Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.3 Release Notes

Date published: 2020-07-28
Date modified: 2026-08-31



https://docs.cloudera.com/

Legal Notice

© Cloudera Inc. 2026. 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

Overview 6

What's new in Platform Support 6

Runtime API Compatibility 6
- Introduction to the API compatibility 6
- Client-server compatibility 7

What's new in Cloudera Runtime 7.3.3 8
- Atlas 8
- What's new in Cloud Connectors 8
- Cruise Control 8
- HBase 8
- HDFS 8
- Hive 9
- Cloudera Data Explorer (Hue) 9
- What's new in Iceberg REST Catalog 9
- Iceberg 9
- Impala 9
- Kafka 10
- Knox 10
- Kudu 10
- Livy 10
- Navigator Encrypt 10
- Oozie 11
- Ozone 11
- Phoenix 11
- Ranger 11
- What's new in Ranger KMS 11
- Schema Registry 11
- Solr 11
- Spark 11
- Spark Atlas Connector 11
- Sqoop 11
- Streams Messaging Manager 11
- Streams Replication Manager 12
- YARN 12
- Zookeeper 12

Fixed issues in Cloudera Runtime 7.3.3 12
- Atlas 12
- Avro 13
- Cloud Connectors 13
- Cruise Control 13
- Fixed issues in Iceberg REST Catalog 13

HDFS....13
HBase....13
Hive....13
Cloudera Data Explorer (Hue)....14
Impala....15
Fixed issues in Apache Iceberg....16
Kafka....16
Kudu....16
Knox....17
Livy....17
Navigator Encrypt....17
Oozie....18
Ozone....18
Parquet....18
Phoenix....18
Ranger....18
Ranger KMS....18
Schema Registry....18
Solr....19
Spark....19
Spark Atlas Connector....19
Sqoop....19
Streams Messaging Manager....19
Streams Replication Manager....19
Yarn and Yarn Queue Manager....19
ZooKeeper....19

Known issues in Cloudera Runtime 7.3.3....19

Atlas....20
Avro....28
Known Issues in Cloud Connectors....28
Cruise Control....28
HBase....29
HDFS....29
Hive....29
Cloudera Data Explorer (Hue)....32
Iceberg....34
Known Issues in Iceberg REST Catalog....38
Impala....38
Kafka....44
Knox....48
Kudu....51
Livy....51
Navigator Encrypt....52
Oozie....52
Ozone....53
Phoenix....62
Ranger....62
Ranger KMS....64
Schema Registry....64
Solr....65
Spark....67
Spark Atlas Connector....68
Sqoop....68
Streams Messaging Manager....69

Streams Replication Manager....70
Yarn and Yarn Queue Manager....72
ZooKeeper....74

Behavioral changes in Cloudera Runtime 7.3.3.... 75
Behavioral Changes in Atlas....75
Behavioral Changes in Avro....75
Behavioral Changes in Cloud Connectors....75
Behavioral Changes in Cruise Control....75
Behavioral changes in HBase....75
Behavioral Changes in HDFS....75
Behavioral Changes in Hive....76
Behavioral Changes in Cloudera Data Explorer (Hue)....76
Behavioral Changes in Impala....76
Behavioral Changes in Apache Iceberg....77
Behavioral Changes in Iceberg REST Catalog....77
Behavioral Changes in Kafka....77
Behavioral Changes in Apache Knox....77
Behavioral Changes in Kudu....77
Behavioral Changes in Livy....77
Behavioral Changes in Oozie....77
Behavioral Changes in Ozone....78
Behavioral Changes in Parquet....78
Behavioral Changes in Phoenix....78
Behavioral Changes in Apache Ranger....78
Schema Registry....78
Behavioral Changes in Solr....79
Behavioral Changes in Spark....79
Behavioral Changes in Spark Atlas Connector....79
Behavioral Changes in Sqoop....79
Behavioral Changes in Streams Messaging Manager....79
Behavioral Changes in Streams Replication Manager....79
Behavioral Changes in YARN....79
Behavioral changes in ZooKeeper....79

Release Matrix....80

Components versions in Cloudera Runtime 7.3.3.... 80

Fixed Common Vulnerabilities and Exposures in Cloudera Runtime 7.3.3....81

Using the Cloudera Runtime Maven repository....143
Cloudera Runtime 7.3.3.0-56....143

Overview

This document provides a summary of the latest updates in Cloudera Runtime 7.3.3. It includes new features, improvements, known and fixed issues, technical previews, and more.

What's new in Platform Support

Cloudera Runtime 7.3.3 release, introduces expanded platform support across Cloudera Manager 7.13.3. Review these additions to plan your infrastructure upgrades and take advantage of newly supported operating systems, databases, Python versions, and Java Development Kits (JDKs).

Cloudera Runtime 7.3.3

Platform support enhancements

None

For more information about the operating system support, see Cloudera Support Matrix.

New database support

None

New JDK support

None

Runtime API Compatibility

Runtime API compatibility defines the architectural standards for communication between software components ensuring backward compatibility and consistent client-server interoperability as systems evolve.

Introduction to the API compatibility

API compatibility ensures seamless interaction between client and server versions by maintaining backward compatibility, and REST protocol consistency to prevent breaking changes across Cloudera Runtime releases.

The API defines a set of rules and definitions for the client-server communication. The client requests services or resources from the server, which processes the request and returns a response. API compatibility allows different client and server versions to interact successfully. Typically, a client supports its corresponding API version and lower backward-compatible versions.

Ensuring API compatibility has the following benefits:

- Changes are not broken: Avoid removing, renaming, or changing mandatory parameters within the software code.
- Versioning: Implement API versioning (for example, /version1/, /version2/) to introduce changes without affecting existing users and their usage patterns.
- Deprecation policy: Call out old APIs as deprecated before removing them.
- Validation opportunities: Verify changes that introducing new code does not break compatibility with lower client versions.

Supported touchpoints for API compatibility

Seamless API integration enables different software systems to interact and exchange data automatically, to ensure real-time synchronization, reduced manual effort, and improved efficiency.

Cloudera supports the following touchpoints to ensure forward and backward- API compatibility across different Cloudera runtime releases:

- Rest APIs
- Command-line interfaces (CLIs)
- Client RPM and Debian packages
- Client-server interactions (Atlas, Ranger, Data Hubs, and Data services)

Client-server compatibility

Client-server API compatibility in Cloudera Enterprise Data Platform ensures seamless interactions between components, even with updates to the client or the server.

Cloudera Enterprise Data Platform involves multiple layers for which API compatibility is critical. Each component in the Cloudera stack follows its own way of exchanging data using the API compatibility touchpoints, and the coverage is only being enhanced for the compatibility options.

Multiple Cloudera Runtime releases complement each other to maintain consistent API interactions.

The following table lists the supported Cloudera Runtime releases that are forward- and backward-compatible:

Table 1: Client-server compatibility

Client	Server
7.3.2.0	7.3.2.0
7.3.1.500	7.3.2.0
CDH: 7.3.1.500-176	7.3.2.0
7.1.9 SP1 CHFx (Latest)	7.3.2.0
	7.3.2.0
	7.3.2.0
	7.3.2.0
	7.3.2.0
	7.3.1.500/400
	7.2.18 latest SP
	7.2.17 latest SP
	7.1.9 Service Pack 1 CHFx
	7.1.7 Service Pack 1 CHFx

For more information about each of these releases, see the Cloudera Product documentation.

Supported Cloudera Runtime components

The following table lists the Runtime components and their supported touchpoints for forward and backward API compatibility across Cloudera runtime releases:

Table 2: Supported Runtime component compatibility

Runtime component	REST / Thrift APIs	Client libraries and SDKs	Client-server interaction (Atlas, Ranger, Data Hubs, and Data services)	CLI
Atlas	Yes	Yes	Yes	N/A
Ranger	Yes	Yes	Yes	N/A
Ranger KMS	Yes	Yes	Yes	N/A
RAZ	Yes	Yes	Yes	N/A
Knox	Yes	Yes	Yes	N/A
HBase	N/A	Yes	N/A	Yes
HDFS		Yes	Yes	Yes
Phoenix	Yes	N/A	N/A	N/A
Ozone	Yes	Yes	Yes	Yes
Hive	Yes	Yes	N/A	Yes
Impala	Yes	Yes	Yes	Yes
Livy	Yes	N/A	N/A	N/A
Spark	Yes	Yes	Yes	Yes

What's new in Cloudera Runtime 7.3.3

This version of Cloudera Runtime provides you with several new capabilities. Learn how the new features and improvements benefit you.

What's new in Apache Atlas

New features and functional updates for Atlas are introduced in Cloudera Runtime 7.3.3.

What's new in Cloud Connectors

Learn about the new features of Cloud Connectors in Cloudera Runtime 7.3.3.

What's new in Cruise Control

New features and functional updates for Cruise Control are introduced in Cloudera Runtime 7.3.3.

What's new in Apache HBase

New features and functional updates for HBase are introduced in Cloudera Runtime 7.3.3.

What's new in HDFS

Learn about the new features of HDFS in Cloudera Runtime 7.3.3.

What's new in Apache Hive

New features and functional updates for Hive are introduced in Cloudera Runtime 7.3.3.

What's new in Cloudera Data Explorer (Hue)

New features and functional updates for Cloudera Data Explorer (Hue) are introduced in Cloudera Runtime 7.3.3.

Cloudera Runtime 7.3.3

SSL certificate verification for Ranger Authorization Service (RAZ)

Data Explorer now supports custom Certificate Authority certificate validation when communicating with the Ranger Authorization Service. You can configure dedicated CA bundles or selectively disable certificate verification for troubleshooting without modifying global cluster security settings.

For more information, see SSL certificate verification for RAZ.

What's new in Iceberg REST Catalog

New features and functional updates for Cloudera Iceberg REST Catalog are introduced in Cloudera Runtime 7.3.3.

What's new in Apache Iceberg

Learn about the new features of Iceberg in Cloudera Runtime 7.3.3.

What's new in Apache Impala

New features and functional updates for Impala are introduced in Cloudera Runtime 7.3.3.

Cloudera Runtime 7.3.3

Cloudera Runtime 7.3.3 introduces enhanced authentication capabilities for JWT tokens and support for multiple identity providers. These updates strengthen security and enable more flexible authentication scenarios.

Impala Executor Groups for workload isolation

You can now assign Impala daemons to executor groups for workload isolation with the Cloudera Manager property Impala Daemon Executor Group (impalad_executor_group). For more information, see Configuring Impala Executor Groups and Isolating workloads with Impala executor groups.

JWT Audience and Issuer Claim Validation with Multiple Identity Provider Support

Impala 7.3.3 introduces enhanced JWT/OAuth authentication with two major capabilities:

- JWT Claim Validation: Validate issuer (iss) and audience (aud) claims per OAuth provider using issuerClaims and audienceClaims arrays within --oauth_servers JSON configuration. Prevents token redirection attacks by ensuring tokens come from trusted authorities and are intended for Impala.
- Multiple Identity Providers: Support configuring and trusting multiple IDPs simultaneously in a single cluster through the unified --oauth_servers framework. Enables federated authentication in multi-cloud, multi-tenant, and hybrid deployments.

For configuration details, see JWT Audience and Issuer Claim Validation and Support for Multiple Token Identity Providers.

Apache Jira: IMPALA-12232

What's new in Apache Kafka

New features and functional updates for Kafka are introduced in Cloudera Runtime 7.3.3.

What's new in Apache Knox

New features and functional updates for Apache Knox are introduced in Cloudera Runtime 7.3.3.

What's new in Apache Kudu

New features and functional updates for Kudu are introduced in Cloudera Runtime 7.3.3.

What's new in Livy

New features and functional updates for Livy are introduced in Cloudera Runtime 7.3.3.

Cloudera Runtime 7.3.3

Livy load balancing with HAProxy

You can run several Livy for Spark 3 service instances behind HAProxy to increase the number of Livy sessions that the cluster can run. Each service instance keeps its own active/passive high availability. HAProxy distributes new sessions between the service instances and uses a cookie to keep each session on the instance that created it. Clients connect to the load balancer URL that you set in the Livy Load Balancer Address property of each service instance.

For more information, see Livy load balancing with HAProxy.

/ha-status endpoint for Livy high availability

Livy provides a GET /ha-status endpoint that reports whether a Livy server is the leader of its service instance. The leader returns HTTP 200 and passive servers return HTTP 503, so a load balancer such as HAProxy can use the endpoint as a health check and send requests only to the leader.

For more information, see Enabling the Livy /ha-status endpoint.

Sticky sessions for Livy Thrift Server

You can connect to Livy Thrift Server through HAProxy. In http transport mode, the JDBC client sends the SRV_ID cookie of the Livy server that owns the session, and HAProxy routes the connection to that server.

For more information, see Connecting to the Apache Livy Thrift Server.

Default YARN queue for Livy sessions

The livy.spark.yarn.queue property sets a default YARN queue for all sessions and batches of a Livy service instance. A request can still select another queue with its queue field.

For more information, see Livy session capacity and YARN queues.

What's new in Navigator Encrypt

Learn about the new features of Navigator Encrypt in Cloudera Runtime 7.3.3.

What's new in Oozie

Learn about the new features of Oozie in Cloudera Runtime 7.3.3.

What's new in Ozone

New features and functional updates for Ozone are introduced in Cloudera Runtime 7.3.3.

What's new in Apache Phoenix

New features and functional updates for Phoenix are introduced in Cloudera Runtime 7.3.3.

What's new in Apache Ranger

Learn about the new features of Apache Ranger in Cloudera Runtime 7.3.3.

What's new in Ranger KMS

Learn about the new features of Ranger KMS in Cloudera Runtime 7.3.3.

What's new in Schema Registry

New features and functional updates for Schema Registry are introduced in Cloudera Runtime 7.3.3.

What's new in Apache Solr

New features and functional updates for Solr are introduced in Cloudera Runtime 7.3.3.

What's new in Spark

Learn about the new features of Spark in Cloudera Runtime 7.3.3.

Related Information

Upgrading Spark

Migrating Spark applications

What's new in Spark Atlas Connector

New features and functional updates for Spark Atlas Connector are introduced in Cloudera Runtime 7.3.3.

What's new in Sqoop

New features and functional updates for Sqoop are introduced in Cloudera Runtime 7.3.3.

What's new in Streams Messaging Manager

New features and functional updates for Streams Messaging Manager are introduced in Cloudera Runtime 7.3.3.

What's new in Streams Replication Manager

New features and functional updates for Streams Replication Manager are introduced in Cloudera Runtime 7.3.3.

What's new in Yarn and Yarn Queue Manager

New features and functional updates for Yarn and Yarn Queue Manager are introduced in Cloudera Runtime 7.3.3.

What's new in Apache ZooKeeper

New features and functional updates for ZooKeeper are introduced in Cloudera Runtime 7.3.3.

Fixed issues in Cloudera Runtime 7.3.3

Fixed issues represent issues reported by Cloudera customers that are addressed in this release.

Fixed issues in Atlas

Fixed issues for Atlas are addressed in Cloudera Runtime 7.3.3.

Cloudera Runtime 7.3.3

CDPD-109807: Async import queue permanently blocked after JanusGraph transaction failure

7.3.3

When a JanusGraph transaction commit failed during asynchronous metadata import, the import queue could remain blocked and later import requests stayed in the WAITING state. Atlas now guarantees that import semaphores are released and requests are completed or failed reliably, even if persisting the failed state throws an exception or the executor is unavailable, ensuring the queue can process subsequent work.

Apache JIRA: ATLAS-5345

CDPD-106349: Atlas UI: Port Apache Atlas React UI upstream fixes

7.3.3

The React-based Atlas user interface now includes upstream fixes for the responsive header and Quick Search layout, glossary nested subcategories and subcategory creation, an upgraded sani tize-html dependency to mitigate cross-site scripting risk, Type System filter and lineage graph improvements, audit pagination and entity payload enrichment, frontend dependency updates, classification tooltip and search results table behavior, classification filter reset without HTTP 400 errors, Create Entity navigation on the entity detail page, text overflow in search filters and modals, audit filter toolbar layout, and build optimizations with lazy route loading.

Apache JIRA: ATLAS-5297, ATLAS-5298, ATLAS-5299, ATLAS-5301, ATLAS-5303, ATLAS-5305, ATLAS-5306, ATLAS-5328, ATLAS-5329, ATLAS-5330, ATLAS-5333, ATLAS-5339, ATLAS-5346

CDPD-110585: Atlas JanusGraph integration fails on JDK 21

7.3.3

Atlas JanusGraph integration failed on JDK 21 because reflective access to modify JanusGraph store manager internals threw an IllegalAccessException due to stricter field access restrictions. The fix restores JDK 21 compatibility by introducing a fallback utility that uses sun.misc.Unsafe to inject custom backend managers when standard reflection is blocked.

OPSAPS-78886: Atlas: Register new Single EKU parameters in the CM upgrade handler

7.13.2.30000

When you upgrade a cluster, Cloudera Manager now initializes new Atlas client keystore and server truststore configuration parameters for single end-entity certificate (Single EKU) support. On upgrade, new keys default to the existing dual end-entity certificate paths so behavior remains unchanged until you configure the split stores.

Fixed issues in Apache Avro

Cloudera Runtime 7.3.3 includes no bug fixes for Avro.

Fixed issues in Cloud Connectors

Cloudera Runtime 7.3.3 includes no bug fixes for Cloud Connectors.

Fixed issues in Cruise Control

Cloudera Runtime 7.3.3 includes no bug fixes for Cruise Control.

Fixed issues in Iceberg REST Catalog

Cloudera Runtime 7.3.3 includes no bug fixes for Iceberg REST Catalog.

Fixed issues in HDFS

Cloudera Runtime 7.3.3 includes no bug fixes for HDFS.

Fixed issues in HBase

Cloudera Runtime 7.3.3 includes no bug fixes for HBase.

Fixed issues in Hive

Fixed issues for Hive are addressed in Cloudera Runtime 7.3.3.

Cloudera Runtime 7.3.3

CDPD-94256: Expand COMMENT column length in COLUMNS_V2 table

Previously, the COMMENT column in the HMS COLUMNS_V2 table was limited to 256 characters, which restricted customers from adding comprehensive context to Hive table columns.

This issue is resolved by expanding the COMMENT column to 4 KB to allow longer column descriptions. This enables customers to add detailed context to each column of their Hive tables for use by AI agents and other tools that generate SQL and reports based on column metadata.

Important: If the backend database is MySQL and the COMMENT column does not use the latin1 character set, you must back up the COLUMNS_V2 table before upgrading. Otherwise, the COMMENT column data might become corrupted and cannot be reverted. In this case, you must modify the MySQL ALTER TABLE statement according to the active character set. For example, if the COMMENT column uses the UTF-8 character set, run the following commands:

```
ALTER TABLE COLUMNS_V2 MODIFY COLUMN COMMENT varchar
(4000) CHARACTER SET utf8mb4 COLLATE utf8mb4_bin DEF
AULT NULL, ALGORITHM=COPY;
                -- Update the schema version
                UPDATE CDH_VERSION SET SCHEMA_VERSION='
3.1.3000.7.3.2.0-Update2', VERSION_COMMENT='Hive rel
ease version 3.1.3000 for CDH 7.3.2.0-Update2' WHERE
 VER_ID=1;
```

Fixed issues in Cloudera Data Explorer (Hue)

Fixed issues for Cloudera Data Explorer (Hue) are addressed in Cloudera Runtime 7.3.3.

Cloudera Runtime 7.3.3

CDPD-125832: Storage Browser file download failure in Data Explorer

7.3.3.0

Previously, file downloads and previews failed when you used the Storage Browser in Data Explorer behind a Knox gateway. This issue occurred because generated download and preview URLs omitted the required gateway proxy prefix, preventing requests from reaching the service. This issue is resolved by appending the proxy path prefix to download and preview URLs.

CDPD-125184: Oozie workflow submission failure in Data Explorer

7.3.3.0

Previously, submitting Oozie workflows from the Storage Browser in Data Explorer failed. This issue occurred because workflow submission omitted external parameter names. This issue is resolved by preserving workflow variable parameters, correcting processing, and refining route query parameters during submission.

CDPD-124478: Error when accessing the Impala tab in the Jobs browser

7.3.3.0

Previously, an error occurred when you clicked a query ID in the editor to open the query profile page if your environment uses an active-passive coordinator setup. This issue exists as a result of changes made to support active-active coordinator configurations, preventing access to detailed query metrics in active-passive environments. This issue is now resolved.

CDPD-124161: TLS hostname validation failure in Data Explorer

7.3.3.0

Previously, Impala connections in Data Explorer running on Python 3.12 failed TLS hostname validation when coordinator fully qualified domain names (FQDNs) exceeded 64 characters or when connections used HAProxy aliases. This issue occurred because hostname validation did not evaluate Subject Alternative Names (SAN) defined in valid certificates. This issue is resolved by updating hostname validation logic in Data Explorer to validate SAN entries on Python 3.12.

CDPD-123398: Gunicorn server log rotation failure in Data Explorer

7.3.3.0

Previously, rungunicornserver.log files in Data Explorer rotated prematurely at irregular file sizes well below the configured 4.8 MB threshold. This behavior caused overlapping timestamps and

fragmented entries across rotated log files. This issue is resolved by improving log file handling across worker processes to ensure log files rotate at the configured file size threshold.

CDPD-101348: Data Explorer administration UI leads to incorrect URL

7.3.3.0

Previously, the Download logs button in the Data Explorer administration UI pointed to an incorrect or outdated URL, which prevented from downloading server logs. This issue is now resolved by updating the log download API endpoints.

CDPD-98406: Excessive Gunicorn debug log generation

7.3.3.0

Previously, Gunicorn log listeners generated debug-level records for access and metrics even when Data Explorer debug mode was disabled. This behavior caused unnecessary log clutter. This issue is now resolved by ensuring Gunicorn log output honors the Data Explorer debug mode setting.

CDPD-96460: Impala query profile row count display error

7.3.3.0

Previously, the Data Explorer Impala editor component displayed incorrect row count values for large tables because it incorrectly parsed the RowsReturned metric property from the EXCHANGE node component inside the Impala query profile log. This issue is fixed by implementing a threshold-based bucketing approach that displays approximate row count indicators, such as 1000+ rows or 10000+ rows, to prevent inaccurate query metadata displays.

CDPD-95948: Orphaned process detection failure in Data Explorer

7.3.3.0

Previously, process cleanup in Data Explorer incorrectly flagged active processes as orphaned, causing valid service processes to stop unexpectedly during cleanup routines. This issue is resolved by improving the detection logic to prevent false positives during process management cleanup.

CDPD-69906: Impala file import failure in Data Explorer

7.3.3.0

Previously, importing files into Impala by using Data Explorer failed on high-speed network connections. This issue occurred because rapid automated query progression stalled during sequential import operations, and result size requests were incorrectly triggered for data definition statements. This issue is resolved by improving query progression handling during automated imports and bypassing result size checks for statements without active query handles.

CDPD-125940: Storage Browser refresh failure in Data Explorer

Previously, refreshing Amazon S3 directories in the Data Explorer Storage Browser failed. This issue occurred because a name collision existed between the fs query parameter in the frontend and a parameter in the backend middleware. This issue is resolved by renaming the frontend query parameter from fs to tab.

CDPD-125833: Zero-byte file upload failure in Data Explorer

Previously, uploading zero-byte files to cloud storage by using the Storage Browser in Data Explorer failed with an XML error. This issue occurred because the upload process initiated multipart upload operations that Amazon S3 rejects when no parts are submitted. This issue is resolved by writing zero-byte files through a single put operation instead of a multipart upload, and removing the zero-byte file size check from the client-side.

Fixed issues in Impala

Fixed issues for Impala are addressed in Cloudera Runtime 7.3.3.

Cloudera Runtime 7.3.3

CDPD-123566: Geospatial functions unavailable when configured via safety valve

Previously, setting --geospatial_library=HIVE_ESRI in the Impala Service Advanced Configuration Snippet (Safety Valve) did not properly register geospatial functions. After restart, geospatial functions like ST_GeomFromText() were reported as unknown and unavailable in the database, and the Impala service entered a bad health state.

This issue is resolved by properly applying the geospatial library configuration when the Impala service restarts.

Cloudera Jira: CDPD-123566

Fixed issues in Apache Iceberg

Cloudera Runtime 7.3.3 includes no bug fixes for Iceberg.

Fixed issues in Kafka

Cloudera Runtime 7.3.3 includes no bug fixes for Kafka.

Fixed issues in Kudu

Fixed issues for Kudu are addressed in Cloudera Runtime 7.3.3.

Cloudera Runtime 7.3.3

CDPD-105447: Buffer overflow during message discard operation

Previously, discarding subprocess messages larger than 4 KB caused a buffer overflow or corrupted subsequent messages on the communication channel. This issue led to channel desynchronization and potential process failures during message handling. This issue is now resolved. The message discard process correctly bounds read operations to prevent memory corruption and maintain proper channel synchronization.

CDPD-124500: Kudu compaction budget degradation during scans

Previously, Kudu included entire UNDO delta sizes in the compaction budget calculation, which caused premature skipping of rowsets during compaction cycles. This behavior is disabled by default and can be enabled at runtime with --rowset_deltas_size_include_undo on clusters that have encountered such failures. Because enabling it can cause large rowsets to be skipped, and scans to slow down as a result, new metrics (rowsets_per_scan, compaction_budget_skip_health, tablets_num_compaction_budget_degraded) and a warning log help detect the problem.

Apache Jira: KUDU-3734

CDPD-124508: Out-of-bounds read error in range predicates

Previously, passing incorrect column types in deprecated range predicates caused out-of-bounds read errors on the server during scan operations. This issue is now resolved by adding column type validation to range predicates to prevent out-of-bounds read attempts.

CDPD-106582: Ranger subprocess logging failure in Kudu

Previously, the Ranger authorization subprocess in Kudu failed to record log entries, which resulted in empty log files during runtime operations. This issue is now fixed.

CDPD-106045: Tablet server crash during count scans

Previously, running COUNT queries with predicates caused tablet servers in Kudu to crash unexpectedly. This issue occurred because scan operations without projected columns encountered an internal memory allocation error when evaluating predicates. This issue is resolved by refining memory allocation calculations and safely handling predicates during counting scan operations.

Apache Jira: KUDU-3782

CDPD-124568: Application crash during range partition writes

Previously, applications using the Kudu C++ client crashed with a segmentation fault when writing data to newly created range partitions that previously belonged to a non-covered range. This issue is resolved by correcting tablet metadata access during verbose logging, preventing client application crashes during write operations.

Apache Jira: KUDU-3704

CDPD-104260: Hive Metastore Thrift client TLS support limitation

Previously, Kudu did not support TLS encryption in the Hive Metastore Thrift client, which prevented secure network communication between Kudu servers and HMS. This issue is now resolved by enabling TLS transport support in the Hive Metastore Thrift client.

Fixed issues in Apache Knox

Fixed issues for Apache Knox are addressed in Cloudera Runtime 7.3.3.

Cloudera Runtime 7.3.3

CDPD-124438: Apache Knox Gateway UIs load the gear and pin icons from the internet, breaking icons in air-gapped environments

7.3.3.0

Fixed an issue where, in air-gapped networks and in environments that block outbound internet access from Knox UI client browsers, Apache Knox Gateway UIs tried to load the gear and pin icons from the internet. The browser could not download the icons, so they appeared as literal ligature text instead of glyphs. Knox now bundles these icons as local inline SVGs so the UIs no longer request them from the internet, and icons render correctly when outbound internet access is blocked.

Apache JIRA: KNOX-3508

CDPD-124561: Apache Knox Gateway default login page flashes before modern-themed login page is displayed

7.3.3.0

Fixed an issue where the Apache Knox Gateway default login page flashed briefly before the modern-themed login page appeared. The modern-themed login page now displays correctly.

Apache JIRA: KNOX-3506

CDPD-97987: AWS SDK for Java v2 log messages are missing from Knox IDBroker logs

7.3.3.0

Fixed an issue where AWS SDK for Java v2 log messages were missing from Knox IDBroker logs when IDBroker performed AWS credential brokering. Knox no longer uses the AWS SDK bundle fat JAR, which included its own SLF4J implementation that overrode Knox logging configuration and suppressed log output from AWS SDK classes. Knox now uses individual AWS SDK modules, and AWS SDK log messages appear in IDBroker logs as expected.

Fixed issues in Livy

Cloudera Runtime 7.3.3 includes no bug fixes for Livy.

Fixed issues in Navigator Encrypt

Cloudera Runtime 7.3.3 includes no bug fixes for Navigator Encrypt.

Fixed issues in Oozie

Cloudera Runtime 7.3.3 includes no bug fixes for Oozie.

Fixed issues in Ozone

Cloudera Runtime 7.3.3 includes no bug fixes for Ozone.

Fixed issues in Apache Parquet

Cloudera Runtime 7.3.3 includes no bug fixes for Parquet.

Fixed issues in Phoenix

Fixed issues for Phoenix are addressed in Cloudera Runtime 7.3.3.

Cloudera Runtime 7.3.3

CDPD-124333: Schema extraction fails for salted tables with reordered primary key columns

Previously, the SHOW CREATE TABLE command and the Phoenix schema extraction tool failed with a NullPointerException error on a salted table if the order of the primary key columns in the CREATE TABLE statement was different from the order in the PRIMARY KEY constraint.

This issue is now resolved.

Apache Jira: PHOENIX-7468

Fixed issues in Ranger

Fixed issues for Ranger are addressed in Cloudera Runtime 7.3.3.

Cloudera Runtime 7.3.3

CDPD-122535: Kafka broker OutOfMemoryError during Streams Replication Manager startup

7.3.3

Previously, when a user or service principal, such as the Streams Replication Manager (SRM) user, had a very large number of Ranger roles assigned, the Kafka broker could run into an OutOfMem oryError during authorization when SRM started. This happened because the Ranger Kafka plugin copied the user's full role set into every access request in a batch authorization call, rather than sharing the resolved roles across the batch.

This issue is now resolved. All access requests in a single Kafka authorization call now share one resolved role set, which greatly reduces the broker's heap memory use.

Fixed issues in Ranger KMS

Cloudera Runtime 7.3.3 includes no bug fixes for Ranger KMS.

Fixed issues in Schema Registry

Cloudera Runtime 7.3.3 includes no bug fixes for Schema Registry.

Fixed issues in Apache Solr

Fixed issues for Solr are addressed in Cloudera Runtime 7.3.3.

Cloudera Runtime 7.3.3

CDPD-125501: Solr log files are empty when Cloudera Manager 7.13.3 event capture is enabled

The Solr launcher now adds the event-publish Log4j 2 library to the Jetty classpath, so the EventCatcher appender loads and Solr writes logs normally.

CDPD-125782: HBase batch indexer and MapReduceIndexerTool jobs fail with GraalVmProcessor not found

HBase batch indexer and MapReduceIndexerTool jobs with a morphline java { } block no longer fail with Provider org.apache.logging.log4j.core.config.plugins.processor.GraalVmProcessor not found when run with Cloudera Manager 7.13.3.0. Cloudera Runtime 7.3.3 ships Log4j 2.25, so the Cloudera Manager event-publish-7.13.3-shaded.jar no longer breaks morphline compilation.

Fixed issues in Spark

Cloudera Runtime 7.3.3 includes no bug fixes for Spark.

Fixed issues in Spark Atlas Connector

Cloudera Runtime 7.3.3 includes no bug fixes for Spark Atlas Connector.

Fixed issues in Apache Sqoop

Cloudera Runtime 7.3.3 includes no bug fixes for Sqoop.

Fixed issues in Streams Messaging Manager

Cloudera Runtime 7.3.3 includes no bug fixes for Streams Messaging Manager.

Fixed issues in Streams Replication Manager

Cloudera Runtime 7.3.3 includes no bug fixes for Streams Replication Manager.

Fixed issues in YARN

Cloudera Runtime 7.3.3 includes no bug fixes for YARN.

Fixed issues in ZooKeeper

Cloudera Runtime 7.3.3 includes no bug fixes for ZooKeeper.

Known issues in Cloudera Runtime 7.3.3

This topic describes known issues and workarounds in this release of Cloudera Runtime 7.3.3.

Known Issues in Apache Atlas

Known issues and technical limitations for Atlas are addressed in Cloudera Runtime 7.3.3.

Known issues identified in Cloudera Runtime 7.3.3

CDPD-123764: Active search list and cancel endpoints do not enforce authorization

7.3.3

The GET /api/atlas/admin/activeSearches and DELETE /api/atlas/admin/activeSearches/{id} REST API endpoints do not perform an authorization check. Any authenticated Atlas user can list the searches that are running for other users and cancel them.

None

Apache JIRA: ATLAS-5414

CDPD-122758: Atlas service upgrade fails after upgrading Ranger from 7.1.9 with OpenJDK 21

7.3.3

When you upgrade from Cloudera Runtime 7.1.9 to Cloudera Runtime 7.3.3 with OpenJDK 21, the Atlas service upgrade fails after the Ranger upgrade. The Atlas control script stops because it cannot find the java or jar command in the OpenJDK 21 installation directory, and reports an error such as: /usr/lib/jvm/java-21-openjdk-21.0.9.0.10-1.el9.x86_64/bin/java and/or /usr/lib/jvm/java-21-ope njdk-21.0.9.0.10-1.el9.x86_64/bin/jar not found on the system. Please make sure java and jar c ommands are available.

None

Known issues identified before Cloudera Runtime 7.3.3

Known issues identified before Cloudera Runtime 7.3.3.0 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.3.0 base release.

CDPD-102546: Rename scenarios are excluded from parallel historical entity resolution

7.3.2 or higher versions

7.3.2.10000

In the distributed notification processing framework, rename scenarios during historical entity resolution are not supported. RENAME events require complex mapping tables and multi-step fallback queries that introduce significant performance overhead. Consequently, lineage adjustments following a rename operation might exhibit discrepancies if the messages are processed out of sequence across parallel topics.

None

CDPD-94210: Lineage linkage failures in entity delete-and-recreate scenarios

7.3.2.10000 or higher versions

In the distributed notification processing framework, a critical misalignment can occur when an entity is deleted and immediately recreated (assigning a new GUID). Lineage messages for the recreated entity are often consumed before metadata for that entity is fully available or committed in the Atlas repository, which can cause lineage linkage to attach incorrectly.

None

OPSAPS-75853: Composite replication (Atlas + Iceberg) history entries recorded with "Partial success"

7.3.2 or higher versions

When running composite replication operations that include both Atlas metadata and Iceberg data replication, the replication history entries are recorded with a status of Partial success even when the replication completes successfully. This can cause confusion about the actual status of replication

operations. The generic multipart upload in Atlas is facing timeout issues while uploading a file of size 291MB.

The workaround is to upload the files to the server and try the async import process.

CDPD-63922: RowTooBigException when reading and writing data to HBase

7.3.2 or higher versions

When Atlas performs operations that read or write large amounts of data to HBase through JanusGraph, a RowTooBigException can occur if the row size exceeds the configured maximum row size limit (hbase.table.max.rowsize). This exception has been observed in multiple escalations during operations such as patch application, advanced search, and entity creation. The error indicates that the row data being processed is larger than the allowed maximum (typically 1GB or 2GB depending on configuration). While increasing the hbase.table.max.rowsize property can help in some cases, it may also cause out-of-memory errors on the HBase side. The root cause is related to how JanusGraph stores data in HBase and requires deeper investigation of Atlas data storage patterns.

As a temporary workaround, you can increase the hbase.table.max.rowsize property value. However, be aware that significantly increasing this value may lead to out-of-memory errors in HBase. Monitor HBase memory usage after making this change.

Apache JIRA: ATLAS-5426

CDPD-87879: Moving an Iceberg table from one database to another database is creating incomplete entity

7.3.2 or higher versions

When renaming an Iceberg table to move it from one database to another using the ALTER TABLE statement with RENAME TO clause (for example, ALTER TABLE test_share.iceberg_table RENAME TO default.iceberg_table), an incomplete entity is created in the target database. This occurs because the Hive Metastore hook incorrectly classifies the operation as an ADDPROPS (add properties) operation instead of a RENAME operation for Iceberg tables. This misclassification causes the HMS hook to process the event incorrectly, passing only the DO_NOT_UPDATE_STATS = true property while omitting critical properties such as old_table_name, alterTableOpType = RENAME, and old_db_name that are required for proper entity handling. As a result, Atlas creates an incomplete entity in the target database rather than properly renaming the existing entity.

None

CDPD-89476: need to enhance the repair_index.py script

7.3.2 or higher versions

When executing the repair_index.py script provided by Cloudera to reindex Atlas metadata, the script fails to run under Java 17, which is now commonly used in customer environments. Customers upgrading their OS or JVM to OpenJDK 17 encounter errors related to java.lang.reflect.InaccessibleObjectException. Even after bypassing this with additional JVM flags, the script later fails with Keystore was tampered with, or password was incorrect. This indicates potential mismatches with the Java version, keytool compatibility, or the way the script selects runtime process directories and handles credentials.

Short-term workaround:

- Add the following JVM options in repair_index.py to bypass JPMS restrictions in Java 17:

```
--add-opens=java.base/java.lang=ALL-UNNAMED \
                --add-opens=java.base/java.lang.reflect=ALL-UNNAMED
```

- Manually verify and specify the correct conf and keystore directory to avoid loading stale credentials.
- Run the script under Java 8 or Java 11 if possible.

CDPD-76686: Concurrent Ingestion: doesn't commit the kafka offset

7.3.2 or higher versions

When concurrent processing is enabled (atlas.notifications.concurrent=true), Atlas processes messages efficiently but may fail to commit the Kafka offset. This results in the current offset getting stuck for an extended period, despite the messages being processed successfully.

None

CDPD-77640: Inconsistent results from global search due to case sensitivity

7.3.2 or higher versions

Global search in Atlas produces inconsistent results based on the case of the search term. Searching with a lowercase term returns the expected entities, whereas searching with the same term beginning with an uppercase character does not return the same results. This issue is related to case-sensitivity handling in the search layer, likely in indexing or query processing.

None

CDPD-80868: Basic search API returns wrong response for attribute-based search with long values

7.3.3.0, 7.1.9 through 7.1.9 SP2 CHF1, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500, 7.3.1.600, 7.3.1.706, 7.3.1.900

7.3.2.200

Basic search using attribute filters, such as qualifiedName, returns no results when the filter values exceed the default maximum token length of 255 characters configured in Solr. This issue affects the startsWith, endsWith, and contains operators.

Use the Advanced Search (DSL) or direct entity lookup by GUID, which are not affected by this Solr token length limitation.

Apache Jira: ATLAS-5032

CDPD-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1.700 through 7.3.1.100, 7.1.9 SP1 CHF15 through 7.1.9 SP1 CHF4

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

CDPD-69150: Unable to add labels or user defined properties in Japanese

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Adding Japanese labels or user defined properties results in the error message: “Invalid label: ###, label should contain alphanumeric characters, _ or -”

None

CDPD-69279: Quick search does not return entities when using Japanese or Chinese characters to search properties

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Entities, such as hdfs_path, are not returned by the search when their indexable properties are searched using partial search terms made of Japanese or Chinese characters. Only exact matches return results when searching the indexed properties made of Japanese or Chinese characters.

None

CDPD-56085: [Impala Iceberg] LOAD DATA INPATH to Iceberg_table creates a temporary hive_table with name <iceberg_table_name>_tmp* and then marks it as DELETED in Atlas

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Running a query like LOAD DATA INPATH to iceberg_table, creates a temporary hive table with name <iceberg_table_name>_tmp* and then marks it as DELETED in Atlas. So in Atlas, a deleted entity is created corresponding to the temporary table <iceberg_table_name>_tmp*.

Tag added to the File system (HDFS) entity will not be propagated to the Iceberg table, user has to manually add to the Iceberg table, since the tag propagation is broken due to the deleted table in the flow.

CDPD-55301: The ddlQueries and ALTERTABLE_* lineage are missing for Spark tables created using spark3-shell

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

The ddlQueries and outputFromProcesses (lineage) is missing for the alter queries.

Apache JIRA: ATLAS-4747

CDPD-75994: Post DL regular upgrade (non ZDU) to 7.3.1, "Exception in getKafkaConsumer ,WakeupException: null" is seen

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

CDPD-76035: Resource lookup for Atlas service is failing

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Once the Atlas configuration snippet atlas.authentication.method.file is enabled and a classification is created, these do not synchronize correctly to the Type Category resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the Type Name.

CDPD-69140: Incremental export: When an entity has tag propagated and is exported, the tag is not propagated to it in the export

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

When creating tables with multiple levels of depth, a tag applied to the first table will be propagated along the lineage. After exporting the last table, it will not have the propagated tag.

Apache JIRA: ATLAS-4889

CDPD-77128: Log out message for Atlas needs to be more informative on Knox enabled cluster

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Logging out of Atlas does not manage the external authentication. Atlas does not communicate a log-out event with the external authentication management, Apache Knox. When you log out of Atlas, you can still open the instance of Atlas from the same web browser without re-authentication.

To prevent additional access to Atlas, close all browser windows and exit the browser.

CDPD-83398: REMOVE_LEGACY_REF_ATTRIBUTES Patch Skipped During Startup, Causing Errors While Processing DML Events

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHF

The REMOVE_LEGACY_REF_ATTRIBUTES patch is skipped during Atlas patch processing, which leads to an error when Data Manipulation Language (DML) events (e.g., "insert data") are processed, events that are typically ignored. This is observed in environments where DML event filtering is not in place or the patch is not applied correctly. As a result, Atlas attempts to process these events and encounters errors related to missing or outdated attribute handling.

CDPD-88230: Bulk import operations currently do not construct lineage as expected

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Bulk import operations do not create lineage consistently, even when the entities are imported successfully.

CDPD-88302: Bulk import operations currently do not create audits for the imported entities

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Bulk import operations do not create audits for the imported entities.

CDPD-7982: HBase bridge stops at HBase table with deleted column family

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Bridge importing metadata from HBase fails when it encounters an HBase table for which a column family was previously dropped. The error indicates:

```
Metadata service API
            org.apache.atlas.AtlasClientV2$API_V2@58112bc4 failed
 with status 404 (Not Found)
            Response Body
              ({""errorCode"":""ATLAS-404-00-007"",""errorMessa
ge"":""Invalid instance
            creation/updation parameters passed :
              hbase_column_family.table: mandatory attribute v
alue missing in type
            hbase_column_family""})
```

None

Apache JIRA: ATLAS-3551

CDPD-6675: Irregular qualifiedName format for Azure storage

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

The qualifiedName for hdfs_path entities created from Azure blog locations (ABFS) doesn't have the clusterName appended to it as do hdfs_path entities in other location types.

None

CDPD-922: IsUnique relationship attribute not honored

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

The Atlas model includes the ability to ensure that an attribute can be set to a specific value in only one relationship entity across the cluster metadata. For example, if you wanted to add metadata tags to relationships that you wanted to make sure were unique in the system, you could design the relationship attribute with the property "IsUnique" equal true. However, in this release, the IsUnique attribute is not enforced.

None

CDPD-1664: Guest users are redirected incorrectly

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Authenticated users logging in to Atlas are redirected to the Cloudera Knox-based login page. However, if a guest user (without Atlas privileges) attempts to log in to Atlas, the user is redirected instead to the Atlas login page.

To avoid this problem, open the Atlas Dashboard in a private or incognito browser window.

CDPD-1823: Queries with ? wildcard return unexpected results

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

DSL queries in Advanced Search return incorrect results when the query text includes a question mark (?) wildcard character. This problem occurs in environments where trusted proxy for Knox is enabled, which is always the case for Cloudera.

None

CDPD-1884: Free text search in Atlas is case sensitive

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

The free text search bar in the top of the screen allows you to search across entity types and through all text attributes for all entities. The search shows the top 5 results that match the search terms at any place in the text (*term* logic). It also shows suggestions that match the search terms that begin with the term (term* logic). However, in this release, the search results are case-sensitive.

If you don't see the results you expect, repeat the search changing the case of the search terms.

CDPD-1892: Ranking of top results in free-text search not intuitive

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

The Free-text search feature ranks results based on which attributes match the search criteria. The attribute ranking is evolving and therefore the choice of top results may not be intuitive in this release.

If you don't find what you need in the top 5 results, use the full results or refine the search.

CDPD-3125: Logging out of Atlas does not manage the external authentication

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

At this time, Atlas does not communicate a log-out event with the external authentication management, Apache Knox. When you log out of Atlas, you can still open the instance of Atlas from the same web browser without re-authentication.

To prevent access to Atlas after logging out, close all browser windows and exit the browser.

CDPD-3160: Hive lineage missing for INSERT OVERWRITE queries

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Lineage is not generated for Hive INSERT OVERWRITE queries on partitioned tables. Lineage is generated as expected for CTAS queries from partitioned tables.

None

CDPD-3208: Table alias values are not found in search

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

When table names are changed, Atlas keeps the old name of the table in a list of aliases. These values are not included in the search index in this release, so after a table name is changed, searching on the old table name will not return the entity for the table.

None

CDPD-4545: Searches for Qualified Names with "@" doesn't fetch the correct results

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

When searching Atlas qualifiedName values that include an "at" character (@), Atlas does not return the expected results or generate appropriate search suggestions.

Consider leaving out the portion of the search string that includes the @ sign, using the wildcard character * instead.

CDPD-4762: Spark metadata order may affect lineage

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Atlas may record unexpected lineage relationships when metadata collection from the Spark Atlas Connector occurs out of sequence from metadata collection from HMS. For example, if an ALTER TABLE operation in Spark changing a table name and is reported to Atlas before HMS has processed the change, Atlas may not show the correct lineage relationships to the altered table.

None

CDPD-10574: Suggestion order doesn't match search weights

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

At this time, the order of search suggestions does not honor the search weight for attributes.

None

Apache JIRA: ATLAS-3081

CDPD-10576: Deleted Business Metadata attributes appear in Search Suggestions

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Atlas search suggestions continue to show Business Metadata attributes even if the attributes have been deleted.

None

Apache JIRA: ATLAS-3739

CDPD-11692: Navigator table creation time not converted to Atlas

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

In converting content from Navigator to Atlas, the create time for Hive tables is not moved to Atlas.

None

Apache JIRA: ATLAS-3739

CDPD-11940: Database audit record misses table delete

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

When a hive_table entity is created, the Atlas audit list for the parent database includes an update audit. However, at this time, the database does not show an audit when the table is deleted.

None

CDPD-11941: Table creation events missed when multiple tables are created in the same Hive command

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

When multiple Hive tables are created in the same database in a single command, the Atlas audit log for the database may not capture all the table creation events. When there is a delay between creation commands, audits are created as expected.

None

CDPD-55122: Any user with ssh access can view the downloaded results

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs

When a user triggers a download from the Atlas user interface, the downloaded data can be viewed by any other user who has SSH access to the host. This occurs because the downloaded files are not sufficiently restricted by file permissions on the local filesystem, allowing unintended access to potentially sensitive data.

None

CDPD-67112: Import transforms do not work as expected when replacing a string which already has ":"

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs

The character “:” is not supported in path replacements. The import succeeds but location remains unchanged. The character “ :” must be avoided.

None

Apache JIRA: ATLAS-4834

CDPD-56590: Create table "like" from Iceberg table creates a hive_table instead of iceberg_table

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

By default, for tables created using the "like" command, lineage is not generated in Atlas. The destination like table should be of the same type as source table. Instead an iceberg_table for source and hive_table for destination are getting created.

CDPD-68191: Suggestions do not return the correct results when searching multiple Chinese characters

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Free text search does not return results when searching Chinese phrases made of multiple characters. Partial searches return the correct results.

None.

CDPD-77434: Search suggestions may fail to return results when using complete qualified names in RAZ-enabled environments

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

When performing searches in the Atlas UI or through the REST API using a complete qualified name as the prefix string (for example, searching for an Iceberg column using its full qualified name), the suggestions API sometimes fails to return the expected entity in the suggestions list, even though the entity exists and would be returned in the actual search results. This issue occurs consistently in RAZ-enabled AWS clusters but has not been observed in IDB (internal data bus) or Cloudera on prem clusters. The behavior is intermittent: the REST API /api/atlas/v2/search/suggestions?prefixString=<complete_qualifiedName> may return an empty suggestions array ({"suggestions": []}), while the same entity appears correctly when searched through the user interface. This inconsistency affects user experience and makes entity discovery less reliable in RAZ-enabled environments.

None

CDPD-81499: Basic search performance degradation when filtering on boolean attributes

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.100 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

When performing basic searches with filters on boolean attributes (for example, isDisposed = true), search operations may experience significant performance degradation, particularly with large datasets. This occurs because boolean datatype attributes are not indexed in Solr and require in-memory filtering in Atlas. When the search API attempts to return the requested number of results (for example, limit of 25), it repeatedly queries Solr for active entities and then filters them in Atlas based on boolean attribute values. This causes Atlas to traverse through potentially large numbers of entities until it finds enough matches, resulting in increased search times as data volume and query offset increase. Atlas logs may display warnings such as: not using index-search for attribute EntityType.booleanAttribute; might cause poor performance.

None

CDPD-69317: Export/Import API: When tag attributes types get modified between the exports, causes issues in import

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

When using the Export/Import API, if a tag's attribute type is modified (for example, from string to double) between incremental exports, the import operation on the target cluster fails due to the difference in attribute types.

None

CDPD-65906: Export: When relationship is updated to block tag , the imported lineage doesn't have the updated relationship

7.3.3.0, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

When using the Export/Import API to migrate data between clusters, if a relationship between entities is modified to block tag propagation (for example, setting propagation from table to process as 'None') and a tag is subsequently added, the tag correctly does not propagate on the source cluster. However, after exporting to the target cluster, the relationship on the target cluster does not reflect the blocked propagation, and the tag incorrectly propagates to the related entities.

None

Apache JIRA: ATLAS-4833

Known Issues in Apache Avro

No new known issues were identified in Cloudera Runtime for Avro.

Known Issues in Cloud Connectors

Known issues and technical limitations for Cloud Connectors are addressed in Cloudera Runtime 7.3.3.

Known issues identified in Cloudera Runtime 7.3.3

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.3

There are no known issues carried forward that affect Cloud Connectors in Cloudera Runtime 7.3.3.

Known issues in Cruise Control

Known issues and technical limitations for Cruise Control are addressed in Cloudera Runtime 7.3.3.

Known issues identified in Cloudera Runtime 7.3.3

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.3

Known issues identified before Cloudera Runtime 7.3.3 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.3 base release.

CDPD-44676: Rebalancing with Cruise Control does not work if the metric reporter fails to report the CPU usage metric

7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

If the CPU usage metric is not reported, the numValidWindows in Cruise Control will be 0 and proposal generation as well as partition rebalancing will not work. If this issue is present, the following message will be included in the Kafka logs:

```
WARN com.linkedin.kafka.cruisecontrol.metricsreporter.CruiseCont
rolMetricsReporter:
       [CruiseControlMetricsReporterRunner]: Failed reporting CPU
 util.
```

```
java.io.IOException: Java Virtual Machine recent CPU usage is not
 available.
```

This issue is only known to affect Kafka broker hosts that have the following specifications:

- CPU: Intel(R) Xeon(R) CPU E5-2699 v4 @ 2.20GHz
- OS: Linux 4.18.5-1.el7.elrepo.x86_64 #1 SMP Fri Aug 24 11:35:05 EDT 2018 x86_64
- Java version: 8-18

Move the broker to a different machine where the CPU is different. This can be done by moving the host to a different cluster. For more information, see Moving a Host Between Clusters

Note: Cluster nodes affected by this issue are not displayed as unhealthy.

Known issues in Apache HBase

Known issues and technical limitations for HBase are addressed in Cloudera Runtime 7.3.3.

Known issues identified in Cloudera Runtime 7.3.3

OPSAPS-81682: HBase clients cannot connect when RPC mTLS client authentication is required

7.3.3

When AutoTLS is enabled and the HBase RPC mTLS client auth mode property is set to Required (NEED), client TLS keystores are not generated. As a result, HBase clients, such as HBase shell, cannot connect to HBase Master, and commands fail with the Received fatal alert: bad_certificate error.

Apply either of the following configuration changes:

- Change the HBase RPC mTLS client auth mode (hbase.server.netty.tls.client.auth.mode) property from Required (NEED) to Disabled (NONE) or Optional (WANT).
- Enable the Master / Region Server TLS Plaintext Enabled (hbase.server.netty.tls.supportplaintext) property.

Known Issues in HDFS

Known issues and technical limitations for HDFS are addressed in Cloudera Runtime 7.3.3.

Known Issues Identified in Cloudera Runtime 7.3.3

There are no new known issues identified in this release.

Known Issues Identified Before Cloudera Runtime 7.3.3

CDPD-101357: Using Windows NFS Client to connect to HDFS NFS gateway is not supported

7.1.7 and its SPs and CHFs, 7.1.8 and its SPs and CHFs, 7.1.9 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.3.2 and its SPs and CHFs, and 7.3.3 and its SPs and CHFs

When attempting to copy files larger than 1.5 MB from a Windows machine to an NFS mount point, the operation fails with the Can't read from the source file or disk error. However, files smaller than 1 MB are copied successfully.

Do the following steps:

1. **Windows to Linux Folder Sharing**: Share a folder from Windows and mount it on the same Linux host, then perform the file copy operation within the Linux environment.
2. **Linux File System Sharing**: Share the Linux file system and mount it on the Windows machine, then perform the file copy operation within the Linux environment.

Known Issues in Apache Hive

Known issues and technical limitations for Hive are addressed in Cloudera Runtime 7.3.3.

Known issues identified in Cloudera Runtime 7.3.3

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.3

CDPD-73781: High operational costs from automatic partition discovery on cloud storage

7.3.1 and its SPs and CHFs, 7.2.18 and its SPs

In cloud environments, the Partition Management Task (PMT) could significantly increase operational costs and performance overhead for large tables with many partitions.

To address this issue, you can perform the following:

- Increase the value of the metastore.partition.management.task.frequency property to 86400 seconds.
- Disable partition discovery manually for specific tables incurring high costs by using the discover.partitions table property.

CDPD-93090: Failure when inserting data into partitions located on a different filesystem than the table

7.3.2

Inserting data into a table partition results in a MoveTask execution error if the partition is stored in a different filesystem than the parent table. For example, if you create an external table in S3 but add a partition with an HDFS location, the `INSERT` operation fails with a Wrong FS error. Although the data is physically inserted into the correct location, the query returns a non-zero exit code and displays an error message.

None

CDPD-92560: Partition insertion errors across multiple buckets in RAZ enabled clusters

7.3.2

Inserting data into a table partition results in a MoveTask execution error when the partition is located in a different S3 bucket than the base table. In Ranger Access Control (RAZ) enabled Private Cloud clusters, the `INSERT` operation fails with a Wrong FS error, even though the data is physically written to the specified partition location.

None

CDPD-74192: Hive SSL certificate error during large record copies

7.3.1 and its SPs and CHFs, 7.2.18.300, 7.2.18.400

When you copy a large number of records (approximately 1 million or more) using Hive, the operation might fail with a

```
javax.net.ssl.SSLException: org.bouncycastle.tls.TlsFatalAlert:
 certificate_unknown(46)
```

error. This issue occurs during the job commit phase when Hive attempts to execute HTTP requests to S3 for file operations such as copyFile.

None

CDPD-68096: Unable to set S3 credentials at the session level in Hive

7.3.2

Accessing S3 buckets from Hive on on premises results in a MetaException or AccessDeniedException when you attempt to create external tables by using session-level credentials. While Spark and other Hadoop jobs allow you to define fs.s3a.access.key and fs.s3a.secret.key at the session level, Hive Metastore (HMS) does not honor these session-specific parameters during path validation.

To access S3 buckets, you must configure the S3 credentials at the global level in the Hive or Hadoop cluster configuration (such as core-site.xml) or use a supported credential provider like the Hadoop Credential Shell to store keys securely. Setting these parameters at the session level is currently not supported for Hive table creation.

HIVE-16913

DWX-22436: DL upgrade recovery fails due to Metastore schema incompatibility

7.3.1.600

When attempting a Data Lake (DL) upgrade recovery from version 7.2.18.1100 to Cloudera Runtime 7.3.1.500, the process fails because the Hive Metastore schema versions are incompatible. The error indicates a mismatch between the Hive version (3.1.3000.7.3.1.500-182) and the database schema version (3.1.3000.7.2.18.0-Update2). This blocks Data Lake recovery if an upgrade fails, impacting customers.

Before you initiate the recovery process, manually update the Hive Metastore schema to match the target version by using the `schematool` utility.

1. Obtain the Hive database password: Run the following command to retrieve the password from the pillar configuration:

```
cat /srv/pillar/postgresql/postgre.sls
```

2. Back up the existing configuration: Move the current configuration directory to a backup location:

```
mv /etc/hive/conf /etc/hive/conf_backup
                  mkdir /etc/hive/conf
```

3. Prepare the temporary configuration: Copy the process files to the new configuration directory:

```
scp /var/run/cloudera-scm-agent/process/<process-id>-hive-me
tastore-create-tables/* /etc/hive/conf/
```

4. Update the connection password: Open the /etc/hive/conf/hive-site.xml file and perform the following modifications:
   - Set the javax.jdo.option.ConnectionPassword property to your Hive database password.
   - Comment out the hadoop.security.credential.provider.path property.
5. Run the schema upgrade tool: Execute the `schematool` to synchronize the version:

```
/opt/cloudera/parcels/CDH/lib/hive/bin/schematool -dbType po
stgres -initOrUpgradeSchema --verbose
```

6. Restore the original configuration: Remove the temporary directory and restore your backup:

```
rm -rf /etc/hive/conf
                  mv /etc/hive/conf_backup /etc/hive/conf
```

7. Restart the cluster: Restart the services to initialize the Hive Metastore with the updated schema.

CDPD-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

7.1.9 SP1 CHF4, 7.1.7 SP3 CHF7, 7.3.1.100, and its higher versions

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

CDPD-74680: DAG not retried after failure

7.3.1 and its higher versions

When executing a Hive query, if the ApplicationMaster container fails, Hive does not retry the DAG if the failure message contains some diagnostic information including a line break, leading to query failure (instead of retry).

None

Known Issues in Cloudera Data Explorer (Hue)

Known issues and technical limitations for Cloudera Data Explorer (Hue) are addressed in Cloudera Runtime 7.3.3.

Known issues identified in Cloudera Runtime 7.3.3

CDPD-123265: Files tab visibility in Data Explorer

7.3.3

In Data Explorer, the Files tab does not appear in the left assist panel when the new Storage Browser is enabled. This occurs because the Storage Browser is active by default and replaces legacy file access elements in the navigation panel.

Enable the old Storage Browser for File tab visibility:

1. In Cloudera Manager, go to Clusters Data Explorer Configuration .
2. In the Data Explorer Service Advanced Configuration Snippet (Safety Valve) field, add the following property:

```
[desktop]
enable_new_storage_browser=false
```

3. Select Save Changes.
4. Restart the Data Explorer service.

Known issues identified Before Cloudera Runtime 7.3.3

CDPD-98754: Data Explorer cannot access HDFS in Data Discovery Template

7.3.2

In the Data Discovery Template, Data Explorer cannot access HDFS. This issue occurs because Cloudera Manager assigns a standby HDFS NameNode to the webhdfs_url property instead of the active NameNode, which prevents the file browser from displaying HDFS data.

1. Log in to Cloudera Manager as an Administrator.
2. Go to Clusters HDFS to identify which NameNode is currently active.
3. Go to Clusters Data Explorer (Hue) Configuration .
4. Change the value of the webhdfs_url property from the standby NameNode to the active NameNode.
5. Click Save Changes.
6. Restart the Data Explorer service.

CDPD-95911: File operations fail for files with special characters in GS RAZ environments

7.3.2

In GS RAZ environments, you cannot rename, move, or copy files that contain special characters such as the percent sign (%). When you attempt these operations, the Failed to receive a valid re sponse from RAZ error message is displayed. Additionally, operations on files containing the hash character (#) do not result in an error, but the action is not performed.

None.

CDPD-94944: File and directory names contain special characters after a rename operation

7.3.2

In Data Explorer, file and directory names that contain special characters, such as the ampersand (&) and the apostrophe ('), are incorrectly transformed into HTML entities after a rename operation. For example, the ampersand is replaced by & and the apostrophe is replaced by '. This issue occurs during rename operations for files with names containing characters such as ~@$&()*!+'=;.txt.

None.

CDPD-95666: Endless progress bar for S3 folders without read permissions on Azure

7.3.2

In Azure environments, a persistent loading bar is displayed when you select a folder or file for which you do not have read permissions. This issue occurs across the file system, including S3 and HDFS. Currently, the user interface does not display an alert message to indicate that permissions are missing. Instead, the UI remains in a loading state.

None.

CDPD-93181: Incorrect Data Explorer server status reporting

7.3.2

Data Explorer server status is incorrectly reported as Green (Healthy) in Cloudera Manager even if the process fails to start. This occurs when the mysqlclient dependency is missing, causing database migrations to fail. The Data Explorer server appears healthy because the loglistener.py sidecar process starts first and maintains a Process ID (PID). Cloudera Manager monitors this PID and reports a successful start, even though the main Data Explorer process has terminated due to the missing dependency.

1. Stop the Hue service in Cloudera Manager.
2. Log in to the affected Hue host and manually terminate the loglistener.py sidecar process.
3. Remove any stale PID files located in the /tmp/ directory, for example, /tmp/hue_*.pid.
4. Ensure that the mysqlclient dependency is installed on the host before attempting to restart the service.

Known issues identified Before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

CDPD-95086: Data Explorer UI fails to load due to a connection timeout error

7.3.2, 7.3.1.706

In a Cloudera Data Hub high availability (HA) environment, the Data Explorer UI can fail to load after the master0 HDFS node is deleted. Data Explorer continues to attempt connections to the deleted node because Data Explorer does not support HDFS HA. As a result, the Data Explorer file browser points to an invalid HDFS endpoint, which causes connection timeouts and renders the UI inaccessible.

You can restore access to the Data Explorer UI by updating the configured HttpFS node to an active HDFS node in Cloudera Manager and then restarting the Data Explorer service.

CDPD-92946: Data Explorer debug logging cannot be disabled

7.3.2, 7.3.1.600, 7.3.1.706

Data Explorer debug logging is always enabled and cannot be turned off.

None.

CDPD-88964: Data Explorer logs missing in Data Explorer UI

7.3.2, 7.3.1.500 through 7.3.1.706

The Data Explorer UI might display only a few lines of logs instead of the complete Data Explorer logs. This can occur due to leftover Gunicorn processes that interfere with the proper logging and display of logs within the Data Explorer interface.

1. Stop the Data Explorer service.

2. Terminate any remaining Gunicorn processes to clear hung or orphan processes that might be causing the issue. Run the following command (use sudo if not running as root):

```
# pids=$(ps -efwww | grep rungunicornserver | grep -v "grep" |
 grep "rungunicornserver"| awk '{ print $2 }') && for i in $pi
ds; do kill -9 $i ; done
```

3. Restart the Data Explorer service.

CDPD-90510: Defunct Data Explorer Gunicorn worker processes accumulate

7.3.2, 7.3.1.500 through 7.3.1.706

Data Explorer on Ubuntu 22 using Oracle Database can accumulate defunct rungunicornserver worker processes due to incomplete process termination and stale database connections. This can lead to a cluttered process table, which does not critically impact service functionality.

