Release Notes for Patches for the MapR 4.0.2 Release

Release Notes for the September 2016 Patch

Released 9/17/2016
These release notes describe the fixes that are included in this patch

Red Hat	Server	mapr-patch-4.0.2.29870.GA-39613.x86_64.rpm
Red Hat	Client	mapr-patch-client-4.0.2.29870.GA-39613.x86_64.rpm
Red Hat	Loopbacknfs	mapr-patch-loopbacknfs-4.0.2.29870.GA-39613.x86_64.rpm
Ubuntu	Server	mapr-patch-4.0.2.29870.GA-39613.x86_64.deb
Ubuntu	Client	mapr-patch-client-4.0.2.29870.GA-39613.x86_64.deb
Ubuntu	Loopbacknfs	mapr-patch-loopbacknfs-4.0.2.29870.GA-39613.x86_64.deb
Win64	Client	mapr-client-4.0.2.39613GA-1.amd64.zip
Mac	Client	mapr-client-4.0.2.39613GA-1.x86_64.tar.gz

Release Notes for the August 2016 Patch

Released 8/27/2016
These release notes describe the fixes that are included in this patch

Red Hat	Server	mapr-patch-4.0.2.29870.GA-39342.x86_64.rpm
Red Hat	Client	mapr-patch-client-4.0.2.29870.GA-39342.x86_64.rpm
Red Hat	Loopbacknfs	mapr-patch-loopbacknfs-4.0.2.29870.GA-39342.x86_64.rpm
Ubuntu	Server	mapr-patch-4.0.2.29870.GA-39342.x86_64.deb
Ubuntu	Client	mapr-patch-client-4.0.2.29870.GA-39342.x86_64.deb
Ubuntu	Loopbacknfs	mapr-patch-loopbacknfs-4.0.2.29870.GA-39342.x86_64.deb
Win64	Client	mapr-client-4.0.2.39342GA-1.amd64.zip
Mac	Client	mapr-client-4.0.2.39342GA-1.x86_64.tar.gz

© 2016 MapR Technologies, Inc.

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 23852

Details

The ResourceManager UI and Warden did not use the same logic to determine the list of running NodeManager nodes. This is because the ResourceManager UI did not consider NodeManagers to be running when they contain the "AsyncDispatcher thread interrupted" warning in the NodeManager log file.

Resolution

With this fix, YARN-2878 has been backported so that the NodeManager will not hang while it is processing the AsynchDispatcher thread.

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.

© 2016 MapR Technologies, Inc.

Bug 24025

Details

When the HistoryServer read the job history file for a job that was not initialized correctly, it read "-" as a delimiter. This caused the job start time to have an empty value. As a result, the following warning displayed:

```
<DATE> <TIME> WARN
org.apache.hadoop.mapreduce.v2.jobhistory.FileNameIndexUtils: Unable
to parse start time from job history file
```

Resolution

With this fix, the default start time is set to 0 instead of -1.

Bug 24050

Details

The output of the MapReduce LineRecordReader function occasionally had the following issues:
-With the multibyte record delimiter, records were dropped due to incorrect split processing.
-With the multibyte record delimiter, duplicate records were produced.
-With the custom delimiter, incomplete records were read.
-With the custom and default delimiter, incorrect key/position information was used for uncompressed input.

Resolution

With this fix, MAPREDUCE-6481, MAPREDUCE-6548, and MAPREDUCE-5948 were backported so that the LineRecordReader function no longer reads records incorrectly.

Bug 24059

Details

There was no way to retrieve the list of containers that were unaware of the rack.

Resolution

With this fix, the `-queue` option in `maprcli dump replicationmanagerqueueinfo` command takes 5 as value to return the list of containers which are not rack aware.

© 2016 MapR Technologies, Inc.

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.

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.

Release Notes for the July 2016 Patch

Released 7/29/2016
These release notes describe the fixes that are included in this patch

Red Hat	Server	mapr-patch-4.0.2.29870.GA-39118.x86_64.rpm
Red Hat	Client	mapr-patch-client-4.0.2.29870.GA-39118.x86_64.rpm

© 2016 MapR Technologies, Inc.

