



Release Notes — Fluid Cache for SAN 2.0.10 for VMware Systems

Build Version: 2.0.10



Contents

PURPOSE OF THIS RELEASE.....	3
FIXED ISSUES IN THIS RELEASE.....	3
SOFTWARE AND HARDWARE COMPATIBILITY	6
KNOWN ISSUES ARISING FROM THIRD PARTY SOFTWARE	8
CANNOT ENABLE PASSTHROUGH FOR NVME DEVICE IN vSPHERE WEB CLIENT	8
KNOWN ISSUES IN FLUID CACHE FOR SAN 2.0.....	8
REMOVING A SERVER FROM A CLUSTER.....	8
REMOVING MULTIPLE SERVERS FROM A CLUSTER.....	8
SHUTTING DOWN MULTIPLE SERVERS FROM AN FLDC CLUSTER	9
REPLACING A FAILED NODE ON A THREE NODE CLUSTER	9
CONFIGURATION ERRORS AFTER MAPPING VOLUME FAILS.....	9
TUI EXITS WHEN PARSING IP ADDRESS WITH ZERO PADDED OCTETS.....	9
ISSUES WITH SAMSUNG NVME AND MICRON SSDS ON SAME SERVER.....	10
VSA AGENT TIMES OUT AND EXITS WHEN CONNECTING TO vCENTER	10
POWER-CYCLING THE ESXi HOST CAUSES CLAIM RULES TO BE REWRITTEN.....	10
CACHED VOLUME MAPPING AND UNMAPPING FAILURES CAN LEAVE STALE CLAIM RULES	11
BACKING LUN OF CACHED VOLUME APPEARS IN vCENTER.....	11
GUEST VM I/O TIMED OUT OR EXPERIENCED VERY LONG DELAYS.....	12
GUEST VMs BECOMES UNRESPONSIVE AND/OR ESXi HOST CANNOT BE MANAGED BY vSPHERE	12
NOT ALL FLUID CACHE VSAs STARTED CORRECTLY AFTER A CLUSTER SHUTDOWN OR ESX HOST POWER FAIL.....	13
UNMAPPING MULTIPLE CACHED VOLUMES CAN LEAVE CLUSTER IN A DEGRADED STATE	13



Purpose of this Release

Fluid Cache for SAN for VMWare 2.0.10 is a maintenance release with the following focus:

1. Increase the robustness of automatically recovering from cluster node failures/reboots.
2. Fixed reported memory leaks within the VSA
3. Add support for 3-9 node clusters and 1 node non-expandable

Note: Installation and Upgrade instructions for 2.0.10 are in the Dell Fluid Cache for SAN Version 2.0.10 Deployment Guide for VMware ESXi Systems

Fixed Issues in this Release

Table 1 Issues fixed in 2.0.10

Issue Summary
An unfortunately timed network outage can leave a CS permanently inactive
CS assert in rna_bsm�_allocate_vsbm_alloc
CS not expelled with san path failure
Block location check failure
cfm assert _rna_free in queued_san_make_data_available
remove duplicate code in tui by replacing them with common function
potential md list corruption in expel_cs() causing CFM crash
cfm logs warnings from queued_san_make_data_available and queued_query_config_cached_lun_callback
Network failure test CS assert(0 == (hp->cbh_flags & RNA_BSMD_HDR_FLAGS_VALID)) failed at bsm�_cache_blocks.c:603
VSA - After ESXi host power fail web page incorrectly showed cache servers down, etc for over 30 mins



2 nodes of 4 didn't shut down cleanly, came up with cachedev_failed and cachedev_failing
CS assertion in empty_entry_cache_req_queue, ce_cache_req_lock already held
CS assert in cs_assign_replica_data_eph() during cache restoration
Network failure test CS assert(0 == (hp->cbh_flags & RNA_BSMD_HDR_FLAGS_RECOVERABLE)) failed at bsmd_cache_blocks.c:604
CS assert in cache_entry_queue_send_replica_block_query(), query-delete case
client MD EAGAIN erroneous tight loop results in hard LOCKUP
Incorrect DNS name configuration in vsphere causes Tui crash
orphan clean replicas created by disconnects can block subsequent shutdowns
Issues with CACHE_DISCARD and replica allocation (Cache Server is stuck processing resilver complete)
block location check failed during node failure test
Mapping a volume took 1hr 20 minutes, exceeding multi full timeout threshold.
rna_service can get into tight re-send loop
ios timeout. maybe deadlock on cache_rw_lock
client kernel com panic with path failure (tcp)
MD sends a CACHE_MASTER_INVLD to the CS if it receives a CACHE_RESPONHSE for a block it doesn't know about
VSA plan B3 4 node Cluster shutdown then restart 3 non-cfm node test failed with vsa0106 wedged on startup and vsa0107 grayed out in vsphere
CS attempted removal of ACTIVE cache device
CS assert in rna_cs_create_existing_entry due to MD RID mismatch