Periodically clean defunct worker processes by running the following command:

```
pgrep -f 'hue rungunicornserver' | xargs -r kill -9
```

OPSAPS-75134: LDAP and Kerberos dual authentication fails with HiveOnTez in HTTP Transport Mode

7.3.2, 7.3.1 and its SPs and CHFs, Cloudera Manager 7.13.1 and its SPs and CHFs

Enabling LDAP for the HiveOnTez service in a Kerberos environment with transport mode set to HTTP, Data Explorer fails to load database information . This failure occurs because Data Explorer does not support the combined KERBEROS and LDAP authentication value required by HiveOnTez. The resulting conflict prevents a successful connection to HiveServer2

1. Log in to Cloudera Manager Hive-On-Tez Configuration Hive Service Advanced Configuration Snippet (Safety Valve) for hive-site.xml and set the following value:

```
hive.server2.authentication=LDAP,KERBEROS
```

2. Go to Hue Configuration Hue Server Advanced Configuration Snippet (Safety Valve) for hive-site.xml and set the following value:

```
hive.server2.authentication=KERBEROS
```

CDPD-58978: Batch query execution using Data Explorer fails with Kerberos error

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

When you run Impala queries in a batch mode, you encounter failures with a Kerberos error even if the keytab is configured correctly. This is because submitting Impala, Sqoop, Pig, or pyspark queries in a batch mode launches a shell script Oozie job from Data Explorer and this is not supported on a secure cluster.

None. Submit the queries individually.

CDPD-59677: Unable to view Phoenix tables on the left assist in Data Explorer

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

On clusters secured with Knox, you might not be able to see Phoenix tables on the left assist that are present under the default database (that is, an empty('') database).

None.

Known Issues in Apache Iceberg

Known issues and technical limitations for Iceberg are addressed in Cloudera Runtime 7.3.3.

Known Issues identified in Cloudera Runtime 7.3.3

CDPD-127140: Hive fails on COUNT(DISTINCT ROW__LINEAGE__ID) on Iceberg format version 3 tables

7.3.3

On Iceberg format version 3 tables, a plain SELECT of row lineage metadata columns such as ROW__LINEAGE__ID and LAST__UPDATED__SEQUENCE__NUMBER succeeds in Hive. When vectorized execution is enabled by default, a query such as SELECT COUNT(DISTINCT ROW__LINEAGE__ID) FROM on the table fails during the vectorized Iceberg scan while reading Parquet data. Hive returns an IndexOutOfBoundsException from the vectorized Parquet record reader. The same distinct-count check on row lineage columns in Impala and Spark completes successfully.

Disable vectorized execution for the session or query by setting hive.vectorized.execution.enabled=false or avoid COUNT(DISTINCT ROW__LINEAGE__ID).

CDPD-127281: Hive applies FOR SYSTEM_VERSION AS OF to every scan of the same Iceberg table in one query

7.3.3, 7.3.2

In Hive, when a single statement scans the same Iceberg table more than once and only one scan uses FOR SYSTEM_VERSION AS OF, Hive can apply the snapshot version to all scans of that table in the statement. The scan that should read the current table then returns the same row set as the time-travel scan. Separate queries that read only the current table or only a snapshot returns correct results. Impala and Spark return correct counts when the current table and a snapshot are compared in one query.

Compare a snapshot to the current table by running separate queries. For example, one SELECT count(*) with FOR SYSTEM_VERSION AS OF and one without instead of combining both scans of the same Iceberg table in one statement.

Apache Jira: HIVE-27190

CDPD-127289: Hive does not expose ROW__LINEAGE__ID on Iceberg format version 3 ORC or Avro tables

7.3.3

On Iceberg format version 3 tables in Hive, ROW__LINEAGE__ID is the metadata column for the Iceberg row ID. When write.format.default is a parquet file format, Hive accepts SELECT references to ROW__LINEAGE__ID. When the table uses ORC or Avro data files, Hive fails at compile time with an SemanticException [Error 10004] invalid column reference ROW__LINEAGE__ID and lists only user-defined table columns. Row lineage metadata is supported only for Parquet file format in Hive in this release.

To query ROW__LINEAGE__ID or LAST__UPDATED__SEQUENCE__NUMBER from Hive, create or migrate the Iceberg format version 3 table with write.format.default=parquet in table properties.

CDPD-127285: Impala returns duplicate _row_id values on Iceberg format version 3 ORC and Avro tables after an update

7.3.3

On Iceberg format version 3 tables, Impala exposes _row_id as row lineage metadata. After a large UPDATE on tables that use ORC or Avro data files, Impala can return duplicate _row_id values across rows. A check such as count(DISTINCT _row_id) = count(*) then returns false even though every row has a non-null _row_id. The same workflow on format version 3 Parquet tables return distinct IDs for every row. Impala reads ORC and Avro Iceberg tables but does not write those formats in the result scenario.

For Iceberg format version 3 tables that require stable, use write.format.default=parquet for unique row IDs in Impala after updates.

Known issues identified before Cloudera Runtime 7.3.3

CDPD-111552: Unsupported Iceberg column type alterations in Hive return unhandled IllegalArgument Exception

7.3.2.20000

Using Hive to change the data type of a column on an Iceberg table to an unsupported conversion, such as BOOLEAN to STRING or a narrowing change such as BIGINT to INT results in the statement failure with an unhandled IllegalArgumentException and a Java stack trace. Hive does not return a concise, user-facing error message (for example, that the conversion is not supported). Supported Iceberg type promotions, such as INT to BIGINT, complete successfully. The same behavior occurs in Beeline.

Use only Iceberg-supported column type promotions when you alter Iceberg table schemas. Unsupported conversions are rejected. If a statement fails, use the Cannot change column type message in the error output to identify the unsupported conversion.

CDPD-101742: Hive Iceberg snapshot expiration fails for ORC LZ4/LZO tables

7.3.2.10000

Hive fails to run queries on Iceberg tables stored in ORC format with LZ4 or LZO compression when LLAP cache is enabled. The failure results in an OWN_TASK_FAILURE status and vertex failures due to an UnsupportedOperationException during ORC decompression. The snapshot expiration functionality itself works as expected and is not impacted.

For Iceberg tables, disable vectorization by setting hive.vectorized.execution.enabled=false.

CDPD-89390/CDPD-83022: Incorrect row count displayed in table metadata after compaction

7.3.2.10000, 7.3.2, 7.3.1.600, 7.3.1.500, 7.3.1.400

After running data compaction operations on large tables, the row count displayed by the DESC RIBE FORMATTED command may be inaccurate. Initially, the count may appear higher than the actual number of rows. Subsequently, after running the ANALYZE command to update table statistics, the count might then appear lower than the actual number of rows.

This issue has been observed in large tables containing a significant number of historical snapshots (exceeding 10000). All these snapshots are primarily generated through UPDATE operations.

It is important to note that this is just a metadata display issue, and there is no loss of data. The underlying table data remains complete and correct.

To obtain an accurate row count, use the SELECT COUNT(*) query.

DWX-18843: Unable to read Iceberg table from Hive Virtual Warehouse

7.3.2

If you have used Apache Flink to insert data into an Iceberg table that is created from Hive, you cannot read the Iceberg table from the Hive Virtual Warehouse.

Add the engine.hive.enabled table property through the Hive beeline and set the value to "true". You can add this table property either while creating the Iceberg table or use the ALTER TABLE statement to add the table property.

DWX-18489: Hive compaction of Iceberg tables results in a failure

7.3.2

When Cloudera Data Warehouse and Cloudera Data Hub are deployed in the same environment and use the same Hive Metastore (HMS) instance, the Cloudera Data Hub compaction workers can inadvertently pick up Iceberg compaction tasks. Since Iceberg compaction is not yet supported in the latest Cloudera Data Hub version, the compaction tasks will fail when they are processed by the Cloudera Data Hub compaction workers.

In such a scenario where both Cloudera Data Warehouse and Cloudera Data Hub share the same HMS instance and there is a requirement to run both Hive ACID and Iceberg compaction jobs, it is recommended that you use the Cloudera Data Warehouse environment for these jobs. If you want to

run only Hive ACID compaction tasks, you can choose to use either the Cloudera Data Warehouse or Cloudera Data Hub environments.

If you want to run the compaction jobs without changing the environment, it is recommended that you use Cloudera Data Warehouse. To avoid interference from Cloudera Data Hub, change the value of the hive.compactor.worker.threads Hive Server (HS2) property to '0'. This ensures that the compaction jobs are not processed by Cloudera Data Hub.

1. In Cloudera Manager, click Clusters Hive Configuration to navigate to the configuration page for HMS.
2. Search for hive.compactor.worker.threads and modify the value to '0'.
3. Save the changes and restart the Hive service.

DWX-17254: Merging Iceberg branches requires a target table alias

7.3.2

Hive supports only one level of qualifier when referencing columns. In other words only one dot is accepted. For example, select table.col from ...; is allowed. select db.table.col is not allowed. Using the merge statement to merge Iceberg branches without a target or source table alias causes an exception:

```
org.apache.hadoop.hive.ql.parse.SemanticException: ... Invalid t
able alias or column reference ...
```

Use an alias, for example t, for the target table.

```
merge into mydb.target.branch_branch1 t using mydb.source.branch
_branch1 s on t.id = s.id when matched then update set value = '
matched';
```

Apache Jira: HIVE-28055

DWX-17210, DWX-13733: Timeout issue querying Iceberg tables from Hive

7.3.2

When querying Iceberg tables from Hive, the queries can faile due to a timeout issue.

1. Add the following configurations to hadoop-core-site for the Database Catalog and the Virtual Warehouse.
   - fs.s3.maxConnections=1000
   - fs.s3a.connection.maximum=1000
2. Restart the Database Catalog and Virtual Warehouse.

DWX-14163: Limitations reading Iceberg tables in Avro file format from Impala

7.3.2

The Avro, Impala, and Iceberg specifications describe some limitations related to Avro, and those limitations exist in Cloudera. In addition to these, the DECIMAL type is not supported in this release.

None.

DEX-7946: Data loss during migration of a Hive table to Iceberg

7.3.2

In this release, by default the table property 'external.table.purge' is set to true, which deletes the table data and metadata if you drop the table during migration from Hive to Iceberg.

Either one of the following workarounds prevents data loss during table migration:

- Set the table property 'external.table.purge'='FALSE'.
- Do not drop a table during migration from Hive to Iceberg.

DWX-13062: Converting a Hive table having CHAR or VARCHAR columns to Iceberg causes an exception

7.3.2

CHAR and VARCHAR data can be shorter than the length specified by the data type. Remaining characters are padded with spaces. Data is converted to a string in Iceberg. This process can yield incorrect results when you query the converted Iceberg table.

Change columns from CHAR or VARCHAR to string types before converting the Hive table to Iceberg.

Apache Jira: HIVE-26507

CDPD-92182: Inserting into Hive Iceberg tables on S3 fails with RazS3ClientCredentialsException

7.3.2, 7.3.1.706, 7.3.1.600, 7.3.1.500

In RAZ-enabled clusters where HDFS is the default file system, attempts to insert data into Hive Iceberg tables pointing explicitly to S3 locations fail with a RazS3ClientCredentialsException

None.

CDPD-89390/CDPD-83022: Incorrect row count displayed in table metadata after compaction

7.3.2, 7.3.1.600, 7.3.1.500, 7.3.1.400

After running data compaction operations on large tables, the row count displayed by the DESC RIBE FORMATTED command may be inaccurate. Initially, the count may appear higher than the actual number of rows. Subsequently, after running the ANALYZE command to update table statistics, the count might then appear lower than the actual number of rows.

This issue has been observed in large tables containing a significant number of historical snapshots (exceeding 10000). All these snapshots are primarily generated through UPDATE operations.

It is important to note that this is just a metadata display issue, and there is no loss of data. The underlying table data remains complete and correct.

To obtain an accurate row count, use the SELECT COUNT(*) query.

Cloudera Lakehouse Optimizer

CDPD-119573: Evaluation API returns an error

7.3.2.10000

The evaluation API returns a Duplicate evaluation request error when stale task metadata exists in the submitted state.

Delete the stale task metadata that is in the submitted state for the affected table from storage (HDFS or S3).

Known Issues in Iceberg REST Catalog

Known issues and technical limitations for Cloudera Iceberg REST Catalog are addressed in Cloudera Runtime 7.3.3.

Known issues identified in Cloudera Runtime 7.3.3

There are no new known issues in this release.

Known Issues in Apache Impala

Known issues and technical limitations for Impala are addressed in Cloudera Runtime 7.3.3.

Known Issues identified in Cloudera Runtime 7.3.3

IMPALA-15408: Executor groups are ignored when an executor has no executor group

7.3.3

If you use custom executor groups, ensure that every executor belongs to a custom executor group. When a cluster contains both custom executor groups and executors in the default group, queries that target a custom group run on the executors in the default group. Clusters that use only the default group are not affected.

Ensure all executors are assigned to a custom executor group. Do not mix custom executor groups with executors in the default group.

Apache Jira: IMPALA-15408

Known issues identified before Cloudera Runtime 7.3.3

Impala Kerberos authentication failures in FIPS clusters

7.3.2.20000 SP2

Impala Kerberos authentication fails through SPNEGO and Impala shell after running for several hours in Red Hat Enterprise Linux (RHEL) 9 FIPS clusters. This behavior occurs because updating krb5 OS packages to the latest version triggers a known Red Hat issue (RHEL-180783). If a user accidentally upgrades krb5 OS packages to the latest version on RHEL 9 FIPS, it triggers a known Red Hat issue (RHEL-180783) and breaks authentication.

Downgrade the krb5 OS packages to a stable version (such as 1.21.1-6) provided by Red Hat.

CDPD-90807: Thrift protocol limitation during Impala zero downtime upgrade (ZDU)

7.3.1.500 through 7.3.1.706, 7.3.2 and its higher versions

Zero Downtime Upgrades (ZDU) for Impala are not supported when upgrading to version 7.3.2. This is due to Thrift protocol incompatibilities that can cause queries to fail during upgrade.

None

CDPD-90250: Incorrect file storage location for Impala tables when using S3 as default filesystem

7.3.2

When you created an external table in Impala by using the LOCATION parameter, the files were stored in S3 instead of HDFS. This occurred even if the provided path was intended for HDFS, provided that S3 was configured as the default filesystem (fs.defaultFS) in the cluster. This inconsistency leads to AccessDeniedException errors if the Impala service does not have the necessary Ranger permissions to write to the S3 bucket.

None

IMPALA-14472: Writing arrays to Kudu tables

7.3.2

Impala does not currently support writing array data into Kudu tables.

None

Apache Jira: IMPALA-14472

DWX-20490: Impala queries fail with "Caught exception The read operation timed out, type=<class 'socket.timeout'> in ExecuteStatement"

7.3.1.500

Queries in impala-shell fail with a socket timeout error in execute statement which submits the query to the coordinator. The error occurs when query execution takes longer to start mainly when query planning is slow due to frequent metadata changes.

Increase the socket timeout on the client side. Set --client_connect_timeout_ms to a higher value, e.g. add --client_connect_timeout_ms=600000 to the impala-shell command line.

DWX-20491: Impala queries fail with EOFException: End of file reached before reading fully

7.3.1.500

Impala queries fail with an EOFException when reading from an HDFS file stored in an S3A location. The error occurs when the file is removed. If the file is removed using SQL commands like `DROP PARTITION`, there may be a significant lag in Hive Metastore event processing. If removed by non-SQL operations, run `REFRESH` or `INVALIDATE METADATA` on the table to resolve the issue.

Run REFRESH/INVALIDATE METADATA <table>;

CDPD-94720: Impala startup failure due to invalid TLS v1.3 ciphers

7.3.1 and its higher versions

When running Impala on a machine with OpenSSL 1.1.1, providing an invalid or TLS v1.2 ciphersuite in the --tls_ciphersuites startup flag causes the process to fail during startup. While OpenSSL 3.x ignores invalid ciphers, OpenSSL 1.1.1 returns an error if any ciphersuite in the list is invalid, even if other valid TLS v1.3 ciphers are present.

Ensure that the list in the --tls_ciphersuites startup flag contains only valid TLS v1.3 ciphersuites and does not contain any TLS v1.2 ciphersuites.

Apache Jira: IMPALA-14625

IMPALA-532: Impala should tolerate bad locale settings

7.3.1 and its higher versions

If the LC_* environment variables specify an unsupported locale, Impala does not start.

Add LC_ALL="C" to the environment settings for both the Impala daemon and the Statestore daemon.

IMPALA-691: Process mem limit does not account for the JVM's memory usage

7.3.1 and its higher versions

Some memory allocated by the JVM used internally by Impala is not counted against the memory limit for the impalad daemon.

To monitor overall memory usage, use the top command, or add the memory figures in the Impala web UI /memz tab to JVM memory usage shown on the **/metrics** tab.

IMPALA-635: Avro Scanner fails to parse some schemas

7.3.1 and its higher versions

The default value in Avro schema must match type of first union type, e.g. if the default value is null, then the first type in the UNION must be "null".

Swap the order of the fields in the schema specification. For example, use ["null", "string"] instead of ["string", "null"]. Note that the files written with the problematic schema must be rewritten with the new schema because Avro files have embedded schemas.

IMPALA-1024: Impala BE cannot parse Avro schema that contains a trailing semi-colon

7.3.1 and its higher versions

If an Avro table has a schema definition with a trailing semicolon, Impala encounters an error when the table is queried.

Remove trailing semicolon from the Avro schema.

IMPALA-1652: Incorrect results with basic predicate on CHAR typed column

7.3.1 and its higher versions

When comparing a CHAR column value to a string literal, the literal value is not blank-padded and so the comparison might fail when it should match.

Use the RPAD() function to blank-pad literals compared with CHAR columns to the expected length.

IMPALA-1821: Casting scenarios with invalid/inconsistent results

7.3.1 and its higher versions

Using a CAST() function to convert large literal values to smaller types, or to convert special values such as NaN or Inf, produces values not consistent with other database systems. This could lead to unexpected results from queries.

None

IMPALA-2005: A failed CTAS does not drop the table if the insert fails

7.3.1 and its higher versions

If a CREATE TABLE AS SELECT operation successfully creates the target table but an error occurs while querying the source table or copying the data, the new table is left behind rather than being dropped.

Drop the new table manually after a failed CREATE TABLE AS SELECT

IMPALA-3509: Breakpad minidumps can be very large when the thread count is high

7.3.1 and its higher versions

The size of the breakpad minidump files grows linearly with the number of threads. By default, each thread adds 8 KB to the minidump size. Minidump files could consume significant disk space when the daemons have a high number of threads.

Add -\-minidump_size_limit_hint_kb=size to set a soft upper limit on the size of each minidump file. If the minidump file would exceed that limit, Impala reduces the amount of information for each thread from 8 KB to 2 KB. (Full thread information is captured for the first 20 threads, then 2 KB per thread after that.) The minidump file can still grow larger than the "hinted" size. For example, if you have 10,000 threads, the minidump file can be more than 20 MB.

IMPALA-4978: Impala requires FQDN from hostname command on Kerberized clusters

7.3.1 and its higher versions

The method Impala uses to retrieve the host name while constructing the Kerberos principal is the gethostname() system call. This function might not always return the fully qualified domain name, depending on the network configuration. If the daemons cannot determine the FQDN, Impala does not start on a Kerberized cluster.

Test if a host is affected by checking whether the output of the `hostname` command includes the FQDN. On hosts where `hostname`, only returns the short name, pass the command-line flag ##hostname=*FULLY_QUALIFIED_DOMAIN_NAME* in the startup options of all Impala-related daemons.

IMPALA-6671: Metadata operations block read-only operations on unrelated tables

7.3.1 and its higher versions

Metadata operations that change the state of a table, like COMPUTE STATS or ALTER RECOVE R PARTITIONS, may delay metadata propagation of unrelated unloaded tables triggered by statements like DESCRIBE or SELECT queries.

None

IMPALA-7072: Impala does not support Heimdal Kerberos

None

CDPD-28139: Set spark.hadoop.hive.stats.autogather to false by default

As an Impala user, if you submit a query against a table containing data ingested using Spark and you are concerned about the quality of the query plan, you must run COMPUTE STATS against such a table in any case after an ETL operation because numRows created by Spark could be incorrect. Also, use other stats computed by COMPUTE STATS, e.g., Number of Distinct Values (NDV) and NULL count for good selectivity estimates.

For example, when a user ingests data from a file into a partition of an existing table using Spark, if spark.hadoop.hive.stats.autogather is not set to false explicitly, numRows associated with this

partition would be 0 even though there is at least one row in the file. To avoid this, the workaround is to set "spark.hadoop.hive.stats.autogather=false" in the "Spark Client Advanced Configuration Snippet (Safety Valve) for spark-conf/spark-defaults.conf" in Spark's CM Configuration section.

IMPALA-2422: % escaping does not work correctly when occurs at the end in a LIKE clause

7.3.1 and its higher versions

If the final character in the RHS argument of a LIKE operator is an escaped \% character, it does not match a % final character of the LHS argument.

None

IMPALA-2603: Crash: impala::Coordinator::ValidateCollectionSlots

A query could encounter a serious error if includes multiple nested levels of INNER JOIN clauses involving subqueries.

None

CDPD-59625: Impala shell in RHEL 9 with Python 2 as default does not work

7.1.9, 7.3.1 and its higher version

If you try to run impala-shell on RHEL 9 by setting the default python executable available in PATH to Python 2, it will fail since RHEL 9 is compatible only with Python 3.

If you run into such issues, set this parameter pointing to Python 3, IMPALA_PYTHON_EXECUTABLE=python3.

Impala cannot update table if the 'external.table.purge' property is not set to true

Impala cannot update a table using DDL statements if the 'external.table.purge' property is FALSE. ALTER TABLE statements return success with no changes to the table.

ALTER TABLE statements should be issued twice if "external.table.purge" was FALSE initially.



Note: For Iceberg tables, Impala users should always use CREATE TABLE since the Impala CREATE TABLE statement sets the flag to TRUE. However, Impala CREATE EXTERNAL TABLE sets the flag to FALSE.

Impala's known limitation when querying compacted tables

When the compaction process deletes the files for a table from the underlying HDFS location, the Impala service does not detect the changes as the compactions does not allocate new write ids. When the same table is queried from Impala it throws a 'File does not exist' exception that looks something like this:

```
Query Status: Disk I/O error on <node>:22000: Failed to open HDF
S file hdfs://nameservice1/warehouse/tablespace/managed/hive/<da
tabase>/<table>/xxxxx
Error(2): No such file or directory Root cause: RemoteException:
 File does not exist: /warehouse/tablespace/managed/hive/<data
base>/<table>/xxxx
```

Use the REFRESH/INVALIDATE statements on the affected table to overcome the 'File does not exist' exception.

Impala api calls via knox require configuration if the knox customized kerberos principal name is a default service user name

To access impala api calls via knox, if the knox customized kerberos principal name is a default service user name, then configure "authorized_proxy_user_config" by clicking Clusters->impala->configuration. Include the knox customized kerberos principal name in the comma separated list of values <knox_custom_kerberos_principal_name>=*" where <knox_custom_kerberos_principal_name> is the value of the Kerberos Principal in the Knox service. Select Clusters>Knox>Configuration and search for Kerberos Principal to display this value.

CDPD-28431: Intermittent errors could be potentially encountered when Impala UI is accessed from multiple Knox nodes.

7.1.7

You must use a single Knox node to access Impala UI.

CDPD-21828: Multiple permission assignment through grant is not working

7.1.7

None

IMPALA-11871: INSERT statement does not respect Ranger policies for HDFS

7.3.1, 7.3.1.300 and its higher version

7.3.1.100, 7.3.1.200

In a cluster with Ranger auth (and with legacy catalog mode), even if you provide RWX to cm_hdfs -> all-path for the user impala, inserting into a table whose HDFS POSIX permissions happen to exclude impala access will result in "AnalysisException: Unable to INSERT into target table (default.t1) because Impala does not have WRITE access to HDFS location: hdfs://XXXXXXXXXXX"

OPSAPS-46641: A single parameter exists in Cloudera Manager for specifying the Impala Daemon Load Balancer. Because BDR and Hue need to use different ports when connecting to the load balancer, it is not possible to configure the load balancer value so that BDR and Hue will work correctly in the same cluster.

The workaround is to use the load balancer configuration either without a port specification, or with the Beeswax port: this will configure BDR. To configure Hue use the "Hue Server Advanced Configuration Snippet (Safety Valve) for impalad_flags" to specify the the load balancer address with the HiveServer2 port.

Impala known limitation when querying compacted tables

7.3.1 and its higher versions

When the compaction process deletes the files for a table from the underlying HDFS location, the Impala service does not detect the changes as the compactions does not allocate new write ids. When the same table is queried from Impala it throws a 'File does not exist' exception that looks something like this:

```
Query Status: Disk I/O error on <node>:22000: Failed to open HDF
S file hdfs://nameservice1/warehouse/tablespace/managed/hive/<da
tabase>/<table>/xxxxx
Error(2): No such file or directory Root cause: RemoteException:
 File does not exist: /warehouse/tablespace/managed/hive/<data
base>/<table>/xxxx
```

Use the REFRESH/INVALIDATE statements on the affected table to overcome the 'File does not exist' exception.

Impala Virtual Warehouses might produce an error when querying transactional (ACID) tables

Problem: If you are querying transactional (ACID) tables with an Impala Virtual Warehouse and compaction is run on the compacting Hive Virtual Warehouse, the query might fail. The compacting process deletes files and the Impala Virtual Warehouse might not be aware of the deletion. Then when the Impala Virtual Warehouse attempts to read the deleted file, an error can occur. This situation occurs randomly.

Run the INVALIDATE METADATA statement on the transactional (ACID) table to refresh the metadata. This fixes the problem until the next compaction occurs.

IMPALA-5605: Configuration to prevent crashes caused by thread resource limits

Impala could encounter a serious error due to resource usage under very high concurrency. The error message is similar to:

```
F0629 08:20:02.956413 29088 llvm-codegen.cc:111] LLVM hit fatal
 error: Unable to allocate section memory!
terminate called after throwing an instance of 'boost::exception_
detail::clone_impl<boost::exception_detail::error_info_injector<
boost::thread_resource_error> >'
```

To prevent such errors, configure each host running an `impalad` daemon with the following settings:

```
echo 2000000 > /proc/sys/kernel/threads-max
echo 2000000 > /proc/sys/kernel/pid_max
echo 8000000 > /proc/sys/vm/max_map_count
```

Add the following lines in /etc/security/limits.conf:

```
impala soft nproc 262144
impala hard nproc 262144
```

IMPALA-9350: Ranger audit logs for applying column masking policies missing

Impala is not producing these logs.

None

IMPALA-1792: ImpalaODBC: Can not get the value in the SQLGetData(m-x th column) after the SQLBindCol(m th column)

If the ODBC SQLGetData is called on a series of columns, the function calls must follow the same order as the columns. For example, if data is fetched from column 2 then column 1, the SQLGetDa ta call for column 1 returns NULL.

Fetch columns in the same order they are defined in the table.

Known Issues in Apache Kafka

Known issues and technical limitations for Kafka are addressed in Cloudera Runtime 7.3.3.

Known issues identified in Cloudera Runtime 7.3.3

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.3

Known issues identified before Cloudera Runtime 7.3.3 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.3 base release.

CDPD-107343: Changing the KRaft listener port with a normal restart breaks the Kafka cluster

7.3.2

When the KRaft listener port is changed for all KRaft roles simultaneously and the cluster is restarted using a normal restart, the Kafka cluster stops working. When all controllers change their listeners at the same time, they only know the old listeners of the other controllers. As a result, controllers cannot send metadata update requests or elect a new active controller, and the metadata

topic is not updated. Brokers become unhealthy because they cannot reach any controller. This issue also affects enabling or disabling AutoTLS because that operation changes the listener port.

Use a rolling restart when changing the KRaft listener port. If a rolling restart is not an option, perform the change in two steps to avoid simultaneous listener changes across all controllers:

1. Add the new listener to the Kafka configuration alongside the existing listener.
2. Restart the cluster.
3. Remove the old listener from the Kafka configuration.
4. Restart the cluster.

You can add multiple listeners in the Kafka Broker Advanced Configuration Snippet in Cloudera Manager.

CDPD-105439: Kafka JVM crashes under high concurrent load when using PAM authentication on RHEL 9 or Rocky Linux 9

7.1.9, 7.2.18, 7.3.1, 7.3.2

On RHEL 9 or Rocky Linux 9 versions 9.1 through 9.7, a concurrency bug in libeconf 0.4.1 causes libpam to crash under high concurrent load. This affects Kafka and Kafka Connect when configured to use PAM authentication. When handling multiple concurrent requests, libpam could crash, causing the JVM to crash.

- Use AD/LDAP authentication instead of PAM authentication.
- Upgrade the OS to RHEL 9.8 or Rocky Linux 9.8.
- Update libeconf to libeconf-0.4.1-7.el9_8 by running dnf update libeconf.

OPSAPS-59553: Streams Messaging Manager bootstrap server config should be updated based on Kafka's listeners

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

Streams Messaging Manager does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

Streams Messaging Manager cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You need to override the bootstrap server URL by performing the following steps:

1. In Cloudera Manager, go to Streams Messaging Manager Configuration Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve)
2. Override bootstrap server URL (hostname:port as set in the listeners for broker) for streams-messaging-manager.yaml.
3. Save your changes.
4. Restart SMM.

KAFKA-2561: Performance degradation when SSL Is enabled

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

In some configuration scenarios, significant performance degradation can occur when SSL is enabled. The impact varies depending on your CPU, JVM version, Kafka configuration, and message size. Consumers are typically more affected than producers.

Configure brokers and clients with ssl.secure.random.implementation = SHA1PRNG. It often reduces this degradation drastically, but its effect is CPU and JVM dependent.

RANGER-3809: Idempotent Kafka producer fails to initialize due to an authorization failure

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

Kafka producers that have idempotence enabled require the Idempotent Write permission to be set on the cluster resource in Ranger. If permission is not given, the client fails to initialize and an error similar to the following is thrown:

```
org.apache.kafka.common.KafkaException: Cannot execute transacti
onal method because we are in an error state
              at org.apache.kafka.clients.producer.internals.Tr
ansactionManager.maybeFailWithError(TransactionManager.java:1125)
              at org.apache.kafka.clients.producer.internals.T
ransactionManager.maybeAddPartition(TransactionManager.java:442)
              at org.apache.kafka.clients.producer.KafkaProduce
r.doSend(KafkaProducer.java:1000)
              at org.apache.kafka.clients.producer.KafkaProduc
er.send(KafkaProducer.java:914)
              at org.apache.kafka.clients.producer.KafkaProducer
.send(KafkaProducer.java:800)
              .
              .
              .
              Caused by: org.apache.kafka.common.errors.Cluste
rAuthorizationException: Cluster authorization failed.
```

Idempotence is enabled by default for clients in Kafka 3.0.1, 3.1.1, and any version after 3.1.1. This means that any client updated to 3.0.1, 3.1.1, or any version after 3.1.1 is affected by this issue.

This issue has two workarounds, do either of the following:

- Explicitly disable idempotence for the producers. This can be done by setting enable.idempoten ce to false.
- Update your policies in Ranger and ensure that producers have Idempotent Write permission on the cluster resource.

CDPD-49304: AvroConverter does not support composite default values

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2.0

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

CFM-3532: The Stateless NiFi Source, Stateless NiFi Sink, and HDFS Stateless Sink connectors cannot use Snappy compression

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

This issue only affects Stateless NiFi Source and Sink connectors if the connector is running a dataflow that uses a processor that uses Hadoop libraries and is configured to use Snappy compression. The HDFS Stateless Sink connector is only affected if the Compression Codec or Compression Codec for Parquet properties are set to SNAPPY.

If you are affected by this issue, errors similar to the following will be present in the logs.

```
Failed to write to HDFS due to java.lang.UnsatisfiedLinkError: o
rg.apache.hadoop.util.NativeCodeLoader.buildSupportsSnappy()
```

```
Failed to write to HDFS due to java.lang.RuntimeException: nativ
e snappy library not available: this version of libhadoop was bu
ilt without snappy support.
```

Download and deploy missing libraries.



Important: Ensure that you complete steps 1-11 on all Kafka Connect hosts. Additionally, ensure that the advanced configuration snippet in step 12 is configured for all Kafka Connect role instances.

1. Create the /opt/nativelibs directory.

```
mkdir /opt/nativelibs
```

2. Change the owner to kafka.

```
chown kafka:kafka /opt/nativelibs
```

3. Locate the directory containing the Hadoop native libraries and copy its contents to the directory you created.

```
cp /opt/cloudera/parcels/CDH/lib/hadoop/lib/native/* /opt/na
tivelibs
```

4. Verify that libsnappy.so was copied to the directory you created.
5. Remove the following from /opt/nativelibs.

```
libhadoop.a
                  libhadoop.so
                  libhadoop.so.1.0.0
```

6. Run the following command.

```
hadoop version
```

The command returns the Hadoop version running in the cluster. Note down the first three digits in the version.

7. Go to https://archive.apache.org/dist/hadoop/common/ and download the Hadoop version that matches the first three digits of the version running in the cluster.

For example, if your Hadoop version is 3.1.1.7.1.9.0-296, then you need to download Hadoop 3.1.1.

8. Extract the downloaded archive.
9. Copy the following libraries from the downloaded archive to /opt/nativelibs on the cluster host.

```
libhadoop.a
                  libhadoop.so.1.0.0
```

The libraries are located in hadoop-*[***VERSION***]*/lib/native.

10. Create a symlink named libhadoop.so and point it to /opt/nativelibs/libhadoop.so.1.0.0.

```
ln -s /opt/nativelibs/libhadoop.so.1.0.0 /opt/nativelibs/lib
hadoop.so
```

11. Change the owner of every entry within /opt/nativelibs to kafka.

```
chown -h kafka:kafka /opt/nativelibs/*
```

12. In Cloudera Manager, go to Kafka service Configuration .
13. Add the following key-value pair to Kafka Connect Environment Advanced Configuration Snippet (Safety Valve).
    - Key: LD_LIBRARY_PATH
    - Value: /opt/nativelibs
14. Click Save Changes.
15. Restart the Kafka service.

Unsupported Features

The following Kafka features are not supported in Cloudera:

- Only Java and .Net based clients are supported. Clients developed with C, C++, Python, and other languages are currently not supported.
- The Kafka default authorizer is not supported. This includes setting ACLs and all related APIs, broker functionality, and command-line tools.
- SASL/SCRAM is only supported for delegation token based authentication. It is not supported as a standalone authentication mechanism.

Limitations

Collection of partition level metrics may cause Cloudera Manager performance to degrade

If the Kafka service operates with a large number of partitions, collection of partition level metrics may cause Cloudera Manager performance to degrade.

If you are observing performance degradation and your cluster is operating with a high number of partitions, you can choose to disable the collection of partition level metrics.



Important: If you are using Streams Messaging Manager to monitor Kafka or Cruise Control for rebalancing Kafka partitions, be aware that both Streams Messaging Manager and Cruise Control rely on partition level metrics. If partition level metric collection is disabled, Streams Messaging Manager will not be able to display information about partitions. In addition, Cruise Control will not operate properly.

Complete the following steps to turn off the collection of partition level metrics:

1. Obtain the Kafka service name.
   a. In Cloudera Manager, Select the Kafka service.
   b. Select any available chart, and select Open in Chart Builder from the configuration icon drop-down.
   c. Find $SERVICENAME= near the top of the display.

      The Kafka service name is the value of $SERVICENAME.
2. Turn off the collection of partition level metrics.
   a. Go to Hosts Hosts Configuration .
   b. Find and configure the Cloudera Manager Agent Monitoring Advanced Configuration Snippet (Safety Valve) configuration property.

      Enter the following to turn off the collection of partition level metrics:

      ```
      [KAFKA_SERVICE_NAME]_feature_send_broker_topic_partition_ent
      ity_update_enabled=false
      ```

      Replace [KAFKA_SERVICE_NAME] with the service name of Kafka obtained in step 1. The service name should always be in lower case.
   c. Click Save Changes.

Known Issues in Apache Knox

Known issues and technical limitations for Apache Knox are addressed in Cloudera Runtime 7.3.3.

Known Issues identified in Cloudera Runtime 7.3.3

CDPD-126571: JDBCTokenStateService fails to initialize on PostgreSQL when multiple Knox instances start at the same time

7.3.3

In a Knox high availability deployment that uses PostgreSQL, JDBCTokenStateService can fail to initialize when multiple Knox Gateway instances start at the same time against an empty database. Startup fails with an error similar to the following:

```
PSQLException: ERROR: duplicate key value violates unique constr
aint "pg_type_typname_nsp_index"
Detail: Key (typname, typnamespace)=(knox_tokens, 2200) already
 exists.
```

Restart the Knox service on the host where startup failed. The KNOX_TOKENS and KNOX_TOKEN_METADATA tables already exist, so the service starts normally. To avoid the issue on the first startup, start Knox on one host at a time, or create the KNOX_TOKENS and KNOX_TOKEN_METADATA tables before you start Knox.

Apache JIRA: KNOX-3536

Known Issues identified before Cloudera Runtime 7.3.3

Known issues identified before Cloudera Runtime 7.3.3 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.3 base release.

CDPD-104007: Tokens on Knox Token Management page might intermittently fail to display

7.3.3, 7.3.2.100 CHF 1 and higher

Tokens created under the Knox Token Management page might intermittently fail to display. The issue occurs randomly upon refreshing the page, where the underlying API call to retrieve user tokens is not triggered, causing the page to appear empty.

Refresh the page again until the tokens are displayed.

OPSAPS-76713: Knox crashes under high concurrent load when using PAM authentication on RHEL 9 or Rocky Linux 9

7.3.3, 7.3.2 and its SPs and CHFs

On RHEL 9 or Rocky Linux 9 versions 9.1 through 9.7, a concurrency bug in libeconf 0.4.1 causes libpam to crash under high concurrent load. This affects Knox when the cdp-proxy-api topology uses PAM authentication. When Knox handles multiple concurrent requests, libpam might crash, causing Knox service to crash and requests to the Knox cdp-proxy-api topology to fail.

Use AD/LDAP authentication instead of PAM authentication for the cdp-proxy-api topology. In Cloudera Manager, go to the Knox service, select Configuration , and configure the gateway_api_authentication_provider parameter to use AD/LDAP authentication. For configuration instructions, see Configure Apache Knox authentication for AD/LDAP.

CDPD-76294: Knox service can not be started in a large size Private Cloud Base Cloudera Manager cluster

7.3.3, 7.3.2 and its SPs and CHFs, 7.3.1.100 through 7.3.1.700

For a large size Private Cloud Base Cloudera Manager cluster installed with Cloudera Runtime 7.3.1.0, you might face the problem that Knox can not be started with the following error message:

```
Wait Until Knox Gateway Can Serve Requests failed on Knox Gateway
```

Increase the Knox configuration parameter Knox Gateway Initial/Max Heapsize from 1 GiB to 2 GiB or 4 GiB, depending on the cluster size. Then save changes and run Restart Stale Services. After these steps, the Knox service can be started.

CDPD-71751: Creation of alias from the Cloudera Manager UI fails on FIPS

7.3.3, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

Users attempting to create aliases through the Cloudera Manager UI face issues in FIPS.

The alias(es) can be created using the Knox CLI:

1. ssh to Knox host.
2. Export these directories: export KNOX_GATEWAY_DATA_DIR="/var/lib/knox/gateway/data"; export KNOX_GATEWAY_CONF_DIR="/var/lib/knox/gateway/conf"
3. Set the FIPS-specific options for the Knox CLI:

```
export KNOX_CLI_MEM_OPTS="--add-exports=java.base/sun.securi
ty.provider=bctls --add-exports=java.base/sun.security.provi
der=com.safelogic.cryptocomply.fips.core --add-modules=com.s
afelogic.cryptocomply.fips.core --add-modules=bctls --module-
path=<BCTLS_JARS_DIR> -Dcom.safelogic.cryptocomply.fips.appro
ved_only=true"
```

   <BCTLS_JARS_DIR> is the directory containing the SafeLogic bctls and fips core jar files.
4. Run the following command to create the alias: /opt/cloudera/parcels/CDH/lib/knox/bin/knoxcli.sh create-alias <ALIAS_NAME> <ALIAS_VALUE>
5. Verify the addition using /opt/cloudera/parcels/CDH/lib/knox/bin/knoxcli.sh list-alias.

For HA deployments, users must do it on every Knox host (whereas the Save Alias command applies the change to all hosts automatically).

CDPD-71305: Concurrent impala shell connection failure

7.3.3, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

If a user makes a concurrent impala-shell connection through Knox, then the connection fails.

Use only one Knox role.

CDPD-3125: Logging out of Atlas does not manage the external authentication

7.3.3, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.2.18, 7.1.9

At this time, Atlas does not communicate a log-out event with the external authentication management, Apache Knox. When you log out of Atlas, you can still open the instance of Atlas from the same web browser without re-authentication.

To prevent additional access to Atlas, close all browser windows and exit the browser.

CDPD-28431: Intermittent errors could be potentially encountered when Impala UI is accessed from multiple Knox nodes

7.3.3, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9

You must use a single Knox node to access Impala UI.

CDPD-22785: Improvements and issues need to be addressed in convert-topology knox cli command

7.3.3, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9

None.

Knox issue with JDK version

7.3.3, 7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9

jdk-1.8.0_391 is not supported.

Cloudera recommends using Cloudera supported JDKs.

CDPD-84236: Token generated by one Knox host fails with Unknown token error on another Knox host in Data Engineering High Availability clusters

7.3.3, 7.3.2 and its SPs and CHFs, 7.3.1.400 through 7.3.1.700

In Data Engineering High Availability clusters, a token generated by one Knox host may fail with an Unknown token error when accessed through another Knox host. This issue occurs due to a race condition in the PostgreSQL database, which prevents one of the Knox instances from properly initializing its configured token state service.

Restart Knox on all hosts.

Known Issues in Apache Kudu

Known issues and technical limitations for Kudu are addressed in Cloudera Runtime 7.3.3.

Known issues identified in Cloudera Runtime 7.3.3

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.3

Known issues identified before Cloudera Runtime 7.3.3 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.3 release.

You get "The user 'kudu' is not part of group 'hive' on the following hosts: " warning by the Host Inspector

If you are using fine grained authorization for Kudu, and you are also using Kudu-HMS integration with HDFS-Sentry sync, then you may get the "The user 'kudu' is not part of group 'hive' on the following hosts: " warning while upgrading.

Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

Known issues in Livy

Known issues and technical limitations for Livy are addressed in Cloudera Runtime 7.3.3.

Known issues identified in Cloudera Runtime 7.3.3

CDPD-99324: Livy gateway role is available only on the first Livy for Spark 3 service

7.3.3

When a cluster has more than one Livy for Spark 3 service, for example to run several Livy for Spark 3 service instances behind HAProxy, Cloudera Manager supports the Livy gateway role only on the first Livy for Spark 3 service. You cannot add the gateway role to the second or any later Livy for Spark 3 service.

Add the Livy gateway role only to the first Livy for Spark 3 service. Point clients to the client-facing URL of the load balancer, or to Knox, instead of to an individual Livy for Spark 3 service. For more information, see Livy load balancing with HAProxy.

CDPD-99323: Knox topologies include the Livy Server roles of every Livy for Spark 3 service

7.3.3

When a cluster has more than one Livy for Spark 3 service, Knox service discovery adds every Livy Server role of every Livy for Spark 3 service to the LIVY_FOR_SPARK3 service in the cdp-proxy, cdp-proxy-api, and cdp-proxy-token topologies. Knox then routes Livy requests to roles that you did not intend to expose through Knox, and the Livy UI can be empty or show sessions from the wrong service.

In Cloudera Manager, use the Knox Safety Valve to define the Livy Server URLs that each topology exposes. The explicit URLs override the URLs that Knox discovers. For more information, see Configuring Knox topologies for multiple Livy for Spark 3 services.

Known issues identified before Cloudera Runtime 7.3.3

Known issues identified before Cloudera Runtime 7.3.3 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.3 release.

CDPD-102241: Missing Livy for Spark 3 logo in Knox UI

7.3.3, 7.3.2.10000 SP1

The Knox Gateway UI does not display the Livy for Spark 3 service logo.

None.

Known Issues in Navigator Encrypt

Known issues and technical limitations for Navigator Encrypt are addressed in Cloudera Runtime 7.3.3.

Known Issues identified in Cloudera Runtime 7.3.3

There are no new known issues identified in this release.

Known Issues identified before Cloudera Runtime 7.3.3

Failed to start navencrypt-mount.service

7.1.9, 7.3.1, 7.3.2.0, 7.3.3.0

After installing NavEncrypt on SLES, if the comand "systemctl status navencrypt-mount" fails with the error: "Failed to start navencrypt-mount.service: Unit navencrypt-mount.service failed to load: No such file or directory", the systemd files need to be reloaded.

Run the command :

```
(sudo) systemctl daemon-reload
```

Known Issues in Oozie

Known issues and technical limitations for Oozie are addressed in Cloudera Runtime 7.3.3.

Known Issues identified in Cloudera Runtime 7.3.3

There are no new known issues identified in this release.

Known Issues identifed before Cloudera Runtime 7.3.3

CDPD-93327: Passwords may be printed in the atlas-application.properties file

7.3.3.0, 7.3.2 and its SPs and CHFs

The /etc/atlas/conf/atlas-application.properties file is a publically readable file, and it may contain the Java trust store password. This file is read by the Oozie server and is propagated to the running actions. This allows the trust store password in the /etc/atlas/conf/atlas-application.properties file to be printed in the Oozie Launcher AM's (Yarn application) console logs, and the whole atlas-application.properties file is copied into the Oozie Launcher AM's (Yarn application) local directory.

None

Oozie jobs fail (gracefully) on secure YARN clusters when JobHistory server is down

7.3.3.0, 7.3.2.0 and its SPs and CHFs

If the JobHistory server is down on a YARN (MRv2) cluster, Oozie attempts to submit a job, by default, three times. If the job fails, Oozie automatically puts the workflow in a SUSPEND state.

Workaround: When the JobHistory server is running again, use the resume command to inform Oozie to continue the workflow from the point at which it left off.

Known Issues in Ozone

Known issues and technical limitations for Ozone are addressed in Cloudera Runtime 7.3.3.

Known issues identified in Cloudera Runtime 7.3.3

CDPD-125823: ListObjectsV2 incorrectly filters CommonPrefixes when StartAfter is before matching prefixes

7.3.3

When a client calls the S3 ListObjectsV2 API with the Prefix, Delimiter, and StartAfter parameters, the Ozone S3 Gateway returns an empty CommonPrefixes list even though matching prefixes exist and sort lexicographically after the specified StartAfter value.

- Omit StartAfter when listing objects with Delimiter and Prefix.
- List objects without Delimiter and derive prefixes on the client side.
- If pagination is required, use the ContinuationToken from a prior listing, or specify a StartAfter value under the same Prefix that does not trigger the issue.

Apache JIRA: HDDS-17308

CDPD-126697: CompleteMultipartUpload accepts duplicate part numbers and creates an incomplete object

7.3.3

The Ozone S3 Gateway does not reject CompleteMultipartUpload requests when the part list contains duplicate part numbers (for example, [1, 1] with the same ETag). The operation returns a success response, but the committed object reflects the bytes of only one part (for example, 5 MiB instead of 10 MiB). Additionally, HeadObject reports a PartsCount that does not match the actual ContentLength, which causes data inconsistency.

Ensure that each part number appears exactly once and in ascending order with no gaps before calling CompleteMultipartUpload.

Apache JIRA: HDDS-17554

CDPD-126757: CompleteMultipartUpload calculates incorrect ETag and PartsCount when completing with a subset of uploaded parts

7.3.3

In Ozone S3, when a user completes a multipart upload using only a subset of uploaded parts, the object metadata does not match its actual content. While the object size and payload correctly reflect only the parts listed in the completion request, HeadObject incorrectly calculates PartsCou nt and the multipart ETag suffix using every part uploaded during the session. For example, if a session contains parts 1, 2, and 3, but completes using only parts 1 and 3, the system reports a Part sCount of 3 and an ETag suffix of -3 instead of the expected values of 2 and -2.

None

Apache JIRA: HDDS-17572

CDPD-126906: Ozone HTTPS endpoints reject TLSv1.3 by default

7.3.3

When TLS is enabled, the Ozone S3 Gateway and the web UIs of the Ozone roles accept only TLSv1.2 by default. Clients that require TLSv1.3 fail the TLS handshake with a protocol_version error. Clients that support both TLSv1.2 and TLSv1.3 are not affected.

In Cloudera Manager, add the following configuration to the Ozone Service Advanced Configuration Snippet (Safety Valve) for ozone-site.xml configuration, and restart the Ozone service:

```
<property>
  <name>ozone.ssl.enabled.protocols</name>
  <value>TLSv1.2,TLSv1.3</value>
</property>
```

Apache JIRA: HDDS-17542

CDPD-126954: S3 CreateMultipartUpload returns 500 InternalError when the key matches an existing directory in an FSO bucket

7.3.3

In a File System Optimized (FSO) bucket, calling CreateMultipartUpload with a key that matches an existing directory (for example, dir1 when dir1/ already exists) returns an HTTP 500 Internal Error. Although the Ozone Manager rejects the request because it cannot write to a directory, the S3 Gateway fails to map the error code correctly and returns a generic internal error. This behavior causes clients, such as the AWS CLI, to retry the request repeatedly.

To work around this issue, perform the following actions:

- Do not start a multipart upload on a key that matches an existing directory.
- Verify whether a key is a directory before starting an upload by running the `ozone sh key info /volume/bucket/key` command and checking whether "file": false.
- If the target path is an existing directory, use PutObject for that path, or write the multipart object under a new key that does not conflict with an existing directory.

Apache JIRA: HDDS-17658

CDPD-126993: S3 CreateMultipartUpload succeeds under a parent path that is a file, and subsequent UploadPart calls fail with 500 InternalError

7.3.3

In a File System Optimized (FSO) bucket, if a parent path is an existing file (for example, a file named parent), a CreateMultipartUpload request for a child path (for example, parent/child) incorrectly succeeds and returns an UploadId. Every subsequent UploadPart request for that upload fails with an HTTP 500 InternalError. Although the Ozone Manager detects the file system violation (NOT_A_FILE), the S3 Gateway returns a generic 500 error instead of a 400-level client error. The multipart upload remains in progress and cannot complete, leaving aborting the upload as the only option.

To work around this issue, perform the following actions:

- Do not start a multipart upload under a key where a parent component is an existing file.
- Before starting an upload, verify the path by using HeadObject or the `ozone sh key info` command.
- If a CreateMultipartUpload request returns an UploadId for such a path, run abort-multipart-uplo ad to clear the incomplete upload, and ensure that all parent components in the target path are directories.

Apache JIRA: HDDS-17644

CDPD-126995: S3 ListParts with max-parts=0 or a negative value returns HTTP 500 with no S3 error body

7.3.3

When a client calls the S3 ListParts operation with max-parts=0 or a negative value (for example, -1), the S3 Gateway returns an HTTP 500 Internal Server Error instead of a standard S3 error response. This issue occurs because the S3 Gateway forwards the invalid value to the Ozone client, which throws an unhandled IllegalArgumentException ("Max Parts Should be greater than zero").

Because the response lacks a valid S3 XML error body, clients such as the AWS CLI retry the request repeatedly, which results in multiple failed requests and generic error messages.

To work around this issue, use one of the following options:

- Specify only positive values for the max-parts parameter (typically between 1 and 1000).
- Omit the --max-parts parameter to use the default page size, and use the part-number-marker parameter for pagination if more parts exist.

Apache JIRA: HDDS-17645

CDPD-127247: UploadPartCopy does not enforce x-amz-copy-source-if-match and x-amz-copy-source-if-none-match preconditions

7.3.3

In Ozone S3, the UploadPartCopy operation fails to enforce ETag-based preconditions (x-amz-co py-source-if-match and x-amz-copy-source-if-none-match). While date-based preconditions work correctly, Ozone S3 ignores ETag conditions. Consequently, the copy operation succeeds and writes the part even when the source ETag does not match, such as after an overwrite of the source object. The completed multipart object then contains mixed data from different versions of the source object instead of returning the expected 412 PreconditionFailed error.

None

Apache JIRA: HDDS-17649

Known Issues identified before Cloudera Runtime 7.3.3

Known issues identified before Cloudera Runtime 7.3.3 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.3 base release.

CDPD-105709: Recon UI displays inconsistent Solr health status

7.3.3 and 7.3.2.10000

The Recon Overview page incorrectly displays SOLR is HEALTHY as an alert when the service functions normally. Additionally, other pages within the Recon UI report SOLR is UNHE ALTHY despite Solr remaining in a healthy state.

None

Apache JIRA: HDDS-15571

CDPD-109005: Storage Container Manager (SCM) fails to update DatanodeDetails when only port configurations change

7.3.2.10000

When a DataNode re-registers with SCM or sends heartbeats after gaining new ports, SCM fails to update the stored DatanodeDetails in the node registry. SCM triggers an update only if the IP address, hostname, or software version changes. Consequently, SCM ignores port-only updates, which leads to stale metadata and potential failures in features that rely on the new ports.

None

CDPD-109002: Ozone DataStream writes fail after enabling Ratis DataStream on existing clusters

7.3.2.10000

When Ratis DataStream is enabled on an existing cluster, open pipelines retain stale DataNode details that lack the required DataStream port information. Consequently, Ozone client writes fail when they exceed the automatic threshold because the client cannot identify the necessary port for pipeline members. This occurs because SCM does not refresh member ports for existing open pipelines after the configuration change.

None

Apache JIRA: HDDS-15799

CDPD-99722: Container data checksum reverts to a stale value after successful reconciliation

7.3.2.0 and 7.3.2.10000

In Ozone, the BackgroundContainerDataScanner service reverts the container data checksum to a stale value even after a successful reconciliation repairs corrupt chunks. This occurs because the scanner builds a checksum tree in memory while reconciliation updates the container data on disk, causing the scanner to persist an outdated tree.

The issue resolves automatically during the next container scan cycle. To resolve the issue immediately, manually trigger container reconciliation from the command line.

Apache JIRA: HDDS-14936

CDPD-105115: DataNodes fail to communicate with SCM after restart

7.3.2.10000

In a SCM High Availability (HA) cluster, a restarting or failing-over follower SCM starts its DataNode protocol server before completing its Ratis log replay. This causes SCM to process DataNode container reports against an incomplete state, which leads to dropped replica locations. Consequently, if that follower becomes the leader, affected containers appear to have missing or empty replicas, and DataNode logs record errors such as Unable to clear disk, Failed to flush DB b efore close, and EOFExceptions.

This issue is temporary and typically resolves during the next report cycle when SCM rebuilds the replica list.

CDPD-119561: Ozone Manager (OM) becomes unresponsive during sustained snapshot operations

7.3.2.10000

The OM leader enters a deadlock state during heavy snapshot creation and deletion loads. This occurs when the Key Deleting Service holds a snapshot database lock while waiting for a transaction that cannot complete because the flush thread is blocked. As a result, the OM stops responding to client requests and snapshot operations time out.

- Restart the affected OM to release the locks and restore service.
- Limit snapshot deletion and heavy write activity to reduce the risk of recurrence.
- Increase the ozone.om.unflushed.transaction.max.count configuration to lower the probability of the deadlock (monitor for increased memory usage).

Apache JIRA: HDDS-16122

CDPD-104229: Delay in deleting snapshot YAML file by OM Snapshot Local Data Manager

7.3.2.10000

A significant delay of approximately 20 to 30 minutes occurs in the deletion of snapshot YAML files by the OM Snapshot Local Data Manager after the system processes a snapshot deletion request. This delay happens because the current design requires the inDegree, which represents the number of descendant snapshots, to reach zero before the system removes the YAML file. This behavior impacts subsequent operations, such as bootstrap test cases, that rely on the timely cleanup of these files.

None

CDPD-119610: OM leader remains in LEADER_AND_NOT_READY state during Snapshot Defragmentation Bootstrap

7.3.2.10000

The Ozone Manager (OM) leader remains in the LEADER_AND_NOT_READY state after a follower node stops for a planned bootstrap. This issue occurs due to a deadlock between the snapshot installation thread and the double-buffer flush thread on the follower node. Consequently, the cluster loses its usable quorum, and all OM client RPC requests fail with an OMLeaderNotR eadyException.

Restart the affected OM follower node to break the deadlock and restore normal cluster operations.

CDPD-119172: OM follower node bootstrap fails when OM leader changes

7.3.2.10000

In an Ozone Manager (OM) High Availability (HA) setup, the bootstrap process for a follower node fails if the OM leader node changes while the bootstrap is in progress. If the original leader stops (for example, through SIGTERM), the bootstrap thread attempts to download the database snapshot from the former leader and fails with a Connection refused error. When the thread subsequently attempts to connect to the newly elected leader, it receives a 503 Service Unavailable error because the new leader is not yet ready to serve requests.

To ensure a successful bootstrap, verify that two of the three Ozone Managers (OMs) are running before bootstrapping the third node.

CDPD-102216: Toggling to Old UI from the Cluster Capacity tab results in a 404 error

7.3.2.10000

The Cluster Capacity feature is available only in the new Recon UI. If you view the Cluster Capacity page in the new UI and click Switch to Old UI, the application attempts to load the same path in the legacy interface. Because this path does not exist in the old UI, a 404 Page Not Found error appears.

Manually navigate to the main dashboard if you encounter a 404 Page Not Found error after switching to the old UI. Use the new UI to view cluster capacity information, as the legacy interface does not support this feature.

Apache JIRA: HDDS-15148

CDPD-106805: Requests to Ozone S3 Gateway through HAProxy fail with 400 Bad Request during directory creation

7.3.2.10000

When using the S3 connector to access Ozone OBS buckets through an S3 Gateway behind HAProxy, requests fail with a 400 Bad Request error. This issue occurs because zero-byte directory-marker PUT requests (which include Expect: 100-continue headers) cause the S3 Gateway to skip reading the request body and return 200 OK immediately. This leaves unconsumed bytes on the wire and desynchronizes keep-alive connections. HAProxy then misinterprets subsequent requests, causing a parsing error.

Disable keep-alive connections in the HAProxy configuration by setting option httpclose and increasing the maxconn value (for example, to 4096).

CDPD-107197: The `ozone admin datanode diskbalancer status` command fails when passing a DataNode ID instead of an address

7.3.2.10000

The `ozone admin datanode diskbalancer status` command accepts only a DataNode network address or host string. When you use a valid datanodeUuid (DataNode ID) to identify the node, the command fails with an UnknownHostException or an invalid argument error. This behavior is inconsistent with other `ozone admin` commands that support both UUIDs and addresses for node identification.

To check the disk balancer status, use the DataNode network address or hostname instead of its UUID. Run the `ozone admin datanode list` command to locate the network address.

Apache JIRA: HDDS-15690

CDPD-100716: Reading keys results in a NO_REPLICA_FOUND error after SCM follower restart and leader transfer

7.3.2.10000

When a Storage Container Manager (SCM) follower restarts, it accepts DataNode container reports before catching up with the Ratis log. Processing these reports against a stale database triggers a NotLeaderException, causing the SCM to skip updating in-memory container replica locations. If

this follower is later promoted to leader, reading keys fails with a NO_REPLICA_FOUND error due to missing replica metadata.

None

Apache JIRA: HDDS-14989

CDPD-95733: Snapshot pruning tests fail due to empty compaction SST backups in Ozone in Cloudera Runtime7.3.2

7.3.2

In Ozone in Cloudera Runtime 7.3.2, snapshot pruning tests fail because the Ozone Manager (OM) RocksDB Compaction Service configuration key requires a new name. The property key changes from ozone.compaction.service.enabled (used in version 7.1.9) to ozone.om.compaction.service.enabled in version 7.3.2. If you use the old configuration key, the compaction service remains disabled by default. This causes empty compaction SST backup directories and fails the pruning operations.

Update the configuration to use the new property key. Set ozone.om.compaction.service.enabled to true to enable the compaction service in version 7.3.2.

OPSAPS-76062: Exposing "ozone.replication" in Cloudera Manager configurations and assigning a default value

7.3.2

Exposing "ozone.replication in Cloudera Manager configurations and assigning a default value to it is causing it to override bucket replication configuration as a client side configuration even when you does not mean to set the client side configurations.

None

CDPD-73792: The new snapshot could not be found after renaming the old snapshot

7.3.2

The `ozone sh snapshot rename` command renames snapshots. It is a feature developed by the Apache Ozone community and is included in Cloudera Base on premises 7.3.1. However, it does not work properly, and Cloudera Base on premises does not support it.

None

Apache JIRA: HDDS-11384

CDPD-63350: Force deleting a FSO bucket and its contents while running `rb --force` from AWS S3 API is failing

7.3.2

Force deleting a File System Optimized (FSO) bucket and its contents while running the `rb --force` command from the AWS S3 API might fail with an error, because the S3 client sends individual delete requests for each key in the bucket. It might delete or fail to delete individual keys or directories, depending on the availability of leaf elements. It can completely delete the bucket only when all the keys or directories are cleaned up. If some keys are not deleted, the bucket will not be deleted.

Sample error message

```
# aws s3 rm s3://buck-fso --recursive
delete: s3://buck-fso/dir1/
delete: s3://buck-fso/dir1/dir2/
delete: s3://buck-fso/dir3/dir4/dir5/
delete: s3://buck-fso/dir3/dir4/
delete failed: s3://buck-fso/dir3/
# Rerun the same command again
# aws s3 rm s3://buck-fso --recursive
delete: s3://buck-fso/dir1/
delete: s3://buck-fso/dir3/
```

```
# To Confirm
# ozone sh key list s3v/buck-fso
[ ]
```

Run the `rb --force` command multiple times to completely clean up the keys and directories.

Apache JIRA: HDDS-9637

CDPD-98751: Mismatched Replicas tab in the Recon UI fails to display containers with inconsistent replica checksums

7.3.2

In the Ozone Recon UI, the Mismatched Replicas tab does not update or display containers when one or more replicas have differing checksums. Instead, they are displayed in the Under-Replicated tab.

The replica can be checked in the Under-Replicated tab, or can be cross-checked against API response.

CDPD-97512: Mismatch in Open Key count between Overview page and OM DB Insight in the Recon UI

7.3.2

In the Ozone Recon UI, the Open Key count on the **Summary** section of the Overview page might not match the count on the OM DB Insight Open Key tab. Users might see different Open Key values for the same cluster across these two views.

Use the OM DB Insight Open Key tab for the most accurate Open Key count.

CDPD-97376: Container replication counts mismatch in Recon UI

7.3.2

In the Ozone Recon UI, container replication counts, including Under-Replicated, Over-Replicated, and Mis-Replicated counts, differ between the updated and the legacy **Container** page for the same cluster.

Refer to the legacy **Container** page to view accurate replication counts.

CDPD-97311: Incorrect Creation Time and Modification Time displayed on the Namespace Usage page in the Recon UI

7.3.2

In the Ozone Recon UI, the **Namespace Usage** page displays incorrect Creation Time and Modification Time timestamps. These values do not accurately reflect the actual creation or last modification times or dates of the namespace, resulting in inaccurate metadata information.

Retrieve the correct timestamp values directly from the API response.

CDPD-97312: Mismatch between cluster State Container count and Container Summary totals

7.3.2

The Ozone Recon UI can display discrepancies between the Storage Container Manager (SCM) container count and the Recon container summary. This occurs because Recon does not synchronize all container states such as QUASI_CLOSED leading to inconsistent totals.

No workaround within the Ozone Recon UI. Use `ozone admin container report` CLI command to obtain the correct container counts for all states in the Ozone cluster.

CDPD-99248: After cdh upgrade, Ozone encounters failure while running the `Finalize Upgrade for SCM on role Storage Container` command

7.3.2

When finalizing an Ozone upgrade for the first time from Cloudera Manager, the `Finalize Upgrade for SCM on role Storage Container` command might fail with the following stderr message:

```
"Invalid response from Storage Container Manager.
```

```
Current finalization status is: FINALIZATION_IN_PROGRESS"
```

This error occurs even though finalization continues to run on the Storage Container Manager (SCM).

Ignore the failure in Cloudera Manager. Use the `ozone admin scm finalizationstatus` command to monitor progress and wait for the process to complete on SCM.

CDPD-98892: File Size Distribution bucket size range calculation is not correct

7.3.2

If large number of buckets exists within a volume, the Ozone Recon UI might display incorrect file size distribution bucket size range calculation in the File Size Distribution chart on the **Insights** page.

