



Release Notes — Fluid Cache for SAN 2.1.0 for VMware Systems

Build Version: 2.1.0

Updated 5/10/16



Contents

PURPOSE OF THIS RELEASE 3

FIXED ISSUES IN THIS RELEASE 3

SOFTWARE AND HARDWARE COMPATIBILITY 4

KNOWN ISSUES ARISING FROM THIRD PARTY SOFTWARE..... 4

 CANNOT ENABLE PASSTHROUGH FOR NVME DEVICE IN vSPHERE WEB CLIENT4

KNOWN ISSUES IN FLUID CACHE FOR SAN 2.0 4

 REMOVING A SERVER FROM A CLUSTER.....4

 REMOVING MULTIPLE SERVERS FROM A CLUSTER.....5

 REPLACING A FAILED NODE ON A THREE NODE CLUSTER5

 CONFIGURATION ERRORS AFTER MAPPING VOLUME FAILS.....5

 TUI EXITS WHEN PARSING IP ADDRESS WITH ZERO PADDED OCTETS.....5

 ISSUES WITH SAMSUNG NVME AND MICRON SSDS ON SAME SERVER.....6

 VSA AGENT TIMES OUT AND EXITS WHEN CONNECTING TO vCENTER6

 POWER-CYCLING THE ESXi HOST CAUSES CLAIM RULES TO BE REWRITTEN.....6

 CACHED VOLUME MAPPING AND UNMAPPING FAILURES CAN LEAVE STALE CLAIM RULES7

 BACKING LUN OF CACHED VOLUME APPEARS IN vCENTER7

 GUEST VM I/O TIMED OUT OR EXPERIENCED VERY LONG DELAYS.....8

 GUEST VMs BECOMES UNRESPONSIVE AND/OR ESXi HOST CANNOT BE MANAGED BY vSPHERE8

 NOT ALL FLUID CACHE VSAs STARTED CORRECTLY AFTER A CLUSTER SHUTDOWN OR ESX HOST POWER FAIL.....9

 UNMAPPING MULTIPLE CACHED VOLUMES CAN LEAVE CLUSTER IN A DEGRADED STATE9



Purpose of this Release

Fluid Cache for SAN for VMWare 2.1.0 is a maintenance release with the following focus:

- Improve diagnostic capability when there are cluster issues
- Improve cluster shutdown robustness
- Fix stability issues and improve cluster operation
- Support rolling upgrades from 2.0.10

Note: Installation and Upgrade instructions for 2.1.0 are in the *Dell Fluid Cache for SAN Version 2.1.0 Admin and Deployment Guide for VMware ESXi Systems* available at dell.com/CacheSolutions.

Fixed Issues in this Release

Improve diagnostic capability

- Reduced log lines
- Eliminate fldccollect filtering
- Reduce max log file size
- Consistent log format
- Stuck SCSI2 logging
- Enable SSH by Default
- RoCE network verification
- Past anomaly detection

Cluster management usability

- CFM Shutdown Intelligence – do not kill all
- Rolling Upgrades automation & validation
- Ctl-Alt-Delete: Prevent accidental cluster shutdown
- Marking blocks truly clean – quicker recovery time
- Protect data when not following process in VSA configuration
- Improve handling of CFM crash during uncaching



Operational improvement

- Issues with shutting down VSA
- Issues with write phase of compare and write
- Shutdown failure scenario
- MD Crash due to unresolved duplicates
- Performance degradation during unaligned I/Os
- Cache device could not be used
- I/O timeouts and errors after cache network failure
- Data loss in multiple failure scenario
- Mellanox updated driver (fixes a PSOD)

Software and Hardware Compatibility

See *Dell Fluid Cache for SAN Compatibility Matrix* available at dell.com/CacheSolutions.

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.

Cluster-wide outage due to corruption in journal

ISSUE: In some circumstances, if you reset the CFM VSA, there is a rare possibility of cluster-wide outage due to corrupted journal.



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

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 that the action is not allowed.

WORKAROUND: Shut down the VSA from vSphere, and then to remove the server from a cluster, refer to the section “Removing a Server from a Cluster” in the *Admin and Deployment Guide*.

Removing Multiple Servers from a Cluster

ISSUE: Under some circumstances, removing more than one server at a time from a cluster may cause I/O to stop responding or leads to data loss.

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

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, and 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 by mistake 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 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 by 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. 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. 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.