



SANtricity® Operating System (Controller Firmware) 8.20.20.00

Product Release Notes

September 2016 | 215-09078_A0_ur011
doccomments@netapp.com



Contents

Deciding whether to use this guide.....	5
New in SANtricity OS 8.20.....	6
E-Series and EF-Series hardware	8
Resolved issues and software dependencies	11
Important issues resolved in SANtricity OS 8.20 maintenance releases	11
Restrictions.....	12
With Veritas, online controller firmware upgrade causes loss of both paths	12
Asynchronous Mirroring volumes are failed during a volume	12
reconfig or Unmap operation if a volume transfer occurs	12
Oracle Java SE Runtime versions later than 8u31 but prior to 8u65 cause random GUI issues	13
Oracle Java SE Runtime versions prior to 8u65 cause random GUI issues on hosts running Solaris	13
Volumes might not stay on the preferred path during an online controller firmware upgrade.....	14
Choosing a custom segment size.....	14
A path failure occurs during a controller firmware upgrade	15
The presence of host I/O during management software upgrade prevents the E-Series Windows DSM from updating properly	15
A host panics during path failover with an Emulex Fibre Channel (FC) host bus adapter (HBA)	16
Emulex Fibre Channel (FC) host bus adapter (HBA) link errors result in a kernel panic.....	16
When the <code>user_friendly_names</code> parameter is enabled in <code>/etc/ multipath.conf</code> , paths are mapped incorrectly	17
A Red Hat Enterprise Linux (RHEL) kernel bug causes a soft lockup under high memory load due to swap thrashing.....	17
I/O error occurs on SUSE Linux Enterprise Server (SLES) VMs with BusLogic SCSI controllers after a controller reboot.....	18
A host with Red Hat Enterprise Linux (RHEL) 7 crashes during a controller firmware upgrade.....	19
Cluster resources cannot be transferred when using network file system (NFS) version 4.....	19
I/O errors occur during an online controller firmware upgrade with I/O activity	20
The <code>multipath-tools</code> command fails to add the physical path to device- mapper multipath map after a failed path is restored	20
Inconsistency in volume ownership transfer occurs during a parallel drive firmware update (PDFU).....	21
I/O errors occur during online controller firmware update	22
I/O errors occur during online controller firmware update	22
A suspect drive is marked as failed unexpectedly.....	23
An early path failover occurs when a host cable is pulled	23

I/O operations fail when SG_IO is used with a block size larger than 512 KB	24
Random read IOPs for cache and solid state disk (SSD) performs below anticipated targets with the data assurance feature	24
An iSCSI session does not persist after a host reboot.....	25
System fails to boot from a local disk with enabled <code>multipathd</code> and the root device not blacklisted.....	25
Java access with speech (JAWS) reads hidden table data in SANtricity Storage Manager	26
iSCSI SAN boot failure	27
While performing a storage VMotion operation, virtual machines (VMs) occasionally fail to transfer properly when under I/O load.....	28
File systems are unmounted during a controller reboot	29
Cannot use <code>dsmUtil.exe</code> to apply load balance policy setting on a per-device basis	30
After a controller tray power failure, some Seagate drives in expansion trays that did not lose power are unresponsive	30
Snapshot image creation using VSS fails.....	31
Drives are bypassed during a controller reboot	31
Snapshot image creation via DiskShadow fails.....	32
I/O errors occur during an online controller firmware update	32
An asynchronous remote volume mirror (ARVM) mirrored pair is marked as failed	33
When an ESM fails or is removed during heavy I/O operations, the controller panics and reboots	33
Hosts continually attempt to rebalance or perform failback on secondary volume of a mirrored pair	34
Excessive unnecessary traffic and alarming host message log entries generated on Linux hosts when a secondary volume in a mirrored pair is mapped to LUN 0	34
Host port creation or deletion is reported unexpectedly	35
Pacemaker node unable to rejoin cluster after node reboot.....	35
VMware ESXi 6.0 host does not detect volumes after upgrade if LUN 0 unused.....	36
Unable to install SLES 10 SP4 as a VM	37
Loss of paths after RHEV hypervisor host reboot.....	38
LUNs are not visible to Solaris 10 hosts with 16GB HBA directly connected to the storage using 16GB speed	38
SMagent and SMdevices not functional on RHEL7.1	39
Windows SAN Boot installation failure.....	39
Solaris 10 may encounter I/O errors when array controller reboots.....	40
RHEL6.7 (IB-iSER) xfs filesystem I/O timeout may occur when storage array is reset	41
RHEL6.6 disk offlined during cable pull	41
Oracle Linux 6.6 disk offlined during cable pull.....	42
iSCSI sessions fail to restart after a OEL6.7 UEK3 U5 host reboot.....	43

Data Assurance - A transient data loss may occasionally occur with 2 failed
 drives in a RAID 6 volume group that is still formatting 43

 Linux devices offlined with multiple ISLs in a Cisco UCS config 44

 ATTO Multipath Director fails to claim LUNs after Windows 10 upgrade 45

Standard features and premium features 46

Copyright information..... 47

Trademark information..... 48

**How to send comments about documentation and receive update notifications
..... 49**

Deciding whether to use this guide

This document includes release notes for SANtricity OS (controller firmware) 8.20.20.00

Where to find the Latest information about the product

You can find information about the latest version of the product, including new features and fixed issues, and a link to the latest documentation at the following address: mysupport.netapp.com.

New in SANtricity OS 8.20

New features

The following table shows features introduced in current and previous.

Release	Feature	What's New
SANtricity OS 8.20.20.00	No new features	
SANtricity OS 8.20.18.00	No new features	
SANtricity OS 8.20.11.00 / SANtricity Storage Manager 11.20.XX00.0015	FDE capable SSDs with SSD Cache	Support for FDE capable SSDs with SSD Cache Feature is added.
	E5600 and EF560 with InfiniBand host interface can use SRP	E5600 and EF560 controller-drive trays with InfiniBand host interfaces can now be converted from iSCSI Extensions for Remote Direct Memory Access (iSER) to SCSI RDMA Protocol (SRP)
SANtricity OS 8.20.09.00 / SANtricity Storage Manager 11.20.XX00.0010	Decrease RAID 1 rebuild time by increasing the rebuild block size	Performance is improved when you rebuild RAID 1 volumes.
SANtricity OS 8.20.08.00 / SANtricity Storage Manager 11.20.XX00.0006	Media scan with redundancy check enabled by default	A media scan with redundancy check will be enabled by default when a volume is created. Please see Media scan with redundancy check enabled by default for a detailed description of this change.
SANtricity OS 8.20.05.00 / SANtricity Storage Manager 11.20.XX00.0002	Reserved drive space	The storage capacity reserved for configuration information (and other system data) on each drive in a volume group or a disk pool is increased from 500 MB to 5.5 GB. Please see Reserved drive space for a detailed description of this change.
	Storage array support	The storage management software now supports the E5600 storage array and the EF560 storage array.