None

Apache JIRA: HDDS-14827

CDPD-93116: Ozone client hangs intermittently when disks are full

7.3.2

This hang is caused by continuous write retries that persist until the pipeline on the Datanode closes.

The Ozone client hangs for approximately five minutes when writing data to a Datanode if a disk full exception or other Datanode error occurs. This hang is caused by continuous write retries that persist until the pipeline on the Datanode closes.

Control the request retry behavior by setting the following configurations on the client side:

Table 3: Client-side retry request configurations

Configuration	Recommended value
hdds.ratis.raft.client.rpc.request.timeout	30s
hdds.ratis.client.multilinear.random.retry.policy	1s, 1
hdds.ratis.client.exponential.backoff.max.sleep	5s
hdds.ratis.client.exponential.backoff.base.sleep	1s
hdds.ratis.client.exponential.backoff.max.retries	2

Apache JIRA: HDDS-14040

CDPD-91562: test_validate_certs_configs configuration is failing with the maximum lifetime validation

7.3.2.0, 7.3.1.600

In daylight saving time zones, the autogenerated Ozone certificate duration might differ from the expected duration. This discrepancy is minor, because the default certificate duration is 365 days or five years, depending on the Ozone component.

CDPD-75954: The ozone debug `ldb` command and `ozone auditparser` fails with java.lang.UnsatisfiedLinkError

7.3.2, 7.3.1.400

For information on workaround, see Changing temporary path for Ozone services and CLI tools.

CDPD-54885: Ozone Prometheus does not work with TLS

7.3.2, 7.3.1 and its SPs and CHFs

The Prometheus service shipped by Ozone does not support TLS mode. So, Prometheus is not able to gather metrics from Ozone endpoints when TLS is enabled.

Go to Ozone Configuration Ozone Prometheus Endpoint Token and in the Ozone Prometheus Endpoint Token property enter any random string. This configuration generates a plaintext token

in the Ozone endpoint process directory allowing Prometheus to authenticate and collect metrics despite the TLS limitation.

CDPD-56684: Keys and buckets get deleted without volume permission

7.3.2, 7.3.1 and its SPs and CHFs

When a volume deletion is initiated, the system recursively deletes all buckets and keys within the volume before attempting to delete the volume itself. Because the ACL check for volume deletion permissions occurs only in the end, all the data within the volume is deleted even without having delete permission on the volume.



Note: This occurs even if no specific bucket or key permissions were granted to the user for recursive deletion.

CDPD-50610: Large file uploads are slow with OPEN and stream data approach

7.3.2, 7.3.1 and its SPs and CHFs

Hue file browser uses the append operation for large files. This API is not supported by Ozone in 7.1.9, therefore large file uploads can be slow or can time out in the browser.

Use native Ozone client to upload large files instead of the Hue file browser.

OPSAPS-66469: Ozone-site.xml is missing if the host does not contain HDFS roles

7.3.2, 7.3.1 and its SPs and CHFs

The client side /etc/hadoop/conf/ozone-site.xml file is not generated by Cloudera Manager if the host does not have any HDFS role. Because of this, issuing Ozone commands from that host fails because it cannot find the service name to hostname mapping. When this issue occurs, an error message is displayed: # ozone sh volume list o3://ozoneabc 23/03/06 18:46:15 WARN ha.OMProxyInfo: OzoneManager address ozoneabc:9862 for serviceID null remains unresolved for node ID null Check your ozone-site.xml file to ensure ozone manager addresses are configured properly.

Add the HDFS gateway role on that host.

CDPD-63144: hadoop.ozone.om.request.key.OMKeyRenameRequestWithFSO: Rename key failed with "failed to get parent dir" error

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs

Key rename inside the FSO bucket fails and discplays the Failed to get parent dir error. This happens when running impala workloads with ozone.

None.

CDPD-74331: Key put fails with "CompletionException: Failed to write chunk"

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs

Key put fails and displays the Failed to write chunk error when there is a volume failure during configuration.

None.

CDPD-74475: YCSB test with Hbase on Ozone degrades performance

7.3.2, 7.3.1 and its SPs and CHFs

HBase on Ozone is currently provided as a Tech Preview feature. Performance characteristics, including throughput, may not yet match those of HBase running on HDFS.

None.

CDPD-74884: Exclusive size of snapshot is always returning 0

7.3.2, 7.3.1 and its SPs and CHFs

The exclusiveSize and exclusiveReplicationSize statistics provided by the snapshot info command return a value of 0 even when the snapshot contains exclusive keys or files that do not exist in other snapshots.

None.

Apache JIRA: HDDS-11528

CDPD-75042: AWS Cli recursive delete only deletes the leaf element

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs

AWS CLI `rm` or `delete` command fails to delete all files and directories on the Ozone FSO bucket. It only deletes the leaf node.

None.

CDPD-75204: Setting file name length limit can cause NN shutdown with FSImage related error

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs

Namenode restart fails after dfs.namenode.fs-limits.max-component-length is set to a lower value and there is existing data present which exceeds the length limit.

Increase the value for the dfs.namenode.fs-limits.max-component-length parameter and restart the namenode.

CDPD-75635: Ozone write fails intermittently because SCM remains in safe mode.

7.3.2, 7.3.1 and its SPs and CHFs

Ozone write operations can fail intermittently after restarting the Storage Container Manager (SCM) leader node (or when stopping the OM leader and an SCM follower). This occurs because SCM may remain in safe mode after the restart, causing write/block allocation prechecks to fail (for example, SafeModePrecheck failed for allocateBlock).

Wait for SCM to exit safe mode automatically after it meets the required thresholds. Alternatively, manually force SCM to exit from safe mode using CLI options.

Known Issues in Apache Phoenix

There are no new known issues identified in Cloudera Runtime 7.3.3 for Phoenix.

Known issues in Apache Ranger

Known issues and technical limitations for Ranger are addressed in Cloudera Runtime 7.3.3.

Known issues identified in Cloudera Runtime 7.3.3

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.3

OPSAPS-75673: Wrong enablement of Ranger RMS Database Full Sync command

7.1.8, 7.1.9, 7.1.9 SP1, 7.2.18, 7.3.1, 7.3.2.0, 7.3.3

The `Ranger RMS Database Full Sync` command should be enabled only when all RMS server instances are stopped. This is required to ensure that the RMS database synchronizes correctly without introducing conflicts or data corruption. However, when HA (High Availability) is enabled on the cluster, the command becomes available from Cloudera Manager Ranger RMS Actions drop-down, even though only one Ranger RMS instance is stopped while the others are still running.

None.

CDPD-68806: The Revoke operation for users belonging to a group or role permission does not function as expected

7.1.9 SP1, 7.3.1, 7.3.2.0, 7.3.3

List command is listing all the tables even when the user permission is revoked. And also the command does not add any deny policy to Ranger for that specific user.

This behavior is currently not supported in HBase shell. Must be handled manually using the Ranger policy change.

CDPD-68739: The revoke command does not work when using the HBase shell

7.1.9 SP1, 7.3.1, 7.3.2.0, 7.3.3

While using the HBase shell, running the revoke command does not cancel the user permission. Users are able to perform actions even after running the revoke command.

None.

CDPD-58704: hadoop roll key/key delete command shows operation failed error when one KMS host is down, even when operation succeeds

7.1.9, 7.1.9 SP1, 7.3.1, 7.3.2.0, 7.3.3

In case of rollover/delete, client sends one more (last after delete request) request to KMS instances to clean their cache and that too to all registered kms instances. if one KMS instance is stopped (not deleted), the client gets a runtime exception.

This simply returns the runtime exception on client end for stopped instances but doesn't break any functionality.

CDPD-56803: When there is no existing policy for user and a revoke request comes from hbase, then will get this error

7.1.9, 7.1.9 SP1, 7.3.1, 7.3.2.0, 7.3.3

None.

CDPD-56741: Improvement in log message when jwtauth not used

7.1.9, 7.1.9 SP1, 7.3.1, 7.3.2.0, 7.3.3

None.

CDPD-56738: Ranger RMS showing FileNotFoundException: /usr/share/java/oraclepki.jar in Oracle 19 setup

7.1.9, 7.1.9 SP1, 7.3.1, 7.3.2.0, 7.3.3

This is a warning log printed in catalina.out file when Ranger RMS server is initialized. The following exception is observed only in Oracle 19 setup: FileNotFoundException: /usr/share/java/oraclepki.jar

None.

CDPD-48975: Ranger KMS KTS to KMS DB migration : keys with the same name but different case are not migrated

7.1.9, 7.1.9 SP1, 7.3.1, 7.3.2.0, 7.3.3

KMS keys are not case sensitive.

No workaround. Such key combinations are very rare and the migration doc was updated to check such keys before starting the migration.

CDPD-42598: Kafka policy creation allowed with incorrect permissions

7.1.8, 7.1.9, 7.1.9 SP1, 7.3.1, 7.3.2.0, 7.3.3

When creating a Kafka policy from the UI, the permissions "Idempotent write"and "Cluster action" are not displayed as they are not applicable for the "topic" resource, but when creating a policy for the "topic" resource with the permissions "Idempotent write" and "Cluster Action", the policy

is created successfully when the expected behaviour is that the policy creation must fail as the permission is not applicable for the Kafka topic resource.

None.

CDPD-41582: Atlas Resource Lookup : Classification for "entity-type" lists only classification for the following payload: {"resourceName": "classification", "userInput": "", "resources": {"classification": []}}

7.1.8, 7.1.9, 7.1.9 SP1, 7.2.16, 7.2.17, 7.2.18, 7.3.1, 7.3.2.0, 7.3.3

Expectation is to return all the classifications. But the response has only "classification". Happens similarly for entity-label, entity-business-metadata.

None.

Known issues in Ranger KMS

There are no new known issues identified in Cloudera Runtime 7.3.3 for Ranger KMS.

Known issues in Schema Registry

Known issues and technical limitations for Schema Registry are addressed in Cloudera Runtime 7.3.3.

Known issues identified in Cloudera Runtime 7.3.3

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.3

Known issues identified before Cloudera Runtime 7.3.3 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.3 base release.

CDPD-40380: Authorization checking issue when Kerberos is disabled

7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

Due to an issue in Ranger, when Kerberos is disabled then it is not possible to check authorization.

1. Open Schema Registry configuration in Cloudera Manager.
2. Find the ranger.plugin.schema-registry.service.name field.
3. Replace GENERATED_RANGER_SERVICE_NAME with the actual name of the service.
4. Restart the Schema Registry service.

CDPD-49304: AvroConverter does not support composite default values

7.1.7 and all its SP and CHF releases, 7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

OPSAPS-70971: Schema Registry does not have permissions to use Atlas after an upgrade

7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

Following an upgrade, Schema Registry might not have the required permissions in Ranger to access Atlas. As a result, Schema Registry's integration with Atlas might not function in secure clusters where Ranger authorization is enabled.

1. Access the Ranger Console (Ranger Admin web UI).
2. Click the cm_atlas resource-based service.
3. Add the schemaregistry user to the all - * policies.

4. Click Manage Service Edit Service .
5. Add the schemaregistry user to the default.policy.users property.

OPSAPS-69317: Kafka Connect Rolling Restart Check fails if SSL Client authentication is required

7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

The rolling restart action does not work in Kafka Connect when the ssl.client.auth option is set to required. The health check fails with a timeout which blocks restarting the subsequent Kafka Connect instances.

You can set ssl.client.auth to requested instead of required and initiate a rolling restart again. Alternatively, you can perform the rolling restart manually by restarting the Kafka Connect instances one-by-one and checking periodically whether the service endpoint is available before starting the next one.

Known Issues in Apache Solr

There are no new known issues identified in Cloudera Runtime 7.3.3 for Solr.

Known issues identified in Cloudera Runtime 7.3.3.0

CDPD-123261: Solr fails on GCP with No FileSystem for scheme "gs"

This issue affects Solr on Cloudera Runtime 7.3.2.10000 SP1 and 7.3.2.20000 SP2 when it accesses Google Cloud Storage. Solr cannot resolve the gs:// file system, and Solr data lake backups and Ranger audit writing fail with UnsupportedFileSystemException: No FileSystem for scheme "gs". Solr uses a small plain gcs-connector JAR without its required dependencies instead of the fat shaded JAR used by other services. Cloudera Runtime 7.3.1 is not affected.

On every host that runs the Solr role, back up the plain gcs-connector JAR under /opt/cloudera/parcels/CDH/lib/solr/server/solr-webapp/webapp/WEB-INF/lib/ and replace it with a symbolic link to the fat shaded JAR under the /opt/cloudera/parcels/CDH-7.3.2*/jars/ directory. For example:

```
SOLR_GCS_JAR=$(find /opt/cloudera/parcels/CDH/lib/solr/server/so
lr-webapp/webapp/WEB-INF/lib -name 'gcs-connector*.jar' -type f |
 head -1)
CDH_GCS_JAR=$(find /opt/cloudera/parcels/CDH-7.3.2*/jars -name 'g
cs-connector*.jar' -type f | head -1)
cp "$SOLR_GCS_JAR" "$SOLR_GCS_JAR.bak"
rm "$SOLR_GCS_JAR"
ln -s "$CDH_GCS_JAR" "$SOLR_GCS_JAR"
```

Restart the Solr service in Cloudera Manager after replacing the JAR. Repeat these steps after any new deployment or redeployment of the Solr service.

Known issues identified before Cloudera Runtime 7.3.3

CDPD-125501: Solr server log files are empty after upgrading Cloudera Manager to 7.13.3

After upgrading to Cloudera Manager 7.13.3.0 on runtimes without the CDPD-125501 launcher fix, the Solr and Infra Solr log files are 0 bytes. The cause is that Cloudera Manager adds the EventCatcher Log4j 2 appender when the Enable Log Event Capture property is on, and the runtime cannot load it.

This issue shows up on the first restart after the upgrade. Solr health checks still pass, and diagnostic bundles contain empty logs.

You have the following options to fix this issue:

- Recommended: Turn off the Enable Log Event Capture (catch_events) property for Solr and Infra Solr, and restart the service. Note that event capture is already non-functional on affected clusters.

- Alternative: Set appenders and rootLogger.appenderRefs without EventCatcher in the Solr Server Log4j 2 Advanced Configuration Snippet (Safety Valve). Note that this workaround is brittle across Cloudera Manager updates.

To permanently resolve the cause of this issue, upgrade the runtime to version 7.3.3.0 or later.

CDPD-125782: HBase batch indexer and MapReduceIndexerTool jobs fail with GraalVmProcessor not found

When you run an HBase batch indexer or MapReduceIndexerTool job on Cloudera Runtime 7.3.2 with Cloudera Manager 7.13.3.0 builds up to -31, the job can fail with Provider org.apache.logging.log4j.core.config.plugins.processor.GraalVmProcessor not found. The failure occurs when the morphline contains a java { } block.

The Cloudera Manager event-publish-7.13.3-shaded.jar contains a service entry for a processor that is not available after the jar is shaded. The Cloudera Manager agent places this jar on the MapReduce task classpath. NRT indexing and Cloudera Runtime 7.3.3 are not affected.

You have the following options to fix this issue:

- Add the Log4j 2.25 jars to the job with the -libjars option:

```
-libjars /opt/cloudera/cm/lib/log4j-core-2.25.5.jar,/opt/clo
udera/cm/lib/log4j-api-2.25.5.jar
```

- Alternatively, append /opt/cloudera/cm/lib/log4j-core-2.25.5.jar and /opt/cloudera/cm/lib/log4j-api-2.25.5.jar to the mapreduce.application.classpath property.
- For a permanent fix, upgrade Cloudera Manager to 7.13.3.0 build 32 or later, and then run Install YARN MapReduce Framework JARs.

Creating Solr collection snapshots using solrctl fails on TLS-enabled clusters

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 and its SPs and CHFs

When you run the `solrctl` command to create a collection snapshot (such as solrctl collection --create-snapshot) on a cluster with TLS and Kerberos enabled, the command fails with an IOException occurred when talking to server at: error message.

This issue occurs because the TLS truststore location and password are not exported in the shell environment, preventing the solrctl client from establishing a secure connection to the Apache Solr server over HTTPS. Using the Solr REST API directly to create or list snapshots is not affected.

To resolve this issue, export the truststore path and password in the *JAVA_TOOL_OPTIONS* environment variable before you run the `solrctl` command:

```
JAVA_TOOL_OPTIONS
=
"-Djavax.net.ssl.trustStore=[***PATH_TO_TRUSTSTORE_FILE***] -Dja
vax.net.ssl.trustStorePassword=[***TRUSTSTORE_PASSWORD***]"
```

After setting the environment variable, run the `solrctl` command:

```
solrctl --jaas
[
***PATH_TO_JAAS_CONF***
]
collection --create-snapshot
[
***SNAPSHOT_NAME***
]
-c
[
***COLLECTION_NAME***
]
```

Known Issues in Spark

Known issues and technical limitations for Spark are addressed in Cloudera Runtime 7.3.3.

Known issues identified in Cloudera Runtime 7.3.3

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.3

Known issues identified before Cloudera Runtime 7.3.3 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.3 base release.

CDPD-111272: MSCK REPAIR TABLE may perform redundant ALTER TABLE when partitions are already tracked

7.3.3, 7.3.2.100 CHF 1 and higher

When you run MSCK REPAIR TABLE on a Hive table, Spark may issue a redundant ALTER TABLE for partitions that are already tracked, which can increase repair time.

None.

CDPD-67806: Migration of ORC table with timestamp column to Iceberg with spark.sql.timestampType=TIMESTAMP_NTZ fails

7.3.3, 7.3.2 and its SPs and CHFs, 7.3.1.x

The Hive External Catalog converts NTZ timestamp types back to TimeStampType, which Hive does not support, leading to migration failures.

Set spark.sql.iceberg.use-timestamp-without-timezone-in-new-tables and spark.sql.iceberg.handle-timestamp-without-timezone to true to migrate ORC tables to Iceberg with timestamp columns without timezone.

CDPD-90447: Table is stored in HDFS when cluster is enabled with raz and S3 is set to default storage system

7.3.3, 7.3.2 and its SPs and CHFs

Tables created on a RAZ-enabled clusters with S3 set as the default storage system are stored in HDFS instead of the specified S3 bucket.

None. S3 was removed as the default storage.

CDPD-94393: RuntimeWarning Failed to add file" message appears even when Spark successfully loads files

7.3.3, 7.3.2 and its SPs and CHFs, 7.3.1.0 and higher

In both Spark 2 and 3, due to an exception when attempting to add files to the Python path, the RuntimeWarning: Failed to add file message appears even when the Python JAR file is successfully loaded.

None. You can safely ignore the message as the file is loaded successfully and the message does not affect job completion.

Spark 3: RAPIDS Accelerator is not available

7.3.3, 7.3.2 and its SPs and CHFs, 7.3.1.0 and higher

The RAPIDS Accelerator for Apache Spark is not available in Cloudera Runtime.

None.

Known Issues in Spark Atlas Connector

There are no new known issues identified in Cloudera Runtime 7.3.3 for Spark Atlas Connector.

Known Issues for Sqoop

Known issues and technical limitations for Sqoop are addressed in Cloudera Runtime 7.3.3.

Known issues identified in Cloudera Runtime 7.3.3

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.3

Known issues identified before Cloudera Runtime 7.3.3 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.3 release.

CDPD-54770: Unable to read Sqoop metastore created by an older HSQLDB version

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

If you have upgraded to Cloudera Base on premises 7.1.8 Cumulative hotfix 4 or higher versions, you may encounter issues in reading the Sqoop metastore that was created using an older version of HyperSQL Database (HSQLDB).

Cloudera upgraded the HSQLDB dependency from 1.8.0.10 to 2.7.1 and this causes incompatibility issues in Sqoop jobs that are stored in HSQLDB.

After upgrading to Cloudera Base on premises 7.1.8 Cumulative hotfix 4, you must upgrade the Sqoop metastore and convert the database files to a format that can easily be read by HSQLDB 2.7.1. For more information, see Troubleshooting Apache Sqoop issues.

CDPD-44431: Using direct mode causes problems

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

Using direct mode has several drawbacks:

- Imports can cause an intermittent and overlapping input split.
- Imports can generate duplicate data.
- Many problems, such as intermittent failures, can occur.
- Additional configuration is required.

Stop using direct mode. Do not use the --direct option in Sqoop import or export commands.

Sqoop direct mode is disabled by default. However, if you still want to use it, enable it by either setting the sqoop.enable.deprecated.direct property globally in Cloudera Manager for Sqoop or by specifying it in the command-line through -Dsqoop.enable.deprecated.direct=true.

CDPD-3089: Avro, S3, and HCat do not work together properly

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs, 7.1.7 SP1 and its CHFs

Importing an Avro file into S3 with HCat fails with Delegation Token not available.

Parquet columns inadvertently renamed

Problem: Column names that start with a number are renamed when you use the --as-parquetfile option to import data.

Prepend column names in Parquet tables with one or more letters or underscore characters.

Importing Parquet files might cause out-of-memory (OOM) errors

Problem: Importing multiple megabytes per row before initial-page-run check (ColumnWriter) can cause OOM. Also, rows that vary significantly by size so that the next-page-size check is based on small rows, and is set very high, followed by many large rows can also cause OOM.

PARQUET-99

Known issues in Streams Messaging Manager

Known issues and technical limitations for Streams Messaging Manager are addressed in Cloudera Runtime 7.3.3.

Known issues identified in Cloudera Runtime 7.3.3

CDPD-119201: Swagger UI API requests fail when Streams Messaging Manager is accessed through Knox

7.3.3

When you access the Streams Messaging Manager Swagger UI through Knox, the generated API request URLs do not include the Knox gateway path. Requests that you send from Swagger UI return an HTTP 404 response.

Use curl commands instead of Swagger UI to send requests to the API. Alternatively, import the OpenAPI specification through Knox from https://{KNOX_HOST}:8443/gateway/cdp-proxy/smm-api/swagger.json or, on newer releases, from https://{KNOX_HOST}:8443/gateway/cdp-proxy/smm-api/openapi.json. Set the base URL in your client to https://{KNOX_HOST}:8443/gateway/cdp-proxy/smm-api.

Known issues identified before Cloudera Runtime 7.3.3

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

Alert policies based on the LEADER ELECTION PER SEC attribute do not trigger

7.3.2

If the monitored Kafka cluster is running in KRaft mode, alert policies based on the **LEADER ELECTION PER SEC** alert policy attribute will not trigger because the underlying metrics are not available from Kafka. The attribute remains visible and selectable in Streams Messaging Manager even if the monitored cluster is running in KRaft mode.

None.

CDPD-99371: Prometheus metrics do not function properly with KRaft mode

7.3.2

If the Kafka cluster monitored by Streams Messaging Manager is running in KRaft mode and Streams Messaging Manager is configured to use Prometheus as its metrics store, alerts related to the active controller do not work as expected. In KRaft mode, the broker_activecontrollercount metric is reported by KRaft controllers rather than brokers. This change in metric reporting can cause aggregation issues and lead to alerts triggering indefinitely. Additionally, adding the /api/prometheus-metrics endpoint from KRaft hosts to the Prometheus configuration can cause anomalies in other alerts.

Use Cloudera Manager as the metrics store for KRaft-based Kafka clusters. Alternatively, if using Prometheus, disable all alerts based on the broker_activecontrollercount metric.

CDPD-96648: zstd compression type is missing from the compression.type drop-down menu on the Edit Topic form

7.3.2

When you edit a topic in the Streams Messaging Manager UI, zstd is not available in the compression.type drop-down menu on the Edit Topic form, even though Kafka supports zstd compression.

Change the topic compression type using the `kafka-topics` CLI tool.

OPSAPS-59597: Streams Messaging Manager UI logs are not supported by Cloudera Manager

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

Cloudera Manager does not display a Log Files menu for Streams Messaging Manager UI role (and Streams Messaging Manager UI logs cannot be displayed in the Cloudera Manager UI) because the logging type used by Streams Messaging Manager UI is not supported by Cloudera Manager.

View the Streams Messaging Manager UI logs on the host.

CDPD-39313: Some numbers are not rendered properly in Streams Messaging Manager UI

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

Very large numbers can be imprecisely represented on the UI. For example, bytes larger than 8 petabytes would lose precision.

None.

OPSAPS-59553: Streams Messaging Manager bootstrap server config should be updated based on Kafka's listeners

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

Streams Messaging Manager does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

Streams Messaging Manager cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You need to override bootstrap server URL (hostname:port as set in the listeners for broker). Add the bootstrap server details in Streams Messaging Manager safety valve in the following path:

1. In Cloudera Manager, go to SMMConfigurationStreams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml.
2. Add the following value for bootstrap servers.

```
streams.messaging.manager.kafka.bootstrap.servers=<comma-sep
arated list of brokers>
```

3. Save your changes.
4. Restart Streams Messaging Manager.

CDPD-82560: Streams Messaging Manager known unprotected endpoint

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

The Streams Messaging Manager Frontend exposes the /cm-configs and /configs public endpoints without requiring authentication. These endpoints do not share any sensitive information.

None.

Known issues in Streams Replication Manager

Known issues and technical limitations for Streams Replication Manager are addressed in Cloudera Runtime 7.3.3.

Known issues identified in Cloudera Runtime 7.3.3

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.3

Known issues identified before Cloudera Runtime 7.3.3 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.3 release.

CDPD-105439: Streams Replication Manager becomes unavailable when Kafka crashes due to PAM authentication on RHEL 9 or Rocky Linux 9

7.1.9, 7.2.18, 7.3.1, 7.3.2

On RHEL 9 or Rocky Linux 9 versions 9.1 through 9.7, a concurrency bug in libeconf 0.4.1 causes libpam to crash under high concurrent load. This affects Kafka when configured to use PAM authentication. When handling multiple concurrent requests, libpam could crash, causing the Kafka JVM to crash. Because Streams Replication Manager depends on Kafka, SRM also becomes unavailable after some time.

- Use AD/LDAP authentication instead of PAM authentication.
- Upgrade the OS to RHEL 9.8 or Rocky Linux 9.8.
- Update libeconf to libeconf-0.4.1-7.el9_8 by running dnf update libeconf.

CDPD-22089: Streams Replication Manager does not sync re-created source topics until the offsets have caught up with target topic

7.1.7 and all its SP and CHF releases, 7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

Messages written to topics that were deleted and re-created are not replicated until the source topic reaches the same offset as the target topic. For example, if at the time of deletion and re-creation there are a 100 messages on the source and target clusters, new messages will only get replicated once the re-created source topic has 100 messages. This leads to messages being lost.

None

CDPD-11079: Blacklisted topics appear in the list of replicated topics

7.1.7 and all its SP and CHF releases, 7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

If a topic was originally replicated but was later disallowed (blacklisted), it will still appear as a replicated topic under the /remote-topics REST API endpoint. As a result, if a call is made to this endpoint, the disallowed topic will be included in the response. Additionally, the disallowed topic will also be visible in the Streams Messaging Manager UI. However, it's Partitions and Consumer Groups will be 0, its Throughput, Replication Latency and Checkpoint Latency will show N/A.

None

CDPD-30275: Streams Replication Manager may automatically re-create deleted topics on target clusters

7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

If auto.create.topics.enable is enabled, deleted topics might get automatically re-created on target clusters. This is a timing issue. It only occurs if remote topics are deleted while the replication of the topic is still ongoing.

1. Remove the topic from the topic allowlist with srm-control. For example:

```
srm-control topics --source [SOURCE_CLUSTER] --target [TARGE
T_CLUSTER] --remove [TOPIC1]
```

2. Wait until Streams Replication Manager is no longer replicating the topic.
3. Delete the remote topic in the target cluster.

CDPD-80872: Streams Replication Manager replication-records-lag is incorrect on empty partitions with non-zero end offset

7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

When a replicated partition is empty and its end offset is non-zero, the replication-records-lag metric is incorrectly reported as the end offset (instead of 0 or NaN).

Wait for a new message to be written into the partition. When it gets replicated, the metric is reported correctly.

Limitations

Streams Replication Manager cannot replicate Ranger authorization policies to or from Kafka clusters

Due to a limitation in the Kafka-Ranger plugin, Streams Replication Manager cannot replicate Ranger policies to or from clusters that are configured to use Ranger for authorization. If you are using Streams Replication Manager to replicate data to or from a cluster that uses Ranger, disable authorization policy synchronization in Streams Replication Manager. This can be achieved by clearing the Sync Topic Acls Enabled (sync.topic.acls.enabled) checkbox.

Known issues in YARN, YARN Queue Manager and MapReduce

Known issues and technical limitations for YARN and YARN Queue Manager are addressed in Cloudera Runtime 7.3.3.

Known issues identified in Cloudera Runtime 7.3.3

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.3

Known issues identified before Cloudera Runtime 7.3.3 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.3 base release.

CDPD-75652: Reverse DNS lookup fails for YARN but works for HDFS

7.3.3, 7.3.2, 7.3.1 and its SPs and CHFs

When submitting a YARN application from a host without the correct DNS configuration (reverse DNS does not work for the YARN ResourceManager's host) Kerberos principal validation error, such as Server has invalid Kerberos principal. This is because the reverse DNS configuration does not work for YARN ResourceManager's host.

Add the following to YARN Service Advanced Configuration Snippet (Safety Valve) for yarn-site.xml:

```
<property>
<name>yarn.resourcemanager.principal.pattern</name>
<value>*</value>
</property>
```

COMPX-14682: Fix health check after Queue Manager restart

7.3.3, 7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

In some cases, after the QM webapp and config-service are restarted Cloudera Manager reports a healthy status after some time. However, the configuration service might not be ready and usable yet when this status is reported. Therefore, QM (webapp) is not ready and usable yet.

Wait for a minute or two after restart and then try QM.

COMPX-14820: Delete Queue and its Children throws "Queue capacity was reduced to zero, but failed to delete queue."

7.3.3, 7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

When trying to perform the operation "Delete Queue and its Children" on a queue that has one or more siblings, the operation fails as YARN has some constraints. If the queue performing the operation "Delete Queue and its Children" is a leaf node, then the operations succeeds.

None.

A fresh install of 7.3.1 or its SPs or CHFs does not allow user to bypass the Setup Database screen for YARN Queue Manager

7.3.3, 7.3.2, 7.3.1 and its SPs and CHFs

YARN Queue Manager in Cloudera Base on premises 7.3.1 does not require you to install a PostGres database, therefore users should not see the Setup Database screen and should be able to skip the Setup Database screen. With this known issue, users who are conducting a fresh install of 7.3.1 or its SPs or CHFs are not able to bypass the Setup Database screen as expected.

1. When conducting a fresh install of YARN Queue Manager in 7.3.1 or its SPs or CHFs, you must ensure that you have both Cloudera and Cloudera Manager upgraded to 7.3.1.
2. When you reach the Setup Database screen in the Cloudera Manager installation wizard for Queue Manager, enter any dummy values for the following fields:
   a. Database name: configstore
   b. Database Username: dbuser
   c. Database Password: dbpassword
   d. Database Hostname: localhost

   YARN Queue Manager will not connect to PostGres with the above details and will fall back to the embedded database.
3. Run the following script command in a browser console to enable the Continue button:

   document.querySelector('.btn.next').removeAttribute('disabled');
4. Click Continue and proceed with the YARN Queue Manager installation.
5. After installation is complete, SSH into the host that has Queue Manager installed, and run this command: sed -i 's/migrationCompleted=true/migrationCompleted=false/' /var/lib/hadoop-yarn/migration.properties

   Note: Enable Queue Manager in the YARN configurations, and restart YARN.
6. Restart YARN Queue Manager.

20202 Database migration after enabling opt-in migration

7.3.3, 7.3.2, 7.3.1 and its SPs and CHFs

When migrating from an H2 database to a PostgreSQL database in YARN Queue Manager after installation or upgrade, you might encounter an issue only when you have followed the following specific scenario:

- New install or upgrade to Cloudera 7.1.9, forcing migration from H2 to PostgreSQL database.
- Upgrade to Cloudera 7.1.9 CHF2, moving back to H2 database.
- Upgrade to Cloudera 7.1.9 SP1 with valid PostgreSQL connection details in Queue Manager configurations.

To avoid any issues during the upgrade to version Cloudera 7.1.9 SP1, ensure that PostgreSQL connection details are removed from the YARN database configuration if you prefer to continue using the H2 database.

Queue Manager does not open on using a custom user with a default Kerberos principal

7.3.3, 7.3.2, 7.3.1 and its SPs and CHFs

If a custom user is used with the default Kerberos principal, the Queue Manager web UI displays an HTTP ERROR 400 error.

Ensure that the Queue Manager process_username property matches the YARN process_username property.

Third-party applications do not launch if MapReduce framework path is not included in the client configuration

7.3.3, 7.3.2, 7.3.1 and its SPs and CHFs

MapReduce application framework is loaded from HDFS instead of being present on the NodeManagers. By default the mapreduce.application.framework.path property is set to the appropriate value, but third-party applications with their own configurations does not launch.

Set the mapreduce.application.framework.path property to the appropriate configuration for third-party applications.

JobHistory URL mismatch after server relocation

7.3.3, 7.3.2, 7.3.1 and its SPs and CHFs

After moving the JobHistory Server to a new host, the URLs listed for the JobHistory Server on the ResourceManager web UI still point to the old JobHistory Server. This affects existing jobs only. New jobs started after the move are not affected.

For any existing jobs that have the incorrect JobHistory Server URL, there is no option other than to allow the jobs to roll off the history over time. For new jobs, make sure that all clients have the updated the mapred-site.xml file that references the correct JobHistory Server.

YARN cannot start if Kerberos principal name is changed

7.3.3, 7.3.2, 7.3.1 and its SPs and CHFs

If the Kerberos principal name is changed in Cloudera Manager after launch, YARN does not start. In such cases, the keytabs can be correctly generated but YARN cannot access ZooKeeper with the new Kerberos principal name and old ACLs.

There are two possible workarounds:

- Delete the znode and restart the YARN service.
- Use the reset ZK ACLs command. This also sets the znodes below /rmstore/ZKRMStateRoot to world:anyone:cdrwa which is less secure.

YARN Resource Manager UI does not display logs

7.3.3, 7.3.2, 7.3.1 and its SPs and CHFs

When the parameter mapreduce.cluster.acl.enabled is set to true, the Yarn RM UI does display and logs and the Logs are not available message is displayed.

Set mapreduce.cluster.acl.enabled to false.

Known Issues in Apache ZooKeeper

Known issues and technical limitations for ZooKeeper are addressed in Cloudera Runtime 7.3.3.

Known issues identified in Cloudera Runtime 7.3.3

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.3

CDPD-121220: OMID TSO fails to start when ZooKeeper plaintext client port 2181 is disabled

7.3.2.10000, 7.3.2.20000, 7.3.3

When ZooKeeper plaintext client port 2181 is disabled by setting zookeeper_enable_client_port to false in secure/TLS-enabled environments, the OMID TSO (Transaction Status Oracle) role fails to start if Omid ZooKeeper Secure Client Enabled omid_enable_secure_zookeeper is set to false. OMID TSO attempts to connect to ZooKeeper using port 2181, resulting in a connection

refused error Can't contact ZK cluster and prevents the service from initializing, thereby causing deployment and cluster startup health checks to time out.

Configure OMID to use the ZooKeeper secure port:

1. In Cloudera Manager, navigate to the OMID service.
2. Go to the Configuration tab.
3. Search for Omid ZooKeeper Secure Client Enabled (omid_enable_secure_zookeeper).
4. Select or set this property to true to allow OMID to connect using ZooKeeper port 2182.
5. Save changes and restart the OMID service, or restart stale services in the cluster.

Behavioral changes in Cloudera Runtime 7.3.3

Behavioral changes denote a marked change in behavior from the previously released version to this version of Cloudera Runtime 7.3.3.

Behavioral Changes in Atlas

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for Atlas.

Behavioral Changes in Avro

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for Avro.

Behavioral Changes in Cloud Connectors

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for Cloud Connectors.

Cloudera Runtime 7.3.3

There are no behavioral changes in this release.

Behavioral Changes in Cruise Control

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for Cruise Control.

Behavioral changes in HBase

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for HBase.

Cloudera Runtime 7.3.3

There are no behavioral changes in this release.

Behavioral Changes in HDFS

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for HDFS.

Cloudera Runtime 7.3.3

There are no behavioral changes in this release.

Behavioral Changes in Hive

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for Hive.

Cloudera Runtime 7.3.3

There are no behavioral changes in this release.

Behavioral Changes in Cloudera Data Explorer (Hue)

Functional adjustments and behavioral updates for Cloudera Data Explorer (Hue) are introduced in Cloudera Runtime 7.3.3.

Summary: Removal of automatic VACUUM from manual cleanup API

Previous behavior:

Previously, the Query Processor manual cleanup API (/cleanup/{startTime}) automatically ran VACUUM after deleting records. This behavior caused compatibility issues with MySQL, unexpected resource consumption, and database locks during disk space recovery.

New behavior:

The manual cleanup API no longer runs VACUUM automatically after deleting records. Database administrators must run VACUUM manually when necessary.

Behavioral Changes in Impala

Functional adjustments and behavioral updates for Impala are introduced in Cloudera Runtime 7.3.3.

Cloudera Runtime 7.3.3

Summary:

AVG_COLLECTION_SIZE query option now configurable

Previous behavior:

The AVG_COLLECTION_SIZE setting was hard-coded to a fixed value and could not be changed during query planning. This prevented optimization of queries that process collection data with different average sizes.

New behavior:

The AVG_COLLECTION_SIZE is now a configurable query option in Cloudera Runtime 7.3.3. You can set this option per query using SET AVG_COLLECTION_SIZE=<value>; to optimize query plans for collections with different average sizes. This allows Impala to generate more accurate cardinality estimates for collection operations.

Apache Jira: IMPALA-15385

Summary:

JWT/OAuth authentication configuration unified under --oauth_servers framework

Previous behavior:

Impala used separate legacy configuration flags for JWT and OAuth token authentication. Each configuration flag controlled a single identity provider setup, requiring multiple independent flags such as --jwt_token_auth, --jwks_url, and their OAuth equivalents (--oauth_token_auth, --oauth_jwks_url, and others). This approach supported only a single provider per cluster configuration and provided no built-in JWT claim validation.

New behavior:

Impala 7.3.3 introduces a unified --oauth_servers JSON configuration framework that replaces the legacy authentication flags. The new framework supports multiple identity providers in a single cluster configuration and enables JWT claim validation (issuer and audience claims) on a per-provider basis using issuerClaims and audienceClaims arrays within each OAuth provider entry. Legacy authentication flags continue to function with logged deprecation warnings, but new deployments and migrations should use --oauth_servers for enhanced security and flexibility. For federated authentication scenarios with multiple trusted IDPs, the unified framework eliminates the need for per-IDP cluster restarts and provides centralized configuration management.

Apache Jira: IMPALA-12232

Behavioral Changes in Apache Iceberg

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for Iceberg.

Behavioral Changes in Iceberg REST Catalog

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for Cloudera Iceberg REST Catalog.

Behavioral Changes in Kafka

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for Kafka.

Cloudera Runtime 7.3.3

There are no behavioral changes in this release.

Behavioral Changes in Apache Knox

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for Apache Knox.

Behavioral Changes in Kudu

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for Kudu.

Cloudera Runtime 7.3.3

There are no behavioral changes in this release.

Behavioral Changes in Livy

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for Livy.

Cloudera Runtime 7.3.3

There are no behavioral changes in this release.

Behavioral Changes in Oozie

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for Oozie.

Cloudera Runtime 7.3.3

There are no behavioral changes in this release.

Behavioral Changes in Ozone

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for Ozone.

Cloudera Runtime 7.3.3

There are no behavioral changes in this release.

Behavioral Changes in Parquet

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for Parquet.

Behavioral Changes in Phoenix

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for Phoenix.

Cloudera Runtime 7.3.3

There are no behavioral changes in this release.

Behavioral Changes in Apache Ranger

Functional adjustments and behavioral updates for Apache Ranger are introduced in Cloudera Runtime 7.3.3.

Cloudera Runtime 7.3.3

Summary:

Disabled Unix Authentication Service by default. The Unix Authentication Service in Ranger Usersync listens on port 5151 with a generic self-signed TLS certificate. Most customers use PAM or LDAP for Ranger authentication, so this change makes the service opt-in instead of always on.

Previous behavior:

The Unix Authentication Service always started with Usersync and listened on port 5151 with a self-signed TLS certificate. It could not be disabled.

New behavior:

The Unix Authentication Service is disabled by default, so port 5151 is not opened. To enable it, set ranger.usersync.unix.auth.enabled=true. The sync source must also be Unix with the backend set to password, and Unix authentication must be enabled in Ranger Admin.

Behavioral Changes in Schema Registry

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for Schema Registry.

Cloudera Runtime 7.3.3

There are no behavioral changes in this release.

Behavioral Changes in Solr

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for Solr.

Behavioral Changes in Spark

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for Spark.

Cloudera Runtime 7.3.3

There are no behavioral changes in this release.

Behavioral Changes in Spark Atlas Connector

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for Spark Atlas Connector.

Cloudera Runtime 7.3.3

There are no behavioral changes in this release.

Behavioral Changes in Sqoop

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for Sqoop.

Cloudera Runtime 7.3.3

There are no behavioral changes in this release.

Behavioral Changes in Streams Messaging Manager

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for Streams Messaging Manager.

Cloudera Runtime 7.3.3

There are no behavioral changes in this release.

Behavioral Changes in Streams Replication Manager

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for Streams Replication Manager.

Cloudera Runtime 7.3.3

There are no behavioral changes in this release.

Behavioral Changes in YARN

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for YARN.

Behavioral changes in ZooKeeper

Existing component behavior remains unchanged in Cloudera Runtime 7.3.3 for ZooKeeper.

Release Matrix

Cloudera Runtime 7.3.3 release matrix lists the runtime release build number, GBN, release dates, and download URLs.

Runtime Release Build Number	GBN	Release Date	Download URL
7.3.3.0-56	84928234	November 13, 2026	https://archive.cloudera.com/p/cdh7/7.3.3/

Components versions in Cloudera Runtime 7.3.3

List of the official component versions for Cloudera Runtime 7.3.3. To know the component versions for compatibility with other applications, you must be familiar with the latest component versions in Cloudera Runtime. You must also be aware of the available Technical Preview components and use them only in a testing environment.

Table 4: Apache components

Component	Version
Apache Arrow	15.0.0.7.3.3.0-56
Apache Atlas	2.4.0.7.3.3.0-56
Apache Avatica	1.26.0.7.3.3.0-56
Apache Avro	1.11.3.7.3.3.0-56
Apache Calcite	1.33.0.7.3.3.0-56
Apache Hadoop (Includes MapReuce, YARN and HDFS)	3.4.2.7.3.3.0-56
Apache HBase	2.6.3.7.3.3.0-56
Apache Hive	3.1.3000.7.3.3.0-56
Apache Iceberg	1.10.1.7.3.3.0-56
Apache Impala	4.5.0.7.3.3.0-56
Apache Kafka	3.9.2.7.3.3.0-56
Apache Knox	2.1.0.7.3.3.0-56
Apache Kudu	1.18.0.7.3.3.0-56
Apache Livy	0.7.23000.7.3.3.0-56
Apache Ozone	2.2.0.7.3.3.0-56
Apache Oozie	5.1.0.7.3.3.0-56
Apache ORC	2.0.3.7.3.3.0-56
Apache Parquet	1.16.0.7.3.3.0-56
Apache Phoenix	5.2.1.7.3.3.0-56
Apache Ranger	2.9.0.7.3.3.0-56
Apache Solr	8.11.2.7.3.3.0-56
Apache Spark	3.5.8.7.3.3.0-56
Apache Sqoop	1.4.7.7.3.3.0-56
Apache Tez	0.9.1.7.3.3.0-56

Component	Version
Apache ZooKeeper	3.8.5.7.3.3.0-56

Table 5: Other components

Component	Version
Cruise Control	2.5.143.7.3.3.0-56
GCS Connector	3.1.3.7.3.3.0-56
Cloudera Data Explorer (Hue)	4.11.0.7.3.3.0-56
Search	1.0.0.7.3.3.0-56
Schema Registry	0.10.0.7.3.3.0-56
Streams Messaging Manager	2.3.0.7.3.3.0-56
Streams Replication Manager	1.1.0.7.3.3.0-56

Table 6: Connectors and encryption components

Component	Version
HBase connectors	1.0.0.7.3.3.0-56
Hive Warehouse Connector	1.0.0.7.3.3.0-56
Spark Atlas Connector	3.5.8.7.3.3.0-56
Spark Schema Registry	3.5.8.7.3.3.0-56

Fixed Common Vulnerabilities and Exposures in Cloudera Runtime 7.3.3

Common vulnerabilities and Exposures (CVEs) fixed in Cloudera Runtime 7.3.3

CVE ID	Description
CVE-2015-9251	jQuery before 3.0.0 is vulnerable to Cross-site Scripting (XSS) attacks when a cross-domain Ajax request is performed without the dataType option, causing text/javascript responses to be executed.
CVE-2016-4947	Cloudera HUE 3.9.0 and earlier allows remote attackers to enumerate user accounts via a request to desktop/api/users/autocomplete.
CVE-2016-10735	In Bootstrap 3.x before 3.4.0 and 4.x-beta before 4.0.0-beta.2, XSS is possible in the data-target attribute, a different vulnerability than CVE-2018-14041.
CVE-2016-10744	In Select2 through 4.0.5, as used in Snipe-IT and other products, rich selectlists allow XSS. This affects use cases with Ajax remote data loading when HTML templates are used to display listbox data.
CVE-2018-8768	In Jupyter Notebook before 5.4.1, a maliciously forged notebook file can bypass sanitization to execute JavaScript in the notebook context. Specifically, invalid HTML is 'fixed' by jQuery after sanitization, making it dangerous.
CVE-2018-14040	In Bootstrap before 4.1.2, XSS is possible in the collapse data-parent attribute.
CVE-2018-14042	In Bootstrap before 4.1.2, XSS is possible in the data-container property of tooltip.
CVE-2018-19351	Jupyter Notebook before 5.7.1 allows XSS via an untrusted notebook because nbconvert responses are considered to have the same origin as the notebook server. In other words, nbconvert endpoints can execute JavaScript with access to the server API. In notebook/nbconvert/handlers.py, NbconvertFileHandler and NbconvertPostHandler do not set a Content Security Policy to prevent this.
CVE-2018-19352	Jupyter Notebook before 5.7.2 allows XSS via a crafted directory name because notebook/static/tree/js/notebooklist.js handles certain URLs unsafely.

CVE ID	Description
CVE-2018-20676	In Bootstrap before 3.4.0, XSS is possible in the tooltip data-viewport attribute.
CVE-2018-20677	In Bootstrap before 3.4.0, XSS is possible in the affix configuration target property.
CVE-2018-21030	Jupyter Notebook before 5.5.0 does not use a CSP header to treat served files as belonging to a separate origin. Thus, for example, an XSS payload can be placed in an SVG document.
CVE-2019-8331	In Bootstrap before 3.4.1 and 4.3.x before 4.3.1, XSS is possible in the tooltip or popover data-template attribute.
CVE-2019-9644	An XSSI (cross-site inclusion) vulnerability in Jupyter Notebook before 5.7.6 allows inclusion of resources on malicious pages when visited by users who are authenticated with a Jupyter server. Access to the content of resources has been demonstrated with Internet Explorer through capturing of error messages, though not reproduced with other browsers. This occurs because Internet Explorer's error messages can include the content of any invalid JavaScript that was encountered.
CVE-2019-10255	An Open Redirect vulnerability for all browsers in Jupyter Notebook before 5.7.7 and some browsers (Chrome, Firefox) in JupyterHub before 0.9.5 allows crafted links to the login page, which will redirect to a malicious site after successful login. Servers running on a base_url prefix are not affected.
CVE-2019-10856	In Jupyter Notebook before 5.7.8, an open redirect can occur via an empty netloc. This issue exists because of an incomplete fix for CVE-2019-10255.
CVE-2019-11358	jQuery before 3.4.0, as used in Drupal, Backdrop CMS, and other products, mishandles jQuery.extend(true, {}, ...) because of Object.prototype pollution. If an unsanitized source object contained an enumerable __proto__ property, it could extend the native Object.prototype.
CVE-2020-11022	In jQuery versions greater than or equal to 1.2 and before 3.5.0, passing HTML from untrusted sources - even after sanitizing it - to one of jQuery's DOM manipulation methods (i.e. .html(), .append(), and others) may execute untrusted code. This problem is patched in jQuery 3.5.0.
CVE-2020-11023	In jQuery versions greater than or equal to 1.0.3 and before 3.5.0, passing HTML containing <option> elements from untrusted sources - even after sanitizing it - to one of jQuery's DOM manipulation methods (i.e. .html(), .append(), and others) may execute untrusted code. This problem is patched in jQuery 3.5.0.
CVE-2020-14338	A flaw was found in Wildfly's implementation of Xerces, specifically in the way the XMLSchemaValidator class in the JAXP component of Wildfly enforced the "use-grammar-pool-only" feature. This flaw allows a specially-crafted XML file to manipulate the validation process in certain cases. This issue is the same flaw as CVE-2020-14621, which affected OpenJDK, and uses a similar code. This flaw affects all Xerces JBoss versions before 2.12.0.SP3.
CVE-2020-26215	Jupyter Notebook before version 6.1.5 has an Open redirect vulnerability. A maliciously crafted link to a notebook server could redirect the browser to a different website. All notebook servers are technically affected, however, these maliciously crafted links can only be reasonably made for known notebook server hosts. A link to your notebook server may appear safe, but ultimately redirect to a spoofed server on the public internet. The issue is patched in version 6.1.5.
CVE-2021-28170	In the Jakarta Expression Language implementation 3.0.3 and earlier, a bug in the ELParserTokenManager enables invalid EL expressions to be evaluated as if they were valid.
CVE-2021-32797	JupyterLab is a user interface for Project Jupyter which will eventually replace the classic Jupyter Notebook. In affected versions untrusted notebook can execute code on load. In particular JupyterLab doesn’t sanitize the action attribute of html `<form>`. Using this it is possible to trigger the form validation outside of the form itself. This is a remote code execution, but requires user action to open a notebook.
CVE-2021-32798	The Jupyter notebook is a web-based notebook environment for interactive computing. In affected versions untrusted notebook can execute code on load. Jupyter Notebook uses a deprecated version of Google Caja to sanitize user inputs. A public Caja bypass can be used to trigger an XSS when a victim opens a malicious ipynb document in Jupyter Notebook. The XSS allows an attacker to execute arbitrary code on the victim computer using Jupyter APIs.
CVE-2021-41182	jQuery-UI is the official jQuery user interface library. Prior to version 1.13.0, accepting the value of the `altField` option of the Datepicker widget from untrusted sources may execute untrusted code. The issue is fixed in jQuery UI 1.13.0. Any string value passed to the `altField` option is now treated as a CSS selector. A workaround is to not accept the value of the `altField` option from untrusted sources.
CVE-2021-41183	jQuery-UI is the official jQuery user interface library. Prior to version 1.13.0, accepting the value of various `*Text` options of the Datepicker widget from untrusted sources may execute untrusted code. The issue is fixed in jQuery UI 1.13.0. The values passed to various `*Text` options are now always treated as pure text, not HTML. A workaround is to not accept the value of the `*Text` options from untrusted sources.

CVE ID	Description
CVE-2021-41184	jQuery-UI is the official jQuery user interface library. Prior to version 1.13.0, accepting the value of the `of` option of the `.position()` util from untrusted sources may execute untrusted code. The issue is fixed in jQuery UI 1.13.0. Any string value passed to the `of` option is now treated as a CSS selector. A workaround is to not accept the value of the `of` option from untrusted sources.
CVE-2022-23437	There's a vulnerability within the Apache Xerces Java (XercesJ) XML parser when handling specially crafted XML document payloads. This causes, the XercesJ XML parser to wait in an infinite loop, which may sometimes consume system resources for prolonged duration. This vulnerability is present within XercesJ version 2.12.1 and the previous versions.
CVE-2022-24758	The Jupyter notebook is a web-based notebook environment for interactive computing. Prior to version 6.4.9, unauthorized actors can access sensitive information from server logs. Anytime a 5xx error is triggered, the auth cookie and other header values are recorded in Jupyter server logs by default. Considering these logs do not require root access, an attacker can monitor these logs, steal sensitive auth/cookie information, and gain access to the Jupyter server. Jupyter notebook version 6.4.x contains a patch for this issue. There are currently no known workarounds.
CVE-2022-29238	Jupyter Notebook is a web-based notebook environment for interactive computing. Prior to version 6.4.12, authenticated requests to the notebook server with `ContentsManager.allow_hidden = False` only prevented listing the contents of hidden directories, not accessing individual hidden files or files in hidden directories (i.e. hidden files were 'hidden' but not 'inaccessible'). This could lead to notebook configurations allowing authenticated access to files that may reasonably be expected to be disallowed. Because fully authenticated requests are required, this is of relatively low impact. But if a server's root directory contains sensitive files whose only protection from the server is being hidden (e.g. `~/.ssh` while serving $HOME), then any authenticated requests could access files if their names are guessable. Such contexts also necessarily have full access to the server and therefore execution permissions, which also generally grants access to all the same files. So this does not generally result in any privilege escalation or increase in information access, only an additional, unintended means by which the files could be accessed. Version 6.4.12 contains a patch for this issue. There are currently no known workarounds.
CVE-2022-31160	jQuery UI is a curated set of user interface interactions, effects, widgets, and themes built on top of jQuery. Versions prior to 1.13.2 are potentially vulnerable to cross-site scripting. Initializing a checkboxradio widget on an input enclosed within a label makes that parent label contents considered as the input label. Calling `.checkboxradio( "refresh" )` on such a widget and the initial HTML contained encoded HTML entities will make them erroneously get decoded. This can lead to potentially executing JavaScript code. The bug has been patched in jQuery UI 1.13.2. To remediate the issue, someone who can change the initial HTML can wrap all the non-input contents of the `label` in a `span`.
CVE-2022-40897	Python Packaging Authority (PyPA) setuptools before 65.5.1 allows remote attackers to cause a denial of service via HTML in a crafted package or custom PackageIndex page. There is a Regular Expression Denial of Service (ReDoS) in package_index.py.
CVE-2023-5752	When installing a package from a Mercurial VCS URL (ie "pip install hg+...") with pip prior to v23.3, the specified Mercurial revision could be used to inject arbitrary configuration options to the "hg clone" call (ie "--config"). Controlling the Mercurial configuration can modify how and which repository is installed. This vulnerability does not affect users who aren't installing from Mercurial.
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-2023-40577	Alertmanager handles alerts sent by client applications such as the Prometheus server. An attacker with the permission to perform POST requests on the /api/v1/alerts endpoint could be able to execute arbitrary JavaScript code on the users of Prometheus Alertmanager. This issue has been fixed in Alertmanager version 0.2.51.
CVE-2023-44483	All versions of Apache Santuario - XML Security for Java prior to 2.2.6, 2.3.4, and 3.0.3, when using the JSR 105 API, are vulnerable to an issue where a private key may be disclosed in log files when generating an XML Signature and logging with debug level is enabled. Users are recommended to upgrade to version 2.2.6, 2.3.4, or 3.0.3, which fixes this issue.
CVE-2023-45288	An attacker may cause an HTTP/2 endpoint to read arbitrary amounts of header data by sending an excessive number of CONTINUATION frames. Maintaining HPACK state requires parsing and processing all HEADERS and CONTINUATION frames on a connection. When a request's headers exceed MaxHeaderBytes, no memory is allocated to store the excess headers, but they are still parsed. This permits an attacker to cause an HTTP/2 endpoint to read arbitrary amounts of header data, all associated with a request which is going to be rejected. These headers can include Huffman-encoded data which is significantly more expensive for the receiver to decode than for an attacker to send. The fix sets a limit on the amount of excess header frames we will process before closing a connection.

CVE ID	Description
CVE-2023-45289	When following an HTTP redirect to a domain which is not a subdomain match or exact match of the initial domain, an http.Client does not forward sensitive headers such as "Authorization" or "Cookie". For example, a redirect from foo.com to www.foo.com will forward the Authorization header, but a redirect to bar.com will not. A maliciously crafted HTTP redirect could cause sensitive headers to be unexpectedly forwarded.
CVE-2023-45290	When parsing a multipart form (either explicitly with Request.ParseMultipartForm or implicitly with Request.FormValue, Request.PostFormValue, or Request.FormFile), limits on the total size of the parsed form were not applied to the memory consumed while reading a single form line. This permits a maliciously crafted input containing very long lines to cause allocation of arbitrarily large amounts of memory, potentially leading to memory exhaustion. With fix, the ParseMultipartForm function now correctly limits the maximum size of form lines.
CVE-2023-50298	Exposure of Sensitive Information to an Unauthorized Actor vulnerability in Apache Solr.This issue affects Apache Solr: from 6.0.0 through 8.11.2, from 9.0.0 before 9.4.1. Solr Streaming Expressions allows users to extract data from other Solr Clouds, using a "zkHost" parameter. When original SolrCloud is setup to use ZooKeeper credentials and ACLs, they will be sent to whatever "zkHost" the user provides. An attacker could setup a server to mock ZooKeeper, that accepts ZooKeeper requests with credentials and ACLs and extracts the sensitive information, then send a streaming expression using the mock server's address in "zkHost". Streaming Expressions are exposed via the "/streaming" handler, with "read" permissions. Users are recommended to upgrade to version 8.11.3 or 9.4.1, which fix the issue. From these versions on, only zkHost values that have the same server address (regardless of chroot), will use the given ZooKeeper credentials and ACLs when connecting.
CVE-2023-51074	json-path v2.8.0 was discovered to contain a stack overflow via the Criteria.parse() method.
CVE-2024-4340	Passing a heavily nested list to sqlparse.parse() leads to a Denial of Service due to RecursionError.
CVE-2024-5206	A sensitive data leakage vulnerability was identified in scikit-learn's TfidfVectorizer, specifically in versions up to and including 1.4.1.post1, which was fixed in version 1.5.0. The vulnerability arises from the unexpected storage of all tokens present in the training data within the `stop_words_` attribute, rather than only storing the subset of tokens required for the TF-IDF technique to function. This behavior leads to the potential leakage of sensitive information, as the `stop_words_` attribute could contain tokens that were meant to be discarded and not stored, such as passwords or keys. The impact of this vulnerability varies based on the nature of the data being processed by the vectorizer.
CVE-2024-6104	go-retryablehttp prior to 0.7.7 did not sanitize urls when writing them to its log file. This could lead to go-retryablehttp writing sensitive HTTP basic auth credentials to its log file. This vulnerability, CVE-2024-6104, was fixed in go-retryablehttp 0.7.7.
CVE-2024-6345	A vulnerability in the package_index module of pypa/setuptools versions up to 69.1.1 allows for remote code execution via its download functions. These functions, which are used to download packages from URLs provided by users or retrieved from package index servers, are susceptible to code injection. If these functions are exposed to user-controlled inputs, such as package URLs, they can execute arbitrary commands on the system. The issue is fixed in version 70.0.
CVE-2024-6485	A security vulnerability has been discovered in bootstrap that could enable Cross-Site Scripting (XSS) attacks. The vulnerability is associated with the data-loading-text attribute within the button plugin. This vulnerability can be exploited by injecting malicious JavaScript code into the attribute, which would then be executed when the button's loading state is triggered.
CVE-2024-6763	Eclipse Jetty is a lightweight, highly scalable, Java-based web server and Servlet engine . It includes a utility class, HttpURI, for URI/URL parsing. The HttpURI class does insufficient validation on the authority segment of a URI. However the behaviour of HttpURI differs from the common browsers in how it handles a URI that would be considered invalid if fully validated against the RRC. Specifically HttpURI and the browser may differ on the value of the host extracted from an invalid URI and thus a combination of Jetty and a vulnerable browser may be vulnerable to a open redirect attack or to a SSRF attack if the URI is used after passing validation checks.
CVE-2024-7254	Any project that parses untrusted Protocol Buffers data containing an arbitrary number of nested groups / series of SGROUP tags can corrupted by exceeding the stack limit i.e. StackOverflow. Parsing nested groups as unknown fields with DiscardUnknownFieldsParser or Java Protobuf Lite parser, or against Protobuf map fields, creates unbounded recursions that can be abused by an attacker.

CVE ID	Description
CVE-2024-12797	Issue summary: Clients using RFC7250 Raw Public Keys (RPKs) to authenticate a server may fail to notice that the server was not authenticated, because handshakes don't abort as expected when the SSL_VERIFY_PEER verification mode is set. Impact summary: TLS and DTLS connections using raw public keys may be vulnerable to man-in-middle attacks when server authentication failure is not detected by clients. RPKs are disabled by default in both TLS clients and TLS servers. The issue only arises when TLS clients explicitly enable RPK use by the server, and the server, likewise, enables sending of an RPK instead of an X.509 certificate chain. The affected clients are those that then rely on the handshake to fail when the server's RPK fails to match one of the expected public keys, by setting the verification mode to SSL_VERIFY_PEER. Clients that enable server-side raw public keys can still find out that raw public key verification failed by calling SSL_get_verify_result(), and those that do, and take appropriate action, are not affected. This issue was introduced in the initial implementation of RPK support in OpenSSL 3.2. The FIPS modules in 3.4, 3.3, 3.2, 3.1 and 3.0 are not affected by this issue.
CVE-2024-12798	ACE vulnerability in JaninoEventEvaluator by QOS.CH logback-core upto including version 0.1 to 1.3.14 and 1.4.0 to 1.5.12 in Java applications allows attacker to execute arbitrary code by compromising an existing logback configuration file or by injecting an environment variable before program execution. Malicious logback configuration files can allow the attacker to execute arbitrary code using the JaninoEventEvaluator extension. A successful attack requires the user to have write access to a configuration file. Alternatively, the attacker could inject a malicious environment variable pointing to a malicious configuration file. In both cases, the attack requires existing privilege.
CVE-2024-12801	Server-Side Request Forgery (SSRF) in SaxEventRecorder by QOS.CH logback version 0.1 to 1.3.14 and 1.4.0 to 1.5.12 on the Java platform, allows an attacker to forge requests by compromising logback configuration files in XML. The attacks involves the modification of DOCTYPE declaration in XML configuration files.
CVE-2024-21742	Improper input validation allows for header injection in MIME4J library when using MIME4J DOM for composing message. This can be exploited by an attacker to add unintended headers to MIME messages.
CVE-2024-23945	Signing cookies is an application security feature that adds a digital signature to cookie data to verify its authenticity and integrity. The signature helps prevent malicious actors from modifying the cookie value, which can lead to security vulnerabilities and exploitation. Apache Hive’s service component accidentally exposes the signed cookie to the end user when there is a mismatch in signature between the current and expected cookie. Exposing the correct cookie signature can lead to further exploitation. The vulnerable CookieSigner logic was introduced in Apache Hive by HIVE-9710 (1.2.0) and in Apache Spark by SPARK-14987 (2.0.0). The affected components are the following: * org.apache.hive:hive-service * org.apache.spark:spark-hive-thriftserver_2.11 * org.apache.spark:spark-hive-thriftserver_2.12
CVE-2024-23953	Use of Arrays.equals() in LlapSignerImpl in Apache Hive to compare message signatures allows attacker to forge a valid signature for an arbitrary message byte by byte. The attacker should be an authorized user of the product to perform this attack. Users are recommended to upgrade to version 4.0.0, which fixes this issue. The problem occurs when an application doesn’t use a constant-time algorithm for validating a signature. The method Arrays.equals() returns false right away when it sees that one of the input’s bytes are different. It means that the comparison time depends on the contents of the arrays. This little thing may allow an attacker to forge a valid signature for an arbitrary message byte by byte. So it might allow malicious users to submit splits/work with selected signatures to LLAP without running as a privileged user, potentially leading to DDoS attack. More details in the reference section.
CVE-2024-24557	Moby is an open-source project created by Docker to enable software containerization. The classic builder cache system is prone to cache poisoning if the image is built FROM scratch. Also, changes to some instructions (most important being HEALTHCHECK and ONBUILD) would not cause a cache miss. An attacker with the knowledge of the Dockerfile someone is using could poison their cache by making them pull a specially crafted image that would be considered as a valid cache candidate for some build steps. 23.0+ users are only affected if they explicitly opted out of Buildkit (DOCKER_BUILDKIT=0 environment variable) or are using the /build API endpoint. All users on versions older than 23.0 could be impacted. Image build API endpoint (/build) and ImageBuild function from github.com/docker/docker/client is also affected as it the uses classic builder by default. Patches are included in 24.0.9 and 25.0.2 releases.
CVE-2024-24783	Verifying a certificate chain which contains a certificate with an unknown public key algorithm will cause Certificate.Verify to panic. This affects all crypto/tls clients, and servers that set Config.ClientAuth to VerifyClientCertIfGiven or RequireAndVerifyClientCert. The default behavior is for TLS servers to not verify client certificates.
CVE-2024-24784	The ParseAddressList function incorrectly handles comments (text within parentheses) within display names. Since this is a misalignment with conforming address parsers, it can result in different trust decisions being made by programs using different parsers.
CVE-2024-24785	If errors returned from MarshalJSON methods contain user controlled data, they may be used to break the contextual auto-escaping behavior of the html/template package, allowing for subsequent actions to inject unexpected content into templates.