unable to map cached volumes since doing a couple of VSA failures
Extra claimrules attached to a LUN in the journal.
claim rule files and path files in /tmp on the VSA are not being cleaned up
incomplete error handling in add_claim.sh when esx operations fail
lost all paths to devices on primary cfm fail/recover
cache server takes a very long time to reconnect to cluster after vsa failure
CS complaining about being unable to open a backing store device even though it is up accessible and available on all nodes
Client should handle TUR directly, don't pass them through to device
the timer for starting the set_preferred_path work isn't long enough and shouldn't be run while mapping
rnablk_generic_cmd accesses ios after com_send, can race with completion
MD asserted due a cache device mis-match in cache response
Block location check in node failure test - MD may have incorrectly declared an entry vestigial and deleted it
Block stuck in INVD_SENT in cache device fail test, ios timeouts result
Activation and connection status for not-activated CS should be displayed
ios timeouts due to stuck replica writes
ESXi 6.0 - "Set logical unit policy" vsphere tasks don't run after mapping LUNs
network failure test hit CS segfault in ping_ctx_register
CFM segfault in free_lun_info
Agent didn't restart CS when requested to do so



kernel OOM due to fldc_cache at 95pct
Initial resilvering doesn't complete for CS on restarted node

Software and Hardware Compatibility

Table 2 Hardware and Software Requirements

Element					
Servers	Dell PowerEdge systems that support Dell Express Flash PCIe SSDs and SR-IOV: M620, M820, R620, R720, R820, R920, or T620 R730, R730XD, R630, or T630				
Operating systems	<ul style="list-style-type: none"> ESXi 5.5 Update 2 (64-bit) Enterprise Plus ESXi 6.0 (64-bit) Enterprise Plus 				
RAM and hard disk space	<ul style="list-style-type: none"> Minimum of 64 GB RAM (the VSA uses 20 GB) A datastore with 30 GB of available disk space 				
Cache Devices	<ul style="list-style-type: none"> Dell Express Flash PCIe SSD of either 175GB, 350GB, 400GB, 800GB, or 1.6TB capacity Micron HDDL PCIe SSDs (700 GB or 1.4 TB) 				
Network adapters	<ul style="list-style-type: none"> Mellanox ConnectX-3 Dual Port 10 GbE SFP+ Adapter Mellanox ConnectX-3 Dual Port 40 GbE QSFP+ Adapter Mellanox ConnectX-3 Dual Port 10 GbE KR Mezzanine Adapter 				
FC Host Bus Adapters	<table border="0"> <tr> <td>Emulex:</td> <td>Qlogic:</td> </tr> <tr> <td> <ul style="list-style-type: none"> LPE12002 LPE12000 LPE16000B LPE16002B </td> <td> <ul style="list-style-type: none"> QLE2562 QLE2662 QLE2560 QLE2660 </td> </tr> </table>	Emulex:	Qlogic:	<ul style="list-style-type: none"> LPE12002 LPE12000 LPE16000B LPE16002B 	<ul style="list-style-type: none"> QLE2562 QLE2662 QLE2560 QLE2660
Emulex:	Qlogic:				
<ul style="list-style-type: none"> LPE12002 LPE12000 LPE16000B LPE16002B 	<ul style="list-style-type: none"> QLE2562 QLE2662 QLE2560 QLE2660 				
Cache network switches	The following switches have been validated for use with Fluid Cache. <ul style="list-style-type: none"> Dell Networking S4810 Dell Networking S6000 Dell Networking Z9000 Dell Networking Z9500 Dell Networking MXL 10/40 GbE Cisco Nexus 5584UP 				
Storage Center software	<ul style="list-style-type: none"> Dell Compellent Storage Center 6.5.2 or higher Dell Compellent Enterprise Manager 2014 R2 or higher 				
Storage Center hardware	<ul style="list-style-type: none"> Dell Compellent SC8000 controller 				

Dell - Internal Use - Confidential



Element	
Optional: Fluid Cache Client Servers (do not have SSDs but participate in a cluster)	<ul style="list-style-type: none">All Dell PowerEdge systems (or non-Dell systems with a supported operating system that can install a supported Mellanox Ethernet adapter and support SR-IOV)

Firmware and Driver Compatibility

Table 3 Firmware and Driver Minimum Versions

Element	Minimum Version
ESXi Mellanox ConnectX-3 Driver	1.9.10.2
ESXi Mellanox ConnectX-3 Firmware	2.30.5118 Rack 2.30.5160 Blades
Express Flash PCIe SSD Drivers	3.3.0
Express Flash PCIe SSD Firmware	B1490908 Micron IPM0DD3Q170 Samsung
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



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.