Release	Feature	What's New
	Online drive firmware update	Adds the ability to update drive firmware while the storage array is operating and online. All selected drives are updated while the volumes stay online and the storage array responds normally to read and write requests.
	Drive support	4-K sector drives are supported through 512 emulation.
	Disk pool enhancements	When a disk pool is being created, configuration candidates that offer tray loss protection or drawer loss protection are now identified and listed first. Tray loss protection and drawer loss protection ensure accessibility to your data if a total loss of communication occurs with a single drive tray or drawer. Your storage array must have six trays or drawers to be capable of this protection. Also, this implementation is focused on initial configuration; the tray loss protection or drawer loss protection is not maintained when additional drives are added to the disk pool.
	Power cycle drives automatically	If a drive is in an unresponsive state, the firmware automatically cycles power to the drive in an attempt to correct the problem. (The individual drive is powered off, and then powered on.) If the drive does not recover, the drive is failed. If the drive does recover, reconstruction completes quickly because only the changed data needs to be reconstructed instead of the entire drive. There is also a new CLI command for user-initiated power cycling of an individual drive in a storage array.
	Data evacuation	If failure of a drive is predicted, the data on that drive is automatically copied from the drive and the drive is powered down. If the affected drive is in a volume group, a hot spare must be available for automatic data copying. Data is not copied if the drive failure is not predicted. Also, if you manually fail a drive, you can copy the data from the drive to enable faster drive reconstruction.

Release	Feature	What's New
	Data Center Bridging Exchange (DCBX)	The iSCSI statistics report now includes data pertaining to DCBX. This data is available only on the E2700 controller or the E5500 controller with specific iSCSI host interface cards. The report is accessible by selecting Monitor > Health > iSCSI Statistics from the SANtricity Storage Manager Array Management Window.
	Microsoft Offloaded Data Transfer (ODX)	You can turn ODX on or off through the CLI. ODX provides a way to transfer data without using buffered read and buffered write operations and does not require direct host involvement with the data transfer operation. ODX is available only on the Windows Server 2012 operating system.

E-Series and EF-Series hardware

Features introduced in previous releases of SANtricity® OS 8.20.

Model Series	What's New
E2624, E2660, E2724, E2760, E5424, E5460, E5524, E5560, E5624, E5660, DE2624, DE2660, DE2724, DE2760, DE5424, DE5460, DE5524, DE5560, DE5624, DE5660	Added support for the HGST Cobra-F 10K 2.5-inch hard disk drive
E2600	Platform support for cabling a mixture of drive tray types (DE5600, DE1600, DE6600) to the controller- drive tray.
E2700	Enhancements to add either a Dual or a Quad Port 10-Gb/s iSCSI HIC.
E5500 and EF5500	- Enhancements to support a Quad Port 12-Gb/s SAS-3 HIC and a Dual Port 56-Gb/s IB HIC - Performance enhancements
E5600 and EF560	New E5600/EF560 platform for increased performance.

Reserved drive space

A portion of each drive in a storage array is reserved for storing configuration information, along with other types of information such as error logs and SANtricity OS (controller firmware) images. This reserved space, referred to as the “DACstore,” reduces the storage capacity available for creating volumes in volume groups and disk pools. Prior to the release of SANtricity OS 8.20, 500 MB of space was reserved on each drive. Beginning with SANtricity OS 8.20, 5.5 GB of space is reserved on each drive. Consequently, the capacity available for storing user data is decreased by 5

GB for each drive in each volume group or disk pool. The additional reserve capacity accommodates the storage of increased amounts of information.

Note: The DACstore extent is increased to 5.5 GB only for drives in volume groups and disk pools. The DACstore extent continues to be 500 MB for drives used for SSD Cache.

The new 5.5 GB DACstore extent (5.5 x 1024 x 1024 x 1024 bytes) is called a "large DACstore extent." The 500 MB DACstore extent (500 x 1024 x 1024 bytes) that was used in previous releases is called a "legacy DACstore extent." The new large DACstore extent applies to all platforms supported by SANtricity 8.20.

- E2600
- E2700
- E5400/EF540
- E5500/EF550
- E5600/EF560 (only if an existing system is upgraded to E5600/EF560 controllers)

The DACstore extent decreases the usable capacity of each drive in the storage array. The usable capacity of each drive is reported in the Storage Array Profile and is used by SANtricity Storage Manager to calculate the total unconfigured capacity reported for the storage array. After a volume group or disk pool is created, the usable capacity from the member drives is used to report the free capacity for the volume group or disk pool. In SANtricity OS 8.20, the usable capacity for unassigned drives is always 500 MB less than the raw capacity reported in the Storage Array Profile for each unassigned drive.

The increased reserve capacity is not applied until a drive is included in a volume group or disk pool. When a volume group or disk pool is created, the usable capacity for each member drive decreases to 5.5 GB less than its raw capacity. So, the free capacity for the volume group will be less than the total of the usable capacity of the unassigned drives. The usable capacity of drives in volume groups and disk pools that were created with pre-8.20 SANtricity OS will be 500MB less than the raw capacity reported for each drive.

If you upgrade from a previous (pre-8.20) release of SANtricity OS, the capacity of any existing volume groups and disk pools is not be reduced. The DACstore extent continues to be 500 MB on drives in volume groups and disk pools that were present in the storage array prior to the upgrade. The DACstore extent is 5.5 GB on drives in volume groups or disk pools that are created after the upgrade occurs. So, it is possible to have a storage array where the DACstore extent is 500 MB for drives in some volume groups and disk pools and 5.5 GB for drives in other volume groups or disk pools. If an existing volume group or disk pool is deleted after the upgrade occurs, the DACstore extent is increased to 5.5 GB when those drives are subsequently added to a new volume group or disk pool.

Media scan with redundancy check enabled by default

On new systems with SANtricity OS 8.20.08.00 media scans are performed using a duration of 30 days. If the NVSRAM file associated with this release is downloaded to an existing system, media scan with redundancy check are enabled by default for newly created volumes. However, if media scans are currently suspended (the default behavior in previous releases), they continue to be suspended unless the **suspend** option is unchecked in the **change media scan settings** option in the SANtricity Storage Manager Array Management window.

For existing systems that upgrade to SANtricity OS 8.20.08.00, the **Change Media Scan Settings** dialog box indicates that Media Scan with Redundancy Check is enabled for newly created volumes, but the media scan does not actually occur unless you have previously unchecked the **suspend** check box. The duration setting for Media Scan functions like a global enable/disable. If the duration is zero, no media scans are performed even though they might be enabled for some volumes. If the **suspend** box is checked, the settings for all volumes are grayed out. If the **suspend**

box was checked and you subsequently uncheck the box, the scan duration will default to 15 days.

Resolved issues and software dependencies

Information in this document is periodically updated between major software and firmware releases. For the most up-to-date information about resolved issues and software dependencies, visit the [NetApp Downloads: Software](#). Log in to the support site, and then look in the list of products for "E- Series SANtricity Storage Manager" or "E-Series SANtricity OS Software." For each product, select the applicable platform from the drop-down list and then click the **GO** button to download the most recent software or firmware updates. The files that you download will include a README file, which includes a complete list of resolved issues for this release. The most essential resolved issues are detailed below.

Important issues resolved in SANtricity OS 8.20 maintenance releases

- **Issues resolved in SANtricity OS 8.20.20.00**
 - Fixed multiple issues with Misbehaving Drive handling, especially in regards to DDP volumes.
 - Fixed critical issues with Online Drive Firmware Upgrade (ODFU). This release should be loaded before attempting any future ODFU operations.
 - Fixed a critical issue with the interaction between E-Series storage systems and Brocade switches that have RDP enabled.
- **Issues resolved in SANtricity OS 8.20.12.00**
 - The following critical security vulnerability (CVE-2015-7599) in the VxWorks portion of SANtricity OS is resolved:
VxWorks RPC (remote procedure call) security vulnerability - A security vulnerability in RPC (remote procedure call) enables a maliciously crafted packet to cause an integer overflow with the possibility of executing remote code. For more detail, refer to [Customer Support Bulletin 9010045](#).
 - In a Windows environment where the Offloaded Data Transfer (ODX) feature is being used to write zeroes to volumes (e.g., to initialize a Hyper-V virtual machine), under certain conditions zeroes can be written incorrectly. The controller may detect this, resulting in both controllers becoming unresponsive. There is no indication to the user that the problem has occurred other than the unresponsive controllers. Note that no other ODX operations are affected by this issue. For more detail, refer to [Customer Support Bulletin 7010218](#).
- **Issues resolved in SANtricity OS 8.20.09.00**
 - For SANtricity OS 8.20.09.00, an issue was resolved where, under certain circumstances, in a Windows clustered environment, Offloaded Data Transfer (ODX) returned a "good" status although the transfer had not completed successfully. For more detail, [Customer Support Bulletin 7010178](#).

