every two jiffies (2 TCP_USER_TIMEOUT has 'expired'. The fix is to make sure tcp_rtx_probe0_timed_out() takes icsk->icsk_user_t the socket would not timeout after icsk->icsk_user_timeout, but would use standard exponential backoff for the ret commit e89688e3e978 ("net: tcp: fix unexcepted socket die when snd_wnd is 0"), the issue would last 2 minutes in
CVE-2024-41009	In the Linux kernel, the following vulnerability has been resolved: bpf: Fix overrunning reservations in ringbuf The as a power-of-2 sized circular buffer, with two logical and ever-increasing counters: consumer_pos is the consume the consumer consumed the data, and producer_pos which is the producer counter denoting the amount of data rese is reserved, the producer that "owns" the record will successfully advance producer counter. In user space each tim advanced the consumer counter once it finished processing. Both counters are stored in separate pages so that from only and the consumer counter is read-write. One aspect that simplifies and thus speeds up the implementation of b data area is mapped twice contiguously back-to-back in the virtual memory, allowing to not take any special measu at the end of the circular buffer data area, because the next page after the last data page would be first data page aga completely contiguous in virtual memory. Each record has a struct bpf_ringbuf_hdr { u32 len; u32 pg_off; } heade is inaccessible to the BPF program. Helpers like bpf_ringbuf_reserve() return `(void *)hdr + BPF_RINGBUF_HD Jhong and Muhammad reported that it is however possible to make a second allocated memory chunk overlapping program is now able to edit first chunk's header. For example, consider the creation of a BPF_MAP_TYPE_RING consumer_pos is modified to 0x3000 /before/ a call to bpf_ringbuf_reserve() is made. This will allocate a chunk A, program is able to edit [0x8,0x3008]. Now, lets allocate a chunk B with size 0x3000. This will succeed because co pass the `new_prod_pos - cons_pos > rb->mask` check. Chunk B will be in range [0x3008,0x6010], and the BPF p Due to the ring buffer memory layout mentioned earlier, the ranges [0x0,0x4000] and [0x4000,0x8000] point to the B at [0x4000,0x4008] is chunk A's header. bpf_ringbuf_submit() / bpf_ringbuf_discard() use the header's pg_off to bpf_ringbuf_restore_from_rec(). Once chunk B modified chunk A's header, then bpf_ringbuf_commit() refers to th by calculating the oldest pending_pos and check whether the range from the oldest outstanding record to the newes that is the case, then reject the request. We've tested with the ring buffer benchmark in BPF selftests (./benchs/run_ while it seems a bit slower on some benchmarks, it is still not significantly enough to matter.
CVE-2024-41034	In the Linux kernel, the following vulnerability has been resolved: nilfs2: fix kernel bug on rename operation of br directory operation on broken directory on nilfs2, __block_write_begin_int() called to prepare block write may fail the folio/page size. This is because nilfs_dotdot(), which gets parent directory reference entry ("..") of the directory consistency enough, and may return location exceeding folio/page size for broken directories. Fix this issue by che in the first chunk of the directory in nilfs_dotdot().

CVE-2024-41035	In the Linux kernel, the following vulnerability has been resolved: USB: core: Fix duplicate endpoint bug by cleari identified a bug in usbcore (see the Closes: tag below) caused by our assumption that the reserved bits in an endpoi always be 0. As a result of the bug, the endpoint_is_duplicate() routine in config.c (and possibly other routines as v for distinct endpoints, even though they have the same direction and endpoint number. This can lead to confusion, descriptors with matching endpoint numbers and directions, where one was interrupt and the other was bulk). To fi bEndpointAddress when we parse the descriptor. (Note that both the USB-2.0 and USB-3.1 specs say these bits are to make a copy of the descriptor earlier in usb_parse_endpoint() and use the copy instead of the original when chec
CVE-2024-41041	In the Linux kernel, the following vulnerability has been resolved: udp: Set SOCK_RCU_FREE earlier in udp_lib_ [0] in udp_v4_early_demux(). In udp_v[46]_early_demux() and sk_lookup(), we do not touch the refcount of the lo skb->destructor, so we check SOCK_RCU_FREE to ensure that the sk is safe to access during the RCU grace perio flagged for a bound socket after being put into the hash table. Moreover, the SOCK_RCU_FREE check is done too and sk_lookup(), so there could be a small race window: CPU1 CPU2 ---- ---- udp_v4_early_demux() udp_lib_get_ __udp4_lib_demux_lookup() | |- DEBUG_NET_WARN_ON_ONCE(sk_is_refcounted(sk)); `- sock_set_flag(sk, S bug in TCP and fixed it in commit 871019b22d1b ("net: set SOCK_RCU_FREE before inserting socket into hashta [0]: WARNING: CPU: 0 PID: 11198 at net/ipv4/udp.c:2599 udp_v4_early_demux+0x481/0xb70 net/ipv4/udp.c:25 11198 Comm: syz-executor.1 Not tainted 6.9.0-g93bda33046e7 #13 Hardware name: QEMU Standard PC (i440FX gd239552ce722-prebuilt.qemu.org 04/01/2014 RIP: 0010:udp_v4_early_demux+0x481/0xb70 net/ipv4/udp.c:2599 e9 31 ff d3 e3 81 e3 bf ef ff ff 89 de e8 2c 74 15 fe 85 db 0f 85 02 06 00 00 e8 9f 7a 15 fe <0f> 0b e8 98 7a 15 fe 4 52 RSP: 0018:ffffc9000ce3fa58 EFLAGS: 00010293 RAX: 0000000000000000 RBX: 0000000000000000 RCX: RSI: ffffffff8318c2f1 RDI: 0000000000000001 RBP: ffff88805a2dd6e0 R08: 0000000000000001 R09: 0000000(R11: 0001ffffffffffff R12: ffff88805a2dd680 R13: 0000000000000007 R14: ffff88800923f900 R15: ffff888054560 GS:ffff88807dc00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 C 000000003de4b002 CR4: 0000000000770ef0 DR0: 0000000000000000 DR1: 0000000000000000 DR2: 0000000(00000000fffe0ff0 DR7: 0000000000000600 PKRU: 55555554 Call Trace: <TASK> ip_rcv_finish_core.constprop ip_rcv_finish+0xda/0x150 net/ipv4/ip_input.c:447 NF_HOOK include/linux/netfilter.h:314 [inline] NF_HOOK inc +0x16c/0x180 net/ipv4/ip_input.c:569 __netif_receive_skb_one_core+0xb3/0xe0 net/core/dev.c:5624 __netif_rece netif_receive_skb_internal net/core/dev.c:5824 [inline] netif_receive_skb+0x271/0x300 net/core/dev.c:5884 tun_rx tun_get_user+0x24db/0x2c50 drivers/net/tun.c:2002 tun_chr_write_iter+0x107/0x1a0 drivers/net/tun.c:2048 new_ vfs_write+0x76f/0x8d0 fs/read_write.c:590 ksys_write+0xbf/0x190 fs/read_write.c:643 __do_sys_write fs/read_w read_write.c:652 [inline] __x64_sys_write+0x41/0x50 fs/read_write.c:652 x64_sys_call+0xe66/0x1990 arch/x86/i do_syscall_x64 arch/x86/entry/common.c:52 [inline] do_syscall_64+0x4b/0x110 arch/x86/entry/common.c:83 entr RIP: 0033:0x7fc44a68bc1f Code: 89 54 24 18 48 89 74 24 10 89 7c 24 08 e8 e9 cf f5 ff 48 8b 54 24 18 48 8b 74 2 05 <48> 3d 00 f0 ff ff 77 31 44 89 c7 48 89 44 24 08 e8 3c d0 f5 ff 48 RSP: 002b:00007fc449126c90 EFLAGS: 00 RAX: ffffffffffffffda RBX: 00000000004bc050 RCX: 00007fc44a68bc1f R ---truncated---
CVE-2024-41044	In the Linux kernel, the following vulnerability has been resolved: ppp: reject claimed-as-LCP but actually malform assumes valid LCP packets (with code from 1 to 7 inclusive), add 'ppp_check_packet()' to ensure that LCP packet h bytes, and reject claimed-as-LCP but actually malformed data otherwise.
CVE-2024-41045	In the Linux kernel, the following vulnerability has been resolved: bpf: Defer work in bpf_timer_cancel_and_free C patch (two timer callbacks trying to cancel each other) can be invoked through bpf_map_update_elem as well, or m containing timers. Since this relies on hrtimer_cancel as well, it is prone to the same deadlock situation as the previ hrtimer_try_to_cancel to fix this problem, as the timer cannot be enqueued after async_cancel_and_free. Once asyn timer must be reinitialized before it can be armed again. The callback running in parallel trying to arm the timer wi waiting is sufficient (given kfree_rcu), and bpf_timer_cb will return HRTIMER_NORESTART, preventing the tim there exists a UAF scenario where the callback arms the timer before entering this function, such that if cancellatio this routine, or the target timer callback running concurrently). In such a case, if the timer expiration is significantly expiration happening before it will free the bpf_hrtimer state and along with it the struct hrtimer, that is enqueued. after async_cancel_and_free, and yet it cannot be done inline due to deadlock issues. We thus modify bpf_timer_ca workqueue, adding a work_struct alongside rcu_head (both used at _different_ points of time, so can share space). new state of affairs.
CVE-2024-41046	In the Linux kernel, the following vulnerability has been resolved: net: ethernet: lantiq_etop: fix double free in deta descriptor is never incremented which results in the same skb being released multiple times.
CVE-2024-41049	In the Linux kernel, the following vulnerability has been resolved: filelock: fix potential use-after-free in posix_loc warning in trace_posix_lock_inode(). The request pointer had been changed earlier to point to a lock entry that was the tracepoint could fire, another task raced in and freed that lock. Fix this by moving the tracepoint inside the spin happen.

CVE-2024-41097	In the Linux kernel, the following vulnerability has been resolved: usb: atm: cxacru: fix endpoint checking in cxacr an old issue [1] that occurs due to incomplete checking of present usb endpoints. As such, wrong endpoints types n in turn triggers a warning in usb_submit_urb(). Fix the issue by verifying that required endpoint types are present f account cmd endpoint type. Unfortunately, this patch has not been tested on real hardware. [1] Syzbot report: usb 1 WARNING: CPU: 0 PID: 8667 at drivers/usb/core/urb.c:502 usb_submit_urb+0xed2/0x18a0 drivers/usb/core/urb. Comm: kworker/0:4 Not tainted 5.14.0-rc4-syzkaller #0 Hardware name: Google Google Compute Engine/Google Workqueue: usb_hub_wq hub_event RIP: 0010:usb_submit_urb+0xed2/0x18a0 drivers/usb/core/urb.c:502 ... Call atm/cxacru.c:649 cxacru_card_status+0x22/0xd0 drivers/usb/atm/cxacru.c:760 cxacru_bind+0x7ac/0x11a0 drivers +0x321/0x1ae0 drivers/usb/atm/usbatm.c:1055 cxacru_usb_probe+0xdf/0x1e0 drivers/usb/atm/cxacru.c:1363 usb_ core/driver.c:396 call_driver_probe drivers/base/dd.c:517 [inline] really_probe+0x23c/0xcd0 drivers/base/dd.c:595 drivers/base/dd.c:747 driver_probe_device+0x4c/0x1a0 drivers/base/dd.c:777 __device_attach_driver+0x20b/0x2f +0x15f/0x1e0 drivers/base/bus.c:427 __device_attach+0x228/0x4a0 drivers/base/dd.c:965 bus_probe_device+0x1e +0xc2f/0x2180 drivers/base/core.c:3354 usb_set_configuration+0x113a/0x1910 drivers/usb/core/message.c:2170 u drivers/usb/core/generic.c:238 usb_probe_device+0xd9/0x2c0 drivers/usb/core/driver.c:293
CVE-2024-41671	Twisted is an event-based framework for internet applications, supporting Python 3.6+. The HTTP 1.0 and 1.1 serv pipelined HTTP requests out-of-order, possibly resulting in information disclosure. This vulnerability is fixed in 24
CVE-2024-41810	Twisted is an event-based framework for internet applications, supporting Python 3.6+. The `twisted.web.util.redir vulnerability. If application code allows an attacker to control the redirect URL this vulnerability may result in Ref redirect response HTML body. This vulnerability is fixed in 24.7.0rc1.
CVE-2024-41957	Vim is an open source command line text editor. Vim < v9.1.0647 has double free in src/alloc.c:616. When closing will be cleared and freed. However a bit later, the quickfix list belonging to that window will also be cleared and if data, Vim will try to free it again, resulting in a double-free/use-after-free access exception. Impact is low since the several non-default flags, but it may cause a crash of Vim. The issue has been fixed as of Vim patch v9.1.0647
CVE-2024-42070	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_tables: fully validate NFT_DATA_ store validation for NFT_DATA_VALUE is conditional, however, the datatype is always either NFT_DATA_VAL requires a new helper function to infer the register type from the set datatype so this conditional check can be remo leaked through the registers.
CVE-2024-42071	In the Linux kernel, the following vulnerability has been resolved: ionic: use dev_consume_skb_any outside of nap need to be careful about how we call napi_consume_skb(), specifically we need to call it with budget==0 to signal was found while running some configuration stress testing of traffic and a change queue config loop running, and t BUG: using smp_processor_id() in preemptible [00000000] code: ethtool/20545 [4371.402897] caller is napi_skb 25 PID: 20545 Comm: ethtool Kdump: loaded Tainted: G OE 6.10.0-rc3-netnext+ #8 [4371.403302] Hardware na DL360 Gen10, BIOS U32 01/23/2021 [4371.403460] Call Trace: [4371.403613] <TASK> [4371.403758] dump_ check_preemption_disabled+0xc1/0xe0 [4371.404051] napi_skb_cache_put+0x16/0x80 [4371.404199] ionic_tx_ ionic_tx_cq_service+0xc4/0x200 [ionic] [4371.404505] ionic_tx_flush+0x15/0x70 [ionic] [4371.404653] ? ionic_ [ionic] [4371.404805] ionic_txrx_deinit+0x71/0x190 [ionic] [4371.404956] ionic_reconfigure_queues+0x5f5/0xf ionic_set_ringparam+0x2e8/0x3e0 [ionic] [4371.405265] ethnl_set_rings+0x1f1/0x300 [4371.405418] ethnl_defa genl_family_rcv_msg_doit+0xff/0x130 [...] I found that ionic_tx_clean() calls napi_consume_skb() which calls nap is the note /* Zero budget indicate non-NAPI context called us, like netpoll */ and DEBUG_NET_WARN_ON_ON that we're doing it wrong. We can pass a context hint down through the calls to let ionic_tx_clean() know what we' correctly.
CVE-2024-42073	In the Linux kernel, the following vulnerability has been resolved: mlxsw: spectrum_buffers: Fix memory corrupti two shared buffer operations make use of the Shared Buffer Status Register (SBSR): # devlink sb occupancy snaps occupancy clearmax pci/0000:01:00.0 The register has two masks of 256 bits to denote on which ingress / egress p Spectrum-4 has more than 256 ports, so the register was extended by cited commit with a new 'port_page' field. Ho payload, the driver specifies the ports as absolute numbers and not relative to the first port of the port page, resultin specifying the ports relative to the first port of the port page. [1] BUG: KASAN: slab-use-after-free in mlxsw_sp_s of size 1 at addr ffff8881068cb00f by task devlink/1566 [...] Call Trace: <TASK> dump_stack_lvl+0xc6/0x120 pri +0xd7/0x110 mlxsw_sp_sb_occ_snapshot+0xb6d/0xbc0 mlxsw_devlink_sb_occ_snapshot+0x75/0xb0 devlink_nl genl_family_rcv_msg_doit+0x20c/0x300 genl_rcv_msg+0x567/0x800 netlink_rcv_skb+0x170/0x450 genl_rcv+0x netlink_sendmsg+0x8d4/0xdb0 __sys_sendto+0x49b/0x510 __x64_sys_sendto+0xe5/0x1c0 do_syscall_64+0xc1/0 +0x77/0x7f [...] Allocated by task 1: kasan_save_stack+0x33/0x60 kasan_save_track+0x14/0x30 __kasan_kmalloc do_check_common+0x2c51/0xc7e0 bpf_check+0x5107/0x9960 bpf_prog_load+0xf0e/0x2690 __sys_bpf+0x1a61 do_syscall_64+0xc1/0x1d0 entry_SYSCALL_64_after_hwframe+0x77/0x7f Freed by task 1: kasan_save_stack+0 kasan_save_free_info+0x3b/0x60 poison_slab_object+0x109/0x170 __kasan_slab_free+0x14/0x30 kfree+0xca/0x do_check_common+0x4828/0xc7e0 bpf_check+0x5107/0x9960 bpf_prog_load+0xf0e/0x2690 __sys_bpf+0x1a61 do_syscall_64+0xc1/0x1d0 entry_SYSCALL_64_after_hwframe+0x77/0x7f

CVE-2024-42076	In the Linux kernel, the following vulnerability has been resolved: net: can: j1939: Initialize unused data in j1939_ in raw_recvmsg() [1]. j1939_send_one() creates full frame including unused data, but it doesn't initialize it. This ca by initializing unused data. [1] BUG: KMSAN: kernel-infoleak in instrument_copy_to_user include/linux/instrume kernel-infoleak in copy_to_user_iter lib/iov_iter.c:24 [inline] BUG: KMSAN: kernel-infoleak in iterate_ubuf inclu KMSAN: kernel-infoleak in iterate_and_advance2 include/linux/iov_iter.h:245 [inline] BUG: KMSAN: kernel-inf linux/iov_iter.h:271 [inline] BUG: KMSAN: kernel-infoleak in _copy_to_iter+0x366/0x2520 lib/iov_iter.c:185 ins instrumented.h:114 [inline] copy_to_user_iter lib/iov_iter.c:24 [inline] iterate_ubuf include/linux/iov_iter.h:29 [inl iov_iter.h:245 [inline] iterate_and_advance include/linux/iov_iter.h:271 [inline] _copy_to_iter+0x366/0x2520 lib/i linux/uio.h:196 [inline] memcpy_to_msg include/linux/skbuff.h:4113 [inline] raw_recvmsg+0x2b8/0x9e0 net/can/ socket.c:1046 [inline] sock_recvmsg+0x2c4/0x340 net/socket.c:1068 ____sys_recvmsg+0x18a/0x620 net/socket.c net/socket.c:2845 do_recvmmsg+0x4fc/0xfd0 net/socket.c:2939 __sys_recvmmsg net/socket.c:3018 [inline] __do_ [inline] __se_sys_recvmmsg net/socket.c:3034 [inline] __x64_sys_recvmmsg+0x397/0x490 net/socket.c:3034 x64 include/generated/asm/syscalls_64.h:300 do_syscall_x64 arch/x86/entry/common.c:52 [inline] do_syscall_64+0xc entry_SYSCALL_64_after_hwframe+0x77/0x7f Uninit was created at: slab_post_alloc_hook mm/slub.c:3804 [inl [inline] kmem_cache_alloc_node+0x613/0xc50 mm/slub.c:3888 kmalloc_reserve+0x13d/0x4a0 net/core/skbuff.c: skbuff.c:668 alloc_skb include/linux/skbuff.h:1313 [inline] alloc_skb_with_frags+0xc8/0xbf0 net/core/skbuff.c:65 core/sock.c:2795 sock_alloc_send_skb include/net/sock.h:1842 [inline] j1939_sk_alloc_skb net/can/j1939/socket.c j1939/socket.c:1142 [inline] j1939_sk_sendmsg+0xc0a/0x2730 net/can/j1939/socket.c:1277 sock_sendmsg_nosec +0x30f/0x380 net/socket.c:745 ____sys_sendmsg+0x877/0xb60 net/socket.c:2584 ___sys_sendmsg+0x28d/0x3c0 socket.c:2667 [inline] __do_sys_sendmsg net/socket.c:2676 [inline] __se_sys_sendmsg net/socket.c:2674 [inline] net/socket.c:2674 x64_sys_call+0xc4b/0x3b50 arch/x86/include/generated/asm/syscalls_64.h:47 do_syscall_x64 a do_syscall_64+0xcf/0x1e0 arch/x86/entry/common.c:83 entry_SYSCALL_64_after_hwframe+0x77/0x7f Bytes 12 of size 16 starts at ffff888120969690 Data copied to user address 00000000200017c0 CPU: 1 PID: 5050 Comm: sy syzkaller-00031-g71b1543c83d6 #0 Hardware name: Google Google Compute Engine/Google Compute Engine, B
CVE-2024-42083	In the Linux kernel, the following vulnerability has been resolved: ionic: fix kernel panic due to multi-buffer handl handle multi-buffer packets properly for XDP_TX and XDP_REDIRECT. When a jumbo frame is received, the io all necessary pages in the rx descriptor. And if the action is either XDP_TX or XDP_REDIRECT, it should unmap NULL for all pages, not only the first page. But it doesn't for SG pages. So, SG pages unexpectedly will be reused. general protection fault, probably for non-canonical address 0x504f4e4dbebc64ff: 0000 [#1] PREEMPT SMP NOP tainted 6.10.0-rc3+ #25 RIP: 0010:xdp_return_frame+0x42/0x90 Code: 01 75 12 5b 4c 89 e6 5d 31 c9 41 5c 31 d2 49 81 ed 68 01 00 00 49 29 c5 49 01 fd <41> 80 7d0 RSP: 0018:ffff99d00122ce08 EFLAGS: 00010202 RAX: 000 RCX: 0000000000000001 RDX: 00000000670e1000 RSI: 000000011f90d000 RDI: 504f4e4d4c4b4a49 RBP: ffff R09: 0000000000000000 R10: 000000011f90d000 R11: 0000000000000000 R12: ffff8d325f904010 R13: 504f4e R15: ffff99d0039077c0 FS: 0000000000000000(0000) GS:ffff8d399f780000(0000) knlGS:0000000000000000 CS 0000000080050033 CR2: 00007f41f6c85e38 CR3: 000000037ac30000 CR4: 00000000007506f0 PKRU: 5555555 exc_general_protection+0x251/0x2f0 ? asm_exc_general_protection+0x22/0x30 ? xdp_return_frame+0x42/0x90 i 15881354510e6a9c655c59c54812b319ed2cd015] ionic_tx_cq_service+0xd3/0x210 [ionic 15881354510e6a9c655c +0x41/0x1b0 [ionic 15881354510e6a9c655c59c54812b319ed2cd015] __napi_poll.constprop.0+0x29/0x1b0 net_r +0xf4/0x320 irq_exit_rcu+0x78/0xa0 common_interrupt+0x77/0x90
CVE-2024-42084	In the Linux kernel, the following vulnerability has been resolved: ftruncate: pass a signed offset The old ftruncate(extension when called in compat mode on 64-bit architectures. As a result, passing a negative length accidentally s 2GiB and 4GiB. Changing the type of the compat syscall to the signed compat_off_t changes the behavior so it ins the truncate() syscall and the corresponding loff_t based variants are all correct already and do not suffer from this
CVE-2024-42085	In the Linux kernel, the following vulnerability has been resolved: usb: dwc3: core: remove lock of otg mode durin deadlock When config CONFIG_USB_DWC3_DUAL_ROLE is selected, and trigger system to enter suspend statu sys/power/state There will be a deadlock issue occurring. Detailed invoking path as below: dwc3_suspend_commo flags); <-- 1st dwc3_gadget_suspend(dwc); dwc3_gadget_soft_disconnect(dwc); spin_lock_irqsave(&dwc->lock, f commit c7ebd8149ee5 ("usb: dwc3: gadget: Fix NULL pointer dereference in dwc3_gadget_suspend") that remove >gadget_driver is NULL or not. It causes the following code is executed and deadlock occurs when trying to get th commit 5265397f9442("usb: dwc3: Remove DWC3 locking during gadget suspend/resume") that forgot to remove redundant lock of otg mode during gadget suspend/resume.
CVE-2024-42086	In the Linux kernel, the following vulnerability has been resolved: iio: chemical: bme680: Fix overflows in compen compensate functions of the driver that there could be overflows of variables due to bit shifting ops. These implicat they were mentioned in log message of Commit 1b3bd8592780 ("iio: chemical: Add support for Bosch BME680 s iio/20180728114028.3c1bbe81@archlinux/
CVE-2024-42087	In the Linux kernel, the following vulnerability has been resolved: drm/panel: ilitek-ili9881c: Fix warning with GP controls the reset GPIO using the non-sleeping gpiod_set_value() function. This complains loudly when the GPIO sleep, use gpiod_set_value_cansleep() to fix the issue.
CVE-2024-42089	In the Linux kernel, the following vulnerability has been resolved: ASoC: fsl-asoc-card: set priv->pdev before usin being used in fsl_asoc_card_audmux_init(). Move this assignment at the start of the probe function, so sub-functio fsl_asoc_card_audmux_init() dereferences priv->pdev to get access to the dev struct, used with dev_err macros. As NULL pointer dereference. Note that if priv->dev is dereferenced before assignment but never used, for example if won't crash probably due to compiler optimisations.

CVE-2024-42090	In the Linux kernel, the following vulnerability has been resolved: pinctrl: fix deadlock in create_pinctrl() when ha pinctrl_maps_mutex is acquired before calling add_setting(). If add_setting() returns -EPROBE_DEFER, create_pi pinctrl_free() attempts to acquire pinctrl_maps_mutex, which is already held by create_pinctrl(), leading to a poten by releasing pinctrl_maps_mutex before calling pinctrl_free(), preventing the deadlock. This bug was discovered a Security Testing (SAST) by Synopsys, Inc.
CVE-2024-42092	In the Linux kernel, the following vulnerability has been resolved: gpio: davinci: Validate the obtained number of I from Device Tree. In case of broken DT due to any error this value can be any. Without this value validation there access in davinci_gpio_probe(). Validate the obtained nirq value so that it won't exceed the maximum number of II Center (linuxtesting.org) with SVACE.
CVE-2024-42093	In the Linux kernel, the following vulnerability has been resolved: net/dpaa2: Avoid explicit cpumask var allocatio CONFIG_CPUMASK_OFFSTACK=y kernel, explicit allocation of cpumask variable on stack is not recommende overflow. Instead, kernel code should always use *cpumask_var API(s) to allocate cpumask var in config-neutral w CONFIG_CPUMASK_OFFSTACK. Use *cpumask_var API(s) to address it.
CVE-2024-42094	In the Linux kernel, the following vulnerability has been resolved: net/iucv: Avoid explicit cpumask var allocation CONFIG_CPUMASK_OFFSTACK=y kernel, explicit allocation of cpumask variable on stack is not recommende overflow. Instead, kernel code should always use *cpumask_var API(s) to allocate cpumask var in config-neutral w CONFIG_CPUMASK_OFFSTACK. Use *cpumask_var API(s) to address it.
CVE-2024-42094	In the Linux kernel, the following vulnerability has been resolved: net/iucv: Avoid explicit cpumask var allocation CONFIG_CPUMASK_OFFSTACK=y kernel, explicit allocation of cpumask variable on stack is not recommende overflow. Instead, kernel code should always use *cpumask_var API(s) to allocate cpumask var in config-neutral w CONFIG_CPUMASK_OFFSTACK. Use *cpumask_var API(s) to address it.
CVE-2024-42096	In the Linux kernel, the following vulnerability has been resolved: x86: stop playing stack games in profile_pc() Th timer-based profiling, which isn't really all that relevant any more to begin with, but it also ends up making assump necessarily valid. Basically, the code tries to account the time spent in spinlocks to the caller rather than the spinloc not worth the code complexity or the KASAN warnings when no serious profiling is done using timers anyway the stack layout that is only true in the simplest of cases. We've lost the comment at some point (I think when the 32-bi to say: Assume the lock function has either no stack frame or a copy of eflags from PUSHF. which explains why it off the stack pointer and then takes a minimal look at the values to just check if they might be eflags or the return p unlike kernel addresses but that basic stack layout assumption assumes that there isn't any lock debugging etc going a stack frame. It causes KASAN unhappiness reported for years by syzkaller [1] and others [2]. With no real practi the code. Just for historical interest, here's some background commits relating to this code from 2006: 0cb91a2293 during profiling for !FP kernels") 31679f38d886 ("Simplify profile_pc on x86-64") and a code unification from 20 profile_pc") but the basics of this thing actually goes back to before the git tree.
CVE-2024-42102	In the Linux kernel, the following vulnerability has been resolved: Revert "mm/writeback: fix possible divide-by-z "mm: Avoid possible overflows in dirty throttling". Dirty throttling logic assumes dirty limits in page units fit into true (see patch 2/2 for more details). This patch (of 2): This reverts commit 9319b647902cbd5cc884ac08a8a6d54c ways. Firstly, the removed (u64) cast from the multiplication will introduce a multiplication overflow on 32-bit arc is actually common - the default settings with 4GB of RAM will trigger this). Secondly, the div64_u64() is unneces div64_ul() in case we want to be safe & cheap. Thirdly, if dirty thresholds are larger than 1<<32 pages, then dirty b spectacular ways anyway so trying to fix one possible overflow is just moot.
CVE-2024-42104	In the Linux kernel, the following vulnerability has been resolved: nilfs2: add missing check for inode numbers on mounting and unmounting a specific pattern of corrupted nilfs2 filesystem images causes a use-after-free of metada lru_add_fn(). As Jan Kara pointed out, this is because the link count of a metadata file gets corrupted to 0, and nilfs tries to delete that inode (ifile inode in this case). The inconsistency occurs because directories containing the inode not be visible in the namespace are read without checking. Fix this issue by treating the inode numbers of these inte when reading directory folios/pages. Also thanks to Hillf Danton and Matthew Wilcox for their initial mm-layer an
CVE-2024-42105	In the Linux kernel, the following vulnerability has been resolved: nilfs2: fix inode number range checks Patch seri to reserved inodes". This series fixes one use-after-free issue reported by syzbot, caused by nilfs2's internal inode b corrupted filesystem, and a couple of flaws that cause problems if the starting number of non-reserved inodes writt (or corruptly) changed from its default value. This patch (of 3): In the current implementation of nilfs2, "nilfs->ns_ reserved inode number, is read from the superblock, but its lower limit is not checked. As a result, if a number that of reserved inodes such as the root directory or metadata files is set in the super block parameter, the inode number NILFS_VALID_INODE) will not function properly. In addition, these test macros use left bit-shift calculations usi via the BIT macro, but the result of a shift calculation that exceeds the bit width of an integer is undefined in the C a large value other than the default value NILFS_USER_INO (=11), the macros may potentially malfunction deper by checking the lower bound of "nilfs->ns_first_ino" and by preventing bit shifts equal to or greater than the NILF test macros. Also, change the type of "ns_first_ino" from signed integer to unsigned integer to avoid the need for ty bound check introduced this time.