CVE ID	Description
CVE-2024-24786	The protojson.Unmarshal function can enter an infinite loop when unmarshaling certain forms of invalid JSON. This condition can occur when unmarshaling into a message which contains a google.protobuf.Any value, or when the UnmarshalOptions.DiscardUnknown option is set.
CVE-2024-24789	The archive/zip package's handling of certain types of invalid zip files differs from the behavior of most zip implementations. This misalignment could be exploited to create an zip file with contents that vary depending on the implementation reading the file. The archive/zip package now rejects files containing these errors.
CVE-2024-24790	The various Is methods (IsPrivate, IsLoopback, etc) did not work as expected for IPv4-mapped IPv6 addresses, returning false for addresses which would return 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-continue" header with a non-informational (200 or higher) status. This mishandling could leave a client connection in an invalid state, where the next request sent on the connection will fail. An attacker sending a request to a net/http/httputil.ReverseProxy proxy can exploit this mishandling to cause a denial of service by sending "Expect: 100-continue" requests which elicit a non-informational response from the backend. Each such request leaves the proxy with an invalid connection, and causes one subsequent request using that connection to fail.
CVE-2024-29869	Hive creates a credentials file to a temporary directory in the file system with permissions 644 by default when the file permissions are not set explicitly. Any unauthorized user having access to the directory can read the sensitive information written into this file. Users are recommended to upgrade to version 4.0.1, which fixes this issue.
CVE-2024-34155	Calling any of the Parse functions on Go source code which contains deeply nested literals can cause a panic due to stack exhaustion.
CVE-2024-34156	Calling Decoder.Decode on a message which contains deeply nested structures can cause a panic due to stack exhaustion. This is a follow-up to CVE-2022-30635.
CVE-2024-34158	Calling Parse on a "// +build" build tag line with deeply nested expressions can cause a panic due to stack exhaustion.
CVE-2024-35255	Azure Identity Libraries and Microsoft Authentication Library Elevation of Privilege Vulnerability
CVE-2024-41110	Moby is an open-source project created by Docker for software containerization. A security vulnerability has been detected in certain versions of Docker Engine, which could allow an attacker to bypass authorization plugins (AuthZ) under specific circumstances. The base likelihood of this being exploited is low. Using a specially-crafted API request, an Engine API client could make the daemon forward the request or response to an authorization plugin without the body. In certain circumstances, the authorization plugin may allow a request which it would have otherwise denied if the body had been forwarded to it. A security issue was discovered In 2018, where an attacker could bypass AuthZ plugins using a specially crafted API request. This could lead to unauthorized actions, including privilege escalation. Although this issue was fixed in Docker Engine v18.09.1 in January 2019, the fix was not carried forward to later major versions, resulting in a regression. Anyone who depends on authorization plugins that introspect the request and/or response body to make access control decisions is potentially impacted. Docker EE v19.03.x and all versions of Mirantis Container Runtime are not vulnerable. docker-ce v27.1.1 containes patches to fix the vulnerability. Patches have also been merged into the master, 19.03, 20.0, 23.0, 24.0, 25.0, 26.0, and 26.1 release branches. If one is unable to upgrade immediately, avoid using AuthZ plugins and/or restrict access to the Docker API to trusted parties, following the principle of least privilege.
CVE-2024-45336	The HTTP client drops sensitive headers after following a cross-domain redirect. For example, a request to a.com/ containing an Authorization header which is redirected to b.com/ will not send that header to b.com. In the event that the client received a subsequent same-domain redirect, however, the sensitive headers would be restored. For example, a chain of redirects from a.com/, to b.com/1, and finally to b.com/2 would incorrectly send the Authorization header to b.com/2.

CVE ID	Description
CVE-2024-45337	Applications and libraries which misuse connection.serverAuthenticate (via callback field ServerConfig.PublicKeyCallback) may be susceptible to an authorization bypass. The documentation for ServerConfig.PublicKeyCallback says that "A call to this function does not guarantee that the key offered is in fact used to authenticate." Specifically, the SSH protocol allows clients to inquire about whether a public key is acceptable before proving control of the corresponding private key. PublicKeyCallback may be called with multiple keys, and the order in which the keys were provided cannot be used to infer which key the client successfully authenticated with, if any. Some applications, which store the key(s) passed to PublicKeyCallback (or derived information) and make security relevant determinations based on it once the connection is established, may make incorrect assumptions. For example, an attacker may send public keys A and B, and then authenticate with A. PublicKeyCallback would be called only twice, first with A and then with B. A vulnerable application may then make authorization decisions based on key B for which the attacker does not actually control the private key. Since this API is widely misused, as a partial mitigation golang.org/x/cry...@v0.31.0 enforces the property that, when successfully authenticating via public key, the last key passed to ServerConfig.PublicKeyCallback will be the key used to authenticate the connection. PublicKeyCallback will now be called multiple times with the same key, if necessary. Note that the client may still not control the last key passed to PublicKeyCallback if the connection is then authenticated with a different method, such as PasswordCallback, KeyboardInteractiveCallback, or NoClientAuth. Users should be using the Extensions field of the Permissions return value from the various authentication callbacks to record data associated with the authentication attempt instead of referencing external state. Once the connection is established the state corresponding to the successful authentication attempt can be retrieved via the ServerConn.Permissions field. Note that some third-party libraries misuse the Permissions type by sharing it across authentication attempts; users of third-party libraries should refer to the relevant projects for guidance.
CVE-2024-45338	An attacker can craft an input to the Parse functions that would be processed non-linearly with respect to its length, resulting in extremely slow parsing. This could cause a denial of service.
CVE-2024-45341	A certificate with a URI which has a IPv6 address with a zone ID may incorrectly satisfy a URI name constraint that applies to the certificate chain. Certificates containing URIs are not permitted in the web PKI, so this only affects users of private PKIs which make use of URIs.
CVE-2024-47561	Schema parsing in the Java SDK of Apache Avro 1.11.3 and previous versions allows bad actors to execute arbitrary code. Users are recommended to upgrade to version 1.11.4 or 1.12.0, which fix this issue.
CVE-2024-51744	golang-jwt is a Go implementation of JSON Web Tokens. Unclear documentation of the error behavior in `ParseWithClaims` can lead to situation where users are potentially not checking errors in the way they should be. Especially, if a token is both expired and invalid, the errors returned by `ParseWithClaims` return both error codes. If users only check for the `jwt.ErrTokenExpired ` using `error.Is`, they will ignore the embedded `jwt.ErrTokenSignatureInvalid` and thus potentially accept invalid tokens. A fix has been back-ported with the error handling logic from the `v5` branch to the `v4` branch. In this logic, the `ParseWithClaims` function will immediately return in "dangerous" situations (e.g., an invalid signature), limiting the combined errors only to situations where the signature is valid, but further validation failed (e.g., if the signature is valid, but is expired AND has the wrong audience). This fix is part of the 4.5.1 release. We are aware that this changes the behaviour of an established function and is not 100 % backwards compatible, so updating to 4.5.1 might break your code. In case you cannot update to 4.5.0, please make sure that you are properly checking for all errors ("dangerous" ones first), so that you are not running in the case detailed above.
CVE-2024-53899	virtualenv before 20.26.6 allows command injection through the activation scripts for a virtual environment. Magic template strings are not quoted correctly when replacing. NOTE: this is not the same as CVE-2024-9287.
CVE-2024-57699	A security issue was found in Netplex Json-smart 2.5.0 through 2.5.1. When loading a specially crafted JSON input, containing a large number of ’{’, a stack exhaustion can be trigger, which could allow an attacker to cause a Denial of Service (DoS). This issue exists because of an incomplete fix for CVE-2023-1370.
CVE-2025-0913	os.OpenFile(path, os.O_CREATE|O_EXCL) behaved differently on Unix and Windows systems when the target path was a dangling symlink. On Unix systems, OpenFile with O_CREATE and O_EXCL flags never follows symlinks. On Windows, when the target path was a symlink to a nonexistent location, OpenFile would create a file in that location. OpenFile now always returns an error when the O_CREATE and O_EXCL flags are both set and the target path is a symlink.
CVE-2025-1647	Improper Neutralization of Input During Web Page Generation (XSS or 'Cross-site Scripting') vulnerability in Bootstrap allows Cross-Site Scripting (XSS).This issue affects Bootstrap: from 3.4.1 before 4.0.0.
CVE-2025-1767	This CVE only affects Kubernetes clusters that utilize the in-tree gitRepo volume to clone git repositories from other pods within the same node. Since the in-tree gitRepo volume feature has been deprecated and will not receive security updates upstream, any cluster still using this feature remains vulnerable.
CVE-2025-2999	A vulnerability was found in PyTorch 2.6.0. It has been rated as critical. Affected by this issue is the function torch.nn.utils.rnn.unpack_sequence. The manipulation leads to memory corruption. Attacking locally is a requirement. The exploit has been disclosed to the public and may be used.

CVE ID	Description
CVE-2025-3000	A vulnerability classified as critical has been found in PyTorch 2.6.0. This affects the function torch.jit.script. The manipulation leads to memory corruption. It is possible to launch the attack on the local host. The exploit has been disclosed to the public and may be used.
CVE-2025-3001	A vulnerability classified as critical was found in PyTorch 2.6.0. This vulnerability affects the function torch.lstm_cell. The manipulation leads to memory corruption. The attack needs to be approached locally. The exploit has been disclosed to the public and may be used.
CVE-2025-3588	A vulnerability, which was classified as problematic, has been found in joelittlejohn jsonschema2pojo 1.2.2. This issue affects the function apply of the file org/jsonschema2pojo/rules/SchemaRule.java of the component JSON File Handler. The manipulation leads to stack-based buffer overflow. Attacking locally is a requirement. The exploit has been disclosed to the public and may be used. The vendor was contacted early about this disclosure but did not respond in any way.
CVE-2025-4565	Any project that uses Protobuf Pure-Python backend to parse untrusted Protocol Buffers data containing an arbitrary number of recursive groups, recursive messages or a series of SGROUP tags can be corrupted by exceeding the Python recursion limit. This can result in a Denial of service by crashing the application with a RecursionError. We recommend upgrading to version =>6.31.1 or beyond commit 17838beda2943d08b8a9d4df5b68f5f04f26d901
CVE-2025-4673	Proxy-Authorization and Proxy-Authenticate headers persisted on cross-origin redirects potentially leaking sensitive information.
CVE-2025-4949	In Eclipse JGit versions 7.2.0.202503040940-r and older, the ManifestParser class used by the repo command and the AmazonS3 class used to implement the experimental amazons3 git transport protocol allowing to store git pack files in an Amazon S3 bucket, are vulnerable to XML External Entity (XXE) attacks when parsing XML files. This vulnerability can lead to information disclosure, denial of service, and other security issues.
CVE-2025-5187	A vulnerability exists in the NodeRestriction admission controller in Kubernetes clusters where node users can delete their corresponding node object by patching themselves with an OwnerReference to a cluster-scoped resource. If the OwnerReference resource does not exist or is subsequently deleted, the given node object will be deleted via garbage collection.
CVE-2025-7962	In Jakarta Mail versions prior to 2.0.2 it is possible to perform an SMTP Injection by utilizing the \r and \n UTF-8 characters to separate different messages.
CVE-2025-8869	When extracting a tar archive pip may not check symbolic links point into the extraction directory if the tarfile module doesn't implement PEP 706. Note that upgrading pip to a "fixed" version for this vulnerability doesn't fix all known vulnerabilities that are remediated by using a Python version that implements PEP 706. Note that this is a vulnerability in pip's fallback implementation of tar extraction for Python versions that don't implement PEP 706 and therefore are not secure to all vulnerabilities in the Python 'tarfile' module. If you're using a Python version that implements PEP 706 then pip doesn't use the "vulnerable" fallback code. Mitigations include upgrading to a version of pip that includes the fix, upgrading to a Python version that implements PEP 706 (Python >=3.9.17, >=3.10.12, >=3.11.4, or >=3.12), applying the linked patch, or inspecting source distributions (sdists) before installation as is already a best-practice.
CVE-2025-8916	Allocation of Resources Without Limits or Throttling vulnerability in Legion of the Bouncy Castle Inc. BC Java bcpkix on All (API modules), Legion of the Bouncy Castle Inc. BC Java bcprov on All (API modules), Legion of the Bouncy Castle Inc. BCPKIX FIPS bcpkix-fips on All (API modules) allows Excessive Allocation. This vulnerability is associated with program files https://github.Com/bcgit/bc-java/blob/main/pkix/src/main/java/org/bouncycastle/pkix/jcajce/PKIXCertPathReviewer.Java, https://github.Com/bcgit/bc-java/blob/main/prov/src/main/java/org/bouncycastle/x509/PKIXCertPathReviewer.Java. This issue affects BC Java: from 1.44 through 1.78; BC Java: from 1.44 through 1.78; BCPKIX FIPS: from 1.0.0 through 1.0.7, from 2.0.0 through 2.0.7.
CVE-2025-11226	ACE vulnerability in conditional configuration file processing by QOS.CH logback-core up to and including version 1.5.18 in Java applications, allows an attacker to execute arbitrary code by compromising an existing logback configuration file or by injecting an environment variable before program execution. A successful attack requires the presence of Janino library and Spring Framework to be present on the user's class path. In addition, the attacker must have write access to a configuration file. Alternatively, the attacker could inject a malicious environment variable pointing to a malicious configuration file. In both cases, the attack requires existing privilege.
CVE-2025-12183	Out-of-bounds memory operations in org.lz4:lz4-java 1.8.0 and earlier allow remote attackers to cause denial of service and read adjacent memory via untrusted compressed input.
CVE-2025-13281	A half-blind Server Side Request Forgery (SSRF) vulnerability exists in kube-controller-manager when using the in-tree Portworx StorageClass. This vulnerability allows authorized users to leak arbitrary information from unprotected endpoints in the control plane’s host network (including link-local or loopback services).
CVE-2025-14813	: Use of a Broken or Risky Cryptographic Algorithm vulnerability in Legion of the Bouncy Castle Inc. BC-JAVA bcprov on all (core modules). This vulnerability is associated with program files G3413CTRBlockCipher. This issue affects BC-JAVA: from 1.59 before 1.80.2, from 1.81 before 1.81.1, from 1.82 before 1.84.

CVE ID	Description
CVE-2025-15599	DOMPurify 3.1.3 through 3.2.6 and 2.5.3 through 2.5.8 contain a cross-site scripting vulnerability that allows attackers to bypass attribute sanitization by exploiting missing textarea rawtext element validation in the SAFE_FOR_XML regex. Attackers can include closing rawtext tags like </textarea> in attribute values to break out of rawtext contexts and execute JavaScript when sanitized output is placed inside rawtext elements. The 3.x branch was fixed in 3.2.7; the 2.x branch was never patched.
CVE-2025-22866	Due to the usage of a variable time instruction in the assembly implementation of an internal function, a small number of bits of secret scalars are leaked on the ppc64le architecture. Due to the way this function is used, we do not believe this leakage is enough to allow recovery of the private key when P-256 is used in any well known protocols.
CVE-2025-22868	An attacker can pass a malicious malformed token which causes unexpected memory to be consumed during parsing.
CVE-2025-22869	SSH servers which implement file transfer protocols are vulnerable to a denial of service attack from clients which complete the key exchange slowly, or not at all, causing pending content to be read into memory, but never transmitted.
CVE-2025-22870	Matching of hosts against proxy patterns can improperly treat an IPv6 zone ID as a hostname component. For example, when the NO_PROXY environment variable is set to "*.example.com", a request to "[::1%25.example.com]:80` will incorrectly match and not be proxied.
CVE-2025-22871	The net/http package improperly accepts a bare LF as a line terminator in chunked data chunk-size lines. This can permit request smuggling if a net/http server is used in conjunction with a server that incorrectly accepts a bare LF as part of a chunk-ext.
CVE-2025-22872	The tokenizer incorrectly interprets tags with unquoted attribute values that end with a solidus character (/) as self-closing. When directly using Tokenizer, this can result in such tags incorrectly being marked as self-closing, and when using the Parse functions, this can result in content following such tags as being placed in the wrong scope during DOM construction, but only when tags are in foreign content (e.g. <math>, <svg>, etc contexts).
CVE-2025-22873	It was possible to improperly access the parent directory of an os.Root by opening a filename ending in "../". For example, Root.Open("../") would open the parent directory of the Root. This escape only permits opening the parent directory itself, not ancestors of the parent or files contained within the parent.
CVE-2025-26796	** UNSUPPORTED WHEN ASSIGNED ** Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') vulnerability in Apache Oozie. This issue affects Apache Oozie: all versions. As this project is retired, we do not plan to release a version that fixes this issue. Users are recommended to find an alternative or restrict access to the instance to trusted users. NOTE: This vulnerability only affects products that are no longer supported by the maintainer.
CVE-2025-27820	A bug in PSL validation logic in Apache HttpClient 5.4.x disables domain checks, affecting cookie management and host name verification. Discovered by the Apache HttpClient team. Fixed in the 5.4.3 release
CVE-2025-30204	golang-jwt is a Go implementation of JSON Web Tokens. Starting in version 3.2.0 and prior to versions 5.2.2 and 4.5.2, the function parse.ParseUnverified splits (via a call to strings.Split) its argument (which is untrusted data) on periods. As a result, in the face of a malicious request whose Authorization header consists of Bearer followed by many period characters, a call to that function incurs allocations to the tune of O(n) bytes (where n stands for the length of the function's argument), with a constant factor of about 16. This issue is fixed in 5.2.2 and 4.5.2.
CVE-2025-33042	Improper Control of Generation of Code ('Code Injection') vulnerability in Apache Avro Java SDK when generating specific records from untrusted Avro schemas. This issue affects Apache Avro Java SDK: all versions through 1.11.4 and version 1.12.0. Users are recommended to upgrade to version 1.12.1 or 1.11.5, which fix the issue.
CVE-2025-46392	Uncontrolled Resource Consumption vulnerability in Apache Commons Configuration 1.x. There are a number of issues in Apache Commons Configuration 1.x that allow excessive resource consumption when loading untrusted configurations or using unexpected usage patterns. The Apache Commons Configuration team does not intend to fix these issues in 1.x. Apache Commons Configuration 1.x is still safe to use in scenario's where you only load trusted configurations. Users that load untrusted configurations or give attackers control over usage patterns are recommended to upgrade to the 2.x version line, which fixes these issues. Apache Commons Configuration 2.x is not a drop-in replacement, but as it uses a separate Maven groupId and Java package namespace they can be loaded side-by-side, making it possible to do a gradual migration.
CVE-2025-47273	setuptools is a package that allows users to download, build, install, upgrade, and uninstall Python packages. A path traversal vulnerability in `PackageIndex` is present in setuptools prior to version 78.1.1. An attacker would be allowed to write files to arbitrary locations on the filesystem with the permissions of the process running the Python code, which could escalate to remote code execution depending on the context. Version 78.1.1 fixes the issue.
CVE-2025-47906	If the PATH environment variable contains paths which are executables (rather than just directories), passing certain strings to LookPath ("", ".", and ".."), can result in the binaries listed in the PATH being unexpectedly returned.

CVE ID	Description
CVE-2025-47907	Cancelling a query (e.g. by cancelling the context passed to one of the query methods) during a call to the Scan method of the returned Rows can result in unexpected results if other queries are being made in parallel. This can result in a race condition that may overwrite the expected results with those of another query, causing the call to Scan to return either unexpected results from the other query or an error.
CVE-2025-47911	The html.Parse function in golang.org/x/net/html has quadratic parsing complexity when processing certain inputs, which can lead to denial of service (DoS) if an attacker provides specially crafted HTML content.
CVE-2025-47912	The Parse function permits values other than IPv6 addresses to be included in square brackets within the host component of a URL. RFC 3986 permits IPv6 addresses to be included within the host component, enclosed within square brackets. For example: "http://[::1]/". IPv4 addresses and hostnames must not appear within square brackets. Parse did not enforce this requirement.
CVE-2025-47913	SSH clients receiving SSH_AGENT_SUCCESS when expecting a typed response will panic and cause early termination of the client process.
CVE-2025-47914	SSH Agent servers do not validate the size of messages when processing new identity requests, which may cause the program to panic if the message is malformed due to an out of bounds read.
CVE-2025-48924	Uncontrolled Recursion vulnerability in Apache Commons Lang. This issue affects Apache Commons Lang: Starting with commons-lang:commons-lang 2.0 to 2.6, and, from org.apache.commons:commons-lang3 3.0 before 3.18.0. The methods ClassUtils.getClass(...) can throw StackOverflowError on very long inputs. Because an Error is usually not handled by applications and libraries, a StackOverflowError could cause an application to stop. Users are recommended to upgrade to version 3.18.0, which fixes the issue.
CVE-2025-50181	urllib3 is a user-friendly HTTP client library for Python. Prior to 2.5.0, it is possible to disable redirects for all requests by instantiating a PoolManager and specifying retries in a way that disable redirects. By default, requests and botocore users are not affected. An application attempting to mitigate SSRF or open redirect vulnerabilities by disabling redirects at the PoolManager level will remain vulnerable. This issue has been patched in version 2.5.0.
CVE-2025-52999	jackson-core contains core low-level incremental ("streaming") parser and generator abstractions used by Jackson Data Processor. In versions prior to 2.15.0, if a user parses an input file and it has deeply nested data, Jackson could end up throwing a StackoverflowError if the depth is particularly large. jackson-core 2.15.0 contains a configurable limit for how deep Jackson will traverse in an input document, defaulting to an allowable depth of 1000. jackson-core will throw a StreamConstraintsException if the limit is reached. jackson-databind also benefits from this change because it uses jackson-core to parse JSON inputs. As a workaround, users should avoid parsing input files from untrusted sources.
CVE-2025-53864	Connect2id Nimbus JOSE + JWT 10.0.x before 10.0.2 and 9.37.x before 9.37.4 allows a remote attacker to cause a denial of service via a deeply nested JSON object supplied in a JWT claim set, because of uncontrolled recursion. NOTE: this is independent of the Gson 2.11.0 issue because the Connect2id product could have checked the JSON object nesting depth, regardless of what limits (if any) were imposed by Gson.
CVE-2025-54410	Moby is an open source container framework developed by Docker Inc. that is distributed as Docker Engine, Mirantis Container Runtime, and various other downstream projects/products. A firewalld vulnerability affects Moby releases before 28.0.0. When firewalld reloads, Docker fails to re-create iptables rules that isolate bridge networks, allowing any container to access all ports on any other container across different bridge networks on the same host. This breaks network segmentation between containers that should be isolated, creating significant risk in multi-tenant environments. Only containers in --internal networks remain protected. Workarounds include reloading firewalld and either restarting the docker daemon, re-creating bridge networks, or using rootless mode. Maintainers anticipate a fix for this issue in version 25.0.13.
CVE-2025-55163	Netty is an asynchronous, event-driven network application framework. Prior to versions 4.1.124.Final and 4.2.4.Final, Netty is vulnerable to MadeYouReset DDoS. This is a logical vulnerability in the HTTP/2 protocol, that uses malformed HTTP/2 control frames in order to break the max concurrent streams limit - which results in resource exhaustion and distributed denial of service. This issue has been patched in versions 4.1.124.Final and 4.2.4.Final.
CVE-2025-58057	Netty is an asynchronous event-driven network application framework for rapid development of maintainable high performance protocol servers & clients. In netty-codec-compression versions 4.1.124.Final and below, and netty-codec versions 4.2.4.Final and below, when supplied with specially crafted input, BrotliDecoder and certain other decompression decoders will allocate a large number of reachable byte buffers, which can lead to denial of service. BrotliDecoder.decompress has no limit in how often it calls pull, decompressing data 64K bytes at a time. The buffers are saved in the output list, and remain reachable until OOM is hit. This is fixed in versions 4.1.125.Final of netty-codec and 4.2.5.Final of netty-codec-compression.
CVE-2025-58181	SSH servers parsing GSSAPI authentication requests do not validate the number of mechanisms specified in the request, allowing an attacker to cause unbounded memory consumption.

CVE ID	Description
CVE-2025-58183	tar.Reader does not set a maximum size on the number of sparse region data blocks in GNU tar pax 1.0 sparse files. A maliciously-crafted archive containing a large number of sparse regions can cause a Reader to read an unbounded amount of data from the archive into memory. When reading from a compressed source, a small compressed input can result in large allocations.
CVE-2025-58185	Parsing a maliciously crafted DER payload could allocate large amounts of memory, causing memory exhaustion.
CVE-2025-58186	Despite HTTP headers having a default limit of 1MB, the number of cookies that can be parsed does not have a limit. By sending a lot of very small cookies such as "a=;", an attacker can make an HTTP server allocate a large amount of structs, causing large memory consumption.
CVE-2025-58187	Due to the design of the name constraint checking algorithm, the processing time of some inputs scale non-linearly with respect to the size of the certificate. This affects programs which validate arbitrary certificate chains.
CVE-2025-58188	Validating certificate chains which contain DSA public keys can cause programs to panic, due to a interface cast that assumes they implement the Equal method. This affects programs which validate arbitrary certificate chains.
CVE-2025-58189	When Conn.Handshake fails during ALPN negotiation the error contains attacker controlled information (the ALPN protocols sent by the client) which is not escaped.
CVE-2025-58190	The html.Parse function in golang.org/x/net/html has an infinite parsing loop when processing certain inputs, which can lead to denial of service (DoS) if an attacker provides specially crafted HTML content.
CVE-2025-60012	Malicious configuration can lead to unauthorized file access in Apache Livy. This issue affects Apache Livy 0.7.0 and 0.8.0 when connecting to Apache Spark 3.1 or later. A request that includes a Spark configuration value supported from Apache Spark version 3.1 can lead to users gaining access to files they do not have permissions to. For the vulnerability to be exploitable, the user needs to have access to Apache Livy's REST or JDBC interface and be able to send requests with arbitrary Spark configuration values. Users are recommended to upgrade to version 0.9.0 or later, which fixes the issue.
CVE-2025-61723	The processing time for parsing some invalid inputs scales non-linearly with respect to the size of the input. This affects programs which parse untrusted PEM inputs.
CVE-2025-61724	The Reader.ReadResponse function constructs a response string through repeated string concatenation of lines. When the number of lines in a response is large, this can cause excessive CPU consumption.
CVE-2025-61725	The ParseAddress function constructs domain-literal address components through repeated string concatenation. When parsing large domain-literal components, this can cause excessive CPU consumption.
CVE-2025-61726	The net/url package does not set a limit on the number of query parameters in a query. While the maximum size of query parameters in URLs is generally limited by the maximum request header size, the net/http.Request.ParseForm method can parse large URL-encoded forms. Parsing a large form containing many unique query parameters can cause excessive memory consumption.
CVE-2025-61727	An excluded subdomain constraint in a certificate chain does not restrict the usage of wildcard SANs in the leaf certificate. For example a constraint that excludes the subdomain test.example.com does not prevent a leaf certificate from claiming the SAN *.example.com.
CVE-2025-61728	archive/zip uses a super-linear file name indexing algorithm that is invoked the first time a file in an archive is opened. This can lead to a denial of service when consuming a maliciously constructed ZIP archive.
CVE-2025-61729	Within HostnameError.Error(), when constructing an error string, there is no limit to the number of hosts that will be printed out. Furthermore, the error string is constructed by repeated string concatenation, leading to quadratic runtime. Therefore, a certificate provided by a malicious actor can result in excessive resource consumption.
CVE-2025-61730	During the TLS 1.3 handshake if multiple messages are sent in records that span encryption level boundaries (for instance the Client Hello and Encrypted Extensions messages), the subsequent messages may be processed before the encryption level changes. This can cause some minor information disclosure if a network-local attacker can inject messages during the handshake.
CVE-2025-61911	python-ldap is a lightweight directory access protocol (LDAP) client API for Python. In versions prior to 3.4.5, the sanitization method `ldap.filter.escape_filter_chars` can be tricked to skip escaping of special characters when a crafted `list` or `dict` is supplied as the `assertion_value` parameter, and the non-default `escape_mode=1` is configured. The method `ldap.filter.escape_filter_chars` supports 3 different escaping modes. `escape_mode=0` (default) and `escape_mode=2` happen to raise exceptions when a `list` or `dict` object is supplied as the `assertion_value` parameter. However, `escape_mode=1` computes without performing adequate logic to ensure a fully escaped return value. If an application relies on the vulnerable method in the `python-ldap` library to escape untrusted user input, an attacker might be able to abuse the vulnerability to launch ldap injection attacks which could potentially disclose or manipulate ldap data meant to be inaccessible to them. Version 3.4.5 fixes the issue by adding a type check at the start of the `ldap.filter.escape_filter_chars` method to raise an exception when the supplied `assertion_value` parameter is not of type `str`.

CVE ID	Description
CVE-2025-61912	python-ldap is a lightweight directory access protocol (LDAP) client API for Python. In versions prior to 3.4.5, ldap.dn.escape_dn_chars() escapes \x00 incorrectly by emitting a backslash followed by a literal NUL byte instead of the RFC-4514 hex form \00. Any application that uses this helper to construct DNs from untrusted input can be made to consistently fail before a request is sent to the LDAP server (e.g., AD), resulting in a client-side denial of service. Version 3.4.5 contains a patch for the issue.
CVE-2025-62198	An authenticated user can perform XSS. This issue affects Apache Atlas versions 2.4.0 and earlier. Users are recommended to upgrade to version 2.5.0, which fixes the issue.
CVE-2025-66249	Improper Limitation of a Pathname to a Restricted Directory ('Path Traversal') vulnerability in Apache Livy. This issue affects Apache Livy: from 0.3.0 before 0.9.0. The vulnerability can only be exploited with non-default Apache Livy Server settings. If the configuration value "livy.file.local-dir-whitelist" is set to a non-default value, the directory checking can be bypassed. Users are recommended to upgrade to version 0.9.0, which fixes the issue.
CVE-2025-66418	urllib3 is a user-friendly HTTP client library for Python. Starting in version 1.24 and prior to 2.6.0, the number of links in the decompression chain was unbounded allowing a malicious server to insert a virtually unlimited number of compression steps leading to high CPU usage and massive memory allocation for the decompressed data. This vulnerability is fixed in 2.6.0.
CVE-2025-66471	urllib3 is a user-friendly HTTP client library for Python. Starting in version 1.0 and prior to 2.6.0, the Streaming API improperly handles highly compressed data. urllib3's streaming API is designed for the efficient handling of large HTTP responses by reading the content in chunks, rather than loading the entire response body into memory at once. When streaming a compressed response, urllib3 can perform decoding or decompression based on the HTTP Content-Encoding header (e.g., gzip, deflate, br, or zstd). The library must read compressed data from the network and decompress it until the requested chunk size is met. Any resulting decompressed data that exceeds the requested amount is held in an internal buffer for the next read operation. The decompression logic could cause urllib3 to fully decode a small amount of highly compressed data in a single operation. This can result in excessive resource consumption (high CPU usage and massive memory allocation for the decompressed data.
CVE-2025-66566	yawkat LZ4 Java provides LZ4 compression for Java. Insufficient clearing of the output buffer in Java-based decompressor implementations in lz4-java 1.10.0 and earlier allows remote attackers to read previous buffer contents via crafted compressed input. In applications where the output buffer is reused without being cleared, this may lead to disclosure of sensitive data. JNI-based implementations are not affected. This vulnerability is fixed in 1.10.1.
CVE-2025-67030	Directory Traversal vulnerability in the extractFile method of org.codehaus.plexus.util.Expand in plexus-utils before 6d780b3378829318ba5c2d29547e0012d5b29642. This allows an attacker to execute arbitrary code
CVE-2025-67721	Aircompressor is a library with ports of the Snappy, LZO, LZ4, and Zstandard compression algorithms to Java. In versions 3.3 and below, incorrect handling of malformed data in Java-based decompressor implementations for Snappy and LZ4 allow remote attackers to read previous buffer contents via crafted compressed input. With certain crafted compressed inputs, elements from the output buffer can end up in the uncompressed output, potentially leaking sensitive data. This is relevant for applications that reuse the same output buffer to uncompress multiple inputs. This can be the case of a web server that allocates a fix-sized buffer for performance purposes. There is similar vulnerability in GHSA-cmp6-m4wj-q63q. This issue is fixed in version 3.4.
CVE-2025-67735	Netty is an asynchronous, event-driven network application framework. In versions prior to 4.1.129.Final and 4.2.8.Final, the `io.netty.handler.codec.http.HttpRequestEncoder` has a CRLF injection with the request URI when constructing a request. This leads to request smuggling when `HttpRequestEncoder` is used without proper sanitization of the URI. Any application / framework using `HttpRequestEncoder` can be subject to be abused to perform request smuggling using CRLF injection. Versions 4.1.129.Final and 4.2.8.Final fix the issue.
CVE-2025-68121	During session resumption in crypto/tls, if the underlying Config has its ClientCAs or RootCAs fields mutated between the initial handshake and the resumed handshake, the resumed handshake may succeed when it should have failed. This may happen when a user calls Config.Clone and mutates the returned Config, or uses Config.GetConfigForClient. This can cause a client to resume a session with a server that it would not have resumed with during the initial handshake, or cause a server to resume a session with a client that it would not have resumed with during the initial handshake.

CVE ID	Description
CVE-2025-68146	filelock is a platform-independent file lock for Python. In versions prior to 3.20.1, a Time-of-Check-Time-of-Use (TOCTOU) race condition allows local attackers to corrupt or truncate arbitrary user files through symlink attacks. The vulnerability exists in both Unix and Windows lock file creation where filelock checks if a file exists before opening it with O_TRUNC. An attacker can create a symlink pointing to a victim file in the time gap between the check and open, causing os.open() to follow the symlink and truncate the target file. All users of filelock on Unix, Linux, macOS, and Windows systems are impacted. The vulnerability cascades to dependent libraries. The attack requires local filesystem access and ability to create symlinks (standard user permissions on Unix; Developer Mode on Windows 10+). Exploitation succeeds within 1-3 attempts when lock file paths are predictable. The issue is fixed in version 3.20.1. If immediate upgrade is not possible, use SoftFileLock instead of UnixFileLock/WindowsFileLock (note: different locking semantics, may not be suitable for all use cases); ensure lock file directories have restrictive permissions (chmod 0700) to prevent untrusted users from creating symlinks; and/or monitor lock file directories for suspicious symlinks before running trusted applications. These workarounds provide only partial mitigation. The race condition remains exploitable. Upgrading to version 3.20.1 is strongly recommended.
CVE-2025-68280	Improper Restriction of XML External Entity Reference vulnerability in Apache SIS. It is possible to write XML files in such a way that, when parsed by Apache SIS, an XML file reveals to the attacker the content of a local file on the server running Apache SIS. This vulnerability impacts the following SIS services: * Reading of GeoTIFF files having the GEO_METADATA tag defined by the Defense Geospatial Information Working Group (DGIWG). * Parsing of ISO 19115 metadata in XML format. * Parsing of Coordinate Reference Systems defined in the GML format. * Parsing of files in GPS Exchange Format (GPX). This issue affects Apache SIS from versions 0.4 through 1.5 inclusive. Users are recommended to upgrade to version 1.6, which will fix the issue. In the meantime, the security vulnerability can be avoided by launching Java with the javax.xml.accessExternalDTD system property sets to a comma-separated list of authorized protocols. For example: java -Djavax.xml.accessExternalDTD="" ...
CVE-2025-69534	Python-Markdown version 3.8 contain a vulnerability where malformed HTML-like sequences can cause html.parser.HTMLParser to raise an unhandled AssertionError during Markdown parsing. Because Python-Markdown does not catch this exception, any application that processes attacker-controlled Markdown may crash. This enables remote, unauthenticated Denial of Service in web applications, documentation systems, CI/CD pipelines, and any service that renders untrusted Markdown. The issue was acknowledged by the vendor and fixed in version 3.8.1. This issue causes a remote Denial of Service in any application parsing untrusted Markdown, and can lead to Information Disclosure through uncaught exceptions.
CVE-2025-71176	pytest through 9.0.2 on UNIX relies on directories with the /tmp/pytest-of-{user} name pattern, which allows local users to cause a denial of service or possibly gain privileges.
CVE-2026-0540	DOMPurify 3.1.3 through 3.3.1 and 2.5.3 through 2.5.8, fixed in commit 729097f, contain a cross-site scripting vulnerability that allows attackers to bypass attribute sanitization by exploiting five missing rawtext elements (noscript, xmp, noembed, noframes, iframe) in the SAFE_FOR_XML regex. Attackers can include payloads like </noscript><img src=x onerror=alert(1)> in attribute values to execute JavaScript when sanitized output is placed inside these unprotected rawtext contexts.
CVE-2026-0636	Improper neutralization of special elements used in an LDAP query ('LDAP injection') vulnerability in Legion of the Bouncy Castle Inc. BC-JAVA bcprov on all (prov modules). This vulnerability is associated with program files LDAPStoreHelper. This issue affects BC-JAVA: from 1.74 before 1.80.2, from 1.81 before 1.81.1, from 1.82 before 1.84.
CVE-2026-0994	A denial-of-service (DoS) vulnerability exists in google.protobuf.json_format.ParseDict() in Python, where the max_recursion_depth limit can be bypassed when parsing nested google.protobuf.Any messages. Due to missing recursion depth accounting inside the internal Any-handling logic, an attacker can supply deeply nested Any structures that bypass the intended recursion limit, eventually exhausting Python’s recursion stack and causing a RecursionError.
CVE-2026-1002	The Vert.x Web static handler component cache can be manipulated to deny the access to static files served by the handler using specifically crafted request URI. The issue comes from an improper implementation of the C. rule of section 5.2.4 of RFC3986 and is fixed in Vert.x Core component (used by Vert.x Web): https://github.com/eclipse-vertx/vert.x/pull/5895 Steps to reproduce Given a file served by the static handler, craft an URI that introduces a string like bar%2F..%2F after the last / char to deny the access to the URI with an HTTP 404 response. For example https://example.com/foo/index.html can be denied with https://example.com/foo/bar%2F..%2Findex.html Mitgation Disabling Static Handler cache fixes the issue. StaticHandler staticHandler = StaticHandler.create().setCachingEnabled(false);
CVE-2026-1225	ACE vulnerability in configuration file processing by QOS.CH logback-core up to and including version 1.5.24 in Java applications, allows an attacker to instantiate classes already present on the class path by compromising an existing logback configuration file. The instantiation of a potentially malicious Java class requires that said class is present on the user's class-path. In addition, the attacker must have write access to a configuration file. However, after successful instantiation, the instance is very likely to be discarded with no further ado.

CVE ID	Description
CVE-2026-1703	When pip is installing and extracting a maliciously crafted wheel archive, files may be extracted outside the installation directory. The path traversal is limited to prefixes of the installation directory, thus isn't able to inject or overwrite executable files in typical situations.
CVE-2026-1839	A vulnerability in the HuggingFace Transformers library, specifically in the `Trainer` class, allows for arbitrary code execution. The `_load_rng_state()` method in `src/transformers/trainer.py` at line 3059 calls `torch.load()` without the `weights_only=True` parameter. This issue affects all versions of the library supporting `torch>=2.2` when used with PyTorch versions below 2.6, as the `safe_globals()` context manager provides no protection in these versions. An attacker can exploit this vulnerability by supplying a malicious checkpoint file, such as `rng_state.pth`, which can execute arbitrary code when loaded. The issue is resolved in version v5.0.0rc3.
CVE-2026-2303	The mongo-go-driver repository contains CGo bindings for GSSAPI (Kerberos) authentication on Linux and macOS. The C wrapper implementation contains a heap out-of-bounds read vulnerability due to incorrect assumptions about string termination in the GSSAPI standard. Since GSSAPI buffers are not guaranteed to be null-terminated or have extra padding, this results in reading one byte past the allocated heap buffer.
CVE-2026-2332	In Eclipse Jetty, the HTTP/1.1 parser is vulnerable to request smuggling when chunk extensions are used, similar to the "funky chunks" techniques outlined here: * https://w4ke.info/2025/06/18/funky-chunks.html * https://w4ke.info/2025/10/29/funky-chunks-2.html Jetty terminates chunk extension parsing at \r\n inside quoted strings instead of treating this as an error. POST / HTTP/1.1 Host: localhost Transfer-Encoding: chunked 1;ext="val X 0 GET /smuggled HTTP/1.1 ... Note how the chunk extension does not close the double quotes, and it is able to inject a smuggled request.
CVE-2026-3219	pip handles concatenated tar and ZIP files as ZIP files regardless of filename or whether a file is both a tar and ZIP file. This behavior could result in confusing installation behavior, such as installing "incorrect" files according to the filename of the archive. New behavior only proceeds with installation if the file identifies uniquely as a ZIP or tar archive, not as both.
CVE-2026-4372	A critical remote code execution vulnerability exists in all versions of the HuggingFace transformers library prior to version 5.3.0. The vulnerability allows an attacker to craft a malicious `config.json` file containing the `_attn_implementation_internal` field set to an attacker-controlled HuggingFace Hub repository ID. When a victim loads this model using the standard `AutoModelForCausalLM.from_pretrained()` API, the library downloads and executes arbitrary Python code from the attacker's repository with the victim's full OS privileges. This issue arises due to unfiltered deserialization of configuration attributes, insufficient sanitization of internal fields, and unsandboxed execution of downloaded kernels. The vulnerability bypasses the `trust_remote_code` security mechanism, is invisible to the victim, and exploits the standard documented usage pattern, making it particularly severe. Users are advised to upgrade to version 5.3.0 or later to mitigate this issue.
CVE-2026-4539	A security flaw has been discovered in pygments up to 2.19.2. The impacted element is the function AdlLexer of the file pygments/lexers/archetype.py. The manipulation results in inefficient regular expression complexity. The attack is only possible with local access. The exploit has been released to the public and may be used for attacks. The project was informed of the problem early through an issue report but has not responded yet.
CVE-2026-5241	A vulnerability in the LightGlue model loading path of huggingface/transformers version 5.2.0 allows an attacker-controlled model repository to execute arbitrary code during model initialization. The issue arises because the `trust_remote_code` parameter, intended to prevent remote code execution, is overridden by untrusted serialized configuration data in a nested code path. Specifically, when loading a LightGlue model using `AutoModel.from_pretrained()` with `trust_remote_code=False`, the `LightGlueConfig` reads the `trust_remote_code` value from the untrusted `config.json` file and propagates it into nested `AutoConfig.from_pretrained()` calls. This results in the execution of attacker-provided Python modules, even when the victim explicitly disables remote code execution. The vulnerability poses a high risk for environments such as API inference servers, research notebooks, CI/CD pipelines, and model evaluation workers, potentially leading to credential theft, lateral movement, or persistence/backdoor deployment.
CVE-2026-5588	Use of a Broken or Risky Cryptographic Algorithm vulnerability in Legion of the Bouncy Castle Inc. BC-JAVA bcpkix on all (pkix modules), Legion of the Bouncy Castle Inc. BCPKIX-FIPS bcpkix on All (pkix modules), Legion of the Bouncy Castle Inc. BCPIX-LTS bcpkix on All (pkix modules). This vulnerability is associated with program files JcaContentVerifierProviderBuilder.Java, JcaContentVerfierProviderBuilder.Java. This issue affects BC-JAVA: from 1.67 before 1.80.2, from 1.81 before 1.81.1, from 1.82 before 1.84; BCPKIX-FIPS: from 2.0.6 before 2.0.11, from 2.1.7 before 2.1.11; BCPIX-LTS: from 2.73.7 before 2.73.11.
CVE-2026-5598	Covert timing channel vulnerability in Legion of the Bouncy Castle Inc. BC-JAVA core on all (core modules). This vulnerability is associated with program files FrodoEngine.Java. This issue affects BC-JAVA: from 1.71 before 1.80.2, from 1.81 before 1.81.1, from 1.82 before 1.84.

CVE ID	Description
CVE-2026-5766	An issue was discovered in 6.0 before 6.0.5 and 5.2 before 5.2.14. ASGI requests with a missing or understated `Content-Length` header can bypass the `FILE_UPLOAD_MAX_MEMORY_SIZE` limit, potentially loading large files into memory and causing service degradation. As a reminder, Django expects a limit to be configured at the web server level rather than solely relying on `FILE_UPLOAD_MAX_MEMORY_SIZE`. Earlier, unsupported Django series (such as 5.0.x, 4.1.x, and 3.2.x) were not evaluated and may also be affected. Django would like to thank Kyle Agronick for reporting this issue.
CVE-2026-6357	pip prior to version 26.1 would run self-update check functionality after installing wheel files which required importing well-known Python modules names. These module imports were intentionally deferred to increase startup time of the pip CLI. The patch changes self-update functionality to run before wheels are installed to prevent newly-installed modules from being imported shortly after the installation of a wheel package. Users should still review package contents prior to installation.
CVE-2026-6790	In Eclipse Jetty, for HTTP/1, HTTP/2 and HTTP/3 requests, there is no strict check that the request authority (host and port) matches what provided in the Host header (if present). This was not enforced in earlier HTTP RFC (for example, in RFC 2616), but it is in the latest RFC (9110 and 9112). This mismatch can cause a number of problems that may be classified as vulnerabilities such as: * URI constructions (for example, for redirects -- this is typical for login pages) * Virtual host selection * Reverse proxying * Misleading logs * Etc. Given that the latest RFCs require that request authority and Host header must match, Jetty should enforce this invariant.
CVE-2026-6860	A TCP client can perform a TLS handshake and present the server name extension with a server name that is accepted by a server wildcard name, e.g. if the server is configured with a certificate accepting *.example.com, any XYZ.example.com where xyz is a valid name can be used.
CVE-2026-6873	An issue was discovered in Django 6.0 before 6.0.6 and 5.2 before 5.2.15. `django.http.HttpRequest.get_signed_cookie` in Django uses a non-injective salt derivation (concatenating the cookie name and salt argument), which allows a remote attacker to use a cookie in a context different from the one where it was signed, via distinct `(name, salt)` pairs that produce the same concatenation. Earlier, unsupported Django series (such as 5.0.x, 4.1.x, and 3.2.x) were not evaluated and may also be affected. Django would like to thank Peng Zhou for reporting this issue.
CVE-2026-6907	An issue was discovered in 6.0 before 6.0.5 and 5.2 before 5.2.14. `django.middleware.cache.UpdateCacheMiddleware` erroneously caches requests where the `Vary` header contained an asterisk (`'*'`). This can lead to private data being stored and served. Earlier, unsupported Django series (such as 5.0.x, 4.1.x, and 3.2.x) were not evaluated and may also be affected. Django would like to thank Ahmad Sadeddin for reporting this issue.
CVE-2026-7666	An issue was discovered in Django 6.0 before 6.0.6 and 5.2 before 5.2.15. `django.core.mail.backends.smtp.EmailBackend` in Django fails to prevent reuse of a partially-initialized connection after a failed `STARTTLS` handshake when `fail_silently=True`, which allows on-path network attackers to read email content via cleartext interception. Earlier, unsupported Django series (such as 5.0.x, 4.1.x, and 3.2.x) were not evaluated and may also be affected. Django would like to thank Kasper Dupont for reporting this issue.
CVE-2026-8384	In Eclipse Jetty, an HTTP URI of this form: /public;/../admin/secret.txt results in an unresolved path of: /public/../admin/secret.txt instead of the expected: /admin/secret.txt Jetty itself is not affected, as it will not serve the secret.txt file because it will not pass the alias checker (only resolved resources are served). However, web applications that rely on resolved paths being provided by Jetty may be confused when receiving an unresolved path.
CVE-2026-8404	An issue was discovered in Django 5.2 before 5.2.15 and 6.0 before 6.0.6. `django.middleware.cache.UpdateCacheMiddleware` in Django does not match `Cache-Control` response directives case-insensitively, which allows remote attackers to read responses that were incorrectly cached because their `Cache-Control` directives used uppercase or mixed-case values. Earlier, unsupported Django series (such as 5.0.x, 4.1.x, and 3.2.x) were not evaluated and may also be affected. Django would like to thank Ahmed Badawe for reporting this issue.
CVE-2026-8643	pip would treat console_scripts and gui_scripts as paths instead of file names without sanitizing the resolved absolute path to the installation directory, leading to entry points being installed outside the installation directory.
CVE-2026-8763	In Bouncy Castle for Java before 1.85, Name Constraints bypass via trailing dot in rfc822Name and URI. This issue also affects Bouncy Castle for Java LTS before 2.73.12, and Bouncy Castle for Java FIPS (BC-FJA) before bc-fips 1.0.2.7 (1.0.X series), 2.0.2 (2.0.X series) and 2.1.3 (2.1.X series).
CVE-2026-9318	tablib prior to 3.10.0 contains a stored cross-site scripting vulnerability in the HTML export functionality that allows attackers to execute arbitrary JavaScript by embedding malicious payloads in dataset titles, which are interpolated unsanitized into HTML output via the export_book method in the _html.py format handler. Attackers can rename worksheet sheets in imported files such as XLSX, ODS, XLS, or YAML with script payloads that are assigned to the Dataset title attribute and rendered unescaped inside an HTML h3 tag, leading to session hijacking, unauthorized administrative actions, and sensitive data exposure when the output is rendered in a browser.

CVE ID	Description
CVE-2026-9828	Deserialization of untrusted data vulnerability in QOS.CH Sarl logback logback-core (HardenedObjectInputStream (logback-core) modules) allows Object Injection albeit heavily restricted. More precisely, an attacker able to influence serialized data sent to SimpleSocketServer or SimpleSSLSocketServer can instantiate objects from classes in the java.lang and java.util packages that are not explicitly blocked. Although deserialization is heavily restricted by HardenedObjectInputStream and no practical way to achieve remote code execution or significant privilege escalation has been identified, this issue constitutes a bypass of the intended security restrictions. This issue affects logback: through 1.5.32 inclusive.
CVE-2026-9856	A vulnerability in huggingface/transformers versions <=5.8.0.dev0 allows an attacker to perform arbitrary file writes via path traversal. The issue resides in the `save_pretrained()` methods of `PreTrainedTokenizerBase` and `ProcessorMixin`, where keys from the `chat_template` dictionary are used directly as filenames without proper validation. An attacker can exploit this by publishing a malicious Hugging Face Hub repository with a crafted `tokenizer_config.json` file. When a victim downloads and saves the tokenizer or processor, the attacker-controlled keys can escape the intended save directory, enabling arbitrary file writes with attacker-controlled content. This vulnerability affects multiple processors inheriting from `ProcessorMixin`, including Idefics, Florence, Gemma, Phi, and Qwen-VL.
CVE-2026-10050	In Eclipse Jetty, the Digest authentication server-side component uses ISO-8859-1 to encode the password as bytes. This was done because the initial specification for HTTP did not specify explicitly a charset, and it was assumed to be ISO-8859-1 for historical reasons. If the password contains characters that cannot be represented in ISO-8859-1, they are silently replaced by `?`. This happens with passwords that contain Chinese, Cyrillic or Greek characters, for example: `##123` converts to `??123`. An attacker can send a request with a digest `Authorization` header crafted with a password made of only `?` characters; the server would match any password of the same length that contains non-ISO-8859-1 characters. Recent HTTP Digest [RFC-7616](https://datatracker.ietf.org/doc/html/rfc7616) supports a `charset` parameters that defaults to UTF-8 that allows for correct encoding/decoding of passwords.
CVE-2026-10051	In Eclipse Jetty, a first HTTP/1.1 request with trailers causes the server to retain the trailers in subsequent requests performed over the same connection. Subsequent request that do not have trailers report the trailers of the first request. Subsequent request that do have trailers report the union of trailers of the first request and the current request.
CVE-2026-10532	Deserialization of untrusted data vulnerability in QOS.CH Sarl logback logback-core (HardenedObjectInputStream (logback-core) modules) allows Object Injection, albeit heavily restricted. More precisely, an attacker able to influence serialized data sent to SimpleSocketServer or SimpleSSLSocketServer can instantiate Proxy objects. Although deserialization is heavily restricted by HardenedObjectInputStream and no practical way to achieve remote code execution or significant privilege escalation has been identified, this issue constitutes a bypass of the intended security restrictions. This issue affects logback: through 1.5.33 inclusive.
CVE-2026-13346	pip would incorrectly handle doubly-encoded package URLs from indexes allowing for files to be installed to arbitrary locations on disk even when installing wheels. This vulnerability requires downloading or installing a package from a malicious package index to succeed, malicious packages alone are not able to exploit this vulnerability. Note that this vulnerability only materially impacts users running `pip download` with the `--only-binary` option as installing source distributions from an untrusted index is already an unsafe operation that executes code during install time.
CVE-2026-13506	In Bouncy Castle for Java before 1.85, Lazy ASN.1 sequence forcing resets nesting-depth guard. This issue also affects Bouncy Castle for Java LTS before 2.73.12, and Bouncy Castle for Java FIPS (BC-FJA) before bc-fips 1.0.2.7 (1.0.X series), 2.0.2 (2.0.X series) and 2.1.3 (2.1.X series).
CVE-2026-15307	An issue was discovered in Django 5.2 before 5.2.17 and 6.0 before 6.0.8. GeoDjango spatial lookups optimistically parse the right-hand-side value as a raster by passing it to the `django.contrib.gis.gdal.GDALRaster` constructor. Any value used in a spatial lookup against a `GeometryField` or `RasterField` reaches this constructor, including untrusted input, for example a spatial-field filter submitted through the Django admin changelist query string by a staff user with view permission. A `dict`, or a `str` holding its JSON representation, is opened in write mode regardless of the constructor's `write=False` default, allowing a file with an attacker-chosen name and contents to be written through a file-backed GDAL driver. Any other `str` is treated as a datasource, allowing an outbound network request through a GDAL virtual filesystem handler. Writing a file to a location later imported by the application can result in remote code execution. Earlier, unsupported Django series (such as 5.1.x, 5.0.x, and 4.2.x) were not evaluated and may also be affected. Django would like to thank Bence Nagy, localhost-detect, and kimchunbok_ for reporting this issue.
CVE-2026-15830	An issue was discovered in Django 6.0 before 6.0.9 and 5.2 before 5.2.18. GeoDjango's `django.contrib.gis.geos.GEOSGeometry` is subject to a potential denial-of-service when parsing deeply nested `GEOMETRYCOLLECTION` objects supplied as well-known text (WKT), well-known binary (WKB), or hex-encoded WKB, which triggers unbounded recursion and a segmentation fault in the underlying GEOS library. Spatial field lookups and the `django.contrib.gis.forms.GeometryField` form field are also affected. Earlier, unsupported Django series (such as 5.1.x, 5.0.x, and 4.2.x) were not evaluated and may also be affected. Django would like to thank Andrew MacPherson and kimchunbok_ for reporting this issue.

CVE ID	Description
CVE-2026-18401	The non-blocking (asynchronous) JSON parser in jackson-core does not enforce the maxNumberLength constraint defined in StreamReadConstraints (default: 1000 characters). An attacker able to submit JSON to an application that uses the async parser API can supply a number token of arbitrary length, leading to excessive memory allocation and potential CPU exhaustion, resulting in a denial of service. The synchronous parser enforces this limit correctly, so the constraint is applied inconsistently depending on which parsing API the application uses. Root cause: the async parsing path in NonBlockingUtf8JsonParserBase and related classes never invokes the number length validation methods. Number parsing methods such as _finishNumberIntegralPart() accumulate digits into the TextBuffer without any length check, then call _valueComplete() to finalize the token. _valueComplete() does not call resetInt() or resetFloat(), which are the methods in ParserBase where validateIntegerLength() and validateFPLength() are performed. Because that validation step is skipped, maxNumberLength is never enforced on the async code path. Impact: an attacker sending a JSON document containing an arbitrarily long number to an application using the async parser (for example a Spring WebFlux or other reactive application) can cause unbounded allocation in the TextBuffer and an OutOfMemoryError. If the application subsequently calls getBigIntegerValue() or getDecimalValue(), the JVM may additionally be tied up in O(n^2) BigInteger parsing, causing CPU-based denial of service. No privileges or user interaction beyond the ability to submit data for parsing are required. This issue affects com.fasterxml.jackson.core:jackson-core from version 2.15.0 through 2.18.5 and from 2.19.0 through 2.21.0, and tools.jackson.core:jackson-core from 3.0.0 through 3.0.x. Versions prior to 2.15.0 are not affected, because StreamReadConstraints -- which defines the maxNumberLength setting -- was first introduced in jackson-core 2.15.0, so no such constraint exists to be bypassed in earlier releases. Note that GHSA-72hv-8253-57qq records the lower bound of the affected 2.x range as 2.0.0.
CVE-2026-19032	jackson-databind's deserializer for java.nio.file.Path resolves an attacker-supplied URI without restricting the URI scheme. In JDKFromStringDeserializer.NioPathHelper.deserialize, a string bound from untrusted JSON is passed to new URI(value) and then to Path.of(uri). When that throws FileSystemNotFoundException, the code enumerates ServiceLoader<FileSystemProvider> and calls provider.getPath(uri) on the first provider whose scheme matches the attacker-chosen scheme. Untrusted JSON can therefore select and drive an arbitrary registered FileSystemProvider during readValue under a default JsonMapper, and forces provider class loading at the same time. With only the JDK built-in providers (file, jar/zipfs) present, the resolved path is inert and no mount or network I/O occurs; further impact requires a side-effecting third-party FileSystemProvider on the classpath. This affects com.fasterxml.jackson.core:jackson-databind from 2.8.0 before 2.18.10, from 2.19.0 before 2.21.6, and from 2.22.0 before 2.22.2, and tools.jackson.core:jackson-databind from 3.0.0 before 3.1.6 and from 3.2.0 before 3.2.2. Users should upgrade to 2.18.10, 2.21.6, 2.22.2, 3.1.6, or 3.2.2. Binding java.nio.file.Path from untrusted JSON should be avoided regardless of version.
CVE-2026-21441	urllib3 is an HTTP client library for Python. urllib3's streaming API is designed for the efficient handling of large HTTP responses by reading the content in chunks, rather than loading the entire response body into memory at once. urllib3 can perform decoding or decompression based on the HTTP `Content-Encoding` header (e.g., `gzip`, `deflate`, `br`, or `zstd`). When using the streaming API, the library decompresses only the necessary bytes, enabling partial content consumption. Starting in version 1.22 and prior to version 2.6.3, for HTTP redirect responses, the library would read the entire response body to drain the connection and decompress the content unnecessarily. This decompression occurred even before any read methods were called, and configured read limits did not restrict the amount of decompressed data. As a result, there was no safeguard against decompression bombs. A malicious server could exploit this to trigger excessive resource consumption on the client. Applications and libraries are affected when they stream content from untrusted sources by setting `preload_content=False` when they do not disable redirects. Users should upgrade to at least urllib3 v2.6.3, in which the library does not decode content of redirect responses when `preload_content=False`. If upgrading is not immediately possible, disable redirects by setting `redirect=False` for requests to untrusted source.
CVE-2026-22701	filelock is a platform-independent file lock for Python. Prior to version 3.20.3, a TOCTOU race condition vulnerability exists in the SoftFileLock implementation of the filelock package. An attacker with local filesystem access and permission to create symlinks can exploit a race condition between the permission validation and file creation to cause lock operations to fail or behave unexpectedly. The vulnerability occurs in the _acquire() method between raise_on_not_writable_file() (permission check) and os.open() (file creation). During this race window, an attacker can create a symlink at the lock file path, potentially causing the lock to operate on an unintended target file or leading to denial of service. This issue has been patched in version 3.20.3.
CVE-2026-22702	virtualenv is a tool for creating isolated virtual python environments. Prior to version 20.36.1, TOCTOU (Time-of-Check-Time-of-Use) vulnerabilities in virtualenv allow local attackers to perform symlink-based attacks on directory creation operations. An attacker with local access can exploit a race condition between directory existence checks and creation to redirect virtualenv's app_data and lock file operations to attacker-controlled locations. This issue has been patched in version 20.36.1.
CVE-2026-22815	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to version 3.13.4, insufficient restrictions in header/trailer handling could cause uncapped memory usage. This issue has been patched in version 3.13.4.

