Hewlett Packard
Enterprise

HPE Ezmeral Data Fabric 6.2 Release Notes

Contents

Version 6.2.0 Release Notes

These release notes contain information about Version 6.2.0 of the Data Fabric Converged Data Platform.

What's New in Release 6.2.0

Release 6.2.0 of the HPE Ezmeral Data Fabric provides substantial new features for the components of the data platform. This page describes and provides links to more information about the new features.

For new features delivered as part of the Ecosystem Pack, see What's New in MEP 7.0.0.

Security and Governance Enhancements

Policy-Based Security

HPE Ezmeral Data Fabric supplies robust security. In release 6.2, you can leverage the new Policy-Based Security features to apply platform-level security across all data.

For an overview of data-fabric security features, see Security.

To learn about policy-based security, see Policy-Based Security.

External KMIP Keystore

Release 6.2 supports external keystores, through KMIP (Key Management Interoperability Protocol), to securely manage the CLDB and DARE master keys at the core platform layer. Release 6.2 supports KMIP versions 1.0 to 1.4 and has been tested with the Utimaco ESKM, Gemalto SafeNet KeySecure, and Vormetric DSM key management servers.

To learn about key management, see External KMIP Keystore Overview.

SSL Encryption for Server-to-Server Communication

In release 6.2.0, the ZooKeeper version is updated to 3.5.6. Version 3.5.6 supports SSL encryption for server-to-server communication. When you install a new release 6.2 secure cluster, SSL between ZooKeeper servers is enabled automatically. For more information, see configure.sh.

Resolved Security Vulnerabilities

For a list of vulnerabilities resolved in this release, see Resolved Issues on page 14.

New Products and Features

Data Tiering Enhancements

Data tiering was introduced in release 6.1. With release 6.2, you have access to enhanced data-tiering functionality, including Storage Labels, Erasure Coding with Local Parity, and parallel offload with the MAST gateway.

You can use an erasure-coding scheme with local parity to reduce erasure-coding storage overhead without incurring high rebuild costs and longer rebuild times, while lowering the probability of data loss. For information about erasure coding, including erasure coding with local parity, and setting the scheme using the Control System, see Erasure Coding Scheme for Data Protection and Recovery.

Storage labels help you confine volumes to specific storage pools. You can establish and apply storage labels to meet desired objectives, such as low latency or easy removability. For examples and an overview, see Using Storage Labels.

You can enable multiple MAST gateways to offload volume data rapidly. See Enabling Tiering.

The following two features have been released through a patch. The earliest patches with these features are: *mapr-apiserver-6.2.0.2.20210210081832-1.noarch.rpm* and *mapr-webserver-6.2.0.2.20210210081832-1.noarch.rpm*. To leverage these features, download the latest patch from sftp.mapr.com (login as maprpatches without a password).

To apply a patch to the Control System, follow the instructions in this Knowledge Article: How do I apply HPE Ezmeral Data Fabric EBFs for mapr-apiserver and mapr-webserver packages?

For additional information on patch nomenclature, see this support article.

- Offload rules support tiering based on Last Access Time.
- The Control System contains a Rule Builder to easily create tiering rules.

Ease of Data Management

Two new features support unlimited scale: Snapshot Restore and Last Access Time.

The Snapshot Restore optimization speeds up the restore operation when compared to manually copying the data to a new volume. See Restoring a Volume From a Snapshot.

Last Access Time (atime) is file metadata that is updated whenever a file is read. You can use atime for file management and governance decisions such as:

- Reviewing access controls of files that have not been accessed recently to prevent data leakage
- Tiering files (to warm or cold tier) that have not been accessed for a while
- Migrating files that have not been accessed frequently

For more information, see Tuning Last Access Time.

The URL Sharing feature has been released through a patch. The earliest patches with this feature are: *mapr-apiserver-6.2.0.2.20210210081832-1.noarch.rpm* and *mapr-webserver-6.2.0.2.20210210081832-1.noarch.rpm*. Using this feature, you can share URLs with filter information to other users on the cluster, who can then login and view the results. This saves time having to repeatedly re-create the filter for any user who wants to use the same filtering.

Hadoop and YARN Delivered with Ecosystem Components

Beginning with core 6.2.0 and MEP 7.0.0, Hadoop and YARN services are no longer included in the HPE Ezmeral Data Fabric repository for core packages. They are provided as ecosystem components in the MEP repository. This enables Hadoop and YARN to be upgraded separately from the HPE Ezmeral Data Fabric. To decouple Hadoop and YARN from core, the mapr-mapreduce2 package was removed, and some new packages were added. For more information, see Installing Hadoop and YARN.

Java / JDK 11 Support

Release 6.2 and the HPE Ezmeral Data Fabric ecosystem components take advantage of the new features and long-term support provided by JDK 11. For more information, see the JDK Support Matrix.

