

Version 1.2.10 Release Notes

Release Information

MapR provides the following packages:

- Apache Hadoop 0.20.2
- Flume 0.9.4
- Hbase 0.90.6
- Hive 0.7.1
- Mahout 0.5
- Oozie 3.0.0
- Pig 0.9.0
- Sqoop 1.3.0
- Whirr 0.3.0

New in This Release

This is a maintenance release. No new features.

Resolved Issues

MapReduce

- (Issue 6132) High-availability JobTracker now logs jstack and automatically restarts when the JT process exists but is unresponsive for a configurable amount of time. The default configuration is ten minutes.
- (Issue 7628) TaskTracker won't run if the node hostname resolves to an IP that is not covered by the value in the `MAPR_SUBNETS` environment variable.
- (Issue 8115) Split calculation now correctly handles a MapR-FS chunk size of zero. The client JVM no longer spins at 100% CPU utilization while submitting a MapReduce job.
- (Issue 7310) MapReduce task preemption now works correctly when tasks are scheduled with ExpressLane.

File System

- (Issue 6021) Cleaning multiple inodes in a container now creates a single update thread.
- (Issue 6717) Potential memory leak in flushing dirty inodes in file system.
- (Issue 7894) Mirroring no longer causes an assert failure by losing track of the most recent replica of a blank source container.
- (Issue 7924) The number of working RPCs reported after a failed `GetData` attempt is now correct.
- (Issue 7925) A node going down no longer creates a large number of containers in a volume.
- (Issue 7931) IOMgr no longer causes assert failures in the filesystem while trying to access the cache pages after failing.
- (Issue 7984) Name containers for large volumes are no longer stuck as under-replicated due to message timeout.
- (Issue 8071) CLDB no longer shows a stale container as master due to messages over 64k not being sent.
- (Issue 8092) fsck now works on Storage Pools over 6TB.
- (Issue 8139) CLDB no longer takes over 20 minutes to come up on clusters with large numbers of containers.
- (Issue 5537) Allocating buffers while buffers are still draining no longer causes segmentation faults.
- (Issue 8179) "disk online" operations no longer crash with an EIO error.
- (Issue 7502) Mirroring: Rollforward operations now correctly update the container epoch, allowing for mirror restarts.
- (Issue 7657) Nodes with a very large number of containers no longer become stuck waiting on a `BECOME_MASTER` command to process.
- (Issue 7881) Write failures now correctly throw exceptions.
- (Issue 7599) MFS cache now populates with correct hostname information for containers.
- (Issue 7685) Disk balancer selects host nodes for new container replicas more evenly.
- (Issue 7722) MFS now supports more than 256 groups.
- (Issue 7752) MFS no longer returns null hostnames on rare occasions.
- (Issue 7691) New command `hadoop mfs -lsrv <path>` recursively lists all paths within a single volume.
- (Issue 7575) Renaming of `mapr.cluster.root` now supported.

Logging

- (Issue 7677) Debug log levels can be set from the command line.
- (Issue 7793) Updating the hostname entry in `serverTab_` before calling `createMapRBlockLocation()` no longer generates garbage

- characters in the JobTracker log.
- (Issue 7653) Excessive MFS logging fixed.
- (Issue 7810) More details are logged, including:
 - Where a job was killed from (MCS or API)
 - User that sent the API call to kill a job
 - IP address sending the API call to kill a job
- (Issue 7892) Balancing classifications for CLDB logs no longer show the same node multiple times.

MCS and CLI

- (Issue 8206) Spurious volume alarms no longer being raised.
- (Issue 7989) The output of the `maprcli dump rereplicationinfo -json` command now includes the start time of the resyncing operation.
- (Issue 8047) Nodes in maintenance mode now display correctly in MCS.
- (Issue 8404) **Volumes** tab in the MCS UI now behaves properly.

NFS

- (Issue 8260) NFS client no longer hangs while accessing files with chunksize set to 0.
- (Issue 8136) The `maprcli nfsmgmt refreshexports` command no longer generates a buffer overflow when `/opt/mapr/conf/exports` is over 1024 bytes in size.

Alerts

- (Issue 5941) Email alerts now dispatch correctly.

Known Issues

(Issue 10598)

When a node's IP number is reassigned to another node, remote procedure calls are misdirected to the wrong node. This can result in widespread errors affecting any system that use the RPC binding layer, including container resync operations that never finish or the creation of several superfluous containers. Restart the cluster to clear this condition.

(Issue 11389)

If you use the latest version of any browser to view the MCS, the CLDB, HBase, JobTracker, and job log pages do not load when you click on the respective links. As a workaround, use the following instructions to disable mixed-content-blocking in your browser:

Internet Explorer

Scroll to the bottom of the screen, and click Show all content. The page refreshes and displays mixed content.

Mozilla Firefox

1. At the top of the page, to the left of the address bar, click (the shield icon).
2. In the pop-up window that appears, click the down arrow next to "Keep Blocking", and select Disable Protection on This Page. The page refreshes and displays mixed content. The shield icon is replaced with a warning symbol to indicate that the page is displaying both secure and non-secure items.

Google Chrome

1. Click (the shield icon) on the right side of the address bar.
2. In the icon dialog box, click Load anyway. The page refreshes and displays mixed content. The URL in the address bar shows https crossed out to indicate that the page is displaying both secure and non-secure items.