Restrictions

With Veritas, online controller firmware upgrade causes loss of both paths

Severity: 2 Functional

Operating System

- RedHat Enterprise Linux RHEL 6.7
- Veritas Storage Foundation 6.0.5 with ASL 6.0.500.610
- DMP failover 6.0.500.300

Hardware/Software/Firmware

All

Problem or Restriction

If you perform an online controller firmware upgrade to NetApp E-Series controllers, that is, with host I/O occurring, after a controller reboots as part of the upgrade process, the paths from the host to the storage may not come back online immediately. When the alternate controller reboots to activate the new firmware, the host may experience a brief all paths down to the storage, causing any subsequent IO to fail. The paths come back online within a few minutes.

Workaround

Halt host I/O before performing a controller firmware upgrade.

Reference Number

200880337

Asynchronous Mirroring volumes are failed during a volume reconfig or Unmap operation if a volume transfer occurs

Severity: 2 Functional

Operating System

All

Hardware/Software/Firmware

All

Problem or Restriction

If a volume ownership transfer between controllers occurs while a secondary volume in an Asynchronous Mirroring pair is undergoing a reconfiguration (volume expansion, RAID migration, etc.) or an Unmap operation, the secondary volume transitions to a failed state.

Workaround

To reduce the likelihood of this issue, disable Unmap while performing a reconfiguration operation. To recover from this issue, resynchronize the mirrored pair

Reference Number

200859174

Oracle Java SE Runtime versions later than 8u31 but prior to 8u65 cause random GUI issues

Severity: 2 Functional

Operating System

All

Hardware/Software/Firmware

- SANtricity Storage Manager version 11.20.XX00.0015
- SANtricity Storage Manager version 11.20.XX00.0017

Problem or Restriction

An issue with the JVM's JNI-compilation in JRE builds later than 8u31 but prior to 8u65 causes random problems in the SANtricity Storage Manager GUI. In some cases, the GUI stops responding to input or the entire application crashes. SANtricity version 11.20.XX00.0015 and 11.20.XX00.0017 include JRE 8u31, which does not have this problem. Earlier SANtricity Storage Manager versions are not affected, as they include JRE 7. A ticket has been opened with Oracle for resolution of this problem.

Workaround

A fix has been provided in JRE version 8u65. Upgrade your JRE to version 8u65 or later. For more information on how to upgrade JRE manually, refer to [Updating the Java Runtime Environment \(JRE\) used by SANtricity Storage Manager 11.10 or later](#).

Reference Number

200770704

Oracle Java SE Runtime versions prior to 8u65 cause random GUI issues on hosts running Solaris

Severity: 2 Functional

Operating System

Solaris

Hardware/Software/Firmware

- SANtricity Storage Manager version 11.20.XX00.0015
- SANtricity Storage Manager version 11.20.XX00.0017

Problem or Restriction

An issue with the JVM's JNI-compilation in JRE builds prior to 8u65 cause random problems in the SANtricity Storage Manager GUI when run on a host with the Solaris operating system. In some cases, the GUI stops responding to input or the entire application crashes. SANtricity versions 11.20.XX00.0015 and 11.20.XX00.017 include JRE 8u31. With other supported operating systems, these problems do not occur with JRE 8u31 (see [reference number 20770704](#)). However, with Solaris,

the problems occur with JRE 8u31 as well as with later versions of JRE. Earlier SANtricity Storage Manager versions are not affected, as they include JRE 7. A ticket has been opened with Oracle for resolution of this problem.

Workaround

If the SANtricity Storage Manager stops responding, close and then reopen the application. A fix has been provided in JRE version 8u65. Upgrade your JRE to version 8u65 or later. For more information on how to upgrade the JRE manually, refer to the [Updating the Java Runtime Environment \(JRE\) used by SANtricity Storage Manager 11.10 or later](#).

Reference Number

200774195

Volumes might not stay on the preferred path during an online controller firmware upgrade

Severity: 2 Functional

Operating System

All

Hardware/Software/Firmware

- E2600 controller-drive trays

Problem or Restriction

In rare circumstances, a volume might transfer to the alternate path during a controller firmware upgrade. If this happens, the storage array is marked **Degraded** and the Recovery Guru in SANtricity Storage Manager reports **Volume not on preferred path**.

Workaround

After completing the controller firmware upgrade process, if the Recovery Guru reports **Volume Not on Preferred Path**, check the volume ownership and redistribute volumes, if necessary. You can redistribute volumes in the SANtricity Storage Manager by selecting **Storage > Volume > Advanced > Redistribute Volumes**.

Reference Number

200642520

Choosing a custom segment size

Severity: 2 Functional

Operating System

All

Hardware/Software/Firmware

- All controllers

Problem or Restriction

When you create a volume, the segment size is automatically chosen based upon your selection of the volume I/O characteristics type (File system, Database, Multimedia, Custom). Large segment sizes (256 KB and larger) should only be chosen if you know that the I/O profile for that volume will be large, sequential reads and writes.

Workaround

If you anticipate an I/O profile with frequent non-sequential single-block writes, choose a smaller (for example 32 KB) segment size for your volume configuration.

Reference Number

200643815

A path failure occurs during a controller firmware upgrade

Severity: 2 Functional

Operating System

Red Hat Enterprise Linux 7

Hardware/Software/Firmware

- SAS or IB HICs
- All controllers

Problem or Restriction

During a controller firmware upgrade, I/O failures occur when the `multipathd` fails to add new paths into the device map after the first controller activation. I/O failure occurs when the second controller reboots during activation. Bugzilla # 1015601 has been created with Red Hat.

Workaround

You can avoid this issue by either halting I/O operations prior to performing the upgrade or enabling the `queue_if_no_path` feature by setting the `no_path_retry` parameter to `queue` in `multipath.conf`. If you make changes to the `queue_if_no_path` setting, make sure to reload `multipathd` service before you initiate the controller firmware upgrade process. With these settings, I/O's will be held up in `multipath` until all paths are restored during the upgrade. You can then run the `rescan-scsi-bus.sh` command after the upgrade to see that all paths are restored correctly.

After the upgrade, you can restore both settings; that is, disable the `queue_if_no_path` feature by command (`dmsetup message mpath-device-name 0 fail_if_no_path`), set the `no_path_retry` parameter to `30` in the `multipath.conf` file and reload `multipathd` service. This workaround will only help if application timeouts are much higher than controller firmware upgrade times. If application timeouts are lower, then applications will error out.

If this issue occurs, you can recover by running the `rescan-scsi-bus.sh` command after the upgrade to confirm that paths are restored correctly. If needed, run the `service multipathd reload` command to allow fresh creation of device maps. Restart I/O.