Controlling Metrics for JMX-Enabled Services

Release 6.2 enables you to control how metrics are collected from JMX-enabled services. For more information, see Controlling Access to JMX Metrics.

Monitoring

The Control System has been modified for Kubernetes environments. You will notice these modifications if you use the Control System with the HPE Ezmeral Container Platform. To learn more about the HPE Ezmeral Container Platform, see the 5.2 documentation.

NFSv4 Updated

Release 6.2 updates the NFSv4 package (mapr-nfsganesha) to version 3.3.

Deprecated Products and Features

MapR Data Science Refinery

The MapR Data Science Refinery is not supported in Releases 6.2.0 and later, but continues to be supported in Releases 6.0.x and 6.1.0. For more information about data science tools and applications, see HPE Ezmeral ML Ops.

Metering Support Removed

Metering is not supported and the documentation for metering has been removed. The Metering feature was added to support consumption-based billing. Monitoring and metrics capabilities continue to be supported.

Documentation Enhancements

Quickly Navigate to the Same Topic in a Different Release

Most documentation pages now include buttons to enable navigation to the same topic in a different release. These buttons appear in the upper right corner of the page:

To display the buttons, the publishing system compares the directory and file name of the current topic with all current releases. If a button is not present, the topic is new or does not exist in other releases.

New Component Versions Matrix

A new matrix has been added to the list of Interoperability Matrices. The Component Versions for Released MEPs matrix lets you compare the versions of ecosystem and Monitoring components across different MEPs.

Revised Product Naming

Release 6.2 introduces a new name for the data platform:

- HPE Ezmeral Data Fabric

In addition to the data platform, many former *MapR* products and features have new names. Product documentation and interfaces are being updated to reflect the new names.

Even though the product names are different:

- The platform works the same.
- Upgrades from *MapR* installations are supported unless noted otherwise.
- HPE Ezmeral features are fully compatible with legacy *MapR* features unless noted otherwise.

- Product versioning remains the same, with the exception of the four-digit versions used for core and patches. See Installation and Upgrade Notes (Release 6.2.0) on page 7.

Some product interfaces continue to use the term *MapR*. These interfaces may or may not be updated. The following table shows key terms that are changing in the product documentation for release 6.2.0:

MapR **Term**	**HPE Ezmeral Data Fabric Term**
MapR <release-number>	release <release-number>
MapR Academy	HPE Training
MapR admin	data-fabric admin
MapR client	data-fabric client
MapR cluster	data-fabric cluster
MapR Container for Developers	development environment for HPE Ezmeral Data Fabric
MapR Control System (MCS)	control system (MCS)
MapR core	core
MapR Data Access Gateway	data-access gateway
MapR Data Platform	HPE Ezmeral Data Fabric
MapR Data Platform for Kubernetes	Kubernetes Interfaces for Data Fabric
MapR Database	HPE Ezmeral Data Fabric Database
MapR Data Science Refinery	Data Science Refinery*
MapR Distribution for Apache Hadoop	HPE Ezmeral Data Fabric Distribution for Apache Hadoop
MapR Ecosystem Pack (MEP)	ecosystem pack (MEP)
MapR Edge	HPE Ezmeral Data Fabric Ege
MapR Event Store for Apache Kafka	HPE Ezmeral Data Fabric Event Store
MapR Filesystem	filesystem
MapR gateway	data-fabric gateway
MapR Hadoop	Hadoop for the HPE Ezmeral Data Fabric
MapR Installer	Installer
MapR Installer Stanza	Installer Stanza
MapR license	HPE Ezmeral Data Fabric license

MapR loopbacknfs POSIX client	loopbacknfs POSIX client
MapR Monitoring	monitoring
MapR NFS	NFS or NFS for the HPE Ezmeral Data Fabric
MapR package	data-fabric package
MapR Object Store with S3-Compatible API	Object Store with S3-Compatible API
MapR patch	patch
MapR Persistent Application Container Client (PACC)	Persistent Application Container Client (PACC)
MapR POSIX Client	POSIX client
MapR Professional Services	HPE Pointnext
MapR Sandbox	sandbox
MapR server ticket	server ticket
MapR services	services
MapR software	data-fabric software or HPE Ezmeral Data Fabric software
MapR Support	HPE Support
MapR Technologies	Hewlett Packard Enterprise Company
MapR ticket	ticket
MapR user	data-fabric user
MapR volume	volume
MapR XD Distributed File and Object Store	HPE Ezmeral Data Fabric XD Distributed File and Object Store
MapR-SASL	data-fabric SASL
MEP (MapR Ecosystem Pack)	MEP (ecosystem pack)

*Deprecated for release 6.2.0. For more information, see "Deprecated Products and Features" earlier on this page.