Cannot Enable Passthrough for NVMe Device in vSphere Web Client

ISSUE: Using the vSphere web client, you cannot enable passthrough for Samsung NVMe SSD devices.

WORKAROUND: Enable passthrough using the desktop version of vSphere.

Known Issues in Fluid Cache for SAN 2.0

The following are known issues in Fluid Cache for SAN 2.0.

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 VSA from vSphere. 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.

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

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.

TUI Exits When Parsing IP address with Zero Padded Octets

ISSUE: When configuring the VSA using the Text-based User Interface (TUI), entering a padded address (e.g., 172.19.2.018) causes the TUI to abruptly exit without a visible error message and



return the user to the VMware console. The TUI may not exit immediately after entry of the padded number, but when the number is first used by the TUI to configure the VSA.

WORKAROUND: Do not enter IP addresses with padded octets when using the TUI to configure the VSA.

Issues with Samsung NVMe and Micron SSDs on Same Server

ISSUE: Installing Samsung Electronics NVMe SSDs and Micron SSDs on the same server may cause heartbeat timeouts and other issues.

WORKAROUND: This mixed configuration is not supported by Dell. Each server in the Fluid Cache cluster may have Micron SSDs or Samsung NVMe SSDs, but not both.

VSA Agent Times Out and Exits When Connecting to vCenter

ISSUE: On some occasions, while attempting to connect to vCenter during startup, the VSA's agent service may timeout and then exit, causing the node to disconnect from the cluster. In this event, the VSA is not fully functional, although it is shown as operational in vSphere. Enterprise Manager shows the VSA as nonfunctional.

WORKAROUND: Restart the VSA.

Power-Cycling the ESXi Host Causes Claim Rules to Be Rewritten

ISSUE: In some circumstances, claim rules may be rewritten when power cycling an ESXi host.

WORKAROUND: Run the following commands on each ESXi host in the cluster to reset the claim rule list to the default values:

```
~ # esxcli storage core claimrule remove --plugin=MASK_PATH
~ # esxcli storage core claimrule add --type=vendor --
  vendor=DELL --model="Universal Xport" --plugin=MASK_PATH --
  rule=101
~ # esxcli storage core claimrule load
~ # esxcli storage core claiming unclaim --type=plugin --
  plugin=MASK_PATH
~ # esxcli storage core claimrule run
```



Then, add back any non-Fluid Cache claim rules, unmap all cached volumes (to clean out the journalled claim list rules), then remap the cached volumes.

Cached Volume Mapping and Unmapping Failures Can Leave Stale Claim Rules

ISSUE: In some circumstances the mapping and unmapping of cached volumes may fail leaving claim rules that hide known volumes on the ESXi host.

WORKAROUND: Run the following commands on each ESXi host in the cluster to reset the claim rule list to the default values:

```
~ # esxcli storage core claimrule remove --plugin=MASK_PATH
~ # esxcli storage core claimrule add --type=vendor --
  vendor=DELL --model="Universal Xport" --plugin=MASK_PATH --
  rule=101
~ # esxcli storage core claimrule load
~ # esxcli storage core claiming unclaim --type=plugin --
  plugin=MASK_PATH
~ # esxcli storage core claimrule run
```

Then, add back any non-Fluid Cache claim rules, unmap all cached volumes (to clean out the journalled claim list rules), then remap the cached volumes.

Backing LUN of cached volume appears in vCenter

ISSUE: When performing a multiple (bulk) cache mapping of LUNs there is the possibility that the backing (uncached) LUN will temporarily appear in vCenter. Important: this LUN should not be used or accessed. It is for internal Fluid Cache usage only.

A backing LUN can be identified by an 8 in the 8th octet from the right of its identifier. See the example below:

36000d31000eea100000000000800000aa – Backing LUN that may temporarily appear in vCenter
– DO NOT USE

36000d31000eea100000000000000000aa – identifier that would appear in Enterprise Manager
and is the real cached LUN.



WORKAROUND: No action is required. The backing LUN will only appear temporarily and will be cleaned up within 24 hours.

Guest VM I/O Timed Out or Experienced Very Long Delays

ISSUE: Under some circumstances when the cache contains a large amount of dirty cache blocks, power cycling a node or removing a cache device in the cluster may cause I/O to hang in one or more of the guest VMs in the cluster.