CVE-2024-42109	In the Linux kernel, the following vulnerability has been resolved: netfilter: nf_tables: unconditionally flush pendir KASAN: slab-uaf in nft_ctx_update include/net/netfilter/nf_tables.h:1831 KASAN: slab-uaf in nft_commit_release slab-uaf int nf_tables_trans_destroy_work+0x152b/0x1750 net/netfilter/nf_tables_api.c:9597 Read of size 2 at addr [..] Workqueue: events nf_tables_trans_destroy_work Call Trace: nft_ctx_update include/net/netfilter/nf_tables.h:1 nf_tables_api.c:9530 [inline] nf_tables_trans_destroy_work+0x152b/0x1750 net/netfilter/nf_tables_api.c:9597 Pro flush, but its possible that the table-to-be-removed is still referenced by transactions being processed by the worker could make the flush_work depend on whether we found a table to delete in nf-next to avoid the flush for most case nf-next, with commit e169285f8c56 ("netfilter: nf_tables: do not store nft_ctx in transaction objects"), with this cor of table->family which is whats triggering the above splat.
CVE-2024-42115	In the Linux kernel, the following vulnerability has been resolved: jffs2: Fix potential illegal address access in jffs2 jffs2 file system,the following abnormal printouts were found: [2430.649000] Unable to handle kernel paging requ [2430.649622] Mem abort info: [2430.649829] ESR = 0x96000004 [2430.650115] EC = 0x25: DABT (current E FnV = 0 [2430.650795] EA = 0, S1PTW = 0 [2430.651032] FSC = 0x04: level 0 translation fault [2430.651446] = 0x00000004 [2430.652001] CM = 0, WnR = 0 [2430.652558] [0069696969696948] address between user and k error: Oops: 96000004 [#1] PREEMPT SMP [2430.654512] CPU: 2 PID: 20919 Comm: cat Not tainted 5.15.25-g name: linux,dummy-virt (DT) [2430.655517] pstate: 20000005 (nzCv daif -PAN -UAO -TCO -DIT -SSBS BTYP [2430.656630] lr : jffs2_free_inode+0x24/0x48 [2430.657051] sp : ffff800009eebd10 [2430.657355] x29: ffff800 0000000000000000 [2430.658327] x26: ffff000038f09d80 x25: 0080000000000000 x24: ffff800009d38000 [243 ffff000038f09d80 x21: ffff8000084f0d14 [2430.659434] x20: ffff0000bf9a6ac0 x19: 0169696969696940 x18: 000 ffff8000b6506000 x16: ffff800009eec000 x15: 0000000000004000 [2430.660637] x14: 0000000000000000 x13: [2430.661345] x11: 0040000800000000 x10: 0000000000000001 x9 : ffff8000084f0d14 [2430.662025] x8 : ffff0 x6 : 0000000003470302 [2430.662695] x5 : ffff00002e41dcc0 x4 : ffff0000bf9aa3b0 x3 : 0000000003470342 [24 x1 : ffff8000084f0d14 x0 : fffffc0000000000 [2430.664217] Call trace: [2430.664528] kfree+0x78/0x348 [2430. [2430.665233] i_callback+0x24/0x50 [2430.665528] rcu_do_batch+0x1ac/0x448 [2430.665892] rcu_core+0x28 +0x18/0x28 [2430.666473] __do_softirq+0x138/0x3cc [2430.666781] irq_exit+0xf0/0x110 [2430.667065] hand gic_handle_irq+0xac/0xe8 [2430.667739] call_on_irq_stack+0x28/0x54 The parameter passed to kfree was 5a5a5 the jffs_inode_info structure. It was found that all variables in the jffs_inode_info structure were 5a5a5a5a, except these variables are not initialized because they were set to 5a5a5a5a during memory testing, which is meant to dete is initialized in the function jffs2_i_init_once, while other members are initialized in the function jffs2_init_inode_ is called after iget_locked, but in the iget_locked function, the destroy_inode process is triggered, which releases th member of the inode is not initialized.In concurrent high pressure scenarios, iget_locked may enter the destroy_ino destroy_inode functionality of jffs2 only releases the target, the fix method is to set target to NULL in jffs2_i_init_
CVE-2024-42119	In the Linux kernel, the following vulnerability has been resolved: drm/amd/display: Skip finding free audio for un ENGINE_ID_UNKNOWN = -1 and can not be used as an array index. Plus, it also means it is uninitialized and do return NULL. This fixes 2 OVERRUN issues reported by Coverity.
CVE-2024-42124	In the Linux kernel, the following vulnerability has been resolved: scsi: qedf: Make qedf_execute_tmf() non-preem preemptible code in qedf_execute_tmf90. This results in BUG_ON() when running an RT kernel. [659.343280] B [00000000] code: sg_reset/3646 [659.343282] caller is qedf_execute_tmf+0x8b/0x360 [qedf]
CVE-2024-42127	In the Linux kernel, the following vulnerability has been resolved: drm/lima: fix shared irq handling on driver remo interrupt handlers must be prepared to be called at any time. At driver removal time, the clocks are disabled early a very end of the remove process due to the devm usage. This is potentially a bug as the interrupts access device regi crash can be triggered by removing the driver in a kernel with CONFIG_DEBUG_SHIRQ enabled. This patch free callback so that the handlers are already unregistered by the time we fully disable clocks.
CVE-2024-42133	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: Ignore too large handle values in BI necessary to filter out cases where the handle value is belonging to ida id range, otherwise ida will be erroneously r
CVE-2024-42138	In the Linux kernel, the following vulnerability has been resolved: mlxsw: core_linecards: Fix double memory deal case of invalid INI file mlxsw_linecard_types_init() deallocates memory but doesn't reset pointer to NULL and retu mlxsw_linecard_types_init() call, mlxsw_linecards_init() calls mlxsw_linecard_types_fini() which performs memo NULL. Found by Linux Verification Center (linuxtesting.org) with SVACE.
CVE-2024-42153	In the Linux kernel, the following vulnerability has been resolved: i2c: pnx: Fix potential deadlock warning from d del_timer_sync() is called in an interrupt context it throws a warning because of potential deadlock. The timer is us after a timeout so replacing the call with wait_for_completion_timeout() allows to remove the problematic timer ar
CVE-2024-42154	In the Linux kernel, the following vulnerability has been resolved: tcp_metrics: validate source addr length I don't s TCP_METRICS_ATTR_SADDR_IPV4 is at least 4 bytes long, and the policy doesn't have an entry for this attribu manually validated).
CVE-2024-42157	In the Linux kernel, the following vulnerability has been resolved: s390/pkey: Wipe sensitive data on failure Wipe copy_to_user() fails.
CVE-2024-42160	In the Linux kernel, the following vulnerability has been resolved: f2fs: check validation of fault attrs in f2fs_build fault attrs in parse_options(), let's fix to add check condition in f2fs_build_fault_attr(). - Use f2fs_build_fault_attr()
CVE-2024-42162	In the Linux kernel, the following vulnerability has been resolved: gve: Account for stopped queues when reading I NIC might send us stats for a subset of queues. Without this change, gve_get_ethtool_stats might make an invalid a

CVE-2024-42224	In the Linux kernel, the following vulnerability has been resolved: net: dsa: mv88e6xxx: Correct check for empty l mv88e6xxx: Support multiple MDIO busses") mv88e6xxx_default_mdio_bus() has checked that the return value o to be intended to guard against the list chip->mdios being empty. However, it is not the correct check as the implem return NULL for empty lists. Instead, use list_first_entry_or_null() which does return NULL if the list is empty. Fl
CVE-2024-42228	In the Linux kernel, the following vulnerability has been resolved: drm/amdgpu: Using uninitialized value *size wh the size before calling amdgpu_vce_cs_reloc, such as case 0x03000001. V2: To really improve the handling we wo 0xffffffff.(Christian)
CVE-2024-42230	In the Linux kernel, the following vulnerability has been resolved: powerpc/pseries: Fix scv instruction crash with (reloc_on_exc), required for scv instruction support, before other CPUs have been shut down. This means they can disabled, which causes an interrupt at an unexpected entry location that crashes the kernel. Change the kexec seque have been brought down. As a refresher, the real-mode scv interrupt vector is 0x17000, and the fixed-location head implementing such high addresses so it was just decided not to support that interrupt at all.
CVE-2024-42236	In the Linux kernel, the following vulnerability has been resolved: usb: gadget: configfs: Prevent OOB read/write i string 's' could trivially have the length zero. Left unchecked this will firstly result in an OOB read in the form `if (s write in the form `str[0 - 1] = '\0'`. There is already a validating check to catch strings that are too long. Let's supply are too short.
CVE-2024-42239	In the Linux kernel, the following vulnerability has been resolved: bpf: Fail bpf_timer_cancel when callback is bei timer2 cb bpf_timer_cancel(timer2); bpf_timer_cancel(timer1); Both bpf_timer_cancel calls would wait for the oth a lockup. Add an atomic_t count named 'cancelling' in bpf_hrtimer. This keeps track of all in-flight cancellation re cancelling a BPF timer, we must check if we have outstanding cancellation requests, and if so, we must fail the ope cancellation is synchronous and waits for the callback to finish executing. This implies that we can enter a deadloc callbacks executing in parallel and attempting to cancel one another. Note that we avoid incrementing the cancellin cancelled) if bpf_timer_cancel is not invoked from a callback, to avoid spurious errors. The whole point of detectin is to not enter a busy wait loop (which may or may not lead to a lockup). This does not apply in case the caller is in continue to cancel as it sees fit without running into errors. Background on prior attempts: Earlier versions of this p following pattern under timer->lock to publish cancellation status. lock(t->lock); t->cancelling = true; mb(); if (cur >lock); hrtimer_cancel(t->timer); t->cancelling = false; The store outside the critical section could overwrite a para to ensure the parallely executing callback observes its cancellation status. It would be necessary to clear this cancel of serialization introduced races. Another option was explored where bpf_timer_start would clear the bit when (re)s would ensure serialized access to the cancelling bit, but may allow it to be cleared before in-flight hrtimer_cancel h occur again. Thus, we choose an atomic counter to keep track of all outstanding cancellation requests and use it to cancel each other while executing in parallel.
CVE-2024-42241	In the Linux kernel, the following vulnerability has been resolved: mm/shmem: disable PMD-sized page cache if n PMD-sized page cache can't be supported by xarray. For example, 512MB page cache on ARM64 when the base p xarray. It leads to errors as the following messages indicate when this sort of xarray entry is split. WARNING: CPU xas_split_alloc+0xf8/0x128 Modules linked in: binfmt_misc nft_fib_inet nft_fib_ipv4 nft_fib_ipv6 \ nft_fib nft_re nft_reject \ nft_ct nft_chain_nat nf_nat nf_conntrack nf_defrag_ipv6 nf_defrag_ipv4 \ ip_set rfkill nf_tables nfnetli \ libcrc32c crct10dif_ce ghash_ce sha2_ce sha256_arm64 sha1_ce virtio_net \ net_failover virtio_console virtio_bl PID: 7578 Comm: test Kdump: loaded Tainted: G W 6.10.0-rc5-gavin+ #9 Hardware name: QEMU KVM Virtual 05/24/2024 pstate: 83400005 (Nzcv daif +PAN -UAO +TCO +DIT -SSBS BTYPE=--) pc : xas_split_alloc+0xf8/0 +0x1c4/0x720 sp : ffff8000882af5f0 x29: ffff8000882af5f0 x28: ffff8000882af650 x27: ffff8000882af768 x26: 00 x24: ffff00010625b858 x23: ffff8000882af650 x22: ffffffdfc0900000 x21: 0000000000000000 x20: 000000000000 0000000000000000 x17: 0000000000000000 x16: 0000018000000000 x15: 52f8004000000000 x14: 0000e00000 x12: 0000000000000020 x11: 52f8000000000000 x10: 52f8e1c0ffff6000 x9 : ffffbeb9619a681c x8 : 000000000000 x6 : ffff00010b02ddb0 x5 : ffffbeb96395e378 x4 : 0000000000000000 x3 : 00000000000cc0 x2 : 000000000000 0000000000000000 Call trace: xas_split_alloc+0xf8/0x128 split_huge_page_to_list_to_order+0x1c4/0x720 trunca shmem_undo_range+0x2bc/0x6a8 shmem_fallocate+0x134/0x430 vfs_fallocate+0x124/0x2e8 ksys_fallocate+0x4 invoke_syscall.constprop.0+0x7c/0xd8 do_el0_svc+0xb4/0xd0 el0_svc+0x44/0x1d8 el0t_64_sync_handler+0x134 disabling PMD-sized page cache when HPAGE_PMD_ORDER is larger than MAX_PAGECACHE_ORDER. As a shmem file isn't represented by a multi-index entry and doesn't have this limitation when the xarry entry is split u index entries in the page cache").

CVE-2024-42243	In the Linux kernel, the following vulnerability has been resolved: mm/filemap: make MAX_PAGECACHE_ORD "mm/filemap: Limit page cache size to that supported by xarray", v2. Currently, xarray can't support arbitrary page from the WARN_ON() statement in xas_split_alloc(). In our test whose code is attached below, we hit the WARN_ base page size is 64KB and huge page size is 512MB. The issue was reported long time ago and some discussions www.spinics.net/lists/linux-xfs/msg75404.html In order to fix the issue, we need to adjust MAX_PAGECACHE_C and avoid PMD-sized page cache if needed. The code changes are suggested by David Hildenbrand. PATCH[1] ad to that supported by xarray PATCH[2-3] avoids PMD-sized page cache in the synchronous readahead path PATCH shmem files if needed Test program ========== # cat test.c #define _GNU_SOURCE #include <stdio.h> #inc #include <string.h> #include <fcntl.h> #include <errno.h> #include <sys/syscall.h> #include <sys/mman.h> #defi #define TEST_SHMEM_FILENAME "/dev/shm/data" #define TEST_MEM_SIZE 0x20000000 int main(int argc, fd = 0; void *buf = (void *)-1, *p; int pgsize = getpagesize(); int ret; if (pgsize != 0x10000) { fprintf(stderr, "64KB -EPERM; } system("echo force > /sys/kernel/mm/transparent_hugepage/shmem_enabled"); system("rm -fr /tmp/da system("echo 1 > /proc/sys/vm/drop_caches"); /* Open xfs or shmem file */ filename = TEST_XFS_FILENAME; filename = TEST_SHMEM_FILENAME; fd = open(filename, O_CREAT | O_RDWR | O_TRUNC); if (fd < 0) { \n", filename); return -EIO; } /* Extend file size */ ret = ftruncate(fd, TEST_MEM_SIZE); if (ret) { fprintf(stderr, " cleanup; } /* Create VMA */ buf = mmap(NULL, TEST_MEM_SIZE, PROT_READ | PROT_WRITE, MAP_SH { fprintf(stderr, "Unable to mmap <%s>\n", filename); goto cleanup; } fprintf(stdout, "mapped buffer at 0x%p\n", MADV_HUGEPAGE); if (ret) { fprintf(stderr, "Unable to madvise(MADV_HUGEPAGE)\n"); goto cleanup; } /* TEST_MEM_SIZE, MADV_POPULATE_WRITE); if (ret) { fprintf(stderr, "Error %d to madvise(MADV_POPU Punch the file to enforce xarray split */ ret = fallocate(fd, FALLOC_FL_KEEP_SIZE | FALLOC_FL_PUNCH_HO if (ret) fprintf(stderr, "Error %d to fallocate()\n", ret); cleanup: if (buf != (void *)-1) munmap(buf, TEST_MEM_SI gcc test.c -o test # cat /proc/1/smaps | grep KernelPageSize | head -n 1 KernelPageSize: 64 kB # ./test shmem : ----- CPU: 17 PID: 5253 at lib/xarray.c:1025 xas_split_alloc+0xf8/0x128 Modules linked in: nft_fib_inet nft_fib_ipv4 n nf_reject_ipv4 nf_reject_ipv6 nft_reject nft_ct \ nft_chain_nat nf_nat nf_conntrack nf_defrag_ipv6 nf_defrag_ipv4 virtio_balloon \ drm fuse xfs libcrc32c crct10dif_ce ghash_ce sha2_ce sha256_arm64 \ virtio_net sha1_ce net_failo virtio_mmio CPU: 17 PID: 5253 Comm: test Kdump: loaded Tainted: G W 6.10.0-rc5-gavin+ #12 Hardware name edk2-20240524-1.el9 05/24/2024 pstate: 83400005 (Nzcv daif +PAN -UAO +TC ---truncated---
CVE-2024-42278	In the Linux kernel, the following vulnerability has been resolved: ASoC: TAS2781: Fix tasdev_load_calibrated_d so it's either a no-op or it leads to a NULL dereference.
CVE-2024-42294	In the Linux kernel, the following vulnerability has been resolved: block: fix deadlock between sd_remove & sd_re task: [2538.459400] INFO: task "kworker/0:0":7 blocked for more than 188 seconds. [2538.459427] Call trace: [[2538.459436] __schedule+0x628/0x9c4 [2538.459442] schedule+0x7c/0xe8 [2538.459447] schedule_preempt_ __mutex_lock+0x3ec/0xf04 [2538.459456] __mutex_lock_slowpath+0x14/0x24 [2538.459459] mutex_lock+0x3 +0xdc/0x350 [2538.459466] sd_remove+0x30/0x60 [2538.459470] device_release_driver_internal+0x1c4/0x2c4 +0x18/0x28 [2538.459478] bus_remove_device+0x15c/0x174 [2538.459483] device_del+0x1d0/0x358 [2538.45 [2538.459493] scsi_forget_host+0x50/0x70 [2538.459497] scsi_remove_host+0x80/0x180 [2538.459502] usb_s usb_unbind_interface+0xd4/0x280 [2538.459510] device_release_driver_internal+0x1c4/0x2c4 [2538.459514] d [2538.459518] bus_remove_device+0x15c/0x174 [2538.459523] device_del+0x1d0/0x358 [2538.459528] usb_d usb_disconnect+0xec/0x300 [2538.459537] hub_event+0xb80/0x1870 [2538.459541] process_scheduled_works+ +0x244/0x334 [2538.459549] kthread+0x114/0x1bc [2538.461001] INFO: task "fsck.":15415 blocked for more t [2538.461016] __switch_to+0x174/0x338 [2538.461021] __schedule+0x628/0x9c4 [2538.461025] schedule+0x +0xc4/0x160 [2538.461034] blk_mq_alloc_request+0x120/0x1d4 [2538.461037] scsi_execute_cmd+0x7c/0x23c +0x5c/0x164 [2538.461046] scsi_set_medium_removal+0x5c/0xb0 [2538.461051] sd_release+0x50/0x94 [2538 [2538.461058] blkdev_release+0x28/0x40 [2538.461063] __fput+0xf8/0x2a8 [2538.461066] __fput_sync+0x28/ +0x84/0xe8 [2538.461073] invoke_syscall+0x58/0x114 [2538.461078] el0_svc_common+0xac/0xe0 [2538.461 el0_svc+0x38/0x68 [2538.461090] el0t_64_sync_handler+0x68/0xbc [2538.461093] el0t_64_sync+0x1a8/0x1ac __blk_mark_disk_dead blk_freeze_queue_start ++q->mq_freeze_depth bdev_release mutex_lock(&disk->open_m blk_queue_enter wait_event(!q->mq_freeze_depth) mutex_lock(&disk->open_mutex) SCSI does not set GD_OWI not set in this scenario. This is a classic ABBA deadlock. To fix the deadlock, make sure we don't try to acquire dis
CVE-2024-42302	In the Linux kernel, the following vulnerability has been resolved: PCI/DPC: Fix use-after-free on concurrent DPC free when a DPC event occurs concurrently to hot-removal of the same portion of the hierarchy: The dpc_handler() the Downstream Port where the DPC event occurred. To do so, it polls the config space of the first child device on concurrently removed, accesses to its struct pci_dev cause the kernel to oops. That's because pci_bridge_wait_for_s on the child device. Before v6.3, the function was only called on resume from system sleep or on runtime resume. I back then because the pciehp IRQ thread could never run concurrently. (On resume from system sleep, IRQs are no phase. And runtime resume is always awaited before a PCI device is removed.) However starting with v6.3, pci_br called on a DPC event. Commit 53b54ad074de ("PCI/DPC: Await readiness of secondary bus after reset"), which i pci_bridge_wait_for_secondary_bus() now needs to hold a reference on the child device because dpc_handler() and The commit was backported to v5.10+ stable kernels, so that's the oldest one affected. Add the missing reference ac unable to handle page fault for address: 00000000091400c0 CPU: 15 PID: 2464 Comm: irq/53-pcie-dpc 6.9.0 RIP: pci_dev_wait() pci_bridge_wait_for_secondary_bus() dpc_reset_link() pcie_do_recovery() dpc_handler()
CVE-2024-42315	In the Linux kernel, the following vulnerability has been resolved: exfat: fix potential deadlock on __exfat_get_den entries than ES_MAX_ENTRY_NUM, the bh-array is allocated in __exfat_get_entry_set. The problem is that the l It does not make sense. In the following cases, a deadlock for sbi->s_lock between the two processes may occur. C lock(fs_reclaim) exfat_iterate lock(&sbi->s_lock) exfat_readdir exfat_get_uniname_from_ext_entry exfat_get_den kmalloc_array ... lock(fs_reclaim) ... evict exfat_evict_inode lock(&sbi->s_lock) To fix this, let's allocate bh-array

CVE-2024-42316	In the Linux kernel, the following vulnerability has been resolved: mm/mglru: fix div-by-zero in vmpressure_calc_ to reclaim folios that have gone through page writeback and become clean before it finishes the first pass, since fol those folios due to the isolation. The second pass tries to avoid potential double counting by deducting scan_contro underflow of nr_scanned, under a condition where shrink_folio_list() does not increment nr_scanned, i.e., when fol the divisor, i.e., scale=scanned+reclaimed in vmpressure_calc_level(), to become zero, resulting in the following cr +101] process_one_work at ffffffffa3313f2b Since scan_control->nr_scanned has no established semantics, the pot Therefore, fix the problem by not deducting scan_control->nr_scanned in evict_folios().
CVE-2024-42317	In the Linux kernel, the following vulnerability has been resolved: mm/huge_memory: avoid PMD-size page cache page cache size. the largest and supported page cache size is defined as MAX_PAGECACHE_ORDER by commit MAX_PAGECACHE_ORDER acceptable to xarray"). However, it's possible to have 512MB page cache in the hu system whose base page size is 64KB. 512MB page cache is breaking the limitation and a warning is raised when t following example. [root@dhcp-10-26-1-207 ~]# cat /proc/1/smaps | grep KernelPageSize KernelPageSize: 64 kB test.c : int main(int argc, char **argv) { const char *filename = TEST_XFS_FILENAME; int fd = 0; void *buf = (v ret = 0; if (pgsize != 0x10000) { fprintf(stdout, "System with 64KB base page size is required!\n"); return -EPERM bdi/253:0/read_ahead_kb"); system("echo 1 > /proc/sys/vm/drop_caches"); /* Open the xfs file */ fd = open(filenar VMA */ buf = mmap(NULL, TEST_MEM_SIZE, PROT_READ, MAP_SHARED, fd, 0); assert(buf != (void *)-1 \n", buf); /* Populate VMA */ ret = madvise(buf, TEST_MEM_SIZE, MADV_NOHUGEPAGE); assert(ret == 0); MADV_POPULATE_READ); assert(ret == 0); /* Collapse VMA */ ret = madvise(buf, TEST_MEM_SIZE, MAD = madvise(buf, TEST_MEM_SIZE, MADV_COLLAPSE); if (ret) { fprintf(stdout, "Error %d to madvise(MADV_ Split xarray entry. Write permission is needed */ munmap(buf, TEST_MEM_SIZE); buf = (void *)-1; close(fd); fd 0); fallocate(fd, FALLOC_FL_KEEP_SIZE | FALLOC_FL_PUNCH_HOLE, TEST_MEM_SIZE - pgsize, pgsize) TEST_MEM_SIZE); if (fd > 0) close(fd); return ret; } [root@dhcp-10-26-1-207 ~]# gcc /tmp/test.c -o /tmp/test [ro ------------[cut here]------------ WARNING: CPU: 25 PID: 7560 at lib/xarray.c:1025 xas_split_alloc+0xf8/0x128 M nft_fib_ipv6 nft_fib \ nft_reject_inet nf_reject_ipv4 nf_reject_ipv6 nft_reject nft_ct \ nft_chain_nat nf_nat nf_conn ip_set rfkill nf_tables nfnetlink vfat fat virtio_balloon drm fuse \ xfs libcrc32c crct10dif_ce ghash_ce sha2_ce sha2 virtio_blk virtio_console failover dimlib virtio_mmio CPU: 25 PID: 7560 Comm: test Kdump: loaded Not tainted 6 KVM Virtual Machine, BIOS edk2-20240524-1.el9 05/24/2024 pstate: 83400005 (Nzcv daif +PAN -UAO +TCO + +0xf8/0x128 lr : split_huge_page_to_list_to_order+0x1c4/0x780 sp : ffff8000ac32f660 x29: ffff8000ac32f660 x28 x26: 0000000000000c40 x25: ffff0000e0969eb0 x24: 000000000000000d x23: ffff8000ac32f6c0 x22: ffffffdfc070 x20: 0000000000000000 x19: ffffffdfc0700000 x18: 0000000000000000 x17: 0000000000000000 x16: ffffd5f370 x14: 0000000000000000 x13: 0000000000000000 x12: 0000000000000000 x11: ffffffffffffffc0 x10: 00000000000 0000000000000003 x7 : 0000000000000000 x6 : ffff0000e0969eb8 x5 : ffffd5f37289e378 x4 : 000000000000000 000000000000000d x1 : 000000000000000c x0 : 0000000000000000 Call trace: xas_split_alloc+0xf8/0x128 split_ truncate_inode_partial_folio+0xdc/0x160 truncate_inode_pages_range+0x1b4/0x4a8 truncate_pagecache_range+0
CVE-2024-43816	In the Linux kernel, the following vulnerability has been resolved: scsi: lpfc: Revise lpfc_prep_embed_io routine w endian architectures, it is possible to run into a memory out of bounds pointer dereference when FCP targets are zo fcp_cmnd, sgl->sge_len) is referencing a little endian formatted sgl->sge_len value. So, the memcpy can cause big ptr as a struct sli4_sge_le to make it clear that we are referring to a little endian formatted data structure. And, upda usages.
CVE-2024-43820	In the Linux kernel, the following vulnerability has been resolved: dm-raid: Fix WARN_ON_ONCE check for syn occasionally trigger the following warning when being resumed after a table load because DM_RECOVERY_RUN at drivers/md/dm-raid.c:4105 raid_resume+0xee/0x100 [dm_raid] The failing check is: WARN_ON_ONCE(test_b >recovery)); This check is designed to make sure that the sync thread isn't registered, but md_check_recovery can s the sync_thread ever getting registered. Instead of checking if MD_RECOVERY_RUNNING is set, check if sync_
CVE-2024-43823	In the Linux kernel, the following vulnerability has been resolved: PCI: keystone: Fix NULL pointer dereference in ks_pcie_setup_rc_app_regs() If IORESOURCE_MEM is not provided in Device Tree due to any error, resource_li pci_parse_request_of_pci_ranges() will just emit a warning. This will cause a NULL pointer dereference. Fix this b Linux Verification Center (linuxtesting.org) with SVACE.
CVE-2024-43828	In the Linux kernel, the following vulnerability has been resolved: ext4: fix infinite loop when replaying fast_comm replay an infinite loop may occur due to an uninitialized extent_status struct. ext4_ext_determine_insert_hole() doe ext4_es_find_extent_range(), which will return immediately without initializing the 'es' variable. Because 'es' conta causing an infinite loop in this function, easily reproducible using fstest generic/039. This commit fixes this issue b function ext4_es_find_extent_range(). Thanks to Zhang Yi, for figuring out the real problem!
CVE-2024-43840	In the Linux kernel, the following vulnerability has been resolved: bpf, arm64: Fix trampoline for BPF_TRAMP_F BPF_TRAMP_F_CALL_ORIG is set, the trampoline calls __bpf_tramp_enter() and __bpf_tramp_exit() functions, *im pointer as an argument in R0. The trampoline generation code uses emit_addr_mov_i64() to emit instructions f R0, but emit_addr_mov_i64() assumes the address to be in the vmalloc() space and uses only 48 bits. Because bpf_ address can use more than 48-bits, in this case the trampoline will pass an invalid address to __bpf_tramp_enter/ex emit_a64_mov_i64() in place of emit_addr_mov_i64() as it can work with addresses that are greater than 48-bits.

CVE-2024-43855	In the Linux kernel, the following vulnerability has been resolved: md: fix deadlock between mddev_suspend and f being suspended while some flush bio is in progress. It is a complex issue. T1. the first flush is at the ending stage, submit data, but is blocked because mddev is suspended by T4. T2. the second flush sets 'mddev->flush_bio', and a which is already running (T1) and won't execute again if on the same CPU as T1. T3. the third flush inc active_io a 'mddev->flush_bio' is not NULL (set by T2). T4. mddev_suspend() is called and waits for active_io dec to 0 which (flush 2) (third 3) (suspend) md_submit_flush_data mddev->flush_bio = NULL; . . md_flush_request . mddev->flu md_handle_request . . active_io + 1 . . md_flush_request . . wait !mddev->flush_bio . . . . mddev_suspend . . wait ! md_submit_flush_data . //md_submit_flush_data is already running (T1) . md_handle_request wait resume The roo during flush process. active_io is dec before md_submit_flush_data is queued, and inc soon after md_submit_flush submit_flushes active_io - 1 md_submit_flush_data md_handle_request active_io + 1 make_request active_io - 1 I instead of within submit_flushes(), make_request() can be called directly intead of md_handle_request() in md_sub inc and dec once in the whole flush process. Deadlock will be fixed. Additionally, the only difference between fixi return error handling of make_request(). But after previous patch cleaned md_write_start(), make_requst() only ret see commit 41425f96d7aa ("dm-raid456, md/raid456: fix a deadlock for dm-raid456 while io concurrent with resha operation into two separate io, io size of flush submitted by dm always is 0, make_request() will not be called in m modifications from introducing issues, add WARN_ON to ensure make_request() no error is returned in this contex
CVE-2024-43873	In the Linux kernel, the following vulnerability has been resolved: vhost/vsock: always initialize seqpacket_allow 1. seqpacket_allow is not initialized when socket is created. Thus if features are never set, it will be read uninitializ is set and then cleared, then seqpacket_allow will not be cleared appropriately (existing apps I know about don't us to be sure no one relies on this). To fix: - initialize seqpacket_allow after allocation - set it unconditionally in set_fe
CVE-2024-43874	In the Linux kernel, the following vulnerability has been resolved: crypto: ccp - Fix null pointer dereference in __s dereference induced by DEBUG_TEST_DRIVER_REMOVE. Return from __sev_snp_shutdown_locked() if the p not initialized. Without the fix, the driver will produce the following splat: ccp 0000:55:00.5: enabling device (000 ccp 0000:55:00.5: psp enabled BUG: kernel NULL pointer dereference, address: 00000000000000f0 #PF: supervis error_code(0x0000) - not-present page PGD 0 P4D 0 Oops: 0000 [#1] PREEMPT SMP DEBUG_PAGEALLOC N Not tainted 6.9.0-rc1+ #29 RIP: 0010:__sev_snp_shutdown_locked+0x2e/0x150 Code: 00 55 48 89 e5 41 57 41 56 48 8b 04 25 28 00 00 00 48 89 45 d8 48 8b 05 6a 5a 7f 06 <4c> 8b a0 f0 00 00 00 41 0f b6 9c 24 a2 00 00 00 48 83 EFLAGS: 00010286 RAX: 0000000000000000 RBX: ffff9e4acd2e0a28 RCX: 0000000000000000 RDX: 0000000 RDI: ffffb2ea4014b808 RBP: ffffb2ea4014b7e8 R08: 0000000000000106 R09: 000000000003d9c0 R10: 0000000 ffff9e49d40590c8 R13: 0000000000000000 R14: ffffb2ea4014b808 R15: 0000000000000000 FS: 0000000000000 knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 00000000000000f0 CR3 0000000000770ef0 PKRU: 55555554 Call Trace: <TASK> ? __die_body+0x6f/0xb0 ? __die+0xcc/0xf0 ? page_fa +0x2a5/0x360 ? do_user_addr_fault+0x583/0x630 ? exc_page_fault+0x81/0x120 ? asm_exc_page_fault+0x2b/0x3 +0x2e/0x150 __sev_firmware_shutdown+0x349/0x5b0 ? pm_runtime_barrier+0x66/0xe0 sev_dev_destroy+0x34/0 +0x39/0x90 sp_pci_remove+0x22/0x60 pci_device_remove+0x4e/0x110 really_probe+0x271/0x4e0 __driver_pro +0x24/0x120 __driver_attach+0xc7/0x280 ? driver_attach+0x30/0x30 bus_for_each_dev+0x10d/0x130 driver_atta +0x171/0x2b0 ? unaccepted_memory_init_kdump+0x20/0x20 driver_register+0x67/0x100 __pci_register_driver+ sp_mod_init+0x13/0x30 do_one_initcall+0xb8/0x290 ? sched_clock_noinstr+0xd/0x10 ? local_clock_noinstr+0x3 +0x21e/0x6a0 ? local_clock+0x1c/0x60 ? stack_depot_save_flags+0x21e/0x6a0 ? sched_clock_noinstr+0xd/0x10 __lock_acquire+0xd90/0xe30 ? sched_clock_noinstr+0xd/0x10 ? local_clock_noinstr+0x3e/0x100 ? __create_obje __create_object+0x66/0x100 ? parameq+0x1b/0x90 ? parse_one+0x6d/0x1d0 ? parse_args+0xd7/0x1f0 ? do_initca +0xb0/0x180 do_initcalls+0x60/0xa0 ? kernel_init+0x1f/0x1d0 do_basic_setup+0x41/0x50 kernel_init_freeable+0 kernel_init+0x1f/0x1d0 ? rest_init+0x1f0/0x1f0 ret_from_fork+0x3d/0x50 ? rest_init+0x1f0/0x1f0 ret_from_fork_ in: CR2: 00000000000000f0 ---[end trace 0000000000000000]--- RIP: 0010:__sev_snp_shutdown_locked+0x2e/ 54 53 48 83 ec 10 41 89 f7 49 89 fe 65 48 8b 04 25 28 00 00 00 48 89 45 d8 48 8b 05 6a 5a 7f 06 <4c> 8b a0 f0 00 02 0f 83 RSP: 0018:ffffb2ea4014b7b8 EFLAGS: 00010286 RAX: 0000000000000000 RBX: ffff9e4acd2e0a28 RC truncated---
CVE-2024-43889	In the Linux kernel, the following vulnerability has been resolved: padata: Fix possible divide-by-0 panic in padata easily reproducible divide-by-0 panic in padata.c at bootup time. [10.017908] Oops: divide error: 0000 1 PREEMP PID: 2627 Comm: kworker/u1666:1 Not tainted 6.10.0-15.el10.x86_64 #1 [10.017908] Hardware name: Lenovo T [7X12CTO1WW], BIOS [PSE140J-2.30] 07/20/2021 [10.017908] Workqueue: events_unbound padata_mt_helpe +0x39/0xb0 : [10.017963] Call Trace: [10.017968] <TASK> [10.018004] ? padata_mt_helper+0x39/0xb0 [10.0 [10.018093] worker_thread+0x266/0x3a0 [10.018111] kthread+0xcf/0x100 [10.018124] ret_from_fork+0x31/0x +0x1a/0x30 [10.018147] </TASK> Looking at the padata_mt_helper() function, the only way a divide-by-0 panic way that chunk_size is initialized in padata_do_multithreaded(), chunk_size can be 0 when the min_chunk in the p divide-by-0 panic by making sure that chunk_size will be at least 1 no matter what the input parameters are.
CVE-2024-43896	In the Linux kernel, the following vulnerability has been resolved: ASoC: cs-amp-lib: Fix NULL pointer crash if ef efi_rt_services_supported() to check that efi.get_variable exists before calling it.
CVE-2024-4453	GStreamer EXIF Metadata Parsing Integer Overflow Remote Code Execution Vulnerability. This vulnerability allo code on affected installations of GStreamer. Interaction with this library is required to exploit this vulnerability but implementation. The specific flaw exists within the parsing of EXIF metadata. The issue results from the lack of pr can result in an integer overflow before allocating a buffer. An attacker can leverage this vulnerability to execute c Was ZDI-CAN-23896.

