

Release Notes for Patches for the MapR 5.1.0 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-5.1.0.37549.GA-40890.x86_64.rpm
Red Hat	Client	mapr-patch-client-5.1.0.37549.GA-40890.x86_64.rpm
Red Hat	Loopbacknfs	mapr-patch-loopbacknfs-5.1.0.37549.GA-40890.x86_64.rpm
Red Hat	Posix-client-basic	mapr-patch-posix-client-basic-5.1.0.37549.GA-40890.x86_64.rpm
Red Hat	Posix-client-platinum	mapr-patch-posix-client-platinum-5.1.0.37549.GA-40890.x86_64.rpm
Ubuntu	Server	mapr-patch-5.1.0.37549.GA-40890.x86_64.deb
Ubuntu	Client	mapr-patch-client-5.1.0.37549.GA-40890.x86_64.deb
Ubuntu	Loopbacknfs	mapr-patch-loopbacknfs-5.1.0.37549.GA-40890.x86_64.deb
Ubuntu	Posix-client-basic	mapr-patch-posix-client-basic-5.1.0.37549.GA-40890.x86_64.deb
Ubuntu	Posix-client-platinum	mapr-patch-posix-client-platinum-5.1.0.37549.GA-40890.x86_64.deb
Win32	Client	mapr-client-5.1.0.40890GA-1.win32.zip
Win64	Client	mapr-client-5.1.0.40890GA-1.amd64.zip
Mac	Client	mapr-client-5.1.0.40890GA-1.x86_64.tar.gz

Fixes

Bug 13187

Description

The `mapcli volume create` command was not setting group ownership to user's primary group.

Resolution

With this fix, the `mapcli volume create` command will set group ownership to user's primary group.

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 23257

Description

In MCS, new NFS VIPs were visible in the **NFS HA > VIP Assignments** tab, but not in the **NFS HA > NFS Setup** tab.

Resolution

With this fix, the NFS VIPs will be available in both the **NFS HA > VIP Assignments** tab and the **NFS HA > NFS Setup** tab.

Bug 23975

Description

In version 5.1, MFS was failing to start on some docker containers as it was trying to figure out number of numa nodes from `/sys/devices/system/node`.

Resolution

With this fix, MFS will work on docker containers.

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 24155

Description

Disk setup was timing out if running trim on flash drives took some time.

Resolution

With this fix, disk setup will complete successfully and the warning message (“Starting Trim of SSD drives, it may take a long time to complete”) is entered in the log file.

Bug 24249

Description

When running map/reduce jobs with older versions of the MapR classes, the system hung because the older classes linked to the native library installed on cluster nodes that have been updated to a newer MapR version.

Resolution

With this fix, the new `fs.mapr.bailout.on.library.mismatch` parameter detects mismatched libraries, fails the map/reduce job, and logs an error message. The parameter is enabled by default. You can disable the parameter on all the TaskTracker nodes and resubmit the job for the task to continue to run. To disable the parameter, you must set it to `false` in the `core-site.xml` file.

Bug 24585

Description

Excessive logging in CLDB audit caused `cldbaudit.log` file to grow to large sizes.

Resolution

With this fix, to reduce the size of `cldbaudit.log` file, the queries to CLDB for ZK string will no longer be logged for auditing.

Bug 24610

Description

In a secure cluster, when there are intermittent connection drops (between MFS-MFS or client-MFS), the client and/or server could crash during authentication.

Resolution

With this fix, the client and/or server will not crash during authentication if there are intermittent connection drops.

Bug 24812

Description

Apache Hadoop could not look up the status of a finished job because `job.xml` was already removed from the search directory. Hive interpreted the job as failing and generated an exception.

Resolution

With this fix, Apache Hadoop correctly reports the status of the finished job.

Bug 24965

Description

On large clusters, sometimes the bind failed with the message indicating unavailability of port when running MR jobs, specifically reducer tasks.

Resolution

With this fix, the new `fs.mapr.bind.retries` configuration parameter in `core-site.xml` file, if set to `true`, will retry to bind during client initialization for 5 minutes before failing. By default, the `fs.mapr.bind.retries` configuration parameter is set to `false`.

Bug 24915

Description

In version 5.1, running the `expandaudit` utility on volumes can result in very large (more than 1GB) audit log files due to incorrect `GETATTR` (get attributes) cache handling.

Resolution

With this fix, the `expandaudit` utility has been updated so that it will not perform subsequent `GETATTR` calls if the original call to the same file identifier failed.

