



## Release Notes— Fluid Cache for SAN 2.0.0 for Linux Systems

**Build Version: 2.0.0**

Updated 1/27/16



## Contents

<b>SOFTWARE AND HARDWARE COMPATIBILITY .....</b>	<b>3</b>
<b>FIRMWARE AND DRIVER COMPATIBILITY .....</b>	<b>4</b>
<b>LINUX OPERATING SYSTEM DEPENDENCIES.....</b>	<b>5</b>
<b>INSTALLATION INSTRUCTIONS.....</b>	<b>6</b>
VERIFYING THE DIGITAL SIGNATURE OF A PACKAGE .....	6
<b>KNOWN ISSUES ARISING FROM THIRD PARTY SOFTWARE.....</b>	<b>7</b>
LOW MEMORY HANDLING ISSUES.....	7
OPENMANAGE INSTALLATION CAUSES RDMA TO FAIL.....	7
JDK INSTALLER GENERATES LOG MESSAGES .....	7
VOLUME MANAGER DISK CREATION.....	7
XFS HANGS DURING FILE SYSTEM CREATION .....	8
XFS HANGS DURING NORMAL SYSTEM OPERATION .....	8
BTRFS HANGS DURING NORMAL SYSTEM OPERATION.....	8
BTRFS HANGS DURING NORMAL SYSTEM OPERATION.....	8
EXT4 DATA INTEGRITY ISSUES .....	8
<b>KNOWN ISSUES IN FLUID CACHE FOR SAN 2.0.0 .....</b>	<b>9</b>
REMOVING A SERVER FROM A CLUSTER.....	9
REMOVING MULTIPLE SERVERS FROM A CLUSTER.....	9
SHUTTING DOWN MULTIPLE SERVERS FROM AN FLDC CLUSTER .....	9
REPLACING A FAILED NODE ON A THREE NODE CLUSTER .....	9
FILE SYSTEMS TABLE (FSTAB) SETTINGS PREVENT SERVER FROM STARTING .....	11
ROOT FILE SYSTEM FILLS UP .....	12
CONFIGURATION ERRORS AFTER MAPPING VOLUME FAILS.....	12
TIMEOUTS ON POWEREDGE R820 SERVERS.....	12
GPT PARTITION TABLE CORRUPTION.....	12
USE OF KPRTX TO PARTITION VOLUMES .....	13
SYSLOG ERROR MESSAGES .....	13
STOPPING THE FLUID CACHE SERVICE CREATES SEGMENTATION FAULT.....	13



## Software and Hardware Compatibility

Table 1 Hardware and Software Requirements

Element	Requirement
Servers	Dell PowerEdge systems that support Dell Express Flash PCIe SSDs: M620, M820, R620, R720, R820, R920, T620, R730, R730XD, R630, or T630
Operating Systems	<ul style="list-style-type: none"> <li>Red Hat Enterprise Linux 6.4 (64-bit)</li> <li>Red Hat Enterprise Linux 6.5 (64-bit)</li> <li>SUSE Linux Enterprise Server (SLES) 11 SP3 (64-bit)</li> <li>Oracle Enterprise Linux 6.4 (64-bit)<sup>1</sup></li> <li>Oracle Enterprise Linux 6.5 (64-bit)<sup>2</sup></li> </ul>
RAM and Hard Disk Space	<ul style="list-style-type: none"> <li>Minimum of 64 GB RAM</li> <li>150 GB of available disk space</li> </ul>
Cache Device	<ul style="list-style-type: none"> <li>Express Flash PCIe SSD of either 175GB, 350GB, 400GB, 800GB, or 1.6TB capacity</li> <li>Micron half-height, half-length SSD of either 700GB or 1.4TB capacity</li> </ul>
Network Adapter	<ul style="list-style-type: none"> <li>Mellanox ConnectX-3 Dual Port 10 GbE SFP+ Adapter</li> <li>Mellanox ConnectX-3 Dual Port 40 GbE QSFP+ Adapter</li> <li>Mellanox ConnectX-3 Dual Port 10 GbE KR Mezzanine Adapter</li> </ul>
Fluid Cache Private Network Switch	<p>The following switches have been validated for use with Fluid Cache.</p> <ul style="list-style-type: none"> <li>Dell Networking S4810</li> <li>Dell Networking S6000</li> <li>Dell Networking MXL 10/40 GbE</li> <li>Cisco Nexus 5584UP</li> </ul>
Storage Center Software	<ul style="list-style-type: none"> <li>Dell Compellent Storage Center 6.5 or higher</li> <li>Dell Compellent Enterprise Manager 2014 R2 or higher</li> </ul>
Storage Center Hardware	<ul style="list-style-type: none"> <li>Dell Compellent SC8000 controller</li> </ul>
Linux Dependencies	<ul style="list-style-type: none"> <li>See list below</li> </ul>
Optional: Cache Client Servers (have no SSDs but participate in the cache cluster)	All Dell PowerEdge systems or non-Dell systems with a supported operating system that can install a supported Mellanox Ethernet adapter

<sup>1</sup> For Oracle Linux 6.4, use the RHEL 6.4 Fluid Cache RPM package.