CVE-2024-44975	In the Linux kernel, the following vulnerability has been resolved: cgroup/cpuset: fix panic caused by partcmd_upd unable to handle page fault for address: 00000003 PGD 0 P4D 0 Oops: 0000 [#1] PREEMPT SMP NOPTI CPU: 3 W I 6.6.0-10893-g60d6 Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 1.15.0-1 04/4 RIP: 00 +0x483/0x600 Code: 01 48 85 d2 74 0d 48 83 05 29 3f f8 03 01 f3 48 0f bc c2 89 c0 48 9 RSP: 0018:ffffc90000fdb 0000000100000003 RBX: ffff888100b3dfa0 RCX: 0000000000000000 RDX: 0000000000000000 RSI: 00000000 RBP: ffff888100b3dfb0 R08: 0000000000000001 R09: 0000000000000000 R10: ffffc90000fdbcb0 R11: 0000000 R13: ffff888100a92b48 R14: 0000000000000000 R15: 0000000000000000 FS: 00007f44a5425740(0000) GS:ffff8 knlGS:0000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 0000000100030973 CR3: 00 00000000000006e0 Call Trace: <TASK> ? show_regs+0x8c/0xa0 ? __die_body+0x23/0xa0 ? __die+0x3a/0x50 ? partition_sched_domains_locked+0x483/0x600 ? search_module_extables+0x2a/0xb0 ? search_exception_tables+0 +0x144/0x1b0 ? __bad_area_nosemaphore+0x211/0x360 ? up_read+0x3b/0x50 ? bad_area_nosemaphore+0x1a/0x __lock_acquire.constprop.0+0x24f/0x8d0 ? __lock_acquire.constprop.0+0x24f/0x8d0 ? asm_exc_page_fault+0x26 +0x483/0x600 ? partition_sched_domains_locked+0xf0/0x600 rebuild_sched_domains_locked+0x806/0xdc0 upda cpuset_write_resmask+0xffc/0x1420 cgroup_file_write+0xb2/0x290 kernfs_fop_write_iter+0x194/0x290 new_syn ksys_write+0x81/0x180 __x64_sys_write+0x21/0x30 x64_sys_call+0x2f25/0x4630 do_syscall_64+0x44/0xb0 ent RIP: 0033:0x7f44a553c887 It can be reproduced with cammands: cd /sys/fs/cgroup/ mkdir test cd test/ echo +cpus cpuset.cpus.partition cat /sys/fs/cgroup/cpuset.cpus.effective 0-3 echo 0-3 > cpuset.cpus // taking away all cpus fro rebuilding of scheduling domains. In this scenario, test/cpuset.cpus.partition should be an invalid root and should n domains. When calling update_parent_effective_cpumask with partcmd_update, if newmask is not null, it should r available for parect/cs that has tasks.
CVE-2024-44983	In the Linux kernel, the following vulnerability has been resolved: netfilter: flowtable: validate vlan header Ensure field of the VLAN header, validate it once before the flowtable lookup. ========== KMSAN: uninit-value in nf_flow_offload_inet_hook+0x45a/0x5f0 net/netfilter/nf_flow_table_inet.c:32 nf_flow_o nf_flow_table_inet.c:32 nf_hook_entry_hookfn include/linux/netfilter.h:154 [inline] nf_hook_slow+0xf4/0x400 ne include/linux/netfilter_netdev.h:34 [inline] nf_ingress net/core/dev.c:5440 [inline]
CVE-2024-44985	In the Linux kernel, the following vulnerability has been resolved: ipv6: prevent possible UAF in ip6_xmit() If skb freed and the associated dst/idev could also have been freed. We must use rcu_read_lock() to prevent a possible UA
CVE-2024-44989	In the Linux kernel, the following vulnerability has been resolved: bonding: fix xfrm real_dev null pointer derefere because packets can be in transit and xfrm might call xdo_dev_offload_ok() in parallel. All callbacks assume real_ unable to handle page fault for address: 0000000000001030 kernel: bond0: (slave eni0np1): making interface the n access in kernel mode kernel: #PF: error_code(0x0002) - not-present page kernel: PGD 0 P4D 0 kernel: Oops: 000 PID: 2237 Comm: ping Not tainted 6.7.7+ #12 kernel: Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), RIP: 0010:nsim_ipsec_offload_ok+0xc/0x20 [netdevsim] kernel: bond0: (slave eni0np1): bond_ipsec_add_sa_all: 48 83 7f 38 00 74 de 0f 0b 48 8b 47 08 48 8b 37 48 8b 78 40 e9 b2 e5 9a d7 66 90 0f 1f 44 00 00 48 8b 86 80 02 0 00 00 c3 0f 1f 80 00 00 00 00 0f 1f kernel: bond0: (slave eni0np1): making interface the new active one kernel: RS 00010246 kernel: bond0: (slave eni0np1): bond_ipsec_add_sa_all: failed to add SA kernel: kernel: RAX: 0000000 ffff9eb403d97c60 kernel: RDX: ffffffffc090de10 RSI: ffff9eb404e74900 RDI: ffff9eb3c5de9e00 kernel: RBP: ffff 0000000000000014 kernel: R10: 7974203030303030 R11: 3030303030303030 R12: 0000000000000000 kernel: R R15: ffff9eb404c53000 kernel: FS: 00007f2a77a3ad00(0000) GS:ffff9eb43bd00000(0000) knlGS:0000000000000 CR0: 0000000080050033 kernel: CR2: 0000000000001030 CR3: 00000001122ab000 CR4: 0000000000350ef0 ke interface the new active one kernel: Call Trace: kernel: <TASK> kernel: ? __die+0x1f/0x60 kernel: bond0: (slave e add SA kernel: ? page_fault_oops+0x142/0x4c0 kernel: ? do_user_addr_fault+0x65/0x670 kernel: ? kvm_read_an (slave eni0np1): making interface the new active one kernel: ? exc_page_fault+0x7b/0x180 kernel: ? asm_exc_pag +0x50/0x50 [netdevsim] kernel: bond0: (slave eni0np1): bond_ipsec_add_sa_all: failed to add SA kernel: ? nsim_i kernel: bond0: (slave eni0np1): making interface the new active one kernel: bond_ipsec_offload_ok+0x7b/0x90 [b kernel: bond0: (slave eni0np1): bond_ipsec_add_sa_all: failed to add SA kernel: ip_push_pending_frames+0x56/0
CVE-2024-44990	In the Linux kernel, the following vulnerability has been resolved: bonding: fix null pointer deref in bond_ipsec_of slave before dereferencing the pointer.
CVE-2024-44994	In the Linux kernel, the following vulnerability has been resolved: iommu: Restore lost return in iommu_report_de iommu_report_device_fault gets called with a partial fault it is supposed to collect the fault into the group and then deleted which results in trying to process the fault and an eventual crash. Deleting the return was a typo, put it back
CVE-2024-44996	In the Linux kernel, the following vulnerability has been resolved: vsock: fix recursive ->recvmsg calls After a vso its prot->recvmsg has been replaced with vsock_bpf_recvmsg(). Thus the following recursiion could happen: vsock -> vsock_connectible_recvmsg() -> prot->recvmsg() -> vsock_bpf_recvmsg() again We need to fix it by calling th sockmap logic in __vsock_recvmsg().

CVE-2024-45000	In the Linux kernel, the following vulnerability has been resolved: fs/netfs/fscache_cookie: add missing "n_accesse dereference bug due to a data race which looks like this: BUG: kernel NULL pointer dereference, address: 000000 in kernel mode #PF: error_code(0x0000) - not-present page PGD 0 P4D 0 Oops: 0000 [#1] SMP PTI CPU: 33 PID tainted 6.8.7-cm4all1-hp+ #43 Hardware name: HP ProLiant DL380 Gen9/ProLiant DL380 Gen9, BIOS P89 10/1 netfs_rreq_write_to_cache_work RIP: 0010:cachefiles_prepare_write+0x30/0xa0 Code: 57 41 56 45 89 ce 41 55 4 ec 08 48 8b 47 08 48 83 7f 10 00 48 89 34 24 48 8b 68 20 <48> 8b 45 08 4c 8b 38 74 45 49 8b 7f 50 e8 4e a9 b0 f EFLAGS: 00010286 RAX: ffff976126be6d10 RBX: ffff97615cdb8438 RCX: 0000000000020000 RDX: ffff97605 ffff97615cdb8438 RBP: 0000000000000000 R08: 0000000000278333 R09: 0000000000000001 R10: ffff97605e6 ffff97605e6c4c68 R13: 0000000000020000 R14: 0000000000000001 R15: ffff976064fe2c00 FS: 0000000000000 knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 0000000000000008 CR3 0000000001706f0 Call Trace: <TASK> ? __die+0x1f/0x70 ? page_fault_oops+0x15d/0x440 ? search_module_ex +0x22/0x2f0 ? exc_page_fault+0x5f/0x100 ? asm_exc_page_fault+0x22/0x30 ? cachefiles_prepare_write+0x30/0x +0x135/0x2e0 process_one_work+0x137/0x2c0 worker_thread+0x2e9/0x400 ? __pfx_worker_thread+0x10/0x10 +0x10/0x10 ret_from_fork+0x30/0x50 ? __pfx_kthread+0x10/0x10 ret_from_fork_asm+0x1b/0x30 </TASK> Mo ---[end trace 0000000000000000]--- This happened because fscache_cookie_state_machine() was slow and was s fscache_unuse_cookie(); this led to a fscache_cookie_lru_do_one() call, setting the FSCACHE_COOKIE_DO_LR by fscache_cookie_state_machine(), withdrawing the cookie via cachefiles_withdraw_cookie(), clearing cookie->c process invoked cachefiles_prepare_write(), which found a NULL pointer in this code line: struct cachefiles_objec The next line crashes, obviously: struct cachefiles_cache *cache = object->volume->cache; During cachefiles_prep non-zero (via fscache_begin_operation()). The cookie must not be withdrawn until it drops to zero. The counter is before switching to FSCACHE_COOKIE_STATE_RELINQUISHING and FSCACHE_COOKIE_STATE_WITH FSCACHE_COOKIE_STATE_FAILED"), but not for FSCACHE_COOKIE_STATE_LRU_DISCARDING ("cas This patch adds the missing check. With a non-zero access counter, the function returns and the next fscache_end_ fscache_cookie_state_machine() call to handle the still-pending FSCACHE_COOKIE_DO_LRU_DISCARD.
CVE-2024-45013	In the Linux kernel, the following vulnerability has been resolved: nvme: move stopping keep-alive into nvme_uni start keep-alive after admin queue setup") moves starting keep-alive from nvme_start_ctrl() into nvme_init_ctrl_fi into nvme_uninit_ctrl(), so keep-alive work can be started and keep pending after failing to start controller, finally driver is unloaded. This patch fixes kernel panic when running nvme/004 in case that connection failure is triggere nvme_uninit_ctrl(). This way is reasonable because keep-alive is now started in nvme_init_ctrl_finish().
CVE-2024-45017	In the Linux kernel, the following vulnerability has been resolved: net/mlx5: Fix IPsec RoCE MPV trace call Preve not allowing IPsec creation over a slave, if master device doesn't support IPsec. WARNING: CPU: 44 PID: 16136 +0x75/0x94 Modules linked in: esp4_offload esp4 act_mirred act_vlan cls_flower sch_ingress mlx5_vdpa vringh v nfs_acl nfs lockd grace fscache netfs xt_CHECKSUM xt_MASQUERADE xt_conntrack ipt_REJECT nf_reject_ip nf_nat nf_conntrack nf_defrag_ipv6 nf_defrag_ipv4 rfkill cuse fuse rpcrdma sunrpc rdma_ucm ib_srpt ib_isert iscs ib_iser libiscsi scsi_transport_iscsi rdma_cm ib_ipoib iw_cm ib_cm ipmi_ssif intel_rapl_msr intel_rapl_common a kvm irqbypass crct10dif_pclmul crc32_pclmul mlx5_ib ghash_clmulni_intel sha1_ssse3 dell_smbios ib_uverbs ae dell_wmi_descriptor cryptd pcspkr ib_core acpi_ipmi sp5100_tco ccp i2c_piix4 ipmi_si ptdma k10temp ipmi_devi acpi_cpufreq ext4 mbcache jbd2 sd_mod t10_pi sg mgag200 drm_kms_helper syscopyarea sysfillrect mlx5_core s drm pci_hyperv_intf libata tg3 sha256_ssse3 tls megaraid_sas i2c_algo_bit psample wmi dm_mirror dm_region_h CPU: 44 PID: 16136 Comm: kworker/44:3 Kdump: loaded Tainted: GOE 5.15.0-20240509.el8uek.uek7_u3_updat Dell Inc. PowerEdge R7525/074H08, BIOS 2.0.3 01/15/2021 Workqueue: events xfrm_state_gc_task RIP: 0010:d 65 48 8b 14 25 80 fc 01 00 83 e0 02 48 09 d0 48 83 c8 01 48 89 45 08 5d 31 c0 89 c2 89 c6 89 c7 e9 cb 88 3b 00 ae 48 89 c2 48 83 ca 02 f0 RSP: 0018:ffffb26387773da8 EFLAGS: 00010282 RAX: 0000000000000000 RBX: fff RDX: 0000000000000000 RSI: ff886bc5e1366f2f RDI: 0000000000000000 RBP: ffffa08b658af940 R08: 000000 R10: 0000000000000000 R11: 0000000000000000 R12: ffffa0a9bfb31540 R13: ffffa0a9bfb37900 R14: 00000000 FS: 0000000000000000(0000) GS:ffffa0a9bfb00000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 00 000055a45ed814e8 CR3: 000000109038a000 CR4: 0000000000350ee0 Call Trace: <TASK> ? show_trace_log_lv +0x1d6/0x2f9 ? mlx5_devcom_for_each_peer_begin+0x29/0x60 [mlx5_core] ? down_read+0x75/0x94 ? __warn+ report_bug+0xa4/0x11d ? handle_bug+0x35/0x8b ? exc_invalid_op+0x14/0x75 ? asm_exc_invalid_op+0x16/0x1b +0xe/0x94 mlx5_devcom_for_each_peer_begin+0x29/0x60 [mlx5_core] mlx5_ipsec_fs_roce_tx_destroy+0xb1/0x [mlx5_core] tx_ft_put+0x53/0xc0 [mlx5_core] mlx5e_xfrm_free_state+0x45/0x90 [mlx5_core] ___xfrm_state_de +0x81/0xa9 process_one_work+0x1f1/0x3c6 worker_thread+0x53/0x3e4 ? process_one_work.cold+0x46/0x3c kth +0x60/0x52 ret_from_fork+0x22/0x2d </TASK> ---[end trace 5ef7896144d398e1]---
CVE-2024-45020	In the Linux kernel, the following vulnerability has been resolved: bpf: Fix a kernel verifier crash in stacksafe() Da crash when playing with sched-ext. Further investigation shows that the crash is due to invalid memory access in st following code: if (exact != NOT_EXACT && old->stack[spi].slot_type[i % BPF_REG_SIZE] != cur->stack[spi]. false; The 'i' iterates old->allocated_stack. If cur->allocated_stack < old->allocated_stack the out-of-bound access >allocated_stack' check such that if the condition is true, stacksafe() should fail. Otherwise, cur->stack[spi].slot_ty legal.

CVE-2024-45022	In the Linux kernel, the following vulnerability has been resolved: mm/vmalloc: fix page mapping if vm_area_allo 0 The __vmap_pages_range_noflush() assumes its argument pages** contains pages with the same page shift. How vmalloc: fix high order __GFP_NOFAIL allocations"), if gfp_flags includes __GFP_NOFAIL with high order in v failed for high order, the pages** may contain two different page shifts (high order and order-0). This could lead _ incorrect mappings, potentially resulting in memory corruption. Users might encounter this as follows (vmap_allow kvmalloc(2M, __GFP_NOFAIL|GFP_X) __vmalloc_node_range_noprof(vm_flags=VM_ALLOW_HUGE_VMAI order-9 allocation failed and fallback to order-0 vmap_pages_range() vmap_pages_range_noflush() __vmap_pages mapping happens We can remove the fallback code because if a high-order allocation fails, __vmalloc_node_range it is unnecessary to fallback to order-0 here. Therefore, fix this by removing the fallback code.
CVE-2024-45027	In the Linux kernel, the following vulnerability has been resolved: usb: xhci: Check for xhci->interrupters being all xhci_mem_init() fails, it calls into xhci_mem_cleanup() to mop up the damage. If it fails early enough, before xhci >max_interrupters has been set, which happens in most (all?) cases, things get uglier, as xhci_mem_cleanup() uncc With prejudice. Gate the interrupt freeing loop with a check on xhci->interrupters being non-NULL. Found while c XHCI driver on this exact path.
CVE-2024-45029	In the Linux kernel, the following vulnerability has been resolved: i2c: tegra: Do not mark ACPI devices as irq safe encounters an issue due to a mutex being called inside a spinlock. This leads to the following bug: BUG: sleeping f locking/mutex.c:585 ... Call trace: __might_sleep __mutex_lock_common mutex_lock_nested acpi_subsys_runtim problem arises because during __pm_runtime_resume(), the spinlock &dev->power.lock is acquired before rpm_re invokes acpi_subsys_runtime_resume(), which relies on mutexes, triggering the error. To address this issue, device considering the dependency of acpi_subsys_runtime_resume() on mutexes.
CVE-2024-46692	In the Linux kernel, the following vulnerability has been resolved: firmware: qcom: scm: Mark get_wq_ctx() as atc configured as a standard call. When two SMC calls are in sleep and one SMC wakes up, it calls get_wq_ctx() to res if get_wq_ctx() is interrupted, goes to sleep and another SMC call is waiting to be allocated a waitq context, it lead must be an atomic call and can't be a standard SMC call. Hence mark get_wq_ctx() as a fast call.
CVE-2024-46697	In the Linux kernel, the following vulnerability has been resolved: nfsd: ensure that nfsd4_fattr_args.context is zero a "goto out" before we get to checking for the security label, then args.context will be set to uninitialized junk on th it early.
CVE-2024-46698	In the Linux kernel, the following vulnerability has been resolved: video/aperture: optionally match the device in sy aperture_remove_conflicting_pci_devices(), we currently only call sysfb_disable() on vga class devices. This leads the pimary device is not VGA compatible: 1. A PCI device with a non-VGA class is the boot display 2. That device VGA device so sysfb_disable() is not called, but the device resources are freed by aperture_detach_platform_devic VGA class and it ends up calling sysfb_disable() 4. NULL pointer dereference via sysfb_disable() since the resourc aperture_detach_platform_device() when it was called by the other device. Fix this by passing a device pointer to s determine if we should execute it or not. v2: Fix build when CONFIG_SCREEN_INFO is not set v3: Move device in aperture_remove_conflicting_pci_devices() Drop __init on pci sysfb_pci_dev_is_enabled()
CVE-2024-46706	In the Linux kernel, the following vulnerability has been resolved: tty: serial: fsl_lpuart: mark last busy before uart initcall_debug=1 loglevel=8" in bootargs, kernel sometimes boot hang. It is because normal console still is not read console putchar will hang in waiting TRDE set in UARTSTAT. The lpuart driver has auto suspend delay set to 300 device serial ctrl will added and probed with its pm runtime enabled(see serial_ctrl.c). The runtime suspend call pa |->device_initial_probe |->__device_attach |-> pm_runtime_get_sync(dev->parent); |-> pm_request_idle(dev); |-> p the end, before normal console ready, the lpuart get runtime suspended. And earlycon putchar will hang. To addres pm_runtime_enable, three seconds is long enough to switch from bootconsole to normal console.
CVE-2024-46709	In the Linux kernel, the following vulnerability has been resolved: drm/vmwgfx: Fix prime with external buffers M goes through the dma_buf interface instead of trying to access pages directly. External buffers might not provide di make sure the bo's created from external dma_bufs can be read dma_buf interface has to be used. Fixes crashes in I usage won't trigger this due to the fact that virtual machines will not have multiple GPUs but it enables better test c
CVE-2024-46711	In the Linux kernel, the following vulnerability has been resolved: mptcp: pm: fix ID 0 endp usage after multiple re 'add_addr_accepted' are decremented for addresses not related to the initial subflow (ID0), because the source and are known from the beginning: they don't count as "additional local address being used" or "ADD_ADDR being ac them when the entrypoint used by the initial subflow is removed and re-added during a connection. Without this m and re-added more than once.
CVE-2024-46736	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix double put of @cfile in smb2_re called with a valid @cfile and returned -EINVAL, we need to call cifs_get_writable_path() again as the reference c smb2_compound_op() call.

CVE-2024-46766	In the Linux kernel, the following vulnerability has been resolved: ice: move netif_queue_set_napi to rtnl-protected is called from ice_vsi_rebuild() that is not rtnl-locked when called from the reset. This creates the need to take the r complicates the synchronization with .ndo_bpf. At the same time, there no actual need to fill napi-to-queue informa information when opening the VSI and clear it when the VSI is being closed. Those routines are already rtnl-locked in a way that prevents inclusion of XDP queues, as this leads to out-of-bounds writes, such as one below. [+0.000 in netif_queue_set_napi+0x1c2/0x1e0 [+0.000012] Write of size 8 at addr ffff889881727c80 by task bash/7047 [bash Not tainted 6.10.0-rc2+ #2 [+0.000004] Hardware name: Intel Corporation S2600WFT/S2600WFT, BIOS SE 08/26/2021 [+0.000003] Call Trace: [+0.000003] <TASK> [+0.000002] dump_stack_lvl+0x60/0x80 [+0.00000 [+0.000007] ? __pfx__raw_spin_lock_irqsave+0x10/0x10 [+0.000007] ? __virt_addr_valid+0x1c9/0x2c0 [+0.00 +0x1c2/0x1e0 [+0.000003] kasan_report+0xe9/0x120 [+0.000004] ? netif_queue_set_napi+0x1c2/0x1e0 [+0.000 [+0.000005] ice_vsi_close+0x161/0x670 [ice] [+0.000114] ice_dis_vsi+0x22f/0x270 [ice] [+0.000095] ice_pf_d [ice] [+0.000086] ice_prepare_for_reset+0x299/0x750 [ice] [+0.000087] pci_dev_save_and_disable+0x82/0xd0 [+0x12d/0x230 [+0.000004] reset_store+0xa0/0x100 [+0.000006] ? __pfx_reset_store+0x10/0x10 [+0.000002] ? [+0.000004] ? __check_object_size+0x4c1/0x640 [+0.000007] kernfs_fop_write_iter+0x30b/0x4a0 [+0.000006] fd_install+0x180/0x350 [+0.000005] ? __pfx_vfs_write+0x10/0xA10 [+0.000004] ? do_fcntl+0x52c/0xcd0 [+0.0 [+0.000003] ? kasan_save_free_info+0x37/0x60 [+0.000006] ksys_write+0xfa/0x1d0 [+0.000003] ? __pfx_ksys __x64_sys_fcntl+0x121/0x180 [+0.000004] ? _raw_spin_lock+0x87/0xe0 [+0.000005] do_syscall_64+0x80/0x1 [+0.000004] ? __pfx__raw_spin_lock+0x10/0x10 [+0.000003] ? file_close_fd_locked+0x167/0x230 [+0.000005 [+0.000005] ? do_syscall_64+0x8c/0x170 [+0.000004] ? do_syscall_64+0x8c/0x170 [+0.000003] ? do_syscall_6 +0x1a/0x2c0 [+0.000004] ? filp_close+0x19/0x30 [+0.000004] ? do_dup2+0x25a/0x4c0 [+0.000004] ? __x64_s syscall_exit_to_user_mode+0x7d/0x220 [+0.000004] ? do_syscall_64+0x8c/0x170 [+0.000003] ? __count_memc handle_mm_fault+0x136/0x820 [+0.000005] ? do_user_addr_fault+0x444/0xa80 [+0.000004] ? clear_bhb_loop+ +0x25/0x80 [+0.000002] entry_SYSCALL_64_after_hwframe+0x76/0x7e [+0.000005] RIP: 0033:0x7f20335931
CVE-2024-46767	In the Linux kernel, the following vulnerability has been resolved: net: phy: Fix missing of_node_put() for leds Th refcount incremented for leds, if it succeeds, it should call of_node_put() to decrease it, fix it.
CVE-2024-46786	In the Linux kernel, the following vulnerability has been resolved: fscache: delete fscache_cookie_lru_timer when avoid UAF The fscache_cookie_lru_timer is initialized when the fscache module is inserted, but is not deleted whe is removed. If timer_reduce() is called before removing the fscache module, the fscache_cookie_lru_timer will be a list of the current cpu. Afterwards, a use-after-free will be triggered in the softIRQ after removing the fscache modu == BUG: unable to hand PF: supervisor read access in kernel mode PF: error_code(0x0000) - not-present page PGD 21ffea067 P4D 21ffea0 PTE 0 Oops: Oops: 0000 [#1] PREEMPT SMP KASAN PTI CPU: 1 UID: 0 PID: 0 Comm: swapper/1 Tainted: G RIP: 0010:__run_timer_base.part.0+0x254/0x8a0 Call Trace: <IRQ> tmigr_handle_remote_up+0x627/0x810 __w tmigr_handle_remote+0x1fa/0x2f0 handle_softirqs+0x180/0x590 irq_exit_rcu+0x84/0xb0 sysvec_apic_timer_inte <TASK> asm_sysvec_apic_timer_interrupt+0x1a/0x20 RIP: 0010:default_idle+0xf/0x20 default_idle_call+0x38/0 cpu_startup_entry+0x54/0x60 start_secondary+0x20d/0x280 common_startup_64+0x13e/0x148 </TASK> Modul == Therefore delete fsca fscahe module.
CVE-2024-46793	In the Linux kernel, the following vulnerability has been resolved: ASoC: Intel: Boards: Fix NULL pointer deref in Since commit 13f58267cda3 ("ASoC: soc.h: don't create dummy Component via COMP_DUMMY()") dummy co SND_SOC_DAILINK_DEF(dummy, DAILINK_COMP_ARRAY(COMP_DUMMY())); expand to: static struct s { }; Which means that dummy is a zero sized array and thus dais[i].codecs should not be dereferenced *at all* sinc stored in the data section as the "dummy" variable has an address but no size, so even dereferencing dais[0] is alrea means that the if (dais[i].codecs->name) check added in commit 7d99a70b6595 ("ASoC: Intel: Boards: Fix NULL on that the part of the next variable which the name member maps to just happens to be NULL. Which apparently s then it results in crashes like this one: [28.795659] BUG: unable to handle page fault for address: 0000000000030 <TASK> ... [28.795862] ? strcmp+0x18/0x40 [28.795872] 0xffffffffc150c605 [28.795887] platform_probe+0x4 +0x10/0x10 [snd_soc_sst_bytcr_wm5102] Really fix things this time around by checking dais.num_codecs != 0.

CVE-2024-46796	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix double put of @cfile in smb2_se called with a valid @cfile and returned -EINVAL, we need to call cifs_get_writable_path() before retrying it as the by previous call. This fixes the following KASAN splat when running fstests generic/013 against Windows Server fs0/scratch run fstests generic/013 at 2024-09-02 19:48:59 ========== BUG: KASAN: slab-use-after-free in detach_if_pending+0xab/0x200 Write of size 8 at addr ffff88811f1a3730 by 176 Comm: kworker/3:2 Not tainted 6.11.0-rc6 #2 Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIC cifsoplockd cifs_oplock_break [cifs] Call Trace: <TASK> dump_stack_lvl+0x5d/0x80 ? detach_if_pending+0xab/ detach_if_pending+0xab/0x200 ? __virt_addr_valid+0x145/0x300 ? __phys_addr+0x46/0x90 ? detach_if_pending detach_if_pending+0xab/0x200 detach_if_pending+0xab/0x200 timer_delete+0x96/0xe0 ? __pfx_timer_delete+0x try_to_grab_pending+0x46/0x3b0 __cancel_work+0x89/0x1b0 ? __pfx___cancel_work+0x10/0x10 ? kasan_save_ +0x110/0x2c0 [cifs] ? __pfx_cifs_close_deferred_file+0x10/0x10 [cifs] ? __pfx_down_read+0x10/0x10 cifs_oploc __pfx_cifs_oplock_break+0x10/0x10 [cifs] ? lock_is_held_type+0x85/0xf0 ? mark_held_locks+0x1a/0x90 process +0x8a/0xa0 ? __pfx_process_one_work+0x10/0x10 ? lock_acquired+0x220/0x550 ? __list_add_valid_or_report+0 __kthread_parkme+0xd1/0xf0 ? __pfx_worker_thread+0x10/0x10 kthread+0x17f/0x1c0 ? kthread+0xda/0x1c0 ? _ +0x31/0x60 ? __pfx_kthread+0x10/0x10 ret_from_fork_asm+0x1a/0x30 </TASK> Allocated by task 1118: kasan_ +0x14/0x30 __kasan_kmalloc+0xaa/0xb0 cifs_new_fileinfo+0xc8/0x9d0 [cifs] cifs_atomic_open+0x467/0x770 [c path_openat+0x4c3/0x1380 do_filp_open+0x167/0x270 do_sys_openat2+0x129/0x160 __x64_sys_creat+0xad/0x entry_SYSCALL_64_after_hwframe+0x77/0x7f Freed by task 83: kasan_save_stack+0x30/0x50 kasan_save_track +0x3b/0x70 poison_slab_object+0xe9/0x160 __kasan_slab_free+0x32/0x50 kfree+0xf2/0x300 process_one_work kthread+0x17f/0x1c0 ret_from_fork+0x31/0x60 ret_from_fork_asm+0x1a/0x30 Last potentially related work creat __kasan_record_aux_stack+0xad/0xc0 insert_work+0x29/0xe0 __queue_work+0x5ea/0x760 queue_work_on+0x6 [cifs] smb2_compound_op+0x911/0x3940 [cifs] smb2_set_path_size+0x228/0x270 [cifs] cifs_set_file_size+0x197 [cifs] notify_change+0x4e3/0x740 do_truncate+0xfa/0x180 vfs_truncate+0x195/0x200 __x64_sys_truncate+0x109 entry_SYSCALL_64_after_hwframe+0x77/0x7f
CVE-2024-46797	In the Linux kernel, the following vulnerability has been resolved: powerpc/qspinlock: Fix deadlock in MCS queue queued_spin_lock_slowpath() after we increment qnodesp->count and before node->lock is initialized, another CP get_tail_qnode(). If the stale lock value happens to match the lock on that CPU, then we write to the "next" pointer a deadlock as the former CPU, once it becomes the head of the MCS queue, will spin indefinitely until it's "next" p the queue. Running stress-ng on a 16 core (16EC/16VP) shared LPAR, results in occasional lockups similar to the vm-bytes 80% --aggressive \ --maximize --oomable --verify --syslog \ --metrics --times --timeout 5m watchdog: CI [c0000000000b78f4] queued_spin_lock_slowpath+0x1184/0x1490 LR [c000000001037c5c] _raw_spin_lock+0x6 (unreliable) _raw_spin_lock+0x6c/0x90 raw_spin_rq_lock_nested.part.135+0x4c/0xd0 sched_ttwu_pending+0x60 +0x1dc/0x670 smp_ipi_demux_relaxed+0xa4/0x100 xive_muxed_ipi_action+0x20/0x40 __handle_irq_event_perc +0x2c/0x80 handle_percpu_irq+0x84/0xd0 generic_handle_irq+0x54/0x80 __do_irq+0xac/0x210 __do_IRQ+0x74 hardware_interrupt_common_virt+0x29c/0x2a0 --- interrupt: 500 at queued_spin_lock_slowpath+0x4b8/0x1490 .. queued_spin_lock_slowpath+0x4b8/0x1490 LR [c000000001037c5c] _raw_spin_lock+0x6c/0x90 --- interrupt: 50 _raw_spin_lock+0x6c/0x90 futex_wake+0x100/0x260 do_futex+0x21c/0x2a0 sys_futex+0x98/0x270 system_call system_call_vectored_common+0x15c/0x2ec The following code flow illustrates how the deadlock occurs. For the (A and B) are contended and we call the queued_spin_lock_slowpath() function. CPU0 CPU1 ---- ---- spin_lock_ir | spin_lock(B) | | ,ñº | id = qnodesp->count++; | (Note that nodes[0].lock == A) | | ,ñº | Interrupt | (happens before spin_lock_irqsave(A) | | ,ñº | id = qnodesp->count++ | nodes[1].lock = A | | ,ñº | Tail of MCS queue | | spin_lock_ CPU0 is previous tail ,ñº | Spin indefinitely ,ñº (until "nodes[1].next != NULL") prev = get_tail_qnode(A, CPU0) | qnodes ---truncated---
CVE-2024-46862	In the Linux kernel, the following vulnerability has been resolved: ASoC: Intel: soc-acpi-intel-mtl-match: add miss in struct snd_soc_acpi_mach {}, and we test !link->num_adr as a condition to end the loop in hda_sdw_machine_s snd_soc_acpi_link_adr array is required.
CVE-2024-46863	In the Linux kernel, the following vulnerability has been resolved: ASoC: Intel: soc-acpi-intel-lnl-match: add missi in struct snd_soc_acpi_mach {}, and we test !link->num_adr as a condition to end the loop in hda_sdw_machine_s snd_soc_acpi_link_adr array is required.
CVE-2024-47675	In the Linux kernel, the following vulnerability has been resolved: bpf: Fix use-after-free in bpf_uprobe_multi_link bpf_uprobe_multi_link_attach() goes to the error_free label and frees the array of bpf_uprobe's without calling bpf >uprobe and worse, this frees bpf_uprobe->consumer without removing it from the uprobe->consumers list.
CVE-2024-47682	In the Linux kernel, the following vulnerability has been resolved: scsi: sd: Fix off-by-one error in sd_read_block_ 0xb1 with length 8 (happens with qemu v2.x, for example), sd_read_block_characteristics() may attempt an out-of- zoned field at offset 8.