Installation and Upgrade Notes (Release 6.2.0)

This page describes considerations for installing or upgrading to release 6.2.0.

Installing Release 6.2.0 (All Methods)

Note these considerations for new installations of release 6.2.0, which apply regardless of the method you use to install the release:

Hadoop and YARN Are Provided as Ecosystem Components	Beginning with core 6.2.0 and MEP 7.0.0, Hadoop and YARN services are no longer included in the repository for core packages. They are provided as ecosystem components in the MEP repository. This enables Hadoop and YARN to be upgraded separately from the HPE Ezmeral Data Fabric. To decouple Hadoop and YARN from core, the mapr-mapreduce2 package was removed, and some new packages were added. For more information, see Installing Hadoop and YARN.
SUSE Support	HPE Ezmeral Data Fabric release 6.2.0 currently is not supported on SUSE Linux.
New MAPR_JMX Environment Variables and configure.sh options	Release 6.2.0 implements new MAPR_JMX environment variables and server configure.sh options to address the following security vulnerabilities: • CVE-2013-0450 • CVE-2013-0431 These environment variables and configure.sh options give you greater control over security aspects of JMX monitoring. Note the following considerations for new installations and upgrades: • On a new installation of release 6.1, remote host is enabled. On a new installation of release 6.2, local binding is enabled. On an upgrade, both local binding and remote JMX are enabled. • On release 6.2, Collectd always uses local binding and ignores remote host or local host settings. If you enable local JMX, Collectd continues to collect metrics, and if you enable remote JMX, Collectd still is able to collect it. But if you disable JMX, then Collectd does NOT collect JMX metrics from the JMX-enabled services. For more information, see Controlling Access to JMX Metrics.
Simplified Installation for Log Monitoring	Release 6.2 simplifies the manual installation of log-monitoring files by adding new certificates. All the log-stack security keys and certificates are generated when you run the server configure.sh command with the -genKeys option during cluster configuration. To configure log monitoring using the manual steps, see Step 9: Install Log Monitoring. To learn more about the new keystore and truststore files, see Understanding the Truststore and Keyfiles.
4-Digit Version Numbers for Core and Patches	Beginning with release 6.2.0, core and patch versions have four digits. The fourth digit adds precision not only for core versions, but also for JAR files and Maven artifacts. For the core version, the fourth digit represents a patch version. Even though newer core versions have four digits, most references to core versions in the data-fabric documentation use two or three digits. For more information, see Understanding Software Versions.

Installing Release 6.2.0 Using the Installer

Note these considerations for new installations of release 6.2.0 using the Installer:

Ubuntu 16.04 Installations Require Manual Installation of JDK Before Using Installer	Before using Installer 1.14 on Ubuntu 16.04 nodes, you must manually install the JDK. If you are using Installer 1.14 on RHEL/CentOS, the Installer installs OpenJDK 11 for you. For more information, see the Installer Prerequisites and Guidelines.

KMIP Installation Is Not Supported Using the Installer	Installer 1.14 does not support installing or configuring the External KMIP Keystore functionality. If you want to use KMIP and you are performing a new installation using the Installer, you must first perform a regular installation. Then you must use the mrhsm utility to import the generated CLDB and DARE keys into the hardware security module (HSM), and copy the contents of the ${MAPR_HOME}/conf/tokens directory to all CLDB nodes in the cluster. For more information, see External KMIP Keystore Overview and Setting Up the External KMIP Keystore

Upgrading to Release 6.2.0 (High-Level Steps)

Depending on your current data-fabric release, upgrading to release 6.2.0 can require multiple steps. Upgrading to release 6.2.0 can require an OS upgrade and always requires a JDK upgrade. The following table summarizes the high-level upgrade steps:

If your cluster is running	Use these steps to upgrade to release 6.2.0	See for more information
6.1.0	**1.** Upgrade your MEP to MEP 6.3.1. **2.** Upgrade your OS to an OS that is supported by release 6.2.0. For example: • RHEL/CentOS 8.2 or 8.3 • Ubuntu 18.04 • Ubuntu 16.04 **3.** Install Java JDK 11 or the equivalent. **4.** Upgrade from release 6.1.0 to 6.2.0 using one of the upgrade workflows.	• Upgrading Ecosystem Packs • Java • Upgrading Your Red Hat, CentOS, or Ubuntu Operating System • Operating System Support Matrix (Release 6.x) • Upgrade Workflows (Release 6.1 to 6.2)
5.2.x or 6.0.x	**1.** Upgrade to release 6.1.0 and MEP 6.3.1 using the release 6.1.0 documentation. See Installation and Upgrade Notes (Release 6.2.0). Then return to the upgrade information here to upgrade from release 6.1.0 to 6.2.0. **2.** Upgrade your OS to an OS that is supported by release 6.2.0. For example: • RHEL/CentOS 8.2 or 8.3 • Ubuntu 18.04 • Ubuntu 16.04 **3.** Install Java JDK 11 or equivalent. **4.** Upgrade from release 6.1.0 to 6.2.0 using one of the upgrade workflows.	

