

Release Notes for Patches for the MapR 4.1.0 Release

Release Notes for the December 2016 Patch

Released 12/09/2016

These release notes describe the fixes that are included in this patch.

Packages

Red Hat	Server	mapr-patch-4.1.0.31175.GA-40888.x86_64.rpm
Red Hat	Client	mapr-patch-client-4.1.0.31175.GA-40888.x86_64.rpm
Red Hat	Loopbacknfs	mapr-patch-loopbacknfs-4.1.0.31175.GA-40888.x86_64.rpm
Ubuntu	Server	mapr-patch-4.1.0.31175.GA-40888.x86_64.deb
Ubuntu	Client	mapr-patch-client-4.1.0.31175.GA-40888.x86_64.deb
Ubuntu	Loopbacknfs	mapr-patch-loopbacknfs-4.1.0.31175.GA-40888.x86_64.deb
Win64	Client	mapr-client-4.1.0.40888GA-1.amd64.zip
Mac	Client	mapr-client-4.1.0.40888GA-1.x86_64.tar.gz

Fixes

Bug 24139

Description

If limit spread was enabled and the nodes were more than 85% full, CLDB did not allocate containers for IOs on non-local volumes.

Resolution

With this fix, CLDB will now allocate new containers to ensure that the IO does not fail.

Bug 24969

Description

The `maprcli volume create` command was not setting group ownership to user's primary group when the user's primary GID was not the first GID in the list of GIDs.

Resolution

With this fix, the primary GID of the user performing the operation will now be the first GID in the list of GIDs.

Bug 24971

Description

When the mirroring operation started after a CLDB failover, sometimes it was sending requests to slave CLDB where data was stale, resulting in the the mirroring operation hanging. If the CLDB failover happened again during this time, the new CLDB master was discarding data resynchronized by the old mirroring operation, but marking the mirroring operation as successful. This resulted in data mismatch between source and destination.

Resolution

With this fix, mirroring requests will be sent to master CLDB node only.

Bug 25041

Description

Whenever a newly added node was made the master of the name container, MFS crashed while deleting files in the background.

Resolution

With this fix, MFS will not crash when a newly added node is made the master of the name container.

Release Notes for the October 2016 Patch

Released 10/24/2016

These release notes describe the fixes that are included in this patch.

Packages

Red Hat	Server	mapr-patch-4.1.0.31175.GA-40098.x86_64.rpm
Red Hat	Client	mapr-patch-client-4.1.0.31175.GA-40098.x86_64.rpm
Red Hat	Loopbacknfs	mapr-patch-loopbacknfs-4.1.0.31175.GA-40098.x86_64.rpm
Ubuntu	Server	mapr-patch-4.1.0.31175.GA-40098.x86_64.deb
Ubuntu	Client	mapr-patch-client-4.1.0.31175.GA-40098.x86_64.deb
Ubuntu	Loopbacknfs	mapr-patch-loopbacknfs-4.1.0.31175.GA-40098.x86_64.deb
Win64	Client	mapr-client-4.1.0.40098GA-1.amd64.zip
Mac	Client	mapr-client-4.1.0.40098GA-1.x86_64.tar.gz

Fixes

Bug 14105

Description

When nodes attempt to register with duplicate IDs, CLDB does not register the nodes and log meaningful error messages.

Resolution

With this fix, when nodes attempt to register with duplicate IDs, CLDB will log appropriate error messages.

Bug 20965

Description

When working with multiple clusters, synchronization issues was causing MapRFileSystem to return NullPointerException.

Resolution

With this fix, MapRFileSystem has been improved to better support working with multiple clusters and MapRFileSystem contains fixes for synchronization issues.

Bug 24562

Description

CLDB (container location database) performance suffered because Warden gave the CLDB service a lower CPU priority.

Resolution

With this fix, Warden uses a new algorithm to set the correct CPU priority for the CLDB service.

Bug 24651

Description

CLDB threw an exception and failed over when the snapshots list was iterated over while snapshots were being created.

Resolution

With this fix, CLDB will no longer fail over when snapshots list is iterated over while new snapshots are being created.

Bug 24656

Description

MFS was churning cpu while taking snapshot because of some debug code in the builds.

Resolution

With this fix, MFS will no longer churn CPU as the debug code has been disabled.

Release Notes for the September 2016 Patch

Released 9/17/2016

These release notes describe the fixes that are included in this patch.