CVE-2024-47687	In the Linux kernel, the following vulnerability has been resolved: vdpa/mlx5: Fix invalid mr resource destroy Cert can end up releasing mr resources which never got initialized in the first place. This patch adds the missing check i to block releasing non-initialized mr resources. Reference trace: mlx5_core 0000:08:00.2: mlx5_vdpa_dev_add:32 provisioned? BUG: kernel NULL pointer dereference, address: 0000000000000000 #PF: supervisor read access in - not-present page PGD 140216067 P4D 0 Oops: 0000 [#1] PREEMPT SMP NOPTI CPU: 8 PID: 2700 Comm: vd 5.14.0-496.el9.x86_64 #1 Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS rel-1.13.0-0-gf21b5a4 RIP: 0010:vhost_iotlb_del_range+0xf/0xe0 [vhost_iotlb] Code: [...] RSP: 0018:ff1c823ac23077f0 EFLAGS: 0001(ffffffff899567a0 RCX: 0000000000000000 RDX: ffffffffffffffff RSI: 0000000000000000 RDI: 000000000000000 0000000000000000 R09: ff1c823ac2307670 R10: ff1c823ac2307668 R11: ffffffff8a9e7b68 R12: 0000000000000(ff1bda1f43e341a0 R15: 00000000fffffea FS: 00007f56eba7c740(0000) GS:ff1bda269f800000(0000) knlGS:0000(CR0: 0000000080050033 CR2: 0000000000000000 CR3: 0000000104d90001 CR4: 0000000000771ef0 DR0: 000 DR2: 0000000000000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 PKRU: 55 +0x1c4/0x2df ? show_trace_log_lvl+0x1c4/0x2df ? mlx5_vdpa_free+0x3d/0x150 [mlx5_vdpa] ? __die_body.cold __irq_work_queue_local+0x2b/0xc0 ? irq_work_queue+0x2c/0x50 ? exc_page_fault+0x62/0x150 ? asm_exc_page +0x10/0x10 [mlx5_vdpa] ? vhost_iotlb_del_range+0xf/0xe0 [vhost_iotlb] mlx5_vdpa_free+0x3d/0x150 [mlx5_vd [vdpa] device_release+0x31/0x90 kobject_cleanup+0x37/0x130 mlx5_vdpa_dev_add+0x2d2/0x7a0 [mlx5_vdpa] v +0x277/0x4c0 [vdpa] genl_family_rcv_msg_doit+0xd9/0x130 genl_family_rcv_msg+0x14d/0x220 ? __pfx_vdpa_ [vdpa] ? _copy_to_user+0x1a/0x30 ? move_addr_to_user+0x4b/0xe0 genl_rcv_msg+0x47/0xa0 ? __import_iovec +0x10/0x10 netlink_rcv_skb+0x54/0x100 genl_rcv+0x24/0x40 netlink_unicast+0x245/0x370 netlink_sendmsg+0x do_read_fault+0x10c/0x1d0 ? do_pte_missing+0x10d/0x190 __x64_sys_sendto+0x20/0x30 do_syscall_64+0x5c/(+0x4f/0xb0 ? mm_account_fault+0x6c/0x100 ? handle_mm_fault+0x116/0x270 ? do_user_addr_fault+0x1d6/0x6a clear_bhb_loop+0x25/0x80 ? clear_bhb_loop+0x25/0x80 ? clear_bhb_loop+0x25/0x80 ? clear_bhb_loop+0x25/0x entry_SYSCALL_64_after_hwframe+0x78/0x80
CVE-2024-47700	In the Linux kernel, the following vulnerability has been resolved: ext4: check stripe size compatibility on remount __ext4_fill_super if it is not a multiple of the cluster ratio however this check is missed when trying to remount. Th cluster_ratio after remount:set making EXT4_B2C(sbi->s_stripe) become 0 that can cause some unforeseen bugs li in remount path as well.
CVE-2024-47702	In the Linux kernel, the following vulnerability has been resolved: bpf: Fail verification for sign-extension of packe a kernel crash due to commit 1f1e864b6555 ("bpf: Handle sign-extenstin ctx member accesses"). The reason is due data/data_end/data_meta uapi field. The original code looks like: r2 = *(s32 *)(r1 + 76) /* load __sk_buff->data */ >data_end */ r0 = r2 r0 += 8 if r3 > r0 goto +1 ... Note that __sk_buff->data load has 32-bit sign extension. After v final asm code looks like: r2 = *(u64 *)(r1 +208) r2 = (s32)r2 r3 = *(u64 *)(r1 +80) r0 = r2 r0 += 8 if r3 > r0 goto kernel __sk_buff->data address invalid which may cause runtime failure. Currently, in C code, typically we have v *data_end = (void *)(long)skb->data_end; ... and it will generate r2 = *(u64 *)(r1 +208) r3 = *(u64 *)(r1 +80) r0 = sign-extension, void *data = (void *)(long)(int)skb->data; void *data_end = (void *)(long)skb->data_end; ... the ge +208) r2 <<= 32 r2 s>>= 32 r3 = *(u64 *)(r1 +80) r0 = r2 r0 += 8 if r3 > r0 goto pc+1 and this will cause verificati as "r2" is a packet pointer. To fix this issue for case r2 = *(s32 *)(r1 + 76) /* load __sk_buff->data */ this patch ad callback function for packet data/data_end/data_meta access. If those accesses are with sign-extenstion, the verifica bpf/000000000000c90eee061d236d37@google.com/
CVE-2024-47715	In the Linux kernel, the following vulnerability has been resolved: wifi: mt76: mt7915: fix oops on non-dbdc mt79 = 1 on the main phy for mt7986 with MT7975_ONE_ADIE or MT7976_ONE_ADIE. Commit 0335c034e726 ("wi checking tx queue fill status") introduced a dereference of the phys array indirectly indexed by band_idx via wcid- This caused the following Oops on affected mt7986 devices: Unable to handle kernel read from unreadable memor Mem abort info: ESR = 0x0000000096000005 EC = 0x25: DABT (current EL), IL = 32 bits SET = 0, FnV = 0 EA level 1 translation fault Data abort info: ISV = 0, ISS = 0x00000005 CM = 0, WnR = 0 user pgtable: 4k pages, 39-b [0000000000000024] pgd=0000000000000000, p4d=0000000000000000, pud=0000000000000000 Internal error: Modules linked in: ... mt7915e mt76_connac_lib mt76 mac80211 cfg80211 ... CPU: 2 PID: 1631 Comm: hostapd N name: ZyXEL EX5700 (Telenor) (DT) pstate: 80400005 (Nzcv daif +PAN -UAO -TCO -DIT -SSBS BTYPE=--) p [mt76] lr : mt76_wcid_cleanup+0x64/0x22c [mt76] sp : fffffc00a803700 x29: fffffc00a803700 x28: ffffff80008f7 ffffff80000a7880 x25: fffffc008c26e00 x24: 0000000000000001 x23: fffffc000a68114 x22: 0000000000000000 x19: ffffff8004152020 x18: 0000000000000000 x17: 00000000000017c0 x16: fffffc008ef5000 x15: 0000000000(x13: ffffff8004172e28 x12: 0000000000000000 x11: 0000000000000000 x10: ffffff8004172e30 x9 : ffffff800417(ffffff8004156020 x6 : 0000000000000000 x5 : 0000000000000031 x4 : 0000000000000000 x3 : 0000000000000(ffffff80008f7300 x0 : 0000000000000024 Call trace: mt76_wcid_cleanup+0x84/0x22c [mt76] __mt76_sta_remove +0x8c/0x1a4 [mt76] mt7915_eeprom_get_power_delta+0x11e4/0x23a0 [mt7915e] drv_sta_state+0x144/0x274 [m [mac80211] sta_set_sinfo+0xaf8/0xc24 [mac80211] sta_info_destroy_addr_bss+0x4c/0x6c [mac80211] ieee80211 [mac80211] cfg80211_check_station_change+0x1360/0x4710 [cfg80211] genl_family_rcv_msg_doit+0xb4/0x110 +0x58/0x120 genl_rcv+0x34/0x50 netlink_unicast+0x1f0/0x2ec netlink_sendmsg+0x198/0x3d0 ____sys_sendmsg +0x80/0xf0 __sys_sendmsg+0x44/0xa0 __arm64_sys_sendmsg+0x20/0x30 invoke_syscall.constprop.0+0x4c/0xe(+0x14/0x4c el0t_64_sync_handler+0x100/0x110 el0t_64_sync+0x15c/0x160 Code: d2800002 910092c0 5280002 7e42dd9a39ed2281]--- Fix by using mt76_dev_phy() which will map band_idx to the correct phy for all hardware

CVE-2024-47719	In the Linux kernel, the following vulnerability has been resolved: iommufd: Protect against overflow of ALIGN() an iova and uptr such that the target iova alignment becomes really big and ALIGN() overflows which corrupts the CONFIG_IOMMUFD_TEST can detect this: WARNING: CPU: 1 PID: 5092 at drivers/iommu/iommufd/io_paget iommu/iommufd/io_pagetable.c:268 [inline] WARNING: CPU: 1 PID: 5092 at drivers/iommu/iommufd/io_pageta drivers/iommu/iommufd/io_pagetable.c:352 Modules linked in: CPU: 1 PID: 5092 Comm: syz-executor294 Not ta g3ffea9a7a6f7 #0 Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google 06/07 drivers/iommu/iommufd/io_pagetable.c:268 [inline] RIP: 0010:iopt_map_pages+0xf95/0x1050 drivers/iommu/iom a4 f3 ff ff e8 1a 8b 4c fc 41 be e4 ff ff ff e9 8a f3 ff ff e8 0a 8b 4c fc 90 0f 0b 90 e9 37 f5 ff ff e8 fc 8a 4c fc 90 <0f ad 8f 80 e1 07 80 c1 03 38 RSP: 0018:ffffc90003ebf9e0 EFLAGS: 00010293 RAX: ffffffff85499fa4 RBX: 00000 RDX: 0000000000000000 RSI: 00000000ffffffef RDI: 0000000000000000 RBP: ffffc90003ebfc50 R08: ffffffff85 0000000000000002 R11: ffff888079b49e00 R12: ffff8880228e0010 R13: 0000000000000000 R14: 1ffff920007d7 000055557d760380(0000) GS:ffff8880b9500000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 0000 00000000005fdeb8 CR3: 000000007404a000 CR4: 00000000003506f0 DR0: 0000000000000000 DR1: 00000000 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 Call Trace: <TASK> iommufd_ioas_copy+ ioas.c:274 iommufd_fops_ioctl+0x4d9/0x5a0 drivers/iommu/iommufd/main.c:421 vfs_ioctl fs/ioctl.c:51 [inline] __ __se_sys_ioctl+0xfc/0x170 fs/ioctl.c:893 do_syscall_x64 arch/x86/entry/common.c:52 [inline] do_syscall_64+0xf entry_SYSCALL_64_after_hwframe+0x77/0x7f Cap the automatic alignment to the huge page size, which is prob alignments can fragment and chew up the available IOVA space without any reason.
CVE-2024-47734	In the Linux kernel, the following vulnerability has been resolved: bonding: Fix unnecessary warnings and logs fro syzbot reported a WARNING in bond_xdp_get_xmit_slave. To reproduce this[1], one bond device (bond1) has xdp bpf_master_redirect_enabled_key. Another bond device (bond0) which is unsupported by XDP but its slave (veth3 This triggers WARN_ON_ONCE() from the xdp_master_redirect(). To reduce unnecessary warnings and improve the WARN_ON_ONCE() and add ratelimit to the netdev_err(). [1] Steps to reproduce: # Needs tx_xdp with return peer veth1 ip l add veth3 type veth peer veth4 ip l add bond0 type bond mode 6 # BOND_MODE_ALB, unsupport # BOND_MODE_ROUNDROBIN by default ip l set veth0 master bond1 ip l set bond1 up # Increases bpf_master bond1 xdpdrv object tx_xdp.o section xdp_tx ip l set veth3 master bond0 ip l set bond0 up ip l set veth4 up # Trigg xdp_master_redirect() ip l set veth3 xdpgeneric object tx_xdp.o section xdp_tx
CVE-2024-49862	In the Linux kernel, the following vulnerability has been resolved: powercap: intel_rapl: Fix off by one in get_rpi() rpi_tpmi which have NR_RAPL_PRIMITIVES number of elements. Thus the > needs to be >= to prevent an off by
CVE-2024-49864	In the Linux kernel, the following vulnerability has been resolved: rxrpc: Fix a race between socket set up and I/O sets up the socket and then sets up the I/O thread that will handle it. This is a problem, however, as there's a gap be come into rxrpc_encap_rcv() from the UDP packet but we oops when trying to wake the not-yet created I/O thread discard the packet if there's no I/O thread yet. A better, but more intrusive fix would perhaps be to rearrange things O thread.
CVE-2024-49885	In the Linux kernel, the following vulnerability has been resolved: mm, slub: avoid zeroing kmalloc redzone Since slub: extend redzone check to extra allocated kmalloc space than requested"), setting orig_size treats the wasted spa redzone. However with init_on_free=1 we clear the full object->size, including the redzone. Additionally we clear stored orig_size, making it zero, which makes check_object() treat the whole object as a redzone. These issues lead "slub_debug=FUZ init_on_free=1": [0.000000] ========== [0.000000] BUG kmalloc-8 (Not tainted): kmalloc Redzone overwritten [0.000000] ---------- [0.000000] [0.000000] 0xffff000010032858-0xffff00001003285f @offset=2136. First byte 0x0 instead of 0xcc [kmalloc Redzone 0xffff000010032858-0xffff00001003285f=0xcc [0.000000] Slab 0xfffffdffc0400c80 objects=36 flags=0x3fffe0000000200(workingset|node=0|zone=0|lastcpupid=0x1ffff) [0.000000] Object 0xffff000010032858 [0.000000] [0.000000] Redzone ffff000010032850: cc cc cc cc cc cc cc cc ..... [0.000000] Object ffff00001003 [0.000000] Redzone ffff000010032860: cc cc cc cc cc cc cc cc ..... [0.000000] Padding ffff0000100328b4: 00 0 [0.000000] CPU: 0 UID: 0 PID: 0 Comm: swapper/0 Not tainted 6.11.0-rc3-next-20240814-00004-g61844c55c3f4 NXP i.MX95 19X19 board (DT) [0.000000] Call trace: [0.000000] dump_backtrace+0x90/0xe8 [0.000000] show dump_stack_lvl+0x74/0x8c [0.000000] dump_stack+0x18/0x24 [0.000000] print_trailer+0x150/0x218 [0.00000 free_to_partial_list+0x2f8/0x5ec To address the issue, use orig_size to clear the used area. And restore the value of When CONFIG_SLUB_DEBUG not defined, (get_orig_size()' directly returns s->object_size. So when using mem orig_size, as orig_size returns object_size when CONFIG_SLUB_DEBUG not enabled. And orig_size can never b
CVE-2024-49951	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: MGMT: Fix possible crash on mgmt called while there are commands queued on cmd_sync it could lead to crashes like the bellow trace: 0x0000053D: 0x0000053D: mgmt_pending_remove+0x18/0x58 [bluetooth] 0x0000053E: mgmt_remove_adv_monitor_complet hci_cmd_sync_work+0xbc/0x164 [bluetooth] So while handling mgmt_index_removed this attempts to dequeue co

CVE-2024-49953	In the Linux kernel, the following vulnerability has been resolved: net/mlx5e: Fix crash caused by calling __xfrm_s is not checked in driver's delayed work. When xfrm_state_check_expire() is called, the state can be reset to XFRM XFRM_STATE_DEAD already. This happens when xfrm state is deleted, but not freed yet. As __xfrm_state_delet following crash occurs. To fix this issue, skip xfrm_state_check_expire() if km.state is not XFRM_STATE_VALI probably for non-canonical address 0xdead000000000108: 0000 [#1] SMP CPU: 5 UID: 0 PID: 7448 Comm: kwo rc2+ #1 Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS rel-1.13.0-0-gf21b5a4aeb02-prebuilt.qe mlx5e_ipsec: eth%d mlx5e_ipsec_handle_sw_limits [mlx5_core] RIP: 0010:__xfrm_state_delete+0x3d/0x1b0 Co 87 c8 00 00 00 05 48 8d bb 40 10 00 00 e8 11 04 1a 00 48 8b 95 b8 00 00 00 48 8b 85 c0 00 00 00 <48> 89 42 08 00 00 00 00 ad de 48 RSP: 0018:ffff88885f945ec8 EFLAGS: 00010246 RAX: dead000000000122 RBX: ffffffff82 RDX: dead000000000100 RSI: 0000000000000000 RDI: ffffffff82afb980 RBP: ffff888109a20340 R08: ffff88885 R10: 0000000000000000 R11: ffff88885f945ff8 R12: 0000000000000246 R13: ffff888109a20340 R14: ffff88885 FS: 0000000000000000(0000) GS:ffff88885f940000(0000) knlGS:0000000000000000 CS: 0010 DS: 0000 ES: 00 00007f2163102430 CR3: 00000001128d6001 CR4: 0000000000370eb0 DR0: 0000000000000000 DR1: 0000000 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000000000000400 Call Trace: <IRQ> ? die_addr+0x33/0x90 asm_exc_general_protection+0x22/0x30 ? __xfrm_state_delete+0x3d/0x1b0 ? __xfrm_state_delete+0x2f/0x1b0 x __xfrm_state_delete+0x1b0/0x1b0 __hrtimer_run_queues+0x121/0x270 hrtimer_run_softirq+0x88/0xd0 handle_s </IRQ> <TASK> __local_bh_enable_ip+0x47/0x50 mlx5e_ipsec_handle_sw_limits+0x7d/0x90 [mlx5_core] proc +0x28d/0x3a0 ? rescuer_thread+0x480/0x480 kthread+0xb8/0xe0 ? kthread_park+0x80/0x80 ret_from_fork+0x2d ret_from_fork_asm+0x11/0x20 </TASK>
CVE-2024-49954	In the Linux kernel, the following vulnerability has been resolved: static_call: Replace pointless WARN_ON() in s static_call_module_notify() triggers a WARN_ON(), when memory allocation fails in __static_call_add_module(). case must be correctly handled by the well known call chain and the error code is passed through to the initiating u fail is not a fatal problem, but the WARN_ON() takes the machine out when panic_on_warn is set. Replace it with
CVE-2024-49976	In the Linux kernel, the following vulnerability has been resolved: tracing/timerlat: Drop interface_lock in stop_kt callback for "trace/osnoise:online", since commit 5bfbcd1ee57b ("tracing/timerlat: Add interface_lock around clea following ABBA deadlock scenario is introduced: T1 | T2 [BP] | T3 [AP] osnoise_hotplug_workfn() | work_for_cp | osnoise_cpu_die() mutex_lock(&interface_lock) | | stop_kthread() | cpus_write_lock() | mutex_lock(&interface_lo As the interface_lock here in just for protecting the "kthread" field of the osn_var, use xchg() instead to fix this issu stop_per_cpu_kthreads() as it can take cpu_read_lock() again.
CVE-2024-49977	In the Linux kernel, the following vulnerability has been resolved: net: stmmac: Fix zero-division error when disab stmmac: No need to calculate speed divider when offload is disabled") allows the "port_transmit_rate_kbps" to be the "div_s64" function when tc-cbs is disabled. This leads to a zero-division error. When tc-cbs is disabled, the idle values are not required to be configured. Therefore, adding a return statement after setting the txQ mode to DCB w division error.
CVE-2024-49983	In the Linux kernel, the following vulnerability has been resolved: ext4: drop ppath from ext4_ext_replay_update_ ext4_force_split_extent_at() in ext4_ext_replay_update_ex(), the 'ppath' is updated but it is the 'path' that is freed, t free in the following process: ext4_ext_replay_update_ex ppath = path ext4_force_split_extent_at(&ppath) ext4_sp ext4_ext_create_new_leaf ext4_ext_grow_indepth ext4_find_extent if (depth > path[0].p_maxdepth) kfree(path) -- ---> null ppath kfree(path) ---> path double-free !!! So drop the unnecessary ppath and use path directly to avoid th directly to update path, avoiding unnecessary memory allocation and freeing. Also, propagate the error returned by error codes.
CVE-2024-49999	In the Linux kernel, the following vulnerability has been resolved: afs: Fix the setting of the server responding flag transcribe the call responded flag to the server record that we used after doing the fileserver iteration loop - but it's response from the server that we've discarded (e.g. it returned an abort or we started receiving data, but the call did might be NULL, but we don't check that before attempting to set the server flag.
CVE-2024-50002	In the Linux kernel, the following vulnerability has been resolved: static_call: Handle module init failure correctly invokes static_call_add_module() to initialize the static calls in a module. static_call_add_module() invokes __stati static_call_mod to either encapsulate the built-in static call sites of the associated key into it so further modules can module chain. If that allocation fails the function returns with an error code and the module core invokes static_cal added static_call_mod entries. This works correctly, when all keys used by the module were converted over to a m then static_call_del_module() causes a #GP as it blindly assumes that key::mods points to a valid struct static_call_ not a individual struct member of struct static_call_key, it's part of a union to save space: union { /* bit 0: 0 = mods static_call_mod *mods; struct static_call_site *sites; }; key::sites is a pointer to the list of built-in usage sites of the differentiated by bit 0. A mods pointer has the bit clear, the sites pointer has the bit set. As static_call_del_module(static_call_mod type, it fails to check for this failure case and dereferences the pointer to the list of built-in call site checking whether the key has a sites or a mods pointer. If it's a sites pointer then the key is not to be touched. As th __static_call_init() the site walk can be terminated because all subsequent sites have not been touched by the init c before the allocation fail, then the inner loop which searches for a module match will find nothing. A fail in the sec harmless and does not require special treatment. The first allocation succeeded and converted the key to a module NULL and mod::next == NULL, so the inner loop of static_call_del_module() will neither find a module match no was either already converted, but can't match the module, or it will exit the outer loop because it has a static_call_s
CVE-2024-50013	In the Linux kernel, the following vulnerability has been resolved: exfat: fix memory leak in exfat_load_bitmap() I is not a bitmap directory entry, 'bh' will not be released and reassigned, which will cause a memory leak.

CVE-2024-50020	In the Linux kernel, the following vulnerability has been resolved: ice: Fix improper handling of refcount in ice_sri addresses an issue with improper reference count handling in the ice_sriov_set_msix_vec_count() function. First, t increments the reference count of the vf pointer. If the subsequent call to ice_get_vf_vsi() fails, the function curren reference count of the vf pointer, leading to a reference count leak. The correct behavior, as implemented in this pa ice_put_vf(vf) before returning an error when vsi is NULL. Second, the function calls ice_sriov_get_irqs(), which a negative value, indicating an error, the function returns an error without decrementing the reference count of the count leak. The patch addresses this by adding a call to ice_put_vf(vf) before returning an error when vf->first_vec experimental static analysis tool developed by our team. The tool specializes in analyzing reference count operatio of reference counts. In this case, the tool flagged the missing decrement operation as a potential issue, leading to th
CVE-2024-50021	In the Linux kernel, the following vulnerability has been resolved: ice: Fix improper handling of refcount in ice_dp a reference count handling issue in the ice_dpll_init_rclk_pins() function. The function calls ice_dpll_get_pins(), w the relevant resources. However, if the condition WARN_ON((!vsi || !vsi->netdev)) is met, the function currently r the resources acquired by ice_dpll_get_pins(), leading to a reference count leak. To resolve this, the check has beer ensures that the function verifies the state before any resources are acquired, avoiding the need for additional resou bug was identified by an experimental static analysis tool developed by our team. The tool specializes in analyzing potential issues where resources are not properly managed. In this case, the tool flagged the missing release operati development of this patch.
CVE-2024-50022	In the Linux kernel, the following vulnerability has been resolved: device-dax: correct pgoff align in dax_set_mapp ALIGN_DOWN() instead of ALIGN(). Otherwise, vmf->address not aligned to fault_size will be aligned to the ne failure getting the wrong address. It's a subtle situation that only can be observed in page_mapped_in_vma() after t dev_dax_huge_fault. Generally, there is little chance to perform page_mapped_in_vma in dev-dax's page unless in to trigger an MCE - memory-failure. In that case, page_mapped_in_vma() will be triggered to determine which tas that task in the end. We used self-developed dax device (which is 2M aligned mapping) , to perform error injection injected to non-2M-aligned address was causing endless MCE until panic. Because page_mapped_in_vma() kept re the failure address was never killed properly: [3783.719419] Memory failure: 0x200c9742: recovery action for da Uncorrected hardware memory error in user-access at 200c9742380 [3784.049190] Memory failure: 0x200c9742: [3784.448042] mce: Uncorrected hardware memory error in user-access at 200c9742380 [3784.448186] Memory dax page: Recovered [3784.792026] mce: Uncorrected hardware memory error in user-access at 200c9742380 [3 recovery action for dax page: Recovered [3785.162502] mce: Uncorrected hardware memory error in user-access failure: 0x200c9742: recovery action for dax page: Recovered [3785.461116] mce: Uncorrected hardware memory [3785.461247] Memory failure: 0x200c9742: recovery action for dax page: Recovered [3785.764730] mce: Unco access at 200c9742380 [3785.764859] Memory failure: 0x200c9742: recovery action for dax page: Recovered [3 memory error in user-access at 200c9742380 [3786.042259] Memory failure: 0x200c9742: recovery action for da Uncorrected hardware memory error in user-access at 200c9742380 [3786.464423] Memory failure: 0x200c9742: [3786.818090] mce: Uncorrected hardware memory error in user-access at 200c9742380 [3786.818217] Memory dax page: Recovered [3787.085297] mce: Uncorrected hardware memory error in user-access at 200c9742380 [3 recovery action for dax page: Recovered It took us several weeks to pinpoint this problem,¬† but we eventually us address and successfully identified the issue. Joao added: ; Likely we never reproduce in production because we al align they provide (Qemu does : similarly with prealloc in hugetlb/file backed memory). I think this : bug requires unaligned to : the device-dax selected alignment (page size i.e. 4K/2M/1G)
CVE-2024-50023	In the Linux kernel, the following vulnerability has been resolved: net: phy: Remove LED entry from LEDs list on Manual remove LEDs to ensure correct ordering") correctly fixed a problem with using devm_ but missed removin cause kernel panic on specific scenario where the port for the PHY is torn down and up and the kmod for the PHY time, the assosiacted LEDs are correctly unregistered. The associated kmod for the PHY is now removed. The kmo up, the associated LED are registered again. On putting the port down again for the second time after these step, th first 2 already unregistered previously and the 2 new one registered again. This cause a kernel panic as the first 2 el correctly removing the element when LED is unregistered.
CVE-2024-50027	In the Linux kernel, the following vulnerability has been resolved: thermal: core: Free tzp copy along with the ther may still be accessed after being freed in thermal_zone_device_unregister(), so move the freeing of it to the point a completed at which it cannot be accessed any more.
CVE-2024-50029	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: hci_conn: Fix UAF in hci_enhanced_ connection remains valid as it could be destroyed while hci_enhanced_setup_sync is pending on cmd_sync leading use-after-free in hci_enhanced_setup_sync+0x91b/0xa60 Read of size 1 at addr ffff888002328ffd by task kworker/ kworker/u5:2 Not tainted 6.11.0-rc6-01300-g810be445d8d6 #7099 Hardware name: QEMU Standard PC (Q35 + I Workqueue: hci0 hci_cmd_sync_work Call Trace: <TASK> dump_stack_lvl+0x5d/0x80 ? hci_enhanced_setup_sy hci_enhanced_setup_sync+0x91b/0xa60 ? __virt_addr_valid+0x1fa/0x420 ? hci_enhanced_setup_sync+0x91b/0xa hci_enhanced_setup_sync+0x91b/0xa60 hci_enhanced_setup_sync+0x91b/0xa60 ? __pfx_hci_enhanced_setup_sy +0x10/0x10 hci_cmd_sync_work+0x1c2/0x330 process_one_work+0x7d9/0x1360 ? __pfx_lock_acquire+0x10/0x assign_work+0x167/0x240 worker_thread+0x5b7/0xf60 ? __kthread_parkme+0xac/0x1c0 ? __pfx_worker_thread +0x10/0x10 kthread+0x293/0x360 ? __pfx_kthread+0x10/0x10 ret_from_fork+0x2f/0x70 ? __pfx_kthread+0x10/ TASK> Allocated by task 34: kasan_save_stack+0x30/0x50 kasan_save_track+0x14/0x30 __kasan_kmalloc+0x8f hci_connect_sco+0x2e1/0xb90 sco_sock_connect+0x2a2/0xb80 __sys_connect+0x227/0x2a0 __x64_sys_connect entry_SYSCALL_64_after_hwframe+0x76/0x7e Freed by task 37: kasan_save_stack+0x30/0x50 kasan_save_trac +0x3b/0x60 __kasan_slab_free+0x101/0x160 kfree+0xd0/0x250 device_release+0x9a/0x210 kobject_put+0x151/ hci_abort_conn_sync+0x46f/0x980 hci_cmd_sync_work+0x1c2/0x330 process_one_work+0x7d9/0x1360 worker_ ret_from_fork+0x2f/0x70 ret_from_fork_asm+0x1a/0x30

CVE-2024-50042	In the Linux kernel, the following vulnerability has been resolved: ice: Fix increasing MSI-X on VF Increasing MS operations. This is caused by not reallocating some arrays. Reproducer: modprobe ice echo 0 > /sys/bus/pci/devices > /sys/bus/pci/devices/$PF_PCI/sriov_numvfs echo 17 > /sys/bus/pci/devices/$VF0_PCI/sriov_vf_msix_count Def this issue. KASAN reports: BUG: KASAN: slab-out-of-bounds in ice_vsi_alloc_ring_stats+0x38d/0x4b0 [ice] Rea task bash/28433 (...) Call Trace: (...) ? ice_vsi_alloc_ring_stats+0x38d/0x4b0 [ice] kasan_report+0xed/0x120 ? ice_ ice_vsi_alloc_ring_stats+0x38d/0x4b0 [ice] ice_vsi_cfg_def+0x3360/0x4770 [ice] ? mutex_unlock+0x83/0xd0 ? __pfx_ice_remove_vsi_lkup_fltr+0x10/0x10 [ice] ice_vsi_cfg+0x7f/0x3b0 [ice] ice_vf_reconfig_vsi+0x114/0x21(+0x3d0/0x960 [ice] sriov_vf_msix_count_store+0x21c/0x300 (...) Allocated by task 28201: (...) ice_vsi_cfg_def+([ice] ice_vsi_setup+0x179/0xa30 [ice] ice_sriov_configure+0xcaa/0x1520 [ice] sriov_numvfs_store+0x212/0x390 instead of ice_vf_reconfig_vsi(). This causes the required arrays to be reallocated taking the new queue count into Set req_txq and req_rxq before ice_vsi_rebuild(), so that realloc uses the newly set queue count. Additionally, ice_ (ice_fltr_remove_all()), so ice_vf_init_host_cfg() is no longer necessary.
CVE-2024-50046	In the Linux kernel, the following vulnerability has been resolved: NFSv4: Prevent NULL-pointer dereference in n node of an NFS client, some files saved in the mountpoint of the NFS server were copied to another location of the nfs42_complete_copies() got a NULL-pointer dereference crash with the following syslog: [232064.838881] NFSv nfs/pvc-12b5200d-cd0f-46a3-b9f0-af8f4fe0ef64.qcow2, error = -116 [232064.839360] NFSv4: state recovery faile cd0f-46a3-b9f0-af8f4fe0ef64.qcow2, error = -116 [232066.588183] Unable to handle kernel NULL pointer derefer [232066.588586] Mem abort info: [232066.588701] ESR = 0x0000000096000007 [232066.588862] EC = 0x25: D [232066.589084] SET = 0, FnV = 0 [232066.589216] EA = 0, S1PTW = 0 [232066.589340] FSC = 0x07: level 3 t Data abort info: [232066.589683] ISV = 0, ISS = 0x00000007 [232066.589842] CM = 0, WnR = 0 [232066.58996 VAs, pgdp=00002000956ff400 [232066.590231] [0000000000000058] pgd=08001100ae100003, p4d=08001100a pmd=08001100b3c00003, pte=0000000000000000 [232066.590757] Internal error: Oops: 96000007 [#1] SMP [23 rpcsec_gss_krb5 auth_rpcgss nfsv4 dns_resolver nfs lockd grace fscache netfs ocfs2_dlmfs ocfs2_stack_o2cb ocfs2 tun ipt_rpfilter xt_multiport ip_set_hash_ip ip_set_hash_net xfrm_interface xfrm6_tunnel tunnel4 tunnel6 esp4 ah4 veth xt_addrtype xt_set nf_conntrack_netlink ip_set_hash_ipportnet ip_set_hash_ipportip ip_set_bitmap_port ip_s ip_vs_wrr ip_vs_rr ip_vs iptable_filter sch_ingress nfnetlink_cttimeout vport_gre ip_gre ip_tunnel gre vport_genev udp_tunnel openvswitch nf_conncount dm_round_robin dm_service_time dm_multipath xt_nat xt_MASQUERAD xt_comment nft_compat nft_counter nf_tables nfnetlink ocfs2 ocfs2_nodemanager ocfs2_stackglue iscsi_tcp libisc nbd overlay 8021q garp mrp bonding tls rfkill sunrpc ext4 mbcache jbd2 [232066.591052] vfat fat cas_cache cas_c acpi_ipmi ipmi_si ipmi_devintf ipmi_msghandler ip_tables vfio_pci vfio_pci_core vfio_virqfd vfio_iommu_type1 dm_mod nf_conntrack nf_defrag_ipv6 nf_defrag_ipv4 br_netfilter bridge stp llc fuse xfs libcrc32c ast drm_vram_h crct10dif_ce sysfillrect ghash_ce sysimgblt sha2_ce fb_sys_fops cec sha256_arm64 sha1_ce drm_ttm_helper ttm n nvme_core i2c_algo_bit i40e scsi_transport_fc megaraid_sas aes_neon_bs [232066.596953] CPU: 6 PID: 4124696 Not tainted 5.15.131-9.cl9_ocfs2.aarch64 #1 [232066.597356] Hardware name: Great Wall .\x93\x8e...RF6260 V5 2024-01-06 [232066.597721] pstate: 20400009 (nzCv daif +PAN -UAO -TCO -DIT -SSBS BTYPE=--) [232066.5 +0x220/0x800 [nfsv4] [232066.598327] lr : nfs4_reclaim_open_state+0x12c/0x800 [nfsv4] [232066.598595] sp : f ffff8000f568fc70 x28: 0000000000001000 x27: ffff21003db33000 [232066.599030] x26: ffff800005521ae0 x25: f [232066.599319] x23: ffff800009920008 x22: ffff21003db33040 x21: ffff21003db33050 [232066.599628] x20: ff 0000000000000000 [232066.599914] x17: 0000000000000000 x16: 0000000000000004 x15: 0000000000000000 x13: ffff800008e685a8 x12: 00000000eac0c6e6 [232066.600498] x11: 00000000000000 ---truncated---
CVE-2024-50064	In the Linux kernel, the following vulnerability has been resolved: zram: free secondary algorithms names We nee when reset zram device that had multi-streams, otherwise we leak memory. [senozhatsky@chromium.org: kfree(N r/20240917013021.868769-1-senozhatsky@chromium.org
CVE-2024-50075	In the Linux kernel, the following vulnerability has been resolved: xhci: tegra: fix checked USB2 port number If U are shared between all Virtual Functions. The USB2 port number owned by an USB2 root hub in a Virtual Functio supported by the Tegra XUSB controller. Using total USB2 phy number as port number to check all PORTSC valu [116.923438] Unable to handle kernel paging request at virtual address 006c622f7665642f ... [117.213640] Call t +0x23c/0x658 [117.222021] tegra_xusb_runtime_suspend+0x40/0x68 [117.227260] pm_generic_runtime_suspe +0x84/0x3c0 [117.237038] rpm_suspend+0x2dc/0x740 [117.241229] pm_runtime_work+0xa0/0xb8 [117.24576 [117.251007] worker_thread+0x23c/0x328 [117.255547] kthread+0x104/0x1b0 [117.259389] ret_from_fork+0x f9461ae8 f8747908 b4ffff48 (f9400100)
CVE-2024-50076	In the Linux kernel, the following vulnerability has been resolved: vt: prevent kernel-infoleak in con_font_get() fo depending on the implementation of vc->vc_sw->con_font_get. This may cause info-leak, so to prevent this, it is s memory space to 0, and it generally does not affect the overall performance of the system.
CVE-2024-50077	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: ISO: Fix multiple init when debugfs successfully, which happens if either CONFIG_DEBUG_FS or CONFIG_DEBUG_FS_ALLOW_ALL is unset, th not set iso_inited to true. This means that a subsequent call to iso_init() will result in duplicate calls to proto_regist CONFIG_LIST_HARDENED and CONFIG_BUG_ON_DATA_CORRUPTION enabled, the duplicate call to pro double add: new=ffffffffc0b280d0, prev=ffffffffbab56250, next=ffffffffc0b280d0. ----------[cut here]---------- k invalid opcode: 0000 [#1] PREEMPT SMP PTI CPU: 2 PID: 887 Comm: bluetoothd Not tainted 6.10.11-1-ao-desk +0x9a/0xa0 ... __list_add_valid_or_report+0x9a/0xa0 proto_register+0x2b5/0x340 iso_init+0x23/0x150 [bluetooth [bluetooth] kmem_cache_free+0x308/0x330 hci_sock_sendmsg+0x990/0x9e0 [bluetooth] __sock_sendmsg+0x7b/ do_iter_readv_writev+0x11d/0x220 vfs_writev+0x180/0x3e0 do_writev+0xca/0x100 ... This change removes the e NULL was unnecessary, it is always NULL when iso_inited is false.