Bug 25003

Description

When a specific queue uses all of its resources, the `UsedResources` tab in the Resource Manager UI might show a greater value than shown in the `MaxResources` tab. This happens when another application is submitted and the application master container size is added in.

Resolution

With this fix, no more containers can be assigned to a queue when its `UsedResource` has reached the `MaxResource` limit.

Bug 25177

Description

When using `FairScheduler` with `maxAMShare` enabled, total `amResourceUsage` per queue is not calculated properly, which may cause applications to hang in `ACCEPTED` state.

Resolution:

AM resource usage is now calculated as expected and YARN jobs no longer get stuck in the `ACCEPTED` state.

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-5.1.0.37549.GA-40163.x86_64.rpm
Red Hat	Client	mapr-patch-client-5.1.0.37549.GA-40163.x86_64.rpm
Red Hat	Loopbacknfs	mapr-patch-loopbacknfs-5.1.0.37549.GA-40163.x86_64.rpm
Red Hat	Posix-client-basic	mapr-patch-posix-client-basic-5.1.0.37549.GA-40163.x86_64.rpm
Red Hat	Posix-client-platnium	mapr-patch-posix-client-platinum-5.1.0.37549.GA-40163.x86_64.rpm
Ubuntu	Server	mapr-patch-5.1.0.37549.GA-40163.x86_64.deb
Ubuntu	Client	mapr-patch-client-5.1.0.37549.GA-40163.x86_64.deb
Ubuntu	Loopbacknfs	mapr-patch-loopbacknfs-5.1.0.37549.GA-40163.x86_64.deb
Ubuntu	Posix-client-basic	mapr-patch-posix-client-basic-5.1.0.37549.GA-40163.x86_64.deb
Ubuntu	Posix-client-platnium	mapr-patch-posix-client-platinum-5.1.0.37549.GA-40163.x86_64.deb
Win32	Client	mapr-client-5.1.0.40163GA-1.win32.zip

Win64	Client	mapr-client-5.1.0.40163GA-1.amd64.zip
Mac	Client	mapr-client-5.1.0.40163GA-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 24408

Description

When running multiple producers as separate threads within a process, with a very small value for buffer.memory (say 1KB), some producers can stall. This is due to a lack of buffer memory.

Resolution

With this fix, the default value for minimum buffer memory is increased to 10kB.

Bug 24477

Description

Jobs failed if a local volume was not available and directories for mapreduce could not be initialized.

Resolution

With this fix, jobs no longer fail, and local volume recovery is enhanced.

Bug 24563

Description

The Resource Manager failed to load or recover after an attempt to read the first byte of a data stream that had no first byte or no data.

Resolution

With this fix, the Resource Manager recovers after attempting to read a data stream with no first byte or no data.

Bug 24566

Description

An older version of the aws-sdk jar was built with MapR.

Resolution

With this fix, MapR upgraded the aws-sdk jar from version 1.7.4 to 1.7.15.

Bug 24630

Description

The time stamp on the files showed negative numbers in nano seconds after applying mapr-patch 5.2 build 39544.

Resolution

With this fix, if the nanosecond timestamp overflows because of the old NFS bug, the nanoseconds value will be 9999.

Bug 24658

Description

CLDB returned “no master” and an empty list for container lookup, which NFS server could not handle, because when multiple servers are down, there can be no master for a container.

Resolution

With this fix, NFS server will handle empty node list for container lookup.

Bug 24505

Description

A job failed when the JvmManager went into an inconsistent state.

Resolution

With this fix, jobs no longer fail as a result of the JvmManager entering an inconsistent state.

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 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.

Bug 24700

Description

The Job Tracker user interface failed with a NullPointerException when a user submitted a Hive job with a null value in a method.

Resolution

With this fix, the Job Tracker interface does not fail when a Hive job is run with a null value in a method.

Bug 24992

Description

Installing a MapR patch caused jar files to be removed from under the drill/drill-1.4.0/jars/ directory

Resolution

Jar files are no longer incorrectly removed

Bug 24618

Description

Remote mirror volumes could not be created on secure clusters using MCS even when the appropriate tickets were present.

Resolution

With this fix, remote mirror volumes can now be created on secure clusters using MCS.