CVE ID	Description
CVE-2026-23901	Observable Timing Discrepancy vulnerability in Apache Shiro. This issue affects Apache Shiro: from 1.*, 2.* before 2.0.7. Users are recommended to upgrade to version 2.0.7 or later, which fixes the issue. Prior to Shiro 2.0.7, code paths for non-existent vs. existing users are different enough, that a brute-force attack may be able to tell, by timing the requests only, determine if the request failed because of a non-existent user vs. wrong password. The most likely attack vector is a local attack only. Shiro security model https://shiro.apache.org/security-model.html#username_enumeration discusses this as well. Typically, brute force attack can be mitigated at the infrastructure level.
CVE-2026-23949	jaraco.context, an open-source software package that provides some useful decorators and context managers, has a Zip Slip path traversal vulnerability in the `jaraco.context.tarball()` function starting in version 5.2.0 and prior to version 6.1.0. The vulnerability may allow attackers to extract files outside the intended extraction directory when malicious tar archives are processed. The strip_first_component filter splits the path on the first `/` and extracts the second component, while allowing `../` sequences. Paths like `dummy_dir/../../etc/passwd` become `../../etc/passwd`. Note that this suffers from a nested tarball attack as well with multi-level tar files such as `dummy_dir/inner.tar.gz`, where the inner.tar.gz includes a traversal `dummy_dir/../../config/.env` that also gets translated to `../../config/.env`. Version 6.1.0 contains a patch for the issue.
CVE-2026-24049	wheel is a command line tool for manipulating Python wheel files, as defined in PEP 427. In versions 0.40.0 through 0.46.1, the unpack function is vulnerable to file permission modification through mishandling of file permissions after extraction. The logic blindly trusts the filename from the archive header for the chmod operation, even though the extraction process itself might have sanitized the path. Attackers can craft a malicious wheel file that, when unpacked, changes the permissions of critical system files (e.g., /etc/passwd, SSH keys, config files), allowing for Privilege Escalation or arbitrary code execution by modifying now-writable scripts. This issue has been fixed in version 0.46.2.
CVE-2026-24051	OpenTelemetry-Go is the Go implementation of OpenTelemetry. The OpenTelemetry Go SDK in version v1.20.0-1.39.0 is vulnerable to Path Hijacking (Untrusted Search Paths) on macOS/Darwin systems. The resource detection code in sdk/resource/host_id.go executes the ioreg system command using a search path. An attacker with the ability to locally modify the PATH environment variable can achieve Arbitrary Code Execution (ACE) within the context of the application. A fix was released with v1.40.0.
CVE-2026-24281	Hostname verification in Apache ZooKeeper ZKTrustManager falls back to reverse DNS (PTR) when IP SAN validation fails, allowing attackers who control or spoof PTR records to impersonate ZooKeeper servers or clients with a valid certificate for the PTR name. It's important to note that attacker must present a certificate which is trusted by ZKTrustManager which makes the attack vector harder to exploit. Users are recommended to upgrade to version 3.8.6 or 3.9.5, which fixes this issue by introducing a new configuration option to disable reverse DNS lookup in client and quorum protocols.
CVE-2026-25087	Use After Free vulnerability in Apache Arrow C++. This issue affects Apache Arrow C++ from 15.0.0 through 23.0.0. It can be triggered when reading an Arrow IPC file (but not an IPC stream) with pre-buffering enabled, if the IPC file contains data with variadic buffers (such as Binary View and String View data). Depending on the number of variadic buffers in a record batch column and on the temporal sequence of multi-threaded IO, a write to a dangling pointer could occur. The value (a `std::shared_ptr<Buffer>` object) that is written to the dangling pointer is not under direct control of the attacker. Pre-buffering is disabled by default but can be enabled using a specific C++ API call (`RecordBatchFileReader::PreBufferMetadata`). The functionality is not exposed in language bindings (Python, Ruby, C GLib), so these bindings are not vulnerable. The most likely consequence of this issue would be random crashes or memory corruption when reading specific kinds of IPC files. If the application allows ingesting IPC files from untrusted sources, this could plausibly be exploited for denial of service. Inducing more targeted kinds of misbehavior (such as confidential data extraction from the running process) depends on memory allocation and multi-threaded IO temporal patterns that are unlikely to be easily controlled by an attacker. Advice for users of Arrow C++: 1. check whether you enable pre-buffering on the IPC file reader (using `RecordBatchFileReader::PreBufferMetadata`) 2. if so, either disable pre-buffering (which may have adverse performance consequences), or switch to Arrow 23.0.1 which is not vulnerable
CVE-2026-25645	Requests is a HTTP library. Prior to version 2.33.0, the `requests.utils.extract_zipped_paths()` utility function uses a predictable filename when extracting files from zip archives into the system temporary directory. If the target file already exists, it is reused without validation. A local attacker with write access to the temp directory could pre-create a malicious file that would be loaded in place of the legitimate one. Standard usage of the Requests library is not affected by this vulnerability. Only applications that call `extract_zipped_paths()` directly are impacted. Starting in version 2.33.0, the library extracts files to a non-deterministic location. If developers are unable to upgrade, they can set `TMPDIR` in their environment to a directory with restricted write access.
CVE-2026-25679	url.Parse insufficiently validated the host/authority component and accepted some invalid URLs.
CVE-2026-25680	Parsing arbitrary HTML can consume excessive CPU time, possibly leading to denial of service.
CVE-2026-25681	Parsing arbitrary HTML which is then rendered using Render can result in an unexpected HTML tree. This can be leveraged to execute XSS attacks in applications that attempt to sanitize input HTML before rendering.

CVE ID	Description
CVE-2026-26007	cryptography is a package designed to expose cryptographic primitives and recipes to Python developers. Prior to 46.0.5, the public_key_from_numbers (or EllipticCurvePublicNumbers.public_key()), EllipticCurvePublicNumbers.public_key(), load_der_public_key() and load_pem_public_key() functions do not verify that the point belongs to the expected prime-order subgroup of the curve. This missing validation allows an attacker to provide a public key point P from a small-order subgroup. This can lead to security issues in various situations, such as the most commonly used signature verification (ECDSA) and shared key negotiation (ECDH). When the victim computes the shared secret as S = [victim_private_key]P via ECDH, this leaks information about victim_private_key mod (small_subgroup_order). For curves with cofactor > 1, this reveals the least significant bits of the private key. When these weak public keys are used in ECDSA , it's easy to forge signatures on the small subgroup. Only SECT curves are impacted by this. This vulnerability is fixed in 46.0.5.
CVE-2026-27136	Parsing arbitrary HTML which is then rendered using Render can result in an unexpected HTML tree. This can be leveraged to execute XSS attacks in applications that attempt to sanitize input HTML before rendering.
CVE-2026-27139	On Unix platforms, when listing the contents of a directory using File.ReadDir or File.Readdir the returned FileInfo could reference a file outside of the Root in which the File was opened. The impact of this escape is limited to reading metadata provided by lstat from arbitrary locations on the filesystem without permitting reading or writing files outside the root.
CVE-2026-27142	Actions which insert URLs into the content attribute of HTML meta tags are not escaped. This can allow XSS if the meta tag also has an http-equiv attribute with the value "refresh". A new GODEBUG setting has been added, htmlmetacontenturlescape, which can be used to disable escaping URLs in actions in the meta content attribute which follow "url=" by setting htmlmetacontenturlescape=0.
CVE-2026-27145	(*x509.Certificate).VerifyHostname previously called matchHostnames in a loop over all DNS Subject Alternative Name (SAN) entries. This caused strings.Split(host, ".") to execute repeatedly on the same input hostname. With a large DNS SAN list, verification costs scaled quadratically based on the number of SAN entries multiplied by the hostname's label count. Because x509.Verify validates hostnames before building the certificate chain, this overhead occurred even for untrusted certificates.
CVE-2026-27448	pyOpenSSL is a Python wrapper around the OpenSSL library. Starting in version 0.14.0 and prior to version 26.0.0, if a user provided callback to `set_tlsext_servername_callback` raised an unhandled exception, this would result in a connection being accepted. If a user was relying on this callback for any security-sensitive behavior, this could allow bypassing it. Starting in version 26.0.0, unhandled exceptions now result in rejecting the connection.
CVE-2026-27459	pyOpenSSL is a Python wrapper around the OpenSSL library. Starting in version 22.0.0 and prior to version 26.0.0, if a user provided callback to `set_cookie_generate_callback` returned a cookie value greater than 256 bytes, pyOpenSSL would overflow an OpenSSL provided buffer. Starting in version 26.0.0, cookie values that are too long are now rejected.
CVE-2026-27830	c3p0, a JDBC Connection pooling library, is vulnerable to attack via maliciously crafted Java-serialized objects and `javax.naming.Reference` instances. Several c3p0 `ConnectionPoolDataSource` implementations have a property called `userOverridesAsString` which conceptually represents a `Map<String,Map<String,String>>`. Prior to v0.12.0, that property was maintained as a hex-encoded serialized object. Any attacker able to reset this property, on an existing `ConnectionPoolDataSource` or via maliciously crafted serialized objects or `javax.naming.Reference` instances could be tailored execute unexpected code on the application's `CLASSPATH`. The danger of this vulnerability was strongly magnified by vulnerabilities in c3p0's main dependency, mchange-commons-java. This library includes code that mirrors early implementations of JNDI functionality, including ungated support for remote `factoryClassLocation` values. Attackers could set c3p0's `userOverridesAsString` hex-encoded serialized objects that include objects "indirectly serialized" via JNDI references. Deserialization of those objects and dereferencing of the embedded `javax.naming.Reference` objects could provoke download and execution of malicious code from a remote `factoryClassLocation`. Although hazard presented by c3p0's vulnerabilites are exarcerbated by vulnerabilities in mchange-commons-java, use of Java-serialized-object hex as the format for a writable Java-Bean property, of objects that may be exposed across JNDI interfaces, represents a serious independent fragility. The `userOverridesAsString` property of c3p0 `ConnectionPoolDataSource` classes has been reimplemented to use a safe CSV-based format, rather than rely upon potentially dangerous Java object deserialization. c3p0-0.12.0+ and above depend upon mchange-commons-java 0.4.0+, which gates support for remote `factoryClassLocation` values by configuration parameters that default to restrictive values. c3p0 additionally enforces the new mchange-commons-java `com.mchange.v2.naming.nameGuardClassName` to prevent injection of unexpected, potentially remote JNDI names. There is no supported workaround for versions of c3p0 prior to 0.12.0.
CVE-2026-28208	Junrar is an open source java RAR archive library. Prior to version 7.5.8, a backslash path traversal vulnerability in `LocalFolderExtractor` allows an attacker to write arbitrary files with attacker-controlled content anywhere on the filesystem when a crafted RAR archive is extracted on Linux/Unix. This can often lead to remote code execution (e.g., overwriting shell profiles, source code, cron jobs, etc). Version 7.5.8 has a fix for the issue.
CVE-2026-32280	During chain building, the amount of work that is done is not correctly limited when a large number of intermediate certificates are passed in VerifyOptions.Intermediates, which can lead to a denial of service. This affects both direct users of crypto/x509 and users of crypto/tls.

CVE ID	Description
CVE-2026-32281	Validating certificate chains which use policies is unexpectedly inefficient when certificates in the chain contain a very large number of policy mappings, possibly causing denial of service. This only affects validation of otherwise trusted certificate chains, issued by a root CA in the VerifyOptions.Roots CertPool, or in the system certificate pool.
CVE-2026-32282	On Linux, if the target of Root.Chmod is replaced with a symlink while the chmod operation is in progress, Chmod can operate on the target of the symlink, even when the target lies outside the root. The Linux fchmodat syscall silently ignores the AT_SYMLINK_NOFOLLOW flag, which Root.Chmod uses to avoid symlink traversal. Root.Chmod checks its target before acting and returns an error if the target is a symlink lying outside the root, so the impact is limited to cases where the target is replaced with a symlink between the check and operation.
CVE-2026-32283	If one side of the TLS connection sends multiple key update messages post-handshake in a single record, the connection can deadlock, causing uncontrolled consumption of resources. This can lead to a denial of service. This only affects TLS 1.3.
CVE-2026-32288	tar.Reader can allocate an unbounded amount of memory when reading a maliciously-crafted archive containing a large number of sparse regions encoded in the "old GNU sparse map" format.
CVE-2026-32289	Context was not properly tracked across template branches for JS template literals, leading to possibly incorrect escaping of content when branches were used. Additionally template actions within JS template literals did not properly track the brace depth, leading to incorrect escaping being applied. These issues could cause actions within JS template literals to be incorrectly or improperly escaped, leading to XSS vulnerabilities.
CVE-2026-33186	gRPC-Go is the Go language implementation of gRPC. Versions prior to 1.79.3 have an authorization bypass resulting from improper input validation of the HTTP/2 `:path` pseudo-header. The gRPC-Go server was too lenient in its routing logic, accepting requests where the `:path` omitted the mandatory leading slash (e.g., `Service/Method` instead of `/Service/Method`). While the server successfully routed these requests to the correct handler, authorization interceptors (including the official `grpc/authz` package) evaluated the raw, non-canonical path string. Consequently, "deny" rules defined using canonical paths (starting with `/`) failed to match the incoming request, allowing it to bypass the policy if a fallback "allow" rule was present. This affects gRPC-Go servers that use path-based authorization interceptors, such as the official RBAC implementation in `google.golang.org/grpc/authz` or custom interceptors relying on `info.FullMethod` or `grpc.Method(ctx)`; AND that have a security policy contains specific "deny" rules for canonical paths but allows other requests by default (a fallback "allow" rule). The vulnerability is exploitable by an attacker who can send raw HTTP/2 frames with malformed `:path` headers directly to the gRPC server. The fix in version 1.79.3 ensures that any request with a `:path` that does not start with a leading slash is immediately rejected with a `codes.Unimplemented` error, preventing it from reaching authorization interceptors or handlers with a non-canonical path string. While upgrading is the most secure and recommended path, users can mitigate the vulnerability using one of the following methods: Use a validating interceptor (recommended mitigation); infrastructure-level normalization; and/or policy hardening.
CVE-2026-33811	When using LookupCNAME with the cgo DNS resolver, a very long CNAME response can trigger a double-free of C memory and a crash.
CVE-2026-33814	When processing HTTP/2 SETTINGS frames, transport will enter an infinite loop of writing CONTINUATION frames if it receives a SETTINGS_MAX_FRAME_SIZE with a value of 0.
CVE-2026-33818	Enforce a recursion limit in Unmarshal to prevent stack exhaustion when parsing deeply-nested, recursive structures.
CVE-2026-33870	Netty is an asynchronous, event-driven network application framework. In versions prior to 4.1.132.Final and 4.2.10.Final, Netty incorrectly parses quoted strings in HTTP/1.1 chunked transfer encoding extension values, enabling request smuggling attacks. Versions 4.1.132.Final and 4.2.10.Final fix the issue.
CVE-2026-33871	Netty is an asynchronous, event-driven network application framework. In versions prior to 4.1.132.Final and 4.2.10.Final, a remote user can trigger a Denial of Service (DoS) against a Netty HTTP/2 server by sending a flood of `CONTINUATION` frames. The server's lack of a limit on the number of `CONTINUATION` frames, combined with a bypass of existing size-based mitigations using zero-byte frames, allows an user to cause excessive CPU consumption with minimal bandwidth, rendering the server unresponsive. Versions 4.1.132.Final and 4.2.10.Final fix the issue.
CVE-2026-33997	Moby is an open source container framework. Prior to version 29.3.1, a security vulnerability has been detected that allows plugins privilege validation to be bypassed during docker plugin install. Due to an error in the daemon's privilege comparison logic, the daemon may incorrectly accept a privilege set that differs from the one approved by the user. Plugins that request exactly one privilege are also affected, because no comparison is performed at all. This issue has been patched in version 29.3.1.
CVE-2026-34073	cryptography is a package designed to expose cryptographic primitives and recipes to Python developers. Prior to version 46.0.6, DNS name constraints were only validated against SANs within child certificates, and not the "peer name" presented during each validation. Consequently, cryptography would allow a peer named bar.example.com to validate against a wildcard leaf certificate for *.example.com, even if the leaf's parent certificate (or upwards) contained an excluded subtree constraint for bar.example.com. This issue has been patched in version 46.0.6.

CVE ID	Description
CVE-2026-34479	The Log4j1XmlLayout from the Apache Log4j 1-to-Log4j 2 bridge fails to escape characters forbidden by the XML 1.0 standard, producing malformed XML output. Conforming XML parsers are required to reject documents containing such characters with a fatal error, which may cause downstream log processing systems to drop or fail to index affected records. Two groups of users are affected: * Those using Log4j1XmlLayout directly in a Log4j Core 2 configuration file. * Those using the Log4j 1 configuration compatibility layer with org.apache.log4j.xml.XMLLayout specified as the layout class. Users are advised to upgrade to Apache Log4j 1-to-Log4j 2 bridge version 2.25.4, which corrects this issue. Note: The Apache Log4j 1-to-Log4j 2 bridge is deprecated and will not be present in Log4j 3. Users are encouraged to consult the Log4j 1 to Log4j 2 migration guide https://logging.apache.org/log4j/2.x/migrate-from-log4j1.html , and specifically the section on eliminating reliance on the bridge.
CVE-2026-34513	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to version 3.13.4, an unbounded DNS cache could result in excessive memory usage possibly resulting in a DoS situation. This issue has been patched in version 3.13.4.
CVE-2026-34514	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to version 3.13.4, an attacker who controls the content_type parameter in aiohttp could use this to inject extra headers or similar exploits. This issue has been patched in version 3.13.4.
CVE-2026-34515	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to version 3.13.4, on Windows the static resource handler may expose information about a NTLMv2 remote path. This issue has been patched in version 3.13.4.
CVE-2026-34516	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to version 3.13.4, a response with an excessive number of multipart headers may be allowed to use more memory than intended, potentially allowing a DoS vulnerability. This issue has been patched in version 3.13.4.
CVE-2026-34517	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to version 3.13.4, for some multipart form fields, aiohttp read the entire field into memory before checking client_max_size. This issue has been patched in version 3.13.4.
CVE-2026-34518	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to version 3.13.4, when following redirects to a different origin, aiohttp drops the Authorization header, but retains the Cookie and Proxy-Authorization headers. This issue has been patched in version 3.13.4.
CVE-2026-34519	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to version 3.13.4, an attacker who controls the reason parameter when creating a Response may be able to inject extra headers or similar exploits. This issue has been patched in version 3.13.4.
CVE-2026-34520	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to version 3.13.4, the C parser (the default for most installs) accepted null bytes and control characters in response headers. This issue has been patched in version 3.13.4.
CVE-2026-34525	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to version 3.13.4, multiple Host headers were allowed in aiohttp. This issue has been patched in version 3.13.4.
CVE-2026-34993	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to version 3.14.0, using ``CookieJar.load()`` with untrusted input may allow arbitrary code execution. Most applications using this function will be doing so with the user's own data, so this is unlikely to affect many applications. Version 3.14.0 patches the issue. If an application does allow attacker controlled files to be loaded, a workaround on older releases would be to sanitize the files before loading.
CVE-2026-35192	An issue was discovered in 6.0 before 6.0.5 and 5.2 before 5.2.14. Response headers do not vary on cookies if a session is not modified, but `SESSION_SAVE_EVERY_REQUEST` is `True`. A remote attacker can steal a user's session after that user visits a cached public page. Earlier, unsupported Django series (such as 5.0.x, 4.1.x, and 3.2.x) were not evaluated and may also be affected. Django would like to thank Cantina for reporting this issue.
CVE-2026-35193	An issue was discovered in Django 5.2 before 5.2.15 and 6.0 before 6.0.6. `django.middleware.cache.UpdateCacheMiddleware` in Django does not add `Authorization` to the `Vary` response header for requests bearing that header without `Cache-Control: public`, which allows remote attackers to read private cached responses via unauthenticated requests to the same URL. Earlier, unsupported Django series (such as 5.0.x, 4.1.x, and 3.2.x) were not evaluated and may also be affected. Django would like to thank Shai Berger for reporting this issue.

CVE ID	Description
CVE-2026-35563	It was identified that the LDAP client implementation in version 2.1.7 does not verify if the server certificate matches the intended LDAP hostname. While the underlying code validates the certificate chain against a trusted authority, the absence of endpoint identification allows a valid certificate issued for an entirely unrelated host to be improperly accepted. This oversight leaves the connection highly vulnerable to server impersonation and complete connection compromise. The root cause of this vulnerability lies in the incomplete TLS server identity verification within the LDAP client implementation. The attacker requires MITM capability on the network to exploit this vulnerability. This attacker must be able to present a certificate trusted by the client's configured trust store. The hostname verification has been enforced in the new version of the LDAP API
CVE-2026-39820	Well-crafted inputs reaching ParseAddress, ParseAddressList, and ParseDate were able to trigger excessive CPU exhaustion and memory allocations.
CVE-2026-39821	The ToASCII and ToUnicode functions incorrectly accept Punycode-encoded labels that decode to an ASCII-only label. For example, ToUnicode("xn--example-.com") incorrectly returns the name "example.com" rather than an error. This behavior can lead to privilege escalation in programs using the idna package. For example, a program which performs privilege checks on the ASCII hostname may reject "example.com" but permit "xn--example-.com". If that program subsequently converts the ASCII hostname to Unicode, it will inadvertently permits access to the Unicode name "example.com".
CVE-2026-39822	On Unix systems, opening a file in an os.Root improperly follows symlinks to locations outside of the Root when the final path component of the a path is a symbolic link and the path ends in /. For example, 'root.Open("symlink/")' will open "symlink" even when "symlink" is a symbolic link pointing outside of the root.
CVE-2026-39823	CVE-2026-27142 fixed a vulnerability in which URLs were not correctly escaped inside of a <meta> tag's <content> attribute. If the URL content were to insert ASCII whitespaces around the '=' rune inside of the <content> attribute, the escaper would fail to similarly escape it, leading to XSS.
CVE-2026-39824	NewNTUnicodeString does not check for string length overflow. When provided with a string that overflows the maximum size of a NTUnicodeString (a 16-bit number of bytes), it returns a truncated string rather than an error.
CVE-2026-39825	ReverseProxy can forward queries containing parameters not visible to Rewrite functions. When used with a Rewrite function, or a Director function which parses query parameters, ReverseProxy sanitizes the forwarded request to remove query parameters which are not parsed by url.ParseQuery. ReverseProxy does not take ParseQuery's limit on the total number of query parameters (controlled by GODEBUG=urlmaxqueryparams=N) into account. This can permit ReverseProxy to forward a request containing a query parameter that is not visible to the Rewrite function. For example, the query "a1=x&a2=x&...&a10000=x&hidden=y" can forward the parameter "hidden=y" while hiding it from the proxy's Rewrite function.
CVE-2026-39826	If a trusted template author were to write a <script> tag containing an empty 'type' attribute or a 'type' attribute with an ASCII whitespace, the execution of the template would incorrectly escape any data passed into the <script> block.
CVE-2026-39827	An authenticated SSH client that repeatedly opened channels which were rejected by the server caused unbounded memory growth, eventually crashing the server process and affecting all connected users. Rejected channels are now properly removed from the connection's internal state and released for garbage collection.
CVE-2026-39828	When an SSH server authentication callback returned PartialSuccessError with non-nil Permissions, those permissions were silently discarded, potentially dropping certificate restrictions such as force-command after a second factor succeeded. Returning non-nil Permissions with PartialSuccessError now results in a connection error.
CVE-2026-39829	The RSA and DSA public key parsers did not enforce size limits on key parameters. A crafted public key with an excessively large modulus or DSA parameter could cause several minutes of CPU consumption during signature verification. This could be triggered by unauthenticated clients during public key authentication. RSA moduli are now limited to 8192 bits, and DSA parameters are validated per FIPS 186-2.
CVE-2026-39830	A malicious SSH peer could send unsolicited global request responses to fill an internal buffer, blocking the connection's read loop. The blocked goroutine could not be released by calling Close(), resulting in a resource leak per connection. Unsolicited global responses are now discarded.
CVE-2026-39831	The Verify() method for FIDO/U2F security key types (sk-ecdsa-sha2-nistp256@openssh.com, sk-ssh-ed25519@openssh.com) did not check the User Presence flag. Signatures generated without physical touch were accepted, allowing unattended use of a hardware security key. To restore the previous behavior, return a "no-touch-required" extension in Permissions.Extensions from PublicKeyCallback.
CVE-2026-39832	When adding a key to a remote agent constraint extensions such as restrict-destination-v00@openssh.com were not serialized in the request. Destination restrictions were silently stripped when forwarding keys, allowing unrestricted use of the key on the remote host. The client now serializes all constraint extensions. Additionally, the in-memory keyring returned by NewKeyring() now rejects keys with unsupported constraint extensions instead of silently ignoring them.

CVE ID	Description
CVE-2026-39833	The in-memory keyring returned by NewKeyring() silently accepted keys with the ConfirmBeforeUse constraint but never enforced it. The key would sign without any confirmation prompt, with no indication to the caller that the constraint was not in effect. NewKeyring() now returns an error when unsupported constraints are requested.
CVE-2026-39834	When writing data larger than 4GB in a single Write call on an SSH channel, an integer overflow in the internal payload size calculation caused the write loop to spin indefinitely, sending empty packets without making progress. The size comparison now uses int64 to prevent truncation.
CVE-2026-39835	SSH servers which use CertChecker as a public key callback without setting IsUserAuthority or IsHostAuthority could be caused to panic by a client presenting a certificate. CertChecker now returns an error instead of panicking when these callbacks are nil.
CVE-2026-39836	The Dial and LookupPort functions panic on Windows when provided with an input containing a NUL (0).
CVE-2026-39882	OpenTelemetry-Go is the Go implementation of OpenTelemetry. Prior to 1.43.0, the otlp HTTP exporters (traces/metrics/logs) read the full HTTP response body into an in-memory bytes.Buffer without a size cap. This is exploitable for memory exhaustion when the configured collector endpoint is attacker-controlled (or a network attacker can mitm the exporter connection). This vulnerability is fixed in 1.43.0.
CVE-2026-39883	OpenTelemetry-Go is the Go implementation of OpenTelemetry. From 1.15.0 to 1.42.0, the fix for CVE-2026-24051 changed the Darwin ioreg command to use an absolute path but left the BSD kenv command using a bare name, allowing the same PATH hijacking attack on BSD and Solaris platforms. This vulnerability is fixed in 1.43.0.
CVE-2026-40192	Pillow is a Python imaging library. Versions 10.3.0 through 12.1.1 did not limit the amount of GZIP-compressed data read when decoding a FITS image, making them vulnerable to decompression bomb attacks. A specially crafted FITS file could cause unbounded memory consumption, leading to denial of service (OOM crash or severe performance degradation). If users are unable to immediately upgrade, they should only open specific image formats, excluding FITS, as a workaround.
CVE-2026-40458	PAC4J is vulnerable to Cross-Site Request Forgery (CSRF). A malicious attacker can craft a specially designed website which, when visited by a user, will automatically submit a forged cross-site request with a token whose hash collides with the victim's legitimate CSRF token. Importantly, the attacker does not need to know the victim's CSRF token or its hash prior to the attack. Collisions in the deterministic String.hashCode() function can be computed directly, reducing the effective token's security space to 32 bits. This bypasses CSRF protection, allowing profile updates, password changes, account linking, and any other state-changing operations to be performed without the victim's consent. This issue was fixed in PAC4J versions 5.7.10 and 6.4.1
CVE-2026-40490	The AsyncHttpClient (AHC) library allows Java applications to easily execute HTTP requests and asynchronously process HTTP responses. When redirect following is enabled (followRedirect(true)), versions of AsyncHttpClient prior to 3.0.9 and 2.14.5 forward Authorization and Proxy-Authorization headers along with Realm credentials to arbitrary redirect targets regardless of domain, scheme, or port changes. This leaks credentials on cross-domain redirects and HTTPS-to-HTTP downgrades. Additionally, even when stripAuthorizationOnRedirect is set to true, the Realm object containing plaintext credentials is still propagated to the redirect request, causing credential re-generation for Basic and Digest authentication schemes via NettyRequestFactory. An attacker who controls a redirect target (via open redirect, DNS rebinding, or MITM on HTTP) can capture Bearer tokens, Basic auth credentials, or any other Authorization header value. The fix in versions 3.0.9 and 2.14.5 automatically strips Authorization and Proxy-Authorization headers and clears Realm credentials whenever a redirect crosses origin boundaries (different scheme, host, or port) or downgrades from HTTPS to HTTP. For users unable to upgrade, set `(stripAuthorizationOnRedirect(true))` in the client config and avoid using Realm-based authentication with redirect following enabled. Note that `(stripAuthorizationOnRedirect(true))` alone is insufficient on versions prior to 3.0.9 and 2.14.5 because the Realm bypass still re-generates credentials. Alternatively, disable redirect following (`followRedirect(false)`) and handle redirects manually with origin validation.
CVE-2026-40984	In Micrometer, it is possible for a user to provide specially crafted HTTP requests that may cause a denial-of-service (DoS) condition. Affected versions: micrometer-core 1.16.0 through 1.16.5; 1.15.0 through 1.15.11; 1.14.0 through 1.14.15; 1.13.0 through 1.13.18; 1.9.0 through 1.9.17. micrometer-jetty11 1.16.0 through 1.16.5; 1.15.0 through 1.15.11; 1.14.0 through 1.14.15; 1.13.0 through 1.13.18. micrometer-jetty12 1.16.0 through 1.16.5; 1.15.0 through 1.15.11; 1.14.0 through 1.14.15; 1.13.0 through 1.13.18.
CVE-2026-41066	lxml is a library for processing XML and HTML in the Python language. Prior to 6.1.0, using either of the two parsers in the default configuration (with resolve_entities=True) allows untrusted XML input to read local files. Setting the resolve_entities option explicitly to resolve_entities='internal' or resolve_entities=False disables the local file access. This vulnerability is fixed in 6.1.0.
CVE-2026-41205	Mako is a template library written in Python. Prior to 1.3.11, TemplateLookup.get_template() is vulnerable to path traversal when a URI starts with // (e.g., //../../../secret.txt). The root cause is an inconsistency between two slash-stripping implementations. Any file readable by the process can be returned as rendered template content when an application passes untrusted input directly to TemplateLookup.get_template(). This vulnerability is fixed in 1.3.11.

CVE ID	Description
CVE-2026-41238	DOMPurify is a DOM-only cross-site scripting sanitizer for HTML, MathML, and SVG. Versions 3.0.1 through 3.3.3 are vulnerable to a prototype pollution-based XSS bypass. When an application uses `DOMPurify.sanitize()` with the default configuration (no `CUSTOM_ELEMENT_HANDLING` option), a prior prototype pollution gadget can inject permissive `tagNameCheck` and `attributeNameCheck` regex values into `Object.prototype`, causing DOMPurify to allow arbitrary custom elements with arbitrary attributes — including event handlers — through sanitization. Version 3.4.0 fixes the issue.
CVE-2026-41239	DOMPurify is a DOM-only cross-site scripting sanitizer for HTML, MathML, and SVG. Starting in version 1.0.10 and prior to version 3.4.0, `SAFE_FOR_TEMPLATES` strips `{{...}}` expressions from untrusted HTML. This works in string mode but not with `RETURN_DOM` or `RETURN_DOM_FRAGMENT`, allowing XSS via template-evaluating frameworks like Vue 2. Version 3.4.0 patches the issue.
CVE-2026-41240	DOMPurify is a DOM-only cross-site scripting sanitizer for HTML, MathML, and SVG. Versions prior to 3.4.0 have an inconsistency between FORBID_TAGS and FORBID_ATTR handling when function-based ADD_TAGS is used. Commit c361baa added an early exit for FORBID_ATTR at line 1214. The same fix was not applied to FORBID_TAGS. At line 1118-1123, when EXTRA_ELEMENT_HANDLING.tagCheck returns true, the short-circuit evaluation skips the FORBID_TAGS check entirely. This allows forbidden elements to survive sanitization with their attributes intact. Version 3.4.0 patches the issue.
CVE-2026-41245	Junrar is an open source java RAR archive library. Prior to version 7.5.10, a path traversal vulnerability in `LocalFolderExtractor` allows an attacker to write arbitrary files with attacker-controlled content into sibling directories when a crafted RAR archive is extracted. Version 7.5.10 fixes the issue.
CVE-2026-41417	Netty allows request-line validation to be bypassed when a `DefaultHttpRequest` or `DefaultFullHttpRequest` is created first and its URI is later changed via `setUri()`. The constructors reject CRLF and whitespace characters that would break the start-line, but `setUri()` does not apply the same validation. `HttpRequestEncoder` and `RtspEncoder` then write the URI into the request line verbatim. If attacker-controlled input reaches `setUri()`, this enables CRLF injection and insertion of additional HTTP or RTSP requests, leading to HTTP request smuggling or desynchronization on the HTTP side and request injection on the RTSP side. This issue is fixed in versions 4.2.13.Final and 4.1.133.Final.
CVE-2026-41567	Moby is an open source container framework. In versions prior to 29.5.1 and in moby/moby v2 prior to v2.0.0-beta.14, when a compressed archive is uploaded to a container via `PUT /containers/{id}/archive` or piped through `docker cp -`, the daemon resolves decompression binaries (such as `xz` or `unpigz`) from the container's filesystem rather than the host's due to incorrect ordering of operations. A malicious container image containing a trojanized decompression binary can achieve arbitrary code execution with full daemon privileges, including host root UID and unrestricted capabilities, when a user uploads a compressed (xz or gzip) archive into that container. This issue is fixed in Docker Engine 29.5.1 and moby/moby v2.0.0-beta.14. Workarounds include only running containers from trusted images, using authorization plugins to restrict access to the `PUT /containers/{id}/archive` endpoint, and avoiding piping compressed archives into containers created from untrusted images
CVE-2026-41568	Moby is an open source container framework. In Docker Engine prior to version 29.5.1, Docker Daemon versions 28.5.2 and prior, and Moby Daemon prior to version 2.0.0-beta.14, a race condition during docker cp mount setup allows a malicious container to create empty files or directories at arbitrary absolute paths on the host filesystem. This issue has been patched in Docker Engine version 29.5.1 and Moby Daemon version 2.0.0-beta.14.
CVE-2026-41608	Improper Handling of Highly Compressed Data (Data Amplification) vulnerability in Apache Thrift Python bindings. This issue affects Apache Thrift: before 0.24.0. Users are recommended to upgrade to version 0.24.0, which fixes the issue.
CVE-2026-41720	Spring LDAP's DirContextAuthenticationStrategy implementations do not reject a bind request where a non-empty username is paired with an empty or null password. Affected versions: Spring LDAP 2.4.0 through 2.4.4; 3.2.0 through 3.2.17; 3.3.0 through 3.3.7; 4.0.0 through 4.0.3.
CVE-2026-41848	Applications may be vulnerable to a Regular Expression Denial of Service (ReDoS) attack if an attacker is able to provide a pattern which is then directly or indirectly supplied to one of the following methods in AntPathMatcher: match(String pattern, String path), matchStart(String pattern, String path), extractUriTemplateVariables(String pattern, String path). Affected versions: Spring Framework 7.0.0 through 7.0.7; 6.2.0 through 6.2.18; 6.1.0 through 6.1.27; 5.3.0 through 5.3.48.
CVE-2026-41850	Applications that evaluate user-supplied Spring Expression Language (SpEL) expressions are vulnerable to an Algorithmic Denial of Service (DoS). By providing a specially crafted expression, an attacker can trigger excessive resource consumption during evaluation, leading to application degradation or unavailability. Affected versions: Spring Framework 7.0.0 through 7.0.7; 6.2.0 through 6.2.18; 6.1.0 through 6.1.27; 5.3.0 through 5.3.48.
CVE-2026-41851	Applications which accept user-supplied Spring Expression Language (SpEL) expressions may be vulnerable to a Denial of Service (DoS) attack if the evaluation of a SpEL expression triggers unbounded cache growth. Affected versions: Spring Framework 7.0.0 through 7.0.7; 6.2.0 through 6.2.18; 6.1.0 through 6.1.27; 5.3.0 through 5.3.48.

CVE ID	Description
CVE-2026-41852	A vulnerability in Spring Expression Language (SpEL) evaluation logic allows for arbitrary zero-argument method invocation, even within restricted or read-only contexts, which may allow an attacker to invoke unintended application logic. Affected versions: Spring Framework 7.0.0 through 7.0.7; 6.2.0 through 6.2.18; 6.1.0 through 6.1.27; 5.3.0 through 5.3.48.
CVE-2026-41854	Due to incorrect host parsing, applications that rely on UriComponentsBuilder to parse and validate an externally provided URL string may be exposed to a server-side request forgery (SSRF) attack. Affected versions: Spring Framework 7.0.0 through 7.0.7; 6.2.0 through 6.2.18.
CVE-2026-42033	Axios is a promise based HTTP client for the browser and Node.js. Prior to 1.15.1 and 0.31.1, when Object.prototype has been polluted by any co-dependency with keys that axios reads without a hasOwnProperty guard, an attacker can (a) silently intercept and modify every JSON response before the application sees it, or (b) fully hijack the underlying HTTP transport, gaining access to request credentials, headers, and body. The precondition is prototype pollution from a separate source in the same process. This vulnerability is fixed in 1.15.1 and 0.31.1.
CVE-2026-42034	Axios is a promise based HTTP client for the browser and Node.js. Prior to 1.15.1 and 0.31.1, for stream request bodies, maxBodyLength is bypassed when maxRedirects is set to 0 (native http/https transport path). Oversized streamed uploads are sent fully even when the caller sets strict body limits. This vulnerability is fixed in 1.15.1 and 0.31.1.
CVE-2026-42035	Axios is a promise based HTTP client for the browser and Node.js. Prior to 1.15.1 and 0.31.1, a prototype pollution gadget exists in the Axios HTTP adapter (lib/adapters/http.js) that allows an attacker to inject arbitrary HTTP headers into outgoing requests. The vulnerability exploits duck-type checking of the data payload, where if Object.prototype is polluted with getHeaders, append, pipe, on, once, and Symbol.toStringTag, Axios misidentifies any plain object payload as a FormData instance and calls the attacker-controlled getHeaders() function, merging the returned headers into the outgoing request. The vulnerable code resides exclusively in lib/adapters/http.js. The prototype pollution source does not need to originate from Axios itself — any prototype pollution primitive in any dependency in the application's dependency tree is sufficient to trigger this gadget. This vulnerability is fixed in 1.15.1 and 0.31.1.
CVE-2026-42036	Axios is a promise based HTTP client for the browser and Node.js. Prior to 1.15.1 and 0.31.1, when responseType: 'stream' is used, Axios returns the response stream without enforcing maxContentLength. This bypasses configured response-size limits and allows unbounded downstream consumption. This vulnerability is fixed in 1.15.1 and 0.31.1.
CVE-2026-42037	Axios is a promise based HTTP client for the browser and Node.js. From 1.0.0 to before 1.15.1, the FormDataPart constructor in lib/helpers/formDataToStream.js interpolates value.type directly into the Content-Type header of each multipart part without sanitizing CRLF (\r\n) sequences. An attacker who controls the .type property of a Blob/File-like object (e.g., via a user-uploaded file in a Node.js proxy service) can inject arbitrary MIME part headers into the multipart form-data body. This bypasses Node.js v18+ built-in header protections because the injection targets the multipart body structure, not HTTP request headers. This vulnerability is fixed in 1.15.1.
CVE-2026-42038	Axios is a promise based HTTP client for the browser and Node.js. Prior to 1.15.1 and 0.31.1, he fix for no_proxy hostname normalization bypass is incomplete. When no_proxy=localhost is set, requests to 127.0.0.1 and [::1] still route through the proxy instead of bypassing it. The shouldBypassProxy() function does pure string matching — it does not resolve IP aliases or loopback equivalents. This vulnerability is fixed in 1.15.1 and 0.31.1.
CVE-2026-42039	Axios is a promise based HTTP client for the browser and Node.js. Prior to 1.15.1 and 0.31.1, toFormData recursively walks nested objects with no depth limit, so a deeply nested value passed as request data crashes the Node.js process with a RangeError. This vulnerability is fixed in 1.15.1 and 0.31.1.
CVE-2026-42040	Axios is a promise based HTTP client for the browser and Node.js. Prior to 1.15.1 and 0.31.1, the encode() function in lib/helpers/AxiosURLSearchParams.js contains a character mapping (charMap) at line 21 that reverses the safe percent-encoding of null bytes. After encodeURIComponent('\x00') correctly produces the safe sequence %00, the charMap entry '%00': '\x00' converts it back to a raw null byte. Primary impact is limited because the standard axios request flow is not affected. This vulnerability is fixed in 1.15.1 and 0.31.1.
CVE-2026-42041	Axios is a promise based HTTP client for the browser and Node.js. Prior to 1.15.1 and 0.31.1, the Axios library is vulnerable to a Prototype Pollution "Gadget" attack that allows any Object.prototype pollution to silently suppress all HTTP error responses (401, 403, 500, etc.), causing them to be treated as successful responses. This completely bypasses application-level authentication and error handling. The root cause is that validateStatus is the only config property using the mergeDirectKeys merge strategy, which uses JavaScript's in operator — an operator that inherently traverses the prototype chain. When Object.prototype.validateStatus is polluted with () => true, all HTTP status codes are accepted as success. This vulnerability is fixed in 1.15.1 and 0.31.1.

CVE ID	Description
CVE-2026-42042	Axios is a promise based HTTP client for the browser and Node.js. Prior to 1.15.1 and 0.31.1, the Axios library's XSRF token protection logic uses JavaScript truthy/falsy semantics instead of strict boolean comparison for the withXSRFToken config property. When this property is set to any truthy non-boolean value (via prototype pollution or misconfiguration), the same-origin check (isURLSameOrigin) is short-circuited, causing XSRF tokens to be sent to all request targets including cross-origin servers controlled by an attacker. This vulnerability is fixed in 1.15.1 and 0.31.1.
CVE-2026-42043	Axios is a promise based HTTP client for the browser and Node.js. Prior to 1.15.1 and 0.31.1, an attacker who can influence the target URL of an Axios request can use any address in the 127.0.0.0/8 range (other than 127.0.0.1) to completely bypass the NO_PROXY protection. This vulnerability is due to an incomplete for CVE-2025-62718, This vulnerability is fixed in 1.15.1 and 0.31.1.
CVE-2026-42044	Axios is a promise based HTTP client for the browser and Node.js. From 1.0.0 to before 1.15.2, he Axios library is vulnerable to a Prototype Pollution "Gadget" attack that allows any Object.prototype pollution in the application's dependency tree to be escalated into surgical, invisible modification of all JSON API responses — including privilege escalation, balance manipulation, and authorization bypass. The default transformResponse function at lib/defaults/index.js:124 calls JSON.parse(data, this.parseReviver), where this is the merged config object. Because parseReviver is not present in Axios defaults, not validated by assertOptions, and not subject to any constraints, a polluted Object.prototype.parseReviver function is called for every key-value pair in every JSON response, allowing the attacker to selectively modify individual values while leaving the rest of the response intact. This vulnerability is fixed in 1.15.2.
CVE-2026-42198	pgjdbc is an open source postgresql JDBC Driver. From version 42.2.0 to before version 42.7.11, pgjdbc is vulnerable to a client-side denial of service during SCRAM-SHA-256 authentication. A malicious server can instruct the driver to perform SCRAM authentication with a very large iteration count. With a large enough value, the client spends an unbounded amount of CPU time inside PBKDF2 before authentication can fail. A single attempt ties up a CPU core. Repeated or concurrent attempts exhaust client CPU and can wedge connection pools. In affected versions, loginTimeout did not fully mitigate this problem. When loginTimeout expired, the caller could stop waiting, but the worker thread performing the connection attempt could continue running and burning CPU inside the SCRAM PBKDF2 computation. This issue has been patched in version 42.7.11.
CVE-2026-42264	Axios is a promise based HTTP client for the browser and Node.js. From version 1.0.0 to before version 1.15.2, fFive config properties (auth, baseURL, socketPath, beforeRedirect, and insecureHTTPParser) in the HTTP adapter are read via direct property access without hasOwnProperty guards, making them exploitable as prototype pollution gadgets. When Object.prototype is polluted by another dependency in the same process, axios silently picks up these polluted values on every outbound HTTP request. This issue has been patched in version 1.15.2.
CVE-2026-42306	Moby is an open source container framework. In Docker Engine prior to version 29.5.1, Docker Daemon versions 28.5.2 and prior, and Moby Daemon prior to version 2.0.0-beta.14, a race condition during docker cp mount setup allows a malicious container to redirect a bind mount target to an arbitrary host path, potentially overwriting host files or causing denial of service. This issue has been patched in Docker Engine version 29.5.1 and Moby Daemon version 2.0.0-beta.14.
CVE-2026-42308	Pillow is a Python imaging library. Prior to version 12.2.0, if a font advances for each glyph by an exceeding large amount, when Pillow keeps track of the current position, it may lead to an integer overflow. This issue has been patched in version 12.2.0.
CVE-2026-42309	Pillow is a Python imaging library. From version 11.2.1 to before version 12.2.0, passing nested lists as coordinates to APIs that accept coordinates such as ImagePath.Path, ImageDraw.ImageDraw.polygon and ImageDraw.ImageDraw.line could cause a heap buffer overflow, as nested lists were recursively unpacked beyond the allocated buffer. Coordinate lists are now validated to contain exactly two numeric coordinates. This issue has been patched in version 12.2.0.
CVE-2026-42310	Pillow is a Python imaging library. From version 4.2.0 to before version 12.2.0, an attacker can supply a malicious PDF that causes the process to hang indefinitely, consuming 100% CPU and making the application unresponsive. This issue has been patched in version 12.2.0.
CVE-2026-42311	Pillow is a Python imaging library. From version 10.3.0 to before version 12.2.0, processing a malicious PSD file could lead to memory corruption, potentially resulting in a crash or arbitrary code execution. This issue has been patched in version 12.2.0.
CVE-2026-42499	Pathological inputs could cause DoS through consumePhrase when parsing an email address according to RFC 5322.
CVE-2026-42502	Parsing arbitrary HTML which is then rendered using Render can result in an unexpected HTML tree. This can be leveraged to execute XSS attacks in applications that attempt to sanitize input HTML before rendering.
CVE-2026-42504	Decoding a maliciously-crafted MIME header containing many invalid encoded-words can consume excessive CPU.

CVE ID	Description
CVE-2026-42505	Handshakes which used Encrypted Client Hello could be de-anonymized by a passive network observer due to a disclosure of pre-shared key identities in the unencrypted client hello.
CVE-2026-42506	Parsing arbitrary HTML which is then rendered using Render can result in an unexpected HTML tree. This can be leveraged to execute XSS attacks in applications that attempt to sanitize input HTML before rendering.
CVE-2026-42507	When returning errors, functions in the net/textproto package would include its input as part of the error. This might allow an attacker to inject misleading content to errors that are printed or logged.
CVE-2026-42508	Previously, a revoked 'SignatureKey' belonging to a CA was not correctly checked for revocation. Now, both the 'key' and 'key.SignatureKey' are checked for @revoked.
CVE-2026-42579	Netty is an asynchronous, event-driven network application framework. Prior to 4.2.13.Final and 4.1.133.Final, Netty's DNS codec does not enforce RFC 1035 domain name constraints during either encoding or decoding. This creates a bidirectional attack surface: malicious DNS responses can exploit the decoder, and user-influenced hostnames can exploit the encoder. This vulnerability is fixed in 4.2.13.Final and 4.1.133.Final.
CVE-2026-42580	Netty is an asynchronous, event-driven network application framework. Prior to 4.2.13.Final and 4.1.133.Final, Netty's chunk size parser silently overflows int, enabling request smuggling attacks. This vulnerability is fixed in 4.2.13.Final and 4.1.133.Final.
CVE-2026-42581	Netty is an asynchronous, event-driven network application framework. Prior to 4.2.13.Final and 4.1.133.Final, HttpObjectDecoder strips a conflicting Content-Length header when a request carries both Transfer-Encoding: chunked and Content-Length, but only for HTTP/1.1 messages. The guard is absent for HTTP/1.0. An attacker that sends an HTTP/1.0 request with both headers causes Netty to decode the body as chunked while leaving Content-Length intact in the forwarded HttpMessage. Any downstream proxy or handler that trusts Content-Length over Transfer-Encoding will disagree on message boundaries, enabling request smuggling. This vulnerability is fixed in 4.2.13.Final and 4.1.133.Final.
CVE-2026-42583	Netty is an asynchronous, event-driven network application framework. Prior to 4.2.13.Final and 4.1.133.Final, Lz4FrameDecoder allocates a ByteBuf of size decompressedLength (up to 32 MB per block) before LZ4 runs. A peer only needs a 21-byte header plus compressedLength payload bytes - 22 bytes if compressedLength == 1 - to force that allocation. This vulnerability is fixed in 4.2.13.Final and 4.1.133.Final.
CVE-2026-42584	Netty is an asynchronous, event-driven network application framework. Prior to 4.2.13.Final and 4.1.133.Final, HttpClientCodec pairs each inbound response with an outbound request by queue.poll() once per response, including for 1xx. If the client pipelines GET then HEAD and the server sends 103, then 200 with GET body, then 200 for HEAD, the queue pairs HEAD with the first 200. The HEAD rule then skips reading that message’s body, so the GET entity bytes stay on the stream and the following 200 is parsed from the wrong offset. This vulnerability is fixed in 4.2.13.Final and 4.1.133.Final.
CVE-2026-42585	Netty is an asynchronous, event-driven network application framework. Prior to 4.2.13.Final and 4.1.133.Final, Netty incorrectly parses malformed Transfer-Encoding, enabling request smuggling attacks. This vulnerability is fixed in 4.2.13.Final and 4.1.133.Final.
CVE-2026-42587	Netty is an asynchronous, event-driven network application framework. Prior to 4.2.13.Final and 4.1.133.Final, HttpContentDecompressor accepts a maxAllocation parameter to limit decompression buffer size and prevent decompression bomb attacks. This limit is correctly enforced for gzip and deflate encodings via ZlibDecoder, but is silently ignored when the content encoding is br (Brotli), zstd, or snappy. An attacker can bypass the configured decompression limit by sending a compressed payload with Content-Encoding: br instead of Content-Encoding: gzip, causing unbounded memory allocation and out-of-memory denial of service. The same vulnerability exists in DelegatingDecompressorFrameListener for HTTP/2 connections. This vulnerability is fixed in 4.2.13.Final and 4.1.133.Final.
CVE-2026-43827	Default configurations of Apache Shiro have a session fixation vulnerability. This issue affects Apache Shiro from 1.0 to 2.1.0, and 3.0.0-alpha-1. Users are recommended to upgrade to version 2.1.1, or 3.0.0-alpha-2 or later, which fixes the issue. In the affected versions, when a session already exists, it is not invalidated upon successful login, nor is a new session being generated with a new ID.
CVE-2026-43828	Default configurations of Apache Shiro send sensitive cookies in HTTPS session without 'Secure' attribute. This issue affects Apache Shiro from 1.0 to 2.1.0, and 3.0.0-alpha-1. Users are recommended to upgrade to version 2.1.1, or 3.0.0-alpha-2 or later, which fixes the issue. In the affected versions, Shiro-native session manager, as well as Remember-Me manager sends JSESSIONID and rememberMe cookies without 'secure' attribute by default.
CVE-2026-43869	Improper Validation of Certificate with Host Mismatch vulnerability in Apache Thrift. This issue affects Apache Thrift: before 0.23.0. Users are recommended to upgrade to version 0.23.0, which fixes the issue.
CVE-2026-43871	Loop with Unreachable Exit Condition ('Infinite Loop') vulnerability in Apache Thrift Python, Go, PHP and Java bindings.This issue affects Apache Thrift: before 0.24.0. Users are recommended to upgrade to version 0.24.0, which fixes the issue.

CVE ID	Description
CVE-2026-44249	Netty is a network application framework for development of protocol servers and clients. In netty-handler prior to versions 4.1.135.Final and 4.2.15.Final, an attacker can bypass IPv6 subnet rules due to an incorrect masking operation in IpSubnetFilterRule.compareTo(). Valid public IP addresses can bypass the restrictions. Versions 4.1.135.Final and 4.2.15.Final patch the issue.
CVE-2026-44307	Mako is a template library written in Python. Prior to 1.3.12, on Windows, a URI using backslash traversal (e.g. \.. \..\ secret.txt) bypasses the directory traversal check in Template.__init__ and the posixpath-based normalization in TemplateLookup.get_template(), allowing reads of files outside the configured template directory. This vulnerability is fixed in 1.3.12.
CVE-2026-44431	urllib3 is an HTTP client library for Python. From 1.23 to before 2.7.0, cross-origin redirects followed from the low-level API via ProxyManager.connection_from_url().urlopen(..., assert_same_host=False) still forward these sensitive headers. This vulnerability is fixed in 2.7.0.
CVE-2026-44432	urllib3 is an HTTP client library for Python. From 2.6.0 to before 2.7.0, urllib3 could decompress the whole response instead of the requested portion (1) during the second HTTPResponse.read(amt=N) call when the response was decompressed using the official Brotli library or (2) when HTTPResponse.drain_conn() was called after the response had been read and decompressed partially (compression algorithm did not matter here). These issues could cause urllib3 to fully decode a small amount of highly compressed data in a single operation. This could result in excessive resource consumption (high CPU usage and massive memory allocation for the decompressed data) on the client side. This vulnerability is fixed in 2.7.0.
CVE-2026-44486	Axios is a promise based HTTP client for the browser and Node.js. Prior to 0.32.0 and 1.16.0, Axios' Node.js HTTP adapter can leak proxy credentials to a redirect target in affected versions. When a request is sent through an authenticated proxy, Axios may add a Proxy-Authorization header. If Axios then follows a redirect and the redirected request is no longer sent through that proxy, the stale Proxy-Authorization header can remain on the redirected request and be sent to the redirect target. This affects Node.js's use of Axios with automatic redirects enabled and an authenticated proxy configuration. Browser adapters are not affected. This vulnerability is fixed in 0.32.0 and 1.16.0.
CVE-2026-44487	Axios is a promise based HTTP client for the browser and Node.js. Prior to 0.32.0 and 1.16.0, Axios's Node.js HTTP adapter may forward a Proxy-Authorization header to a redirected origin during specific proxy-to-direct redirect flows. This affects Node.js usage, where an initial HTTP request is sent through an authenticated HTTP proxy, redirects are followed, and the redirected URL is no longer proxied. Under affected redirect shapes, the final origin can receive the proxy credential that was intended only for the outbound proxy. This vulnerability is fixed in 0.32.0 and 1.16.0.
CVE-2026-44488	Axios is a promise based HTTP client for the browser and Node.js. Axios versions 1.7.0 through 1.15.x did not enforce configured request and response size limits when requests were sent with the fetch adapter. Applications that selected adapter: 'fetch', or ran in environments where axios resolved to the fetch adapter, could receive or send bodies larger than maxContentLength or maxBodyLength despite those limits being explicitly configured. This can cause resource exhaustion in server-side usage when a malicious or compromised server returns an oversized response, when an attacker can supply a large data: URL, or when an application forwards attacker-controlled request bodies through axios while relying on maxBodyLength as a boundary. This vulnerability is fixed in 0.32.0 and 1.16.0.
CVE-2026-44490	Axios is a promise based HTTP client for the browser and Node.js. Prior to 0.32.0 and 1.16.0, axios exposes two read-side prototype-pollution gadgets. When Object.prototype is polluted by an upstream dependency in the same process (e.g. lodash _.merge / CVE-2018-16487), axios silently picks up the polluted values. (1) lib/utils.js line 406 builds merge()'s accumulator as result = {}, so result[targetKey] (line 414) walks Object.prototype and the polluted bucket's own keys are copied into the merged headers and ride out on the wire. (2) lib/core/mergeConfig.js line 26 builds the hasOwnProperty descriptor as a plain-object literal. Object.defineProperty reads descriptor.get/ descriptor.set via the prototype chain, so a polluted Object.prototype.get or Object.prototype.set makes the call throw TypeError synchronously on every axios request. This vulnerability is fixed in 0.32.0 and 1.16.0.
CVE-2026-44492	Axios is a promise based HTTP client for the browser and Node.js. Prior to 0.32.0 and 1.16.0, Axios does not normalise IPv4-mapped IPv6 addresses. When NO_PROXY lists an IPv4 address such as 127.0.0.1 or 169.254.169.254, a request URL using the IPv4-mapped IPv6 form (::ffff:7f00:1, ::ffff:a9fe:a9fe) still routes through the configured proxy. Node.js resolves these addresses to the underlying IPv4 host, so the request reaches the internal service via the proxy rather than being blocked. This vulnerability is fixed in 0.32.0 and 1.16.0.
CVE-2026-44494	Axios is a promise based HTTP client for the browser and Node.js. From 1.0.0 to before 1.16.0, the Axios library is vulnerable to a Prototype Pollution "Gadget" attack that allows any Object.prototype pollution in the application's dependency tree to be escalated into a full Man-in-the-Middle (MITM) attack — intercepting, reading, and modifying all HTTP traffic including authentication credentials. The HTTP adapter at lib/adapters/http.js:670 reads config.proxy via standard property access, which traverses the prototype chain. Because proxy is not present in Axios defaults, the merged config object has no own proxy property, making it trivially injectable via prototype pollution. Once injected, setProxy() routes all HTTP requests through the attacker's proxy server. This vulnerability is fixed in 1.16.0.

CVE ID	Description
CVE-2026-44495	Axios is a promise based HTTP client for the browser and Node.js. From 0.19.0 to before 0.31.1 and 1.15.2, Axios contains prototype-pollution gadgets in request config processing. If another vulnerability in the same JavaScript process has already polluted Object.prototype.transformResponse, affected Axios versions may treat that inherited value as request configuration or as an option validator. Axios does not itself create the prototype pollution. Exploitability requires a separate prototype-pollution vulnerability or equivalent attacker control over Object.prototype before Axios creates a request. This vulnerability is fixed in 0.31.1 and 1.15.2.
CVE-2026-44496	Axios is a promise based HTTP client for the browser and Node.js. Axios versions before 0.32.0 on the 0.x line and before 1.16.0 on the 1.x line build a regular expression from the configured XSRF cookie name without escaping regex metacharacters. In standard browser environments, an attacker who can influence the cookie name passed to axios can cause expensive regex backtracking while axios reads document.cookie. The practical impact is client-side availability degradation, such as freezing the affected browser tab while axios prepares a request. The issue does not affect ordinary Node.js HTTP adapter usage, React Native, or web workers, where axios does not read document.cookie. This vulnerability is fixed in 0.32.0 and 1.16.0.
CVE-2026-45205	Uncontrolled Recursion vulnerability in Apache Commons. When processing an untrusted configuration file, Commons Configuration will throw a StackOverflowError for YAML input with cycles. This issue affects Apache Commons: from 2.2 before 2.15.0. Users are recommended to upgrade to version 2.15.0, which fixes the issue.
CVE-2026-45292	opentelemetry-java is the Java implementation of the OpenTelemetry API for recording telemetry, and SDK for managing telemetry recorded by the API. Prior to 1.62.0, a vulnerability affects the baggage propagation implementation in opentelemetry-api and opentelemetry-extension-trace-propagators. Parsing oversized baggage causes unbounded memory allocation and CPU consumption. Because baggage is automatically re-injected into every outgoing request, the effect can fan out to downstream services that never received the original malicious request. This vulnerability is fixed in 1.62.0.
CVE-2026-45300	The AsyncHttpClient (AHC) library allows Java applications to easily execute HTTP requests and asynchronously process HTTP responses. Versions on the 2.x branch prior to 2.15.0 and the 3.x branch prior to 3.0.10 leak `Cookie` headers to cross-origin redirect targets. When following a redirect to a different origin, the `propagatedHeaders()` method in `Redirect30xInterceptor.java` strips `Authorization` and `Proxy-Authorization` headers but does not strip the `Cookie` header, causing session cookies and other sensitive cookie values to be sent to attacker-controlled servers. Versions 2.15.0 and 3.0.10 patch the issue.
CVE-2026-45409	Internationalized Domain Names in Applications (IDNA) for Python provides support for Internationalized Domain Names in Applications (IDNA) and Unicode IDNA Compatibility Processing. In versions prior to 3.15, payloads such as `"\u0660" * N` or `"\u30fb" * N + "\u6f22"` utilize the `valid_contexto` function prior to length rejection, and for high values of `N` will take a long time to process. This is the same issue as CVE-2024-3651, however the original remediation in 2024 was not a complete fix. A specially crafted argument to the `idna.encode()` function could consume significant resources. This may lead to a denial-of-service. Starting in version 3.14, the function rejects long inputs as soon as practicable prior to any further processing to minimize resource consumption. In version 3.15, this approach was extended to lesser used alternate functions (i.e. per-label conversions and codec support). A workaround is available. Domain names cannot exceed 253 characters in length. If this length limit is enforced prior to passing the domain to the `idna.encode()` function, it should no longer consume significant resources. This is triggered by arbitrarily large inputs that would not occur in normal usage, but may be passed to the library assuming there is no preliminary input validation by the higher-level application.
CVE-2026-45416	Netty is a network application framework for development of protocol servers and clients. Prior to versions 4.1.135.Final and 4.2.15.Final, SslClientHelloHandler.decode() reads the 24-bit TLS handshake length and, when the ClientHello does not fit in the first record, eagerly allocates `ctx.alloc().buffer(handshakeLength)` (line 161). The guard at line 140 is `handshakeLength > maxClientHelloLength && maxClientHelloLength != 0`, and the commonly-used SniHandler/AbstractSniHandler constructors (SniHandler(Mapping), SniHandler(AsyncMapping), AbstractSniHandler()) pass maxClientHelloLength=0 and handshakeTimeoutMillis=0, so the length guard is disabled and no timeout is scheduled. A 16 MiB request exceeds the default pooled chunk size and becomes a huge/unpooled allocation performed immediately. The buffer is retained in the handler until the channel closes. Versions 4.1.135.Final and 4.2.15.Final patch the issue.
CVE-2026-45536	Netty is a network application framework for development of protocol servers and clients. Prior to versions 4.1.135.Final and 4.2.15.Final, netty_unix_socket_recvFd sets msg_control to `char control[CMSG_SPACE(sizeof(int))]` (line 940) — 24 bytes on 64-bit Linux. A peer-sent SCM_RIGHTS cmsg carrying two ints has cmsg_len = CMSG_LEN(8) = 24, which fits exactly with no MSG_CTRUNC, so the kernel installs both fds in the receiving process. The subsequent check `cmsg->cmsg_len == CMSG_LEN(sizeof(int))` (line 972, expected 20) fails, the branch that would read the fd is skipped, and neither installed fd is closed. The for(;;) loop calls recvmsg again (non-blocking # EAGAIN # Java maps to 0 # read loop exits normally), leaving two leaked fds per message. There is no MSG_CTRUNC handling. Reachable via Epoll/KQueue DomainSocketChannel when the application opts into DomainSocketReadMode.FILE_DESCRIPTORS (non-default). Versions 4.1.135.Final and 4.2.15.Final patch the issue.