CVE-2024-50078	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: Call iso_exit() on module unload If i be called on module unload. Without that, the struct proto that iso_init() registered with proto_register() becomes in problems later. In my case, with CONFIG_LIST_HARDENED and CONFIG_BUG_ON_DATA_CORRUPTION triggers this BUG(): list_add corruption. next->prev should be prev (ffffffffb5355fd0), but was 0000000000000068 here]---------- kernel BUG at lib/list_debug.c:29! Oops: invalid opcode: 0000 [#1] PREEMPT SMP PTI CPU: 1 6.10.11-4+bt2-ao-desktop #1 RIP: 0010:__list_add_valid_or_report+0x61/0xa0 ... __list_add_valid_or_report+0x6 hci_sock_init+0x16/0xc0 [bluetooth] bt_init+0x68/0xd0 [bluetooth] __pfx_bt_init+0x10/0x10 [bluetooth] do_one_ +0x8b/0x230 __do_sys_init_module+0x15f/0x190 do_syscall_64+0x68/0x110 ...
CVE-2024-50081	In the Linux kernel, the following vulnerability has been resolved: blk-mq: setup queue ->tag_set before initializing helper for checking if one CPU is mapped to specified hctx") needs to check queue mapping via tag set in hctx's cp setup yet when the cpuhp handler is enabled, then kernel oops is triggered. Fix the issue by setup queue tag_set bef
CVE-2024-50101	In the Linux kernel, the following vulnerability has been resolved: iommu/vt-d: Fix incorrect pci_for_each_dma_al domain_context_clear() function incorrectly called pci_for_each_dma_alias() to set up context entries for non-PCI other unexpected behavior. Add a check to only call pci_for_each_dma_alias() for PCI devices. For non-PCI devic directly.
CVE-2024-50102	In the Linux kernel, the following vulnerability has been resolved: x86: fix user address masking non-canonical spe has a "Meltdown Lite(tm)" issue with non-canonical accesses in kernel space. And so using just the high bit to deci kernel space ends up with the good old "leak speculative data" if you have the right gadget using the result: CVE-2 Non-Canonical Accesses,Äú Now, the kernel surrounds the access with a STAC/CLAC pair, and those instructions Zen architectures, which closes the speculation window. But that was true only up until Zen 5, which renames the A of STAC/CLAC a lot, but also means that the speculation window is now open. Note that this affects not just the n valid_user_address() check used by access_ok(), and the asm version of the sign bit check in the get_user() helpers variants, since there's no speculative result to be used in a gadget for those operations.
CVE-2024-50107	In the Linux kernel, the following vulnerability has been resolved: platform/x86/intel/pmc: Fix pmc_core_iounmap 50c6dbdfd16e ("x86/ioremap: Improve iounmap() address range checks") introduces a WARN when adrress range P1 Gen 7 (Meteor Lake-P) this caused the following warning to appear: WARNING: CPU: 7 PID: 713 at arch/x86/ Modules linked in: rfkill(+) snd_timer(+) fjes(+) snd soundcore intel_pmc_core(+) int3403_thermal(+) int340x_the acpi_pad pmt_class acpi_tad int3400_thermal acpi_thermal_rel joydev loop nfnetlink zram xe drm_suballoc_helpe gpu_sched drm_gpuvm drm_exec drm_buddy i2c_algo_bit crct10dif_pclmul crc32_pclmul ttm crc32c_intel polyva polyval_generic mmc_core hid_multitouch drm_display_helper ghash_clmulni_intel typec_ucsi nvme sha512_ssse sha1_ssse3 rtsx_pci cec typec nvme_auth i2c_hid_acpi i2c_hid wmi pinctrl_meteorlake serio_raw ip6_tables ip_tab (udev-worker) Not tainted 6.12.0-rc2iounmap+ #42 Hardware name: LENOVO 21KWCTO1WW/21KWCTO1WW RIP: 0010:iounmap+0x58/0x1f0 Code: 85 6a 01 00 00 48 8b 05 e6 e2 28 04 48 39 c5 72 19 eb 26 cc cc cc 48 ba 00 39 c5 72 23 <0f> 0b 48 83 c4 08 5b 5d 41 5c c3 cc cc cc cc 48 ba 00 00 00 00 00 RSP: 0018:ffff888131eff038 EFL RBX: 0000000000000000 RCX: ffff888e33b80000 RDX: dffffc0000000000 RSI: ffff888e33bc29c0 RDI: 000000 R08: ffff8881598a8000 R09: ffff888e2ccedc10 R10: 0000000000000003 R11: ffffffffb3367634 R12: 00000000fe0 ffffffffc2e437e0 R15: ffff888110b03b28 FS: 00007f3c1d4b3980(0000) GS:ffff888e33b80000(0000) knlGS:00000 CR0: 0000000080050033 CR2: 00005651cfc93578 CR3: 0000000124e4c002 CR4: 0000000000f70ef0 DR0: 0000 DR2: 0000000000000000 DR3: 0000000000000000 DR6: 00000000ffff07f0 DR7: 0000000000000400 PKRU: 55 +0xb6/0x176 ? iounmap+0x58/0x1f0 ? report_bug+0x1f4/0x2b0 ? handle_bug+0x58/0x90 ? exc_invalid_op+0x17 iounmap+0x58/0x1f0 pmc_core_ssram_get_pmc+0x477/0x6c0 [intel_pmc_core] ? __pfx_pmc_core_ssram_get_pm __pfx_do_pci_enable_device+0x10/0x10 ? pci_wait_for_pending+0x60/0x110 ? pci_enable_device_flags+0x1e3/0 [intel_pmc_core] pmc_core_ssram_init+0x7f/0x110 [intel_pmc_core] mtl_core_init+0xda/0x130 [intel_pmc_core] pmc_core_probe+0x27e/0x10b0 [intel_pmc_core] ? _raw_spin_lock_irqsave+0x96/0xf0 ? __pfx_pmc_core_probe __pfx_mutex_unlock+0x10/0x10 ? __pfx_mutex_lock+0x10/0x10 ? device_pm_check_callbacks+0x82/0x370 ? ac platform_probe+0x9f/0x150 really_probe+0x1e0/0x8a0 __driver_probe_device+0x18c/0x370 ? __pfx___driver_at +0x4a/0x120 __driver_attach+0x190/0x4a0 ? __pfx___driver_attach+0x10/0x10 bus_for_each_dev+0x103/0x180 klist_add_tail+0x136/0x270 bus_add_driver+0x2fc/0x540 driver_register+0x1a5/0x360 ? __pfx_pmc_core_driver do_one_initcall+0xa4/0x380 ? __pfx_do_one_initcall+0x10/0x10 ? kasan_unpoison+0x44/0x70 do_init_module+0 __pfx_load_module+0x10/0x10 ? ima_post_read_file+0x193/0x200 ? __pfx_ima_post_read_file+0x10/0x10 ? rw_ +0x257/0x750 ? __pfx_kernel_read_file+0x10/0x10 ? __pfx_filemap_get_read_batch+0x10/0x10 ? init_module_f +0xd1/0x130 ? __pfx_init_module_from_file+0x10/0 ---truncated---
CVE-2024-50109	In the Linux kernel, the following vulnerability has been resolved: md/raid10: fix null ptr dereference in raid10_siz raid10_set_queue_limits() succeed, the return value is set to zero, and if following procedures failed raid10_run() w NULL, causing null ptr dereference in raid10_size(). Fix the problem by only overwrite the return value if raid10_s

CVE-2024-50114	In the Linux kernel, the following vulnerability has been resolved: KVM: arm64: Unregister redistributor for failed syzkaller has managed to trigger a use-after-free when tearing down a VM: BUG: KASAN: slab-use-after-free in k kvm_main.c:5769 Read of size 8 at addr ffffff801c6890d0 by task syz.3.2219/10758 CPU: 3 UID: 0 PID: 10758 C dirty #64 Hardware name: linux,dummy-virt (DT) Call trace: dump_backtrace+0x17c/0x1a8 arch/arm64/kernel/sta arm64/kernel/stacktrace.c:324 __dump_stack lib/dump_stack.c:93 [inline] dump_stack_lvl+0x94/0xc0 lib/dump_s kasan/report.c:377 kasan_report+0xcc/0x128 mm/kasan/report.c:601 __asan_report_load8_noabort+0x20/0x2c mn +0x300/0xe68 virt/kvm/kvm_main.c:5769 kvm_vm_release+0x4c/0x60 virt/kvm/kvm_main.c:1409 __fput+0x198 +0x20/0x30 fs/file_table.c:450 task_work_run+0x1cc/0x23c kernel/task_work.c:228 do_notify_resume+0x144/0x1 el0_svc+0x64/0x68 arch/arm64/kernel/entry-common.c:169 el0t_64_sync_handler+0x90/0xfc arch/arm64/kernel/e +0x190/0x194 arch/arm64/kernel/entry.S:598 Upon closer inspection, it appears that we do not properly tear down fails creation late in the game, e.g. a vCPU w/ the same ID already exists in the VM. It is important to consider the by moving the unregistration out of __kvm_vgic_vcpu_destroy(). That change correctly sought to avoid an srcu v. vCPU teardown into two parts, one guarded by the config_lock. Fix the use-after-free while avoiding lock inversio __kvm_vgic_vcpu_destroy(). This is safe because failed vCPUs are torn down outside of the config_lock.
CVE-2024-50120	In the Linux kernel, the following vulnerability has been resolved: smb: client: Handle kstrdup failures for passwor ctx->password and ctx->password2 with kstrdup(), we need to check for allocation failures. If ses->password alloc >password2 allocation fails, free ses->password, set it to NULL, and return -ENOMEM.
CVE-2024-50124	In the Linux kernel, the following vulnerability has been resolved: Bluetooth: ISO: Fix UAF on iso_sock_timeout c while waiting for iso_conn_lock so this checks if the conn->sk is still valid by checking if it part of iso_sk_list.
CVE-2024-50126	In the Linux kernel, the following vulnerability has been resolved: net: sched: use RCU read-side critical section in in 'taprio_dump()' by adding RCU read-side critical section there. Never seen on x86 but found on a KASAN-enabl syzkaller.appspot.com/bug?extid=b65e0af58423fc8a73aa: [T15862] BUG: KASAN: slab-use-after-free in taprio_c 4 at addr ffff0000d4bb88f8 by task repro/15862 [T15862] [T15862] CPU: 0 UID: 0 PID: 15862 Comm: repro Not dirty #2 [T15862] Hardware name: QEMU QEMU Virtual Machine, BIOS edk2-20240524-5.fc40 05/24/2024 [T1 +0x20c/0x220 [T15862] show_stack+0x2c/0x40 [T15862] dump_stack_lvl+0xf8/0x174 [T15862] print_report+0x +0xb8/0x1d4 [T15862] __asan_report_load4_noabort+0x20/0x2c [T15862] taprio_dump+0xa0c/0xbb0 [T15862] t qdisc_notify.isra.0+0x330/0x3a0 [T15862] tc_modify_qdisc+0x7b8/0x1838 [T15862] rtnetlink_rcv_msg+0x3c8/0 [T15862] rtnetlink_rcv+0x28/0x40 [T15862] netlink_unicast+0x51c/0x790 [T15862] netlink_sendmsg+0x79c/0xc [T15862] ____sys_sendmsg+0x6c0/0x840 [T15862] ___sys_sendmsg+0x1ac/0x1f0 [T15862] __sys_sendmsg+0x +0x74/0xb0 [T15862] invoke_syscall+0x88/0x2e0 [T15862] el0_svc_common.constprop.0+0xe4/0x2a0 [T15862] el0_svc+0x50/0x184 [T15862] el0t_64_sync_handler+0x120/0x12c [T15862] el0t_64_sync+0x190/0x194 [T1586 [T15862] kasan_save_stack+0x3c/0x70 [T15862] kasan_save_track+0x20/0x3c [T15862] kasan_save_alloc_info+ +0xd4/0xe0 [T15862] __kmalloc_cache_noprof+0x194/0x334 [T15862] taprio_change+0x45c/0x2fe0 [T15862] tc rtnetlink_rcv_msg+0x3c8/0xc20 [T15862] netlink_rcv_skb+0x1f8/0x3d4 [T15862] rtnetlink_rcv+0x28/0x40 [T15 [T15862] netlink_sendmsg+0x79c/0xc20 [T15862] __sock_sendmsg+0xe0/0x1a0 [T15862] ____sys_sendmsg+0x +0x1ac/0x1f0 [T15862] __sys_sendmsg+0x110/0x1d0 [T15862] __arm64_sys_sendmsg+0x74/0xb0 [T15862] inv el0_svc_common.constprop.0+0xe4/0x2a0 [T15862] do_el0_svc+0x44/0x60 [T15862] el0_svc+0x50/0x184 [T158 [T15862] el0t_64_sync+0x190/0x194 [T15862] [T15862] Freed by task 6192: [T15862] kasan_save_stack+0x3c/0 [T15862] kasan_save_free_info+0x4c/0x80 [T15862] poison_slab_object+0x110/0x160 [T15862] __kasan_slab_fr [T15862] taprio_free_sched_cb+0x18c/0x220 [T15862] rcu_core+0x920/0x1b7c [T15862] rcu_core_si+0x10/0x1c [T15862] __do_softirq+0x14/0x20
CVE-2024-50128	In the Linux kernel, the following vulnerability has been resolved: net: wwan: fix global oob in wwan_rtnl_policy wwan_rtnl_link_ops assign a *bigger* maxtype which leads to a global out-of-bounds read when parsing the netlin Exactly same bug cause as the oob fixed in commit b33fb5b801c6 ("net: qualcomm: rmnet: fix global oob in rmnet ========== BUG: KASAN: globa nlattr.c:388 [inline] BUG: KASAN: global-out-of-bounds in __nla_validate_parse+0x19d7/0x29a0 lib/nlattr.c:603 by task syz.1.66276/323862 CPU: 0 PID: 323862 Comm: syz.1.66276 Not tainted 6.1.70 #1 Hardware name: QEM BIOS 1.13.0-1ubuntu1.1 04/01/2014 Call Trace: <TASK> __dump_stack lib/dump_stack.c:88 [inline] dump_stack print_address_description mm/kasan/report.c:284 [inline] print_report+0x14f/0x750 mm/kasan/report.c:395 kasan_ validate_nla lib/nlattr.c:388 [inline] __nla_validate_parse+0x19d7/0x29a0 lib/nlattr.c:603 __nla_parse+0x3c/0x50 include/net/netlink.h:1269 [inline] __rtnl_newlink net/core/rtnetlink.c:3514 [inline] rtnl_newlink+0x7bc/0x1fd0 ne +0x794/0xef0 net/core/rtnetlink.c:6122 netlink_rcv_skb+0x1de/0x420 net/netlink/af_netlink.c:2508 netlink_unicas [inline] netlink_unicast+0x74b/0x8c0 net/netlink/af_netlink.c:1352 netlink_sendmsg+0x882/0xb90 net/netlink/af_ net/socket.c:716 [inline] __sock_sendmsg net/socket.c:728 [inline] ____sys_sendmsg+0x5cc/0x8f0 net/socket.c:24 socket.c:2553 __sys_sendmsg net/socket.c:2582 [inline] __do_sys_sendmsg net/socket.c:2591 [inline] __se_sys_se do_syscall_x64 arch/x86/entry/common.c:51 [inline] do_syscall_64+0x45/0x90 arch/x86/entry/common.c:81 entry +0x63/0xcd RIP: 0033:0x7f67b19a24ad RSP: 002b:00007f67b17febb8 EFLAGS: 00000246 ORIG_RAX: 000000 00007f67b1b45f80 RCX: 00007f67b19a24ad RDX: 0000000000000000 RSI: 0000000020005e40 RDI: 00000000 R08: 0000000000000000 R09: 0000000000000000 R10: 0000000000000000 R11: 0000000000000246 R12: 0000 R14: 00007ffd251376e0 R15: 00007f67b17fed40 </TASK> The buggy address belongs to the variable: wwan_rtnl address belongs to the physical page: page:ffffea00002c2700 refcount:1 mapcount:0 mapping:0000000000000000 0xfff00000001000(reserved|node=0|zone=1|lastcpupid=0x7ff) raw: 00fff00000001000 ffffea00002c2708 ffffea000 0000000000000000 0000000000000000 00000001ffffffff 0000000000000000 page dumped because: kasan: bad a (never set?) Memory state around the buggy address: ffffffff8b09ca00: 05 f9 f9 f9 05 f9 f9 f9 00 01 f9 f9 00 01 f9 f 00 00 03 f9 f9 f9 f9 f9 >ffffffff8b09cb00: 00 00 00 00 05 f9 f9 f9 00 00 00 00 f9 f9 f9 f9 ^ ffffffff8b09cb80: 00 00 ========== According to the com correct size `IFLA_WWAN_MAX` here to fix this issue.

CVE-2024-50130	In the Linux kernel, the following vulnerability has been resolved: netfilter: bpf: must hold reference on net namesp free in __nf_unregister_net_hook+0x640/0x6b0 Read of size 8 at addr ffff8880106fe400 by task repro/72= bpf_nf_ +0x139/0x2d0 bpf_link_release+0x68/0x80 __fput+0x414/0xb60 Eric says: It seems that bpf was able to defer the close() time. Perhaps a netns reference is missing, because the netns has been dismantled/freed already. bpf_nf_link not see a reference being taken on net. Add such a reference and release it after hook unreg. Note that I was unable know if this resolves this splat.
CVE-2024-50139	In the Linux kernel, the following vulnerability has been resolved: KVM: arm64: Fix shift-out-of-bounds bug Fix a UBSAN when running VM with MTE enabled host kernel. UBSAN: shift-out-of-bounds in arch/arm64/kvm/sys_r too large for 32-bit type 'int' CPU: 26 UID: 0 PID: 7629 Comm: qemu-kvm Not tainted 6.12.0-rc2 #34 Hardware n BIOS 00.00. 2024-10-12 09:28:54 10/14/2024 Call trace: dump_backtrace+0xa0/0x128 show_stack+0x20/0x38 du +0x18/0x28 __ubsan_handle_shift_out_of_bounds+0xf8/0x1e0 reset_clidr+0x10c/0x1c8 kvm_reset_sys_regs+0x5 __kvm_vcpu_set_target+0x84/0x158 kvm_vcpu_set_target+0x138/0x168 kvm_arch_vcpu_ioctl_vcpu_init+0x40/0 kvm_vcpu_ioctl+0x4bc/0x7a8 __arm64_sys_ioctl+0xb4/0x100 invoke_syscall+0x70/0x100 el0_svc_common.con el0_svc+0x3c/0x158 el0t_64_sync_handler+0x120/0x130 el0t_64_sync+0x194/0x198
CVE-2024-50145	In the Linux kernel, the following vulnerability has been resolved: octeon_ep: Add SKB allocation failures handlin returns NULL in case of a memory allocation failure so handle it inside __octep_oq_process_rx() to avoid NULL p is called during NAPI polling by the driver. If skb allocation fails, keep on pulling packets out of the Rx DMA que immediately and thus falsely indicate to the octep_napi_poll() that the Rx pressure is going down. As there is no as packets and don't push them up the network stack - they are skipped. Helper function is implemented to unmmap/fl dropped packet. 'alloc_failures' counter is incremented to mark the skb allocation error in driver statistics. Found by with SVACE.
CVE-2024-50152	In the Linux kernel, the following vulnerability has been resolved: smb: client: fix possible double free in smb2_se warningÔºö fs/smb/client/smb2ops.c:1304:2: Attempt to free released memory. 1304 | kfree(ea); | ^~~~~~~~~ The initialized to NULL' -> 'first successful memory allocation for ea' -> 'something failed, goto sea_exit' -> 'first memo 'second goto sea_exit before allocate memory for ea' -> 'second memory release for ea resulted in double free'. Re-i label, it can fix this double free problem.
CVE-2024-50155	In the Linux kernel, the following vulnerability has been resolved: netdevsim: use cond_resched() in nsim_dev_tra syzbot reports hinting that syzbot might fool nsim_dev_trap_report_work() with hundreds of ports [1] Lets use con instead of implicit system_wq. [1] INFO: task syz-executor:20633 blocked for more than 143 seconds. Not tainted #0 "echo 0 > /proc/sys/kernel/hung_task_timeout_secs" disables this message. task:syz-executor state:D stack:2585 flags:0x00004006 ... NMI backtrace for cpu 1 CPU: 1 UID: 0 PID: 16760 Comm: kworker/1:0 Not tainted 6.12.0-r Hardware name: Google Google Compute Engine/Google Compute Engine, BIOS Google 09/13/2024 Workqueue 0010:__sanitizer_cov_trace_pc+0x0/0x70 kernel/kcov.c:210 Code: 89 fb e8 23 00 00 00 48 8b 3d 04 fb 9c 0c 48 8 90 90 90 90 90 90 90 90 90 90 90 <f3> 0f 1e fa 48 8b 04 24 65 48 8b 0c 25 c0 d7 03 00 65 8b 15 60 f0 RSP: 00 RAX: 0000000000000100 RBX: ffffc90000a188e0 RCX: ffff888027d3bc00 RDX: ffff888027d3bc00 RSI: 00000 ffff88804a2e6000 R08: ffffffff8a4bc495 R09: ffffffff89da3577 R10: 0000000000000004 R11: ffffffff8a4bc2b0 R1 R14: dffffc0000000000 R15: ffff8880663cca00 FS: 0000000000000000(0000) GS:ffff8880b8700000(0000) knlGS 0000 ES: 0000 CR0: 0000000080050033 CR2: 00007fc90a747f98 CR3: 000000000e734000 CR4: 0000000000352 000000000000002b DR2: 0000000000000000 DR3: 0000000000000000 DR6: 00000000ffff0ff0 DR7: 0000000000 <TASK> __local_bh_enable_ip+0x1bb/0x200 kernel/softirq.c:382 spin_unlock_bh include/linux/spinlock.h:396 [i netdevsim/dev.c:820 [inline] nsim_dev_trap_report_work+0x75d/0xaa0 drivers/net/netdevsim/dev.c:850 process_c process_scheduled_works+0xa63/0x1850 kernel/workqueue.c:3310 worker_thread+0x870/0xd30 kernel/workqueu kthread.c:389 ret_from_fork+0x4b/0x80 arch/x86/kernel/process.c:147 ret_from_fork_asm+0x1a/0x30 arch/x86/en
CVE-2024-50157	In the Linux kernel, the following vulnerability has been resolved: RDMA/bnxt_re: Avoid CPU lockups due fifo oc for the fifo occupancy to go below a threshold as soon as the pacing interrupt is received. This can cause soft locku is very high. Add a loop count for FPGA and exit the __wait_for_fifo_occupancy_below_th if the loop is taking m the occupancy is below the threshold. This is ensured by the checks in bnxt_re_pacing_timer_exp and further sched occupancy.
CVE-2024-50158	In the Linux kernel, the following vulnerability has been resolved: RDMA/bnxt_re: Fix out of bound check Driver adapters. But while parsing the pacing stats, driver has a check for "rdev->dbr_pacing". This caused a trace when K of-bounds in bnxt_re_get_hw_stats+0x2b6a/0x2e00 [bnxt_re] Write of size 8 at addr ffff8885942a6340 by task mo
CVE-2024-50164	In the Linux kernel, the following vulnerability has been resolved: bpf: Fix overloading of MEM_UNINIT's meani verifier where check_mem_size_reg() has the following code: if (!tnum_is_const(reg->var_off)) /* For unprivilege so that the program is required to * initialize all the memory that the helper could * just partially fill up. */ meta = N checked when the register containing the size of the passed buffer has not a fixed size. Through this bug, a BPF pro as read-only, for example, .rodata global maps. The problem is that MEM_UNINIT's initial meaning that "the passe to be initialized" which was added back in commit 435faee1aae9 ("bpf, verifier: add ARG_PTR_TO_RAW_STAC "the passed buffer is being written to". The problem however is that checks such as the above which were added la access to variable memory") set meta to NULL in order force the user to always initialize the passed buffer to the h MEM_UNINIT, this bypasses verifier write checks to the memory (not boundary checks though) and only assumes by reverting MEM_UNINIT back to its original meaning, and having MEM_WRITE as an annotation to BPF helpe checks for writing to memory. Some notes: check_arg_pair_ok() ensures that for ARG_CONST_SIZE{,_OR_ZER it must contain a preceding ARG_PTR_TO_MEM. For check_mem_reg() the meta argument can be removed altog and BPF_WRITE. Same for the equivalent check_kfunc_mem_size_reg().

CVE-2024-50165	In the Linux kernel, the following vulnerability has been resolved: bpf: Preserve param->string when parsing mour keep the value of param->string intact so it can be freed later. Otherwise, the kmalloc area pointed to by param->st unreferenced object 0xffff888118c46d20 (size 8): comm "new_name", pid 12109, jiffies 4295580214 hex dump (fi \~ backtrace (crc e1b7f876): [<00000000c6848ac7>] kmemleak_alloc+0x4b/0x80 [<00000000de9f7d00>] __kmal [<000000003e29b886>] memdup_user+0x32/0xa0 [<0000000007248326>] strndup_user+0x46/0x60 [<00000000 +0x368/0x3d0 [<0000000018657927>] x64_sys_call+0xff/0x9f0 [<00000000c0cabc95>] do_syscall_64+0x3b/0xc entry_SYSCALL_64_after_hwframe+0x4b/0x53
CVE-2024-50169	In the Linux kernel, the following vulnerability has been resolved: vsock: Update rx_bytes on read_skb() Make sur virtio_transport_dec_rx_pkt() calls are balanced (i.e. virtio_vsock_sock::rx_bytes doesn't lie) after vsock_transport peer that we've freed up space and it has more credit. Failing to update rx_bytes after packet is dequeued leads to a [233.396654] rx_queue is empty, but rx_bytes is non-zero [233.396702] WARNING: CPU: 11 PID: 40601 at net/
CVE-2024-50172	In the Linux kernel, the following vulnerability has been resolved: RDMA/bnxt_re: Fix a possible memory leak In bnxt_qplib_map_db_bar() fails driver is not freeing the memory allocated for "rdev->chip_ctx".
CVE-2024-50182	In the Linux kernel, the following vulnerability has been resolved: secretmem: disable memfd_secret() if arch cann memfd_secret() syscall if !can_set_direct_map(). This is the case for example on some arm64 configurations, wher present can only be done if the direct map is set up at 4k granularity in the first place (as ARM's break-before-make apart large/gigantic pages). More precisely, on arm64 systems with !can_set_direct_map(), set_direct_map_invalid success (0) instead of an error. This means that memfd_secret will seemingly "work" (e.g. syscall succeeds, you ca it does not actually achieve its goal of removing its memory from the direct map. Note that with this patch, memfd_ where can_set_direct_map() returns false (arm64 with CONFIG_RODATA_FULL_DEFAULT_ENABLED=n, CO CONFIG_KFENCE=n), but that still seems better than the current silent failure. Since CONFIG_RODATA_FULL arm64 systems actually have a working memfd_secret() and aren't be affected. From going through the iterations o seems that disabling the syscall in these scenarios was the intended behavior [1] (preferred over having set_direct_ would result in SIGBUSes at page-fault time), however the check for it got dropped between v16 [2] and v17 [3], v allocations. [1]: https://lore.kernel.org/lkml/20201124164930.GK8537@kernel.org/ [2]: https://lore.kernel.org/lkm #t [3]: https://lore.kernel.org/lkml/20201125092208.12544-10-rppt@kernel.org/
CVE-2024-50197	In the Linux kernel, the following vulnerability has been resolved: pinctrl: intel: platform: fix error path in device_ device_for_each_child_node() loop requires calls to fwnode_handle_put() upon early returns to decrement the refc memory if that error path is triggered. There is one early returns within that loop in intel_platform_pinctrl_prepare_ missing. Instead of adding the missing call, the scoped version of the loop can be used to simplify the code and avo are added, as the child node is only used for parsing, and it is never assigned.
CVE-2024-50200	In the Linux kernel, the following vulnerability has been resolved: maple_tree: correct tree corruption on spanning corruption on spanning store", v3. There has been a nasty yet subtle maple tree corruption bug that appears to have the algorithm. This bug seems far more likely to happen since commit f8d112a4e657 ("mm/mmap: avoid zeroing v point at which reports started to be submitted concerning this bug. We were made definitely aware of the bug thank helped enormously in my being able to track this down and identify the cause of it. The bug arises when an attempt two leaf nodes, where the right leaf node is the rightmost child of the shared parent, AND the store completely con mas_wr_spanning_store() mitakenly duplicating the new and existing entries at the maximum pivot within the rang patch corrects this by detecting this scenario and disallowing the mistaken duplicate copy. The fix patch commit m occurs. This series also includes a test which reliably reproduces the issue, and asserts that the fix works correctly. it resolved his issues. Also Mikhail Gavrilov kindly reported what appears to be precisely the same bug, which this There has been a subtle bug present in the maple tree implementation from its inception. This arises from how store overwrite overlapping ranges and adjust the tree as necessary to accommodate this. A range may always ultimately the two leaf nodes, determine which elements are not overwritten to the left and to the right of the start and end of t tree to contain these entries and the newly inserted one. This kind of store is dubbed a 'spanning store' and is implem to reach this stage, mas_store_gfp() invokes mas_wr_preallocate(), mas_wr_store_type() and mas_wr_walk() in tur to traverse to the location where the write should be performed, determining its store type. When a spanning store i at the parent node which contains the target range, and mas_wr_store_type() marks the mas->store_type as wr_spa go to perform the store in mas_wr_spanning_store(), we first determine the elements AFTER the END of the range entry to be inserted) - we do this by walking to the NEXT pivot in the tree (i.e. r_mas.last + 1), starting at the node over which we intend to write. We then turn our attention to the entries to the left of the entry we are inserting, who these into a 'big node', which is a special node which contains enough slots to contain two leaf node's worth of data immediately after this - the copy and the insertion of the new entry is performed by mas_store_b_node(). After this of the range which we are inserting, if we have not exceeded the length of the node (i.e. r_mas.offset <= r_mas.end circumstances, this logic can break and corrupt the maple tree. Consider the following tree: Height 0 Root Node / \ truncated---
CVE-2024-50215	In the Linux kernel, the following vulnerability has been resolved: nvmet-auth: assign dh_key to NULL after kfree_ across multiple calls to nvmet_setup_dhgroup() for the same controller. So it's better to nullify it after release on er nvmet_destroy_auth(). Found by Linux Verification Center (linuxtesting.org) with Svace.
CVE-2024-50216	In the Linux kernel, the following vulnerability has been resolved: xfs: fix finding a last resort AG in xfs_filestream xfs_filestream_pick_ag fails to find a suitable AG it tries to just pick the online AG. But the loop for that uses args- expects pag to be set. Fix this by reusing the max_pag case for this last resort, and also add a check for impossible uninitialized pag doesn't even escape in theory.