Reference Number

200445426

The presence of host I/O during management software upgrade prevents the E-Series Windows DSM from updating properly

Severity: 2 Functional

Operating System

Windows

Hardware/Software/Firmware

- All E-Series Storage Arrays

Problem or Restriction

This issue only occurs when the Windows DSM for Windows MPIO is upgraded as part of the SANtricity Storage Manager upgrade process. During an update of the DSM for Windows MPIO, the driver stacks are torn down and rebuilt. The presence of host I/O during the upgrade can extend the time needed to tear down and rebuild the stacks, thereby extending the time to complete the upgrade operation beyond the time allowed by the Windows system API. This results in a timeout and failure of the DSM to complete the update operation.

Workaround

To avoid the issue, halt host I/O during the SANtricity Storage Manager upgrade process.

To recover from the issue, if the upgrade process fails, reboot the host and verify that the correct DSM version is installed. If the correct version is not installed, re-execute the appropriate upgrade process.

Reference Number

200493844

A host panics during path failover with an Emulex Fibre Channel (FC) host bus adapter (HBA)

Severity: 2 Functional**Operating System**

Red Hat Enterprise Linux 6.5 Kernel 2.6.32-431.el6.x86_64

Hardware/Software/Firmware

- Emulex Fibre Channel HBA 1600x with the inbox driver 8.3.7.21.4p

Problem or Restriction

A path failover due to a controller reboot might cause the HBA to return an error of unsafe-linked-list and might ultimately cause the server to panic and reboot. Bug 1071656 has been submitted to Red Hat.

Workaround

The host reboot will resolve the issue and the host will return to optimal mode.

Reference Number

200509398

Emulex Fibre Channel (FC) host bus adapter (HBA) link errors result in a kernel panic

Severity: 2 Functional**Operating System**

Red Hat Enterprise Linux 6.5 Kernel 2.6.32-431.el6.x86_64

Hardware/Software/Firmware

- All Fibre Channel systems
- Emulex HBA 1600x with the in-box driver 8.3.7.21.4p.

Problem or Restriction

The Emulex driver causes a RHEL OS kernel panic because the Emulex FC HBA hardware produces link errors.

Workaround

None

Several patches for Emulex have been submitted through Red Hat BugZilla 1054369.

Reference Number

200513972

When the `user_friendly_names` parameter is enabled in `/etc/multipath.conf`, paths are mapped incorrectly

Severity: 2 Functional**Operating System**

Red Hat Enterprise Linux 7

Hardware/Software/Firmware

- Device Mapper Multipath with `user_friendly_names` enabled.
- SAN boot configuration

Problem or Restriction

A condition can occur where one user-friendly name is given to a disk during boot, but multipathd tries to allocate a different name after boot. This might occur, for example, after a path failure. If the second alias is already used by another device, multipathd cannot rename it. Multipathd then has incorrect information about the alias and wwid relationships, which can result in paths being added to the wrong map. If the problem does occur, paths from two volumes would be mixed into one multipath device, potentially leading to data corruption. Bug 1086825 has been submitted to Red Hat.

Workaround

Use one of the following methods to prevent the problem:

- Disable `user_friendly_names` completely.
- Edit `/etc/multipath.conf` and the `multipath { }` section definitions to manually define aliases: See https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/DM_Multipath/config_file_multipath.html for instructions

Reference Number

200517798

A Red Hat Enterprise Linux (RHEL) kernel bug causes a soft lockup under high memory load due to swap thrashing

Severity: 2 Functional**Operating System**

Red Hat Enterprise Linux 6.4 kernel 2.6.32-358.2.1.el6

Hardware/Software/Firmware

- All platforms

Problem or Restriction

This known issue occurs on Red Hat 6.4 (kernel 2.6.32-358.2.1.el6) server under high memory load conditions. It produces the error message `kernel:BUG: soft lockup`. The following conditions might cause the server to operate under high memory load:

1. A high number of KVM virtual machines which are using vhost networking.
2. A system is under high memory pressure (for example, several GBs of swap are in use).
3. A system is under high network load.

According to Red Hat, “there is a bug in `shrink_zone()` which can enter an indefinite loop when attempting to reclaim memory for a network buffer allocation.” Bugzilla number 9851155 has been opened with Red Hat. This problem is resolved in errata kernel 2.6.32-358.32.3.el6 and later versions.

Workaround

Reduce memory pressure by adding RAM to the host system, reduce the number of VMs, or reduce the number of tasks that consume memory running on the system.

Reference Number

200526690

I/O error occurs on SUSE Linux Enterprise Server (SLES) VMs with BusLogic SCSI controllers after a controller reboot

Severity: 2 Functional**Operating System**

VMWare ESXi 5.5 U1 with SLES 10 and SLES 11 VMs

Hardware/Software/Firmware

- BusLogic virtual SCSI adapters

Problem or Restriction

When a controller reboot occurs, the underlying BusLogic adapter initializes as the pRDM device resets. During this time, pRDMs on BusLogic-enabled devices may take a long time to come back on line and may be left in a device offlined state. Messages such as `Device offlined - not ready after error recovery` or `Result: hostbyte=DID_OK driverbyte=DRIVER_TIMEOUT` appear in the VM's log file for the specific pRDM device. This issue results in loss of data access on specific volumes owned by the rebooting controller until the workaround is applied or issue is fixed.

Workaround

Use an LSI Logic virtual SCSI adapter as recommended by VMware. Refer to the following web

sites for additional information:

http://partnerweb.vmware.com/GOSIG/SLES_10.html

http://partnerweb.vmware.com/GOSIG/SLES_11.html

Reference Number

200527398

A host with Red Hat Enterprise Linux (RHEL) 7 crashes during a controller firmware upgrade

Severity: 2 Functional

Operating System

Red Hat Enterprise Linux 7.0 PC64 or Kernel 3.10.0-123.el7.x86_64

Hardware/Software/Firmware

- Host Bus Adapter (HBA) Model 1: 10N9824, IBM variant LPe12K (with inbox driver version 8.3.7.34.3p)
- HBA Model 2: 74Y2221, IBM variant LPe16K (with inbox driver version 8.3.7.34.3p)
- HBA Model 3: Emulex LPe1200x (with inbox driver version 8.3.7.34.3p)

Problem or Restriction

A host running RHEL 7 crashed during an online controller firmware upgrade. The problem with the in-box Emulex driver occurs at the driver level for the HBAs listed above. You lose access to data for the duration of the upgrade.

Workaround

Stop I/O activity between the RHEL 7 host and the storage array before you perform a controller firmware upgrade.

Red Hat 7.0 z stream (Errata) kernel will have a fix for this issue according to Red Hat Bugzilla items 1081718, 1103881, and 1122575.

Reference Number

200528430

Cluster resources cannot be transferred when using network file system (NFS) version 4

Severity: 2 Functional

Operating System

Red Hat Enterprise Linux (RHEL) 7.0

NFS version 4

Hardware/Software/Firmware

- All controllers

Problem or Restriction

If resources are manually transferred, the resources cannot be transferred to the expected cluster node. A reboot of any node in the cluster should force the resources to be transferred to the other node while I/O operations are running through a remote client. However, the node reboot process hangs because the mount points on the rebooting node fail to be unmounted. This causes the service to go to the **failed** state.

Workaround

To avoid the issue, use NFS version 3 instead of NFS version 4.

To recover from the issue, stop cluster services running on all nodes, then restart the cluster. You lose redundancy at the server level while the services are stopped.

Reference Number

200533838

I/O errors occur during an online controller firmware upgrade with I/O activity

Severity: 2 Functional

Operating System

Windows

Hardware/Software/Firmware

- All controllers

Problem or Restriction

Rarely, during a controller firmware upgrade with I/O operations running, a Microsoft Windows server detects that both paths to its volumes have been lost for some time interval because of the way controllers reboot during the upgrade. The length of the interval when both paths are lost depends on the number of volumes and on the I/O load. I/O errors might occur if the paths to both controllers are unavailable for more than 150 seconds.