Bug 24971

Description

When the mirroring operation started after a CLDB failover, sometimes it was sending request 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 September 2016 Patch

Released 9/23/2016

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

Packages

Red Hat	Server	mapr-patch-5.1.0.37549.GA-39728.x86_64.rpm
Red Hat	Client	mapr-patch-client-5.1.0.37549.GA-39728.x86_64.rpm
Red Hat	Loopbacknfs	mapr-patch-loopbacknfs-5.1.0.37549.GA-39728.x86_64.rpm
Red Hat	Posix-client-basic	mapr-patch-posix-client-basic-5.1.0.37549.GA-39728.x86_64.rpm
Red Hat	Posix-client-platinum	mapr-patch-posix-client-platinum-5.1.0.37549.GA-39728.x86_64.rpm
Ubuntu	Server	mapr-patch-5.1.0.37549.GA-39728.x86_64.deb

Ubuntu	Client	mapr-patch-client-5.1.0.37549.GA-39728.x86_64.deb
Ubuntu	Loopbacknfs	mapr-patch-loopbacknfs-5.1.0.37549.GA-39728.x86_64.deb
Ubuntu	Posix-client-basic	mapr-patch-posix-client-basic-5.1.0.37549.GA-39728.x86_64.deb
Ubuntu	Posix-client-platinum	mapr-patch-posix-client-platinum-5.1.0.37549.GA-39728.x86_64.deb
Win32	Client	mapr-client-5.1.0.39728GA-1.win32.zip
Win64	Client	mapr-client-5.1.0.39728GA-1.amd64.zip
Mac	Client	mapr-client-5.1.0.39728GA-1.x86_64.tar.gz

Fixes

Bug 23652

Description

The POSIX loopbacknfs client did not automatically refresh renewed service tickets.

Resolution

With this fix, the POSIX loopbacknfs client will:

- Automatically use the renewed service ticket without requiring a restart if the ticket is replaced before expiration (ticket expiry time + grace period of 55 minutes). If the ticket is replaced after expiration (which is ticket expiry time + grace period of 55 minutes), the POSIX loopbacknfs client will not refresh the ticket as the mount will become stale.
- Allow impersonation if a service ticket is replaced before ticket expiration (which is ticket expiry time + grace period of 55 minutes) with a servicewithimpersonation ticket.
- Honor all changes to user/group IDs of the renewed ticket.

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.

Bug 24057

Description

Only SPs in average and above buckets were considered for balancing because disk balancer had to spread the containers for optimal SP utilization.

Resolution

With this fix, disk balancer balances SPs in all bins giving preference to those SPs that are highly utilized. SPs in lower bins are balanced periodically (configurable), when there is not too much balancing activity (of SPs that are highly utilized). By default, SPs in lower bins are balanced every 60 minutes. To set it to a lower value, reset the value for `dbal.below.avg.bins.balancing.frequency` using the `maprcli config save` command.

The default size of above average, average, and below average bins is 20 (percentage). To allow more granular and aggressive balancing of storage pools across bins, reduce the size of each bin. To reduce the size of each bin, specify the value for the following parameters using the `maprcli config save` command:

<code>dbal.above.avg.bin.size</code>	specifies the bin size (%) of SPs whose usage is above cluster average
<code>dbal.avg.bin.size</code>	specifies the bin size (%) of SPs whose usage is in the average range
<code>dbal.below.avg.bin.size</code>	specifies the bin size (%) of SPs whose usage is below average range

For example, to reduce the size of the below average bin to 10%, run the following command:

```
maprcli config save -values {"dbal.below.avg.bin.size":"10"}
```

Periodically, disk balancer logs information about all the bins and the storage pools in those bins in the file `/opt/mapr/logs/clbdbdiskbalancer.log`. By default, the information is logged every 10 minutes and can be configured using the parameter `dbal.loadtracker.info.log.frequency`.

Bug 24119

Details

Warden adjusts the FileServer (MFS) and Node Manager (NM) memory incorrectly when NM and TaskTracker (TT) are on the same node. This can result in too much memory being allocated to MFS.

Resolution

With this fix, Warden does not adjust MFS memory when NM and TT are on the same node. Memory adjustment is implemented only when TT and MapR-FS (but no NM) are on the same node.

Bug 24159

Description

The mtime was updated whenever a hard link was created. Also, when a hard link was created from the FUSE mount point, although the ctime was updated, the update timestamp only showed the minutes and seconds and not the nanoseconds.

Resolution

With this fix, mtime will not change on the hard link and when a hard link is created from the FUSE mount point, the timestamp for ctime will include nanoseconds.

Bug 24232

Description

In certain cases, files were created with stale chunk IDs, which prevented users from accessing files in the parent directory.