Upgrade Considerations (All Upgrade Methods)

Core and MEP Requirements for Upgrading	Clusters to be upgraded to release 6.2.0 must be installed with release 6.1.0 and MEP 6.3.1 before upgrading, as indicated in the high-level upgrade steps earlier on this page. If your cluster is on release 5.2.x or 6.0.x, you must first upgrade to release 6.1.0 and MEP 6.3.1.
Upgrade Support for SUSE Is Pending	Some upgrade topics include SUSE commands; however, release 6.2 is not currently supported on SUSE platforms. Therefore, upgrades from those platforms to release 6.2.0 are not supported. This page will be updated when SUSE support is announced.

Upgrades to Red Hat / CentOS 8.1	Upgrades from Red Hat / CentOS 7.x to Red Hat / CentOS 8.1 are not supported. To take advantage of the new features in release 6.2.0 by way of an upgrade, you must upgrade to Red Hat / CentOS 8.2 or 8.3. For more information about OS upgrade paths, see the Red Hat documentation.
NFSv4 Upgrade Restrictions	NFSv4 is not supported on release 6.1.0 clusters on CentOS 8.2. Because of this restriction, a fresh install of release 6.1.0 on CentOS 8.2 is not supported for clusters using NFSv4. An OS-only upgrade of release 6.1.0 from CentOS 7.x to release 6.1.0 on CentOS 8.2 also is not supported. To support NFSv4 on CentOS 8.2, a 6.1.0 cluster must be upgraded to release 6.2 using one of the following procedures: • **NFSv4 Gateway-Only Node** **1.** Stop NFSv4 on the node. **2.** Upgrade the OS to CentOS 8.2 and the DF to release 6.2. **3.** Restart NFSv4 on the node. • **NFSv4 Data Node** **1.** Stop everything on the node. **2.** Upgrade the OS to CentOS 8.2 and the DF to release 6.2. **3.** Bring the node back up.
Hadoop and YARN Are Provided as Ecosystem Components	Beginning with core 6.2.0 and MEP 7.0.0, Hadoop and YARN services are no longer included in the repository for core packages. They are provided as ecosystem components in the MEP repository. If you are upgrading and need Hadoop and YARN services, you must install the packages as ecosystem components after upgrading. For more information, see Installing Hadoop and YARN.
Data-on-wire-encryption	Beginning with release 6.1, data-on-wire encryption is enabled by default for newly created volumes on secure clusters. Data-on-wire encryption encrypts data in a volume during transmission over the wire. Data-on-wire encryption is *not* supported for non-secure clusters. When you upgrade a cluster to a release 6.1.0 secure cluster, data-on-wire encryption is enabled by default. You can disable data-on-wire encryption for individual volumes using MCS, the maprcli, or REST API commands.
Release 6.2.0 and MEP 7.0.0	Release 6.2.0 requires MEP 7.0.0 or later. For MEP compatibility information, see MEP Support by Core Version.
Regenerating the mapruserticket File	Changes to the CanImpersonate parameter of the mapruserticket file in release 6.1.0 require users who upgrade manually to regenerate the file before restarting Warden. See Step 1: Restart and Check Cluster Services. The file needs to be regenerated to ensure that impersonation works correctly for non-mapr users. Prior to release 6.1.0, all mapruserticket files were generated with CanImpersonate = false. Release 6.1.0 and later enforce the CanImpersonate parameter and sets the parameter to true for freshly installed clusters. For upgraded clusters, if CanImpersonate is not set to true, some services will not be able to impersonate.
Enabling PBS	If you are upgrading from a data-fabric version that does not support extended attributes, you must enable extended attributes before you enable policy-based security (PBS): `/opt/mapr/bin/maprcli cluster feature enable -name mfs.feature.fileace.support`

	To enable PBS: `/opt/mapr/bin/maprcli cluster feature enable -name mfs.feature.pbs` For more information, see Step 4: Enable New Features and Policy-Based Security.
Professional Support for Upgrades	Upgrading can be time-consuming and complicated. Consider engaging HPE professional support services to assist in planning and executing your upgrade. For more information, contact your support representative.