Workaround

If I/O errors are encountered, initiate a host-side device discovery action to re-discover the array and re-establish access to the volumes. You will lose access to data until the workaround is completed. To avoid this problem, perform controller firmware upgrades during planned maintenance periods.

Reference Number

200545493

The `multipath-tools` command fails to add the physical path to device-mapper multipath map after a failed path is restored

Severity: 2 Functional

Operating System

SUSE Linux Enterprise System (SLES) 11 SP3 with the base in-box multipath-tools-0.4.9-0.83.2

Hardware/Software/Firmware

- All Controllers

Problem or Restriction

Very rarely, after a path failure, when the connection is restored, the `multipath-tools` command fails to add the physical path to device-mapper multipath map. This results in the path remaining an orphan path that is not used for I/O operations.

In host logs the following error messages appear repeatedly.

```
Apr 17 09:17:10 terrier multipathd: sdax: failed to get path uid Apr 17
09:17:11 terrier multipathd: sdax: add missing path
Apr 17 09:17:12 terrier multipathd: sdax: failed to get path uid Apr 17
09:17:13 terrier multipathd: sdax: add missing path
```

Bug 893664 has been submitted to Novell.

Workaround

To avoid this issue, in the device section for "LSI|NETAPP" devices, provide the following `getuid_callout` parameter.

```
getuid_callout "/lib/udev/scsi_id --whitelisted --device=/dev/%n"
```

After this, run `multipathd reconfigure` to apply the new configuration.

To recover from this issue, immediately recover the path that is unusable while I/Os are still flowing through the remaining path(s). Follow any one of the following recovery options:

- Add the `getuid_callout` parameter to the devices section of "LSI|NETAPP" devices and run the `multipathd reconfigure` command or the service `multipathd reload` command to recover paths.
- Run the `echo add > /sys/block/sd{x}/uevent` command for all those `sd{x}` devices that appear as Orphan in the output from the `multipathd show paths` command.

Reference Number

200545556

Inconsistency in volume ownership transfer occurs during a parallel drive firmware update (PDFU)

Severity: 2 Functional

Operating System

All

Hardware/Software/Firmware

- All controllers

Problem or Restriction

During a PDFU, you can transfer all volumes in a volume group or DDP, but you cannot transfer a single volume.

Workaround

Wait until PDFU completes before attempting a single volume ownership transfer.

Reference Number

200554087

I/O errors occur during online controller firmware update

Severity: 2 Functional

Operating System

Red Hat Linux 6.5 and, possibly, other Red Hat versions

Hardware/Software/Firmware

- All controllers
- Device Mapper Multi-Path failover configurations

Problem or Restriction

Rarely, during an online controller firmware upgrade with a large number of volumes (for example, 256 volumes per host), random file system(s) might get unmounted resulting in I/O errors. This problem occurs because `udisks-daemon` might automatically unmount the file systems during controller firmware upgrades. You lose access to data until the workaround is applied. Bugzilla number 1103362 has been opened with Red Hat.

Workaround

If the I/O errors happen on a volume that is not used for SAN boot, unmount the read-only file system and remount it as a read/write file system. If the problem occurs on a SAN Boot volume, reboot the host.

Reference Number

200564542

I/O errors occur during online controller firmware update

Severity: 2 Functional

Operating System

Red Hat Linux operating systems

Hardware/Software/Firmware

- All controllers

Problem or Restriction

During an online controller firmware upgrade with a large number of volumes (for example, 256 volumes per host), random file system(s) may get unmounted resulting in I/O errors. This problem occurs because `udisks-daemon` automatically unmounts the file systems during controller firmware upgrades. Bugzilla number 1103362 has been opened with Red Hat.

Workaround

If an application error occurs or root filesystem runs out of space or is mounted as read-only, you can

perform following recovery options:

- If root device is not on SAN, run the command `#remount -o remount,ro /dev/{device_name} /{mount_point}`
- If the boot device is on the storage array, reboot the host.

After you can access the root file system again, identify any new files that were created under the mount directory while actual file system was unmounted. Based on the application needs, restore these files to a temporary location before the actual file system is mounted again at the same location. This enables you to free up space on root file system and to restore the application data to original file system device again. Uninstall the `udisks` package from the system.

Reference Number

200567225

A suspect drive is marked as failed unexpectedly

Severity: 2 Functional

Operating System

All

Hardware/Software/Firmware

- All controllers

Problem or Restriction

Rarely, if a controller initiates a forced volume transfer while a drive is powered down, the drive might be unexpectedly marked as failed. This will only happen during forced volume transfers where the owning controller fails or is removed during the power cycle. Host requested volume transfers do not cause this issue because I/O is quiesced before transferring the volume. You might see a degraded or failed volume depending on data protection scheme and redundancy level at that point in time.

Workaround

Replace the failed drive.

Reference Number

200577809

An early path failover occurs when a host cable is pulled

Severity: 3 Functional

Operating System

SUSE Linux Enterprise Server (SLES) 12.0

Hardware/Software/Firmware

- All Controllers
- Device Mapper Multipath `scsi_dh_rdac` path checker `mpt2sas` driver version 18.100.00.00

Problem or Restriction

This problem occurs when a path fails to a host running SLES 12, with a SAS host adapter, the `mpt2sas` driver version 18.100.00.00 and the `scsi_dh_rdac` driver. The host immediately fails SCSI inquiry commands under certain conditions. The `scsi_dh_rdac` driver does not retry when `DID_SOFT_ERROR` is returned from a SCSI inquiry. When an inquiry fails in this way, and does not retry, Device Mapper Multipath (DM-MP) will fail the path. The return status can cause DM-MP to fail the path after a cable is pulled, even if it is replaced before the failover trigger.

Workaround

Reconnect the cable, if necessary. The path will then recover normally.

Reference Number

200579172

I/O operations fail when SG_IO is used with a block size larger than 512 KB

Severity: 2 Functional

Operating System

- Red Hat Enterprise Linux 6.5 or prior releases
- SuSE Enterprise Linux 11.3 or prior releases
- Mellanox OFED 2.2 or prior releases
- iSER initiator with an InfiniBand HCA

Hardware/Software/Firmware

- 56-Gb/s InfiniBand host interface card

Problem or Restriction

Block devices that are presented through iSCSI Extensions for RDMA (iSER) connections will not accept I/O block sizes greater than 512 KB. Under normal conditions, where I/O goes through the linux block layer, large I/O requests will be split into smaller sizes without any errors occurring. However, an I/O request that does not go through the block layer, like `SG_IO ioctl`, that tries to send block I/O larger than the value of the `max_sectors_kb` parameter (512 KB for the identified iSER drivers) will experience an error.

Workaround

Issue I/O requests with a size less than or equal to 512 KB when using `SG_IO`.

Reference Number

200588987

Random read IOPs for cache and solid state disk (SSD) performs below anticipated targets with the data assurance feature

Severity: 2 Functional

Operating System

All

Hardware/Software/Firmware

- E5500 and EF550
- 56-Gb/s InfiniBand host connections

Problem or Restriction

This issue applies to both 4K random reads and 512b random cache reads, because both are bound by the CPU overhead of dealing with data assurance (PI).

Workaround

None.