Resolution

With this fix, files will not be created with stale chunk IDs.

Bug 24324

Description

The “disk not found” error was thrown because the script used to list disks was looking up disks in every instance of the fileserver process.

Resolution

With this fix, the script will look for disks only in the specific instance of the fileserver process.

Bug 24280

Description

Running the `maprcli dashboard info` command occasionally throws a `TimeoutException` error.

Resolution

With this fix, the internal timeout command was increased to provide more allowance for command processing.

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-5.1.0.37549.GA-39353.x86_64.rpm
Red Hat	Client	mapr-patch-client-5.1.0.37549.GA-39353.x86_64.rpm
Red Hat	Loopbacknfs	mapr-patch-loopbacknfs-5.1.0.37549.GA-39353.x86_64.rpm
Ubuntu	Server	mapr-patch-5.1.0.37549.GA-39353.x86_64.deb
Ubuntu	Client	/mapr-patch-client-5.1.0.37549.GA-39353.x86_64.deb
Ubuntu	Loopbacknfs	mapr-patch-loopbacknfs-5.1.0.37549.GA-39353.x86_64.deb
Win32	Client	mapr-client-5.1.0.39353GA-1.win32.zip
Win64	Client	mapr-client-5.1.0.39353GA-1.amd64.zip
Mac OS X	Client	mapr-client-5.1.0.39353GA-1.x86_64.tar.gz

Fixes

Bug 20498 and 24143

Details

JobTracker attempts to restart TaskTrackers resulted in the loading of the job configuration object multiple times. Loading the job configuration object multiple times caused a JobTracker lock contention and it also caused the JobTracker service to become unresponsive.

Resolution

With this fix, the JobTracker caches the job configuration object for each job. Then, it uses the cached configuration object associated with each job for all of the job's task completion events.

Bug 23776

Details

The client accessing using the Java API experienced a crash because the `closed_flag` in `Inode.close()` method was not set to `true` when there was an error (like `EACCES`).

Resolution

With this fix, the `closed_flag` will be set to `true` irrespective of the error.

Bug 23931

Details

Nested user queues did not inherit labels from their parent queue. As a result, no labels were configured for nested user queues.

Resolution

With this fix, each nested user queue inherits the label and label policy from its parent queue.

Bug 23933

Details

When an OJAI Document is created by parsing a JSON String that contains a List of Maps, for example [{} , {} , ...], an extraneous null element is added to this list,for example [{} , null, {} , ...] in the resultant document.

Resolution

With this fix, the extraneous null element does not occur.

Bug 24018

Details

YARN jobs failed with "Rename cannot overwrite non empty destination directory"

Resolution

With this fix, failures no longer occur.

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 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 24054

Details

When MapR-DB is used for time series data, a large number of tablets are generated with empty partitions. This can be merged into single tablet but, in this case, the single tablet will have thousands of empty partitions. Having a single tablet with thousands of empty partitions may result in a degradation of MapR-DB performance.

Resolution

With this fix, bucket flush is optimized for the tablet when a tablet has thousands of partitions.

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 24097

Details

Retrieval of zookeeper address for every hadoop command required a scan of the complete cluster information, which made cldb threads busy and unresponsive for new requests causing hadoop commands to take more time.

Resolution

With this fix, the complete cluster will not be scanned to retrieve zookeeper information.

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 24170

Details

If the new ($\geq 5.1.0$) FileClient talks to older clusters, applications like YARN, which requires the ZooKeeper string, might fail because of the new way of retrieving the ZooKeeper address (which is available on clusters $\geq 5.1.0$).

Resolution

With this fix, the FileClient will fallback to the older way of retrieving the ZooKeeper information from CLDB if the new way of retrieving the information does not work.

Bug 24190

Details

During table replication on 10,000 tables, the destination table takes a long time (30 minutes or so) to catch up with the primary table. This is due to a bug in idleWorker which resulted in open buckets being flushed only during bucketflush instead of by the idleWorker.

Resolution

With this fix, idleWorker flushes idle open buckets for replication every 2 seconds.