Upgrade Considerations (Installer Upgrades)

Note these considerations for using the MapR Installer to upgrade to release 6.2.0:

Installer 1.14.0	Before upgrading, update the MapR Installer to version 1.14.0.x or later. Installer 1.14 is compatible with earlier data-fabric releases and MEPs, and only Installer 1.14 or later can be used with data-fabric release 6.2.0 and MEP 7.0.0.
Most Upgrades to Release 6.2.0 Require an OS Upgrade	The Installer version upgrade feature can upgrade the data-fabric core version, the Java version, and ecosystem components. However, the Installer does not upgrade the OS version. In most cases, upgrading to core 6.2.0 requires an OS upgrade. For example, if your cluster is installed with core 6.1.0 running on CentOS 7.8, you must upgrade the OS to CentOS 8.2 before using the Installer to upgrade data-fabric software to core 6.2.0. See When Upgrading Core with the Installer Requires an OS Upgrade. You can use the Installer to upgrade directly to core 6.2.0 ONLY if your cluster is running an OS version that core 6.2 supports. These OS versions currently include: • Red Hat / CentOS 8.2 • Ubuntu 18.04 • Ubuntu 16.04 For information about supported OS versions, see Operating System Support Matrix (MapR 6.x).
Java Version Upgrade	All upgrades to core 6.2.0 require upgrading the Java version to JDK 11. The Installer upgrades the Java version if it detects that an upgrade is needed. For more information about Java requirements, see the JDK Support Matrix.

Operational Changes (Release 6.2.0)

Lists the functional changes made to existing commands in HPE Ezmeral Data Fabric release 6.2.0.

In release 6.2.0, there are no changes to commands that prevent existing scripts to run.

Known Issues at Release (Release 6.2.0)

You may encounter the following known issues after upgrading to Release 6.2. This list is current as of the release date. For a dynamic list of all known issues for all HPE Ezmeral Data Fabric product releases, refer to EZMERAL DATA FABRIC SUPPORT PORTAL

Where available, the workaround for an issue is also documented in this topic. HPE regularly releases maintenance releases and patches to fix issues. We recommend checking the release notes for any subsequent maintenance releases to see if one or more of these issues are fixed.

Installation and Configuration Issues

You can see generic installation issues here: Installer Known Issues.

SPYG-1136

During a manual installation or upgrade, Collectd provided in core 6.1.0 won't start on RHEL / CentOS 8.2 because it expects the Python 2 libraries to be installed, and RHEL / CentOS 8.2 provides the Python 3 libraries instead. This issue does not affect installations or upgrades performed using the Installer.

Workaround: Before installing the monitoring components, check to see if Python 2 is installed. If the following error is generated, try installing Python 2 on RHEL / CentOS 8.2:

```
failed: libpython2.7.so.1.0: cannot open shared object file
```

IN-2637

After a manual installation, Oozie and Hive services can fail to connect to a MySQL or MariaDB database because the server time-zone value is unrecognized or represents more than one time zone. This issue affects manual installations but is fixed in Installer 1.14.0.0.

Workaround: For manual installations, you must configure either the server or JDBC driver (using the serverTimezone configuration property) to use a more specific time-zone value if you want to utilize time-zone support. After running configure.sh but before starting the Oozie or Hive services, update the serverTimezone parameter in the hive-site.xml or oozie-site.xml. For more information, see MySQL Bug #95036.

Performance Issues

MFS-10838

In release 6.2.0, the default setting for Java garbage collection (GC) is ParallelGC. Garbage-collection settings can influence performance. If you notice a difference in the performance of Spark or MapReduce jobs as compared to release 6.1.0, try reverting to the G1GC setting, which is the Java 11 default. You can do this by overwriting the Java properties for the am, map, and reduce containers in the mapred-site.xml file:

```
<property>
  <name>yarn.app.mapreduce.am.command-opts</
name>
  <value>-Xmx1024m --add-opens java.base/
java.lang=ALL-UNNAMED</value>
</property>

<property>
  <name>mapreduce.map.java.opts</name>
  <value>-Xmx900m --add-opens java.base/
java.lang=ALL-UNNAMED</value>
</property>

<property>
  <name>mapreduce.reduce.java.opts</name>
  <value>-Xmx2560m --add-opens java.base/
java.lang=ALL-UNNAMED</value>
</property>
```

To change the garbage-collection (GC) default setting:

1. Comment out the following lines in hadoop_home/etc/hadoop/yarn-env.sh and mapred-env.sh:

   ```
   YARN_OPTS="$YARN_OPTS -XX:+UseParallelGC"

   HADOOP_OPTS="$HADOOP_OPTS -XX:+UseParallelGC"
   ```

2. Restart Hadoop services, as described in node services.

For more information about garbage collection, see Types of Java Garbage Collectors.

Timeline Server

See [].

Permission Issues

MFS-6776

Teragen jobs fail to complete when run as the root user, because of a permission issue on the application directory. The directory should have permission-mode bits 750 and not 700 as it is at present.

Workaround: Set the permission-mode bits on this directory to 750.