<sup>2</sup> For Oracle Linux 6.5, use the RHEL 6.5 Fluid Cache RPM package.



## Firmware and Driver Compatibility

Table 2 Firmware and Driver Minimum Versions

Element	Minimum Version
Mellanox ConnectX-3 Drivers	SLES 11 SP3: 2.0-2.6.8 RHEL 6.4: 2.0-2.6.8 RHEL 6.5: 2.3-1.0.1 OEL 6.4: 2.0-2.6.8 <sup>3</sup> OEL 6.5: 2.3-1.0.1
Mellanox ConnectX-3 Firmware	2.30.5118
Express Flash PCIe SSD Drivers	Micron: 3.3.0 Samsung: 1.11
Express Flash PCIe SSD Firmware	Micron: B1490908 Samsung: IPMODD3Q170
PowerEdge M620 BIOS	2.1.6
PowerEdge M820 BIOS	1.7.3
PowerEdge R620 BIOS	2.1.3
PowerEdge R720 BIOS	2.1.3
PowerEdge R820 BIOS	1.7.2
PowerEdge R920 BIOS	1.2.2
PowerEdge T620 BIOS	2.1.2
iDRAC	1.51.51
Dell Lifecycle Controller	1.3.0.850
Java JRE for EM/SC Managers (both 32-bit and 64-bit required)	Jre-7u25
Compellent Storage Center 6.5	R06.05.02.012
Compellent Enterprise Manager 2014 R2	14.2.2.6

Table 3 Driver Required Versions

Element	Required Version
Mellanox ConnectX-3 OFED Drivers	SLES 11 SP3: 2.0-2.6.8 <a href="#">[Download]</a> RHEL 6.4: 2.0-2.6.8 <a href="#">[Download]</a> RHEL 6.5: 2.3-1.0.1 <a href="#">[Download]</a> OEL 6.4: 2.0-2.6.8 <a href="#">[Download]</a> OEL 6.5: 2.3-1.0.1 <a href="#">[Download]</a>

<sup>3</sup> The Mellanox ConnectX-3 Driver version 2.0-2.6.8 will not by default install on OEL6.4 systems. Workaround: edit the distro file for RHEL6.4 and change the RHEL6.4 to OEL6.4.



## Linux Operating System Dependencies

To check which dependencies are currently installed on your system, run the following command: `rpm -qa`

Ensure that the following open source component packages are available in your system before you install Fluid Cache for SAN. See the operating system distribution media for the supported version of each package.

The following dependencies are required for RHEL, except as noted. The dependencies for SLES are similar. For Oracle Linux, please contact Dell Customer Support.

- avahi-libs
- avahi-tools (RHEL )
- avahi-utils (SLES)
- bash
- glib2
- glibc
- kernel
- libaio
- libibverbs\*
- librdmacm\*
- libuuid
- libxml2
- libxslt
- openssl
- pam
- perl
- perl-XML-LibXML
- perl-XML-Namespacesupport
- perl-XML-SAX
- per-YAML-LibYAML
- python
- sg3\_utils-libs
- sg3\_utils
- shadow-utils
- xmlsec1-openssl
- xmlsec1
- zlib

\* Installed when you install the Mellanox OFED package.



## Installation Instructions

---

Refer to the *Deployment Guide* for complete installation instructions. However, before deploying Fluid Cache you must verify the digital signature of the Fluid Cache package.

### Verifying the Digital Signature of a Package

To verify a digital signature on a Linux system, you must have the Gnu Privacy Guard (GPG) package installed.

1. Get the Dell Linux public GnuPG key by searching <http://pgp.mit.edu/> for this footprint:  
0x1285491434d8786f

2. Import the public key by running the following command:

```
gpg --import <public_key>
```

A message appears, informing you that the public key was processed but was unchanged.

3. Validate the Dell public key, if you haven't done so previously, by entering the following command: `gpg --list-keys`

4. Trust the Dell public key with the command `gpg -edit-key 23B66A9D` and then enter the command `trust`.

5. Decide how far you trust this user to correctly verify other users' keys (by looking at passports, checking fingerprints from different sources, and so on). Select the highest number/level of trust, and then enter the command `quit`.

6. Verify the file's digital signature by entering the following command:

```
gpg --verify <file_name>.sign <file_name>
```

A message appears, telling you that the verification resulted in a good signature.

**NOTE:** If you have not validated the key as shown in the previous step, additional messages appear, warning you that the key was not certified with a trusted signature.