Reference Number

200601026

An iSCSI session does not persist after a host reboot

Severity: 3 Minor

Operating System

- Red Hat Enterprise Virtualization (RHEV) Hypervisor - 6.5 - 20140725.0.el6ev
- RHEV Hypervisor - 6.4 - 20131016.0.el6
- RHEV Hypervisor - 6.3 - 20130129.0.el6_3 with RHEV-M 3.4.1

Hardware/Software/Firmware

- RHEL software iSCSI initiator
- All controllers

Problem or Restriction

This problem occurs on an RHEV-H host with the software iSCSI initiator running and four sessions (four paths, two to each controller). Sessions are lost after a host reboot. When the host completes a reboot, it comes up with only two new sessions, one to each controller. RedHat bugzilla 1132731 has been filed for this issue.

Workaround

After the host completes the reboot, use the `iscsiadm` tool to bring up all of the lost sessions manually.

Reference Number

200601223

System fails to boot from a local disk with enabled `multipathd` and the root device not blacklisted

Severity: 2 Functional

Operating System

SUSE Linux Enterprise Server (SLES)12

Hardware/Software/Firmware

- All controllers
- Device Mapper MultiPath

Problem or Restriction

This problem occurs only on local disk installations and not on SAN boot systems, and only when the local root disk is not blacklisted in the `/etc/multipath.conf` file after installation. The system enters into “Emergency Mode” on reboot after enabling `multipathd` service on the host. Bug 890284 has been submitted to SuSE.

This is a problem with `multipath` integration with `systemd`. During boot time, `dracut` mounts the physical path (for example, `/dev/sda2` etc) as the root device. But the `udev` process will think that the `multipath` device will be created on top of it, as its not blacklisted, and defer creation of `/dev/disk/by-uuid` symbolic links to sd devices. Because the `/dev/sda2` device is already mounted, `multipathd` fails to create a multipath device (for example, `/dev/dm-x`) for the local disk and `local-fs.target` dependencies will fail waiting on the symlinks `/dev/disk/by-uuid` to `/dev/sd` devices. The following message appears during boot-up on the system screen or console.

```
Timed out waiting for device dev-disk-by\x2duuid-c4a...cfef77d.device.
[DEPEND] Dependency failed for /opt. [DEPEND] Dependency failed for
Local File Systems. [DEPEND] Dependency failed for Postfix Mail
Transport Agent. Welcome to emergency shell Give root password for
maintenance (or press Control-D to continue):
```

Workaround

To prevent the system from entering into “Emergency Mode”, perform the following steps before host reboot but after installation.

1. Blacklist the root device in `/etc/multipath.conf`.
2. Run the `multipath -v2` command. The system output looks like the following:

```
Jul 29 10:10:03 | 3600508b1001030343841423043300400: ignoring map
```

3. Verify that the string between `'|'` and `':'` is the `wwid` of the root device, which should be added to `/etc/multipath.conf` as

```
blacklist{
  wwid
  "3600508b1001030343841423043300400"
}
```

With this setting, the system will not try to create a multipath device for the local root device.

4. Run `systemctl reboot` to reboot the system.

Reference Number

200601561

Java access with speech (JAWS) reads hidden table data in SANtricity Storage Manager

Severity: 3 Minor**Operating System**

Windows

Hardware/Software/Firmware

- Freedom Scientific JAWS
- All controllers

Problem or Restriction

SANtricity Storage Manager sometimes uses tables in multiple locations. In some cases, certain columns of data are not supposed to be displayed to the user; however, JAWS reads this data anyway.

Workaround

None.

Reference Number

200602538

iSCSI SAN boot failure

Severity: 3 Minor**Operating System**

Red Hat Enterprise Linux (RHEL) 7.0

Hardware/Software/Firmware

- Cisco Unified Computing System (UCS)
- Virtual Interface Card (VIC) 1240/1280
- All Controllers

Problem or Restriction

The IP parameter and Subnet Address are not properly set in the `/etc/grub2.cfg` file. This causes iSCSI SANboot to fail because no IPs are assigned and no iSCSI sessions can be created.

RHEL 7 does not create a proper `/etc/grub2.cfg` file after iSCSI installation. Red Hat Bugzilla link: https://bugzilla.redhat.com/show_bug.cgi?id=1114464.

Workaround

To avoid this issues, modify the `/etc/grub2.cfg` file manually during the boot process. Press **e** when prompted to edit the grub file and add the necessary IPs to the desired ports. This is done by going to the kernel line and modifying the entries for `ip=: : :HBA:port:none` to `ip=10.10.10.127:::255.255.255.0:HBA:port:none`.

Below is an example of the grub file that will be given to you, and that same file modified to add the necessary IPs.

Unmodified Kernel Line:

```
linux16 /vmlinuz-3.10.0-123.el7.x86_64 root=UUID=2bc9e0c8-31b7-4c74-b755-
ba8f645ce741 ro ifname=ens1f0:00:00:c9:b0:47:98 vconsole.keymap=us
ip=:::IBMx3550-210-127:ens1f0:none ip=:::IBMx3550-210-127:ens1f1:none
ifname=ens1f1:00:00:c9:b0:47:9a netroot=iscsi:@10.10.10.10::3260::iqn.
1992-08.com.netapp:sn.5d35f5e7ed9711e3ba53123478563412:vs.10
netroot=iscsi:@11.11.11.11::3260::iqn.1992-08.com.netapp:sn.
5d35f5e7ed9711e3ba53123478563412:vs.10 crashkernel=auto
rd.lvm.lv=rhel_210-127/root rd.lvm.lv=rhel_210-127/swap
netroot=iscsi:@11.11.11.10::3260::iqn.1992-08.com.netapp:sn.
5d35f5e7ed9711e3ba53123478563412:vs.10 iscsi_initiator=iqn.
1994-05.com.redhat:127 netroot=iscsi:@10.10.10.11::3260::iqn.
1992-08.com.netapp:sn.5d35f5e7ed9711e3ba53123478563412:vs.10
vconsole.font=latarcyrheb-sun16 rhgb quiet LANG=en_US.UTF-8
```

Modified Kernel Line:

```
linux16 /vmlinuz-3.10.0-123.el7.x86_64 root=UUID=2bc9e0c8-31b7-4c74-b755-
ba8f645ce741 ro ifname=ens1f0:00:00:c9:b0:47:98 vconsole.keymap=us
ip=10.10.10.127::255.255.255.0:IBMx3550-210-127:ens1f0:none
ip=11.11.11.127::255.255.255.0:IBMx3550-210-127:ens1f1:none
ifname=ens1f1:00:00:c9:b0:47:9a netroot=iscsi:@10.10.10.10::3260::iqn.
1992-08.com.netapp:sn.5d35f5e7ed9711e3ba53123478563412:vs.10
netroot=iscsi:@11.11.11.11::3260::iqn.1992-08.com.netapp:sn.
5d35f5e7ed9711e3ba53123478563412:vs.10 crashkernel=auto
rd.lvm.lv=rhel_210-127/root rd.lvm.lv=rhel_210-127/swap
netroot=iscsi:@11.11.11.10::3260::iqn.1992-08.com.netapp:sn.
5d35f5e7ed9711e3ba53123478563412:vs.10 iscsi_initiator=iqn.
1994-05.com.redhat:127 netroot=iscsi:@10.10.10.11::3260::iqn.
1992-08.com.netapp:sn.5d35f5e7ed9711e3ba53123478563412:vs.10
vconsole.font=latarcyrheb-sun16 rhgb quiet LANG=en_US.UTF-8
```

Reference Number

200602974

While performing a storage VMotion operation, virtual machines (VMs) occasionally fail to transfer properly when under I/O load

Severity: 3 Minor

Operating System

VMware ESXi 5.5, storage vMotion

Hardware/Software/Firmware

- All Controllers

Problem or Restriction

This VMware problem occurs when a host that currently owns resources fails to remove its lock on those resources, causing other hosts to fail to gain ownership. As a result, the virtual machine (VM) might not power up. The issue is seen when a storage vMotion operation is used to migrate a vmdk disk file to a different datastore. It occurs when the operation is initiated manually by the user and when it is initiated automatically by the storage distributed resource scheduler (SDRS). The issue does not occur every time a vmdk disk is migrated, but does occur fairly often. SR 14506568307 has been filed with VMware to further track this issue. The root cause appears to be the same as SR 13383440910.