CVE ID	Description
CVE-2026-45673	Netty is a network application framework for development of protocol servers and clients. Prior to versions 4.1.135.Final and 4.2.15.Final, Netty's DNS resolver uses a predictable PRNG for generating DNS transaction IDs and defaults to a static UDP source port. This combination reduces the entropy of DNS queries, enabling DNS Cache Poisoning (Kaminsky attack). Versions 4.1.135.Final and 4.2.15.Final patch the issue.
CVE-2026-45674	Netty is a network application framework for development of protocol servers and clients. Prior to versions 4.1.135.Final and 4.2.15.Final, Netty's DnsResolveContext fails to validate the origin (bailiwick) of CNAME records in DNS responses. Versions 4.1.135.Final and 4.2.15.Final patch the issue.
CVE-2026-46595	Previously, CVE-2024-45337 fixed an authorization bypass for misused ssh server configurations; if any other type of callback is passed other than public key, then the source-address validation would be skipped.
CVE-2026-46597	An incorrectly placed cast from bytes to int allowed for server-side panic in the AES-GCM packet decoder for well-crafted inputs.
CVE-2026-46598	For certain crafted inputs, a 'ed25519.PrivateKey' was created by casting malformed wire bytes, leading to a panic when used.
CVE-2026-46600	Parsing an invalid SVCB or HTTPS RR can panic when the size of a parameter value overflows the message buffer.
CVE-2026-46718	Use of Externally-Controlled Input to Select Classes or Code ('Unsafe Reflection') vulnerability in Apache Calcite. This issue affects Apache Calcite: from 1.5.0 before 1.42. Users are recommended to upgrade to version 1.42, which fixes the issue.
CVE-2026-47244	Netty is a network application framework for development of protocol servers and clients. Prior to versions 4.1.135.Final and 4.2.15.Final, DefaultHttp2Connection.DefaultEndpoint initialises maxActiveStreams/maxStreams to Integer.MAX_VALUE, and Http2Settings never inserts SETTINGS_MAX_CONCURRENT_STREAMS by default (Http2Settings.java:305-307 only clamps a user-supplied value). Unless the application explicitly calls initialSettings().maxConcurrentStreams(n), a Netty HTTP/2 server advertises no limit and enforces none locally. Each open stream allocates a DefaultStream object, PropertyMap slots, flow-controller state and IntObjectHashMap entry; with ~2^30 permissible odd stream IDs a single TCP connection can create hundreds of thousands of long-lived stream objects. This is also the precondition for CVE-2023-44487-style Rapid-Reset amplification, where the absence of a low concurrent cap multiplies backend work. Versions 4.1.135.Final and 4.2.15.Final patch the issue.
CVE-2026-47265	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to version 3.14.0, cookies set with the `cookies` parameter on requests are sent after following a cross-origin redirect. If a developer uses the `cookies` parameter on a per-request basis then sensitive data might be leaked to an attacker if they manage to control a redirect. Version 3.14.0 patches the issue. If unable to upgrade, using a `Cookie` header in the `headers` parameter is not vulnerable.
CVE-2026-47691	Netty is a network application framework for development of protocol servers and clients. Prior to versions 4.1.135.Final and 4.2.15.Final, Netty's `DnsResolveContext` insufficiently validates the bailiwick of NS records, enabling DNS Cache Poisoning. An attacker controlling an authoritative name server for a subdomain can poison the cache for parent domains (like `.co.uk`). In `io.netty.resolver.dns.DnsResolveContext.AuthoritativeNameServerList#add` method accepts any NS record from the AUTHORITY section as long as the record's name is a suffix of the questionName. Subsequently, the `handleWithAdditional` method caches the associated A records from the ADDITIONAL section directly into the `authoritativeDnsServerCache` under the parent domain's key. This bypasses standard bailiwick rules, where a server authoritative for a subdomain should not be trusted to provide authoritative records for its parent. The poisoned cache is then used for all future resolutions under the parent domain's key. Versions 4.1.135.Final and 4.2.15.Final patch the issue.
CVE-2026-47883	UrlHandlerFilter can be vulnerable to an open redirect when configured with very broadly matching patterns. The issue applies to the filter variants in both Spring MVC and Spring WebFlux. Spring Framework 7.0.0 - 7.0.8 Spring Framework 6.2.0 - 6.2.19
CVE-2026-47889	A WebFlux application running on the Jetty 12 Core reactive adapter serializes response cookies without the sameSite attribute. Spring Framework 7.0.0 - 7.0.8 Spring Framework 6.2.0 - 6.2.19
CVE-2026-47890	Spring MVC and WebFlux applications are vulnerable to stream corruption when using Server-Sent Events (SSE) with view fragments. Spring Framework 7.0.0 - 7.0.8 Spring Framework 6.2.0 - 6.2.19

CVE ID	Description
CVE-2026-48043	Netty is a network application framework for development of protocol servers and clients. In netty-codec-http2 prior to versions 4.1.135.Final and 4.2.15.Final, the `DelegatingDecompressorFrameListener` class orchestrates HTTP/2 decompression by embedding a per-stream `EmbeddedChannel` that runs the appropriate decompression codec (gzip, deflate, zstd) and forwards decompressed chunks to a wrapped listener. Each decompressed chunk is a pooled `ByteBuf` handed to an anonymous `ChannelInboundHandlerAdapter` tail handler, which becomes the sole owner responsible for releasing it. A remote peer could send frames that would result in the flow-controller throwing and so trigger a resource leak which at the end might take down the whole JVM due OOME. Versions 4.1.135.Final and 4.2.15.Final patch the issue.
CVE-2026-48587	An issue was discovered in Django 5.2 before 5.2.15 and 6.0 before 6.0.6. `django.utils.cache.has_vary_header()` in Django does not strip leading or trailing whitespace from `Vary` response header values before comparison, which allows remote attackers to read cached responses via requests to URLs whose responses contain whitespace-padded Vary header values. Earlier, unsupported Django series (such as 5.0.x, 4.1.x, and 3.2.x) were not evaluated and may also be affected. Django would like to thank Navid Rezazadeh for reporting this issue.
CVE-2026-48588	An issue was discovered in Django 6.0 before 6.0.7 and 5.2 before 5.2.16. `UpdateCacheMiddleware` and the `cache_page()` decorator cache responses that vary on cookies when the incoming request carries unrelated cookies, which allows remote attackers to read private data from the shared cache. Earlier, unsupported Django series (such as 5.0.x, 4.1.x, and 3.2.x) were not evaluated and may also be affected. Django would like to thank Chris Whyland for reporting this issue.
CVE-2026-49264	OAuthLib is an implementation of the OAuth request-signing logic for Python 3.6+. From 0.6.1 until 4.0.0, OAuthLib's RevocationEndpoint in oauthlib/oauth2/rfc6749/endpoints/revocation.py reflects an unvalidated callback request parameter into success and error response bodies when enable_jsonp=True. The callback can come from the query string or request body and is used without restricting it to a legal JavaScript function name. A browser application that loads the revocation response as trusted JSONP and permits attacker influence over callback can therefore execute attacker-controlled JavaScript in the including application's origin, with access to its DOM and script-accessible data. The feature is disabled by default, so only deployments that explicitly enable JSONP are affected. This issue is fixed in version 4.0.0.
CVE-2026-49265	OAuthLib is an implementation of the OAuth request-signing logic for Python 3.6+. From 3.0.0 until 4.0.0, OAuthLib's PKCE checks in oauthlib/oauth2/rfc6749/grant_types/authorization_code.py compare attacker-supplied code_verifier values with stored challenges using ordinary string equality in code_challenge_method_plain and code_challenge_method_s256. Because the comparison can return after a length difference or the first mismatching character, repeated token endpoint measurements can reveal common-prefix timing information that could potentially allow recovery of the code_verifier, unauthorized access-token acquisition, and account takeover. Practical exploitation requires an intercepted, single-use authorization code and enough measurements to overcome real-world network noise; the local timing benchmark does not establish practical remote account takeover, so the weakness remains a defense-in-depth concern. This issue is fixed in version 4.0.0.
CVE-2026-49268	A remote attacker can inject LDAP special characters into the Distinguished Name (DN) construction in DefaultLdapRealm class. User-supplied username input is directly concatenated into the LDAP DN template without any escaping of RFC 2253 special characters. This allows an attacker to manipulate the DN structure used for LDAP bind authentication, potentially bypassing authentication or impersonating other users. This issue affects all Apache Shiro versions through 2.2.0, and 3.0.0-alpha-1 when using DefaultLdapRealm Upgrade to Apache Shiro 2.2.1 or 3.0.0-alpha-2 or later, which fixes the issue.
CVE-2026-49272	OAuthLib is an implementation of the OAuth request-signing logic for Python 3.6+. From 3.3.0 until 4.0.1, OAuthLib's RFC 8628 device workflow does not require DeviceCodeGrant.validate_token_request to check that the device_code parameter is present and valid. A caller whose client request is accepted by the configured validator and who knows a valid client ID enabled for the device workflow can submit a missing or bogus device code to the token endpoint and obtain a bearer token without completing device authorization. This issue is fixed in version 4.0.1.
CVE-2026-49458	DOMPurify is a DOM-only cross-site scripting sanitizer for HTML, MathML, and SVG. Prior to 3.4.6, DOMPurify.sanitize(node, { IN_PLACE: true }) accepted same-origin foreign-realm DOM nodes while follow-on checks used parent-realm constructors, causing instanceof checks for forms, named node maps, document fragments, and elements to fail and skip clobber, template-content, and shadow-DOM sanitization branches so executable markup could survive. This issue is fixed in version 3.4.6.
CVE-2026-49459	DOMPurify is a DOM-only cross-site scripting sanitizer for HTML, MathML, and SVG. Prior to 3.4.6, DOMPurify.sanitize(root, { IN_PLACE: true }) could preserve event-handler attributes on an attacker-controlled <form> root when a descendant name clobbered properties checked by _isClobbered, because _forceRemove no-opped on the parent-less root and _sanitizeAttributes returned early. This issue is fixed in version 3.4.6.

CVE ID	Description
CVE-2026-49844	Improper encoding of non-finite floating-point values during MapMessage JSON serialization in Apache Log4j API produces output that is not valid JSON. This issue affects Apache Log4j API versions 2.13.1 through 2.25.4 and version 2.26.0. The fix for CVE-2026-34481 did not cover all code paths: when a MapMessage contains a non-finite IEEE 754 value (NaN, Infinity, or -Infinity), MapMessage.asJson() emits the corresponding bare token. RFC 8259 does not permit these tokens, so a conformant parser rejects the resulting document. The defect is reachable only when both of the following conditions hold: * The application uses the message resolver https://logging.apache.org/log4j/2.x/manual/json-template-layout.html#event-template-resolver-message of JsonTemplateLayout or any other layout that relies on MapMessage.asJson() or MapMessage.getFormattedMessage(new String[]{"JSON"}). * The application logs a MapMessage that contains an attacker-controlled floating-point value. An attacker who can supply a non-finite value can cause the affected layout to emit malformed JSON, which may corrupt the enclosing log record or disrupt downstream log ingestion and parsing. Users are advised to upgrade to Apache Log4j API 2.25.5 or 2.26.1, both of which emit RFC 8259-compliant JSON for non-finite values.
CVE-2026-49978	DOMPurify is a DOM-only cross-site scripting sanitizer for HTML, MathML, and SVG. Prior to 3.4.7, DOMPurify IN_PLACE sanitization could skip shadow contents attached to an element inside <template>.content, allowing attacker-controlled markup such as event handlers, JavaScript URLs, or scripts to survive and execute when an application cloned and inserted the sanitized template. This issue is fixed in version 3.4.7.
CVE-2026-50010	Netty is a network application framework for development of protocol servers and clients. Prior to versions 4.1.135.Final and 4.2.15.Final, SimpleTrustManagerFactory.engineGetTrustManagers() and related paths wrap any user-supplied plain X509TrustManager in X509TrustManagerWrapper, which extends X509ExtendedTrustManager but implements the 3-arg checkServerTrusted(chain, authType, SSLEngine) by discarding the SSLEngine and calling the 2-arg delegate. Because the object now IS an X509ExtendedTrustManager, neither SunJSSE's internal AbstractTrustManagerWrapper nor Netty's own OpenSslX509TrustManagerWrapper will re-wrap it to add endpoint-identification. Consequently, even though Netty 4.2 sets endpointIdentificationAlgorithm="HTTPS" by default, a client built with `SslContextBuilder.forClient().trustManager(somePlainX509TrustManager)` performs no hostname verification at all. Versions 4.1.135.Final and 4.2.15.Final patch the issue.
CVE-2026-50020	Netty is a network application framework for development of protocol servers and clients. Prior to versions 4.1.135.Final and 4.2.15.Final, before reading the first request-line, `HttpObjectDecoder` skips every byte for which `Character.isISOControl(b)` is `true` (0x00–0x1F and 0x7F) as well as all whitespace. RFC 9112 §2.2 only asks servers to ignore empty CRLF lines preceding the request-line — a carefully scoped robustness allowance intended to handle HTTP/1.0 POST workarounds. Silently absorbing NUL bytes, SOH, STX, and other non-CRLF control characters goes significantly beyond this, and can be exploited for request-boundary confusion in pipelined or multiplexed transports where a front-end component treats those bytes differently. Versions 4.1.135.Final and 4.2.15.Final patch the issue.
CVE-2026-50269	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to 3.14.0, attacker-controlled input included into multipart/payload headers can be used to modify a request to inject additional headers or similar. In the unlikely situation that an application is passing user-controlled strings into MultipartWriter.append(headers=...) or Payload.headers, then an attacker may be able to modify the request to inject headers or change the contents of the request. This vulnerability is fixed in 3.14.0.
CVE-2026-50560	Netty is a network application framework for development of protocol servers and clients. Prior to versions 4.1.135.Final and 4.2.15.Final, Netty HTTP/2 max header size handling produces an attack similar to HTTP/2 Rapid Reset. There is a setting in the http2 specification called `SETTINGS_MAX_HEADER_LIST_SIZE`. When a client sends that setting to Netty, it appears that Netty will behave as follows: read the request; proxy the request to the origin; attempt to produce a response; and create an exception while writing the headers for the response. Functionally, this should be similar to the http2 reset attack, but with a different on-the-wire signature. Versions 4.1.135.Final and 4.2.15.Final patch the issue.
CVE-2026-53434	Detection of Error Condition Without Action vulnerability in Apache Tomcat when configuring CRLs for a FFM based connector. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.22, from 10.1.0-M7 through 10.1.55, from 9.0.83 through 9.0.118. Users are recommended to upgrade to version 11.0.23, 10.1.56 or 9.0.119, which fixes the issue.
CVE-2026-53877	An issue was discovered in Django 6.0 before 6.0.7 and 5.2 before 5.2.16. `django.contrib.gis.gdal.GDALRaster` over-reads its in-memory buffer when constructed from a bytes object, which can disclose adjacent memory or cause service degradation via a potential segmentation fault when the `vsi_buffer` property is accessed. Earlier, unsupported Django series (such as 5.0.x, 4.1.x, and 3.2.x) were not evaluated and may also be affected. Django would like to thank Bence Nagy for reporting this issue.
CVE-2026-53878	An issue was discovered in Django 6.0 before 6.0.7 and 5.2 before 5.2.16. `DomainNameValidator` does not prohibit newlines in domain names (unless used via a form field, since `CharField` strips newlines). If an application uses values with newlines in an HTTP response, header injection can occur. Django itself is unaffected because `HttpResponse` prohibits newlines in HTTP headers. Earlier, unsupported Django series (such as 5.0.x, 4.1.x, and 3.2.x) were not evaluated and may also be affected. Django would like to thank Bence Nagy for reporting this issue.

CVE ID	Description
CVE-2026-54058	Pillow is a Python imaging library. Prior to 12.3.0, when Pillow loads an uncompressed McIdas AREA image from a filename through the mmap raw codec path, attacker-controlled header words can set a row stride smaller than the natural row width, causing pixel access such as Image.tobytes(), getpixel, convert, or save to read beyond the mapped region and disclose adjacent process memory or fault. This issue is fixed in version 12.3.0.
CVE-2026-54059	Pillow is a Python imaging library. Prior to 12.3.0, PIL/PcfFontFile.py _load_bitmaps() read glyph dimensions from the PCF METRICS section and passed them directly to Image.frombytes() without calling Image._decompression_bomb_check(), allowing crafted PCF font data to cause excessive memory allocation. This issue is fixed in version 12.3.0.
CVE-2026-54060	Pillow is a Python imaging library. Prior to 12.3.0, PIL/FontFile.py FontFile.compile() assembled per-glyph images into a combined bitmap with Image.new("1", (xsize, ysize)) without calling Image._decompression_bomb_check(), allowing a font to trigger excessive allocation during conversion or saving. This issue is fixed in version 12.3.0.
CVE-2026-54273	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to 3.14.1, no limit was present on the number of pipelined requests that could be queued. An attacker may be able to use pipelined requests to use excessive amounts of memory, potentially leading to DoS. This vulnerability is fixed in 3.14.1.
CVE-2026-54274	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to 3.14.1, if an attacker sends large incomplete websocket frame payloads, it may be possible to bypass the usual size limits on memory use. This vulnerability is fixed in 3.14.1.
CVE-2026-54275	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to 3.14.1, the server_hostname TLS SNI check can be bypassed when an existing connection is reused. If an application makes multiple requests to the same domain, but with different per-request server_hostname parameters, then the later calls may succeed by reusing the existing connection when they should have been rejected due to the TLS SNI check. This vulnerability is fixed in 3.14.1.
CVE-2026-54276	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to 3.14.1, DigestAuthMiddleware can send an authentication response after following a cross-origin redirect. This likely requires an open redirect vulnerability or similar on the target domain for an attacker to be able to execute. Further, the attacker is only receiving the digest, so should only be able to extract the user's credentials if the cryptography is weak or there is some kind of password reuse. This vulnerability is fixed in 3.14.1.
CVE-2026-54277	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to 3.14.1, it is possible to bypass the max_line_size check in parts of an HTTP request in the C parser. If using the optimised C parser (the default in pre-built wheels), then an attacker may be able to send oversized lines through the HTTP parser and use an excessive amount of memory, potentially leading to DoS. This vulnerability is fixed in 3.14.1.
CVE-2026-54278	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to 3.14.1, during cleanup it is possible for a compressed request body to be decompressed into memory in one chunk. An attacker may be able to send a compressed payload in specific situations that could be decompressed into memory, potentially leading to DoS (a zip bomb edge case). This vulnerability is fixed in 3.14.1.
CVE-2026-54279	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to 3.14.1, host-only cookies that are saved with CookieJar.save() and then restored later with CookieJar.load() lose their host-only status. This vulnerability is fixed in 3.14.1.
CVE-2026-54280	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to 3.14.1, payload resources are not closed correctly when a client disconnects in the middle of a write. If a payload is using an open file or similar limited resource, then an attacker may be able to cause resource starvation temporarily until garbage collection or similar closes the file. This vulnerability is fixed in 3.14.1.
CVE-2026-54284	sqlparse is a non-validating SQL parser module for Python. Prior to 0.6.0, TokenList construction and string conversion in sqlparse/sql.py repeatedly flatten nested token subtrees constructed by group_parenthesis and group_case, causing quadratic CPU consumption through sqlparse.parse(), sqlparse.format(), and sqlparse.split() before depth and token limits terminate processing. This issue is fixed in version 0.6.0.
CVE-2026-54291	pgjdbc is an open source postgresql JDBC Driver. In releases 42.7.4 through 42.7.11, channelBinding=require connections can be silently downgraded from SCRAM-SHA-256-PLUS with channel binding to plain SCRAM-SHA-256 without it, losing the man-in-the-middle protection the setting is meant to guarantee. An attacker who can intercept the TLS connection can trigger the downgrade with a certificate whose signature algorithm has no tls-server-end-point channel-binding hash, because the bundled com.ongres.scram:scram-client returns an empty byte array instead of failing and pgJDBC ScramAuthenticator checks only that the server advertised a PLUS mechanism, without rejecting the empty binding or checking that the negotiated mechanism uses channel binding. This issue is fixed in version 42.7.12.
CVE-2026-54399	Uncontrolled Resource Consumption vulnerability in the HTTP/1.1 message parser in Apache HttpComponents Core (5.4.2 and earlier, 5.5-beta1 and earlier) allows an remote attacker to cause a denial of service through memory exhaustion by sending messages with excessive number of headers / excessive header length

CVE ID	Description
CVE-2026-54428	Allocation of resources without limits or throttling in the HTTP/2 HPACK decoder in Apache HttpComponents Core (5.4.2 and earlier, 5.5-beta1 and earlier) allows an remote attacker to cause a denial of service through memory exhaustion by sending oversized compressed header blocks before the HTTP/2 SETTINGS acknowledgement causes the configured header list size limit to be applied.
CVE-2026-54512	jackson-databind contains the general-purpose data-binding functionality and tree-model for Jackson Data Processor. From 2.10.0 until 2.18.8, 2.21.4, and 3.1.4, jackson-databind's PolymorphicTypeValidator (PTV) is the primary safety mechanism guarding polymorphic deserialization. When polymorphic typing is enabled and a type identifier contains generic parameters (i.e. the type ID string contains <), DatabindContext._resolveAndValidateGeneric() validates only the raw container class name (the substring before <) against the configured PTV. If the container type is approved, the method parses the full canonical type string via TypeFactory.constructFromCanonical() and returns the fully parameterized type without ever validating the nested type arguments against the PTV. The nested type arguments are then resolved, instantiated, and populated as beans during deserialization. An attacker who controls the type ID can therefore place a denied class as a generic type parameter of an allowed container — for example java.util.ArrayList<com.evil.Gadget> when only java.util.ArrayList is allow-listed. The container passes the PTV check; com.evil.Gadget is loaded via Class.forName(name, true, loader), instantiated, and its properties are set from attacker-controlled JSON. This completely bypasses an explicitly configured PTV allow-list. This vulnerability is fixed in 2.18.8, 2.21.4, and 3.1.4.
CVE-2026-54513	jackson-databind contains the general-purpose data-binding functionality and tree-model for Jackson Data Processor. From 2.10.0 until 2.18.8, 2.21.4, and 3.1.4, BasicPolymorphicTypeValidator.Builder.allowIfSubTypeIsArray() allowlists any array type based only on clazz.isArray(), without validating the array's component (element) type against the configured allowlist. A PTV built with allowIfSubTypeIsArray() plus an explicit concrete-type allowlist therefore still permits EvilType[] even though EvilType is not allowlisted. When Jackson deserializes the elements and no per-element type IDs are present, it instantiates the component type directly with no further PTV check, bypassing the allowlist. This vulnerability is fixed in 2.18.8, 2.21.4, and 3.1.4.
CVE-2026-54514	jackson-databind contains the general-purpose data-binding functionality and tree-model for Jackson Data Processor. From 2.0.0 until 2.18.8, 2.21.4, and 3.1.4, JDKFromStringDeserializer constructed InetSocketAddress with new InetSocketAddress(host, port), which performs eager DNS name resolution for hostname inputs at deserialization time. An application that binds untrusted JSON into a type containing an InetSocketAddress field issues an attacker-chosen DNS query during readValue, before any application-level validation or connect logic. The fix uses InetSocketAddress.createUnresolved(host, port), deferring DNS to an explicit connect. This vulnerability is fixed in 2.18.8, 2.21.4, and 3.1.4.
CVE-2026-54515	jackson-databind contains the general-purpose data-binding functionality and tree-model for Jackson Data Processor. From 2.8.0 until 2.18.9, 2.21.5, and 3.1.4, in BeanDeserializerBase.createContextual(), per-property @JsonIgnoreProperties exclusions are applied by _handleByNameInclusion(), producing a contextual deserializer whose BeanPropertyMap has the ignored properties removed. The subsequent per-property case-insensitivity block (triggered by @JsonFormat(ACCEPT_CASE_INSENSITIVE_PROPERTIES)) rebuilds from this._beanProperties (the original, unfiltered map) instead of contextual._beanProperties, then overwrites the filtered map — restoring every property _handleByNameInclusion had just removed. The ignored property becomes writable again. This vulnerability is fixed in 2.18.9, 2.21.5, and 3.1.4.
CVE-2026-54516	jackson-databind contains the general-purpose data-binding functionality and tree-model for Jackson Data Processor. From 2.21.0 until 2.21.4 and 3.1.4, POJOPropertiesCollector._renameProperties() allows a property with @JsonProperty("renamed") on the getter and @JsonIgnore on the setter to be renamed rather than dropped. With MapperFeature.INFER_PROPERTY_MUTATORS enabled (default), the private backing field is retained; during deserialization BeanDeserializerFactory.addBeanProps() sees hasField()==true, builds a FieldProperty, and makes the backing field writable. An attacker supplying the renamed JSON key writes the backing field directly, bypassing the @JsonIgnore on the setter. This vulnerability is fixed in 3.1.4.
CVE-2026-54517	jackson-databind contains the general-purpose data-binding functionality and tree-model for Jackson Data Processor. From 2.21.0 until 2.21.4 and 3.1.4, in BeanDeserializer._deserializeUsingPropertyBased, the active-view (@JsonView) filter was applied only to creator properties; the regular property-buffering branch performed no prop.visibleInView(activeView) check. A change making SetterlessProperty.isMerging() return true routed setterless Collection/Map properties through this unguarded path, so a setterless collection annotated with a restricted @JsonView is populated from attacker JSON even when the active view excludes it. This vulnerability is fixed in 2.21.4 and 3.1.4.
CVE-2026-54518	jackson-databind contains the general-purpose data-binding functionality and tree-model for Jackson Data Processor. From 2.21.0 until 2.21.4 and 3.1.4, UnwrappedPropertyHandler.processUnwrappedCreatorProperties() replays buffered JSON into creator parameters but never consults prop.visibleInView(activeView). The normal property-based creator path gates creator properties on the active view, but this unwrapped-creator replay path bypasses that check, so a constructor parameter annotated with both @JsonView(AdminView.class) and @JsonUnwrapped is populated from attacker JSON even when a more restrictive view is active. This vulnerability is fixed in 2.21.4 and 3.1.4.

CVE ID	Description
CVE-2026-54704	OpenTelemetry Java Instrumentation provides OpenTelemetry auto-instrumentation and instrumentation libraries for Java. In versions prior to 2.28.0, the JDBC auto-instrumentation may fail to sanitize passwords in SQL CONNECT statements when the password is double-quoted. As a result, clear-text database passwords can be added to trace span attributes and exported to observability backends. This issue has been fixed in version 2.28.0.
CVE-2026-54712	OpenTelemetry Java Instrumentation provides OpenTelemetry auto-instrumentation and instrumentation libraries for Java. In versions prior to 2.27.0, the RMI context propagation payload reader limits the number of context entries but does not limit the aggregate size of the strings read from the stream. An attacker who can reach an RMI endpoint on an instrumented JVM can send an oversized context propagation payload. This can cause excessive memory allocation while the JVM reads the payload, potentially leading to denial of service. The issue affects only deployments where RMI instrumentation is enabled and an RMI endpoint is network-reachable. This issue has been fixed in version 2.27.0.
CVE-2026-55153	mchange-commons-java is a Java library of shared utility classes used by mchange projects like the c3p0 connection pool. Prior to version 0.6.0, its JNDI ObjectFactory implementation (com.mchange.v2.naming.JavaBeanObjectFactory) will construct objects of arbitrary classes and initialize "JavaBean"-style properties, which for certain classes enables JNDI injection and "deserialization gadgets." Such initialization is unsafe for some classes: for example, setting the contentType property of a Swing JEditorPane to text/html and its text property to HTML containing a stylesheet <link> will provoke an HTTP GET on an arbitrary URL, potentially from within a trusted security domain. The problem is aggravated by the library's ReferenceIndirector, through which malicious JNDI Reference objects can be smuggled in for dereferencing wherever an application reads a Java-serialized object. This has been resolved in version 0.6.0.
CVE-2026-55223	c3p0 is a JDBC Connection pooling library. In versions prior to 0.14.0, c3p0 in combination with other libraries, can compose to a "sink" for deserialization gadgets. The JDBC spec's DataSource.getConnection() and ConnectionPoolDataSource.getPooledConnection() match the getXXX() form, so JavaBean libraries treat them as "properties" assumed safe while they actually call into JDBC drivers. Attackers can thus craft malicious DataSource objects whose property lookups invoke vulnerable drivers, then smuggle them in serialized form to where an application deserializes and auto-resolves bean properties — triggering the attack. This requires a susceptible DataSource/ConnectionPoolDataSource and JDBC driver on the CLASSPATH, plus a carrier that auto-looks-up JavaBean properties on = deserialization, most commonly a collection paired with an Apache commons-beanutils Comparator that sorts by bean properties. c3p0 supplied that susceptible DataSource/ConnectionPoolDataSource, which was an essential component of the trigger. This issue has been fixed in version 0.14.0.
CVE-2026-55379	Pillow is a Python imaging library. Prior to 12.3.0, PIL/BdfFontFile.py bdf_char() read the BBX width and height field from a BDF font file and passed attacker-controlled dimensions to Image.new() without calling Image._decompression_bomb_check(), bypassing Pillow's documented decompression bomb protection and allowing excessive memory allocation. This issue is fixed in version 12.3.0.
CVE-2026-55380	Pillow is a Python imaging library. Prior to 12.3.0, PIL/GdImageFile.py GdImageFile._open() read image dimensions from the GD 2.x header and stored them in self._size without calling Image._decompression_bomb_check(), allowing a crafted .gd file to trigger excessive C-heap allocation when loaded. This issue is fixed in version 12.3.0.
CVE-2026-55688	The AsyncHttpClient (AHC) library allows Java applications to easily execute HTTP requests and asynchronously process HTTP responses. In versions from 2.0.0 prior to 2.16.0 and from 3.0.0.Beta1 prior to 3.0.11, ThreadSafeCookieStore stored a cookie under the value of its Domain attribute without verifying that the responding host is allowed to set a cookie for that domain, leading to a cookie tossing / cookie injection issue. A host the client connects to can therefore plant a cookie scoped to an unrelated domain, and the client will then send that cookie on later requests to that domain. Applications that use a single AsyncHttpClient instance - and thus the default, shared CookieStore - to reach both an attacker-influenced host and a trusted host are impacted. This issue has been fixed in versions 2.16.0 and 3.0.11.
CVE-2026-55798	Pillow is a Python imaging library. Prior to 12.3.0, WindowsViewer.get_command() constructed a cmd.exe shell command by directly embedding a file path into an f-string without escaping and passed the result to subprocess.Popen(..., shell=True), allowing shell metacharacters in the file path to inject arbitrary cmd.exe commands. This issue is fixed in version 12.3.0.
CVE-2026-55831	Netty is a network application framework for development of protocol servers and clients. Prior to 4.1.136.Final and 4.2.16.Final, Netty's SPDY SETTINGS decoder accepts a peer-declared SETTINGS entry count up to the 24-bit frame-length limit and materializes every unique setting ID in `DefaultSpdySettingsFrame`, allowing a remote SPDY/3.1 peer to send a syntactically valid roughly 2 MiB SETTINGS frame that creates 262144 map entries and amplifies network input into heap growth and ordered-map insertion work. This issue is fixed in versions 4.1.136.Final and 4.2.16.Final.

CVE ID	Description
CVE-2026-55833	Netty is a network application framework for development of protocol servers and clients. Prior to 4.1.136.Final and 4.2.16.Final, Netty SPDY header decoding continues inflating zlib-compressed header blocks after the raw header parser has exceeded `maxHeaderSize` and marked the frame truncated in `SpdyFrameCodec`, allowing a remote peer to send a small compressed `HEADERS` block that expands into much larger raw header data and causes compression-amplified CPU and allocation churn. This issue is fixed in versions 4.1.136.Final and 4.2.16.Final.
CVE-2026-55856	MariaDB Connector/J is used to connect applications developed in Java to MariaDB and MySQL databases. Prior to 2.7.14, 3.3.5, 3.4.3, and 3.5.9, when a Java application connects with sslMode=verify-full or sslMode=verify-ca, supplies a password, and does not configure serverSslCert or trustStore, Connector/J can accept an untrusted self-signed certificate through the fallbackToSystemTrustStore=true ephemeral trust manager and record its certFingerprint for later identity binding. The OK-packet and authentication-switch paths enforce the certificate fingerprint, but the initial-handshake path does not. HandshakeResponse.encode() can therefore build and send a mysql_clear_password response before checking certFingerprint != null && !isMitMProof(), sslMode, or whether the authentication plugin is resistant to a man-in-the-middle, and the initial path also bypasses restrictedAuth. An active man-in-the-middle or hostile server can present a self-signed certificate, claim to be MariaDB, select mysql_clear_password as the initial authentication plugin, and receive the full database password before the connection is rejected. This issue is fixed in versions 2.7.14, 3.3.5, 3.4.3, and 3.5.9.
CVE-2026-55857	MariaDB Connector/J is used to connect applications developed in Java to MariaDB and MySQL databases. Prior to 2.7.14, 3.3.5, 3.4.3, and 3.5.9, PAM dialog authentication can be coerced into transmitting the account password over an insecure connection. The mysql_clear_password plugin is gated behind a secure transport, but the sibling PAM handler SendPamAuthPacketFactory, named dialog by the server, does not declare that requirement and inherits the default secure-required value false; older branches implement the same affected behavior in SendPamAuthPacket. A hostile or man-in-the-middle server can send an Authentication Switch Request for dialog over plain TCP, causing the driver to return the user's password in cleartext when sslMode=DISABLE and restrictedAuth=null, which is the default configuration. Properly verified TLS and local Unix sockets are not exposed to this transport vector. This issue is fixed in versions 2.7.14, 3.3.5, 3.4.3, and 3.5.9.
CVE-2026-55858	MariaDB Connector/J is used to connect applications developed in Java to MariaDB and MySQL databases. Prior to 2.7.14, 3.3.5, 3.4.3, and 3.5.9, the connector encodes and decodes protocol text and performs client-side escaping under the assumption that the connection character set is UTF-8. The server can report a mid-session change to character_set_client through OK-packet session-state tracking, including a change caused by SET NAMES, a stored routine or trigger, server configuration, or a hostile server. If character_set_client changes to a non-UTF-8 value, the driver continues to read and write UTF-8 while the server interprets the same bytes under another encoding, causing silent data corruption and a client/server charset-confusion mismatch that can defeat byte-wise quoting or escaping. The fix accepts only utf8, utf8mb3, or utf8mb4 after initialization; any other value causes SQLException with SQLState 08000 and closes the connection. This issue is fixed in versions 2.7.14, 3.3.5, 3.4.3, and 3.5.9.
CVE-2026-56130	"Remember me" cookie age is not verified on the server. This potentially allows an attacker to intercept a valid cookie and reuse it indefinitely, even after the configured expiration time has passed. This issue affects all Apache Shiro versions from 1.2.4 through 2.x, and 3.0.0-alpha-1, only when RememberMe functionality is enabled. Upgrade to version 3.0.0 or later, which fixes the issue.
CVE-2026-56740	JLine is a Java library for handling console input. Prior to 3.30.14, 4.0.16, and 4.2.1, the JLine3 Telnet server remote-telnet module does not limit the number of environment variables a client may inject via the Telnet NEW-ENVIRON option, and TelnetIO.readNEVariables() in TelnetIO.java:1127-1180 stores each variable pair in a HashMap held by ConnectionData, allowing an unauthenticated attacker to flood unique variable pairs before the terminating IAC SE byte and exhaust JVM heap memory with an OutOfMemoryError. This issue is fixed in versions 3.30.14, 4.0.16, and 4.2.1.
CVE-2026-56741	JLine is a Java library for handling console input. Prior to 3.30.14, 4.0.16, and 4.2.1, the JLine3 Telnet server remote-telnet module does not apply an upper bound to terminal dimensions received via the Telnet NAWS option, and TelnetIO.handleNAWS() in TelnetIO.java:856-879 reads client-supplied width and height as 16-bit unsigned integers and passes values such as 65535x65535 to setTerminalGeometry(), allowing an unauthenticated remote attacker to repeatedly alternate values and trigger continuous expensive rendering work that causes CPU exhaustion and denial of service. This issue is fixed in versions 3.30.14, 4.0.16, and 4.2.1.
CVE-2026-56745	Netty is a network application framework for development of protocol servers and clients. In versions 4.2.0.Final through 4.2.15.Final and 4.1.0.Final through 4.1.135.Final, the `SpdyHttpDecoder` handler in Netty's SPDY-to-HTTP codec allocates a pooled `ByteBuf` when processing a client-initiated `SYN_STREAM` frame with `FLAG_FIN=0` and stores the partially constructed `FullHttpRequest` in `messageMap`; when the remote peer sends `RST_STREAM` for that stream or the accumulated content exceeds `maxContentLength`, the decoder removes the entry but does not release the pooled `ByteBuf`, causing native memory exhaustion. This issue is fixed in versions 4.1.136.Final and 4.2.16.Final.

CVE ID	Description
CVE-2026-56746	Netty is a network application framework for development of protocol servers and clients. Versions 4.2.0.Final through 4.2.15.Final and 4.1.0.Final through 4.1.135.Final, are vulnerable to security control bypass during the origin evaluation process. CorsHandler provides a shortCircuit() configuration designed to reject unauthorized cross-origin requests immediately, acting as a security control before requests reach the application. However, due to a logical operator error in the origin evaluation process, this protection can be entirely bypassed. An attacker can bypass the short-circuit mechanism by sending a request with an Origin: null header. This failure forwards unauthorized requests to the backend application, bypassing intended access controls. This issue is fixed in versions 4.1.136.Final and 4.2.16.Final.
CVE-2026-56819	Netty is a network application framework for development of protocol servers and clients. In versions 4.2.0.Final through 4.2.15.Final and 4.1.0.Final through 4.1.135.Final, a remote unauthenticated peer can leak one direct `ByteBuf` per HTTP/2 `DATA` frame in applications that enable HTTP/2 content decompression via `DelegatingDecompressorFrameListener`. When a `DATA` frame is processed for a stream whose decompressor has already been closed, `Http2Decompressor.decompress(...)` calls `decompressor.writeInbound(data.retain())` and does not release the retained buffer on the error path, eventually exhausting direct memory and crashing the JVM. This issue is fixed in versions 4.1.136.Final and 4.2.16.Final.
CVE-2026-56851	The Nickname profile can panic with an out-of-bounds slice error when transforming crafted input into a short destination buffer.
CVE-2026-56852	A norm.Iter can enter an infinite loop when handling input containing invalid UTF-8 bytes.
CVE-2026-56853	When a server is configured to support unencrypted HTTP/2, it reads a few bytes from each new connection to see if they contain the HTTP/2 client preface. ReadHeaderTimeout is unexpectedly not being applied when doing this.
CVE-2026-56854	The source-address critical option in the Permissions returned by an authentication callback was only enforced for the PublicKeyCallback and VerifiedPublicKeyCallback paths, extending the fix for CVE-2026-46595. Permissions returned by the PasswordCallback, KeyboardInteractiveCallback, NoClientAuthCallback, and GSSAPIWithMICConfig.AllowLogin callbacks were not validated against the client's remote address, so a source-address restriction set by those callbacks was silently ignored. The check is now applied to the Permissions returned by any authentication callback.
CVE-2026-56855	Previously, after a channel has been established, a malicious peer could send crafted messages that would deadlock the entire connection. Now, we handle all RFC 4254 channel messages; global requests are handled explicitly. Then, treat all other messages as a protocol error and tear the connection down instead of buffering and blocking.
CVE-2026-56857	On Windows, when the target of Root.Mkdir or Root.MkdirAll is a junction pointing to an empty location, the operation can create a directory at the junction target even when that target is located outside the root. This only applies to operations where the last path component is a junction (path/to/junction, but not path/junction/target).
CVE-2026-56858	Previously, pathological inputs could close an unescaped '/' early, allowing for attack-controlled data to inject arbitrary content, potentially leading to XSS.
CVE-2026-56859	Previously, DecodeElement would reset the depth counter causing it to never fire; this could lead to stack exhaustion.
CVE-2026-56860	Previously, resolving relative paths containing parent directory ('..') segments performed string conversions and buffer rewrites on each step, resulting in quadratic time complexity and high memory allocation overhead. Now, path resolution operates on a byte buffer using index-based backtracking for '..' segments, eliminating the quadratic time complexity and significantly reducing memory allocations.
CVE-2026-56862	Handshake messages, such as KeyUpdate, are always considered as state-advancing, regardless of whether a handshake has been completed or not. As a result, a malicious client can keep sending KeyUpdate messages to force the server to keep performing key derivation operations indefinitely.
CVE-2026-56866	When http.Transport sends an HTTP/1 CONNECT request with a non-empty Request.Body, it writes the body directly to the connection without framing after the request headers. If the server rejects the CONNECT request with a non-2xx keep-alive response, Transport returns the connection to the idle pool. Because CONNECT requests do not have a request body, the server may interpret the trailing body bytes as a subsequent pipelined HTTP/1.1 request on the connection, leaving the pooled connection desynchronized and causing the next caller that reuses it to read the response to the injected request. In reverse proxies (including httputil.ReverseProxy) that forward CONNECT requests through a shared Transport, this can lead to cross-user response poisoning.
CVE-2026-57914	By sending a deeply nested ASN1 structure to a Apache Kerby client or service, it's possible to trigger a StackOverFlow Exception which can lead to denial of service issues. Users are recommended to upgrade to version 2.1.2, which fixes this issue.

CVE ID	Description
CVE-2026-59171	JSON-java is a reference implementation of a JSON package in Java. From version 20150729 through 20260522, JSONTokener.nextSimpleValue can accumulate unbounded numeric tokens before JSONObject.stringToNumber passes them to BigInteger or BigDecimal conversion without a length limit. An application that parses attacker-controlled JSON or XML through JSONObject, JSONArray, or XML.toJSONObject can spend super-linear CPU time converting a sufficiently long numeric literal, blocking request-handling threads and causing denial of service. Strict mode and keep-strings settings in JSONParserConfiguration do not prevent the affected conversion path. Version 20260719 limits numeric conversion length.
CVE-2026-59197	Pillow is a Python imaging library. Prior to 12.3.0, Pillow's public rank-filter API can trigger a native heap out-of-bounds write when given a very large odd filter size because ImageFilter.RankFilter.filter() calls image.expand(size // 2, size // 2) before rank-filter size validation and ImagingExpand() computes output dimensions with unchecked signed int arithmetic. This issue is fixed in version 12.3.0.
CVE-2026-59198	Pillow is a Python imaging library. From 5.2.0 until 12.3.0, Pillow's TGA RLE encoder reads past its packed row buffer when saving a mode 1 image with TGA RLE compression, allowing adjacent process heap bytes to be copied into the generated TGA file. This issue is fixed in version 12.3.0.
CVE-2026-59199	Pillow is a Python imaging library. Prior to 12.3.0, Pillow public image coordinate APIs can trigger a native heap out-of-bounds write when given coordinates near the signed 32-bit integer limits in Image.paste(), Image.crop(), or Image.alpha_composite(). This issue is fixed in version 12.3.0.
CVE-2026-59200	Pillow is a Python imaging library. From 5.1.0 until 12.3.0, PdfParser.PdfStream.decode() in PIL/PdfParser.py calls zlib.decompress() with bufsize set to the PDF stream Length field without bounding the decompressed output size, allowing a crafted FlateDecode PDF stream to exhaust memory from a small file. This issue is fixed in version 12.3.0.
CVE-2026-59203	Pillow is a Python imaging library. From 12.0.0 through 12.2.0, Pillow's EPS parser in PIL/EpsImagePlugin.py accepts a negative byte count in the %%BeginBinary directive, allowing a crafted EPS file to cause Image.open() to seek backwards to the same directive and parse it repeatedly in an infinite loop. This issue is fixed in version 12.3.0.
CVE-2026-59204	Pillow is a Python imaging library. From 8.2.0 through 12.2.0, src/libImaging/Jpeg2KDecode.c accumulates total_component_width across every tile in a JPEG2000 image instead of recomputing it per tile, allowing a crafted tiled JPEG2000 file to force substantially higher transient memory usage and trigger out-of-memory failures during decoding. This issue is fixed in version 12.3.0.
CVE-2026-59205	Pillow is a Python imaging library. Prior to 12.3.0, Pillow's ImageCms.ImageCmsTransform.apply(im, imOut) API can trigger controlled native heap corruption when the caller supplies an output image whose mode does not match the transform's declared output mode. This issue is fixed in version 12.3.0.
CVE-2026-59881	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to 3.14.2, the WebSocket client accepts and decompresses frames with the RSV1 bit set even when the permessage-deflate extension was not negotiated, allowing a malicious server to cause unexpected CPU and memory consumption. This issue is fixed in version 3.14.2.
CVE-2026-59888	jackson-databind contains the general-purpose data-binding functionality and tree-model for Jackson Data Processor. From 2.15.0 until 2.18.8, 2.21.4, and 3.1.4, Java Records using a PropertyNamingStrategy can bypass @JsonIgnore because POJOPropertiesCollector._removeUnwantedIgnorals() records an ignored component under its original implicit name before _renameUsing() applies the naming strategy, allowing the renamed JSON key to be assigned to the Record constructor parameter. This issue is fixed in versions 2.18.8, 2.21.4, and 3.1.4.
CVE-2026-59889	jackson-databind contains the general-purpose data-binding functionality and tree-model for Jackson Data Processor. From 2.18.0 until 2.18.9, 2.21.5, 2.22.1, 3.1.5, and 3.2.1, UnwrappedPropertyHandler.processUnwrapped() replays buffered JSON for a @JsonUnwrapped property and calls prop.deserializeAndSet() without a prop.visibleInView(ctxt.getActiveView()) guard, allowing a property annotated with both @JsonView and @JsonUnwrapped to be written from attacker JSON under a less-privileged active view. This issue is fixed in versions 2.18.9, 2.21.5, 2.22.1, 3.1.5, and 3.2.1.
CVE-2026-59890	setuptools is a package that allows users to download, build, install, upgrade, and uninstall Python packages. Prior to 83.0.0, FileList applied MANIFEST.in exclude, global-exclude, recursive-exclude, and prune directives by matching compiled glob patterns against on-disk file names without Unicode normalization, so on macOS APFS or HFS+ an NFD file name could bypass an NFC exclusion rule and be packed into a source distribution. This issue is fixed in version 83.0.0.
CVE-2026-59893	sqlparse is a non-validating SQL parser module for Python. Prior to 0.6.0, SQL_REGEX in sqlparse/keywords.py and the per-position loop in sqlparse/lexer.py repeatedly scan unmatched dollar-quoted literal and multiline-comment delimiters, causing quadratic CPU consumption through sqlparse.parse(), sqlparse.format(), and sqlparse.split(). This issue is fixed in version 0.6.0.

CVE ID	Description
CVE-2026-59894	sqlparse is a non-validating SQL parser module for Python. Prior to 0.6.0, sqlparse/filters/output.py fails to escape existing backslashes before quotes in sqlparse.format output_format='python' and output_format='php' and the corresponding sqlformat -l modes, allowing crafted SQL to terminate the generated string and inject Python or PHP code when a downstream consumer executes or imports the generated source. This issue is fixed in version 0.6.0.
CVE-2026-59898	Netty is an asynchronous, event-driven network application framework. Prior to versions 4.1.136.Final and 4.2.16.Final, ab attacker can force WebSocket upgrade via the lax V07 (or V08) handshaker by sending `Sec-WebSocket-Version: 7` and omitting `Connection: Upgrade` / `Upgrade: websocket` headers, completing a protocol switch that a proxy would not recognize as an Upgrade request and enabling HTTP request smuggling / protocol-confusion attacks. This issue has been fixed in versions 4.1.136.Final and 4.2.16.Final.
CVE-2026-59899	Netty is an asynchronous, event-driven network application framework. Prior to versions 4.1.136.Final and 4.2.16.Final, `HttpContentEncoder` (the superclass of the production handler `HttpContentCompressor`) maintains a per-channel `ArrayDeque<CharSequence>` named `acceptEncodingQueue` that accumulates attacker-controlled data without any size limit. The queue is filled on the I/O thread for every inbound HTTP request and drained only when the application later writes a non-1xx response. This creates a resource exhaustion vulnerability when an attacker exploits HTTP/1.1 pipelining to flood the connection with requests faster than the application produces responses. This issue has been fixed in versions 4.1.136.Final and 4.2.16.Final.
CVE-2026-59900	Netty is an asynchronous, event-driven network application framework. Prior to versions 4.1.136.Final and 4.2.16.Final, Netty's HTTP/2-to-HTTP/1.x translation layer (`Http2StreamFrameToHttpObjectCodec` and `InboundHttp2ToHttpAdapter`) fails to deduplicate or validate `Host` headers when an HTTP/2 client supplies both the `:authority` pseudo-header and a literal `host` header in a single HEADERS frame. The translator maps `:authority` to `Host` and separately copies the literal `host` header, producing an `HttpRequest` object containing two `Host` headers with attacker-controlled differing values. This issue has been fixed in versions 4.1.136.Final and 4.2.16.Final.
CVE-2026-59901	Netty is an asynchronous, event-driven network application framework. Prior to versions 4.1.136.Final and 4.2.16.Final, the `Bzip2Decoder` handler in Netty's compression codec pipeline is vulnerable to a denial-of-service attack through a malformed bzip2 stream that permanently captures the event-loop thread in an infinite loop. The vulnerability exists in the run-length encoding (RLE) state machine within [`Bzip2BlockDecompressor.read()`]. This issue has been fixed in versions 4.1.136.Final and 4.2.16.Final.
CVE-2026-59902	Netty is an asynchronous, event-driven network application framework. Prior to 4.1.137.Final and 4.2.17.Final, io.netty.handler.codec.sctp.SctpMessageCompletionHandler limits incomplete messages and fragment counts but not maxBufferedBytes, allowing unauthenticated peers to exhaust memory with large SCTP fragments. This issue is fixed in versions 4.1.137.Final and 4.2.17.Final.
CVE-2026-59903	Netty is an asynchronous, event-driven network application framework. Prior to 4.1.137.Final and 4.2.17.Final, io.netty.handler.codec.http.cors.CorsHandler setVaryHeader replaces application Vary headers such as Authorization or Cookie with Origin, allowing a caching proxy or CDN to reuse authenticated responses across users and disclose sensitive information. This issue is fixed in versions 4.1.137.Final and 4.2.17.Final.
CVE-2026-59921	Netty is an asynchronous, event-driven network application framework. Prior to versions 4.1.136.Final and 4.2.16.Final, HttpPostRequestEncoder constructs multipart HTTP request bodies by directly concatenating user-supplied filenames and field names into Content-Disposition MIME headers without validating or sanitizing CRLF characters (\r\n). Since MIME headers are delimited by CRLF, an attacker who controls the filename can inject arbitrary MIME headers into the multipart body part. The root cause is that neither the encoder nor the FileUpload implementations' setFilename() methods, which only check for null, neutralize CRLF characters before the filename is embedded into the header. This issue has been fixed in versions 4.1.136.Final and 4.2.16.Final.
CVE-2026-59949	yawkat LZ4 Java provides LZ4 compression for Java. Prior to 1.11.1, JNI-backed XXHash implementations fail to validate the byte array object and the off and len arguments in XXHashFactory.nativeInstance().hash32().hash(), XXHashFactory.nativeInstance().hash64().hash(), XXHashFactory.nativeInstance().newStreamingHash32().update(), and XXHashFactory.nativeInstance().newStreamingHash64().update(), allowing null arrays or oversized ranges to reach native code, read outside the Java array, and fatally terminate the JVM. This issue is fixed in version 1.11.1.
CVE-2026-61634	The RabbitMQ Java client library allows Java and JVM-based applications to connect to and interact with RabbitMQ nodes. Prior to 5.33.0, the AMQP connection tuning path records the negotiated AMQP frame_max value, but src/main/java/com/rabbitmq/client/impl/SocketFrameHandler.java and NettyFrameHandlerFactory continue to validate broker-controlled frame payload lengths against maxInboundMessageBodySize because the negotiated limit is not applied consistently through setMaxInboundFramePayloadSize. A malicious or compromised broker can send a method frame larger than the negotiated frame_max during or after connection establishment, causing the client to allocate and decode a protocol-invalid frame instead of rejecting it with MalformedFrameException. The protocol violation can disrupt the affected connection and cause client-side denial of service. This issue is fixed in version 5.33.0.

CVE ID	Description
CVE-2026-61700	MariaDB Connector/J is used to connect applications developed in Java to MariaDB and MySQL databases. Prior to 2.7.14, 3.3.5, 3.4.3, and 3.5.9, ClientMessage.readPacket processes a server-initiated LOCAL INFILE protocol packet 0xfb without enforcing allowLocalInfile=false. When an application sends a LOAD DATA LOCAL INFILE COM_QUERY, a rogue or man-in-the-middle server can echo the same filename and cause the connector to transmit that file despite the disabled option. The server cannot redirect the request to an arbitrary path and can receive only the exact file already selected by the application, so exploitation requires an application that actively loads sensitive data over an untrusted connection. This issue is fixed in versions 2.7.14, 3.3.5, 3.4.3, and 3.5.9.
CVE-2026-63335	The RabbitMQ Java client library allows Java and JVM-based applications to connect to and interact with RabbitMQ nodes. Prior to 5.31.0, inbound AMQP command assembly in src/main/java/com/rabbitmq/client/impl/CommandAssembler.java processes a content-bearing method and header whose remainingBodyBytes value is smaller than a following AMQP.FRAME_BODY payload. CommandAssembler.consumeBodyFrame subtracts the peer-controlled payload length before validating that it fits, drives remainingBodyBytes negative, and throws a raw UnsupportedOperationException instead of MalformedFrameException. A malicious or compromised broker peer can send this malformed sequence on an open nonzero channel to terminate frame processing and close the client connection, causing denial of service for work using that connection. This issue is fixed in version 5.31.0.
CVE-2026-63336	The RabbitMQ Java client library allows Java and JVM-based applications to connect to and interact with RabbitMQ nodes. Prior to 5.33.0, com.rabbitmq.client.ConnectionFactory.useSslProtocol() and ConnectionFactory.useSslProtocol(String) configure com.rabbitmq.client.TrustEverythingTrustManager and leave hostname verification disabled, causing arbitrary server certificates, including self-signed certificates, to be accepted. A network attacker able to intercept a TLS connection can impersonate the RabbitMQ broker, read protected AMQP traffic, and modify traffic without certificate or hostname validation. The fix changes the production TLS helpers to use the JVM default trust store and enables hostname verification, while retaining an explicitly named development-only no-verification helper. This issue is fixed in version 5.33.0.
CVE-2026-63337	The RabbitMQ Java client library allows Java and JVM-based applications to connect to and interact with RabbitMQ nodes. Prior to 5.33.0, com.rabbitmq.tools.jsonrpc.ProcedureDescription receives a javaReturnType value in an untrusted system.describe response and passes it through JSONUtil.tryFill, setJavaReturnType, and computeReturnTypeAsJavaClass to Class.forName(javaReturnType) with initialization enabled. An attacker able to answer the JsonRpcClient request through a shared broker or network interception can select a class already present in the victim JVM and trigger its static initializer, while JsonRpcClient.java later passes getReturnType output to mapper.parse and may also create type confusion. Successful exploitation can affect confidentiality, integrity, and availability in the client process. This issue is fixed in version 5.33.0.
CVE-2026-63374	AnyIO is a high level asynchronous concurrency and networking framework that works on top of either Trio or asyncio. Prior to 4.14.2, connect_tcp() and TLSStream.wrap() can validate internationalized host names after the standard library converts them with IDNA 2003 instead of IDNA 2008. When a connection to a non-ASCII domain is hijacked or redirected, an attacker can obtain a legitimate certificate for the different ASCII hostname produced by IDNA 2003 and present it to the client, causing the malicious endpoint's certificate to validate. This issue is fixed in version 4.14.2.
CVE-2026-64607	HttpClient based on the classic i/o model fails to correctly release the underlying connection back to the connection manager if it encounters an invalid or unsupported `Content-Encoding` header value in the response message. Please note this defect does not affect HttpClient based on the async i/o model. This issue affects Apache HttpComponents Client: from 5.0-alpha1 through 5.6.2.
CVE-2026-64847	AnyIO is a high level asynchronous concurrency and networking framework that works on top of either Trio or asyncio. Prior to 4.14.2, AnyIO starts process-pool workers with standard error connected to a pipe that the parent never drains, even though the documented behavior redirects all three standard streams. Worker code that writes enough attacker-influenced data to sys.stderr can fill the pipe and block before returning the standard-output protocol response, causing the awaiting process-pool call to remain blocked indefinitely. Applications that run untrusted or faulty worker code capable of producing substantial standard-error output are affected. This issue is fixed in version 4.14.2.
CVE-2026-65182	Improper Access Control, Incorrect Authorization vulnerability in Apache Tomcat leads to security constraint bypass if a constraint for a longer path is specified before a more restrictive constraint for a shorter sub-path. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.24, from 10.1.0-M1 through 10.1.57, from 9.0.0.M1 through 9.0.120, from 8.5.0 through 8.5.100, from 7.0.0 through 7.0.109. Users are recommended to upgrade to version 11.0.25, 10.1.58, 9.0.121, which fixes the issue.
CVE-2026-65637	Improper Input Validation vulnerability in Apache Tomcat due to incomplete fix for CVE-2026-32990. This issue affects Apache Tomcat: from 11.0.20 through 11.0.24, from 10.1.53 through 10.1.57, from 9.0.115 through 9.0.120. Users are recommended to upgrade to version 11.0.25, 10.1.58 or 9.0.121, which fix the issue.
CVE-2026-65898	DOMPurify before 3.4.11 fails to clone the ALLOWED_ATTR allowlist when setConfig() is used with an uponSanitizeAttribute hook, allowing the hook to permanently mutate the shared allowlist. Attackers can register a hook that conditionally allows dangerous attributes like onerror for trusted elements, then submit untrusted content that inherits the polluted allowlist and executes event handlers as stored XSS.

CVE ID	Description
CVE-2026-65899	DOMPurify 3.0.0 before 3.4.9 does not reset the retained Trusted Types policy when clearConfig() is called, so a DOMPurify instance reused across trust boundaries stays bound to a previously supplied TRUSTED_TYPES_POLICY. A later caller that requests RETURN_TRUSTED_TYPE output receives a TrustedHTML object created by the old (potentially unsafe) policy rather than a clean default, which can lead to script execution at a Trusted Types sink. Passing TRUSTED_TYPES_POLICY: null on the later call also does not clear the retained policy.
CVE-2026-65900	DOMPurify versions >=3.0.0 and before 3.4.8, when configured with SAFE_FOR_TEMPLATES together with a DOM output mode (RETURN_DOM, RETURN_DOM_FRAGMENT, or IN_PLACE), fail to strip template expressions (e.g. ${evil}, {{evil}}, <%evil%>) inside <template> element content. The final normalization/ scrub pass (_scrubTemplateExpressions) uses a NodeIterator and node.normalize() that do not descend into template.content, so expressions that only form after adjacent text nodes merge survive sanitization. This bypasses SAFE_FOR_TEMPLATES and can allow a downstream template engine to evaluate attacker-supplied expressions. The string output path is not affected.
CVE-2026-65901	DOMPurify through 3.4.6 contains a cross-site scripting vulnerability in IN_PLACE mode that trusts attacker-controlled nodeName on live non-form nodes. Attackers can supply hostile live DOM objects with real script children whose observable nodeName is clobbered to appear as allowed elements, causing scripts to execute when the sanitized tree is inserted into a live document.
CVE-2026-65902	DOMPurify before 3.4.7 passes direct references to the module-level DEFAULT_ALLOWED_TAGS and DEFAULT_ALLOWED_ATTR sets to the uponSanitizeElement and uponSanitizeAttribute hooks via data.allowedTags / data.allowedAttributes when sanitize is called without an explicit cfg.ALLOWED_TAGS / cfg.ALLOWED_ATTR array. A hook that mutates these fields permanently widens the default allow-lists for the lifetime of the DOMPurify instance, so all subsequent default-config sanitize calls inherit the widened defaults and attacker payloads using the poisoned tag/attribute name survive sanitization. removeAllHooks(), clearConfig(), and passing a fresh cfg do not recover the state; only constructing a new DOMPurify instance does.
CVE-2026-65903	DOMPurify before 3.4.0 contains a logic error in the ADD_TAGS function where short-circuit evaluation allows forbidden tags to bypass FORBID_TAGS restrictions. Attackers can craft input containing tags listed in FORBID_TAGS that are also added via ADD_TAGS function, causing them to be retained in sanitized output.
CVE-2026-65905	Authentication Bypass by Capture-replay vulnerability in Apache Tomcat's DIGEST authenticator. If, before windowSize requests have been made, a client makes a DIGEST authenticated request with a nonceCount on the upper boundary of the replay window then that request is replayable once only while the associated nonceCount remains within the replay window. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.24, from 10.1.0-M1 through 10.1.57, from 9.0.0.M1 through 9.0.120. The following versions were EOL at the time the CVE was created but are known to be affected: from 8.5.0 through 8.5.100, from 7.0.30 through 7.0.109. Other unsupported versions may also be affected. Users are recommended to upgrade to version 11.0.25, 10.1.58 or 9.0.121, which fix the issue.
CVE-2026-65912	DOMPurify before 3.3.2 contains a URI validation bypass vulnerability when ADD_ATTR is provided as a predicate function via EXTRA_ELEMENT_HANDLING.attributeCheck. Attackers can supply a predicate that accepts specific attribute and tag combinations to bypass URI-safe validation, allowing unsafe protocols like javascript: to survive sanitization and execute as DOM-based XSS when the link is activated.
CVE-2026-65913	DOMPurify before 3.3.2 contains a prototype pollution vulnerability in USE_PROFILES mode that allows attackers to bypass attribute filtering by polluting Array.prototype properties. Attackers can set Array.prototype properties like onclick to true, causing DOMPurify to accept event handlers as allowlisted attributes and resulting in DOM-based XSS when sanitized markup is rendered.
CVE-2026-65914	DOMPurify before 3.3.2 contains a mutation-XSS vulnerability when sanitized HTML is reinserted into special parsing contexts using innerHTML with wrappers like script, xmp, iframe, noembed, noframes, or noscript. Attackers can craft payloads with closing sequences that break out of the wrapper context during reparsing, reactivating dangerous markup with event handlers to execute JavaScript.
CVE-2026-66010	DOMPurify before 3.4.12 fails to execute afterSanitizeElements hook for custom elements allowed via CUSTOM_ELEMENT_HANDLING.tagNameCheck, allowing attributes to bypass application security policies. Attackers can preserve sensitive attributes on custom elements that later re-inject them into innerHTML sinks, creating second-order XSS gadgets.
CVE-2026-66053	Improper Validation of Certificate with Host Mismatch vulnerability in Apache Thrift Python bindings. This issue affects Apache Thrift: before 0.24.0. Users are recommended to upgrade to version 0.24.0, which fixes the issue. This replaces CVE-2026-41603
CVE-2026-67312	axios versions from 0.28.0 before 0.33.0 and from 1.0.0 before 1.18.0 contain uncontrolled recursion in formDataToJSON (exposed as axios.formToJSON() and used internally when serializing FormData with Content-Type: application/json). When an application passes attacker-controlled FormData field names, a field name with thousands of nested bracket-delimited segments causes unbounded recursion in buildPath(), exhausting the JavaScript call stack (RangeError: Maximum call stack size exceeded) and causing denial of service for that request, or process termination in applications without appropriate error handling.

CVE ID	Description
CVE-2026-67313	axios versions 0.28.0 and later contain uncontrolled recursion in formDataToJSON when processing FormData field names with deeply nested bracket segments. Attackers can supply FormData with field names containing thousands of nested brackets to exhaust the JavaScript call stack and trigger RangeError, causing request failure or process termination in applications that do not handle the exception.
CVE-2026-67315	axios versions 0.31.0 before 0.33.0 and 1.15.0 before 1.18.0 fail to recognize 0.0.0.0 as a loopback address in shouldBypassProxy.js, allowing requests to 0.0.0.0 to bypass NO_PROXY rules. Attackers can supply 0.0.0.0 URLs to route requests through configured proxies, potentially exposing local services when the proxy can reach the destination.
CVE-2026-67316	axios is vulnerable to read-side prototype-pollution gadgets that can alter request construction when Object.prototype has already been polluted by a separate vulnerability or dependency. In the bodyless method aliases (axios.get(), axios.delete(), axios.head(), axios.options()), inherited data is read via (config || {}).data before config normalization, causing an attacker-controlled body to be sent on requests that did not set one. Additional low-level paths, only reachable when calling exported adapters/helpers (e.g. lib/adapters/http.js, unsafe/helpers/resolveConfig.js) directly with plain configs and no own proxy or paramsSerializer, can inherit polluted proxy values (routing requests through an attacker-controlled proxy) or paramsSerializer values (attacker-controlled URL serialization). These low-level gadgets do not reproduce through normal high-level axios calls on 1.15.2+. The issue is fixed in axios 1.18.0 and 0.33.0.
CVE-2026-67317	axios versions 1.7.0 before 1.18.0 fail to enforce maxBodyLength for WHATWG ReadableStream request bodies in the fetch adapter when Content-Length cannot be determined. Attackers can supply unknown-length stream data to bypass upload size limits and cause uncontrolled network egress or resource exhaustion.
CVE-2026-67318	axios versions >=1.13.0 (Node.js HTTP adapter) fail to enforce the configured maxBodyLength limit on streamed request bodies when requests are sent with httpVersion: 2. Because Node's HTTP/2 request API does not honor the maxBodyLength option and axios's byte-counting stream wrapper is gated on maxRedirects === 0, an attacker who controls a stream passed to axios can cause the application to transmit outbound data exceeding the configured finite maxBodyLength. Impact is limited to resource consumption and policy bypass (excess egress, upstream quota consumption, limited availability); it does not enable code execution, credential disclosure, or request-destination control. Calls using the default maxBodyLength: -1 and browser adapters are not affected.
CVE-2026-67319	axios before 0.33.0 (and 1.x before 1.18.0) can consume inherited properties from nested request option objects when the JavaScript process's Object.prototype has already been polluted by another component. While the top-level merged config uses a null prototype, nested plain objects such as auth and paramsSerializer are cloned into ordinary objects and read without own-property checks. When an application passes placeholder nested objects such as auth: {} or paramsSerializer: {}, inherited username/password values can cause silent injection of an Authorization: Basic header, and inherited encode/serialize values can alter query-string serialization (full serializer replacement requires a function-valued pollution primitive). This is exploitable only in the presence of pre-existing prototype pollution.