## Known Issues Arising from Third Party Software

---

The following issues are not caused by Fluid Cache itself, but arise from known issues in third party software.

### Low Memory Handling Issues

**ISSUE:** When the system is low on memory, memory allocation errors, I/O timeouts, and EAGAIN errors (resource temporarily unavailable) may appear.

**WORKAROUND:** Add or free up RAM memory.

### OpenManage Installation Causes RDMA to Fail

**ISSUE:** On SLES 11 SP3, if you install the Mellanox drivers and afterward install OpenManage 7.4, the installation of libmlx4-rdmav2 breaks the already installed RDMA configuration and Fluid Cache runs in TCP mode instead of RDMA.

**WORKAROUND:** Reinstall the Mellanox drivers after installing OpenManage and before installing Fluid Cache.

### JDK Installer Generates Log Messages

**ISSUE:** Installing Fluid Cache after installing jdk-7u6-linux-x64.rpm causes the JDK installer to send harmless messages to log. This is a known Oracle and Java bug that occurs only on SLES.

**WORKAROUND:** Add a Required-Stop section to the jexec init script to stop the messages from being generated. In the file `/etc/init.d/jexec`, enter the following line:

```
# Required_Stop:
```

### Volume Manager Disk Creation

**ISSUE:** Attempts to use Volume Manager to create a volume on a Fluid Cache disk fail with the following error: `Device /dev/fldc0 not found (or ignored by filtering)`

**WORKAROUND:** Volume Manager is not supported by Fluid Cache, and should not be used to create volumes out of Fluid Cache disks. Instead, create a volume on the backend store and then cache that volume.



## XFS Hangs During File System Creation

**ISSUE:** On RHEL 6.4, the XFS file system may become unresponsive. During creation of the file system, this may occur because of a deadlock between the xfsyncd and xfsbufd daemons. Consequently, several kernel threads become unresponsive and the XFS file system cannot be successfully created, leading to a kernel oops. A patch is available to prevent this situation by forcing the active XFS log onto a disk. For more details, see the bug report at [https://bugzilla.redhat.com/show\\_bug.cgi?id=1014867](https://bugzilla.redhat.com/show_bug.cgi?id=1014867).

**WORKAROUND:** A patch that addresses this issue is available at <http://rhn.redhat.com/errata/RHSA-2013-1645.html>.

## XFS Hangs During Normal System Operation

**ISSUE:** On RHEL 6.4 and SLES 11 SP3, the XFS file system may become unresponsive when running a large number of file operations. For more details, see the bug report at [http://oss.sgi.com/bugzilla/show\\_bug.cgi?id=922](http://oss.sgi.com/bugzilla/show_bug.cgi?id=922).

**WORKAROUND:** N/A

## Btrfs Hangs During Normal System Operation

**ISSUE:** Under some circumstances, the Btrfs file system becomes unstable, goes into read-only mode, and becomes unresponsive.

**WORKAROUND:** N/A

## Btrfs Hangs During Normal System Operation

**ISSUE:** When using the Btrfs file system on SLES 11 SP3, under some circumstances the kernel can become so busy dumping stack traces that the operating system starts sending non-maskable interrupts (NMI) and becomes unresponsive.

**WORKAROUND:** N/A

## Ext4 Data Integrity Issues

**ISSUE:** When using the Ext4 file system, under some circumstances data at the end of a cached file may be lost, or data within a file may be lost.

**WORKAROUND:** N/A





---

## Known Issues in Fluid Cache for SAN 2.0.0

---

The following are known issues in Fluid Cache for SAN 2.0.0

### Removing a Server from a Cluster

**ISSUE:** If you try to remove a server from a cluster before shutting down the server, Enterprise Manager displays a message saying that the action is not allowed.

**WORKAROUND:** Shut down the server: In Enterprise Manager, locate the server, right-click, and select **Shut Down**. Then continue with the *Deployment Guide* procedure “Removing a Server from a Cluster.”

### Removing Multiple Servers from a Cluster

**ISSUE:** Under some circumstances, remove more than one server at a time from a cluster may cause I/O to hang or data to be lost.

**WORKAROUND:** Avoid removing more than one server at a time from a cluster.

### Shutting Down Multiple Servers from an FLDC cluster

**ISSUE:** Shutting down more than one server (host) from an FLDC cluster for maintenance purposes may cause potential performance issues and data loss.

**WORKAROUND:** Dell recommends not shutting down multiple FLDC Cluster hosts at the same time to perform the updates.

To add new hardware, update BIOS, firmware version, OS level patches, or driver version on an FLDC cluster host, you must perform one of the following steps:

- Shut down a single FLDC Cluster host at a time and perform the updates
- Shut down the entire cluster during a scheduled maintenance period and perform the updates

### Replacing a Failed Node on a Three Node Cluster

**ISSUE:** If one of the nodes in a three node cluster fails, Enterprise Manager does not allow removal of the failed node, and displays a message that removing the node would result in fewer than the minimum of three nodes required. If the node is manually removed and returned to an operational state, Enterprise Manager does not allow it to rejoin the cluster, displaying a message that the node already belongs to the cluster.