Workaround

To avoid this issue, either disable the SDRS feature for datastore clusters or set the SDRS automation level to No Automation to prevent the SDRS from automatically performing vMotion operations. To recover from this issue, power the VM back on. If the VM fails to power on because it cannot find a vmdk, backup and then download the VM's vmx configuration file. Then, modify the entry for the vmdk file path with the correct location. Upload the file to the datastore and power the VM back on.

Reference Number

200603079

File systems are unmounted during a controller reboot

Severity: 2 Functional**Operating System**

Red Hat Enterprise Linux (RHEL) Operating Systems

Hardware/Software/Firmware

- All Controllers

Problem or Restriction

If a controller exception causes a controller to reboot, random file system(s) may get unmounted, resulting in I/O errors. The root cause of the problem is the RHEL `udisks-daemon`, which automatically unmounts the file systems during a controller reboot. This is the same issue as 200567225. Red Hat Bugzilla 1103282 has been opened to investigate this.

Workaround

To avoid this issue, uninstall the `udisks` package from the system. Red Hat suggests removing this package, as it was designed only for desktop systems.

1. Determine which `udisks` package is installed on your system by running the command: `#rpm -qa | grep udisks`
2. Uninstall the `udisks` package from your system by running the command: `#rpm -e <udisks package name>`

To recover from this issue if an application error occurs or the root file system goes out of space or is mounted read-only, perform the following steps:

1. If root device is not on SAN, run the command: `#rpm -qa | grep udisks`
2. If root device is on SAN, reboot the system.
3. Once you are able to access the root file system again, identify any new files created under the mount directory while the actual file system was unmounted. Based on the application needs, restore the files to a temporary location before the actual file system is mounted again at the same location. This enables you to free up space on root file system and restore the application data to the original file system device.
4. Uninstall the `udisks` package from the system by running the command: `#rpm -e <udisks package name>`

Reference Number

200603904

Cannot use dsmUtil.exe to apply load balance policy setting on a per-device basis

Severity: 2 Functional

Operating System

Windows

Hardware/Software/Firmware

- All controllers

Problem or Restriction

The Device Specific Module (DSM) 2.0 data model supports a number of configurable settings, which are managed via the `dsmUtil` utility. Some of these settings are global, meaning that they apply to all storage managed by the DSM, while others can be applied on a per-device basis. Currently, the `dsmUtil` supports global settings only; it does not let you change settings for an individual device. Specifically, **load balance policy**, which is initially applied as a global setting, can only be changed on a per-device basis. To change the load balance policy, you must use an interface that supports per-device configuration, such as the **MPIO** tab, which is available for any MPIO disk device that appears in disk manager.

Workaround

Use the **MPIO** tab for the MPIO disk device to apply a new load balance policy setting.

Reference Number

200626230

After a controller tray power failure, some Seagate drives in expansion trays that did not lose power are unresponsive

Severity: 2 Functional

Operating System

All

Hardware/Software/Firmware

- Drive model ST900MM0006 with drive firmware MS02
- Drive model ST91000641SS with drive firmware MS04
- Drive model ST450MM0006 with drive firmware MS02
- Drive model ST9600105SS with drive firmware MS07

Problem or Restriction

Certain Seagate drives become unresponsive in the expansion trays when the controller enclosure is power cycled without the expansion enclosure(s) being power cycled.

Workaround

To recover from this issue, power cycle the expansion tray(s) if all drives are unresponsive, or remove and re-insert any unresponsive drives.

To avoid this issue, ensure controller and expansion enclosures are on the same power circuit. This is the most commonly used configuration.

Reference Number

200608933

Snapshot image creation using VSS fails

Severity: 2 Functional

Operating System

Windows 2008 and 2012

Hardware/Software/Firmware

- E-Series SANtricity VSS Provider

Problem or Restriction

Creating snapshot images with the E-Series SANtricity VSS Provider on a Windows 2008 or 2012 host can occasionally fail with the VSS_E_NO_SNAPSHOTS_IMPORTED error. Once this occurs, subsequent attempts to create snapshot images will also fail.

Workaround

If a snapshot image failure occurs, use SANtricity Storage Manager to clean up the system by deleting the failed snapshot image. Verify that there are no stray snapshot images or snapshot image groups, and resume operations. A controller reboot might be required. Then, request a new snapshot image.

Reference Number

200612160

Drives are bypassed during a controller reboot

Severity: 2 Functional

Operating System

All

Hardware/Software/Firmware

- DE6600 drive enclosures
- ST31000425SS
- ST6000NM0104

Problem or Restriction

Rarely, during a controller reboot, certain Seagate drives are bypassed.

Workaround

Remove and reinsert the drive after the controller reboots. Make sure to wait the specified amount of time to remove or insert the drive, as documented in the appropriate drive replacement instruction.

Reference Number

200615723

Snapshot image creation via DiskShadow fails

Severity: 2 Functional

Operating System

Windows 2008 and Windows 2012

Hardware/Software/Firmware

- E-Series SANtricity VSS Provider

Problem or Restriction

Occasionally, creating snapshot images using the Microsoft DiskShadow utility on the host fails with the VSS Unexpected Provider Error error. Once this occurs, subsequent attempts to create snapshot images also fail.

Workaround

If a snapshot image failure occurs, use SANtricity Storage Manager to clean up the system by deleting the failed snapshot image. Verify that there are no stray snapshot images or snapshot image groups, and resume operations. A controller reboot might be required. Then, request a new snapshot image.

Reference Number

200618868

I/O errors occur during an online controller firmware update

Severity: 2 Functional

Operating System

Windows

Hardware/Software/Firmware

- iSCSI software initiator
- iSCSI host interface cards (HICs)
- All platforms

Problem or Restriction

This problem can occur with large configurations consisting of multiple hosts and arrays with large block I/O (2MB or 4MB) and a large number of volumes (more than 80). I/O errors are observed during an online controller firmware download.

Workaround

To avoid this problem:

- Do not configure large configurations for large block I/O.
- Or, when running large configurations and large block I/O, perform offline controller firmware upgrade during maintenance hours.

Reference Number

200632197

An asynchronous remote volume mirror (ARVM) mirrored pair is marked as failed

Severity: 2 Functional

Operating System

All

Hardware/Software/Firmware

- All controllers
- Asynchronous Remote Volume Mirroring (ARVM)

Problem or Restriction

Under extremely rare circumstances, an ARVM mirrored pair is marked as failed due to a metadata corruption of the repository. The asynchronous mirror group is suspended. User data is not lost (that is, the data in the volume is still available on the primary array).

Workaround

Remove the mirror from the asynchronous mirror group and re-add it. This will cause a full synchronization process.

Reference Number

200639882

When an ESM fails or is removed during heavy I/O operations, the controller panics and reboots

Severity: 2 Functional

Operating System

All

Hardware/Software/Firmware

- E5500
- EF550

Problem or Restriction

During heavy I/O, an ESM failure or ESM pull causes a SAS discovery error, which results in a

controller reboot.