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-5.1.0.37549.GA-39127.x86_64.rpm
Red Hat	Client	mapr-patch-client-5.1.0.37549.GA-39127.x86_64.rpm
Red Hat	Posix-client-basic	mapr-patch-posix-client-basic-5.1.0.37549.GA-39127.x86_64.rpm

Red Hat	Posix-client-platinum	mapr-patch-posix-client-platinum-5.1.0.37549.GA-39127.x86_64.rpm
Red Hat	Loopbacknfs	mapr-patch-loopbacknfs-5.1.0.37549.GA-39127.x86_64.rpm
Ubuntu	Server	mapr-patch-5.1.0.37549.GA-39127.x86_64.deb
Ubuntu	Client	mapr-patch-client-5.1.0.37549.GA-39127.x86_64.deb
Ubuntu	Posix-client-basic	mapr-patch-posix-client-basic-5.1.0.37549.GA-39127.x86_64.deb
Ubuntu	Posix-client-platinum	mapr-patch-posix-client-platinum-5.1.0.37549.GA-39127.x86_64.deb
Ubuntu	Loopbacknfs	mapr-patch-loopbacknfs-5.1.0.37549.GA-39127.x86_64.deb
Win32	Client	mapr-client-5.1.0.39127GA-1.win32.zip
Win64	Client	mapr-client-5.1.0.39127GA-1.amd64.zip
Mac OS X	Client	mapr-client-5.1.0.39127GA-1.x86_64.tar.gz

Fixes

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 23331

Details

When re-reading large files, some cache misses were seen despite warm cache.

Resolution

With this fix, re-reading large files will not result in cache misses.

Bug 23353

Details

Spark 1.5.2 job server core dumps when running with MapR 5.1 client.

Resolution

Fixed an internal double free error in MapR client when handling file attribute objects.

Bug 23488

Details

When a volume with a replication factor below 3 and no minimum number of replication was created, the replication manager set the minimum replication value to 1 and containers with only one copy were not re-replicated.

Resolution

With this fix, by default, the minimum number of copies is 1 if the replication factor is less than or equal to 2.

Bug 23544

Details

Incorrect values are returned for Name Container Replication Factor. Namespace and Dataspace replication can be independently specified. However, if Namespace replica factors are not set explicitly, they assume the same values as that of data space for both 'desired' and 'min'.

If the 'min' and 'desired' repl factors of Dataspace are 2, and 3 respectively, when the replication manager checks for the desired replication of the namespace container, it would see '2' instead of '3'. Hence, it might fail to create an additional copy of the namespace containers (since it thinks the desired is 2, when in fact it should have been 3).

Volume info might show the desired and min correctly, but the available copies might fall short of the desired replica factor.

Resolution

This issue is fixed by invoking the correct function to return the namespace replica factor when is is not different from the data replica factor.

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.

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 23596

Details

In the scenario where Fastfailover is enabled and star replica is implemented, (for example, with a star replica chain of A-B A-C), if B goes down, the error causes the replication operation to not send the replica operation to C. A does not send the operation to C because it assumes that it received a reply from C.

Resolution

With this fix, the operation is actually sent from A to C.

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 23676

Details

When File ACEs (Access Control Expressions) were set on a symbolic link, flags were set to indicate aces were set on the symlink file itself.

Resolution

With this fix, flags will not be set on symbolic link file and ACEs on the target file will be honored. However, for file ACEs on existing symbolic links, this fix is not available; delete and recreate the symbolic links for setting file ACEs.

Bug 23687

Details

When a user tried to use `hdfsOpenFile()`, `hdfsWrite()`, `hdfsFlush()`, and/or `hdfsCloseFile()` on a file on which the user did not have the right permissions for the operation, 0 was returned instead of the right error code.

Resolution

With this fix, when `hdfsOpenFile()`, `hdfsWrite()`, `hdfsFlush()`, and/or `hdfsCloseFile()` is used by a user without the right permissions for the operation on the file, the operation will fail and the appropriate error code will be returned.

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 23725

Details

On clusters with more than 50K volumes, when CLDB fails over, the CPU spin was observed on the new CLDB master because of ACEs (even when there were no ACEs on the volumes).

Resolution

With this fix, the CLDB threads will no longer spin CPU.

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 23762

Details

When compiling code with `hdfsExists2()`, an error was returned as `hdfs.h` did not expose `hdfsExists2()`.

Resolution

The `hdfs.h` file has been updated to include `hdfsExists2()`.

Bug 23791

Details

The Resource Manager UI does not show details about task attempts. Instead, it only shows that a task is "Processing..." until the job completes.

Resolution

With this fix, the task attempts page works as expected.

Bug 23795

Details

Some storage pools were going offline frequently with CRC errors because:

- NFS server was encrypting files even when network-encryption bit was not set on files. This happens when `mfs.feature.filecipherbit.support` feature is not enabled.
Note: This happens only on clusters that were upgraded from 3.0, or prior to 3.1 or later and configured as secure.
- During replication, if replica reported a CRC error for replicated writes, the source was retransmitting without verifying CRC at source. This caused the replica to go stale unnecessarily.