CVE-2024-50223	In the Linux kernel, the following vulnerability has been resolved: sched/numa: Fix the potential null pointer deref stress-ng-vm-segv test, we found a null pointer dereference error in task_numa_work(). Here is the backtrace: [323 pointer dereference at virtual address 0000000000000020 ..... [323676.067108] CPU: 35 PID: 2694524 Comm: str 23401009 (nzCv daif +PAN -UAO +TCO +DIT +SSBS BTYPE=--) [323676.067115] pc : vma_migratable+0x1c/ +0x1ec/0x4e0 [323676.067127] sp : ffff8000ada73d20 [323676.067128] x29: ffff8000ada73d20 x28: 0000000000 [323676.067130] x26: 0000000000080000 x25: ffff800081b5c0d8 x24: ffff800081b27000 [323676.067133] x23: 0 x21: ffff0009f7158000 [323676.067135] x20: 0000000000000000 x19: 0000000000000000 x18: ffff8000ada73db x16: ffff800080df40b0 x15: 0000000000000035 [323676.067140] x14: ffff8000ada73cc8 x13: 1fffe0017cc72001 x11: ffff80008001160c x10: ffff000be639000c x9 : ffff8000800f4ba4 [323676.067145] x8 : ffff000810375000 x7 [323676.067147] x5 : 0068000b33e26707 x4 : 0000000000000001 x3 : ffff0009f7158000 [323676.067149] x2 : 00 x0 : 0000000000000000 [323676.067152] Call trace: [323676.067153] vma_migratable+0x1c/0xd0 [323676.0671 [323676.067157] task_work_run+0x78/0xd8 [323676.067161] do_notify_resume+0x1ec/0x290 [323676.067163] el0t_64_sync_handler+0xf8/0x128 [323676.067170] el0t_64_sync+0x17c/0x180 [323676.067173] Code: d288800 [323676.067177] SMP: stopping secondary CPUs [323676.070184] Starting crashdump kernel... stress-ng-vm-segv SIGSEGV error handling function of the system, which tries to cause a SIGSEGV error on return from unmapping process. Normally this program will not cause kernel crashes. But before the munmap system call returns to user m balancing could be added and executed. In this scenario, since the child process has no vma after munmap, the vma null pointer even if the vma iterator restarts from 0. Recheck the vma pointer before dereferencing it in task_numa_
CVE-2024-50235	In the Linux kernel, the following vulnerability has been resolved: wifi: cfg80211: clear wdev->cqm_config pointe when unregistering, we also need to clear out the pointer since the same wdev/netdev may get re-registered in anot running this code again, which results in a double-free.
CVE-2024-50250	In the Linux kernel, the following vulnerability has been resolved: fsdax: dax_unshare_iter needs to copy entire blo to iomap in dax_unshare_iter is very very broken, which bfoster's recent fsx changes have exposed. If the pos and l to an fsblock boundary, the iter pos and length in the _iter function will reflect this unalignment. dax_iomap_direct start of the kmapped fsdax page, even if its pos argument is in the middle of that page. This is catastrophic for data a page, because daddr/saddr do not point to the same byte in the file as iter->pos. Hence we corrupt user data by co iomap_length() in the _iter function not aligned to a page, then we fail to copy a full block, and only partially popu for data confidentiality because we expose stale pmem contents. Fix both of these issues by aligning copy_pos/cop fsdax so 1 fsblock == 1 base page) so that we always copy full blocks. We're not done yet -- there's no call to inval have the file range mmap'd will continue accessing the old memory mapping after the file metadata updates have c if the unshare succeeds, we still need to return the number of bytes that the iomap iter thinks we're operating on.
CVE-2024-50252	In the Linux kernel, the following vulnerability has been resolved: mlxsw: spectrum_ipip: Fix memory leak when c stores IPv6 addresses that are used for encapsulation in linear memory that is managed by the driver. Changing the never worked properly, but since cited commit the following reproducer [1] would result in a warning [2] and a me remote address is never added by the driver to its hash table (and therefore the device) and the old address is never new address when the configuration of the ip6gre net device changes and removing the old one. If the address did r increasing the reference count of the address and then decreasing it. [1] # ip link add name bla up type ip6gre local inherit ttl inherit # ip link set dev bla type ip6gre remote 2001:db8:3::1 # ip link del dev bla # devlink dev reload pc 1682 at drivers/net/ethernet/mellanox/mlxsw/spectrum.c:3002 mlxsw_sp_ipv6_addr_put+0x140/0x1d0 Modules li Not tainted 6.12.0-rc3-custom-g86b5b55bc835 #151 Hardware name: Nvidia SN5600/VMOD0013, BIOS 5.13 05/ +0x140/0x1d0 [...] Call Trace: <TASK> mlxsw_sp_router_netdevice_event+0x55f/0x1240 notifier_call_chain+0x +0x39/0x90 unregister_netdevice_many_notify+0x63e/0x9d0 rtnl_dellink+0x16b/0x3a0 rtnetlink_rcv_msg+0x142 netlink_unicast+0x242/0x390 netlink_sendmsg+0x1de/0x420 ____sys_sendmsg+0x2bd/0x320 ___sys_sendmsg+0 do_syscall_64+0x9e/0x1a0 entry_SYSCALL_64_after_hwframe+0x77/0x7f [3] unreferenced object 0xffff898081 4294719324 hex dump (first 32 bytes): 20 01 0d b8 00 02 00 00 00 00 00 00 00 00 01 ..... 21 49 61 83 80 backtrace (crc fd9be911): [<00000000df89c55d>] __kmalloc_cache_noprof+0x1da/0x260 [<00000000ff2a1ddb>] +0x281/0x340 [<000000009ddd445d>] mlxsw_sp_router_netdevice_event+0x47b/0x1240 [<00000000743e7757> [<000000007c7b9e13>] call_netdevice_notifiers_info+0x39/0x90 [<000000002509645d>] register_netdevice+0x5 ip6gre_newlink_common.isra.0+0x65/0x130 [<0000000087cd6d8d>] ip6gre_newlink+0x72/0x120 [<000000004d [<0000000057ed632a>] rtnetlink_rcv_msg+0x142/0x3f0 [<0000000032e0d5b5>] netlink_rcv_skb+0x50/0x100 [< +0x242/0x390 [<00000000cdbe1c87>] netlink_sendmsg+0x1de/0x420 [<0000000011db153e>] ____sys_sendmsg ___sys_sendmsg+0x9a/0xe0 [<00000000cae27c62>] __sys_sendmsg+0x7a/0xd0
CVE-2024-50259	In the Linux kernel, the following vulnerability has been resolved: netdevsim: Add trailing zero to terminate the str This was found by a static analyzer. We should not forget the trailing zero after copy_from_user() if we will furthe case. Adding a trailing zero will ensure that the function performs properly.

CVE-2024-50261	In the Linux kernel, the following vulnerability has been resolved: macsec: Fix use-after-free while sending the off UAF. The metadata_dst, which is used to store the SCI value for macsec offload, is already freed by metadata_dst_ still use it for sending the packet. To fix this issue, dst_release() is used instead to release metadata_dst. So it is not if still referenced by skb. BUG: KASAN: slab-use-after-free in mlx5e_xmit+0x1e8f/0x4190 [mlx5_core] Read of s kworker/7:2/714 [...] Workqueue: mld mld_ifc_work Call Trace: <TASK> dump_stack_lvl+0x51/0x60 print_repor mlx5e_xmit+0x1e8f/0x4190 [mlx5_core] dev_hard_start_xmit+0x120/0x530 sch_direct_xmit+0x149/0x11e0 __qd +0x1196/0x2ed0 vlan_dev_hard_start_xmit+0x32e/0x510 [8021q] dev_hard_start_xmit+0x120/0x530 __dev_queu +0x13e9/0x2340 dev_hard_start_xmit+0x120/0x530 __dev_queue_xmit+0x14a7/0x2ed0 ip6_finish_output2+0x92 ip6_output+0x1ce/0x3a0 NF_HOOK.constprop.0+0x15f/0x190 mld_sendpack+0x59a/0xbd0 mld_ifc_work+0x48a worker_thread+0x79c/0x1290 kthread+0x28f/0x350 ret_from_fork+0x2d/0x70 ret_from_fork_asm+0x11/0x20 </T kasan_save_stack+0x20/0x40 kasan_save_track+0x10/0x30 __kasan_kmalloc+0x77/0x90 __kmalloc_noprof+0x18 macsec_newlink+0x914/0x1410 __rtnl_newlink+0xe08/0x15b0 rtnl_newlink+0x5f/0x90 rtnetlink_rcv_msg+0x667 netlink_unicast+0x551/0x770 netlink_sendmsg+0x72d/0xbd0 __sock_sendmsg+0xc5/0x190 ____sys_sendmsg+0x __sys_sendmsg+0xb5/0x140 do_syscall_64+0x4c/0x100 entry_SYSCALL_64_after_hwframe+0x4b/0x53 Freed b kasan_save_track+0x10/0x30 kasan_save_free_info+0x37/0x50 poison_slab_object+0x10c/0x190 __kasan_slab_f macsec_free_netdev+0x3f/0x140 netdev_run_todo+0x450/0xc70 rtnetlink_rcv_msg+0x66f/0xa80 netlink_rcv_skb netlink_sendmsg+0x72d/0xbd0 __sock_sendmsg+0xc5/0x190 ____sys_sendmsg+0x52e/0x6a0 ___sys_sendmsg+0 do_syscall_64+0x4c/0x100 entry_SYSCALL_64_after_hwframe+0x4b/0x53
CVE-2024-50268	In the Linux kernel, the following vulnerability has been resolved: usb: typec: fix potential out of bounds in ucsi_c The "*cmd" variable can be controlled by the user via debugfs. That means "new_cam" can be as high as 255 while array is UCSI_MAX_ALTMODES (30). The call tree is: ucsi_cmd() // val comes from simple_attr_write_xsigned(ucsi_send_command_common() -> ucsi_run_command() // calls ucsi->ops->sync_control() -> ucsi_ccg_sync_con
CVE-2024-50271	In the Linux kernel, the following vulnerability has been resolved: signal: restore the override_rlimit logic Prior to RLIMIT_SIGPENDING on top of ucounts") UCOUNT_RLIMIT_SIGPENDING rlimit was not enforced for a clas unconditionally, even if override_rlimit is set. This behavior change caused production issues. For example, if the l SIGSEGV signal, sigqueue_alloc fails to allocate the necessary resources for the signal delivery, preventing the sig prevents the process from correctly identifying the fault address and handling the error. From the user-space perspe has been reached and that the siginfo is effectively 'corrupted'. This can lead to unpredictable behavior and crashes, this by passing override_rlimit into inc_rlimit_get_ucounts() and skip the comparison to max there if override_rlim behavior.
CVE-2024-50294	In the Linux kernel, the following vulnerability has been resolved: rxrpc: Fix missing locking causing hanging calls a signal) between it being queued for connection and the I/O thread picking up the call, the abort will be prioritised from local->new_client_calls by rxrpc_disconnect_client_call() without a lock being held. This may cause other ca this by taking the client_call_lock when removing a call from whatever list its ->wait_link happens to be on.
CVE-2024-50303	In the Linux kernel, the following vulnerability has been resolved: resource,kexec: walk_system_ram_res_rev mus walk_system_ram_res_rev() erroneously discards resource flags when passing the information to the callback. This IORESOURCE_SYSRAM_DRIVER_MANAGED memory to have these resources selected during kexec to store to be at placed above normal system ram. This leads to undefined behavior after reboot. If the kexec buffer is never kexec buffer is touched, it could lead to a crash (like below) or undefined behavior. Tested on a system with CXL r managed memory, TPM enabled, and CONFIG_IMA_KEXEC=y. Adding printk's showed the flags were being dis IORESOURCE_SYSRAM_DRIVER_MANAGED passes. find_next_iomem_res: name(System RAM (kmem)) st flags(83000200) locate_mem_hole_top_down: start(10000000000) end(1034fffffff) flags(0) [.] BUG: unable to han [.] #PF: supervisor read access in kernel mode [.] #PF: error_code(0x0000) - not-present page [.] PGD c04c8bf067 PMD 0 [.] Oops: 0000 [#1] SMP [.] RIP: 0010:ima_restore_measurement_list+0x95/0x4b0 [.] RSP: 0018:ffffc900(0000000000001000 RBX: 0000000000000000 RCX: ffff89834ffff000 [.] RDX: 0000000000000018 RSI: ffff8983 ffffc900000d3ba0 R08: 0000000000000020 R09: ffff888132b8a900 [.] R10: 4000000000000000 R11: 000000003 ffffffff8404ac28 R14: 0000000000000000 R15: ffff89834ffff000 [.] FS: 0000000000000000(0000) GS:ffff893d44 [.] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 [.] ata5: SATA link down (SStatus 0 SControl 300) [.] C 000001034d00f001 CR4: 0000000000770ef0 [.] PKRU: 55555554 [.] Call Trace: [.] <TASK> [.] ? __die+0x78/0x [.] ? exc_page_fault+0x84/0x130 [.] ? asm_exc_page_fault+0x22/0x30 [.] ? ima_restore_measurement_list+0x95/0 +0x317/0x410 [.] ? crypto_alloc_tfm_node+0x9c/0xc0 [.] ? init_ima_lsm+0x30/0x30 [.] ima_load_kexec_buffer+0 [.] __initstub__kmod_ima__373_1201_init_ima7+0x1e/0xb0 [.] ? init_ima_lsm+0x30/0x30 [.] do_one_initcall+0x +0xaa/0x110 [.] ? new_slab+0x12c/0x420 [.] ? new_slab+0x12c/0x420 [.] ? number+0x12a/0x430 [.] ? sysvec_api asm_sysvec_apic_timer_interrupt+0x16/0x20 [.] ? parse_args+0xd4/0x380 [.] ? parse_args+0x14b/0x380 [.] kerne +0xb0/0xb0 [.] kernel_init+0x16/0x1a0 [.] ret_from_fork+0x2f/0x40 [.] ? rest_init+0xb0/0xb0 [.] ret_from_fork_as
CVE-2024-5197	There exists interger overflows in libvpx in versions prior to 1.14.1.¬†Calling vpx_img_alloc()¬†with a large value may result in integer overflows in the calculations of buffer sizes and offsets and some fields of the returned vpx_i vpx_img_wrap()¬†with a large value of the d_w, d_h, or stride_align¬†parameter may result in integer overflows i and some fields of the returned vpx_image_t¬†struct may be invalid. We recommend upgrading to version 1.14.1 o

CVE-2024-53044	In the Linux kernel, the following vulnerability has been resolved: net/sched: sch_api: fix xa_insert() error path in t replace dev eth0 ingress_block 1 egress_block 1 clsact Error: block dev insert failed: -EBUSY. fails because user s for both ingress and egress. [side note, I don't think it even failed prior to commit 913b47d3424e ("net/sched: Intro this is a command from an old set of notes of mine which used to work, but alas, I did not scientifically bisect this that the second time around, it fails differently (and irrecoverably): $ tc qdisc replace dev eth0 ingress_block 1 egre cb is busy. [another note: the extack is added by me for illustration purposes. the context of the problem is that clsa pointer as &q->egress_block, and since we call tcf_block_get_ext() on both of them, "dev" will be added to the blo operation: once through the ingress block pointer, and once again through the egress block pointer. the problem itse emitted a FLOW_BLOCK_BIND command through ndo_setup_tc(), but the offload never sees a corresponding FL the bad user input, we still cannot recover: $ tc qdisc replace dev swp3 ingress_block 1 egress_block 2 clsact Error the only way to recover is to reboot the system, or unbind and rebind the net device driver. To fix the bug, we need was missed during code movement, and call tcf_block_offload_unbind() when xa_insert() fails. [last note, fundam in tcf_block_get_ext() for the bug. The labels should be named after what they do, not after the error path that jump that two labels pointing to the same code mean something is wrong, and checking the code correctness at the goto s
CVE-2024-53045	In the Linux kernel, the following vulnerability has been resolved: ASoC: dapm: fix bounds checker error in dapm_ snd_soc_dapm_widget_list has a __counted_by attribute attached to it, which points to the num_widgets variable. T if it is not set before the array is filled, then the bounds sanitizer will issue a warning or a kernel panic if CONFIG_ size of the widgets list calculated with list_for_each as the initial value for num_widgets as it is used for allocating actual number of added elements after the array is filled.
CVE-2024-53047	In the Linux kernel, the following vulnerability has been resolved: mptcp: init: protect sched with rcu_read_lock Er with its dependence CONFIG_RCU_EXPERT creates this splat when an MPTCP socket is created: ========== suspicious RCU usage 6.12.0-rc2+ #11 Not tainted ----------------------------- net/mptcp/sched.c:44 RCU-list traverse that might help us debug this: rcu_scheduler_active = 2, debug_locks = 1 no locks held by mptcp_connect/176. sta Comm: mptcp_connect Not tainted 6.12.0-rc2+ #11 Hardware name: Bochs Bochs, BIOS Bochs 01/01/2011 Call T dump_stack.c:123) lockdep_rcu_suspicious (kernel/locking/lockdep.c:6822) mptcp_sched_find (net/mptcp/sched.c mptcp/protocol.c:2867 (discriminator 1)) ? sock_init_data_uid (arch/x86/include/asm/atomic.h:28) inet_create.part __sock_create (include/linux/rcupdate.h:347 (discriminator 1)) __sock_create (net/socket.c:1576) __sys_socket (ne (net/socket.c:1712) ? do_user_addr_fault (arch/x86/mm/fault.c:1419 (discriminator 1)) __x64_sys_socket (net/sock entry/common.c:52 (discriminator 1)) entry_SYSCALL_64_after_hwframe (arch/x86/entry/entry_64.S:130) That's rcu_read_lock() is not used despite the explicit comment written above the declaration of mptcp_sched_find() in sc the warning.
CVE-2024-53048	In the Linux kernel, the following vulnerability has been resolved: ice: fix crash on probe for DPLL enabled E810 design is vendor specific. Intel provides the reference design, but it is up to vendor on the final product design. For the static values defined in the driver does not reflect the actual LOM design. Current implementation of dpll pins i driver for such DPLL enabled E810 LOM designs: WARNING: (...) at drivers/dpll/dpll_core.c:495 dpll_pin_get+0 __warn+0x83/0x130 ? dpll_pin_get+0x2c4/0x330 ? report_bug+0x1b7/0x1d0 ? handle_bug+0x42/0x70 ? exc_inva +0x1a/0x20 ? dpll_pin_get+0x117/0x330 ? dpll_pin_get+0x2c4/0x330 ? dpll_pin_get+0x117/0x330 ice_dpll_get_ dpll pins enabled by LOM vendor is greater than expected and defined in the driver for Intel designed NICs, which generic pin initialization within Linux DPLL subsystem for DPLL enabled E810 LOM designs. Newly designed so "per HW design" pin initialization. It requires pin information dynamically acquired from the firmware and is alrea
CVE-2024-53070	In the Linux kernel, the following vulnerability has been resolved: usb: dwc3: fix fault at system suspend if device was already runtime suspended then during system suspend we cannot access the device registers else it will crash. dwc3_core_exit() on some platforms so move the dwc3_enable_susphy() call to the top.
CVE-2024-53072	In the Linux kernel, the following vulnerability has been resolved: platform/x86/amd/pmc: Detect when STB is not as: amd_pmc enable_stb=1 ...can result in the following messages in the kernel ring buffer: amd_pmc AMDI0009:0 RAM at 0x0000000000000000 - 0x00000000000fffff WARNING: CPU: 10 PID: 2151 at arch/x86/mm/ioremap.c:1 debugging reveals that this occurs when the requests for S2D_PHYS_ADDR_LOW and S2D_PHYS_ADDR_HIG is inaccessible. To prevent the ioremap warning and provide clarity to the user, handle the invalid address and displ
CVE-2024-53074	In the Linux kernel, the following vulnerability has been resolved: wifi: iwlwifi: mvm: don't leak a link on AP remo AP removal. This impacted devices that do not support the MLD API (9260 and down). On those devices, we could already started and stopped.
CVE-2024-53082	In the Linux kernel, the following vulnerability has been resolved: virtio_net: Add hash_key_length check Add has avoid possible out of bound errors when setting/reading the hash key.
CVE-2024-53085	In the Linux kernel, the following vulnerability has been resolved: tpm: Lock TPM chip in tpm_pm_suspend() first in the end of tpm_pm_suspend() can be racy according, as this leaves window for tpm_hwrng_read() to be called w recent bug report gives also evidence of this behaviour. Aadress this by locking the TPM chip before checking any tpm_hwrng_read(). Move TPM_CHIP_FLAG_SUSPENDED check inside tpm_get_random() so that it will be alw

CVE-2024-53091	In the Linux kernel, the following vulnerability has been resolved: bpf: Add sk_is_inet and IS_ICSK check in tls_s of the support for vsock and unix sockets in sockmap, tls_sw_has_ctx_tx/rx cannot presume the socket passed in m sockets have vsock_sock and unix_sock instead of inet_connection_sock. For these sockets, tls_get_ctx may return in function tls_sw_ctx_rx. BUG: unable to handle page fault for address: 0000000000040030 Workqueue: vsock-lo 0010:sk_psock_strp_data_ready+0x23/0x60 Call Trace: ? __die+0x81/0xc3 ? no_context+0x194/0x350 ? do_page +0x3e/0x50 ? sk_psock_strp_data_ready+0x23/0x60 virtio_transport_recv_pkt+0x750/0x800 ? update_load_avg+(+0xd0/0x100 process_one_work+0x1a7/0x360 worker_thread+0x30/0x390 ? create_worker+0x1a0/0x1a0 kthread +0x40/0x40 ret_from_fork+0x1f/0x40 v2: - Add IS_ICSK check v3: - Update the commits in Fixes
CVE-2024-53096	In the Linux kernel, the following vulnerability has been resolved: mm: resolve faulty mmap_region() error path be is somewhat terrifying, with spaghetti-like control flow and numerous means by which issues can arise and incomp unpleasantness can occur. A large amount of the complexity arises from trying to handle errors late in the process c of recently observed issues with resource leaks and observable inconsistent state. Taking advantage of previous pat checks earlier in the code, simplifying things by moving the core of the logic into a static internal function __mmap a number of checks up front before we do any real work, and allows us to unwind the writable unmap check uncon CONFIG_DEBUG_VM_MAPLE_TREE validation unconditionally also. We move a number of things here: 1. We we call the file-backed memory hook, allowing us to exit early and avoid having to perform complicated and error- iterator state on both success and error paths. 2. The enclosing mmap_region() function handles the mapping_map_ had the mapping_map_writable() at the point of mapping a newly allocated file-backed VMA, and a matching map paths. We now do this unconditionally if this is a file-backed, shared writable mapping. If a driver changes the flag doing so does not invalidate the seal check we just performed, and we in any case always decrement the counter in ensure a driver does not attempt to do the opposite. 3. We also move arch_validate_flags() up into the mmap_regio and sparc64, and the check is only meaningful for SPARC with ADI enabled. We explicitly add a warning for this the code ought eventually to be fixed to eliminate the need for this. With all of these measures in place, we no long paths, as we place all checks which might fail prior to a call to any driver mmap hook. This eliminates an entire cla about and more robust.
CVE-2024-53097	In the Linux kernel, the following vulnerability has been resolved: mm: krealloc: Fix MTE false alarm in __do_kre introduced by commit 1a83a716ec233 ("mm: krealloc: consider spare memory for __GFP_ZERO") which causes M falsely report a slab-out-of-bounds error. The problem occurs when zeroing out spare memory in __do_krealloc. Th based KASAN and did not account for MTE. It does not reset the KASAN tag before calling memset, leading to a memory tag, resulting in a false positive. Example of the error: ================================== swapper/0: BUG: KASAN: slab-out-of-bounds in __memset+0x84/0x188 swapper/0: Write at addr f4ffff8005f0fd tag: [f4], memory tag: [fe] swapper/0: swapper/0: CPU: 4 UID: 0 PID: 1 Comm: swapper/0 Not tainted 6.12. swapp (DT) swapper/0: Call trace: swapper/0: dump_backtrace+0xfc/0x17c swapper/0: show_stack+0x18/0x28 swapper/(print_report+0x1b8/0x71c swapper/0: kasan_report+0xec/0x14c swapper/0: __do_kernel_fault+0x60/0x29c swapp swapper/0: do_tag_check_fault+0x20/0x34 swapper/0: do_mem_abort+0x58/0x104 swapper/0: el1_abort+0x3c/0x +0x80/0xcc swapper/0: el1h_64_sync+0x68/0x6c swapper/0: __memset+0x84/0x188 swapper/0: btf_populate_kfu __register_btf_kfunc_id_set+0x43c/0x468 swapper/0: register_btf_kfunc_id_set+0x48/0x60 swapper/0: register_n nf_nat_init+0xc0/0x128 swapper/0: do_one_initcall+0x184/0x464 swapper/0: do_initcall_level+0xdc/0x1b0 swapp do_basic_setup+0x1c/0x28 swapper/0: kernel_init_freeable+0x144/0x1b8 swapper/0: kernel_init+0x20/0x1a8 swa ==
CVE-2024-53105	In the Linux kernel, the following vulnerability has been resolved: mm: page_alloc: move mlocked flag clearance i a bad page state problem caused by a page being freed using free_page() still having a mlocked flag at free_pages_ process syz.5.504 pfn:61f45 page: refcount:0 mapcount:0 mapping:0000000000000000 index:0x0 pfn:0x61f45 flag mlocked|node=0|zone=1|lastcpupid=0x7ff) raw: 00fff00000080204 0000000000000000 dead000000000122 00000 0000000000000000 00000000ffffffff 0000000000000000 page dumped because: PAGE_FLAGS_CHECK_AT_FI page as allocated page last allocated via order 0, migratetype Unmovable, gfp_mask 0x400dc0(GFP_KERNEL_AC 8442 (syz.5.504), ts 201884660643, free_ts 201499827394 set_page_owner include/linux/page_owner.h:32 [inline page_alloc.c:1537 prep_new_page mm/page_alloc.c:1545 [inline] get_page_from_freelist+0x303f/0x3190 mm/pag +0x292/0x710 mm/page_alloc.c:4733 alloc_pages_mpol_noprof+0x3e8/0x680 mm/mempolicy.c:2265 kvm_coales coalesced_mmio.c:99 kvm_create_vm virt/kvm/kvm_main.c:1235 [inline] kvm_dev_ioctl_create_vm virt/kvm/kvm +0x12dc/0x2240 virt/kvm/kvm_main.c:5530 __do_compat_sys_ioctl fs/ioctl.c:1007 [inline] __se_compat_sys_ioc do_syscall_32_irqs_on arch/x86/entry/common.c:165 [inline] __do_fast_syscall_32+0xb4/0x110 arch/x86/entry/co arch/x86/entry/common.c:411 entry_SYSENTER_compat_after_hwframe+0x84/0x8e page last free pid 8399 tgid include/linux/page_owner.h:25 [inline] free_pages_prepare mm/page_alloc.c:1108 [inline] free_unref_folios+0xf1 folios_put_refs+0x76c/0x860 mm/swap.c:1007 free_pages_and_swap_cache+0x5c8/0x690 mm/swap_state.c:335 mmu_gather.c:136 [inline] tlb_batch_pages_flush mm/mmu_gather.c:149 [inline] tlb_flush_mmu_free mm/mmu_g +0x3a3/0x680 mm/mmu_gather.c:373 tlb_finish_mmu+0xd4/0x200 mm/mmu_gather.c:465 exit_mmap+0x496/0x +0x115/0x390 kernel/fork.c:1348 exit_mm+0x220/0x310 kernel/exit.c:571 do_exit+0x9b2/0x28e0 kernel/exit.c:92 exit.c:1088 __do_sys_exit_group kernel/exit.c:1099 [inline] __se_sys_exit_group kernel/exit.c:1097 [inline] __x64 exit.c:1097 x64_sys_call+0x2634/0x2640 arch/x86/include/generated/asm/syscalls_64.h:232 do_syscall_x64 arch/ do_syscall_64+0xf3/0x230 arch/x86/entry/common.c:83 entry_SYSCALL_64_after_hwframe+0x77/0x7f Modules syz.5.504 Not tainted 6.12.0-rc6-syzkaller #0 Hardware name: Google Google Compute Engine/Google Compute E <TASK> __dump_stack lib/dump_stack.c:94 [inline] dump_stack_lvl+0x241/0x360 lib/dump_stack.c:120 bad_pag free_page_is_bad mm/page_alloc.c:918 [inline] free_pages_prepare mm/page_alloc.c:1100 [inline] free_unref_pag kvm_destroy_vm virt/kvm/kvm_main.c:1327 [inline] kvm_put_kvm+0xc75/0x1350 virt/kvm/kvm_main.c:1386 kv kvm_main.c:4143 __fput+0x23f/0x880 fs/file_table.c:431 task_work_run+0x24f/0x310 kernel/task_work.c:239 ex [inline] do_exit+0xa2f/0x28e0 kernel/exit.c:939 do_group_exit+0x207/0x2c0 kernel/exit.c:1088 __do_sys_exit_gr