Packages

Red Hat	Server	mapr-patch-4.1.0.31175.GA-39743.x86_64.rpm
Red Hat	Client	mapr-patch-client-4.1.0.31175.GA-39743.x86_64.rpm
Red Hat	Loopbacknfs	mapr-patch-loopbacknfs-4.1.0.31175.GA-39743.x86_64.rpm
Ubuntu	Server	mapr-patch-4.1.0.31175.GA-39743.x86_64.deb
Ubuntu	Client	mapr-patch-client-4.1.0.31175.GA-39743.x86_64.deb
Ubuntu	Loopbacknfs	mapr-patch-loopbacknfs-4.1.0.31175.GA-39743.x86_64.deb
Win64	Client	mapr-client-4.1.0.39743GA-1.amd64.zip
Mac	Client	mapr-client-4.1.0.39743GA-1.x86_64.tar.gz

Fixes

Bug 24053

Description

During client initialization, the client crashed if there was an error during initialization.

Resolution

With this fix, the client will not crash if there is an error during initialization.

Release Notes for the August 2016 Patch

Released 8/27/2016

These release notes describe the fixes that are included in this patch.

Packages

Red Hat	Server	mapr-patch-4.1.0.31175.GA-39341.x86_64.rpm
Red Hat	Client	mapr-patch-client-4.1.0.31175.GA-39341.x86_64.rpm
Red Hat	Loopbacknfs	mapr-patch-loopbacknfs-4.1.0.31175.GA-39341.x86_64.rpm
Ubuntu	Server	mapr-patch-4.1.0.31175.GA-39341.x86_64.deb
Ubuntu	Client	mapr-patch-client-4.1.0.31175.GA-39341.x86_64.deb
Ubuntu	Loopbacknfs	mapr-patch-loopbacknfs-4.1.0.31175.GA-39341.x86_64.deb
Win64	Client	mapr-client-4.1.0.39341GA-1.amd64.zip
Mac	Client	mapr-client-4.1.0.39341GA-1.x86_64.tar.gz

Fixes

Bug 21930

Details

hadoop-2.5.1 RM GUI showed values that are too high for reserved memory and cores; actual values were substantially lower.

Resolution:

With this fix, the GUI shows accurate values.

Bug 24022

Details

Mirroring of a volume on a container which does not have a master container caused the mirror thread to hang.

Resolution

With this fix, mirroring will not hang when the container associated with the volume has no master.

Bug 24063

Details

During mirroring, the volume property update messages were logged (in cldb.log) every 5 seconds because the log level was set to INFO.

Resolution

With this fix, the log level is now DEBUG and the log will not contain multiple volume property update messages from mirroring.

Bug 24140

Details

While trying to access secure cluster from Windows, if the MAPR_TICKETFILE_LOCATION pointed to an incorrect location, the MapRClient threw an assert and the application crashed.

Resolution

With this fix, the application will not crash if the MAPR_TICKETFILE_LOCATION points to an incorrect location. Instead, the application will exit gracefully.

Release Notes for the July 2016 Patch

Released 7/29/2016

These release notes describe the fixes that are included in this patch.

Packages

Red Hat	Server	mapr-patch-4.1.0.31175.GA-39105.x86_64.rpm
Red Hat	Client	mapr-patch-client-4.1.0.31175.GA-39105.x86_64.rpm
Red Hat	Loopbacknfs	mapr-patch-loopbacknfs-4.1.0.31175.GA-39105.x86_64.rpm
Ubuntu	Server	mapr-patch-4.1.0.31175.GA-39105.x86_64.deb
Ubuntu	Client	mapr-patch-client-4.1.0.31175.GA-39105.x86_64.deb
Ubuntu	Loopbacknfs	mapr-patch-loopbacknfs-4.1.0.31175.GA-39105.x86_64.deb

Win64	Client	mapr-client-4.1.0.39105GA-1.amd64.zip
Mac	Client	mapr-client-4.1.0.39105GA-1.x86_64.tar.gz

Fixes

Bug 22124

Details

The CLDB was crashing because the loopback NFS server and POSIX client were participating in VIP rebalancing. When the loopback NFS server registered with CLDB, it was participating in VIP balancing, which could result in crash or incorrect assignment of VIP.

Resolution

With this fix, the loopback NFS server and POSIX client will no longer be included in VIP rebalancing.

Bug 22368

Details

Sometimes the `mrconfig info dumpcontainers` command was looping its output indefinitely.

Resolution

With this fix, the command will not loop its output indefinitely.

Bug 23629

Details

While allocating large number of inodes during resynchronization of containers, the source container would timeout if destination container did not respond within 5 minutes.

Resolution

With this fix, instead of sending large number of inodes during resynchronization, multiple commands with a fixed number of inodes per command will be sent to allocate the required number of inodes.

Bug 23715

Details

The MFS C and Java APIs did not return the requested number of bytes.

Resolution

With this fix, both C and Java APIs will return the requested number of bytes if present.