Control System Issues

MON-5653

The system displays an incorrect number of failed nodes for the NFSv4 service on the Services page when the failed node count is clicked.

Workaround: Use the maprcli node list command to get the node count.

MAPRDB-2309

Setting permissions on Column Families from the Control System causes issues with table fields.

Workaround: Set permissions using the maprcli table cf colperm set command directly.

CORE-538

ACL is not removed when using the Control System to delete an ACL from a volume with multiple ACLs.

Workaround: None

MON-5933

Table creation from the Control System fails when both public and a user are selected for permissions.

Workaround: None

MapR-DB Issues

MAPRDB-2092

In the HPE Ezmeral Data Fabric Database, adding a table index or replicating a table fails if the cluster administrator (MAPR_USER) does not have write access to the parent volume of the table.

Workaround: For more information, see this support advisory. If a patch is not available or the available patch has not been applied, another workaround is to add the MAPR_USER to the writeACE for the table parent volume.

Resolved Issues

Lists the issues that were resolved in HPE Ezmeral Data Fabric release 6.2.0.

The following issues, which were reported by customers, are resolved in release 6.2.0.

Filesystem Issues

CORE-290	Fixed permissions on files and directories in the /opt/mapr directory.
CORE-293	The mapr-zookeeper and mapr-warden processes were not properly started with systemd when the OS security updates were installed.
CORE-387	MapR user ticket got overwritten when Kerberos authentication is enabled.
CORE-472	Fixed an issue where the disklist.sh script was slow on nodes with a high number of udev symlinks.
CORE-476	Fixed the empty MAPR_JMXAUTH value in the env_override.sh script that prevented NodeManager from starting on custom secure clusters.
MFS-1984	The maprcli dashboard info command returned incorrect statistics about compressed and uncompressed data.
MFS-2019	apiserver crashed intermittently when there were multiple networks in MapR subnets.
MFS-2055	CLDB crashed while processing alarms due to an error when opening one of the clearedAlarm tables in kvstore.
MFS-2079	All update operations for the main volume Name container failed due to incorrect disk permissions.
MFS-2143	Excluded audit data operations were not preserved on MFS restart.
MFS-2209	Fixed a Java exception in the getVolName() function.
MFS-2211	Primary container instance was frozen during resync of orphan entries.
MFS-2266	Remote wire-level secure Read RPC of unaligned compressed data failed.
MFS-2294	Fixed a CLDB exception that occurred when adding NFS version 4 nodes.
MFS-2306	Fixed a CLDB exception that occurred when nodes were removed.
MFS-2307	Added a cluster level flag to not offline SP when Read CRC error is encountered.
MFS-2392	GFSCK failed on a secure cluster due to missing jackson-core jar file.
MFS-2444	FUSE failed to remove shared memory segments when running on a Kubernetes cluster.
MFS-2462	Fixed DBworker thread crashes when looking up ACEs at the column level.

MFS-2551	Fixed a Local Privilege Escalation vulnerability in maprexecute.
MFS-2553	Ecosystem jobs using ZooKeeper failed with the java.lang.NoSuchMethod error after applying the 6.1-EBF patch.
MFS-2554	Fixed a vulnerability where the cldbautdit.log file stored the ldap.bindpassword in cleartext.
MFS-2608	Priorities of child processes were not in sync with that of the parent process when the priority of the parent process is changed with the renice Linux command.
MFS-2610	Fixed hangs with Persistent Volume Mounts when the MapR ticket expired.
MFS-2623	Resolved a CLDB deadlock between two threads in the ListSorterPurgeTask resource.
MFS-2631	Fixed CLDB shut down with ConcurrentModificationException when accessing Server.nfsIds.
MFS-2638	Store alarms sorted by alarm type in CLDB, to avoid sorting each time the maprcli alarm list -sortby alarmtype command is used.
MFS-2691	Optimized HashTable walk for enhanced performance when fetching Muted and RaisedAlarms.
MFS-2695	Cross cluster mirror failed with an *Operation not Supported* error after enabling the snapshot lite feature.
MFS-2700	The FUSE kernel sent the wrong credentials to the MapR FUSE Process.
MFS-2708	Disk failure related log files were world writable causing a potential security vulnerability.
MFS-2711	Freed Up CLDB CPU from background activity. Now expired snapshots are removed only after the cluster is stable.
MFS-2720	Added a maprcli dump command to print in-progress container replication.
MFS-2732	RPC connections between MFS and CLDB were failing.
MFS-2757	MapR services displayed incorrect services when systemd was updated.
MFS-2767	FileClient retried the same CLDB node on failure.
MFS-2968	Inactive alarms returned an incorrect alarm value.
MFS-3291	The number of flusher threads for MapR Streams was not configurable. The default value of 64 threads was insufficient.
MFS-3310	Added alert to warn customers about expiring SSL certificates.
MFS-4480	Fixed intermittent NFS version 4 server crashes.
MFS-4531	Mirror source snapshots were not deleted after mirroring completed.
MFS-4532	Jobs failed with IO error as MFS processes could not establish connections with each other, resulting in data remaining under replicated for long periods of time until manual intervention was done.
MFS-4597	Dependent services failed to start on a Warden restart.