CVE-2024-53109	In the Linux kernel, the following vulnerability has been resolved: nommu: pass NULL argument to vma_iter_prea maple tree, it has to pass NULL to vma_iter_prealloc() in order to calculate internal state of the tree, but it passed a crashed upon accessing a vma iterator, such as acct_collect() reading the size of vma entries after do_munmap(). T argument to the preallocation call.
CVE-2024-53110	In the Linux kernel, the following vulnerability has been resolved: vp_vdpa: fix id_table array not null terminated e null terminator, otherwise vdpa_mgmtdev_get_classes() may iterate multiple times and visit undefined memory.
CVE-2024-53113	In the Linux kernel, the following vulnerability has been resolved: mm: fix NULL pointer dereference in alloc_pag pointer dereference for ac.preferred_zoneref->zone in alloc_pages_bulk_noprof() when the task is migrated betwee in prepare_alloc_pages(), ac->nodemask may be &current->mems_allowed. when first_zones_zonelist() is called t >nodemask may be modified concurrently if the task is migrated between different cpusets. Assuming we have 2 N ac->zonelist, the nodemask is 2, and when traversing Node2 in ac->zonelist, the nodemask is 1. As a result, the ac- In alloc_pages_bulk_noprof(), for_each_zone_zonelist_nodemask() finds a allowable zone and calls zonelist_node NULL pointer dereference. __alloc_pages_noprof() fixes this issue by checking NULL pointer in commit ea57485 NULL preferred_zone") and commit df76cee6bbeb ("mm, page_alloc: remove redundant checks from alloc fastpat preferred_zoneref->zone.
CVE-2024-53117	In the Linux kernel, the following vulnerability has been resolved: virtio/vsock: Improve MSG_ZEROCOPY error prevent memory leaks.
CVE-2024-53118	In the Linux kernel, the following vulnerability has been resolved: vsock: Fix sk_error_queue memory leak Kernel completion notifications on the error queue. Where they remain, until explicitly recv()ed. To prevent memory leaks the socket is destroyed. unreferenced object 0xffff8881028beb00 (size 224): comm "vsock_test", pid 1218, jiffies 4 32 bytes): 90 b0 21 17 81 88 ff ff 90 b0 21 17 81 88 ff ff ..!.....!..... 00 00 00 00 00 00 00 00 00 b0 21 17 81 88 ff f 6c7031ca): [<ffffffff81418ef7>] kmem_cache_alloc_node_noprof+0x2f7/0x370 [<ffffffff81d35882>] __alloc_skb sock_omalloc+0x4b/0x80 [<ffffffff81d3a8ae>] msg_zerocopy_realloc+0x9e/0x240 [<ffffffff81fe5cb2>] virtio_tra [<ffffffff81fe6183>] virtio_transport_stream_enqueue+0x43/0x50 [<ffffffff81fe0813>] vsock_connectible_sendms ____sys_sendmsg+0x365/0x3a0 [<ffffffff81d246f4>] ___sys_sendmsg+0x84/0xd0 [<ffffffff81d26f47>] __sys_se do_syscall_64+0x93/0x180 [<ffffffff8220012b>] entry_SYSCALL_64_after_hwframe+0x76/0x7e
CVE-2024-53120	In the Linux kernel, the following vulnerability has been resolved: net/mlx5e: CT: Fix null-ptr-deref in add rule err mlx5_tc_ct_entry_add_rule(), in case ct_rule_add() callback returns error, zone_rule->attr is used uninitiated. Fix i pointer value. Kernel log: BUG: kernel NULL pointer dereference, address: 0000000000000110 RIP: 0010:mlx5_t [mlx5_core] ,Ä¶ Call Trace: <TASK> ? __die+0x20/0x70 ? page_fault_oops+0x150/0x3e0 ? exc_page_fault+0x7 mlx5_tc_ct_entry_add_rule+0x2b1/0x2f0 [mlx5_core] ? mlx5_tc_ct_entry_add_rule+0x1d5/0x2f0 [mlx5_core] ml [mlx5_core] ? nf_flow_offload_tuple+0xd8/0x190 [nf_flow_table] nf_flow_offload_tuple+0xd8/0x190 [nf_flow_t +0x142/0x320 [nf_flow_table] ? finish_task_switch.isra.0+0x15b/0x2b0 process_one_work+0x16c/0x320 worker_ +0x10/0x10 kthread+0xb8/0xf0 ? __pfx_kthread+0x10/0x10 ret_from_fork+0x2d/0x50 ? __pfx_kthread+0x10/0x
CVE-2024-53122	In the Linux kernel, the following vulnerability has been resolved: mptcp: cope racing subflow creation in mptcp_r - i.e. created by the in kernel path manager - are included into the subflow list before starting the 3whs. A racing re established subflow would unconditionally call tcp_cleanup_rbuf() on all the current subflows, potentially hitting a ones. Explicitly check that the subflow is in a suitable state before invoking tcp_cleanup_rbuf().
CVE-2024-53123	In the Linux kernel, the following vulnerability has been resolved: mptcp: error out earlier on disconnect Eric repor MPTCP protocol: Oops: divide error: 0000 [#1] PREEMPT SMP KASAN PTI CPU: 1 UID: 0 PID: 6094 Comm: s rc5-syzkaller-00291-g05b92660cdfe #0 Hardware name: Google Google Compute Engine/Google Compute Engin 0010:__tcp_select_window+0x5b4/0x1310 net/ipv4/tcp_output.c:3163 Code: f6 44 01 e3 89 df e8 9b 75 09 f8 44 3 e9 04 ff ff ff e8 00 74 09 f8 44 89 f0 99 <f7> 7c 24 14 41 29 d6 45 89 f4 e9 ec fe ff ff e8 e8 73 09 f8 48 89 RSP: 0 RAX: 0000000000017e67 RBX: 0000000000017e67 RCX: ffffffff8983314b RDX: 0000000000000000 RSI: ffffff 00000000005d6000 R08: 0000000000000004 R09: 0000000000017e67 R10: 0000000000003e80 R11: 000000000 ffff888031d9b440 R14: 0000000000017e67 R15: 00000000002eb000 FS: 00007feb5d7f16c0(0000) GS:ffff8880b8 CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033 CR2: 00007feb5d8adbb8 CR3: 0000000074e4c000 CR4: 0 DR1: 0000000000000000 DR2: 0000000000000000 DR3: 0000000000000000 DR6: 00000000fffe0ff0 DR7: 0000 __tcp_cleanup_rbuf+0x3e7/0x4b0 net/ipv4/tcp.c:1493 mptcp_rcv_space_adjust net/mptcp/protocol.c:2085 [inline] mptcp/protocol.c:2289 inet_recvmsg+0x469/0x6a0 net/ipv4/af_inet.c:885 sock_recvmsg_nosec net/socket.c:1051 net/socket.c:1073 __sys_recvfrom+0x1a5/0x2e0 net/socket.c:2265 __do_sys_recvfrom net/socket.c:2283 [inline] _ [inline] __x64_sys_recvfrom+0xe0/0x1c0 net/socket.c:2279 do_syscall_x64 arch/x86/entry/common.c:52 [inline] entry/common.c:83 entry_SYSCALL_64_after_hwframe+0x77/0x7f RIP: 0033:0x7feb5d857559 Code: 28 00 00 0 00 90 48 89 f8 48 89 f7 48 89 d6 48 89 ca 4d 89 c2 4d 89 c8 4c 8b 4c 24 08 0f 05 <48> 3d 01 f0 ff ff 73 01 c3 48 c 002b:00007feb5d7f1208 EFLAGS: 00000246 ORIG_RAX: 000000000000002d RAX: ffffffffffffffda RBX: 00007 RDX: 000000800000000e RSI: 0000000000000000 RDI: 0000000000000003 RBP: 00007feb5d8e1310 R08: 0000 0000000000000100 R11: 0000000000000246 R12: 00007feb5d8e131c R13: 00007feb5d8ae074 R14: 0000008000 a nice reproducer. The root cause is the current bad handling of racing disconnect. After the blamed commit below, the underlying socket disconnected and a zero rcv_mss. Catch the error and return without performing any addition
CVE-2024-53134	In the Linux kernel, the following vulnerability has been resolved: pmdomain: imx93-blk-ctrl: correct remove path >onecell_data.num_domains', not 'bc->onecell_data.num_domains' which will make the look never finish and caus address "imx93-blk-ctrl 4ac10000.system-controller: Unbalanced pm_runtime_enable!"

CVE-2024-5742	A vulnerability was found in GNU Nano that allows a possible privilege escalation through an insecure temporary saves to an emergency file with the permissions of the running user provides a window of opportunity for attackers symlink.
CVE-2024-7348	Time-of-check Time-of-use (TOCTOU) race condition in pg_dump in PostgreSQL allows an object creator to exec running pg_dump, which is often a superuser. The attack involves replacing another relation type with a view or fo pg_dump to start, but winning the race condition is trivial if the attacker retains an open transaction. Versions befor 12.20 are affected.
CVE-2024-7348	Time-of-check Time-of-use (TOCTOU) race condition in pg_dump in PostgreSQL allows an object creator to exec running pg_dump, which is often a superuser. The attack involves replacing another relation type with a view or fo pg_dump to start, but winning the race condition is trivial if the attacker retains an open transaction. Versions befor 12.20 are affected.
DLA-3859-1	systemd - security update
DLA-3875-1	gnutls28 - security update
DLA-3893-1	expat - security update
DLA-3898-1	nghttp2 - security update
DLA-3904-1	cups - security update
DLA-3907-1	sqlite3 - security update
DLA-3910-1	e2fsprogs - security update
DLA-3910-1	e2fsprogs - security update
DLA-3926-1	perl - security update
DLA-3926-1	perl - security update
DLA-3926-1	perl - security update
DLA-3926-1	perl - security update
DLA-3930-1	libsepol - security update
DLA-3931-1	ghostscript - security update
DLA-3931-1	ghostscript - security update
DLA-3931-1	ghostscript - security update
DLA-3934-1	libheif - security update
DLA-3945-1	libheif - security update
DLA-3951-1	curl - security update
DLA-3954-1	postgresql-13 - security update
DLA-3954-2	
DLA-3962-1	glib2.0 - security update
DLA-3962-1	glib2.0 - security update
DLA-3965-1	ghostscript - security update
DLA-3965-1	ghostscript - security update
DLA-3965-1	ghostscript - security update
DLA-3980-1	python3.9 - security update
DLA-3980-1	python3.9 - security update
DLA-3980-1	python3.9 - security update
DLA-3980-1	python3.9 - security update
DLA-3990-1	avahi - security update
DLA-3990-1	avahi - security update
DLA-3990-1	avahi - security update

DSA-5650-1	util-linux - security update
DSA-5650-1	util-linux - security update
DSA-5650-1	util-linux - security update
DSA-5650-1	util-linux - security update
DSA-5650-1	util-linux - security update
DSA-5650-1	util-linux - security update
DSA-5650-1	util-linux - security update
DSA-5726-1	krb5 - security update
DSA-5726-1	krb5 - security update
DSA-5726-1	krb5 - security update
DSA-5726-1	krb5 - security update
GHSA-jm77-qphf-c4w8	pyca/cryptography's wheels include a statically linked copy of OpenSSL. The versions of OpenSSL included in cry to several security issues. More details about the vulnerabilities themselves can be found in https://www.openssl.or www.openssl.org/news/secadv/20230719.txt, and https://www.openssl.org/news/secadv/20230714.txt. If you are b you are responsible for upgrading your copy of OpenSSL. Only users installing from wheels built by the cryptogra need to update their cryptography versions.
GHSA-v8gr-m533-ghj9	pyca/cryptography's wheels include a statically linked copy of OpenSSL. The versions of OpenSSL included in cry several security issues. More details about the vulnerabilities themselves can be found in https://www.openssl.org/ cryptography source ("sdist") then you are responsible for upgrading your copy of OpenSSL. Only users installing (i.e., those distributed on PyPI) need to update their cryptography versions.
RHSA-2024:0889	Oniguruma is a regular expressions library that supports a variety of character encodings.
RHSA-2024:10244	Pluggable Authentication Modules (PAM) provide a system to set up authentication policies without the need to re
RHSA-2024:10472	WebKitGTK is the port of the portable web rendering engine WebKit to the GTK platform.
RHSA-2024:10791	PostgreSQL is an advanced object-relational database management system (DBMS).
RHSA-2024:10791	PostgreSQL is an advanced object-relational database management system (DBMS).
RHSA-2024:5815	Node.js is a software development platform for building fast and scalable network applications in the JavaScript pr
RHSA-2024:5815	Node.js is a software development platform for building fast and scalable network applications in the JavaScript pr
RHSA-2024:5815	Node.js is a software development platform for building fast and scalable network applications in the JavaScript pr
RHSA-2024:5815	Node.js is a software development platform for building fast and scalable network applications in the JavaScript pr
RHSA-2024:6148	Node.js is a software development platform for building fast and scalable network applications in the JavaScript pr
RHSA-2024:6148	Node.js is a software development platform for building fast and scalable network applications in the JavaScript pr
RHSA-2024:6148	Node.js is a software development platform for building fast and scalable network applications in the JavaScript pr
RHSA-2024:6148	Node.js is a software development platform for building fast and scalable network applications in the JavaScript pr
RHSA-2024:6162	The python-urllib3 package provides the Python HTTP module with connection pooling and file POST abilities.
RHSA-2024:6163	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2024:6163	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2024:6179	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2024:6179	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2024:6179	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2024:6192	The wget packages provide the GNU Wget file retrieval utility for HTTP, HTTPS, and FTP protocols.

RHSA-2024:6464	GLib provides the core application building blocks for libraries and applications written in C. It provides the core o implementation, and a large set of utility functions for strings and common data structures.
RHSA-2024:6464	GLib provides the core application building blocks for libraries and applications written in C. It provides the core o implementation, and a large set of utility functions for strings and common data structures.
RHSA-2024:6510	GNU Emacs is a powerful, customizable, self-documenting text editor. It provides special code editing features, a s to read e-mail and news.
RHSA-2024:6754	Expat is a C library for parsing XML documents.
RHSA-2024:6754	Expat is a C library for parsing XML documents.
RHSA-2024:6783	OpenSSL is a toolkit that implements the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protoco cryptography library.
RHSA-2024:6783	OpenSSL is a toolkit that implements the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protoco cryptography library.
RHSA-2024:6783	OpenSSL is a toolkit that implements the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protoco cryptography library.
RHSA-2024:6963	The GTK+ library provides a multi-platform toolkit for creating graphical user interfaces. The gtk3 packages conta
RHSA-2024:6963	The GTK+ library provides a multi-platform toolkit for creating graphical user interfaces. The gtk3 packages conta
RHSA-2024:8038	The container-tools module contains tools for working with containers, notably podman, buildah, skopeo, and runc
RHSA-2024:8038	The container-tools module contains tools for working with containers, notably podman, buildah, skopeo, and runc
RHSA-2024:8038	The container-tools module contains tools for working with containers, notably podman, buildah, skopeo, and runc
RHSA-2024:8162	The kernel packages contain the Linux kernel, the core of any Linux operating system.
RHSA-2024:8180	WebKitGTK is the port of the portable web rendering engine WebKit to the GTK platform.
RHSA-2024:8374	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2024:8374	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2024:8374	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2024:8446	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2024:8446	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2024:8680	The mod_h2 Apache httpd module implements the HTTP2 protocol (h2+h2c) on top of libnghttp2 for httpd 2.4 ser
RHSA-2024:8914	The libtiff packages contain a library of functions for manipulating Tagged Image File Format (TIFF) files.
RHSA-2024:8914	The libtiff packages contain a library of functions for manipulating Tagged Image File Format (TIFF) files.
RHSA-2024:9144	WebKitGTK is the port of the portable web rendering engine WebKit to the GTK platform.
RHSA-2024:9184	The GTK+ library provides a multi-platform toolkit for creating graphical user interfaces. The gtk3 packages conta
RHSA-2024:9184	The GTK+ library provides a multi-platform toolkit for creating graphical user interfaces. The gtk3 packages conta
RHSA-2024:9192	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2024:9192	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2024:9192	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2024:9302	GNU Emacs is a powerful, customizable, self-documenting text editor. It provides special code editing features, a s to read e-mail and news.
RHSA-2024:9306	The httpd packages provide the Apache HTTP Server, a powerful, efficient, and extensible web server.
RHSA-2024:9306	The httpd packages provide the Apache HTTP Server, a powerful, efficient, and extensible web server.

RHSA-2024:9306	The httpd packages provide the Apache HTTP Server, a powerful, efficient, and extensible web server.
RHSA-2024:9306	The httpd packages provide the Apache HTTP Server, a powerful, efficient, and extensible web server.
RHSA-2024:9306	The httpd packages provide the Apache HTTP Server, a powerful, efficient, and extensible web server.
RHSA-2024:9306	The httpd packages provide the Apache HTTP Server, a powerful, efficient, and extensible web server.
RHSA-2024:9306	The httpd packages provide the Apache HTTP Server, a powerful, efficient, and extensible web server.
RHSA-2024:9306	The httpd packages provide the Apache HTTP Server, a powerful, efficient, and extensible web server.
RHSA-2024:9306	The httpd packages provide the Apache HTTP Server, a powerful, efficient, and extensible web server.
RHSA-2024:9315	The kernel packages contain the Linux kernel, the core of any Linux operating system.
RHSA-2024:9331	Kerberos is a network authentication system, which can improve the security of your network by eliminating the in network in unencrypted form. It allows clients and servers to authenticate to each other with the help of a trusted th (KDC).
RHSA-2024:9331	Kerberos is a network authentication system, which can improve the security of your network by eliminating the in network in unencrypted form. It allows clients and servers to authenticate to each other with the help of a trusted th (KDC).
RHSA-2024:9331	Kerberos is a network authentication system, which can improve the security of your network by eliminating the in network in unencrypted form. It allows clients and servers to authenticate to each other with the help of a trusted th (KDC).
RHSA-2024:9331	Kerberos is a network authentication system, which can improve the security of your network by eliminating the in network in unencrypted form. It allows clients and servers to authenticate to each other with the help of a trusted th (KDC).
RHSA-2024:9331	Kerberos is a network authentication system, which can improve the security of your network by eliminating the in network in unencrypted form. It allows clients and servers to authenticate to each other with the help of a trusted th (KDC).
RHSA-2024:9333	OpenSSL is a toolkit that implements the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protoco cryptography library.
RHSA-2024:9333	OpenSSL is a toolkit that implements the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protoco cryptography library.
RHSA-2024:9333	OpenSSL is a toolkit that implements the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protoco cryptography library.
RHSA-2024:9371	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2024:9371	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2024:9404	The libgcrypt library provides general-purpose implementations of various cryptographic algorithms.
RHSA-2024:9405	Vim (Vi IMproved) is an updated and improved version of the vi editor.
RHSA-2024:9405	Vim (Vi IMproved) is an updated and improved version of the vi editor.
RHSA-2024:9470	The Common UNIX Printing System (CUPS) provides a portable printing layer for Linux, UNIX, and similar oper
RHSA-2024:9474	Kerberos is a network authentication system, which can improve the security of your network by eliminating the in network in unencrypted form. It allows clients and servers to authenticate to each other with the help of a trusted th (KDC).
RHSA-2024:9474	Kerberos is a network authentication system, which can improve the security of your network by eliminating the in network in unencrypted form. It allows clients and servers to authenticate to each other with the help of a trusted th (KDC).
RHSA-2024:9474	Kerberos is a network authentication system, which can improve the security of your network by eliminating the in network in unencrypted form. It allows clients and servers to authenticate to each other with the help of a trusted th (KDC).
RHSA-2024:9474	Kerberos is a network authentication system, which can improve the security of your network by eliminating the in network in unencrypted form. It allows clients and servers to authenticate to each other with the help of a trusted th (KDC).

RHSA-2024:9474	Kerberos is a network authentication system, which can improve the security of your network by eliminating the in network in unencrypted form. It allows clients and servers to authenticate to each other with the help of a trusted th (KDC).
RHSA-2024:9541	Expat is a C library for parsing XML documents.
RHSA-2024:9541	Expat is a C library for parsing XML documents.
RHSA-2024:9553	WebKitGTK is the port of the portable web rendering engine WebKit to the GTK platform.
RHSA-2024:9559	The libsoup packages provide an HTTP client and server library for GNOME.
RHSA-2024:9605	The kernel packages contain the Linux kernel, the core of any Linux operating system.
TEMP-0290435-0B57B5	tar's rmt command may have undesired side effects
TEMP-0517018-A83CE6	sysvinit: no-root option in expert installer exposes locally exploitable security flaw
TEMP-0601525-BEBB65	libgd2: gdImageColorTransparent can write outside buffer
TEMP-0628843-DBAD28	more related to CVE-2005-4890
TEMP-0841856-B18BAF	Privilege escalation possible to other user than root

Cloudera Data Services on premises 1.5.4 SP2

The service packs for new features, known issues, and fixed issues for 1.5.4 SP2.

Note: ECS Customers: Direct upgrade path is not available for customers currently on Cloudera Data Services on premises 1.5.2. Customers must upgrade to Cloudera Data Services on premises 1.5.4 prior to consuming any Service Packs (SPs) built on top of 1.5.4.

Note:

OCP Customers: Direct upgrade path is available. Customers can directly upgrade from Cloudera Data Services on premises 1.5.2 to any 1.5.4 Cumulative Hotfixes (CHFs) or Service Packs (SPs).

Attention: Cloudera Data Services on premises 1.5.4 is not supported on Cloudera Base on premises 7.3.1. You must not install or upgrade to Cloudera Base on premises 7.3.1, if you are using Cloudera Data Services on premises 1.5.4 on your cluster as it is incompatible.

Note: You must use Cloudera Manager 7.13.1 CHF1 only to install or upgrade to Cloudera Data Services on premises 1.5.4 SP2. That is, ensure to install Cloudera Manager 7.13.1 CHF1 before upgrading or installing Cloudera Data Services on premises 1.5.4 SP2.

Certifications in 1.5.4 SP2

The following are the certifications supported in 1.5.4 SP2:

- Cloudera Base on premises 7.1.9 SP1, 7.1.9 CHF7 (7.1.9.14), 7.1.7 SP3 (7.1.7.3000)
- Cloudera Manager 7.13.1 CHF1 and later
- Iceberg v2 GA on Cloudera Data Warehouse, Cloudera Data Engineering, & Cloudera AI with Ozone
- OEL (RHCK Kernel Only) 8.7, 8.8, 8.9, 8.10, 9.1, 9.2, 9.3, 9.4
- RHEL 8.7, 8.8, 8.9, 8.10, 9.1, 9.2, 9.3, 9.4, 9.5
- K8s 1.29 and OCP 4.16
- Longhorn version 1.6.2

What's new in Cloudera Data Services on premises 1.5.4 SP2

New features introduced in this service pack release of Cloudera Data Services on premises 1.5.4 SP2.

Pre-upgrade checklist

There is a list of pre-upgrade checks that will run after the upgrade version has been chosen. This checklist verifies if your cluster is ready for upgrade. For more information, see Pre-Upgrade Checklist .

Note: Cloudera Manager 7.13.1 CHF1 support Cloudera Data Services on premises 1.5.4 SP2 release.

Cloudera Manager 7.11.3 CHF8 does not support any Cloudera Data Services on premises release.

Kubernetes Ingress NGINX Controller vulnerabilities

Kubernetes Ingress NGINX Controller vulnerabilities.

Five vulnerabilities affecting the Ingress Nginx Controller for Kubernetes were publicly disclosed on March 24, 2025, and were given the nickname "IngressNightmare."

The ‘IngressNightmare’ vulnerabilities may allow Remote Code Execution (RCE) and potentially expose Kubernetes clusters to malicious configuration modifications. Exploitation requires specially crafted HTTP requests that bypass security measures, such as a Web Application Firewall (WAF). Successful exploitation may lead to complete cluster compromise, data exfiltration, and denial of service.

Details of the CVEs:

- CVE-2025-1974 (CVSS score: 9.8) – An unauthenticated attacker with access to the pod network can achieve arbitrary code execution in the context of the ingress-nginx controller under certain conditions
- CVE-2025-24514 (CVSS score: 8.8) – The auth-url Ingress annotation can be used to inject configuration into NGINX, resulting in arbitrary code execution in the context of the ingress-nginx controller and disclosure of secrets accessible to the controller
- CVE-2025-1097 (CVSS score: 8.8) – The auth-tls-match-cn Ingress annotation can be used to inject configuration into NGINX, resulting in arbitrary code execution in the context of the ingress-nginx controller and disclosure of secrets accessible to the controller
- CVE-2025-1098 (CVSS score: 8.8) – The mirror-target and mirror-host Ingress annotations can be used to inject arbitrary configuration into NGINX, resulting in arbitrary code execution in the context of the ingress-nginx controller and disclosure of secrets accessible to the controller
- CVE-2025-24513 (CVSS score: 4.8) – An improper input validation vulnerability that could result in directory traversal within the container, leading to denial-of-service (DoS) or limited disclosure of secret objects from the cluster when combined with other vulnerabilities

Note: Cloudera will not provide remediation options for unsupported versions, and has not tested mitigations on unsupported versions. Customers are advised to upgrade to a supported product version. For more information, refer to the Support Lifecycle Policy page.

Action required

Multiple mitigation scenarios are possible depending on the version of Cloudera Data Services deployed.

1. Check your Data Services version.

   To determine which version of Data Services is running, log into Cloudera Manager UI, click “Data Services” in the left sidebar, and find each Data Services cluster's version(s).
2. Check your Cloudera Manager version.

   To determine which version of Cloudera Manager is running, log into Cloudera Manager UI and click on the version number at the lower-left corner of the home page. The Cloudera Manager 7.11.3 CHF8 Build number shows 56304673.

Mitigation steps for Cloudera Data Services On Premises:

1. No action is needed in the following scenarios:
   - Running Data Services on OpenShift Container Platform (OCP) and using the default HAproxy ingress controller. If you have customized OCP to use NGINX as the ingress controller, please contact RedHat support for further information
   - Running Data Services 1.5.4 or later with Cloudera Manager 7.11.3 CHF8 or later. These versions already disable the affected admission webhook component.

2. Data Services before 1.5.4, or Cloudera Manager prior to 7.11.3 CHF8, perform one of the following:
   - (Recommended) Upgrade Cloudera Manager to version 7.13.1 CHF1 or higher and all Embedded Container Service clusters to 1.5.4 SP2 or higher. This will not upgrade the nginx version but will automatically disable the webhook via Cloudera Manager, which mitigates the CVEs.
   - If an upgrade is not possible, follow the below steps to disable the webhook to mitigate the CVEs manually.

     Note: These steps should be run on each ECS cluster. If you create a new ECS cluster while on an affected release, reapply these steps on the new cluster(s).

1. SSH into the ECS server host.
2. Create the following YAML file to patch the built-in Nginx Ingress Controller Helm Chart:

```
apiVersion: helm.cattle.io/v1
kind: HelmChartConfig
metadata:
 name: rke2-ingress-nginx
 namespace: kube-system
spec:
 valuesContent: |-
  controller:
   admissionWebhooks:
    enabled: false
```

3. Run the following commands to apply the patch:

```
alias kubectl=/var/lib/rancher/rke2/bin/kubectl
export KUBECONFIG=/etc/rancher/rke2/rke2.yaml
kubectl apply -f <PATCH_FILENAME>.yaml
```

Note: No cluster restart or ECS refresh is required.

Note: Cloudera will not provide remediation options for unsupported versions, and has not tested mitigations on unsupported versions. Customers are advised to upgrade to a supported product version. For more information, refer to the Support Lifecycle Policy page

Cloudera Data Services on premises version 1.5.4 SP2 CHF1 will contain fixes to upgrade the Ingress NGINX controller. Target release date: April 25. 2025.

For the latest update on this issue, see the corresponding Knowledge article: TSB 2025-839: Critical Kubernetes Ingress NGINX Controller Vulnerability Allows RCE Without Authentication

For the latest update on this issue, see the corresponding Data Services Knowledge article: https://my.cloudera.com/knowledge/TSB-2025-839-Mitigation-steps-for-Cloudera-Data-Services-on?id=405006

Apache Parquet CVE-2025-30065

On April 1, 2025, a critical vulnerability in the parquet-avro module of Apache Parquet (CVE-2025-30065, CVSS score 10.0) was announced.

Cloudera has determined the list of affected products, and is issuing this TSB to provide details of remediation for affected versions.

Upgraded versions are being released for all currently affected supported releases of Cloudera products. Customers using older versions are advised to upgrade to a supported release that has the remediation, once it becomes available.

Vulnerability Details

Exploiting this vulnerability is only possible by modifying the accepted schema used for translating Parquet files and subsequently submitting a specifically crafted malicious file.

CVE-2025-30065 | **Schema parsing in the parquet-avro module of Apache Parquet 1.15.0 and previous versions allows bad actors to execute arbitrary code.**

CVE: NVD - CVE-2025-30065

Severity (Critical): CVSS:4.0/AV:N/AC:L/AT:N/PR:N/UI:N/VC:H/VI:H/VA:H/SC:H/SI:H/SA:H

Impact

Schema parsing in the parquet-avro module of Apache Parquet 1.15.0 and previous versions allows bad actors to execute arbitrary code. Attackers may be able to modify unexpected objects or data that was assumed to be safe from modification. Deserialized data or code could be modified without using the provided accessor functions, or unexpected functions could be invoked.

Deserialization vulnerabilities most commonly lead to undefined behavior, such as memory modification or remote code execution.

Releases affected

Cloudera Data Services on premises

- All versions

Mitigation

Until Cloudera has released product version with the Apache Parquet vulnerability fix, please continue to use the the mitigations listed below:

Customers with their own FIM Solution:

1. Utilize a File Integrity Monitoring (FIM) solution. This allows administrators to monitor files at the filesystem level and receive alerts on any unexpected or suspicious activity in the schema configuration.

General advisory:

1. Use network segmentation and traffic monitoring with a device capable of deep packet inspection, such as a network firewall or web application firewall, to inspect all traffic sent to the affected endpoints.
2. Configure alerts for any suspicious or unexpected activity. You may also configure sample analysis parameters to include:
   - Parquet file format “magic bytes” = PAR1
   - Connections from sending hosts that are not expected source IP ranges.
3. Be cautious with Parquet files from unknown or untrusted sources. If possible, do not process files with uncertain origins or that can be ingested from outside the organization.
4. Ensure that only authorized users have access to endpoints that ingest Parquet files.

For the latest update on this issue, see the corresponding Knowledge article: Cloudera Customer Advisory 2025-847: Cloudera's remediation actions for Apache Parquet CVE-2025-30065

Known Issues in Cloudera Data Services on premises 1.5.4 SP2

The following are the known issues in the 1.5.4 SP2 release of Cloudera Data Services on premises.

DWX-20809: Cloudera Data Services on premises installations on RHEL 8.9 or lower versions may encounter issues

You may notice issues when installing Cloudera Data Services on premises on Cloudera Embedded Container Service clusters running on RHEL 8.9 or lower versions. Pod crashloops are noticed with the following error:

```
Warning  FailedCreatePodSandBox          1s (x2 over 4s)  kubel
et   Failed to create pod sandbox: rpc error: code = Unknown des
c = failed to create containerd task: failed to create shim task:
 OCI runtime create
 failed: runc create failed: unable to start container process:
 unable to init seccomp: error loading seccomp filter into kerne
l: error loading seccomp filter: errno 524: unknown
```

The issue is due to a memory leak with 'seccomp' (Secure Computing Mode) in the Linux kernel. If your kernel version is not on 6.2 or higher verisons or if it is not part of the list of versions mentioned here, you may face issues during installation.

To avoid this issue, increase the value of net.core.bpf_jit_limit by running the following command on all ECS hosts:

```
[root@host ~]# sysctl net.core.bpf_jit_limit=528482304
```

However, Cloudera recommends upgrading the Linux kernel to an appropriate version that contains a patch for the memory leak issue. For a list of versions that contain this patch, see this link.

OPSX-5776 and OPSX-5747: ECS - Some of the rke2-canal DaemonSet pods in the kube-system namespace are stuck in Init state causing longhorn volume attach issues

In a few cases, after upgrading from Cloudera Data Services on premises 1.5.3 or 1.5.3-CHF1 to 1.5.4 SP2, a pod that belongs to `rke2-canal DaemonSet` is stuck in Init status. This causes some pods in kube-system and longhorn-system namespaces to be in `Init` or `CrashLoopBackOff` status. This manifests as volume attach failure in the embedded-db-0 pod in CDP namespace, and causes some pods in CDP namespace to be in `CrashLoopBackOff` state.

1. Perform a rolling restart of `rke2-canal DaemonSet` by running the following command:

   ```
   kubectl rollout restart ds rke2-canal -n kube-system
   ```

2. Monitor the `DaemonSet` restart status by running the following command:

   ```
   kubectl get ds -n kube-system
   ```

3. After the `rke2-canal DaemonSet` restart is complete, if any pods in `DaemonSets` within the longhorn-system namespace remain in Init or CrashLoopBackOff state, perform a rolling restart of those `DaemonSets`. Choose the appropriate command based on the specific `DaemonSet` that is failing. If more than one `DaemonSet` requires a restart, restart them sequentially, one at a time.

   ```
   kubectl rollout restart ds longhorn-csi-plugin -n longhorn-s
   ystem
   kubectl rollout restart ds longhorn-manager -n longhorn-system
   kubectl rollout restart ds engine-image-ei-6b4330bf -n long
   horn-system
   kubectl rollout restart ds engine-image-ei-ea8e2e58 -n longh
   orn-system
   ```

4. Monitor the `DaemonSet` restart status by running the following command:

   ```
   kubectl get ds -n longhorn-system
   ```

   Note: After the `DaemonSet` restart completes, it can take another ten minutes for pods in `longhorn-system` and `CDP namespaces` to come back to Running status.

OPSX-5903: 1.5.3 to 1.5.4 SP2 ECS - The upgrade fails when the `rke2-ingress-nginx-controller` system exceeds its progress deadline.

Upgrade fails while running the following command:

```
kubectl rollout status deployment/rke2-ingress-nginx-controller -
n kube-system --timeout=5m
```

Run the following command:

```
kubectl rollout status deployment/rke2-ingress-nginx-controller -
n kube-system --timeout=5m
```

Run the command for `Refresh ECS` by performing the following steps:

1. Find the number of replicas for the deployment:

   ```
   kubectl get deployment -n kube-system rke2-ingress-nginx-con
   troller
   ```

2. Scale it down to 0:

   ```
   kubectl scale deployment rke2-ingress-nginx-controller -n ku
   be-system --replicas=0
   ```

3. Once all associated pods are terminated scale it back:

   ```
   kubectl scale deployment rke2-ingress-nginx-controller -n ku
   be-system --replicas=<n>
   ```

4. Resume upgrade.

OPSAPS-72270: Start ECS command fails on uncordon nodes step

In an ECS HA cluster sometimes, the server node restarts during start up. This causes the uncordon step to fail.

Run the following command on the same node to verify whether the kube-apiserver is ready:

```
kubectl get pods -n kube-system | grep kube-apiserver
```

Resume the command from the Cloudera Manager UI.

OPSAPS-72964 and OPSAPS-72769: Unseal Vault command fails after restarting the ECS service

Unseal Vault command fails sometimes, after restarting the ECS service.

It may take sometime for the ECS cluster to be up and running after a restart operation. In case `Unseal Vault` fails after the restart operation please follow the below steps:

1. Verify that the pod `vault-0` in the vault-system namespace is running.
2. Once it is in Running state, initiate the `Unseal Vault` command from the ECS service's Action menu.