Bug 23745

Details

On a secure cluster, Pig jobs failed because zero-configuration Resource Manager HA did not handle the case where the filesystem set in the job configuration object is not the MapR-FS.

Resolution

With this fix, zero-configuration Resource Manager HA now handles the case where the filesystem set in the job configuration object is not the MapR-FS.

Bug 23799

Details

When there is an error, the container resync work area was freed, but the inode resync work area was still referring to the container resync work area.

Resolution

With this fix, the container work area will wait till the completion of all node resync operations before releasing the work area.

Bug 23876

Details

Sometimes, the same node was getting added to the replica chain twice and this was blocking the next resync request.

Resolution

With this fix, the same the node will not get added twice as a check has been included to verify if a node is already in the chain before adding a node.

Bug 23946

Details

When the Xmx parameter for `mapreduce.reduce.java.opts` was increased to a value greater than 2GB, the updated memory limit did not take effect.

Resolution

With this fix, MAPREDUCE-5649 was backported so that the `MergeManagerImpl` can be configured to use a memory limit greater than 2GB.

Release Notes for the June 2016 Patch

Released 6/24/2016

These release notes describe the fixes that are included in this patch.

Packages

Red Hat	Server	mapr-patch-4.1.0.31175.GA-38656.x86_64.rpm
Red Hat	Client	mapr-patch-client-4.1.0.31175.GA-38656.x86_64.rpm
Red Hat	Loopbacknfs	mapr-patch-loopbacknfs-4.1.0.31175.GA-38656.x86_64.rpm
Ubuntu	Server	mapr-patch-4.1.0.31175.GA-38656.x86_64.deb
Ubuntu	Client	mapr-patch-client-4.1.0.31175.GA-38656.x86_64.deb
Ubuntu	Loopbacknfs	mapr-patch-loopbacknfs-4.1.0.31175.GA-38656.x86_64.deb
Windows 64-bit	Client	mapr-client-4.1.0.38656GA-1.amd64.zip
Mac OS X	Client	mapr-client-4.1.0.38656GA-1.x86_64.tar.gz

Fixes

Bug 20186

Details

The MapR dbclient logged ESTALE (Stale file handle) errors for successful operations on binary tables when MapR-DB client applications followed these steps:

1. User X initiated the truncation of a binary table.
2. User X issued an operation (get/put/scan, etc) against the table while MapR-DB was in the process of deleting the table, recreating it, and then refreshing inode information.

The attempt to modify the table was unsuccessful and the dbclient logged an ESTALE error. However, the dbclient retried the operation after the truncation process was over and successfully modified the table.

Resolution

With this fix, the dbclient no longer logs the ESTALE error before retrying an action on a table after the truncation process is finished.

Bug 23228

Details

In Hadoop 2.5.1, a MRv2 job is not committed until all tasks are complete. Therefore, large MRv2 Mapreduce jobs can take long time to commit even when the job itself completes after a much shorter amount of time.

Resolution

With this fix, MAPREDUCE-4815 is backported into Hadoop 2.5.1. This patch includes a new fileoutputcommitter algorithm that improves performance. To enable the new algorithm, set `mapreduce.fileoutputcommitter.algorithm.version` to 2 in the `mapred-site.xml` or during job submission.

Bug 23382

Details

CLDB fails over with an exception when a node with stale containers is removed.

Resolution

With this fix, a node with stale containers can be removed successfully from the cluster and CLDB exceptions are not thrown.

Bug 23473

Details

In this type of situation in MapR-DB, the first of a series of puts for a row would succeed, while the remaining puts in the series would fail without errors:

1. A tablet T is split into T1 and T. The dbclient still has tablet T cached with the original key range.

2. The dbclient issues a series of puts against a rowkey that used to be in T, but which is now in T1.
3. The server returns an ERANGE error for the first put, but not for the remaining puts in the series.
4. The dbclient retries the first put and succeeds, but does not retry the remaining puts because the dbclient never received the ERANGE error for those puts.

This problem could occur for different types of errors that applied to all of the puts issued together for a single row.

Resolution

With this fix, the server returns the relevant error message for all of the puts in a series for a single row.

Bug 23541

Details

A ddlopen of libmapr_pam.so using immediate symbol resolution throws an undefined symbol error.

Resolution

An updated libmapr_pam.so has been provided that links to libpam.so.

Bug 23545

Details

MapReduce Application Master does not honor the binding port range specified through the `yarn.app.mapreduce.am.job.client.port-range` property.

Resolution

With the fix, MapReduce Application Master honors the binding port range specified by the `yarn.app.mapreduce.am.job.client.port-range` property.

Release Notes for the May 2016 Patch

Released 5/20/2016

These release notes describe the fixes that are included in this patch.