WORKAROUND: Set the cache mode to Write-Through to avoid filling the cache with dirty blocks.

Guest VMs Becomes Unresponsive and/or ESXi host cannot be managed by vSphere

ISSUE: Under some circumstances if the back end storage experiences performance issues, the TGT stack on the ESXi hosts can timeout causing guest VMs to become unresponsive and/or ESXi host to become unmanageable.

WORKAROUND: To recover from this situation reboot the ESXi host that is experiencing the issue. To help avoid future issues:

1. Review state of back end storage to identify performance issue and resolve.
2. Manually change the Queue Depth settings inside ESX using the following steps:
 - a. Use the vSphere Client or vSphere Web Client to navigate to the Configuration tab of the VMware ESX host you want to modify.
 - b. Click **Advanced Settings** under the Software section.
 - c. Click **Disk** in the left side pane.
 - d. Set QFullSampleSize to a value greater than zero. The usable range is 0 to 64.
 - i. Set the QFullSampleSize value to 32.
 - e. Set QFullThreshold to a value lesser than or equal to QFullSampleSize. The usable range is 1 to 16.
 - i. Set the QFullThreshold value to 8.



- f. The settings take effect immediately. You do not need to reboot the ESX/ESXi host.

Not all Fluid Cache VSAs Started Correctly After a Cluster Shutdown or ESX Host Power Fail

ISSUE: Under some circumstances VSAs are not started correctly following a cluster shutdown and restart, or an ESXi host power fail. This would appear in the EM console as a server that has not joined the cluster or cache devices are missing on a server.

WORKAROUND: Using vSphere, reboot the VSA on the server that is reporting issues.

Unmapping Multiple Cached Volumes Can Leave Cluster in a Degraded State

ISSUE: In some circumstances, during the unmapping of multiple cached volumes fluid cache servers can be left in a degraded state. This would appear in the EM console as a server that has left the cluster or a server that is missing cache devices.

WORKAROUND: Using vSphere, reboot the VSA on the server that is reporting issues.

Unable to add a Fluid Cache ESXi 6.0 host to an existing server cluster

ISSUE: Using Enterprise Manager GUI, if you try to add an ESXi 6.0 Fluid Cache host to an existing server cluster within a Storage Center, the following error message is displayed:

SC servers must all have the same operating system.

WORKAROUND: When adding a new Fluid Cache ESXi6.0 server host to an existing sub server cluster, in the **Edit Server Settings** dialog box, select **VMware ESXi 5.5** for the Operating System type instead of **VMware ESXi 6.0**. After adding the server to the sub server cluster, when Enterprise Manager prompts you to change the version, change the operating system type to VMware ESXi 6.0.

This method enables you to add an ESXi 6.0 host to an existing sub server cluster on the SC8000 without any issue.



Unable to add a Fluid Cache ESXi 6.0 host to an existing server cluster

ISSUE: When configuring Fluid Cache VSA iSCSI interface on the Fluid Cache Cluster ESXi5.5U3 and ESXi6.0U1 platforms during VSA installation, the following error message is displayed.

Cannot configure iSCSI initiator already configured for a different ESX host.

WORKAROUND 1 (Preferred):

Perform the following steps and configure the iSCSI section in the VSA setup.

1. In the <is it VMware vSphere Web client?> Web Client, open the VSA console, and enable SSH.
2. Log in to VSA through SSH.
3. Run the following command to navigate to the iscsi directory:

```
#cd /etc/iscsi
```

4. Open the initiatorname.iscsi file and verify if the value for the InitiatorName is in asterisk (*). For example:

```
root@localhost iscsi|# cat initiatorname.iscsi
InitiatorName=*****
```

5. Run the following command to edit the initiatorname.iscsi file.

```
#vi initiatorname.iscsi
```

6. Change the characters with asterisk (*) with the proper ESXi Hosts vmhba## software IQN that the VSA is associated. For example:

```
root@localhost iscsi|# cat initiatorname.iscsi
InitiatorName=iqn.1998-01.com.vmware:FLDC-QAESXi55U1-Srv1-48a43810
```

7. Run the following command to restart iscsid service.

```
#service iscsid restart
```

WORKAROUND 2:

Perform the steps provided in the section “Configuring the VSA” of the deployment guide in the following sequence:

Dell - Internal Use - Confidential



1. Perform the steps 1 through 5 as described in the deployment guide.
2. Select **Configure This Appliance** and press <Enter>.
3. Configure your Backend Storage connection.
4. Perform the steps 6 and 7 and continue with step 10 as provided in the deployment guide.