MFS-4562	The stat command on FUSE mounted files reported an incorrect IO block size.
MFS-4605	CDB crashed when ACLs larger than 2K bytes are set.
MFS-4670	The CLDB process consumed massive amounts of memory at periodic intervals and caused CLDB to fail.
MFS-4752	Fixed assert in RpcBinding::GetCredentials.
MFS-4776	FUSE RPCs failed when a node's hostid is changed.
MFS-4805	Implemented a memory tracker for the NFS version 3 server.
MFS-5236	Fixed a Remote Code Execution vulnerability in the Warden Java JMX server.
MFS-5235	Fixed a Remote Code Execution vulnerability in the Gateway Java JMX server.
MFS-5234	Fixed a Remote Code Execution vulnerability in the CLDB Java JMX server.
MFS-5229	Fixed a Remote Code Execution vulnerability in the MAST Gateway Java JMX server.
MFS-5356	The getAces() method returned a Null Pointer Exception when called on a non-existent object.
MFS-5422	The create() method was creating files with the wrong permissions.
MFS-5430	NFS server could not parse export files with line lengths greater than 8192 characters.
MFS-5482	Fixed memory leak in CLDB.
MFS-5488	Fixed intermittent failing of FileClient.
MFS-5502	PathWalk error occurred when a client node is upgraded but CLDB is of an older version.
MFS-5710	Added log rotation for the mfs.err and the mfs.out log files to avoid filling local OS partition.
MFS-5724	The MFS configuration parameter mfs.max.restore.count was not being honored causing mirror resync operations to be delayed due to the lack of sufficient restore slots.
MFS-5732	Change the severity of log messages so as to not unduly alarm customers.
MFS-6666	Implement a throttle to prevent the NFS server from overwhelming CLDB with too many RPCs.
MFS-6717	Added lazy unmount option to allow NFS mounts to be unmounted when they are not busy.
MFS-6748	Automatic Offload failed for EC Volumes.
MFS-6785	The mrconfig info containers rw command was too slow to respond on a cluster with a large number of volumes.
MFS-6873	File operations across clusters failed when using FUSE.
MFS-7126	The maprcli disk remove command timed out on a cluster with a large number of volumes.
MFS-8454	Volume creation failed with the error: *Could not get status of mount path*.

MFS-8459	Fixed volume access problems for volumes that reused the volume ID of deleted volumes.
MFS-8475	The createsystemvolumes.sh script took hours to complete when adding a new node to a cluster with a large number of volumes.
MFS-10328	The maprlogin renew command failed to refresh updated group membership details.
MFS-11109	Core dumps were generated frequently due to a memory leak issue.
MFS-11221	Fixed Drill query crashes.

HPE Ezmeral Data Fabric Streams Issues

MS-811	Segmentation fault occurred when deleting topics from monitoring streams.

HPE Ezmeral Data Fabric YARN Issues

MAPRYARN-161	The History server deletion thread stopped once it found an invalid application directory.

HPE Ezmeral Data Fabric Hadoop Issues

MAPRHADOOP-102	Fixed an error that occurred when copying ACES from the MapR-FS to the local FS.

HPE Ezmeral Data Fabric Database Issues

MAPRDB-1520	OR conditions on array field path in correlation tracker were not handled correctly.
MAPRDB-1985	Fixed a segmentation fault on the mapr dbshell utility.
MAPRDB-2062	Failed to scan table on a remote secure cluster using the mapr dbshell utility because of a wrong ticket that was sent to ZooKeeper.
MAPRDB-2096	Enhanced the performance of HPE Ezmeral Data Fabric Database for workloads with a large number of threads.
MAPRDB-2113	Fixed the database logic to select the index appropriate to the query being executed, when there are multiple indexes over the same field in a table.
MAPRDB-2125	OJAI APIs failed to connect to ZooKeeper after applying the 6.1-EBF patch.
MAPRDB-2156	When running queries with set timeout, the number of threads on the MapR client increased exponentially and caused the client to stop responding and become completely blocked.
MAPRDB-2201	Fixed a memory leak caused by a dangling reference of MetaTable in the BaseJson table.
MAPRDB-2250	Added a throttle for internal DB operations to prevent the cluster from becoming unstable.

Upgrade Issues

ES-58 Upgrading to MEP 6.0 on MapR 6.1.0 from MapR 5.2.1 with MEP 3.0.3 broke Elasticsearch and FluentD.