Red Hat	Loopbacknfs	mapr-patch-loopbacknfs-4.0.2.29870.GA-39118.x86_64.rpm
Ubuntu	Server	mapr-patch-4.0.2.29870.GA-39118.x86_64.deb
Ubuntu	Client	mapr-patch-client-4.0.2.29870.GA-39118.x86_64.deb
Ubuntu	Loopbacknfs	mapr-patch-loopbacknfs-4.0.2.29870.GA-39118.x86_64.debb
Win64	Client	mapr-client-4.0.2.39118GA-1.amd64.zip
Mac	Client	mapr-client-4.0.2.39118GA-1.x86_64.tar.gz

Fixes

Bug 17642

Details

Resync failed with ENOSPACE error as it tried to reserve more space than required for resync to complete.

Resolution

With this fix, MFS reserves the correct amount of space for resync operation to complete.

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.

© 2016 MapR Technologies, Inc.

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 23541

Details

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

Resolution

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

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.

Bug 23569

Details

The files under the /tmp/mapr-hadoop/mapred/local/toBeDeleted/ directory were not deleted when the TaskTracker service was restarted or when it reconnected with Jobtracker.

© 2016 MapR Technologies, Inc.

Resolution

With this fix, the TaskController deletes the contents of /tmp/mapr-hadoop/mapred/local/toBeDeleted/ whenever the TaskTracker service restarts or when it reconnects with Jobtracker.

Bug 23591

Details

The deletion of snapshots could be delayed by tens of minutes after commands to delete snapshots were issued. Also, deletions were recorded in the CLDB log before they were actually performed.

Resolution

To avoid such delays, increase the number of threads that is set by the cleanup.pool.threads.count parameter in the file /opt/mapr/conf/cldb.conf . After changing the value, stop and restart the CLDB node. Set the value to 4 or 8, depending on the size of the delays and the number of schedules for the creation of mirror volumes and snapshots.
CLDB log files no longer record the deletion of snapshots before deletions have in fact taken place.

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.

© 2016 MapR Technologies, Inc.

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 23922

© 2016 MapR Technologies, Inc.

Details

The ResourceManager failed to start when FileSystemRMStateStore contained invalid application states.

Resolution

With this fix, before the ResourceManager tries to recover applications, it verifies the integrity of FileSystemRMStateStore and removes the invalid entries.

Bug 23944

Details

In some cases, when a local write times out with ETIMEDOUT error, the NFS server re-uses shared pages, before mfs releases those pages, resulting in mfs crash.

Resolution

With this fix, on ETIMEDOUT error for local writes in NFS server, NFS server will not reuse those pages.

Bug 24034

Details

RACK violation messages are continuously printed to cldb.log. This issue occurs since the Replication Manager prints queue stats every two minutes and Containers having Rack Violation are fixed after 12 hours.

Resolution

This fix reduces logging by printing the status of Rack Violation Queue after 10 minutes. For other queues, logging is kept at the current 2 minutes.

Release Notes for the June 2016 Patch

Released 6/24/2016
These release notes describe the fixes that are included in this patch.

© 2016 MapR Technologies, Inc.

Packages

Red Hat	Server	mapr-patch-4.0.2.29870.GA-38573.x86_64.rpm
Red Hat	Client	mapr-patch-client-4.0.2.29870.GA-38573.x86_64.rpm
Red Hat	Loopbacknfs	mapr-patch-loopbacknfs-4.0.2.29870.GA-38573.x86_64.rpm
Ubuntu	Server	mapr-patch-4.0.2.29870.GA-38573.x86_64.deb
Ubuntu	Client	mapr-patch-client-4.0.2.29870.GA-38573.x86_64.deb
Ubuntu	Loopbacknfs	mapr-patch-loopbacknfs-4.0.2.29870.GA-38573.x86_64.deb
Windows 64-bit	Client	mapr-client-4.0.2.38573GA-1.amd64.zip
Mac OS X	Client	mapr-client-4.0.2.38573GA-1.x86_64.tar.gz

Fixes

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 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.

© 2016 MapR Technologies, Inc.