OPSX-5986: ECS fresh install failing with helm-install-rke2-ingress-nginx pod failing to come into Completed state

ECS fresh install fails at the "Execute command Reapply All Settings to Cluster on service ECS" step due to a timeout waiting for `helm-install`. To confirm the issue, run the following `kubectl` command on the ECS server host to check if the pod is stuck in a running state:

```
kubectl get pods -n kube-system | grep helm-install-rke2-ingress-
nginx
```

To resolve the issue, manually delete the pod by running the following command:

```
kubectl delete pod <helm-install-rke2-ingress-nginx-pod-name> -n
 kube-system
```

Then, click Resume to proceed with the fresh install process on the Cloudera Manager UI.

Repository Locations for 1.5.4 SP2

The URLs for Cloudera Data Services on premises 1.5.4 SP2 are listed in the following table:

URL Type	Repository Location
Index	`https://username:password@archive.cloudera.com/p/cdp-pvc-ds/1.5.4-h15/`
Manifest	Repository: `https://username:password@archive.cloudera.com/p/cdp-pvc-ds/1.5.4-h15/manifest.json`
Parcels	Repository: `https://username:password@archive.cloudera.com/p/cdp-pvc-ds/1.5.4-h15/parcels/`

Fixed Common Vulnerabilities and Exposures in 1.5.4 SP2

Review the Common vulnerabilities and Exposures (CVEs) that were fixed in 1.5.4 SP2 release of Cloudera Data Services on premises.

Issue ID	Description
CVE-2024-21236	Vulnerability in the MySQL Server product of Oracle MySQL (component: InnoDB). Supported versions that are a and 9.0.1 and prior. Easily exploitable vulnerability allows high privileged attacker with network access via multip Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repeatable cr CVSS 3.1 Base Score 4.9 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C:N/I:N/
CVE-2024-21237	Vulnerability in the MySQL Server product of Oracle MySQL (component: Server: Group Replication GCS). Supp and prior, 8.4.2 and prior and 9.0.1 and prior. Difficult to exploit vulnerability allows high privileged attacker with compromise MySQL Server. Successful attacks of this vulnerability can result in unauthorized ability to cause a pa MySQL Server. CVSS 3.1 Base Score 2.2 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:H/PR:H/UI
CVE-2024-21238	Vulnerability in the MySQL Server product of Oracle MySQL (component: Server: Thread Pooling). Supported ve 8.4.1 and prior and 9.0.1 and prior. Difficult to exploit vulnerability allows low privileged attacker with network ac MySQL Server. Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequer MySQL Server. CVSS 3.1 Base Score 5.3 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:H/PR:L/UI
CVE-2024-21239	Vulnerability in the MySQL Server product of Oracle MySQL (component: InnoDB). Supported versions that are a and 9.0.1 and prior. Easily exploitable vulnerability allows high privileged attacker with network access via multip Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repeatable cr CVSS 3.1 Base Score 4.9 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C:N/I:N/
CVE-2024-21241	Vulnerability in the MySQL Server product of Oracle MySQL (component: Server: Optimizer). Supported version and prior and 9.0.1 and prior. Easily exploitable vulnerability allows high privileged attacker with network access v Server. Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repea Server. CVSS 3.1 Base Score 4.9 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C
CVE-2024-21247	Vulnerability in the MySQL Client product of Oracle MySQL (component: Client: mysqldump). Supported versio 8.4.2 and prior and 9.0.1 and prior. Easily exploitable vulnerability allows high privileged attacker with network ac MySQL Client. Successful attacks of this vulnerability can result in unauthorized update, insert or delete access to well as unauthorized read access to a subset of MySQL Client accessible data. CVSS 3.1 Base Score 3.8 (Confiden (CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C:L/I:L/A:N).

CVE-2024-23454	Apache Hadoop,Äôs RunJar.run()¬†does not set permissions for temporary directory¬†by default. If sensitive data users may be able to view the content. This is because, on unix-like systems, the system temporary directory is shar written in this directory, without setting the correct posix permissions explicitly, may be viewable by all other local
CVE-2024-27289	pgx is a PostgreSQL driver and toolkit for Go. Prior to version 4.18.2, SQL injection can occur when all of the follo simple protocol is used; a placeholder for a numeric value must be immediately preceded by a minus; there must be the first placeholder; both must be on the same line; and both parameter values must be user-controlled. The proble do not use the simple protocol or do not place a minus directly before a placeholder.
CVE-2024-27304	pgx is a PostgreSQL driver and toolkit for Go. SQL injection can occur if an attacker can cause a single query or bi integer overflow in the calculated message size can cause the one large message to be sent as multiple messages un resolved in v4.18.2 and v5.5.4. As a workaround, reject user input large enough to cause a single query or bind mes
CVE-2024-28863	node-tar is a Tar for Node.js. node-tar prior to version 6.2.1 has no limit on the number of sub-folders created in the generates a large number of sub-folders can consume memory on the system running node-tar and even crash the N using a path with too many sub-folders inside. Version 6.2.1 fixes this issue by preventing extraction in excessively
CVE-2024-29018	Moby is an open source container framework that is a key component of Docker Engine, Docker Desktop, and othe runtimes. Moby's networking implementation allows for many networks, each with their own IP address range and frequently referred to as custom networks, as each network can have a different driver, set of parameters and thus b internal` flag is used to designate a network as _internal_. The `internal` attribute in a docker-compose.yml file may and other API clients may specify the `internal` parameter as well. When containers with networking are created, th and IP addresses. The host serves as a router for non-internal networks, with a gateway IP that provides SNAT/DN internal network may communicate between each other, but are precluded from communicating with any networks no default route is configured, and firewall rules are set up to drop all outgoing traffic. Communication with the ga configured host services) is possible, and the host may communicate with any container IP directly. In addition to c networking features to enable container networking, `dockerd` directly provides some services to container network resolver, enabling service discovery, and resolution of names from an upstream resolver. When a DNS request for a is received, the request is forwarded to the configured upstream resolver. This request is made from the container's routing of traffic is the same as if the request was made by the container itself. As a consequence of this design, con will be unable to resolve names using the upstream resolver, as the container itself is unable to communicate with t also attached to the internal network are able to be resolved. Many systems run a local forwarding DNS resolver. A loopback devices, a consequence of the design described above is that containers are unable to resolve names from reach these addresses on the host loopback device. To bridge this gap, and to allow containers to properly resolve n is used on a loopback address, `dockerd` detects this scenario and instead forward DNS requests from the host nam then forwards the requests to its configured upstream resolvers, as expected. Because `dockerd` forwards DNS requ the container network namespace's normal routing semantics entirely, internal networks can unexpectedly forward registering a domain for which they control the authoritative nameservers, an attacker could arrange for a compromi it in DNS queries that will eventually be answered by their nameservers. Docker Desktop is not affected, as Docker a RFC 1918 address. Moby releases 26.0.0, 25.0.4, and 23.0.11 are patched to prevent forwarding any DNS reques run containers intended to be solely attached to internal networks with a custom upstream address, which will force from the container's network namespace.
CVE-2024-29415	The ip package through 2.0.1 for Node.js might allow SSRF because some IP addresses (such as 127.1, 012000345 and ::fFFf:127.0.0.1) are improperly categorized as globally routable via isPublic. NOTE: this issue exists because
CVE-2024-33655	The DNS protocol in RFC 1035 and updates allows remote attackers to cause a denial of service (resource consump accumulated for seconds, such that responses are later sent in a pulsing burst (which can be considered traffic ampl issue.
CVE-2024-34750	Improper Handling of Exceptional Conditions, Uncontrolled Resource Consumption vulnerability in Apache Tomc Tomcat did not handle some cases of excessive HTTP headers correctly. This led to a miscounting of active HTTP/ incorrect infinite timeout which allowed connections to remain open which should have been closed. This issue aff 11.0.0-M20, from 10.1.0-M1 through 10.1.24, from 9.0.0-M1 through 9.0.89. Users are recommended to upgrade t which fixes the issue.
CVE-2024-36129	The OpenTelemetry Collector offers a vendor-agnostic implementation on how to receive, process and export telem vulnerability allows unauthenticated attackers to crash the collector via excessive memory consumption. OTel Coll also fixed in the confighttp module version 0.102.0 and configgrpc module version 0.102.1.
CVE-2024-38821	Spring WebFlux applications that have Spring Security authorization rules on static resources can be bypassed und an application, all of the following must be true: * It must be a WebFlux application * It must be using Spring's sta permitAll authorization rule applied to the static resources support
CVE-2024-42368	OpenTelemetry, also known as OTel, is a vendor-neutral open source Observability framework for instrumenting, g telemetry data such as traces, metrics, and logs. The bearertokenauth extension's server authenticator performs a sim of the received & configured bearer tokens. This impacts anyone using the `bearertokenauth` server authenticator. I collector may perform a timing attack against a collector with this authenticator to guess the configured token, by it response time. This would allow an attacker to introduce fabricated or bad data into the collector's telemetry pipelin fixed by using constant-time comparison in 0.107.0

CVE-2024-43167	DISPUTE NOTE: this issue does not pose a security risk as it (according to analysis by the original software develo functionality and security controls of the application. Red Hat has made a claim that there is a security risk within t further information about the claim, and suggests that affected Red Hat customers refer to available Red Hat docun DESCRIPTION: A NULL pointer dereference flaw was found in the ub_ctx_set_fwd function in Unbound. This is specific sequences of API calls to cause a segmentation fault. When certain API functions such as ub_ctx_set_fwd particular order, the program attempts to read from a NULL pointer, leading to a crash. This issue can result in a de terminate unexpectedly.
CVE-2024-43168	DISPUTE NOTE: this issue does not pose a security risk as it (according to analysis by the original software develo functionality and security controls of the application. Red Hat has made a claim that there is a security risk within t further information about the claim, and suggests that affected Red Hat customers refer to available Red Hat docun DESCRIPTION: A heap-buffer-overflow flaw was found in the cfg_mark_ports function within Unbound's config_ This issue could allow an attacker with local access to provide specially crafted input, potentially causing the applic execution. This could result in a denial of service or unauthorized actions on the system.
CVE-2024-45043	The OpenTelemetry Collector module AWS firehose receiver is for ingesting AWS Kinesis Data Firehose delivery records received based on the configured record type. `awsfirehosereceiver` allows unauthenticated remote request OpenTelemetry Collector can be configured to receive CloudWatch metrics via an AWS Firehose Stream. Firehose Key` with an arbitrary configured string. The OpenTelemetry Collector awsfirehosereceiver can optionally be conf requests. However, when this is configured it **still accepts incoming requests with no key**. Only OpenTelemetr ,Äúalpha,Äù `awsfirehosereceiver` module are affected. This module was added in version v0.49.0 of the ,ÄúCont custom builds). There is a risk of unauthorized users writing metrics. Carefully crafted metrics could hide other ma data. It,Äôs likely these endpoints will be exposed to the public internet, as Firehose does not support private HTTP #34847 and released with v0.108.0. All users are advised to upgrade. There are no known workarounds for this vul
CVE-2024-5206	A sensitive data leakage vulnerability was identified in scikit-learn's TfidfVectorizer, specifically in versions up to in version 1.5.0. The vulnerability arises from the unexpected storage of all tokens present in the training data with only storing the subset of tokens required for the TF-IDF technique to function. This behavior leads to the potential `stop_words_` attribute could contain tokens that were meant to be discarded and not stored, such as passwords or based on the nature of the data being processed by the vectorizer.
CVE-2024-52303	aiohttp is an asynchronous HTTP client/server framework for asyncio and Python. In versions starting with 3.10.6 occur when a request produces a MatchInfoError. This was caused by adding an entry to a cache on each request, d producing a unique cache entry. An attacker may be able to exhaust the memory resources of a server by sending a such requests. Those who use any middlewares with aiohttp.web should upgrade to version 3.10.11 to receive a pat
CVE-2024-54661	readline.sh in socat before1.8.0.2 relies on the /tmp/$USER/stderr2 file.
CVE-2024-9341	A flaw was found in Go. When FIPS mode is enabled on a system, container runtimes may incorrectly handle certa the containers/common Go library. This flaw allows an attacker to exploit symbolic links and trick the system into container. This issue also allows attackers to access critical host files, bypassing the intended isolation between con
CVE-2024-9823	There exists a security vulnerability in Jetty's DosFilter which can be exploited by unauthorized users to cause remo server using DosFilter. By repeatedly sending crafted requests, attackers can trigger OutofMemory errors and exha
CVE-2025-0291	Type Confusion in V8 in Google Chrome prior to 131.0.6778.264 allowed a remote attacker to execute arbitrary co page. (Chromium security severity: High)
CVE-2025-0434	Out of bounds memory access in V8 in Google Chrome prior to 132.0.6834.83 allowed a remote attacker to potenti HTML page. (Chromium security severity: High)
CVE-2025-0435	Inappropriate implementation in Navigation in Google Chrome on Android prior to 132.0.6834.83 allowed a remot crafted HTML page. (Chromium security severity: High)
CVE-2025-0436	Integer overflow in Skia in Google Chrome prior to 132.0.6834.83 allowed a remote attacker to potentially exploit (Chromium security severity: High)
CVE-2025-0437	Out of bounds read in Metrics in Google Chrome prior to 132.0.6834.83 allowed a remote attacker to potentially ex page. (Chromium security severity: High)
CVE-2025-0438	Stack buffer overflow in Tracing in Google Chrome prior to 132.0.6834.83 allowed a remote attacker to potentially page. (Chromium security severity: High)
CVE-2025-0439	Race in Frames in Google Chrome prior to 132.0.6834.83 allowed a remote attacker who convinced a user to engag spoofing via a crafted HTML page. (Chromium security severity: Medium)
CVE-2025-0440	Inappropriate implementation in Fullscreen in Google Chrome on Windows prior to 132.0.6834.83 allowed a remo crafted HTML page. (Chromium security severity: Medium)
CVE-2025-0441	Inappropriate implementation in Fenced Frames in Google Chrome prior to 132.0.6834.83 allowed a remote attack from the system via a crafted HTML page. (Chromium security severity: Medium)
CVE-2025-0442	Inappropriate implementation in Payments in Google Chrome prior to 132.0.6834.83 allowed a remote attacker wh gestures to perform UI spoofing via a crafted HTML page. (Chromium security severity: Medium)

CVE-2025-0443	Insufficient data validation in Extensions in Google Chrome prior to 132.0.6834.83 allowed a remote attacker who gestures to perform privilege escalation via a crafted HTML page. (Chromium security severity: Medium)
CVE-2025-0444	Use after free in Skia in Google Chrome prior to 133.0.6943.53 allowed a remote attacker to potentially exploit hea (Chromium security severity: High)
CVE-2025-0445	Use after free in V8 in Google Chrome prior to 133.0.6943.53 allowed a remote attacker to potentially exploit heap (Chromium security severity: High)
CVE-2025-0446	Inappropriate implementation in Extensions in Google Chrome prior to 132.0.6834.83 allowed a remote attacker w gestures to perform UI spoofing via a crafted Chrome Extension. (Chromium security severity: Low)
CVE-2025-0447	Inappropriate implementation in Navigation in Google Chrome prior to 132.0.6834.83 allowed a remote attacker to HTML page. (Chromium security severity: Low)
CVE-2025-0448	Inappropriate implementation in Compositing in Google Chrome prior to 132.0.6834.83 allowed a remote attacker page. (Chromium security severity: Low)
CVE-2025-0451	Inappropriate implementation in Extensions API in Google Chrome prior to 133.0.6943.53 allowed a remote attack UI gestures to perform UI spoofing via a crafted Chrome Extension. (Chromium security severity: Medium)
CVE-2025-0611	Object corruption in V8 in Google Chrome prior to 132.0.6834.110 allowed a remote attacker to potentially exploit (Chromium security severity: High)
CVE-2025-0612	Out of bounds memory access in V8 in Google Chrome prior to 132.0.6834.110 allowed a remote attacker to poten HTML page. (Chromium security severity: High)
CVE-2025-0762	Use after free in DevTools in Google Chrome prior to 132.0.6834.159 allowed a remote attacker to potentially expl Extension. (Chromium security severity: Medium)
CVE-2025-21490	Vulnerability in the MySQL Server product of Oracle MySQL (component: InnoDB). Supported versions that are a and 9.1.0 and prior. Easily exploitable vulnerability allows high privileged attacker with network access via multip Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repeatable cr CVSS 3.1 Base Score 4.9 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C:N/I:N/
CVE-2025-21491	Vulnerability in the MySQL Server product of Oracle MySQL (component: InnoDB). Supported versions that are a and 9.1.0 and prior. Easily exploitable vulnerability allows high privileged attacker with network access via multip Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repeatable cr CVSS 3.1 Base Score 4.9 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C:N/I:N/
CVE-2025-21492	Vulnerability in the MySQL Server product of Oracle MySQL (component: Server: Optimizer). Supported version 8.4.0. Easily exploitable vulnerability allows high privileged attacker with network access via multiple protocols to attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repeatable crash (compl Score 4.9 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C:N/I:N/A:H).
CVE-2025-21494	Vulnerability in the MySQL Server product of Oracle MySQL (component: Server: Security: Privileges). Supporte prior, 8.4.2 and prior and 9.0.1 and prior. Difficult to exploit vulnerability allows high privileged attacker with logo executes to compromise MySQL Server. Successful attacks of this vulnerability can result in unauthorized ability t (complete DOS) of MySQL Server. CVSS 3.1 Base Score 4.1 (Availability impacts). CVSS Vector: (CVSS:3.1/AV
CVE-2025-21497	Vulnerability in the MySQL Server product of Oracle MySQL (component: InnoDB). Supported versions that are a and 9.1.0 and prior. Easily exploitable vulnerability allows high privileged attacker with network access via multip Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repeatable cr well as unauthorized update, insert or delete access to some of MySQL Server accessible data. CVSS 3.1 Base Sco CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C:N/I:L/A:H).
CVE-2025-21500	Vulnerability in the MySQL Server product of Oracle MySQL (component: Server: Optimizer). Supported version and prior and 9.1.0 and prior. Easily exploitable vulnerability allows low privileged attacker with network access v Server. Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repea Server. CVSS 3.1 Base Score 6.5 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:U/C
CVE-2025-21501	Vulnerability in the MySQL Server product of Oracle MySQL (component: Server: Optimizer). Supported version and prior and 9.1.0 and prior. Easily exploitable vulnerability allows low privileged attacker with network access v Server. Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repea Server. CVSS 3.1 Base Score 6.5 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:U/C
CVE-2025-21503	Vulnerability in the MySQL Server product of Oracle MySQL (component: InnoDB). Supported versions that are a and 9.1.0 and prior. Easily exploitable vulnerability allows high privileged attacker with network access via multip Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repeatable cr CVSS 3.1 Base Score 4.9 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C:N/I:N/

CVE-2025-21504	Vulnerability in the MySQL Server product of Oracle MySQL (component: Server: Optimizer). Supported version and prior and 9.0.1 and prior. Easily exploitable vulnerability allows high privileged attacker with network access v Server. Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repea Server. CVSS 3.1 Base Score 4.9 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C
CVE-2025-21505	Vulnerability in the MySQL Server product of Oracle MySQL (component: Server: Components Services). Suppor and prior, 8.4.3 and prior and 9.1.0 and prior. Easily exploitable vulnerability allows high privileged attacker with r compromise MySQL Server. Successful attacks of this vulnerability can result in unauthorized ability to cause a ha DOS) of MySQL Server. CVSS 3.1 Base Score 4.9 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/
CVE-2025-21518	Vulnerability in the MySQL Server product of Oracle MySQL (component: Server: Optimizer). Supported version and prior and 9.1.0 and prior. Easily exploitable vulnerability allows low privileged attacker with network access v Server. Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repea Server. CVSS 3.1 Base Score 6.5 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:U/C
CVE-2025-21519	Vulnerability in the MySQL Server product of Oracle MySQL (component: Server: Security: Privileges). Supporte and prior, 8.4.3 and prior and 9.1.0 and prior. Difficult to exploit vulnerability allows high privileged attacker with compromise MySQL Server. Successful attacks of this vulnerability can result in unauthorized ability to cause a ha DOS) of MySQL Server. CVSS 3.1 Base Score 4.4 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:H/
CVE-2025-21520	Vulnerability in the MySQL Server product of Oracle MySQL (component: Server: Options). Supported versions t and prior and 9.1.0 and prior. Difficult to exploit vulnerability allows high privileged attacker with logon to the infr to compromise MySQL Server. Successful attacks require human interaction from a person other than the attacker. result in unauthorized read access to a subset of MySQL Server accessible data. CVSS 3.1 Base Score 1.8 (Confide AV:L/AC:H/PR:H/UI:R/S:U/C:L/I:N/A:N).
CVE-2025-21521	Vulnerability in the MySQL Server product of Oracle MySQL (component: Server: Thread Pooling). Supported ve 8.4.2 and prior and 9.0.1 and prior. Easily exploitable vulnerability allows unauthenticated attacker with network ac MySQL Server. Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequer MySQL Server. CVSS 3.1 Base Score 7.5 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:N/UI
CVE-2025-21522	Vulnerability in the MySQL Server product of Oracle MySQL (component: Server: Parser). Supported versions tha prior and 9.1.0 and prior. Easily exploitable vulnerability allows low privileged attacker with network access via m Server. Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repea Server. CVSS 3.1 Base Score 6.5 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:U/C
CVE-2025-21523	Vulnerability in the MySQL Server product of Oracle MySQL (component: InnoDB). Supported versions that are a and 9.1.0 and prior. Easily exploitable vulnerability allows high privileged attacker with network access via multip Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repeatable cr CVSS 3.1 Base Score 4.9 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C:N/I:N/
CVE-2025-21525	Vulnerability in the MySQL Server product of Oracle MySQL (component: Server: DDL). Supported versions that prior and 9.0.1 and prior. Easily exploitable vulnerability allows high privileged attacker with network access via n Server. Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repea Server. CVSS 3.1 Base Score 4.9 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C
CVE-2025-21529	Vulnerability in the MySQL Server product of Oracle MySQL (component: Server: Information Schema). Support and prior, 8.4.3 and prior and 9.1.0 and prior. Easily exploitable vulnerability allows high privileged attacker with r compromise MySQL Server. Successful attacks of this vulnerability can result in unauthorized ability to cause a ha DOS) of MySQL Server. CVSS 3.1 Base Score 4.9 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/
CVE-2025-21531	Vulnerability in the MySQL Server product of Oracle MySQL (component: InnoDB). Supported versions that are a and 9.1.0 and prior. Easily exploitable vulnerability allows high privileged attacker with network access via multip Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repeatable cr CVSS 3.1 Base Score 4.9 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C:N/I:N/
CVE-2025-21534	Vulnerability in the MySQL Server product of Oracle MySQL (component: Server: Performance Schema). Suppor and prior, 8.4.2 and prior and 9.0.1 and prior. Easily exploitable vulnerability allows high privileged attacker with r compromise MySQL Server. Successful attacks of this vulnerability can result in unauthorized ability to cause a ha DOS) of MySQL Server. CVSS 3.1 Base Score 4.9 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/
CVE-2025-21536	Vulnerability in the MySQL Server product of Oracle MySQL (component: Server: Optimizer). Supported version and prior and 9.0.1 and prior. Easily exploitable vulnerability allows high privileged attacker with network access v Server. Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repea Server. CVSS 3.1 Base Score 4.9 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C
CVE-2025-21540	Vulnerability in the MySQL Server product of Oracle MySQL (component: Server: Security: Privileges). Supporte and prior, 8.4.3 and prior and 9.1.0 and prior. Easily exploitable vulnerability allows low privileged attacker with n compromise MySQL Server. Successful attacks of this vulnerability can result in unauthorized update, insert or del accessible data as well as unauthorized read access to a subset of MySQL Server accessible data. CVSS 3.1 Base S impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:U/C:L/I:L/A:N).

CVE-2025-21543	Vulnerability in the MySQL Server product of Oracle MySQL (component: Server: Packaging). Supported version and prior and 9.1.0 and prior. Easily exploitable vulnerability allows high privileged attacker with network access v Server. Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repea Server. CVSS 3.1 Base Score 4.9 (Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C
CVE-2025-21546	Vulnerability in the MySQL Server product of Oracle MySQL (component: Server: Security: Privileges). Supporte and prior, 8.4.3 and prior and 9.1.0 and prior. Easily exploitable vulnerability allows high privileged attacker with r to compromise MySQL Server. Successful attacks of this vulnerability can result in unauthorized update, insert or accessible data as well as unauthorized read access to a subset of MySQL Server accessible data. CVSS 3.1 Base S impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C:L/I:L/A:N).
CVE-2025-21555	Vulnerability in the MySQL Server product of Oracle MySQL (component: InnoDB). Supported versions that are a and 9.1.0 and prior. Easily exploitable vulnerability allows high privileged attacker with network access via multip Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repeatable cr well as unauthorized update, insert or delete access to some of MySQL Server accessible data. CVSS 3.1 Base Sco CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C:N/I:L/A:H).
CVE-2025-21559	Vulnerability in the MySQL Server product of Oracle MySQL (component: InnoDB). Supported versions that are a and 9.1.0 and prior. Easily exploitable vulnerability allows high privileged attacker with network access via multip Successful attacks of this vulnerability can result in unauthorized ability to cause a hang or frequently repeatable cr well as unauthorized update, insert or delete access to some of MySQL Server accessible data. CVSS 3.1 Base Sco CVSS Vector: (CVSS:3.1/AV:N/AC:L/PR:H/UI:N/S:U/C:N/I:L/A:H).
CVE-2025-23083	With the aid of the diagnostics_channel utility, an event can be hooked into whenever a worker thread is created. T exposes internal workers, where an instance of them can be fetched, and its constructor can be grabbed and reinstat affects Permission Model users (--permission) on Node.js v20, v22, and v23.
CVE-2025-23090	With the aid of the diagnostics_channel utility, an event can be hooked into whenever a worker thread is created. T exposes internal workers, where an instance of them can be fetched, and its constructor can be grabbed and reinstat affects Permission Model users (--permission) on Node.js v20, v22, and v23.
GHSA-2w8w-qhg4-f78j	Related UI vulnerability advisory: https://github.com/jaegertracing/jaeger-ui/security/advisories/GHSA-vv24-rm95 `json-markup` dependency to display span attributes and resources. This dependency is not sanitising keys of an ob vulnerable to XSS. ### Details
GHSA-74fp-r6jw-h4mp	CVE-2019-11253 is a denial of service vulnerability in the kube-apiserver, allowing authorized users sending malic kube-apiserver to consume excessive CPU or memory, potentially crashing and becoming unavailable. When creat references contained in it, excessive CPU usage can occur. This appears to be an instance of a "Billion Laughs" atta parsing issue. Applying this manifest to a cluster causes the client to hang for some time with considerable CPU us
GHSA-7jwh-3vrq-q3m8	### Impact SQL injection can occur if an attacker can cause a single query or bind message to exceed 4 GB in size message size can cause the one large message to be sent as multiple messages under the attacker's control. ### Patc
GHSA-7ww5-4wqc-m92c	# /sys/devices/virtual/powercap accessible by default to containers Intel's RAPL (Running Average Power Limit) f microarchitecture, provides software insights into hardware energy consumption. To facilitate this, Intel introduced 3.13, which reads values via relevant MSRs (model specific registers) and provides unprivileged userspace access access a hardware feature, it is only available when running on bare metal with the module compiled into the kerne cases unprivileged access to RAPL readings could be exploited as a power-based side-channel against security feat SGX enclave) and KASLR (kernel address space layout randomization). Also known as the [PLATYPUS attack](h [CVE-2020-8694](https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2020-8694) and [CVE-2020-8695](https name=CVE-2020-8695), and AMD assigned [CVE-2020-12912](https://cve.mitre.org/cgi-bin/cvename.cgi?name=
GHSA-87m9-rv8p-rgmg	### Impact A malicious user could cause a denial of service (DoS) when using a specially crafted gRPC request. T respect the limits imposed by gRPC, allowing rapid memory usage increases. Versions v1.1.4 through to v1.2.2 ma github.com/klauspost/compress/zstd to decompress data provided by the peer. The vulnerability is exploitable only users of github.com/mostynb/go-grpc-compression/zstd or github.com/mostynb/go-grpc-compression/nonclobberin
GHSA-c9cp-9c75-9v8c	### Impact A bug was found in containerd where containers were incorrectly started with non-empty inheritable Li Linux environment and enabling programs with inheritable file capabilities to elevate those capabilities to the perm executable programs have specified permitted file capabilities, otherwise unprivileged users and processes can exec file capabilities up to the bounding set. Due to this bug, containers which included executable programs with inheri unprivileged users and processes to additionally gain these inheritable file capabilities up to the container's boundin groups to perform privilege separation inside the container are most directly impacted. This bug did not affect the c set never contained more capabilities than were included in the container's bounding set.

GHSA-jq35-85cj-fj4p	Intel's RAPL (Running Average Power Limit) feature, introduced by the Sandy Bridge microarchitecture, provides consumption. To facilitate this, Intel introduced the powercap framework in Linux kernel 3.13, which reads values registers) and provides unprivileged userspace access via `sysfs`. As RAPL is an interface to access a hardware fea on bare metal with the module compiled into the kernel. By 2019, it was realized that in some cases unprivileged a exploited as a power-based side-channel against security features including AES-NI (potentially inside a SGX encl layout randomization). Also known as the [PLATYPUS attack](https://platypusattack.com/), Intel assigned [CVE-2 bin/cvename.cgi?name=CVE-2020-8694) and [CVE-2020-8695](https://cve.mitre.org/cgi-bin/cvename.cgi?name= [CVE-2020-12912](https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2020-12912). Several mitigations were via a microcode update, and the Linux kernel [prevents access by non-root users](https://git.kernel.org/pub/scm/lin id=949dd0104c496fa7c14991a23c03c62e44637e71) since 5.10. However, this kernel-based mitigation does not ap Unless using user namespaces, root inside a container has the same level of privilege as root outside the container, system
GHSA-mh55-gqvf-xfwm	Middleware causes a prohibitive amount of heap allocations when processing malicious preflight requests that incl (ACRH) header whose value contains many commas. This behavior can be abused by attackers to produce undue l cause a denial of service.
GHSA-mhpq-9638-x6pw	An attacker controlled input of a PBES2 encrypted JWE blob can have a very large p2c value that, when decrypted
GHSA-qq97-vm5h-rrhg	### Impact Systems that rely on digest equivalence for image attestations may be vulnerable to type confusion. ###
GHSA-vp35-85q5-9f25	### Description Moby is the open source Linux container runtime and set of components used to build a variety of Docker CE, Mirantis Container Runtime (formerly Docker EE), and Docker Desktop. Moby allows for building co instructions (usually named and referred to as a "Dockerfile"), and a build context, which is not unlike the CWD in executed. Containers may be built using a variety of tools and build backends available in the Moby ecosystem; in of the build context (such as using absolute or relative-parent paths). This is enforced through both checks in the bu build process itself.
RHSA-2024:10289	The container-tools module contains tools for working with containers, notably podman, buildah, skopeo, and runc
RHSA-2024:10289	The container-tools module contains tools for working with containers, notably podman, buildah, skopeo, and runc
RHSA-2024:10289	The container-tools module contains tools for working with containers, notably podman, buildah, skopeo, and runc
RHSA-2024:10953	Python is an interpreted, interactive, object-oriented programming language, which includes modules, classes, exce and dynamic typing. Python supports interfaces to many system calls and libraries, as well as to various windowing
RHSA-2024:4563	The java-1.8.0-openjdk packages provide the OpenJDK 8 Java Runtime Environment and the OpenJDK 8 Java Sof
RHSA-2024:8846	The container-tools module contains tools for working with containers, notably podman, buildah, skopeo, and runc
RHSA-2024:8846	The container-tools module contains tools for working with containers, notably podman, buildah, skopeo, and runc
RHSA-2024:8846	The container-tools module contains tools for working with containers, notably podman, buildah, skopeo, and runc
RHSA-2024:9573	The libsoup packages provide an HTTP client and server library for GNOME.
RHSA-2025:0837	The unbound packages provide a validating, recursive, and caching DNS or DNSSEC resolver.
RHSA-2025:0837	The unbound packages provide a validating, recursive, and caching DNS or DNSSEC resolver.
RHSA-2025:0838	The libsoup packages provide an HTTP client and server library for GNOME.
RHSA-2025:1372	The container-tools module contains tools for working with containers, notably podman, buildah, skopeo, and runc
RHSA-2025:1372	The container-tools module contains tools for working with containers, notably podman, buildah, skopeo, and runc
RHSA-2025:1372	The container-tools module contains tools for working with containers, notably podman, buildah, skopeo, and runc

Cumulative hotfixes: Cloudera Data Services on premises 1.5.4 SP2 CHF1

Review the fixes in the Cloudera Data Services on premises 1.5.4 Service Pack 2 Cumulative hotfix 1 release.

To know more about fixes in this release, see these release notes:

- Cloudera Data Warehouse on premise 1.5.4 SP2 CHF1
- Cloudera AI on premises 1.5.4 SP2 CHF1

Repository locations for 1.5.4 SP2 CHF1

The URLs for Cloudera Data Services on premises 1.5.4 SP2 Cumulative hotfix 1 release are listed in the following table:

URL type	Repository location
Index	`https://username:password@archive.cloudera.com/p/cdp-pvc-ds/1.5.4-h30/`
Manifest	`https://username:password@archive.cloudera.com/p/cdp-pvc-ds/1.5.4-h30/manifest.json`
Parcels	`https://username:password@archive.cloudera.com/p/cdp-pvc-ds/1.5.4-h30/parcels/`