MS-925 Upgrading to MEP 6.2 (Spark 2.4.0) broke consumption of Kafka / HPE Ezmeral Data Fabric Streams.

MFS-2469 CLDB secondary node was stuck in BECOMING_SLAVE mode during rolling upgrade from MapR 5.2.2 to MapR 6.1.0.

MFS-2561 Increased the value of the cldb.datacontainer.block.updates.max.sec parameter to avoid overwhelming the cluster with the VOLUME_ALARM_DATA_UNDER_REPLICATED alarm after upgrade from MapR version 3.0.1.

Control System Issues

MON-1231 Pagination functionality was not working as expected when changes to the number of rows are made in pages other than page 1.

MON-1968 Was not able to edit user email property in Volume User Disk Usage page.

MON-2018 The Add Field option was disabled on the Add Change Log page.

MON-2076 Removing the last stream topic from the Remove Topic Window list caused the remove topic window to be closed, which was not the expected behaviour.

MON-2321 On the Service List page, the node pop-over did not display the list of nodes if Elasticsearch had failed.

MON-2374 The Node Details page did not display logs for nodes that have two IPs.

MON-2409 The Control System did not display any feedback about page loading when navigating from one cluster to another.

MON-3404 The Search field for a stream topic was not displayed properly on the Firefox browser (Firefox version 58).

MON-3438 Page response was slow between clicks on the Alarms page.

MON-3452 Volumes Set Quota action cleared the system volumes but kept their checkboxes selected.

MON-3476 Users with admin permission were not able to modify table ACLs.

MON-3709 The Service tab did not auto-refresh as per the refresh configuration settings.

MON-3817 The Topology view on the Nodes page did not refresh until one switched to another tab and back.

MON-3922 Volumes prior to MapR 6.0 that lack volume ACEs were not displayed after being upgraded to version 6.x.

MON-4701 Data was not returned for percentiles in the activity grid if the end time of the query was in the future.

MON-4776 Oops error was displayed when mounting/unmounting multiple volumes that were selected across volume grid pages.

MON-4845 Could not start the apiserver as a user other than the mapr user.

MON-4862	The Control System did not start after an upgrade to MapR 6.x if the installed license did not have table support enabled.
MON-4874	When the MAPR_EXTERNAL environment variable was set, the Control System was not accessible and returned the following error: HTTP ERROR 503.
MON-4892	The Snapshot tab on the Control System indicated that a license upgrade is needed after upgrading from MapR 5.2.1 to MapR 6.1 even with an M5 enterprise license (license permits all operations except tables/streams handling) already installed.
MON-5011	The Control System reported an incorrect value for CPU utilization when Drill queries were executed on the cluster.
MON-5041	Added a security mechanism to prevent Cross Site Request Forgery (CSRF) attacks on the Control System.
MON-5057	Table properties could not be modified from the Control System due to insufficient permissions.
MON-5075	Was unable to modify existing volume quota using the Control System UI due to insufficient permissions.
MON-5116	Fixed multiple Control System issues pointed out by a customer.
MON-5139	Accessing the REST endpoints caused the apiserver to freeze after a certain number of requests.
MON-5194	The CPU utilization graph on the Control System had an incorrect scale on the Y-axis.
MON-5201	The Control System failed to display Drill queries that were completed within the last 10 minutes.
MON-5401	The Control System crashed on one of the nodes.
MON-5275	The Control System failed to display MAC addresses on a multi-NIC environment.
MON-5412	The Control System failed to display the resourcemanager service under the MEP services list due to an unsupported shell command.
MON-5774	The apiserver crashed when tmp directories were mounted with the noexec flag.
MON-5488	An IO exception occurred when loading metadata of huge tables in the Control System.

Packages and Dependencies for Data Fabric Software

This section describes package and dependency details for the Release 6.2.0 core and ecosystem components.

For downloadable packages, see these links:

- Core Packages
- MEP Packages
- Installer Packages

For core package dependencies for the supported OS distributions, see:

- Package Dependencies

For Installer package dependencies, see:

- Installer Prerequisites and Guidelines

Patches and Documentation

Describes important considerations for patches and patch documentation.

Whenever possible, keep your software up to date by applying the latest patches available on the Support Portal. This practice can help you to resolve issues and minimize downtime.

Some patches enable new features or behaviors that are described in the documentation. However, the data-fabric documentation does not typically include patch numbers or identify the features or behaviors that are delivered by specific patches. If you see a fix or feature in the documentation that is not available on your platform, you might need to apply a patch in order to use the fix or feature.

To understand which patches apply to your platform, contact your support representative, or visit the EZMERAL DATA FABRIC SUPPORT PORTAL. For information about applying a patch, see Applying a Patch.