CVE ID	Description
CVE-2026-68494	The fix released in jackson-core 2.18.6 and 2.21.1 for CVE-2026-18401 (GHSA-72hv-8253-57qq, number length constraint bypass in the non-blocking parser) is incomplete. This record covers the remaining bypass. The earlier fix wired validateIntegerLength() into a new _setIntLength() helper and invoked it wherever the integer portion of a number is decided: a terminator byte arrives, a '.' or 'e'/'E' is seen, or input ends inside a fully buffered value. It was not invoked on the attacker-relevant path where the parser runs out of input while still inside the MINOR_NUMBER_INTEGER_DIGITS minor state and returns NOT_AVAILABLE to the caller. As a result, an attacker who streams JSON to a non-blocking parser in many small chunks, without ever sending a terminator byte, keeps the parser inside MINOR_NUMBER_INTEGER_DIGITS indefinitely. _textBuffer.expandCurrentSegment() grows the accumulator on every chunk while validateIntegerLength() is never called. The accumulator is bounded only by maxStringLength (20 MiB by default) rather than by maxNumberLength (1000 by default), an amplification of roughly 20,000x over the documented limit. Because Java char values occupy two bytes, a single connection can be driven to approximately 40 MiB of heap before the validator finally fires when the value completes. The equivalent fraction-path code is correct: _finishFloatFraction() calls _setFractLength() before its NOT_AVAILABLE return. The missing call affects the integer-digit paths in _startPositiveNumber(), _startNegativeNumber() and _finishNumberIntegralPart() in NonBlockingUtf8JsonParserBase. Impact: reactive frameworks such as Spring WebFlux/Reactor, Quarkus, Helidon and Vert.x feed inbound HTTP or gRPC bytes to the async parser as they arrive, which is precisely the chunked-feed shape required. Operators who set StreamReadConstraints.maxNumberLength expecting it to cap memory per number value do not get that guarantee; memory accumulates per concurrent connection and attacker-controlled concurrency can exhaust the JVM heap. The synchronous parsers (UTF8StreamJsonParser, ReaderBasedJsonParser) and the async parser operating on complete input are not affected. Exploitation requires only the ability to stream data to a parsing endpoint; no privileges or user interaction are needed. This issue affects com.fasterxml.jackson.core:jackson-core from version 2.15.0 through 2.18.7, and from 2.19.0 through 2.21.3, and tools.jackson.core:jackson-core from 3.0.0 through 3.1.3. Versions prior to 2.15.0 are not affected, because StreamReadConstraints -- which defines the maxNumberLength setting -- was first introduced in jackson-core 2.15.0, so no such constraint exists to be bypassed in earlier releases. Note that GHSA-r7wm-3cxj-wff9 states the affected 2.x range without a lower bound. The 2.22.x and 3.2.x release lines are not affected: those branches were created after the fix commit landed on 2026-05-21 and therefore contain it from their initial releases (2.22.0, tagged 2026-06-03, and 3.2.0, tagged 2026-06-08).
CVE-2026-68497	jackson-databind binds a JSON string to a javax.xml.datatype.Duration or javax.xml.datatype.XMLGregorianCalendar field by passing the raw string verbatim to DatatypeFactory.newDuration(value) or newXMLGregorianCalendar(value) in CoreXMLDeserializers.Std._deserialize. These deserializers are registered by default with no opt-in, so a plain ObjectMapper or JsonMapper with no polymorphic typing and no special configuration reaches this path. The XML Schema lexical grammar permits numeric components of arbitrary length, which the JDK materializes through the native BigInteger(String) and BigDecimal(String) constructors, both quadratic in digit count. Because the digits sit inside a JSON string token rather than a JSON number token, jackson-core's StreamReadConstraints.maxNumberLength guard never applies; jackson's own NumberDeserializers call validateIntegerLength or validateFPLength before parsing a stringified number, but the XML datatype deserializer omits that pre-check. An unauthenticated attacker can therefore submit a single request of a few megabytes, such as a Duration value consisting of the letter P followed by several million digits and the letter Y, and force tens of seconds to several minutes of single-threaded CPU work; a handful of concurrent requests can saturate a server's worker threads. This affects com.fasterxml.jackson.core:jackson-databind from 2.0.0 before 2.18.10, from 2.19.0 before 2.21.6, and from 2.22.0 before 2.22.2, and tools.jackson.core:jackson-databind from 3.0.0 before 3.1.6 and from 3.2.0 before 3.2.2. Users should upgrade to 2.18.10, 2.21.6, 2.22.2, 3.1.6, or 3.2.2.
CVE-2026-68525	Incorrect Authorization vulnerability in Apache Tomcat's FORM authentication process allows the bypassing of a security constraint that limits user has access to a resource POST but not GET. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.24, from 10.1.0-M1 through 10.1.57, from 9.0.0.M1 through 9.0.120. The following versions were EOL at the time the CVE was created but are known to be affected: from 8.5.0 through 8.5.100, from 7.0.0 through 7.0.109. Other unsupported versions may also be affected. Users are recommended to upgrade to version 11.0.25, 10.1.58 or 9.0.121, which fixes the issue.
CVE-2026-68770	sentence-transformers contains a security control bypass vulnerability that allows attackers to achieve arbitrary code execution by exploiting a logic flaw in the import_module_class helper within sentence_transformers/util/misc.py, where the guard condition includes an 'or os.path.exists(model_name_or_path)' clause that satisfies the trust gate whenever the supplied path exists on the local filesystem, regardless of the trust_remote_code=False argument. Attackers who can control or influence the contents of a model directory on disk can place malicious Python files such as modeling_*.py referenced via modules.json, causing the code to execute at import time when an application loads the model with SentenceTransformer(path, trust_remote_code=False), bypassing the documented security contract and achieving code execution within the loading process.

CVE ID	Description
CVE-2026-69219	The RabbitMQ Java client library allows Java and JVM-based applications to connect to and interact with RabbitMQ nodes. Prior to 5.33.1, src/main/java/com/rabbitmq/client/impl/ValueReader.java uses ValueReader.readBytes to accept a wire-declared contentLength below Integer.MAX_VALUE and allocate a byte array before checking the bytes available in the frame. A malicious AMQP peer can send a LongString or byte-array field with type tag S and a declared length such as 0x7FFFFFFE during the pre-authentication connection.start server-properties table, causing an approximately 2 GB allocation and OutOfMemoryError before readFully consumes data. The resulting memory exhaustion can terminate the JVM and cause denial of service. This issue is fixed in version 5.33.1.
CVE-2026-69220	The RabbitMQ Java client library allows Java and JVM-based applications to connect to and interact with RabbitMQ nodes. Prior to 5.33.1, src/main/java/com/rabbitmq/client/impl/ValueReader.java permits ValueReader.readTable and ValueReader.readArray to call ValueReader.readFieldValue recursively for AMQP table type F and AMQP array type A values without a nesting-depth limit. A malicious AMQP server or network intermediary can send approximately 580 nested table levels in the pre-authentication connection.start frame, fitting within the default 131072-byte frame maximum, to trigger StackOverflowError. The error terminates the client input processing thread and causes denial of service. This issue is fixed in version 5.33.1.
CVE-2026-69243	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to 3.14.2, the HTTP parsers were vulnerable to a request smuggling attack relating to WebSocket upgrades. If using the server-side component, an attacker may be able to execute a request smuggling vulnerability using an edge case in the WebSocket upgrade procedure. A WebSocket upgrade request with a body could cause the parser to switch protocols before the complete request body was received, leaving trailing bytes to be handled as upgraded-protocol or pipelined data rather than normal HTTP body data. This issue is fixed in version 3.14.2.
CVE-2026-69244	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to 3.14.3, an out-of-bounds heap read could occur in the C response parser while building an error message for a malformed response. An attacker controlled server, or possibly an accidental response, could trigger a DoS in the client. The vulnerable path was error message construction in aiohttp/_http_parser.pyx, where an llhttp error-position pointer was used to build a snippet for malformed chunked responses and malformed request or response bytes at the buffer end. This issue is fixed in version 3.14.3.
CVE-2026-69249	python-cryptography is a package designed to expose cryptographic primitives and recipes to Python developers. In versions 42.0.0 through 48.0.0, when resolving invalid certificate chains that include duplicate copies of self-signed certificates, the processing recursively invokes the same candidate, leading to an exponential blowup. Although the limitation that the chain depth cannot exceed a specified maximum depth prevents unbounded recursion and guarantees termination, an attacker-controlled certificate chain can lead the processing to easily take more than 5s to reject in testing. This amplification could form the basis for a resource exhaustion denial of service attack. The core issue arises in the recursive nature of build_chain_inner, which does not de-duplicate against previously analyzed candidates. As the correctness of validation is not affected, the integrity of a system cannot be compromised through this vector, only its availability. This issue is fixed in 49.0.0.
CVE-2026-71491	sqlparse is a non-validating SQL parser module for Python. Prior to 0.6.0, group_comments in sqlparse/engine/grouping.py repeatedly rescans comment-only statements before the MAX_GROUPING_TOKENS guard, causing quadratic CPU consumption through sqlparse.parse() and sqlparse.format(sql, strip_comments=True). This issue is fixed in version 0.6.0.
CVE-2026-73228	Django REST framework is a toolkit for building Web APIs. Prior to 3.17.2, Django REST Framework's request.data parsing in rest_framework/request.py Request._parse() passes the underlying HttpRequest stream to JSONParser and FormParser for application/json and application/x-www-form-urlencoded bodies, bypassing Django's DATA_UPLOAD_MAX_MEMORY_SIZE protection and allowing oversized request bodies to consume additional memory and CPU. This issue is fixed in version 3.17.2.
CVE-2026-73229	Django REST framework is a powerful and flexible toolkit for building Web APIs. Prior to 3.17.2, Django REST Framework's rest_framework/renderers.py AdminRenderer.render() uses override_method() to simulate GET and directly invokes view.get() without view.check_permissions() while rendering an invalid write request, allowing a 400 Bad Request HTML response to disclose data from a GET representation that the requester is not permitted to access. This issue is fixed in version 3.17.2.
CVE-2026-73508	Netty is an asynchronous, event-driven network application framework. Prior to 4.1.136.Final and 4.2.16.Final, io.netty.handler.codec.dns.AbstractDnsRecord, io.netty.handler.codec.dns.DefaultDnsRecordDecoder.decodeRecord(), and io.netty.handler.codec.dns.DnsCodecUtil.decompressDomainName() failed to release retained or newly allocated ByteBuf objects when IDN.toASCII() or encodeDomainName() rejected a malformed domain name, allowing unauthenticated remote DNS packets to leak direct memory incrementally until denial of service. This issue is fixed in versions 4.1.136.Final and 4.2.16.Final.

CVE ID	Description
CVE-2026-75516	The RabbitMQ Java client library allows Java and JVM-based applications to connect to and interact with RabbitMQ nodes. Prior to 5.34.0, AMQConnection.start() applies Math.min(maxInboundMessageBodySize, frameMax) after Connection.Tune negotiation even though AMQP defines frameMax value zero as unlimited and ConnectionFactory.DEFAULT_FRAME_MAX is zero. When the client default and server-negotiated value are both zero, the result is passed to Utils.framePayloadLimit(int), which interprets zero as Integer.MAX_VALUE and disables the configured maxInboundMessageBodySize cap. A malicious AMQP server, or a man-in-the-middle attacker able to modify Connection.Tune and inject frames into the connection, can then send an oversized frame of any frame type, causing Frame.readFrom() to allocate a large byte array before content-level validation and potentially terminate the client process through memory exhaustion. This issue is fixed in version 5.34.0.
CVE-2026-75595	Netty is an asynchronous, event-driven network application framework. Prior to 4.1.137.Fina and 4.2.17.Final, io.netty.handler.ssl.SslClientHelloHandler#decode checks the wrong offset before reading the four-byte TLS handshake header, so a ClientHello whose handshake header spans records can cause an IndexOutOfBoundsException and invoke select(ctx, null). This selects the default SslContext instead of the SNI-specific context. In deployments where per-SNI clientAuth=REQUIRE is the sole mutual TLS gate, the default SslContext uses clientAuth=NONE or clientAuth=OPTIONAL, and no application-layer certificate verification exists, an unauthenticated remote attacker can bypass the protected route's mutual TLS requirement. This issue is fixed in versions 4.1.137.Final and 4.2.17.Final.
CVE-2026-75596	Netty is an asynchronous, event-driven network application framework. Prior to 4.1.137.Final and 4.2.17.Final, the default io.netty.handler.ssl.SniHandler constructors use the pre-handshake ClientHello aggregation path in handler/src/main/java/io/netty/handler/ssl/SslClientHelloHandler.java at io.netty.handler.ssl.SslClientHelloHandler#decode, where handshakeBuffer.clear() and writeBytes() recopy all previously received body bytes for every additional TLS record. An unauthenticated remote peer can advertise a large ClientHello and deliver its body in thousands of tiny records, causing quadratic CPU work on the event loop before the TLS handshake completes and degrading TLS handling for other clients. This issue is fixed in versions 4.1.137.Final and 4.2.17.Final.
CVE-2026-77310	jackson-databind contains the general-purpose data-binding functionality and tree-model for Jackson Data Processor. Prior to versions 2.18.9, 2.21.5, 2.22.1, 3.1.5, and 3.2.1 on their respective release lines, the java.net.InetAddress branch of FromStringDeserializer.Std._deserialize() calls InetAddress.getByName() on attacker-controlled input, causing eager DNS resolution during deserialization and enabling DNS-based server-side request forgery and internal-host enumeration. This issue is fixed in versions 2.18.9, 2.21.5, 2.22.1, 3.1.5, and 3.2.1.
CVE-2026-77420	JLine is a Java library for handling console input. From 3.0.0 until 3.30.15 and 4.3.1, DefaultHistory.matchPatterns(String patterns, String line) in reader/src/main/java/org/jline/reader/impl/history/DefaultHistory.java converts the HISTORY_IGNORE configuration value into a Java regular expression while escaping only part of its syntax, allowing other regex metacharacters to reach the backtracking engine. An attacker who can control application or user configuration can supply a nested-quantifier expression that is reevaluated whenever a command is added to history, consuming excessive CPU and indefinitely blocking the reader thread. This issue is fixed in versions 3.30.15 and 4.3.1.
CVE-2026-77421	JLine is a Java library for handling console input. From 3.0.0 until 3.30.15 and 4.3.1, the JLine built-in nano editor's regex search mode passes a user-controlled search term from doSearch(String text) in builtins/src/main/java/org/jline/builtins/Nano.java to Java's backtracking regular expression engine without a timeout or backtracking bound. A nested-quantifier expression evaluated against non-matching buffer content can consume excessive CPU and indefinitely block the editor session thread, and remote multi-user deployments can lose a worker thread for each affected session. This issue is fixed in versions 3.30.15 and 4.3.1.
CVE-2026-77422	JLine is a Java library for handling console input. From 3.0.0 until 3.30.15 and 4.3.1, the JLine built-in grep command in builtins/src/main/java/org/jline/builtins/PosixCommands.java accepts a user-controlled regular expression in grep(...) and, unless line-regexp mode is used, automatically adds a dot-star prefix and suffix before compiling it with Java's backtracking regular expression engine. The wrapping expands the backtracking search space, so a short nested-quantifier expression evaluated against non-matching input can consume excessive CPU and indefinitely block a command worker, including in remotely exposed shell sessions. This issue is fixed in versions 3.30.15 and 4.3.1.
CVE-2026-78659	When "Trailer" headers are sent by a client, the HTTP server internally uses the header values to populate the Request.Trailer map passed to the server handler. Because Request.Trailer is a map, each entry incurs memory overhead. For HTTP/2 servers, a malicious client can exploit this by sending a "Trailer" header that declares a large number of fields, causing the server to allocate a disproportionate amount of memory while bypassing Server.MaxHeaderValueCount and Server.MaxHeaderBytes limits. This exploit is not applicable for HTTP/1 servers, which do not support multiplexing a large number of requests over one TCP connection, and whose Server.MaxHeaderBytes are calculated differently.
CVE-2026-78660	Historically, we have been rather lax about malformed framing-related headers in our HTTP/2 implementation, as they cannot interfere with HTTP/2 framing. However, this makes it possible for our HTTP/2 implementation to forward responses containing such headers to an HTTP/1 client when acting as a reverse proxy. If the HTTP/1 client also does not behave strictly enough, this can result in response smuggling.

CVE ID	Description
CVE-2026-78661	When coalescing adjacent text nodes in addText and fosterParent, appending to an existing TextNode's Data via string concatenation (n.Data += text) reallocated and copied the accumulated string on every token. Inputs containing many short text runs separated by ignored end tags (such as "a</x>") caused O(n^2) allocations and byte copies. Accumulate coalesced text in per-node byte slices during parsing and materialize Data strings once at the end of parse.
CVE-2026-78662	Previously, a channel registered in the mux's chanList is not usable until it is established. A malicious peer was able flood the channel's incomingRequests, deadlocking the entire connection. Now, we add an atomic established state, set when a channel becomes usable. Until such a time, handlePacket drops every packet other than the open confirmation/failure, without blocking and without tearing down the connection.
CVE-2026-78663	The HTTP/2 server can refund connection-level flow control twice for the same data: Once when a client resets a stream (refunding data for any sent-but-unread portion of the stream), and again when a request handler reads the buffered data. A malicious client can exploit this to bypass the configured connection-level flow control limit (MaxReceiveBufferPerConnection). Total buffered data is still limited by the concurrent stream limit and stream-level flow control.
CVE-2026-78667	When parsing a Range header containing a large number of small ranges, FileServer(FS), ServeContent, and ServeFile(FS) can consume an excessive amount of CPU.
CVE-2026-78669	A malicious HTTP/2 peer can cause excessive CPU consumption in the client or server by opening a large number of streams and then sending many small SETTINGS frames containing SETTINGS_INITIAL_WINDOW_SIZE values.
CVE-2026-80047	A vulnerability in Hugging Face Transformers (versions 4.57.0 to 5.16.1) allows remote Python files to be written to local disk without user consent when using GenerativePreTrainedModel.load_custom_generate(). The function fetches and caches a remote module file before performing the required trust_remote_code consent check, inverting the security model enforced by other code-loading paths (such as AutoConfig, AutoModel, and AutoTokenizer). As a result, attacker#controlled Python code from custom_generate/generate.py is copied into the user's ~/.cache/huggingface/modules directory even if the user declines the trust prompt. Although execution is correctly gated, the file write is not reversible and can persist across sessions. This can lead to persistent, unauthorized files on disk and stale cache collisions where cached attacker code may later be executed during trusted model loads. The issue stems from an unconditional file write in dynamic_module_utils.py prior to any trust verification.
CVE-2026-81869	OpenTelemetry-Go is the Go implementation of OpenTelemetry. From version 1.10.0 until 1.33.0, the sdk/trace/span.go attribute truncation path can fail to enforce AttributeValueLengthLimit for string and string-slice attributes containing the valid Unicode replacement character U+FFFD. safeTruncateValidUTF8 treats the valid replacement rune as invalid UTF-8 and returns the original input, while strings.ToValidUTF8 leaves that valid rune unchanged, so a second safeTruncate attempt can also return the oversized value. An attacker who controls span attribute content can retain values longer than the configured limit, increasing per-span memory use and weakening denial-of-service protection in the instrumented process. This issue is fixed in version 1.33.0.
CVE-2026-81870	OpenTelemetry-Go is the Go implementation of OpenTelemetry. From version 1.5.0 to 1.44.0, sdk/trace.NewTracerProvider emits a TracerProvider created internal Info-level diagnostic event whose MarshalLog implementations recursively include span processor, exporter, and client configuration. Applications that call otel.SetLogger to enable OpenTelemetry internal Info logging can therefore record OTLP gRPC and HTTP collector endpoints, the OTLP HTTP Insecure flag, and complete Zipkin collector URLs. A person or system with access to those logs can learn internal collector topology and can recover credentials or tokens embedded in Zipkin URL user information or query strings. The default OpenTelemetry logger does not emit the event, and this path does not log OTLP authentication headers, TLS key material, or span payloads. This issue is fixed in version 1.45.0.
CVE-2026-83557	DefaultBaseTypeLimitingValidator is the PolymorphicTypeValidator applied automatically whenever @JsonTypeInfo is used without an explicitly configured custom validator. It denies polymorphic resolution only for a fixed set of "unsafe base types", and its isSafeSubType method returns true unconditionally for every base type outside that set. java.lang.Comparable was absent from the list despite being implemented by a very large fraction of JDK and application classes, comparable in breadth to java.io.Serializable, which is on the list for that reason. An application declaring an @JsonTypeInfo-annotated property or class with Comparable as its base type, and no custom PolymorphicTypeValidator, will accept a type identifier for essentially any class implementing Comparable. This yields an attacker-controlled object instantiation primitive; a demonstrated case constructs a java.io.File for an arbitrary attacker-chosen path, which becomes path-traversal-adjacent if the application subsequently calls path-sensitive methods on the value. No class implementing Comparable has been identified that yields code execution through deserialization alone. Global Default Typing via activateDefaultTyping is not affected, because that method structurally requires an explicit PolymorphicTypeValidator argument. This affects com.fasterxml.jackson.core:jackson-databind from 2.11.0 before 2.18.10, from 2.19.0 before 2.21.6, and from 2.22.0 before 2.22.2, and tools.jackson.core:jackson-databind from 3.0.0 before 3.1.6 and from 3.2.0 before 3.2.2. Users should upgrade to 2.18.10, 2.21.6, 2.22.2, 3.1.6, or 3.2.2.

CVE ID	Description
CVE-2026-84303	gRPC-Go is the Go language implementation of gRPC. Prior to 1.83.1, the xDS RBAC HTTP filter in internal/xds/httpfilter/rbac/rbac.go does not lowercase header matcher names in normalizeHeaderMatcher even though incoming metadata keys are lowercase. A DENY policy using a mixed-case name such as X-Role or User-Agent therefore does not match and fails open, allowing requests that should be rejected. The same case mismatch permits :Scheme or Grpc-Status to evade gRFC A41 validation and prevents Host from being rewritten to :authority. This issue is fixed in version 1.83.1.
CVE-2026-84304	gRPC-Go is the Go language implementation of gRPC. Prior to 1.83.1, internal/transport/transport.go stores each fragmented HTTP/2 DATA frame as a separate recvMsg in recvBuffer, so millions of one-byte frames can consume disproportionate heap memory even when payload bytes remain within connection and stream flow-control windows. An unauthenticated remote attacker can use concurrent multiplexed streams to exhaust process memory and cause a runtime panic or out-of-memory termination. Receive-buffer compaction is enabled by default and can be controlled temporarily with GRPC_GO_EXPERIMENTAL_ENABLE_RECEIVE_BUFFER_COMPACTION. This issue is fixed in version 1.83.1.
CVE-2026-84305	sqlparse is a non-validating SQL parser module for Python. Prior to 0.6.0, sqlparse.format(sql, reindent=True) and sqlformat --reindent route attacker-controlled parenthesized tuple lists through ReindentFilter._get_offset() in sqlparse/filters/reindent.py, where _flatten_up_to_token() repeatedly rebuilds and joins the statement prefix. Thousands of offset calculations walk an expanding token tree, producing quadratic CPU consumption for inputs that remain below MAX_GROUPING_TOKENS and causing request delays, reduced throughput, or worker starvation. This issue is fixed in version 0.6.0.
CVE-2026-84445	gRPC-Go is the Go language implementation of gRPC. Prior to 1.82.2 and 1.83.2, servers created with xds.NewGRPCServer() allow internal/transport/http2_server.go to accept an RPC containing neither the :authority header nor the Host header, while RouteAndProcess in internal/xds/server/routing.go assumes that an authority value exists and indexes the empty slice. A remote client that can complete transport connection establishment can trigger an index-out-of-bounds panic that is not recovered by the per-RPC goroutine and terminates the entire server process. In insecure or ordinary TLS deployments the request can be unauthenticated, while strict mTLS or ALTS deployments require valid transport credentials before the malformed RPC can reach the interceptor. This issue is fixed in versions 1.82.2 and 1.83.2.
CVE-2026-84939	Path traversal vulnerability in Apache FreeMarker template loading mechanism, if the attacker can specify an arbitrary malformed locale identifier to FreeMarker, and the localized lookup configuration setting is enabled (it's by default enabled). This issue affects Apache FreeMarker from 2.2.0 through 2.3.34. Users are recommended to upgrade to version 2.3.35. Disabling localized lookup in previous versions also mitigates this. Note that even in versions affected by this vulnerability, the files that can be loaded remain restricted by the TemplateLoader that FreeMarker is configured to use. In particular, FileTemplateLoader prevents attempts to traverse outside the baseDir specified in its constructor. Other TemplateLoader implementations may allow access outside their designated base directory, but they are still constrained by the underlying storage mechanism—for example, a loader wrapping a Java class loader can only access resources that the class loader can load, while one wrapping a web application context can only access resources available through that context.
CVE-2026-85720	The AsyncHttpClient (AHC) library allows Java applications to easily execute HTTP requests and asynchronously process HTTP responses. From 2.0.0 until 2.16.1 and 3.0.12, a request using an HTTP proxy to reach an HTTPS origin can expose preemptive origin credentials because NettyRequestFactory and NettyRequestSender.sendRequestWithNewChannel attach Authorization to the plaintext CONNECT request before the TLS tunnel exists. Basic or Digest credentials and per-connection NTLM, Kerberos, or SPNEGO tokens intended for the origin are therefore visible to the proxy and to observers on the client-to-proxy hop. The tunneled request still receives origin Authorization after the tunnel is established, while Proxy-Authorization remains on CONNECT for its intended proxy recipient. This issue is fixed in versions 2.16.1 and 3.0.12.
CVE-2026-85721	The AsyncHttpClient (AHC) library allows Java applications to easily execute HTTP requests and asynchronously process HTTP responses. From 2.0.0 until 2.16.1 and 3.0.12, automatic response decompression on the HTTP/1.1 path uses ChannelManager.newHttpContentDecompressor() to install Http1ContentDecompressor without a cumulative output-size limit. A hostile or compromised server, or an attacker who can alter a response in transit, can send a small gzip, deflate, or snappy response that expands across chunks until the client exhausts its heap and raises OutOfMemoryError; brotli and zstd are also affected when their optional codecs are present. In versions 3.0.8 through 3.0.10, the HTTP/2 decompressor is also unbounded, so switching protocols does not mitigate the issue on those releases. A limit applied to each decode call is insufficient because the response can be delivered as many small chunks, so the fixed implementation tracks total decompressed bytes for the whole response. This issue is fixed in versions 2.16.1 and 3.0.12.
CVE-2026-85747	Distribution is a toolkit to pack, ship, store, and deliver container content. Prior to 3.1.2, registry/client.(*manifests).Get() reads a manifest returned for a digest-pinned request and passes it to distribution.UnmarshalManifest without hashing the body against the requested digest, allowing a malicious or intercepted registry response to substitute attacker-chosen manifest content even when Docker-Content-Digest claims the requested value. The bundled pull-through proxy in registry/proxy/proxymanifeststore.go can cache and re-serve the substituted manifest. This issue is fixed in version 3.1.2.

CVE ID	Description
CVE-2026-86071	Junrar is an open source Java RAR archive library. Prior to version 7.6.1, LocalFolderExtractor in src/main/java/com/github/junrar/LocalFolderExtractor.java can create directories outside the intended extraction root when processing a crafted archive entry. LocalFolderExtractor.createFile() validates only the final canonical file path, while LocalFolderExtractor.makeFile() creates intermediate path segments with unchecked mkdir() calls. An entry can therefore make the final path resolve inside the destination while causing intermediate directory creation outside it, enabling filesystem pollution or file-versus-directory squatting that can make later security-sensitive writes fail. The demonstrated impact is directory creation, not unconditional arbitrary file-content write. This issue is fixed in version 7.6.1.
CVE-2026-89407	NumberInput.looksLikeValidNumber() in FasterXML jackson-core pre-validates "stringified numbers" with two regular expressions: PATTERN_FLOAT ([+-]?[0-9]*[\.]?[0-9]+([eE][+-]?[0-9]+)?), present since 2.17.0, and PATTERN_FLOAT_TRAILING_DOT, added in 2.17.2. PATTERN_FLOAT places adjacent quantifiers over the same character class -- an optional [0-9]* run, an optional dot, then a required [0-9]+ run -- so input that ultimately fails to match forces Java's backtracking engine to retry every possible split point of the digit run. Matching cost therefore grows with the square of the input length. An attacker who can supply JSON that an application deserializes into a numeric target type reaches this method through jackson-databind's default String-to-number coercion (StdDeserializer and NumberDeserializers for BigDecimal, BigInteger, Double and Float). Because StreamReadConstraints.maxStringLength defaults to 20,000,000 characters, no constraint bounds the input before it reaches the regex. Testing by the reporter confirmed O(n^2) growth across five consecutive input-size doublings, with a single 160,000-character string consuming roughly 74 seconds in one call; a small number of concurrent requests of ordinary body size can therefore exhaust a server's request-handling thread pool. The affected method does not exist before 2.17.0, so 2.16.x and earlier releases are not affected. The fix replaces both regular expressions with a hand-rolled single-pass scan.
CVE-2026-89425	UTF8DataInputJsonParser._reportInvalidToken() in FasterXML jackson-core builds the offending-token text for its error message by appending Java identifier characters to a StringBuilder in a loop that has no upper bound. Unlike the three sibling parser implementations, including UTF8StreamJsonParser, it never consults ErrorReportConfiguration.getMaxErrorTokenLength() (default 256). A malformed token supplied to a parser created through JsonFactory.createParser(DataInput) is therefore accumulated in full. No StreamReadConstraints setting mitigates this: maxDocumentLength cannot be applied to DataInput sources at all, and maxStringLength does not cover this path because the accumulation bypasses ReadConstrainedTextBuffer. The reporter measured a 20,000,109-character exception message from a 20-million-character malformed token on the DataInput path, against 367 characters for identical input on the InputStream path. Scaling the payload drives the StringBuilder, which also incurs byte-to-char expansion and internal array doubling, to many times the raw payload size and can trigger OutOfMemoryError for the whole JVM. UTF8DataInputJsonParser was introduced in 2.8.0 together with createParser(DataInput); releases before 2.8.0 do not contain the affected class.
CVE-2026-91776	TypeDeserializerBase._findDeserializer() in FasterXML jackson-databind caches the resolved deserializer under the raw, attacker-supplied type ID. When name-based polymorphism is configured with a fallback, for example @JsonTypeInfo(use = Id.NAME, defaultImpl = ...), every distinct unrecognized type ID resolves to the same fallback deserializer but is retained as its own key in the _deserializers map. That map has no configurable bound and lives for the lifetime of the type deserializer, so an attacker who can repeatedly supply fresh unknown type IDs causes monotonic memory retention across requests. The reporter observed 10,000 retained entries from 10,000 distinct unknown IDs, against a single entry for a control that repeated one unknown ID the same number of times, isolating attacker-controlled key cardinality from request volume. Exploitation requires an application that enables name-based polymorphism with a defaultImpl or equivalent fallback, accepts attacker-influenced type IDs, and reuses a long-lived ObjectMapper across requests. The fix stops caching fallback resolutions for unrecognized IDs and bounds both the number of cached entries and the length of a cacheable type ID.
CVE-2026-91777	Forward-reference completion for @JsonIdentityInfo object IDs in FasterXML jackson-databind performs a linear scan of the pending-reference accumulator for every resolved ID. The affected paths are CollectionDeserializer.CollectionReferringAccumulator.resolveForwardReference() and the equivalent implementation in MapDeserializer. When a document first creates N unresolved object-ID references in an identity-enabled collection or map and then defines those same IDs in reverse order, completion performs on the order of N * (N + 1) / 2 identity comparisons, so a shallow document whose size grows linearly causes quadratic CPU work during deserialization. The reporter instrumented equals() calls on the ID class and measured exactly 2,003,000 comparisons at N = 2,000, against zero comparisons in the pending-reference lookup path for an equally sized control in which every reference was already resolved. The input requires no deep nesting and no syntactically unusual JSON. Exploitation requires an application that deserializes attacker-influenced JSON into an identity-enabled collection or map. The fix replaces the repeated linear lookup with a keyed pending-reference structure.
CVE-2026-93451	snappy-java through 1.1.10.8 contains a buffer overflow vulnerability in typed Snappy.uncompress*Array methods that allocate output arrays by dividing uncompressed length by element size but pass the undivided length to native code. Attackers controlling compressed input can cause misaligned length values to write past array bounds with attacker-controlled bytes, corrupting heap memory.

CVE ID	Description
CVE-2026-94439	When an HTTP server handler sends a 2xx response to an HTTP/1 CONNECT request and returns without hijacking the connection, the server improperly continues to read and serve requests from the connection. Since a 2xx response to an HTTP/1 CONNECT converts the connection into a tunnel, the server should not treat the connection as continuing to contain HTTP. The impact of this misbehavior is mostly limited to potential request smuggling, where an intermediate proxy considers the data on the connection to be tunneled and the server considers it to be HTTP.
CVE-2026-94440	Parsing a multipart form can bypass memory limits and read an arbitrarily long line into memory when the remaining limit at the start of a part is less than 400 bytes.
CVE-2026-94448	When a JavaScript template literal contains consecutive expressions, the context tracking state was not properly reset upon entering a new expression. We now ensure that template-literal expression entries correctly reset context variables so all subsequent regular expression literals are accurately recognized and escaped.
CVE-2026-96636	Pillow is a Python imaging library. Prior to 13.0.0, MpoImageFile.seek() and DcxImageFile.seek() reparse later MPO or DCX frames and replace self._size without rerunning Image._decompression_bomb_check(). A crafted multi-frame image can use a tiny first frame to pass Image.open() and then declare an over-limit later frame. Normal frame iteration through seek(), load(), or ImageSequence.Iterator can decode that frame and commit hundreds of megabytes or more, causing memory-exhaustion denial of service. This issue is fixed in version 13.0.0.
CVE-2026-96637	Pillow is a Python imaging library. Prior to 13.0.0, SpiderImageFile.seek() in src/PIL/SpiderImagePlugin.py reparses a later SPIDER stack frame and replaces self._size without rerunning Image._decompression_bomb_check(). A tiny file can use a safe first frame to pass Image.open() and then declare very large dimensions in a later mode F frame. Loading that frame allocates approximately four bytes per attacker-declared pixel and can consume multiple gigabytes of memory. This issue is fixed in version 13.0.0.
CVE-2026-96638	Pillow is a Python imaging library. From 9.1.0 until 13.0.0, BmpRleDecoder.decode() in src/PIL/BmpImagePlugin.py permits crafted RLE BMP delta escapes and absolute-mode runs to append beyond the remaining dest_length without a per-branch cap. A delta can add roughly 255 image rows in one iteration, so a kilobyte-sized file with a large but permitted width can expand the data bytearray to hundreds of megabytes or more, causing memory-exhaustion denial of service. This issue is fixed in version 13.0.0.
CVE-2026-96639	Pillow is a Python imaging library. From 11.0.0 until 13.0.0, multi-frame loaders can retain an ImagingCore allocation created for a narrow frame mode while a later frame of the same dimensions changes the logical mode and decoder tiles to wider RGB or CMYK data. ImageFile.load_prepare() reuses the stale native storage because _im is non-null without verifying that its mode and size still match. Crafted PSD, MPO, DCX, or MIC frame transitions can then write wider pixels into the narrow allocation, causing native out-of-bounds writes and crashes. The MPO path can also cause a native over-read and, in demonstrated runtime-specific profiles, command execution. This issue is fixed in version 13.0.0.
CVE-2026-96640	Pillow is a Python imaging library. From 5.1.0 until 13.0.0, PdfParser.linearize_page_tree() in src/PIL/PdfParser.py treats a PDF /Pages graph as a tree without tracking duplicate indirect references. When Image.save(..., format="PDF", append=True) processes an attacker-controlled existing PDF whose /Kids arrays repeatedly reference the same child /Pages node, Pillow expands the shared subtree as distinct pages. A compact malformed PDF can therefore create exponentially many page references and consume disproportionate CPU before the append completes. This issue is fixed in version 13.0.0.
CVE-2026-96641	Pillow is a Python imaging library. Prior to 13.0.0, PILfont text-width accumulation in src/_imaging.c checks positive glyph advances against INT_MAX but does not check negative advances against INT_MIN. Repeating a signed 16-bit advance of -32768 can therefore execute a signed C integer overflow in xsize += dx, producing a large positive width from a negative mathematical result. Applications that use getlength(), getsize(), or derived layout, quota, serialization, or allocation calculations can experience logic-integrity errors or resource consumption, while the high-level mask path commonly retains a decompression-bomb check. This issue is fixed in version 13.0.0.
CVE-2026-96642	Pillow is a Python imaging library. Prior to 13.0.0, the TarIO initializer in src/PIL/TarIO.py accepts a negative octal member size and computes a backward seek that returns to the same 512-byte header. When the requested member name does not match, the parser repeatedly reads that header, consumes a full CPU core, and never returns. TarIO is reached only when an application directly uses PIL.TarIO.TarIO rather than the default Image.open() path. This issue is fixed in version 13.0.0.
CVE-2026-96643	Pillow is a Python imaging library. From 11.0.0 until 13.0.0, Exif.load() in src/PIL/Image.py repeatedly slices the remaining EXIF data while it begins with the six-byte Exif header. src/PIL/PngImagePlugin.py accepts an uncapped PNG eXIf chunk, so a file filled with repeated headers causes a new shrinking copy on every iteration and quadratic CPU work. The condition is reachable through Image.getexif(), Image._getexif(), and ImageOps.exif_transpose(), including common photo-upload rotation workflows. This issue is fixed in version 13.0.0.

CVE ID	Description
CVE-2026-96644	Pillow is a Python imaging library. Prior to 13.0.0 on Windows, src/PIL/WmfImagePlugin.py passes unchecked 32-bit WMF or EMF dimensions to the native drawwmf handler in src/display.c. The handler sizes a DIB allocation from 16-bit-truncated bcWidth and bcHeight fields but copies pixel data using the original full dimensions. A crafted wide, short metafile can pass the pixel-count decompression-bomb check and make the Windows-only renderer read beyond the heap buffer, potentially disclosing memory or crashing the process. This issue is fixed in version 13.0.0.
CVE-2026-96778	Pillow is a Python imaging library. From 9.4.0 until 13.0.0, ImagingJpegEncode() in src/libImaging/JpegEncode.c can call jpeg_write_marker() for a JPEG COM comment when an ICC profile has exactly filled the 65536-byte encoder buffer. libjpeg writes one byte before checking the exhausted free_in_buffer counter, which then wraps and permits the attacker-controlled comment and subsequent compressed image data to overwrite adjacent heap memory. Re-saving an untrusted image as JPEG while preserving its ICC profile and forwarded comment can crash the process and may enable code execution in suitable runtime and heap layouts. This issue is fixed in version 13.0.0.
CVE-2026-97030	A trusted template author may have previously written a valid template wherein the use of the 'yield' keyword would not be correctly escaped. We now ensure that valid keyword uses are escaped and non-keyword uses are not escaped.
CVE-2026-97031	Multiple ECH outer extension references are not permitted under RFC 9849; previously, a client could send a well-crafted packet that could trigger memory exhaustion in the server process by specifying multiple references. We now reject these as malformed and curb the memory amplification vector as a result.
CVE-2026-97032	HTTP/2 servers could end up crashing due to inadvertently modifying its HPACK encoder concurrently. This happens because the server modifies the HPACK encoder from two goroutines without synchronization: one uses the encoder to encode a HEADERS frame as part of a response sent to a client and the other modifies the encoder's table size when handling a SETTINGS frame containing SETTINGS_HEADER_TABLE_SIZE that a client sends. A malicious client can repeatedly send a request while changing the header table size to crash the server.
CVE-2026-97045	In childTextNodesAreLiteral, walking up the tree previously stopped at the nearest foreign ancestor once an HTML integration point was found. However, that integration point may itself be a descendant of an element that was foster-parented into a foreign namespace (such as <mglyph> or <malignmark> fostered under <mtext>), causing a parser to resolve the subtree in a different foreign namespace upon re-parsing the rendered output. Continue walking the full ancestor chain and verify that every transition from a non-empty namespace occurs at an HTML integration point before emitting literal text.
CVE-2026-97370	Junrar is an open source Java RAR archive library. Prior to 7.6.1, RarVM.ExecuteCode in src/main/java/com/github/junrar/unpack/vm/RarVM.java can repeatedly execute a VM_JMP instruction whose target is greater than or equal to the prepared program's codeSize. The setIP() function returns success without changing the instruction pointer or decrementing maxOpCount, and the jump handler continues before loop-end accounting, leaving both values unchanged. A crafted RAR4 archive can therefore keep an extraction thread in an infinite CPU loop and exhaust archive-processing worker capacity. The RAR virtual machine does not expose native process execution, so the supported impact is availability loss rather than code execution. This issue is fixed in version 7.6.1.
CVE-2026-97687	urllib3 is an HTTP client library for Python. From 1.26.0 until 2.8.0, the proxy_ssl_context, proxy_assert_hostname, proxy_assert_fingerprint, ssl_context, cert_reqs, verify_mode, use_forwarding_for_https=True, and CERT_NONE configuration paths fail to remain separated because target-server TLS settings are incorrectly applied to the HTTPS proxy connection. The trigger is that an application uses an HTTPS proxy and configures target-server TLS settings that must remain separate from the proxy TLS handshake, including HTTPS forwarding with target-specific identity or credentials. Applying cert_reqs=CERT_NONE can overwrite proxy_ssl_context.verify_mode in place, and the mutation persists so later connections reusing the same context may connect to the HTTPS proxy without certificate verification. The attack mechanism is that an attacker intercepts and impersonates the HTTPS proxy after the effective proxy policy accepts the attacker's certificate. The impact is that the attacker can observe or modify forwarded traffic or receive a target TLS client certificate, while CONNECT tunneling still preserves the separate end-to-end target TLS connection. This issue is fixed in version 2.8.0.
CVE-2026-97688	urllib3 is an HTTP client library for Python. From 2.6.2 until 2.8.0, HTTPResponse.stream and HTTPResponse.read_chunked can enter an infinite loop because the Deflate decoder retains trailing bytes as unconsumed input after reaching end-of-stream and repeatedly decodes them without progress. The issue occurs when an untrusted server sends a chunked Deflate response whose decoded body exceeds a positive finite chunk size and whose encoded body has trailing bytes, specifically a response with Transfer-Encoding: chunked and Content-Encoding: deflate, content decoding enabled, and the positive finite amt=N streaming chunk size. The attack mechanism is that a malicious server returns a compressed chunked response with trailing bytes after the Deflate stream. The impact is excessive CPU usage and a request that does not complete, and network read timeouts do not interrupt the loop because no further socket read occurs. This issue is fixed in version 2.8.0.

CVE ID	Description
CVE-2026-97689	urllib3 is an HTTP client library for Python. From 1.10.3 until 2.8.0, the HTTPResponse.read_chunked and HTTPResponse.stream methods can allocate unbounded memory because the streaming chunk parser buffers the chunk-size field until newline or EOF without a length bound. The trigger is that a malicious server returns Transfer-Encoding: chunked followed by a very long run of bytes without a newline. The attack mechanism is that a malicious HTTP server sends a very long unterminated chunk-size line. The impact is that unbounded memory allocation can exhaust the client process. This issue is fixed in version 2.8.0.
CVE-2026-101898	Axios is a promise-based HTTP client for the browser and Node.js. From 1.13.0 until 1.20.0, Axios HTTP/2 request setup does not consistently apply proxy settings and caller-supplied DNS lookup policy. An HTTPS request uses httpVersion: 2 with explicit config.proxy or environment-derived proxy settings, or relies on caller-supplied config.lookup DNS policy. The HTTP/2 path can connect without the configured proxy behavior or without applying the caller-supplied config.lookup policy before http2.connect(). Requests can bypass the intended proxy route or the caller-supplied DNS resolution policy. This issue is fixed in version 1.20.0.
CVE-2026-101899	Axios is a promise-based HTTP client for the browser and Node.js. From 1.15.0 until 1.20.0, Axios shouldBypassProxy treats the NO_PROXY CIDR form as a textual hostname suffix instead of evaluating the network range. A request targets an address covered by a CIDR entry in NO_PROXY while a proxy is configured. The CIDR bypass rule is not matched, so the request is sent through the configured proxy. For plaintext HTTP, the proxy can observe and modify URLs, headers, and bodies; for HTTPS, it can observe connection metadata and CONNECT requests. This issue is fixed in version 1.20.0.
CVE-2026-101900	Axios is a promise-based HTTP client for the browser and Node.js. From 1.12.0 until 1.20.0, ResolveConfig reads inherited Symbol.toStringTag, append, and getHeaders properties while resolving FormData headers. A separate same-process prototype-pollution flaw supplies an array or non-plain class instance whose inherited properties make it appear FormData-like; plain objects are blocked. The inherited getHeaders function can return attacker-controlled headers that resolveConfig merges into a fetch adapter request. Attacker-controlled headers can alter authorization, cache, metadata-service, or application-specific request behavior. This issue is fixed in version 1.20.0.
CVE-2026-101901	Axios is a promise-based HTTP client for the browser and Node.js. From 1.13.0 until 1.20.0, Http2Sessions does not install adequate error handling for a ClientHttp2Session during Axios HTTP/2 session initialization or reuse. A request uses httpVersion: 2 and the ClientHttp2Session emits an error during session initialization or reuse. The unhandled session error escapes normal Promise rejection handling. The uncaught error can terminate the Node.js process and cause denial of service. This issue is fixed in version 1.20.0.
CVE-2026-101902	Axios is a promise-based HTTP client for the browser and Node.js. From 0.27.2 until 0.34.0 and 1.20.0, Axios default-instance requests that omit an explicit method can read an inherited method value from Object.prototype. If another vulnerability in the same process pollutes Object.prototype.method, calls such as axios.request({ url }) and axios({ url }) can send a state-changing HTTP method instead of the expected default GET. Axios does not create the prototype pollution source. This is a read-side gadget in axios request dispatch. This issue is fixed in version 0.34.0 and 1.20.0.
CVE-2026-101904	Axios is a promise-based HTTP client for the browser and Node.js. From 1.0.0 until 1.20.0, the dispatchRequest function normalizes inherited Object.prototype.headers from a replacement request configuration. A separate same-process prototype-pollution flaw sets Object.prototype.headers, and trusted request interceptors return a new ordinary configuration without an own headers property. After the interceptor chain, dispatchRequest resolves the inherited headers during normalization. Downstream request processing can observe attacker-controlled headers, including authorization-related values. This issue is fixed in version 1.20.0.
CVE-2026-101906	Axios is a promise-based HTTP client for the browser and Node.js. From 1.15.0 until 1.20.0, Axios shouldBypassProxy applies a quadratic trailing-dot regular expression to redirect hostnames. HTTP_PROXY or HTTPS_PROXY is configured, NO_PROXY or no_proxy is non-empty, redirects are followed, and a crafted redirect Location contains many dots followed by a non-dot character. Hostname.replace(/.+$/, '') backtracks quadratically while processing the crafted redirect hostname. Synchronous regular-expression processing can block the Node.js event loop and cause denial of service. This issue is fixed in version 1.20.0.
CVE-2026-101908	Axios is a promise-based HTTP client for the browser and Node.js. From 1.7.0 until 1.20.0, the fetch adapter constructs a Request with sanitized resolvedOptions but then calls fetch with the original fetchOptions. A separate same-process prototype-pollution flaw populates Object.prototype.headers so fetchOptions.headers resolves through inheritance. The inherited fetchOptions.headers value overrides the sanitized Request headers through the second argument to fetch after Request construction. Attacker-controlled request headers can alter authorization, caching, metadata-service access, or application-specific behavior. This issue is fixed in version 1.20.0.
CVE-2026-101917	PyJWT is a Python implementation of JSON Web Token standards. Prior to 2.14.0, PyJWT get_signing_key_from_jwt is affected because unknown kid misses force refreshes without a negative cache or minimum refresh interval. This occurs when unauthenticated tokens repeatedly use the same unknown kid or varying kid values absent from the cached JWKS. As a result, each cache miss causes PyJWKClient to refresh the JWKS. Consequently, attacker traffic can amplify outbound requests to the configured JWKS endpoint. This issue is fixed in version 2.14.0.

CVE ID	Description
CVE-2026-101918	PyJWT is a Python implementation of JSON Web Token standards. From 2.0.0a1 until 2.15.0, PyJWT PyJWKClient.get_signing_key_from_jwt is affected because payload parser catches ValueError but not RecursionError. This occurs when an attacker-controlled recursively nested payload reaches json.loads. As a result, documented PyJWT exception handling does not contain the failure. Consequently, an unauthenticated request can raise an exception that may produce an HTTP 500 response. The advisory-defined affected implementation also includes jwt/api_jwt.py, verify_signature=False. This issue is fixed in version 2.15.0.
CVE-2026-102265	PyJWT is a Python implementation of JSON Web Token standards. From 2.13.0 until 2.14.0, PyJWS._load in jwt/api_jws.py is affected because parser catches ValueError but not RecursionError. This occurs when a deeply nested token header reaches json.loads. As a result, RecursionError escapes the documented PyJWT error hierarchy. Consequently, an unauthenticated malformed token can cause a request-level failure and HTTP 500. This issue is fixed in version 2.14.0.
CVE-2026-102266	PyJWT is a Python implementation of JSON Web Token standards. From 2.13.0 until 2.14.0, HMACAlgorithm.from_jwk is affected because PyJWK verification path used the decoded key without applying prepare_key validation. This occurs when a trusted JWK Set contains an oct entry with an empty k value. As a result, an attacker signs an HMAC token with the same zero-length key accepted by PyJWT. Consequently, forged token can carry arbitrary authenticated claims. This issue is fixed in version 2.14.0.
CVE-2026-102267	PyJWT is a Python implementation of JSON Web Token standards. Prior to 2.14.0, PyJWT PyJWKClient is affected because redirect destinations are not revalidated against the JWKS trust boundary. This occurs when a configured trusted JWKS endpoint returns an attacker-influenced redirect. As a result, PyJWKClient follows the redirect and consumes the redirected response as key material. Consequently, forwarded credentials may be disclosed or verification keys may be substituted. This issue is fixed in version 2.14.0.
CVE-2026-102268	PyJWT is a Python implementation of JSON Web Token standards. Prior to 2.14.0, is_pem_format in jwt/utils.py is affected because is_pem_format does not recognize every PEM representation accepted by the cryptography loader. This occurs when an application mixes HMAC and asymmetric algorithms and supplies a mutated public-key PEM as raw key bytes. As a result, HMACAlgorithm.prepare_key treats the unrecognized asymmetric public key as an HMAC secret. Consequently, an attacker who knows the public key can forge authenticated HMAC tokens. This issue is fixed in version 2.14.0.
CVE-2026-102269	PyJWT is a Python implementation of JSON Web Token standards. Prior to 2.14.0, PyJWT signature segment is affected because signature segment decoding accepts characters outside the canonical Base64URL representation. This occurs when non-Base64URL characters are appended to a valid compact JWS signature segment. As a result, base64url_decode produces the same signature bytes for different serialized segments. Consequently, raw-token revocation checks can fail to recognize an equivalent modified token. This issue is fixed in version 2.14.0.
CVE-2026-102270	PyJWT is a Python implementation of JSON Web Token standards. Prior to 2.14.0, PyJWT is_pem_format is affected because lazy PEM regular expression backtracks extensively. This occurs when a certificate-like input contains repeated BEGIN markers without a matching END marker. As a result, is_pem_format performs unbounded backtracking while searching for a PEM end marker. Consequently, an attacker can cause intensive CPU consumption. This issue is fixed in version 2.14.0.
CVE-2026-102271	PyJWT is a Python implementation of JSON Web Token standards. From 2.4.0 until 2.14.0, PyJWT HMACAlgorithm.prepare_key is affected because asymmetric-key guard relies on textual markers that are absent from DER encoding. This occurs when an application mixes HMAC and asymmetric algorithms and supplies a DER public key as the shared verification key. As a result, PyJWT uses public DER bytes as an HMAC secret. Consequently, an attacker who knows the public key can forge authenticated HMAC tokens. This issue is fixed in version 2.14.0.
CVE-2026-102272	PyJWT is a Python implementation of JSON Web Token standards. From 2.13.0 until 2.14.0, HMACAlgorithm.prepare_key in jwt/algorithms.py is affected because raw-JWK detector does not normalize accepted Unicode byte-order marks before checking for JSON. This occurs when a public JWK is prefixed with a UTF-8 BOM and used in a mixed-algorithm verification path. As a result, public JWK bypasses asymmetric-key detection and becomes the HMAC secret. Consequently, an attacker who knows the public key can forge authenticated tokens. This issue is fixed in version 2.14.0.
CVE-2026-102273	PyJWT is a Python implementation of JSON Web Token standards. From 2.13.0 until 2.14.0, PyJWT HMACAlgorithm.prepare_key is affected because HMAC key guard only recognizes top-level public JWK forms and misses container representations. This occurs when an application allows HMAC and asymmetric algorithms and passes a public JWK container as the raw key. As a result, public asymmetric key material is accepted as the HMAC secret. Consequently, an attacker who knows the public key can forge a token with arbitrary authenticated claims. This issue is fixed in version 2.14.0.
CVE-2026-102274	PyJWT is a Python implementation of JSON Web Token standards. From 2.9.0 until 2.14.0, PyJWKSet does not catch the plain ValueError raised for malformed RSA JWK components by RSAAlgorithm.from_jwk in jwt/api_jwk.py. This occurs when a JWK Set contains a malformed RSA key alongside otherwise usable keys. As a result, one malformed member aborts construction of the entire PyJWKSet. Consequently, applications can experience authentication failures or request-level denial of service. This issue is fixed in version 2.14.0.

CVE ID	Description
CVE-2026-102275	PyJWT is a Python implementation of JSON Web Token standards. From 2.1.0 until 2.15.0, PyJWT OKPAlgorithm.from_jwk in jwt/algorithms.py is affected because private-JWK import path does not compare the public key derived from d with x. This occurs when an OKP private JWK supplies non-corresponding x and d components. As a result, identity derived from x can differ from operations performed with d. Consequently, if an integration also accepts private key parameters from a proof header without rejecting them, an attacker may use a stolen sender-constrained token without the legitimate private key. This issue is fixed in version 2.15.0.
CVE-2026-102554	Allocation of resources without limits or throttling (CWE-770) during Java object deserialization in Google Guava versions 4.0 through 33.7.1 allows an attacker to cause a Denial of Service via OutOfMemoryError. When deserializing CompactHashMap, CompactHashSet, or MapMakerInternalMap instances, Guava eagerly allocates an array based on a caller-specified size parameter without throttling, permitting memory exhaustion from crafted serialization streams.
CVE-2026-102925	virtualenv is a tool for creating isolated virtual python environments. Prior to 21.7.13, the generated activate (bash and zsh) and activate.fish scripts place values already escaped by shlex.quote inside an additional quoted context. In the bash and zsh script, a crafted virtual environment path reaches __VIRTUAL_ENV__ when a relocated environment's recorded directory is absent; in the fish script, crafted Tcl or Tk library paths reach __TCL_LIBRARY__ or __TK_LIBRARY__. The surplus quotes can terminate the data-only quoted run and leave shell metacharacters parsed as commands when a user sources the activation script, allowing code execution with that user's privileges. This issue is fixed in version 21.7.13.
CVE-2026-102930	virtualenv is a tool for creating isolated virtual python environments. Prior to 21.7.12, download_wheel() accepts pip and setuptools seed wheels fetched for periodic updates or the --download option without checking their bytes against an authoritative digest equivalent to the embedded wheels' BUNDLE_SHA256 verification. A compromised index, stale mirror, or intercepted TLS connection can substitute a different wheel under the requested distribution, version, and filename, after which virtualenv caches and seeds the attacker-controlled wheel into subsequently created environments. The verification applies to the default PyPI path and is intentionally skipped when PIP_INDEX_URL, PIP_EXTRA_INDEX_URL, or PIP_INDEX configures a custom index that may legitimately publish rebuilt wheels. This issue is fixed in version 21.7.12.
CVE-2026-102937	virtualenv is a tool for creating isolated virtual python environments. Prior to 21.7.12, BatchActivator.quote() returns prompt text unchanged before activate.bat inserts it into a cmd.exe set "VAR=value" statement. An attacker who influences --prompt, VIRTUALENV_PROMPT, or the corresponding configuration value can include a double quote that closes the assignment and leaves following cmd.exe operators as executable syntax. When a user activates the generated Windows environment, the injected commands run with that user's privileges. This issue is fixed in version 21.7.12.
CVE-2026-102938	virtualenv is a tool for creating isolated virtual python environments. Prior to 21.7.11, PyEnvCfg.write() writes prompt values verbatim to the line-oriented pyvenv.cfg format while PyEnvCfg._read_values() parses the file with str.splitlines() and accepts the last value for duplicate keys. An attacker who influences --prompt, VIRTUALENV_PROMPT, or configuration input can insert a recognized line boundary and additional keys, including home, causing consumers to use an attacker-selected base interpreter or corrupted environment metadata. The security impact requires prompt input from outside the operator's trust boundary; directly supplied prompt content primarily corrupts the operator's own environment. This issue is fixed in version 21.7.11.
CVE-2026-102991	Mako is a template library written in Python. Prior to 1.4.2, on Windows, TemplateLookup.get_template() in mako/lookup.py resolves template URIs with posixpath, while Template.__init__() in mako/template.py validates them with os.path, which uses ntpath. A URI beginning with a drive designator causes ntpath to absorb the traversal segments before the leading dot-dot check, while posixpath resolution can escape the configured template directory. An application that passes attacker-controlled template names or include paths can disclose process-readable files on the same volume, and a targeted file containing Mako template syntax may also be parsed and executed as a template. Raw URL paths are generally normalized before reaching this form, but query strings, form or JSON bodies, route parameters, and dynamic include expressions can preserve it. This issue is fixed in version 1.4.2.
CVE-2026-103001	PyJWT is a Python implementation of JSON Web Token standards. From 2.11.0 through 2.13.0, PyJWT's PyJWT._merge_options() method can modify a caller-supplied mutable options mapping when verify_signature is false. If an application reuses that same mapping for a later decode() or decode_complete() call and changes verify_signature to true, the mapping can retain false values for expiration, not-before, issued-at, audience, issuer, subject, and JWT ID checks. A signed token with invalid registered claims can then be accepted without disabling signature verification, but applications that create a fresh options mapping for each call are not affected.
CVE-2026-104851	fsspec is a specification and Python implementation framework for filesystem interfaces. From 0.9.0 until 2026.6.0, fsspec.implementations.reference.ReferenceFileSystem evaluates fields from Kerchunk reference JSON documents through unrestricted jinja2.Template(...).render(...) calls in _process_references1._render_jinja, _process_templates, and _process_gen in fsspec/implementations/reference.py. A document supplied inline or fetched from an attacker-controlled URL can provide template expressions that execute Python code when the reference filesystem is opened, including through consumers such as xarray, before referenced data is read. The _process_gen path is reached whenever a document includes a gen array, while the other paths depend on template-related options and values. This issue is fixed in version 2026.6.0.

CVE ID	Description
CVE-2026-104874	Multidict is an implementation of a multidict data structure. From 6.7.0 until 6.9.1, the C extension's items-view reflected union operation, operand | d.items(), in multidict_itemsview_or2_impl and subtraction operation, d.items() - operand, in multidict_itemsview_sub1_impl fail to release new key-identity and value references returned for each operand element. Applications that perform these operations over attacker-influenced sequences can leak two strong references per element, and garbage collection cannot reclaim them, so repeated operations can cause unbounded process memory growth and denial of service. Forward union, intersection, non-tuple operand elements, and pure-Python builds are not affected by this reference leak. This issue is fixed in version 6.9.1.
CVE-2026-106121	The RabbitMQ Java client library allows Java and JVM-based applications to connect to and interact with RabbitMQ nodes. Prior to 5.37.0, com.rabbitmq.tools.json.JSONReader.read() fails to terminate when input ends inside a quoted string or a line comment because its string and whitespace scanners do not stop at CharacterIterator.DONE. The default DefaultJsonRpcMapper passes JSON-RPC message bodies to this parser for JsonRpcServer and client replies. A truncated string causes the parser to append replacement end markers until heap exhaustion, while a line comment without a terminating newline can keep a thread consuming CPU indefinitely, resulting in denial of service. This issue is fixed in version 5.37.0.
CVE-2026-106122	The RabbitMQ Java client library allows Java and JVM-based applications to connect to and interact with RabbitMQ nodes. Prior to 5.36.0, ValueReader.readShortstr decodes malformed UTF-8 bytes into replacement characters that can re-encode beyond the AMQP shortstr limit enforced by ValueWriter.writeShortstr. An attacker who can submit an RPC message with a malformed echoed property can cause reply publication in RpcServer.mainloop() or tutorial-style consumers to throw an unchecked exception before acknowledgement. The broker requeues the message, allowing the same message to disable replacement consumers until the queue is purged. This issue is fixed in version 5.36.0.
CVE-2026-106123	The RabbitMQ Java client library allows Java and JVM-based applications to connect to and interact with RabbitMQ nodes. Prior to 5.35.0, ConnectionFactoryConfigurator.load() includes the raw uri value in wrapped exceptions when AMQP URI parsing fails. Because the URI may contain a plaintext username and password, startup logs, application performance monitoring systems, CI logs, and copied stack traces can disclose broker credentials to users who should not have access to them. This issue is fixed in version 5.35.0.
CVE-2026-106449	yawkat LZ4 Java provides LZ4 compression for Java. Prior to 1.11.4, net.jpountz.lz4.LZ4BlockInputStream configured with stopOnEmptyBlock set to false handles each well-formed empty LZ4Block by recursively calling refill(), allowing a long sequence of empty blocks in an attacker-controlled compressed stream to exhaust the decoding thread's stack and throw StackOverflowError. The default stopOnEmptyBlock setting is true and is not affected, and the issue does not cause memory corruption. This issue is fixed in version 1.11.4.
CVE-2026-106450	yawkat LZ4 Java provides LZ4 compression for Java. Prior to 1.11.4, net.jpountz.lz4.LZ4FrameInputStream readHeader() allocates two new 4 MiB block buffers whenever a maximum-block-size frame header is read, and the default concatenated-frame mode allows attacker-controlled streams containing many minimal empty frames to trigger roughly 8 MiB of allocation for every 11 input bytes. The stream produces no decompressed output while consuming CPU and garbage-collection time, so decompressed-size limits do not mitigate the issue; readSingleFrame mode is not affected. This issue is fixed in version 1.11.4.
CVE-2026-106451	yawkat LZ4 Java provides LZ4 compression for Java. From 1.7.0 until 1.11.4, net.jpountz.util.Native.load() uses File.createTempFile to create an exclusive temporary .lck file but derives the native-library path by removing the suffix, then FileOutputStream opens that predictable path without exclusive creation, allowing another local user with access to the same shared temporary directory to create or replace the library file before System.load() uses it. Successful exploitation depends on shared-directory permissions, host protections, and winning the race, and can execute native code as the victim; hardened systems may instead cause library loading to fail and fall back to Java implementations. Configurations using a system library, a private java.io.tmpdir, or Java-only implementations are not affected. This issue is fixed in version 1.11.4.
CVE-2026-106452	yawkat LZ4 Java provides LZ4 compression for Java. Prior to 1.11.2, net.jpountz.lz4.LZ4BlockInputStream refill() validates that the compressedLen field in a legacy LZ4Block header is nonnegative but allocates a compressed-input buffer of that attacker-controlled size before reading payload data, allowing a header-only stream to request a near-2 GiB allocation and exhaust the JVM heap. Canonical writers emit raw blocks when compression is not smaller than the original block, but vulnerable readers accept non-canonical oversized compressed blocks. This issue is fixed in version 1.11.2.
CVE-2026-106453	yawkat LZ4 Java provides LZ4 compression for Java. Prior to 1.11.2, LZ4DecompressorWithLength uses getDecompressedLength to trust the four-byte decompressed-length header before validating the compressed input, allowing a five-byte attacker-supplied input whose header declares a large output size to request up to approximately 2 GiB and exhaust the JVM heap. Convenience overloads backed by LZ4FastDecompressor or LZ4SafeDecompressor allocate the untrusted size, while overloads that write to a caller-provided destination buffer are not affected because the caller controls the destination size. This issue is fixed in version 1.11.2.