Dell - Internal Use - Confidential



**WORKAROUND:** Image a fourth node and add it to the cluster, then remove the failed node from the cluster.



## File Systems Table (fstab) Settings Prevent Server from Starting

**ISSUE:** The server can fail to boot because of errors related to setting non-zero `fs_passno` field in the `/etc/fstab` file. Setting the force `fsck` option causes the system boot to fail. This is because the devices do not exist at boot time when `fstab` is processed. When the error occurs, the following message is displayed:

```
/dev/fldc0 is mounted. e2fsck: Cannot continue, aborting.  
/dev/fldc1 is mounted. e2fsck: Cannot continue, aborting.  
*** An error occurred during the file system check.  
*** Dropping you to a shell; the system will reboot  
*** when you leave the shell.
```

The sixth field (`fsck` option) is used by the `fsck(8)` program to determine the order in which file system checks are done at reboot time. If the sixth field is set to zero, these system check errors do not occur during reboot. Therefore it is recommended that for any `fldc`-related user entry in `/etc/fstab`, the fifth field (`dump` option) and the sixth field (`fsck` option) are both set to zero. In case the sixth field is not zero, it is very likely the above Linux error message occurs during reboot.

**WORKAROUND:** Remount the root file system as read/write (`mount -n -o remount, rw /`) and then comment out the Fluid Cache disk device entries from `fstab`, as shown:

```
#/dev/fldc0 /root/fldc/mnt ext2 defaults 1 3
```

Then change:

```
/dev/fldc1 /root/fldc/mnt2 ext2 defaults 1 4
```

either to:

```
/dev/fldc1 /root/fldc/mnt2 ext2 defaults 1 0
```

or to:

```
#/dev/fldc1 /root/fldc/mnt2 ext2 defaults 1 4
```



## Root File System Fills Up

**ISSUE:** Fluid Cache may fail and require manual intervention if the root file system on a node becomes full. The cluster may still be operational, but error messages are generated and Fluid Cache is not fully functional.

**WORKAROUND:** Please contact Dell Customer Support for the steps required to return the node to full functionality.

## Configuration Errors after Mapping Volume Fails

**ISSUE:** When mapping a volume to a Fluid Cache node, if the operation times out or fails to complete normally, under some circumstances a partial configuration is created on the node, and Enterprise Manager mistakenly shows an apparently normal volume mapping. When the mapping is later completed normally, this partial configuration remains, and may interfere with administrative actions. For instance, deleting the cluster may fail because Storage Center has a record of a volume still in use by the cluster.

**WORKAROUND:** Contact Dell Customer Support for assistance in resolving this issue.

## Timeouts on PowerEdge R820 Servers

**ISSUE:** On SLES 11 SP3, under some circumstances, I/O loads could cause CPU usage spikes that decrease performance on PowerEdge R820 Servers.

**WORKAROUND:** None available.

## GPT Partition Table Corruption

**ISSUE:** Partprobe reports GUID Partition Table (GPT) corruption and device mapper errors on all other nodes after making partition table changes to cached LUNs on one node.

**WORKAROUND:** Do not modify partition tables on cached LUNs. Modifying partition tables on shared volumes is not supported.



## Use of Kpartx to Partition Volumes

**ISSUE:** When using the `kpartx` command to partition volumes, the partition table on shared block devices is in an inconsistent state across the cluster. For instance, partitions created on a LUN on one node may not be seen by other nodes.

**WORKAROUND:** `Kpartx` is not supported for partitioning a cached volume. Use the `partprobe` command instead.

## Syslog Error Messages

**ISSUE:** The `avahi-daemon` writes numerous messages to `syslog` about receiving a packet from an invalid interface when there is another mDNS stack (such as the one that Oracle RAC installs) installed and running on the same node as Fluid Cache.

**WORKAROUND:** Move the other application's mDNS stack to another node, if possible. Otherwise, it is possible to patch and recompile the `avahi-daemon` to stop sending this message. Contact Dell Customer Support for instructions on patching the `avahi-daemon` or to obtain a replacement `avahi-daemon`.

## Stopping the Fluid Cache Service Creates Segmentation Fault

**ISSUE:** On SLES 11 SP3, stopping the Fluid Cache service results in a segmentation fault because of a conflict in the Avahi binary rpm.

**WORKAROUND:** Contact Dell Customer Support to obtain a patch with which to recompile the Avahi binary rpm, or to obtain a pre-compiled version of the Avahi binary rpm. You may also contact SLES to see whether a version of Avahi is available that is based on source version 0.6.24 or greater.