Resolution

With this fix:

- NFS server will no longer encrypt data unless "network-encryption" bit is set on the files.
- The source for the replica will check CRC and if it doesn't match, the source will crash. This will allow one of the replicas to become master.

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 24008

Details

On a NFS node, CLDB was trying to get the port from the IP list of the server that did not send any IPs and this caused CLDB to crash.

Resolution

With this fix, CLDB will no longer crash as it won't try to get the port.

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-5.1.0.37549.GA-38657.x86_64.rpm
Red Hat	Client	mapr-patch-client-5.1.0.37549.GA-38657.x86_64.rpm
Red Hat	Posix-client-basic	mapr-patch-posix-client-basic-5.1.0.37549.GA-38657.x86_64.rpm
Red Hat	Posix-client-platinum	mapr-patch-posix-client-platinum-5.1.0.37549.GA-38657.x86_64.rpm
Red Hat	Loopbacknfs	mapr-patch-loopbacknfs-5.1.0.37549.GA-38657.x86_64.rpm
Ubuntu	Server	mapr-patch-5.1.0.37549.GA-38657.x86_64.deb
Ubuntu	Client	mapr-patch-client-5.1.0.37549.GA-38657.x86_64.deb
Ubuntu	Posix-client-basic	mapr-patch-posix-client-basic-5.1.0.37549.GA-38657.x86_64.deb
Ubuntu	Posix-client-platinum	mapr-patch-posix-client-platinum-5.1.0.37549.GA-38657.x86_64.deb
Ubuntu	Loopbacknfs	mapr-patch-loopbacknfs-5.1.0.37549.GA-38657.x86_64.deb
Windows 32-bit	Client	mapr-client-5.1.0.38657GA-1.win32.zip
Windows 64-bit	Client	mapr-client-5.1.0.38657GA-1.amd64.zip
Mac OS X	Client	mapr-client-5.1.0.38657GA-1.x86_64.tar.gz

Fixes

Bug 23106

Details

On an unsecure cluster, C client applications were able to impersonate even when the:

- MAPR_IMPERSONATION_ENABLED environment variable was not set to `true`

- Impersonating user did not have a file under `/opt/mapr/conf/proxy`

Resolution

With this fix, to enable impersonation on an unsecure cluster for C client applications, the `MAPR_IMPERSONATION_ENABLED` environment variable must be set to `true` and the impersonating user must have a file under `/opt/mapr/conf/proxy`.

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.

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-5.1.0.37549.GA-38290.x86_64.rpm
Red Hat	Client	mapr-patch-client-5.1.0.37549.GA-38290.x86_64.rpm
Red Hat	Posix-client-basic	mapr-patch-posix-client-basic-5.1.0.37549.GA-38290.x86_64.rpm
Red Hat	Posix-client-platinum	mapr-patch-posix-client-platinum-5.1.0.37549.GA-38290.x86_64.rpm
Red Hat	Loopbacknfs	mapr-patch-loopbacknfs-5.1.0.37549.GA-38290.x86_64.rpm
Ubuntu	Server	mapr-patch-5.1.0.37549.GA-38290.x86_64.deb
Ubuntu	Client	mapr-patch-client-5.1.0.37549.GA-38290.x86_64.deb
Ubuntu	Posix-client-basic	mapr-patch-posix-client-basic-5.1.0.37549.GA-38290.x86_64.deb
Ubuntu	Posix-client-platinum	mapr-patch-posix-client-platinum-5.1.0.37549.GA-38290.x86_64.deb
Ubuntu	Loopbacknfs	mapr-patch-loopbacknfs-5.1.0.37549.GA-38290.x86_64.deb
Windows 32-bit	Client	mapr-client-5.1.0.38258GA-1.win32.zip
Mac OS X	Client	mapr-client-5.1.0.38258GA-1.x86_64.tar.gz

Fixes

Bug 22491

Details

A typographical error in log messages prevented MapR support from accurately confirming that container resync errors were causing mirroring failures.

Resolution

With this fix, the typographical errors are corrected.

Bug 22829

Details

On installations with two MFS instances per node on (SSD-based) clusters, instead of assigning one license per node, MFS was using one license per instance. For example, a 10 node cluster with SSDs required 20 licenses instead of 10 licenses.