CVE ID	Description
CVE-2026-107226	The AsyncHttpClient (AHC) library allows Java applications to easily execute HTTP requests and asynchronously process HTTP responses. From 2.1.0 until 3.0.14, the default ThreadSafeCookieStore does not retain whether a Set-Cookie response used a secure transport. It accepts Secure cookies received over plaintext and permits insecure cookies to overlay or delete existing Secure cookies with matching scope. An attacker controlling a plaintext same-site response can consequently plant, replace, or remove cookie state later sent to an HTTPS origin, enabling session fixation, CSRF-token replacement, or loss of security-relevant state. This issue is fixed in version 3.0.14.
CVE-2026-107227	The AsyncHttpClient (AHC) library allows Java applications to easily execute HTTP requests and asynchronously process HTTP responses. From 2.2.0 until 3.0.14, WebSocket permessage-deflate decompression is unbounded when compression is enabled. The inbound pipeline aggregates compressed frames before WebSocketClientCompressionHandler inflates them, so webSocketMaxFrameSize and webSocketMaxBufferSize do not bound decompressed output. A malicious WebSocket peer can send a small compressed message that expands to a very large Netty buffer and exhausts JVM heap. This issue is fixed in version 3.0.14.
CVE-2026-107228	The AsyncHttpClient (AHC) library allows Java applications to easily execute HTTP requests and asynchronously process HTTP responses. From 2.1.0 until 3.0.14, the enabled-by-default cookie store replaces a Cookie header explicitly supplied through setHeader or addHeader whenever the store contributes any cookie for the origin. In a shared client, stored cookies originating from one user can replace a different user's request cookie, causing the request to execute under the wrong session. This bypasses the earlier CVE-2024-53990 remediation, which covered cookies supplied through addCookie but not a directly supplied header. This issue is fixed in version 3.0.14.
CVE-2026-107230	The AsyncHttpClient (AHC) library allows Java applications to easily execute HTTP requests and asynchronously process HTTP responses. From 2.0.0 until 3.0.14, connection-pool partitioning still omits identity-defining fields for Kerberos, SPNEGO, NTLM, and authenticated proxy connections. Logins without a configured principal, proxy realms, identities sharing a user name, and SOCKS or CONNECT proxy logins can reuse a socket authenticated as a different identity. A later request is then executed under the first identity and can expose that identity's data or authority to another caller. In the affected execution path, SpnegoEngine, NTLM, Kerberos, SPNEGO, SOCKS, and CONNECT control or expose the vulnerable behavior. This issue is fixed in version 3.0.14.
CVE-2026-107231	The AsyncHttpClient (AHC) library allows Java applications to easily execute HTTP requests and asynchronously process HTTP responses. Prior to 3.0.13 and 2.16.1, Realm.Builder treats a Digest challenge that yields no usable nonce as a Basic challenge. A malicious origin or proxy can label a challenge Digest while omitting or emptying the nonce, causing the client to resend the username and password using reversible Basic authentication. Both origin and proxy challenge parsers are affected. This issue is fixed in versions 3.0.13 and 2.16.1.
CVE-2026-107280	The AsyncHttpClient (AHC) library allows Java applications to easily execute HTTP requests and asynchronously process HTTP responses. Prior to 3.0.13 and 2.16.1, ThreadSafeCookieStore validates Domain attributes with domain matching but does not reject public suffixes. A host beneath a suffix such as co.uk can set a cookie for that suffix, after which the shared cookie store sends it to unrelated hosts under the suffix. This can inject or overwrite session-relevant cookie values across origins. This issue is fixed in versions 3.0.13 and 2.16.1.
CVE-2026-107282	The AsyncHttpClient (AHC) library allows Java applications to easily execute HTTP requests and asynchronously process HTTP responses. Prior to 3.0.13 and 2.16.1, cross-host request replay updates the current request but leaves the target request and related proxy context pointing at the original origin. Connection-pool selection, CONNECT handling, realm selection, and TLS setup can consequently send the original host's path, Host header, Authorization credentials, or plaintext request to the replay destination. Documented ResponseFilter failover and retry paths can trigger the replay. This issue is fixed in versions 3.0.13 and 2.16.1.
CVE-2026-107283	The AsyncHttpClient (AHC) library allows Java applications to easily execute HTTP requests and asynchronously process HTTP responses. Prior to 3.0.12 and 2.16.1, Realm.Builder generates the HTTP Digest client nonce with ThreadLocalRandom rather than a cryptographically secure random source. Digest relies on an unpredictable cnonce to resist chosen-plaintext and credential precomputation attacks, so an observer able to infer generator state can reduce the protection of the authentication exchange. This issue is fixed in versions 3.0.12 and 2.16.1.
CVE-2026-107285	The AsyncHttpClient (AHC) library allows Java applications to easily execute HTTP requests and asynchronously process HTTP responses. Prior to 3.0.12 and 2.16.1, a proxied ws request is carried through CONNECT, but NettyRequestFactory.newNettyRequest and requestUri decide whether to attach proxy authentication and an absolute-form target only from whether the URI is secure. Because ws is not marked secure, the tunneled WebSocket upgrade sent to the origin includes the proxy's Proxy-Authorization value. Basic credentials are directly recoverable and Digest responses can be replayed or cracked offline. This issue is fixed in versions 3.0.12 and 2.16.1.
CVE-2026-108106	Xerial snappy-java before 1.1.10.9 contains an unbounded memory allocation vulnerability that allows attackers to exhaust JVM memory by declaring a large uncompressed length in compressed input. Attackers can supply a few crafted bytes to Snappy.uncompress, uncompressString, SnappyInputStream or SnappyFramedInputStream to force allocations up to 2 GB, causing OutOfMemoryError and denial of service.
CVE-2026-111063	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to 3.14.4, the Secure attribute for a shared cookie is not enforced. An attacker can consequently cause a shared cookie marked Secure to be sent with an unencrypted request. This issue is fixed in version 3.14.4.

CVE ID	Description
CVE-2026-111064	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to 3.14.4, keepalive timeout handling is absent during the initial portion of a request. An attacker can keep an idle connection open indefinitely before sending any request data, consuming server connection capacity. This issue is fixed in version 3.14.4.
CVE-2026-111065	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to 3.14.4, Max-Age processing can overflow and skip closing a connection. An attacker can consequently exhaust the connection pool temporarily until garbage collection resolves the connections. This issue is fixed in version 3.14.4.
CVE-2026-111066	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. From 2.0.0 until 3.14.4, multipart request processing can retain an unbounded number of open file descriptors. An unauthenticated attacker can consequently exhaust file descriptors with specially crafted multipart requests when a handler uses Request.post(). This issue is fixed in version 3.14.4.
CVE-2026-111497	A buffer overflow vulnerability in the Brotli decoder (c/dec/decode.c) in Google Brotli versions 1.0.3 through 1.2.0 on all platforms allows a remote unauthenticated attacker to corrupt adjacent decoder state memory, disclose sensitive memory contents, cause a denial of service, or potentially execute arbitrary code using a crafted Brotli-compressed stream. Users need to update any Brotli libraries or binaries between 1.0.3 through 1.2.0 to the latest available version: 1.0.10 or 1.1.1 or 1.2.1.
CVE-2026-111748	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to 3.14.4, an attacker can send an absolute-form request that causes Request.scheme to use attacker-controlled request-target scheme data instead of the connection protocol. Applications that use Request.scheme for security decisions can consequently be tricked into treating an unencrypted connection as encrypted or an encrypted connection as unencrypted. This issue is fixed in version 3.14.4.
CVE-2026-111749	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to 3.14.4, Host-Only scope on an existing parent-domain cookie can be removed by a subdomain. A user who accesses an untrusted subdomain can consequently have host-only protections stripped from cookies, although impact is limited when applications do not reflect those cookies to a browser. This issue is fixed in version 3.14.4.
CVE-2026-111750	AIOHTTP is an asynchronous HTTP client/server framework for asyncio and Python. Prior to 3.14.4, the pure Python parser accepts ambiguous HTTP request framing. An unauthenticated attacker can consequently smuggle requests when an application uses AIOHTTP_NO_EXTENSIONS or otherwise runs without a standard wheel and an intermediary interprets the framing differently. This issue is fixed in version 3.14.4.
CVE-2026-111943	On most platforms, Go uses the given system root certificates. However, in some narrower cases, the system trust stores may not be updated without further action. To defend against the narrower set of cases that may remain affected, a denylist was added in crypto/x509 to return an error when calling Verify with any of the known intermediate certificates that are no longer considered to be trustworthy.
CVE-2026-112390	Pillow is a Python imaging library. From 7.2.0 until 13.0.0, Pillow can enter an infinite loop while determining the frame count of a MIC image. MicImageFile inherits TIFF frame traversal, but MicImageFile.tell reports the OLE stream index and its seek() resets TIFF frame state, so an embedded TIFF with at least three IFDs never advances. When an application with olefile installed reads n_frames or saves all frames from a crafted MIC file, the non-advancing traversal can hold a worker at full CPU indefinitely. This issue is fixed in version 13.0.0.
CVE-2026-112391	Pillow is a Python imaging library. From 5.1.0 until 13.0.0, Pillow's PdfParser.get_value() can recurse without a limit while reading nested PDF objects. The parser follows attacker-controlled arrays and dictionaries when an existing PDF is opened for append. An attacker can provide a deeply nested PDF to an application that saves an image with PDF append enabled. This can raise an uncaught RecursionError or terminate the worker, causing denial of service. This issue is fixed in version 13.0.0.
CVE-2026-112392	Pillow is a Python imaging library. Prior to 13.0.0, Pillow's rank-based filters can dereference invalid image storage for a zero-size image. ImagingExpand() allocates a non-empty output and then faults while replicating edges by reading the first pixel from empty input storage; ImagingRankFilter is never reached. An attacker can supply or cause creation of a zero-size image to an application that applies a rank, median, minimum, or maximum filter. This causes a native process crash. This issue is fixed in version 13.0.0.
CVE-2026-112393	Pillow is a Python imaging library. From 2.0.0 until 13.0.0, Pillow has a heap buffer overflow and stale-heap disclosure in _tiffWriteProc when saving compressed TIFF data through libtiff to a target without a usable fileno(), such as BytesIO. The callback sizes its allocation from the prior buffer size instead of the current write position, allowing an attacker-controlled out-of-line TIFF tag or retained ICC profile to be copied beyond the allocation after libtiff seeks forward. This overflow can crash the process and may allow code execution. Independently, unwritten alignment gaps in compressed TIFF output larger than 64 KiB can disclose stale heap bytes without a crafted tag. File paths and targets with real file descriptors are unaffected. This issue is fixed in version 13.0.0.

CVE ID	Description
CVE-2026-112394	Pillow is a Python imaging library. Prior to 13.0.0, Pillow's P;2L and P;4L raw decoders can read past the end of an input buffer. ImagingRawDecode calculates a packed-pixel row size, while the unpackers read separately padded bit planes. A caller can supply a buffer and image dimensions whose accepted size is smaller than the unpacker's actual footprint. This can disclose up to three adjacent heap bytes in output pixels or crash the process. This issue is fixed in version 13.0.0.
CVE-2026-112395	Pillow is a Python imaging library. From 8.2.0 until 13.0.0, Pillow's tiled-TIFF decoder can read past the tile buffer while copying four-byte pixels. In _decodeTile, libtiff can return a tile allocation smaller than the required row stride multiplied by the tile length, and Pillow's stride-and-copy loop walks beyond that allocation. An attacker can supply a crafted tiled TIFF to an application that loads it. This can expose adjacent heap data in pixels or crash the process. This issue is fixed in version 13.0.0.
CVE-2026-112396	Pillow is a Python imaging library. From 11.3.0 until 13.0.0, Pillow has a heap use-after-free when libtiff rejects a tag while encoding TIFF data to an in-memory buffer. When an application saves attacker-influenced TIFF tags to BytesIO, libtiff can reject a tag during encoding in ImagingLibTiffEncode; the error path then frees the output buffer without closing or clearing the TIFF handle, and later cleanup flushes through the write callback into freed memory, which can crash the process and may enable code execution. This issue is fixed in version 13.0.0.
CVE-2026-112397	Pillow is a Python imaging library. Prior to 13.0.0, Pillow's ImageMorph.MorphOp.apply() can access memory before a row buffer for a zero-width image. PPM and PGM input can create an image with width zero, and the native morphology routine indexes outrow[width - 1] without rejecting that geometry. An attacker can supply such an image to an application that applies a morphology operation. This can crash the process. This issue is fixed in version 13.0.0.
CVE-2026-112398	Pillow is a Python imaging library. Prior to 13.0.0, Pillow's FlashPix JPEG-tile decoder can write decoded scanlines beyond the allocated tile buffer. FlashPix (FPX) JPEG decoding selects raw mode LA with two bytes per pixel, while an unrecognized JPEG mode can be decoded as JCS_UNKNOWN with four components, causing ImagingJpegDecode to write four-component scanlines into the two-byte tile buffer. An attacker can supply a crafted OLE2 FlashPix file to an application with the optional olefile dependency installed that opens or thumbnails untrusted input. This can corrupt heap memory and may allow arbitrary code execution. This issue is fixed in version 13.0.0.
CVE-2026-112399	Pillow is a Python imaging library. From 8.1.0 until 13.0.0, Pillow can allocate an APNG disposal background at the full canvas size before applying the decompression-bomb check. A crafted APNG with a small first frame, a very large declared canvas, and background disposal or first-frame OP_PREVIOUS reaches PngImageFile._seek while Image.open() positions frame zero. The full-canvas allocation occurs inside Image.open() without requiring pixel loading or frame iteration. This can exhaust memory and deny service. This issue is fixed in version 13.0.0.
CVE-2026-112400	Pillow is a Python imaging library. From 9.2.0 until 13.0.0, Pillow can retain a rejected GIF frame size after a decompression-bomb exception. In GifImageFile._seek, a crafted GIF can leave rejected dimensions stored after Image.DecompressionBombError; if an application catches that exception and continues processing the same image object, later seek or load operations allocate according to those dimensions and can exhaust memory, denying service. This issue is fixed in version 13.0.0.
CVE-2026-112401	Pillow is a Python imaging library. From 5.4.0 until 13.0.0, Pillow performs quadratic copying while parsing GIF comments. GifImageFile._seek() repeatedly appends immutable bytes for comment sub-blocks and multiple comment extensions, and frame zero is parsed during Image.open(). An attacker can submit a small-dimension GIF containing many comment extensions to any application that opens the file. This can consume excessive CPU and deny service before the application can inspect the image. This issue is fixed in version 13.0.0.
CVE-2026-112402	Pillow is a Python imaging library. Prior to 13.0.0, Pillow can double free image memory when ImagingScaleAffine encounters a selective allocation failure. The affected fast path handles Transform.EXTENT or pure scale-and-translate Transform.AFFINE operations with the default NEAREST resampling. The output image allocation must succeed while the additional per-column lookup allocation fails, after which the error path deletes image storage that later cleanup releases again. An attacker who can drive such a large Image.transform() operation under memory pressure may trigger the failure. This can crash the process and may corrupt memory. This issue is fixed in version 13.0.0.
CVE-2026-112403	Pillow is a Python imaging library. Prior to 13.0.0, Pillow's JPEG decoder can write a full-width scanline into a buffer sized for a requested draft scale. ImagingJpegDecode does not verify libjpeg's actual output width and component count after scaling is requested, and an 8-bit lossless SOF3 JPEG decoded by a build linked to libjpeg-turbo 3.x can ignore the requested downscaling geometry. An attacker can supply such a JPEG to an application that uses draft() or thumbnail() to request downscaling. This can corrupt heap memory, crash the process, or potentially allow code execution. Older libjpeg-turbo 2.x and IJG builds reject the SOF3 input. This issue is fixed in version 13.0.0.

CVE ID	Description
CVE-2026-112404	Pillow is a Python imaging library. From 8.1.0 until 13.0.0, Pillow's TIFF tile and RGBA decode paths can read beyond row data produced by libtiff. In _decodeTile and the RGBA unpacker path, when Pillow and libtiff interpret duplicated PhotometricInterpretation or BitsPerSample tags differently, Pillow does not verify that the required input row fits inside the row libtiff produced, while the strip path has this validation. Loading pixels from such a crafted TIFF can expose adjacent heap data in output pixels or crash the process. This issue is fixed in version 13.0.0.
CVE-2026-112405	Pillow is a Python imaging library. From 12.0.0 until 13.0.0, Pillow's Image.merge() can dereference a null image8 plane. Validation skips the first input band, allowing a 32-bit integer or floating-point image backed by image32 storage to be merged as a narrower multiband image. A crafted IPTC file can reach this path when an application opens and loads it, or a caller can invoke Image.merge() directly with such a first band. This causes a native process crash. This issue is fixed in version 13.0.0.
CVE-2026-112406	Pillow is a Python imaging library. From 5.1.0 until 13.0.0, Pillow can enter an infinite loop while rewriting PDF pages. PdfParser.rewrite_pages() follows Parent relationships without detecting a cycle in the page tree. An attacker can supply a cyclic PDF to an application that appends an image to that PDF. This can hold the worker at full CPU and deny service. This issue is fixed in version 13.0.0.
CVE-2026-112891	Pillow is a Python imaging library. From 9.2.0 until 13.0.0, Pillow performs quadratic copying when saving a GIF with a long comment. _get_global_header() repeatedly copies portions of comment data while assembling the global header, including comments retained from an untrusted input image. An attacker can supply an image with a large GIF comment to an application that saves it as GIF. This can consume excessive CPU and delay or deny service. This issue is fixed in version 13.0.0.
CVE-2026-112892	Pillow is a Python imaging library. From 4.1.0 until 13.0.0, Pillow's MSP RLE decoder can expand a small file into an unbounded in-memory buffer. The decoder writes run-length output to BytesIO without limiting expansion to the declared image size. An attacker can supply a crafted MSP image to an application that loads its pixels. This can exhaust memory and deny service even when the declared pixel count passes the decompression-bomb check. This issue is fixed in version 13.0.0.
CVE-2026-112893	Pillow is a Python imaging library. From 10.0.0 until 13.0.0, Pillow can free an invalid glyph pointer during FreeType rendering with an active stroker. If FT_Load_Glyph fails, cleanup may call FT_Done_Glyph on an uninitialized or stale glyph variable. An attacker can supply a malformed font as untrusted bytes to an application that draws text with a nonzero stroke width, forcing the glyph-load failure path. This can corrupt memory or crash the process. This issue is fixed in version 13.0.0.
CVE-2026-112894	Pillow is a Python imaging library. From 5.1.0 until 13.0.0, Pillow takes quadratic time to reserve PDF object identifiers when saving many frames with save_all enabled. PdfParser.next_object_id() repeatedly rebuilds and scans all previously reserved identifiers for every object. An attacker can supply a multi-frame image, including a small APNG with a large declared frame count, to an application that converts it to a multipage PDF. This can consume excessive CPU and deny service. This issue is fixed in version 13.0.0.
CVE-2026-112895	Pillow is a Python imaging library. From 5.1.0 until 13.0.0, Pillow can exhaust resources while parsing an existing PDF during an append operation. PdfParser follows page-tree child references without cycle or duplication tracking and follows previous-trailer chains with only an exact repeated-offset guard, so deep or cyclic page trees and long chains of unique trailer offsets can exhaust the Python stack. Repeated non-cyclic /Kids references can also cause super-linear CPU and memory use. An attacker can supply such a PDF to an application that calls Image.save with PDF append enabled. This can raise an uncaught RecursionError or consume excessive CPU and memory before output is written, causing denial of service. This issue is fixed in version 13.0.0.
CVE-2026-112896	Pillow is a Python imaging library. Prior to 13.0.0, Pillow's TIFF encoder can read beyond a ColorMap buffer when retained TIFF metadata has a plane size inconsistent with BitsPerSample. A small TIFF can be opened normally and then retain a crafted ColorMap that is reused when the image is saved. An application that re-encodes the attacker's TIFF with libtiff can trigger the out-of-bounds read. This can disclose heap memory in the output or crash the process. This issue is fixed in version 13.0.0.
CVE-2026-112897	Pillow is a Python imaging library. From 5.3.0 until 13.0.0, Pillow can return uninitialized memory as pixels when libtiff performs a short decode of JPEG-compressed or Group 4 fax TIFF data. The _decodeStrip and _decodeTile buffers can be larger than the bytes written by libtiff, but Pillow converts the expected rows without clearing the unwritten tail. The tile path is affected from 5.3.0 and the strip path from 6.0.0. An attacker can supply a crafted TIFF to an application that loads and returns or re-encodes its pixels. This can disclose stale heap contents without memory corruption or a process crash. This issue is fixed in version 13.0.0.
CVE-2026-112898	Pillow is a Python imaging library. From 9.1.0 until 13.0.0, Pillow can disclose uninitialized stack memory when Image.quantize() uses an empty palette. ImagingPaletteCacheUpdate() leaves a local cache array uninitialized when the palette has zero entries, and the conversion path emits those bytes as output pixel indices. An attacker who controls the palette image can obtain stack contents from a returned PNG or GIF. This can disclose addresses and stack-canary data without an out-of-bounds access, memory corruption, or process crash. This issue is fixed in version 13.0.0.

CVE ID	Description
CVE-2026-112899	Pillow is a Python imaging library. In versions 2.3.0 through 12.3.0, EPS rendering does not validate the scaled image dimensions before invoking Ghostscript. Applications that process untrusted EPS or PostScript files using an application-selected scale greater than one can incur excessive Ghostscript CPU and temporary disk consumption. In versions with decompression-bomb protection, the scaled output is checked only after Ghostscript has completed rendering, allowing resource consumption beyond the configured pixel limit. The amplification depends on the scale selected by the application, not an attacker-controlled scale value. This issue is fixed in version 13.0.0.
CVE-2026-112900	Pillow is a Python imaging library. From 2.6.0 until 13.0.0, Pillow is vulnerable to argument injection when invoking Ghostscript for EPS input. EpsImagePlugin.Ghostscript passes a caller-controlled input filename after -f without converting it to a safe absolute path, so a name beginning with a hyphen can be interpreted as a Ghostscript option. An attacker who controls the filename can inject Ghostscript arguments. Arbitrary command execution specifically requires a leading @ argument-file filename and control of the referenced text file containing additional Ghostscript arguments; the reported chain reproduces with the default -dSAFER invocation. This issue is fixed in version 13.0.0.
CVE-2026-112901	Pillow is a Python imaging library. Prior to 13.0.0, Pillow's XBM header parser is vulnerable to regular-expression denial of service. In XbmImageFile._open, when an application calls Image.open() on a small crafted XBM file, XbmImagePlugin applies a backtracking expression to the attacker-controlled header before image dimensions are accepted, stalling a worker for minutes and denying service. This issue is fixed in version 13.0.0.
CVE-2026-112902	Pillow is a Python imaging library. In versions through 12.3.0, the libtiff TIFF encoder can pass arguments to TIFFSetField() that do not match libtiff's definition of the corresponding tag. When an application saves attacker-controlled image metadata through this encoder, crafted tag values can cause invalid memory reads, copying process memory into the encoded output or crashing the process. Applications that return the encoded output to an attacker can disclose that memory. The issue is reachable when re-saving compressed TIFF images and through certain conversions that preserve EXIF metadata into compressed TIFF output. This vulnerability is fixed in version 13.0.0.
CVE-2026-112909	filelock is a platform-independent file lock for Python. Prior to 4.2.1, UnixFileLock can truncate a regular file renamed onto the lock path when another local user has actual rename permission, or hard-linked there under a permissive hard-link policy; when an explicit mode is requested, affected versions from 3.10.2 can also change that file's mode. This issue is fixed in version 4.2.1.
CVE-2026-112910	filelock is a platform-independent file lock for Python. From 3.10.0 until 4.2.1, filelock on Windows ignores the mode argument for WindowsFileLock files and the ReadWriteLock SQLite database. A lock created with mode=0o600 receives the parent directory's inheritable ACL because _open_non_reparse_fd and _nt_open do not create an owner-only security descriptor. Another local account with inherited access can read holder metadata or the lock database, modify those files where the ACL permits writes, or take the byte-range lock after release and delay later acquisitions. UnixFileLock is not affected, and mutual exclusion between cooperating processes continues to hold. This issue is fixed in version 4.2.1.
CVE-2026-113714	virtualenv is a tool for creating isolated virtual python environments. Prior to 21.14.9, download_wheel() writes package-index downloads directly into the reusable wheel/house cache before selecting a verification target from pip's textual Saved output. With default-index digest comparison enabled, a controlled package index can inject a misleading Saved line through a discarded link's data-requires-python attribute; this forged-line variant also requires knowledge of an absolute path to the genuine bundled wheel. The wrong path can then be hashed, while a digest mismatch leaves the rejected archive in the cache for later selection without verification. A separate rejected-cache-file variant requires control over downloaded data or archive selection but not the forged-line prerequisites. Explicit custom-index exceptions are outside the promised PyPI comparison, and the issue does not independently bypass TLS or establish control of default PyPI. This issue is fixed in version 21.14.9.
CVE-2026-113715	virtualenv is a tool for creating isolated virtual python environments. Prior to 21.14.9, the stdlib venv creator reads generated pyvenv.cfg with str.splitlines() and merges parsed keys over virtualenv's selected configuration. An untrusted caller who influences the destination pathname can include a line-boundary character such as U+2028 to inject keys including home and include-system-site-packages, causing path data to change configuration structure before final output. Setting include-system-site-packages to true weakens environment isolation, while an altered home can change which standard-library files the interpreter loads if the attacker also controls that location. The builtin creator does not use this readback path. This issue is fixed in version 21.14.9.
CVE-2026-113716	virtualenv is a tool for creating isolated virtual python environments. Prior to 21.14.9, the builtin creator calls ensure_dir() for executable, library, and package directories, and ensure_dir() accepts an existing symlink to a directory because exists() follows it. Interpreter creation, seed installation, and activation generation then write through the linked parent outside the requested destination when another local account prepared the tree and the invoking user can access the target. The --clear option removes the planted tree first, while ordinary creation reaches it; arbitrary attacker-selected output bytes are not established. This issue is fixed in version 21.14.9.

CVE ID	Description
CVE-2026-113717	virtualenv is a tool for creating isolated virtual python environments. Prior to 21.14.9, PipInvoke.get_pip_install_cmd() handles --seeder pip --download by calling get_wheel(download=False) and then allowing pip to consult the configured index, bypassing virtualenv's download_wheel() and verify_wheel_digest() path. When virtualenv's default-index PyPI digest comparison remains enabled, an index candidate matching the bundled package and version can therefore be selected without comparison to PyPI SHA-256 metadata. Explicit custom-index settings already opt out of that comparison and remain outside this verification-bypass claim; the demonstrated configuration used a non-PyPI index while leaving the index environment variables unset so the comparison policy remained active. This condition does not independently defeat TLS or establish control of the index response. This issue is fixed in version 21.14.9.
CVE-2026-113718	virtualenv is a tool for creating isolated virtual python environments. Prior to 21.14.9, POSIX xonsh activation quote() uses repr() to write printable non-ASCII prompt or destination data into UTF-8 activate.xsh, while xonsh can read that script using $XONSH_ENCODING. Under a legacy multibyte codec such as GBK or Big5, a UTF-8 trail byte can consume the backslash before an adjacent quote, allowing the literal to close and attacker-influenced statements to execute when the user sources the generated script. This issue is fixed in version 21.14.9.
CVE-2026-113719	virtualenv is a tool for creating isolated virtual python environments. Prior to 21.14.9, PipInstall._uninstall_previous_version() accepts existing seed-package metadata and _uninstall_dist() joins entries from top_level.txt and RECORD to site-packages without checking containment. When a reused destination contains distribution metadata influenced by another account, relative traversal, absolute entries, or linked parents can cause app-data reseeding to delete files outside the environment or the environment root with the invoking user's permissions, provided that user has deletion rights over the named paths. The deletion occurs during environment creation before any later use of the environment. Fresh environments without existing package metadata and creation with --clear do not enter this cleanup path. This issue is fixed in version 21.14.9.
CVE-2026-113720	virtualenv is a tool for creating isolated virtual python environments. Prior to 21.14.9, the app-data copy seeder uses shutil.copy() to copy cached package files into later environments, causing cached source modes to override the current run's umask. A cache populated under a permissive umask can therefore produce files writable by another local account in a later traversable environment, allowing package code such as pip/__init__.py to be modified and executed by the environment owner. The attacker cannot force the earlier umask, and a private destination parent prevents access. This issue is fixed in version 21.14.9.
CVE-2026-113721	virtualenv is a tool for creating isolated virtual python environments. Prior to 21.14.9, ViaTemplateActivator._generate() removes an existing output only when Path.exists() is true, allowing a dangling symlink to survive and redirect generated activation content, while PyEnvCfg.write() opens pyvenv.cfg without first replacing an existing symbolic or hard link. Another local account that prepares final entries inside the destination can therefore cause environment creation to create or overwrite files outside the requested tree with constrained virtualenv-generated content and the invoking user's permissions. The destination parent must remain controlled by the environment owner during creation; concurrent parent-directory replacement is a separate concern. This issue is fixed in version 21.14.9.
CVE-2026-113722	virtualenv is a tool for creating isolated virtual python environments. Prior to 21.14.9, virtualenv's ScriptMakerCustom and pip pass an incompletely quoted interpreter path to distlib ScriptMaker when generating POSIX console scripts. When the path contains a space or produces a shebang longer than 127 bytes, distlib emits a /bin/sh wrapper whose double-quoted interpreter path still expands command substitutions and backticks. An attacker who influences the environment destination can therefore execute commands with the user's privileges when the user later runs a generated launcher; environment creation and activation do not trigger the command. Existing launchers are not repaired by upgrading virtualenv alone and must be regenerated, while a later independent pip installation or upgrade can reintroduce an affected launcher because pip's vendored distlib is unchanged. The virtualenv fix is released in version 21.14.9.
CVE-2026-113723	virtualenv is a tool for creating isolated virtual python environments. Prior to 21.14.9, caller-supplied prompt text can become executable syntax in the bash activation prompt. An attacker must influence the value passed to virtualenv, and a user must subsequently activate the environment and draw the prompt. This vulnerability is fixed in 21.14.9.
CVE-2026-113724	virtualenv is a tool for creating isolated virtual python environments. Prior to 21.14.9, the safe_delete() deletion error handler responds to a failed removal by reading mode bits with os.lstat(path) and applying owner-write permission with os.chmod(path, ...), which follows a symlink. A planted link inside a read-only directory can therefore cause a victim-owned target outside the destination that is initially not writable by the invoking user to become broadly writable before the retry fails, provided that user owns the target and can change its mode. Another account can use the widened mode only if it can traverse the target's parent directories; root does not reach the demonstrated failure path, and the operation does not itself read the target. This issue is fixed in version 21.14.9.
CVE-2026-113725	virtualenv is a tool for creating isolated virtual python environments. Prior to 21.14.7, the bash activate script escapes the environment name for only one of the two passes bash makes over PS1, so a backslash in the name undoes the escaping and a name containing \$(...) runs as commands in the user's shell each time bash draws the prompt. This issue is fixed in version 21.14.7.

CVE ID	Description
CVE-2026-113726	virtualenv is a tool for creating isolated virtual python environments. Prior to 21.13.0, the bash activate script that zsh also sources leaves the environment name unescaped; from 21.13.0 until 21.14.4, it escapes backslashes, dollar signs, and backticks only when BASH_VERSION is set, so zsh receives the raw value. The script assigns that value to PS1. In zsh with PROMPT_SUBST, an attacker-influenced environment directory, --prompt, or VIRTUALENV_PROMPT value containing command, parameter, or arithmetic expansion syntax can therefore execute commands in the user's shell each time the prompt is drawn. A percent sign is separately interpreted as a zsh prompt escape but does not by itself execute commands. Plain zsh with PROMPT_SUBST disabled and bash are not affected by this mechanism. This issue is fixed in version 21.14.4.
CVE-2026-113732	virtualenv is a tool for creating isolated virtual python environments. Prior to 21.14.2, FishActivator inherits ViaTemplateActivator.quote(), which uses POSIX shlex.quote() even though fish interprets backslashes and single quotes differently. An attacker-influenced environment path, --prompt, or VIRTUALENV_PROMPT value containing a backslash before a quote can therefore terminate the generated fish literal and execute command substitution when the user sources activate.fish. This issue is fixed in version 21.14.2.
CVE-2026-114136	Django REST framework is a toolkit for building Web APIs. From 3.6.4 until 3.18.2, JSONParser.parse in rest_framework/parsers.py uses the client-supplied Content-Type charset value with codecs.getreader, which accepts byte-transform codecs such as bz2_codec rather than only text encodings. A remote unauthenticated client can therefore cause a compressed JSON request body to be expanded after Django has enforced DATA_UPLOAD_MAX_MEMORY_SIZE against only the wire bytes, allowing a very small request to consume substantial per-worker memory. The default parser set is sufficient, and any reachable endpoint that accepts a JSON body is affected. This issue is fixed in version 3.18.2.
CVE-2026-114137	Django REST framework is a toolkit for building Web APIs. From 3.12.0 until 3.18.2, TokenProxy.delete() in rest_framework/authtoken/models.py uses the proxy's user-id pk while the concrete Token token-string primary key identifies rows, so Django issues a deletion query that matches no row and reports zero deletions without raising an error. Applications, administrative integrations, or operational scripts that revoke credentials through a TokenProxy instance can therefore believe revocation succeeded while the token remains live and continues authenticating. DRF's own TokenAdmin.delete_model workaround is not affected, but other callers can leave access available indefinitely because DRF tokens do not expire automatically. This issue is fixed in version 3.18.2.
CVE-2026-114138	Django REST framework is a toolkit for building Web APIs. Prior to 3.18.2, DefaultContentNegotiation.get_accept_list in rest_framework/negotiation.py accepts an unbounded number and aggregate size of Accept tokens, while select_renderer in that file repeatedly invokes _MediaType parsing from rest_framework/utils/mediatypes.py for each attacker-controlled token, every renderer, and forced 406 renegotiation. A remote unauthenticated client can send many distinct media types containing large quoted parameter values with numerous semicolons, causing Django's parameter parser and DRF's repeated parsing structure to consume disproportionate CPU. Requests to any reachable endpoint can pin synchronous workers for prolonged periods, and concurrent requests can make the service unavailable until the requests stop. This issue is fixed in version 3.18.2.
CVE-2026-114501	Pillow is a Python imaging library. From 11.2.1 until 13.0.0, Arrow export assumes that an image's rows are tightly contiguous and exports from im->blocks[0].ptr even when PILLOW_ALIGNMENT makes aligned_linesize include row padding or an aligned offset before the first row. The resulting Arrow buffer can include uninitialized padding, bytes preceding the first row, or stale heap contents from previously freed images. Applications that configure alignment and export images to Arrow can therefore disclose one user's pixels or other process memory to another user. This issue is fixed in version 13.0.0.
CVE-2026-114613	MessagePack is the serializer implementation for Python msgpack.org. Prior to 1.2.3, the msgpack-python C extension (msgpack._cmsgpack) leaks a decoded map value when deserialization rejects a key with strict_map_key=True (the default). Repeated malformed inputs can cause unbounded process memory growth and denial of service in applications that deserialize attacker-controlled MessagePack data. Both unpackb() and the streaming Unpacker API are affected. The pure-Python fallback does not exhibit this leak in the reported tests. This issue is fixed in version 1.2.3.
CVE-2026-115927	Tornado is a Python web framework and asynchronous networking library. Prior to 6.6.0, Tornado decodes HTTP/1.1 chunked transfer encoding by processing body data that has already arrived without yielding to the event loop. A remote unauthenticated peer can send millions of one-byte chunks to monopolize the process event loop and stall every other connection served by that process, while lowering max_body_size only limits the duration of each request. For handlers that do not use @stream_request_body, retaining each chunk as a separate object can consume roughly twenty times the received body size, and SimpleAsyncHTTPClient has the same event-loop behavior when reading a response from a malicious server. This issue is fixed in version 6.6.0.
CVE-2026-116577	Tornado is a Python web framework and asynchronous networking library. Prior to 6.6, Tornado's HTTP header parser accepts obsolete line folding and repeatedly concatenates short continuation lines, causing quadratic CPU consumption on the IOLoop thread. An unauthenticated client can submit a header block containing many obsolete line-folding continuations before application code runs, blocking other connections served by the process. The cost grows with max_header_size, and only availability is affected. This issue is fixed in version 6.6.

CVE ID	Description
CVE-2026-116578	Tornado is a Python web framework and asynchronous networking library. Prior to 6.6, tornado.web.RedirectHandler URL templates can place regex-capture substitutions where crafted request paths change the redirect scheme or network location. Affected placements include substitutions at the start of a path, directly after a fixed host, or at the beginning of the template, including protocol-relative, absolute, backslash, encoded, host-extension, and user-info variants interpreted by browsers as attacker-controlled destinations. Static templates and substitutions after a fixed host followed by a slash are not affected. An attacker can use the resulting open redirect on the application's domain for phishing. This issue is fixed in version 6.6.
CVE-2026-116590	snappy-java is a Java port of the snappy, a fast C++ compresser/decompresser developed by Google. Prior to 1.1.10.9, Snappy.arrayCopy(Object src, int offset, int byteLength, Object dest, int destOffset) passes caller-supplied ranges to native memcpy without validating either array, allowing negative or oversized byteLength and offset values to read or write outside array bounds and crash the JVM. SnappyOutputStream.rawWrite and SnappyInputStream.rawRead, including their typed-array write and read methods, pass caller-controlled ranges through this primitive, and typed-array byte-size calculations can overflow int. This issue is fixed in version 1.1.10.9.
CVE-2026-118259	opentelemetry-java is the Java implementation of the OpenTelemetry API for recording telemetry, and SDK for managing telemetry recorded by the API. Prior to 1.67.0, W3CTraceContextPropagator.extract() uses regex delimiter processing that can exhibit quadratic backtracking while parsing tracestate headers containing long runs of spaces or tabs without a following comma. An unauthenticated attacker whose headers reach extraction can repeatedly block request or event-loop threads and cause excessive CPU consumption, with practical impact bounded by ingress header-size limits, proxy configuration, and threading. This issue is fixed in version 1.67.0.
CVE-2026-119210	MariaDB Connector/J is used to connect applications developed in Java to MariaDB and MySQL databases. From 3.0.0 until 3.3.7, 3.4.5, and 3.5.11, the ClientMessage.validateLocalFileName() check used a regular expression that could match LOAD DATA LOCAL INFILE text inside quoted identifiers, string literals, or comments instead of requiring a real statement. A malicious or compromised server, or a machine-in-the-middle attacker, could return crafted database or table identifiers that the connector reused in metadata or updatable-result queries, then answer with a LOCAL INFILE request for a client file. With allowLocalInfile enabled, this caused the connector to stream any readable client-host file named by the crafted text, including through DatabaseMetaData.getImportedKeys(), DatabaseMetaData.getExportedKeys(), or ResultSet.updateRow() paths. This issue is fixed in versions 3.3.7, 3.4.5, and 3.5.11.
CVE-2026-119276	MariaDB Connector/J is used to connect applications developed in Java to MariaDB and MySQL databases. From 3.0.0 until 3.3.7, 3.4.5, and 3.5.11, when CLIENT_SESSION_TRACK was negotiated, the OkPacket.parse() and OkPacket.parseWithInfo() paths allocated byte arrays using server-declared lengths for the OK packet info field and session-state variable names and values before confirming that the bytes remained in the packet. A malicious or compromised server, or a machine-in-the-middle attacker, could send a small OK packet declaring a much larger length and repeatedly force attacker-controlled heap allocations before the subsequent copy failed. This could create garbage-collection pressure or an OutOfMemoryError in the client application. This issue is fixed in versions 3.3.7, 3.4.5, and 3.5.11.
CVE-2026-119279	MariaDB Connector/J is used to connect applications developed in Java to MariaDB and MySQL databases. From 3.0.0 until 3.3.7, 3.4.5, and 3.5.10, the ReadableByteBuf.readLength() and ReadableByteBuf.readIntLengthEncodedNotNull() methods narrowed the protocol's unsigned 8-byte length-encoded integer to a Java int without checking that the value fit a non-negative integer. Values above Integer.MAX_VALUE could wrap to a negative or unrelated length that downstream allocation and read paths used. A malicious or compromised server, or a machine-in-the-middle attacker, could trigger NegativeArraySizeException, out-of-bounds failures, or oversized allocation behavior in client connection handling. This issue is fixed in versions 3.3.7, 3.4.5, and 3.5.10.
CVE-2026-119280	MariaDB Connector/J is used to connect applications developed in Java to MariaDB and MySQL databases. From 3.0.0 until 3.3.7, 3.4.5, and 3.5.11, the MariaDbPoolConnection.recover() and MariaDbPoolPinnedConnection.recover() implementations trusted the gtrid_length and bqual_length columns returned by XA RECOVER. The methods allocated and split Xid byte arrays without validating the lengths against Xid.MAXGTRIDSIZE, Xid.MAXBQUALSIZE, or the actual data payload. A malicious or compromised server, or a machine-in-the-middle attacker, could cause large allocations, NegativeArraySizeException, or ArrayIndexOutOfBoundsException during XA recovery. This issue is fixed in versions 3.3.7, 3.4.5, and 3.5.11.
CVE-2026-119281	MariaDB Connector/J is used to connect applications developed in Java to MariaDB and MySQL databases. From 3.0.0 until 3.3.7, 3.4.5, and 3.5.11, the connector assumed result bytes were UTF-8 but guarded only changes to character_set_client. It did not reject non-UTF-8 character_set_results changes and relied on the server's global session_track_system_variables configuration to report character-set changes. A malicious or compromised server, or a machine-in-the-middle attacker, could suppress tracking or change the result charset so application strings differed from the returned bytes, potentially defeating application-side validation or escaping that depended on the decoded values. This issue is fixed in versions 3.3.7, 3.4.5, and 3.5.11.

Using the Cloudera Runtime Maven repository

Information about using Maven to build applications with Cloudera Runtime components.

If you want to build applications or tools for use with Cloudera Runtime components and you are using Maven or Ivy for dependency management, you can pull the Cloudera Runtime artifacts from the Cloudera Maven repository. The repository is available at https://repository.cloudera.com/artifactory/cloudera-repos/.

Important: When you build an application JAR, do not include CDH JARs, because they are already provided. If you do, upgrading CDH can break your application. To avoid this situation, set the Maven dependency scope to provided. If you have already built applications which include the CDH JARs, update the dependency to set scope to provided and recompile.

The following is a sample POM (pom.xml) file:

```
<project xmlns="http://maven.apache.org/POM/4.0.0" xmlns:xsi="http://www.w3.
org/2001/XMLSchema-instance" xsi:schemaLocation="http://maven.apache.org/POM
/4.0.0 http://maven.apache.org/maven-v4_0_0.xsd">
  <repositories>
    <repository>
      <id>cloudera</id>
      <url>https://repository.cloudera.com/artifactory/cloudera-repos/</url>
    </repository>
  </repositories>
</project>
```

Cloudera Runtime 7.3.3.0-56

The following table lists the project name, groupId, artifactId, and version required to access each RUNTIME artifact.

Project	groupId	artifactId	version
Atlas	org.apache.atlas	atlas-authn	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-authorization	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-aws-s3-bridge	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-azure-adls-bridge	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-classification-updater	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-client-common	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-client-v1	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-client-v2	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-client-v2-shaded	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-common	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-distro	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-docs	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-graphdb-api	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-graphdb-common	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-graphdb-janus	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-hdfs-bridge	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-index-repair-tool	2.4.0.7.3.3.0-56

Project	groupId	artifactId	version
Atlas	org.apache.atlas	atlas-intg	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-notification	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-notification-analyzer	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-plugin-classloader	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-repository	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-server-api	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-server-common	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	atlas-testtools	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	hbase-bridge	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	hbase-bridge-shim	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	hbase-testing-util	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	hdfs-model	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	hive-bridge	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	hive-bridge-shim	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	impala-bridge	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	impala-bridge-shim	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	impala-hook-api	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	janusgraph-rdbms	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	kafka-bridge	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	kafka-bridge-shim	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	navigator-to-atlas	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	sample-app	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	sqoop-bridge	2.4.0.7.3.3.0-56
Atlas	org.apache.atlas	sqoop-bridge-shim	2.4.0.7.3.3.0-56
Avro	org.apache.avro	avro	1.11.3.7.3.3.0-56
Avro	org.apache.avro	avro-android	1.11.3.7.3.3.0-56
Avro	org.apache.avro	avro-codegen-test	1.11.3.7.3.3.0-56
Avro	org.apache.avro	avro-compiler	1.11.3.7.3.3.0-56
Avro	org.apache.avro	avro-grpc	1.11.3.7.3.3.0-56
Avro	org.apache.avro	avro-ipc	1.11.3.7.3.3.0-56
Avro	org.apache.avro	avro-ipc-jetty	1.11.3.7.3.3.0-56
Avro	org.apache.avro	avro-ipc-netty	1.11.3.7.3.3.0-56
Avro	org.apache.avro	avro-mapred	1.11.3.7.3.3.0-56
Avro	org.apache.avro	avro-maven-plugin	1.11.3.7.3.3.0-56
Avro	org.apache.avro	avro-perf	1.11.3.7.3.3.0-56
Avro	org.apache.avro	avro-protobuf	1.11.3.7.3.3.0-56
Avro	org.apache.avro	avro-service-archetype	1.11.3.7.3.3.0-56
Avro	org.apache.avro	avro-test-custom-conversions	1.11.3.7.3.3.0-56

Project	groupId	artifactId	version
Avro	org.apache.avro	avro-thrift	1.11.3.7.3.3.0-56
Avro	org.apache.avro	avro-tools	1.11.3.7.3.3.0-56
Avro	org.apache.avro	trevni-avro	1.11.3.7.3.3.0-56
Avro	org.apache.avro	trevni-core	1.11.3.7.3.3.0-56
Calcite	org.apache.calcite	calcite-babel	1.33.0.7.3.3.0-56
Calcite	org.apache.calcite	calcite-core	1.33.0.7.3.3.0-56
Calcite	org.apache.calcite	calcite-druid	1.33.0.7.3.3.0-56
Calcite	org.apache.calcite	calcite-kafka	1.33.0.7.3.3.0-56
Calcite	org.apache.calcite	calcite-linq4j	1.33.0.7.3.3.0-56
Calcite	org.apache.calcite	calcite-plus	1.33.0.7.3.3.0-56
Calcite	org.apache.calcite	calcite-server	1.33.0.7.3.3.0-56
Calcite	org.apache.calcite	calcite-testkit	1.33.0.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-aliyun	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-annotations	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-archives	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-assemblies	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-auth	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-aws	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-azure	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-benchmark	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-build-tools	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-client	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-client-api	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-common	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-datajoin	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-distcp	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-dynamometer-blockgen	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-dynamometer-dist	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-dynamometer-infra	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-dynamometer-workload	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-extras	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-federation-balance	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-fs2img	3.4.2.7.3.3.0-56

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-gridmix	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-hdfs	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-kafka	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-kms	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-minicluster	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-minikdc	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-nfs	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-openstack	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-registry	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-rumen	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-sls	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-streaming	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-csi	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.4.2.7.3.3.0-56

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-yarn-server-globalpolicygenerator	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-documentstore	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.4.2.7.3.3.0-56
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.4.2.7.3.3.0-56
HBase	org.apache.hbase	hbase-annotations	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-asyncfs	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-backup	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-checkstyle	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-client	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-client-project	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-common	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-compression-aircompressor	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-compression-brotli	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-compression-lz4	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-compression-snappy	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-compression-zstd	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-endpoint	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-examples	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-external-blockcache	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-hadoop-compat	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-hadoop2-compat	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-hbtop	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-http	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-it	2.6.3.7.3.3.0-56

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-logging	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-mapreduce	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-metrics	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-metrics-api	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-openssl	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-procedure	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-protocol	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-protocol-shaded	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-replication	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-resource-bundle	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-rest	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-rsgroup	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-server	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-shaded-client	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-shaded-client-project	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-shaded-testing-util	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-shell	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-testing-util	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-thrift	2.6.3.7.3.3.0-56
HBase	org.apache.hbase	hbase-zookeeper	2.6.3.7.3.3.0-56
Hive	org.apache.hive	catalogd-unit	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-beeline	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-classification	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-cli	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-common	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-contrib	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-exec	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-hms-catalog	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-iceberg-catalog	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-iceberg-handler	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-iceberg-shading	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-impala	3.1.3000.7.3.3.0-56

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-it-iceberg	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-it-iceberg-integration	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-it-util	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-jmh	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-metastore	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-parser	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-serde	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-service	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-shims	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-streaming	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-testutils	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-udf	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	kafka-handler	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	patched-iceberg-api	3.1.3000.7.3.3.0-56
Hive	org.apache.hive	patched-iceberg-core	3.1.3000.7.3.3.0-56
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector-spark3_2.12	1.0.0.7.3.3.0-56
Kafka	org.apache.kafka	connect	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	connect-api	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	connect-basic-auth-extension	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.9.2.7.3.3.0-56

Project	groupId	artifactId	version
Kafka	org.apache.kafka	connect-cloudera-common	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	connect-file	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	connect-json	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	connect-mirror	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	connect-mirror-client	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	connect-runtime	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	connect-test-plugins	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	connect-transforms	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	generator	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-clients	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-examples	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-group-coordinator	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-group-coordinator-api	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-log4j-appender	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-metadata	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-raft	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-server	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-server-common	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-shell	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-storage	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-storage-api	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-examples	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-test-utils	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.9.2.7.3.3.0-56

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-34	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-35	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-36	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-37	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-38	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-tools	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-tools-api	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka-transaction-coordinator	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka_2.12	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	kafka_2.13	3.9.2.7.3.3.0-56
Kafka	org.apache.kafka	trogdor	3.9.2.7.3.3.0-56
Knox	org.apache.knox	gateway-adapter	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-admin-ui	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-applications	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-cloud-bindings	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-demo-ldap	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-demo-ldap-launcher	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-discovery-ambari	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-discovery-cm	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-docker	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-i18n	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-i18n-logging-log4j	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-i18n-logging-sl4j	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-openapi-ui	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-performance-test	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-ha	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-identity-assertion-common	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	2.1.0.7.3.3.0-56

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-identity-assertion-no-doas	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-jersey	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-rewrite	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-rewrite-common	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-security-authc-anon	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-security-authc-remote	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-security-authz-acls	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-security-authz-composite	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-security-authz-path-acls	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-security-clientcert	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-security-hadoopauth	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-security-jwt	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-security-pac4j	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-security-preauth	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-security-shiro	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-provider-security-webappsec	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-release	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-server	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-server-launcher	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-server-xforwarded-filter	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-admin	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-as	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-auth	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-definitions	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-hashicorp-vault	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-hbase	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-health	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-hive	2.1.0.7.3.3.0-56

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-service-idbroker	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-idbroker-plugins	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-impala	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-jkg	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-knoxsso	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-knoxssout	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-knoxtoken	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-livy	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-metadata	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-nifi	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-nifi-registry	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-remoteconfig	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-restcatalog	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-rm	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-session	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-storm	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-test	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-tgs	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-vault	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-service-webhdfs	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-shell	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-shell-launcher	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-shell-release	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-shell-samples	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-spi	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-spi-common	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-test	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-test-idbroker	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-test-release-utils	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-test-utils	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-topology-hadoop-xml	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-topology-simple	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-util-common	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-util-configinjector	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-util-launcher	2.1.0.7.3.3.0-56
Knox	org.apache.knox	gateway-util-urltemplate	2.1.0.7.3.3.0-56
Knox	org.apache.knox	hadoop-examples	2.1.0.7.3.3.0-56
Knox	org.apache.knox	knox-cli-launcher	2.1.0.7.3.3.0-56
Knox	org.apache.knox	knox-homepage-ui	2.1.0.7.3.3.0-56

Project	groupId	artifactId	version
Knox	org.apache.knox	knox-token-generation-ui	2.1.0.7.3.3.0-56
Knox	org.apache.knox	knox-token-management-ui	2.1.0.7.3.3.0-56
Knox	org.apache.knox	knox-webshell-ui	2.1.0.7.3.3.0-56
Knox	org.apache.knox	webhdfs-kerb-test	2.1.0.7.3.3.0-56
Knox	org.apache.knox	webhdfs-test	2.1.0.7.3.3.0-56
Kudu	org.apache.kudu	kudu-backup-common	1.18.0.7.3.3.0-56
Kudu	org.apache.kudu	kudu-backup-tools	1.18.0.7.3.3.0-56
Kudu	org.apache.kudu	kudu-backup3_2.12	1.18.0.7.3.3.0-56
Kudu	org.apache.kudu	kudu-client	1.18.0.7.3.3.0-56
Kudu	org.apache.kudu	kudu-flatbuffers	1.18.0.7.3.3.0-56
Kudu	org.apache.kudu	kudu-hive	1.18.0.7.3.3.0-56
Kudu	org.apache.kudu	kudu-proto	1.18.0.7.3.3.0-56
Kudu	org.apache.kudu	kudu-replication	1.18.0.7.3.3.0-56
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.18.0.7.3.3.0-56
Kudu	org.apache.kudu	kudu-spark3_2.12	1.18.0.7.3.3.0-56
Kudu	org.apache.kudu	kudu-subprocess	1.18.0.7.3.3.0-56
Kudu	org.apache.kudu	kudu-test-utils	1.18.0.7.3.3.0-56
Livy	org.apache.livy	livy-api	0.7.23000.7.3.3.0-56
Livy	org.apache.livy	livy-client-common	0.7.23000.7.3.3.0-56
Livy	org.apache.livy	livy-client-http	0.7.23000.7.3.3.0-56
Livy	org.apache.livy	livy-core_2.12	0.7.23000.7.3.3.0-56
Livy	org.apache.livy	livy-examples	0.7.23000.7.3.3.0-56
Livy	org.apache.livy	livy-integration-test	0.7.23000.7.3.3.0-56
Livy	org.apache.livy	livy-repl_2.12	0.7.23000.7.3.3.0-56
Livy	org.apache.livy	livy-rsc	0.7.23000.7.3.3.0-56
Livy	org.apache.livy	livy-scala-api_2.12	0.7.23000.7.3.3.0-56
Livy	org.apache.livy	livy-server	0.7.23000.7.3.3.0-56
Livy	org.apache.livy	livy-test-lib	0.7.23000.7.3.3.0-56
Livy	org.apache.livy	livy-thriftserver	0.7.23000.7.3.3.0-56
Livy	org.apache.livy	livy-thriftserver-session	0.7.23000.7.3.3.0-56
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-analyzers-opennlp	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.3.3.0-56

Project	groupId	artifactId	version
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-core	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-join	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.3.3.0-56
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.3.3.0-56
Oozie	org.apache.oozie	oozie-client	5.1.0.7.3.3.0-56
Oozie	org.apache.oozie	oozie-core	5.1.0.7.3.3.0-56
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.3.3.0-56
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.3.3.0-56
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.3.3.0-56
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.3.3.0-56
Oozie	org.apache.oozie	oozie-server	5.1.0.7.3.3.0-56
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.3.3.0-56
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.3.3.0-56
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.3.3.0-56
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.3.3.0-56
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.3.3.0-56
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.3.3.0-56
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.3.3.0-56
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.3.3.0-56
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.3.3.0-56
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.3.3.0-56

Project	groupId	artifactId	version
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.3.3.0-56
ORC	org.apache.orc	orc-core	2.0.3.7.3.3.0-56
ORC	org.apache.orc	orc-mapreduce	2.0.3.7.3.3.0-56
ORC	org.apache.orc	orc-shims	2.0.3.7.3.3.0-56
ORC	org.apache.orc	orc-tools	2.0.3.7.3.3.0-56
Ozone	org.apache.ozone	hdds-annotation-processing	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	hdds-cli-common	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	hdds-client	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	hdds-common	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	hdds-config	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	hdds-container-service	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	hdds-crypto-api	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	hdds-crypto-default	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	hdds-docs	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	hdds-erasurecode	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	hdds-interface-admin	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	hdds-interface-client	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	hdds-interface-server	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	hdds-managed-rocksdb	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	hdds-rocks-native	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	hdds-server-framework	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	hdds-server-scm	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	hdds-test-utils	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	mini-chaos-tests	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-cli-admin	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-cli-debug	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-cli-repair	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-cli-shell	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-client	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-common	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-csi	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-datanode	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-dev-support	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-dist	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-filesystem	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-filesystem-common	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-filesystem-hadoop2	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-filesystem-hadoop3	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-filesystem-shaded	2.2.0.7.3.3.0-56

Project	groupId	artifactId	version
Ozone	org.apache.ozone	ozone-freon	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-httpfsgateway	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-iceberg	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-insight	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-integration-test	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-integration-test-recon	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-integration-test-s3	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-interface-client	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-interface-storage	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-manager	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-mini-cluster	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-multitenancy-ranger	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-network-tests	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-recon	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-reconcodegen	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-s3-secret-store	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-s3gateway	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-tools	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	ozone-vapor	2.2.0.7.3.3.0-56
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	2.2.0.7.3.3.0-56
Parquet	org.apache.parquet	parquet-avro	1.16.0.7.3.3.0-56
Parquet	org.apache.parquet	parquet-column	1.16.0.7.3.3.0-56
Parquet	org.apache.parquet	parquet-common	1.16.0.7.3.3.0-56
Parquet	org.apache.parquet	parquet-encoding	1.16.0.7.3.3.0-56
Parquet	org.apache.parquet	parquet-format-structures	1.16.0.7.3.3.0-56
Parquet	org.apache.parquet	parquet-generator	1.16.0.7.3.3.0-56
Parquet	org.apache.parquet	parquet-hadoop	1.16.0.7.3.3.0-56
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.16.0.7.3.3.0-56
Parquet	org.apache.parquet	parquet-jackson	1.16.0.7.3.3.0-56
Parquet	org.apache.parquet	parquet-protobuf	1.16.0.7.3.3.0-56
Parquet	org.apache.parquet	parquet-thrift	1.16.0.7.3.3.0-56
Parquet	org.apache.parquet	parquet-tools	1.16.0.7.3.3.0-56
Parquet	org.apache.parquet	parquet-variant	1.16.0.7.3.3.0-56
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.6	5.2.1.7.3.3.0-56
Phoenix	org.apache.phoenix	phoenix-client-lite-hbase-2.6	5.2.1.7.3.3.0-56
Phoenix	org.apache.phoenix	phoenix-core	5.2.1.7.3.3.0-56
Phoenix	org.apache.phoenix	phoenix-core-client	5.2.1.7.3.3.0-56
Phoenix	org.apache.phoenix	phoenix-core-server	5.2.1.7.3.3.0-56
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.1	5.2.1.7.3.3.0-56

Project	groupId	artifactId	version
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.5.0	5.2.1.7.3.3.0-56
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.5.4	5.2.1.7.3.3.0-56
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.6.0	5.2.1.7.3.3.0-56
Phoenix	org.apache.phoenix	phoenix-mapreduce-byo-shaded-hbase-hbase-2.6	5.2.1.7.3.3.0-56
Phoenix	org.apache.phoenix	phoenix-pherf	5.2.1.7.3.3.0-56
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.3.3.0-56
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.3.3.0-56
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.3.3.0-56
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.3.3.0-56
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.3.3.0-56
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.6	5.2.1.7.3.3.0-56
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.2.1.7.3.3.0-56
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.3.3.0-56
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.3.3.0-56
Phoenix	org.apache.phoenix	phoenix5-spark3	6.0.0.7.3.3.0-56
Phoenix	org.apache.phoenix	phoenix5-spark3-shaded	6.0.0.7.3.3.0-56
Ranger	org.apache.ranger	authz-embedded	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	authz-remote	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	conditions-enrichers	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	credentialbuilder	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	embeddedwebserver	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	hdfs-nn-perf	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	jisql	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ldapconfigcheck	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-adls-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-atlas-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-audit-core	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-audit-dest-cloudwatch	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-audit-dest-hdfs	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-audit-dest-kafka	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-audit-dest-log4j	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-audit-dest-os	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-audit-dest-solr	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-audit-orc-util	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-authn	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-authz-api	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-caii-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-clo-plugin	2.8.0.7.3.3.0-56

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-common-ha	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-common-utils	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-distro	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-examples-distro	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-gs-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-hbase-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-hive-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-intg	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-kafka-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-kms	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-kms-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-knox-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-kudu-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-kylin-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-liquibase-upgrade	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-metrics	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-nestedstructure-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-nifi-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-opensearch-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-ozone-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-pdp	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-plugin-classloader	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-plugin-perf	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-plugins-common	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-plugins-cred	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-plugins-installer	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-policymigration	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-raz-adls	2.8.0.7.3.3.0-56

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-raz-hook-s3	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-raz-intg	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-raz-processor	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-raz-s3	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-raz-s3-lib	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-rms-common	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-rms-hive	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-rms-tools	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-rms-webapp	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-s3-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-solr-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-storm-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-tagsync	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-tools	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-trino-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-util	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-yarn-plugin	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	sample-client	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	sampleapp	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	ugsync-util	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	unixauthclient	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	unixauthservice	2.8.0.7.3.3.0-56
Ranger	org.apache.ranger	unixusersync	2.8.0.7.3.3.0-56
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.3.3.0-56
Solr	org.apache.solr	solr-analytics	8.11.2.7.3.3.0-56
Solr	org.apache.solr	solr-cell	8.11.2.7.3.3.0-56
Solr	org.apache.solr	solr-core	8.11.2.7.3.3.0-56
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.3.3.0-56
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.3.3.0-56

Project	groupId	artifactId	version
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.3.3.0-56
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.3.3.0-56
Solr	org.apache.solr	solr-langid	8.11.2.7.3.3.0-56
Solr	org.apache.solr	solr-ltr	8.11.2.7.3.3.0-56
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.3.3.0-56
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.3.3.0-56
Solr	org.apache.solr	solr-security-util	8.11.2.7.3.3.0-56
Solr	org.apache.solr	solr-solrj	8.11.2.7.3.3.0-56
Solr	org.apache.solr	solr-test-framework	8.11.2.7.3.3.0-56
Solr	org.apache.solr	solr-velocity	8.11.2.7.3.3.0-56
Spark	org.apache.spark	spark-avro_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-catalyst_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-common-utils_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-connect-client-jvm_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-connect-common_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-connect_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-core_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-graphx_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-hadoop-cloud_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-hive_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-kubernetes_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-kvstore_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-launcher_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-mllib-local_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-mllib_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-network-common_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-network-shuffle_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-network-yarn_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-protobuf_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-repl_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-shaded-raz	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-sketch_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-sql-api_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-sql-kafka-0-10_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-sql_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-streaming_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-tags_2.12	3.5.8.7.3.3.0-56

Project	groupId	artifactId	version
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-unsafe_2.12	3.5.8.7.3.3.0-56
Spark	org.apache.spark	spark-yarn_2.12	3.5.8.7.3.3.0-56
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.3.3.0-56
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.3.3.0-56
Tez	org.apache.tez	hadoop-shim	0.9.1.7.3.3.0-56
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.3.3.0-56
Tez	org.apache.tez	tez-api	0.9.1.7.3.3.0-56
Tez	org.apache.tez	tez-aux-services	0.9.1.7.3.3.0-56
Tez	org.apache.tez	tez-build-tools	0.9.1.7.3.3.0-56
Tez	org.apache.tez	tez-common	0.9.1.7.3.3.0-56
Tez	org.apache.tez	tez-dag	0.9.1.7.3.3.0-56
Tez	org.apache.tez	tez-examples	0.9.1.7.3.3.0-56
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.3.3.0-56
Tez	org.apache.tez	tez-history-parser	0.9.1.7.3.3.0-56
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.3.3.0-56
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.3.3.0-56
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.3.3.0-56
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.3.3.0-56
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.3.3.0-56
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.3.3.0-56
Tez	org.apache.tez	tez-tests	0.9.1.7.3.3.0-56
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.3.3.0-56
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.3.3.0-56
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.3.3.0-56
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.3.3.0-56
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.5.7.3.3.0-56
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.5.7.3.3.0-56
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.5.7.3.3.0-56
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.5.7.3.3.0-56
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.5.7.3.3.0-56
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.5.7.3.3.0-56
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.5.7.3.3.0-56
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.5.7.3.3.0-56
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.5.7.3.3.0-56
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.5.7.3.3.0-56
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.5.7.3.3.0-56