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 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-4.0.2.29870.GA-37933.x86_64.rpm
Red Hat	Client	mapr-patch-client-4.0.2.29870.GA-37933.x86_64.rpm
Red Hat	Loopbacknfs	mapr-patch-loopbacknfs-4.0.2.29870.GA-37933.x86_64.rpm
Ubuntu	Server	mapr-patch-4.0.2.29870.GA-37933.x86_64.deb
Ubuntu	Client	mapr-patch-client-4.0.2.29870.GA-37933.x86_64.deb
Ubuntu	Loopbacknfs	/mapr-patch-loopbacknfs-4.0.2.29870.GA-37933.x86_64.deb
Windows 64-bit	Client	mapr-client-4.0.2.37933GA-1.amd64.zip
Mac OS X	Client	mapr-client-4.0.2.37933GA-1.x86_64.tar.gz

Fixes

Bug 20161

Details

HBase region servers no longer freeze due to high memstore flushes and high compaction activity in this type of situation:

An HBase region server used a MapR client to open two secure sockets to an MFS node and sent requests using both secure sockets to the MFS node. When using secure sockets, however, an issue caused RPC to send replies on only one socket. If the socket that RPC used to reply is closed, the reply was never sent. This situation resulted in the HBase region server continuing to send requests using the other open socket.

© 2016 MapR Technologies, Inc.

Resolution

With this fix, RPC can reply on a secure socket if the request has been received or sent on the secure socket.

Bug 21259

Details

The volume mirror module created two different containers for the same source container because of RPC timeout. Sometimes, when volume mirror requested CLDB to create a container, it received an error although CLDB created a container. So, volume mirror resent a container create request to CLDB resulting in duplicate containers.

Resolution

With this fix, after creating containers, volume mirror checks for duplicate containers and fails mirroring if there are any duplicate containers.

Bug 22303

Details

If a replica container was ahead of the primary container during container resync operation, a cleanup of the replica is performed. If the source node failed during this cleanup, resync messages failed. Also, if resync was resumed before the cleanup was completed, resync resumed incorrectly.

Resolution

With this fix, when the source node is restarted after a failure, resync will be restarted if resync had been cleaning up a replica container at the time of failure.

Bug 22316

Details

After a node restart, the replica performs a fast resync to match versions with the primary. During a fast resync (that is, when versions are matching on primary and replica containers), there should not be any resync data. But if a new snapshot is created on the primary container and if the replica attempts to resync with primary, because the replica does not expect any resync data from the primary, the replica hits an assert.

© 2016 MapR Technologies, Inc.

Resolution

With this fix, if there is a need for a snapshot resync when a replica is performing a fast resync, the primary will return EAGAIN so that the replica retries with a slow resync, which allows for data transfer when there is data mismatch between the containers.

Bug 22318

Details

When volume mirroring was in progress, intra-volume resynchronization of data was not waiting until the mirroring process was complete.

When mirroring volumes, temporary snapshots are created to resynchronize the data. If replicas and primary replicas have snapshots, then resynchronization is allowed. If the primary replica has a snapshot but the replica does not, intra-volume resynchronization should wait until the mirroring process is complete.

Resolution

With this fix, if the primary replica has a snapshot but the replica does not, intra-volume resynchronization waits until the mirroring process is complete, meaning until each replica has a temporary snapshot.

Bug 22389

Details

When volume mirroring was in progress and some nodes in the cluster (where the mirror volume resides) failed and were restarted, some containers of the mirror volume would have old data and some would have new data.

Because of multiple node failures, a temporary snapshot of the container is deleted even though mirroring completes. The snapshot is used to resynchronize the data. As a result, containers, associated with deleted snapshots, have old data. In addition, the replica sends snapshot versions, which, if different, triggers the primary replica to try to resynchronize the data. Because snapshots are deleted, due to node failure and restart, the snapshot versions sent by the replica are different.

Resolution

With this fix, the replica sends the correct snapshot version, the primary replica verifies that both the replica and primary replica are at the same state, and MFS establishes the replica chain.

© 2016 MapR Technologies, Inc.

Bug 22491

Details

A log message typographical error prevents support from accurately confirming that container resync errors are causing mirroring failures.

Resolution

With this fix, the log message is corrected.

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.

© 2016 MapR Technologies, Inc.

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.

© 2016 MapR Technologies, Inc.