Resolution

With this fix, MFS will no longer assign multiple licenses to a node if more than one MFS instance is running on that node.

Bug 22873

Details

Occasionally, writes to MapR-FS by C client applications were not visible immediately after a flush. Therefore, these writes appeared to be missing.

Resolution

With this fix, the data inconsistency issue is resolved.

Bug 22949

Details

Occasionally, C client applications connected to MapR-FS experienced short reads during concurrent reads or concurrent reads and writes.

Resolution

With this fix, the short read issue is resolved.

Bug 22978

Details

If posix-client packages (mapr-posix-client-basic/platinum) were installed on both a MapR core cluster and a client node, the mapr-patch-posix-client patch did not work as expected.

Resolution

With this fix, the posix client patch can be applied to both the MapR cluster and the client node.

Bug 23178

Details

Consumer applications reading messages from topics in MapR Streams generate globally unique identifiers (GUIDs) that the server can use to identify individual consumers. Consumers running on OS X could occasionally generate GUIDs in formats that the server would not recognize.

Resolution

With this fix, consumers running on OS X always generate GUIDs in the correct format.

Bug 23186

Details

When a MapR-FS client application written in C called `hdfsDisconnect()`, the corresponding file system handle was not deleted, resulting in memory leaks.

Resolution

With this fix, file handles are now deleted when `hdfsDisconnect()` is called.

Bug 23192

Description

User credentials were not set explicitly whenever a client application accessed a MapR-DB table, which caused an EACCES error in the following type of situation:

For connecting to a single MapR cluster, a client application written in C creates two connection objects (connA and connB), using separate user credentials (for userA and userB) for each connection object.

A single application thread is used for table operations with both connA and connB. This thread performs these operations:

1. As userA, the thread creates table X via connA. The MapR dbclient, which mediates connections between client applications and MapR clusters running MapR-DB, caches the credentials for userA in thread local storage.
2. As userB, the thread deletes table Y via connB. The dbclient overwrites userA's credentials in thread local storage with the credentials for userB.
3. As userA, the thread attempts a put operation on table X. Before attempting to access the table with userA's credentials, the dbclient does not first overwrite the credentials in thread local storage. Because the stored credentials are for userB, the attempt to access the table fails.

Resolution

With this fix, the dbclient overwrites the credentials in thread local storage with the current user's credentials before attempting to access tables.

Bug 23215

Description

When a client application written in C successfully deleted a table in MapR-DB, error messages such as the following would be logged:

```
2016-04-27 14:30:58,0688 ERROR Client fs/client/fileclient/cc/client.cc:2372
Thread: 1120 Unlink failed for file /user/temp/test.table, error Invalid
argument(22)
2016-04-27 14:30:58,0688 ERROR Inode fs/client/fileclient/cc/inode.cc:485
Thread: 1120 Unlink failed on file /user/temp/test.table with error 22
```

Resolution

With this fix, such error messages are no longer logged.

Bug 23303

Details

In C client applications for MapR-FS, calling `hdfsConnectAsUserNewInstance()` with an invalid user and then calling the same function call with a valid user caused mfs to core.

Resolution

With this fix, when all further calls with a valid user is made, the memory corruption (because of failed user resolution during the first call) will not result in core dump.

Bug 23304

Details

Running the `ls` command from a valid (working) cluster returned the following error when the `clusters.conf` file included an invalid (non-working) cluster entry:

```
ls: cannot open directory .: No such file or directory
```

Resolution

With this fix, the `ls` command no longer returns the aforementioned error if the `clusters.conf` file includes an invalid (non-working) cluster entry.

Bug 23311

Details

If a table was a source in table replication, an incremental bulk load of one or more non-replicated column families would cause the mfs service to core.

For example, suppose a source table contained the column families `cf1`, `cf2`, and `cf3`. Only `cf3` was being replicated. If an incremental bulk load was started for `cf1` and `cf2`, the mfs service cored.

Resolution

With this fix, the mfs service no longer cores in this type of situation.

Bug 23312

Details

The duration specified while generating cross-cluster tickets was not being set and the default duration of 14 days was being applied instead.

Resolution

With this fix, for:

- Admin generated cross-cluster and service tickets:
 - The duration specified while generating a ticket will now be honored.
 - The default duration is now LIFETIME.

Note: Service and cross-cluster tickets are now not bounded by the CLDB duration properties.
- Password authenticated tickets, if duration is:
 - Specified, the specified duration will be honored.
 - Not specified, the default duration configured as CLDB properties will be used.