Packages

Red Hat	Server	mapr-patch-4.1.0.31175.GA-38213.x86_64.rpm
Red Hat	Client	mapr-patch-client-4.1.0.31175.GA-38213.x86_64.rpm
Red Hat	Loopbacknfs	mapr-patch-loopbacknfs-4.1.0.31175.GA-38213.x86_64.rpm
Ubuntu	Server	mapr-patch-4.1.0.31175.GA-38212.x86_64.deb
Ubuntu	Client	mapr-patch-client-4.1.0.31175.GA-38213.x86_64.deb
Ubuntu	Loopbacknfs	mapr-patch-loopbacknfs-4.1.0.31175.GA-38212.x86_64.deb
Windows 64-bit	Client	mapr-client-4.1.0.38213GA-1.amd64.zip
Mac OS X	Client	mapr-client-4.1.0.38213GA-1.x86_64.tar.gz

Fixes

Bug 21975

Details

Mirroring failed during source volume snapshot creation and returned a -1 error when a cross-cluster security ticket was used for mirroring.

Resolution

With this fix, mirroring no longer fails during source volume snapshot creation when a cross-cluster security ticket is used for mirroring.

Bug 22024

Details

Secure cross-cluster mirroring failed to complete because a container resynchronization procedure was rejected by the MapR-FS.

Resolution

With this fix, the MapR-FS allows all the resynchronization procedures that are required when cross-cluster mirroring occurs on secure clusters.

Release Notes for the April 2016 Patch

Released 4/22/2016

These release notes describe the fixes that are included in this patch.

Packages

Red Hat	Server	mapr-patch-4.1.0.31175.GA-37836.x86_64.rpm
Red Hat	Client	mapr-patch-client-4.1.0.31175.GA-37836.x86_64.rpm
Red Hat	Loopbacknfs	mapr-patch-loopbacknfs-4.1.0.31175.GA-37836.x86_64.rpm
Ubuntu	Server	mapr-patch-4.1.0.31175.GA-37836.x86_64.deb
Ubuntu	Client	mapr-patch-client-4.1.0.31175.GA-37836.x86_64.deb
Ubuntu	Loopbacknfs	mapr-patch-loopbacknfs-4.1.0.31175.GA-37836.x86_64.deb
Windows 64-bit	Client	mapr-client-4.1.0.37836GA-1.amd64.zip
Mac OS X	Client	mapr-client-4.1.0.37836GA-1.x86_64.tar.gz

Fixes

Bug 21372

Details

Snapshot delete could cause performance drops because the cache invalidation process walked over a global list of cache pages in the fileserver.

The container delete process needs to invalidate all pages belonging to the container. The process of invalidation was walking over the global list of cache pages causing high CPU usage in the fileserver. High CPU usage in the fileserver results in a performance drop.

Resolution

With this fix, a list of cache pages is maintained per container and used for invalidating container pages, avoiding the need to walk over a global list.

Bug 22534

Details

In situations where a client application looped between creating and deleting the same MapR-DB table, either of the following two circumstances could lead to a fileserver deadlock, preventing any other MapR filesystem operations in the volume hosting the table:

- The creation of a snapshot of the volume was triggered.
- A node hosting one of the containers of the table data failed.

Resolution

With this fix, fileserver deadlocks are no longer possible in these situations.

Bug 22808

Details

The calculation of the preemption utilization threshold of the Fair Scheduler's Dominant Resource Fairness (drf) scheduling policy did not consider disk usage as a resource. Instead, the preemption utilization threshold was calculated based on memory and CPU alone.

Resolution

With this fix, the drf scheduling policy considers memory, CPU, and disk usage when allocating resources to applications. For example, because MapReduce jobs require disk resources, preemption will now occur when the disk resources are at capacity.

Bug 22860

Details

Client applications holding two or more connections to the server could experience RPC timeouts in the following type of situation: After one connection establishes a session key with

the server, all of the connections remain idle long enough to trigger a session key renewal on the server. Two or more requests are then sent in parallel on different connections. The first request processed on the server triggers a change of the previous session key to the new session key. The remaining requests subsequently reaching the server on the other connections have the old session key, rather than the new session key.

Resolution

With this fix, the requests with the old session key are now discarded by the server and the client retransmits the requests with the new session key after a timeout that generally lasts from one to two minutes.

Bug 22881

Details

When mirroring was started for a volume, a new container, if not present, was created for each container in the source volume and the new containers were deleted if the mirroring was stopped. While deleting the new containers, the volume mirror module missed the last container in each iteration because the volume mirror module was incrementing the start key container ID (CID) during each iteration.

Resolution

With this fix, the volume mirror module will query the list of containers without missing a container and delete them.