The maplogin print command will now print ticket of any type.

Bug 23315

Details

Many GetXAttr (get extended attribute) calls (for file ACEs) were made on the NFS mount irrespective of the file type. This resulted in a lot of GetXAttr calls on the NFS mount for normal file operations.

Resolution

With this fix, GetXAttr calls will only be made for a special file (`.dfs_attributes`).

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-5.1.0.37549.GA-38007.x86_64.rpm
Red Hat	Client	mapr-patch-client-5.1.0.37549.GA-38007.x86_64.rpm
Red Hat	Posix-client-basic	mapr-patch-posix-client-basic-5.1.0.37549.GA-38007.x86_64.rpm
Red Hat	Posix-client-platinum	mapr-patch-posix-client-platinum-5.1.0.37549.GA-38007.x86_64.rpm
Red Hat	Loopbacknfs	mapr-patch-loopbacknfs-5.1.0.37549.GA-38007.x86_64.rpm
Ubuntu	Server	mapr-patch-5.1.0.37549.GA-38007.x86_64.deb
Ubuntu	Client	mapr-patch-client-5.1.0.37549.GA-38007.x86_64.deb
Ubuntu	Posix-client-basic	mapr-patch-posix-client-basic-5.1.0.37549.GA-38007.x86_64.deb
Ubuntu	Posix-client-platinum	mapr-patch-posix-client-platinum-5.1.0.37549.GA-38007.x86_64.deb
Ubuntu	Loopbacknfs	mapr-patch-loopbacknfs-5.1.0.37549.GA-38007.x86_64.deb
Windows 64-bit	Client	mapr-client-5.1.0.37962GA-1.amd64.zip
Mac OS X	Client	mapr-client-5.1.0.37962GA-1.x86_64.tar.gz

Fixes

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 22574

Details

The CLDB DiskBalancer thread monopolized the LoadTracker lock, resulting in delayed heartbeat processing and slow heartbeat alarms.

In every heartbeat from fileserver, CLDB receives the list of storage pools and the associated usage information. The heartbeat processing thread then puts the storage pools into different bins/buckets based on their utilization. The bucketing is used by disk balancer and the replication manager to pick a good storage pool to create a replica on. However, all of them take the LoadTracker object lock. This can potentially delay heartbeat processing by as long as 10 seconds.

Resolution

With this fix, the bucketing is now moved out of heartbeat processing code as a delayed operation by queuing a task to a separate thread pool. To avoid queuing a task for every heartbeat, a delayed task is checked to see if it has already been scheduled for a fileserver. If so, the task is not queued and bucketing is allowed to complete. When the delayed task is done, another delayed task, to bucket the storage pool, will be scheduled during a subsequent heartbeat.

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.

Bug 22883

Details

The MFS did not handle different formats in the CPU list for setting affinity on NUMA nodes. The MFS was only handling consecutive hyphen separated bit ranges. For example, MFS would only handle "0-3,8-11" and would not handle "0,4,8,12".

Resolution

With this fix, MFS will now support both comma and hyphen separated CPU list for setting NUMA affinity.

Bug 22898

Details

On nodes with more than one MFS instance, if the node topology changed, the local volume topology did not change and the topology of the local volume continued to show the old topology. For example, on a node with 2 MFS instances, if the node topology was pointing to `/data/default-rack/nodeIP` and then changed to `/data/rack/nodeIP`, the local topology continued to point to `/data/default-rack/nodeIP` instead of `/data/rack/nodeIP`.

Resolution

With this fix, on nodes with more than one MFS instance, local volume topology will change if the node topology changes.

Bug 22924

Details

MapR-DB client applications calling the `hb_get_add_column()` C function were required to use qualifier names in lexicographic order when creating more than one column in a column family.

Resolution

With this fix, the requirement is removed.

Bug 22948

Details

The `mfs` service would terminate unexpectedly in the following situation: the primary disk in a storage pool was removed, thereby removing the metadata stored on that disk about the storage pool; an attempt was then made to add a secondary disk to the storage pool.

Resolution

With this fix, the `mfs` service no longer terminated unexpectedly in this situation.

Bug 23131

Details

The NFS server running on 5.1 cluster was hanging when it tried to access a cluster running older versions. The new features introduced in 5.1 are not available on older releases and the NFS server, instead of returning an error, was hanging while attempting to access the cluster running the older version.

Resolution

With this fix, the NFS server will return a "feature not present" error instead of hanging.