Workaround

Do not pull an ESM during heavy I/O. Quiesce the system before performing this operation.

Reference Number

200414272

Hosts continually attempt to rebalance or perform failback on secondary volume of a mirrored pair

Severity: 3 Minor

Operating System

All OSs with multipath drivers that perform failback or rebalance

Hardware/Software/Firmware

- All E-Series Storage Arrays

Problem or Restriction

Hosts continually try to rebalance or perform failback on secondary volumes of Asynchronous Mirroring or Synchronous Mirroring mirrored pairs. This causes unnecessary I/O traffic and logging in the host system log and storage array logs. The storage array doesn't present any information that host multipath drivers can use when the LUN is initially discovered to determine whether it is non-transferable. Secondary mirrored volumes are non-transferable from a host standpoint. The host is allowed to send transfer requests to the controllers for these LUNs and will receive a good status from the storage array for the request but the storage array will not actually take any action since the secondary volumes are non-transferable. This behavior applies to both the Asynchronous Mirroring and Synchronous Mirroring. Most host multipath drivers have re-balance/fail-back mechanisms which attempt to periodically put LUNs back on their preferred path. This results in unnecessary transfer requests from the host which will result in no change in volume ownership.

Workaround

None. However, once the primary volumes are moved back to their preferred path, the secondary volumes will follow. This should cause the host multipath drivers to stop attempting to rebalance or perform fail-back on the secondary LUNs.

Reference Number

200432679

Excessive unnecessary traffic and alarming host message log entries generated on Linux hosts when a secondary volume in a mirrored pair is mapped to LUN 0

Severity: 3 Minor

Operating System

All Linux OSs

Hardware/Software/Firmware

- All E-Series Storage Arrays

Problem or Restriction

With Synchronous Mirroring and/or Asynchronous Mirroring, the host sends a SEND DIAGNOSTIC page 0 SCSI command to LUN 0, which is rejected with 07/27/00 (Write Protected) check condition, because a secondary volume was mapped to LUN 0 and secondary volumes are write-protected. When this occurs a lot of unnecessary traffic is generated and alarming entries appear in the host messages log.

Workaround

To avoid this problem, do not map a secondary volume of a mirrored pair to LUN 0.

Reference Number

200432665

Host port creation or deletion is reported unexpectedly

Severity: 2 Functional

Operating System

All

Hardware/Software/Firmware

- Controller firmware release 8.20
- SMI-S version 10.28.x.x

Problem or Restriction

This problem occurs when you move a host port or initiator from one host to another. The SMI-S provider generates unexpected creation and deletion events.

Workaround

None.

Reference Number

200321628

Pacemaker node unable to rejoin cluster after node reboot

Severity: 2 Functional

Operating System

- Red Hat Enterprise Linux (RHEL) Server release 6.6, x86_64. Base Kernel 2.6.32-502.el6.x86_64
- Active/Active Pacemaker setup built using in-box tools (pacemaker-cli-1.1.12-4.el6.x86_64, pacemaker-libs-1.1.12-4.el6.x86_64, pacemaker-1.1.12-4.el6.x86_64, pacemaker-cluster-libs-1.1.12-4.el6.x86_64)

Hardware/Software/Firmware

- All E-Series Storage Arrays

Problem or Restriction

A shutdown request for a rebooted node does not get cleared from that node's transient attributes fast enough. The result is that the node comes back online and pacemaker determines it should still be shutdown. If this particular node reboot causes the cluster to have a minimum number of nodes the cluster can handle, then another node reboot may result in a resource failure. A Bug 1198638 has been entered into the RedHat Bugzilla system.

Workaround

Restart cluster in the node which was unable to successfully join the cluster after reboot by issuing the command: `pcs cluster start`

Reference Number

200686900

VMware ESXi 6.0 host does not detect volumes after upgrade if LUN 0 unused

Severity: 2 Functional

Operating System

VMware ESXi 6.0

Hardware/Software/Firmware

- All E-Series Storage Arrays
- 12 Gb SAS HBA based on Avago/LSI 9300-8e, using inbox driver solution (lsi_msgpt3)

Problem or Restriction

VMware ESXi 6.0 hosts using 12 Gb SAS HBAs based on Avago/LSI 9300-8e and inbox lsi_msgpt3 driver, will not see volumes from storage arrays if LUN identifier of 0 is not used. For example, if volumes are mapped from storage array with LUN identifiers 1-100, the host will not detect any of the 100 LUNs if inbox lsi_msgpt3 driver is being used.

Possible customer impact scenarios:

1. If customer chooses to upgrade to ESXi 6.0 from host installed with ESXi 5.x and mpt3sas driver, then no failure should occur as long as customer selects "Upgrade ESXi, preserve VMFS datastore" option during installation. In this scenario, the mpt3sas driver used previously on ESXi 5.x will be retained after the upgrade.
2. If customer chooses to install ESXi 6.0 on host previously installed with ESXi 5.x and mpt3sas driver, this failure will occur if customer selects "Install ESXi, preserve VMFS datastore" option during installation. In this scenario, the mpt3sas driver used previously on ESXi 5.x will be not retained after installation, and the lsi_msgpt3 inbox driver will be used. If a storage array is connected to host without a corresponding LUN 0, then no devices will be seen from that storage array by the host.
3. If customer chooses to install ESXi 6.0 on host previously installed with ESXi 5.x and mpt3sas driver, this failure will occur if customer selects "Install ESXi, overwrite VMFS datastore" option during installation. In this scenario, the mpt3sas driver used previously on ESXi 5.x will be overwritten during installation, and the lsi_msgpt3 inbox driver will be used.

If a storage array is connected to host without a corresponding LUN 0, then no devices will be seen from that storage array by the host.

Workaround

This issue may be avoided with the following options:

1. If upgrading from previous ESXi 5.x release that utilizes the mpt3sas driver, select "Upgrade ESXi, preserve VMFS datastore" option during ESXi 6.0 installation to retain driver.
OR
2. If performing clean install of ESXi 6.0 release:
 - a. Enter maintenance mode.
 - b. Install mpt3sas driver.
 - c. Disable inbox `lsi_msgpt3` driver by executing `esxcli system module set -e false - m lsi_msgpt3;`
 - d. Reboot host

Reference Number

200695587

Unable to install SLES 10 SP4 as a VM

Severity: 2 Functional

Operating System

- SuSe Linux Enterprise Server 10 SP4 as a Virtual Machine
- RedHat Enterprise Virtualization Manager (RHEV-M): 3.5.0-0.29.el6ev
- RedHat Enterprise Virtualization Hypervisor (RHEV-H): 7.0 - 20150114.0.el7ev
- RedHat Enterprise Virtualization Hypervisor (RHEV-H): 6.6 - 20150114.0.el6ev
- RedHat Enterprise Virtualization Hypervisor (RHEV-H): 6.5 - 20140725.0.el6ev

Hardware/Software/Firmware

- All E-Series Storage Arrays

Problem or Restriction

During SLES 10 SP4 OS installation to a VM, it tries to create LVM on the virtual harddisk (vda), but pvcreate process ignores the volume with error code:253 (Unrecognized LVM device type). The result is that SLES 10 SP4 is unable to be installed as a VM using RHEV-M. Bugzilla 1204573 has been filed with RedHat.

Workaround

None

Reference Number

200699376

Loss of paths after RHEV hypervisor host reboot

Severity: 2 Functional

Operating System

- RedHat Enterprise Virtualization RHEV Hypervisor - 6.6 (Version 20150114.0.el6ev)
- RHEV manager: 3.5.0-0.29.el6ev

Hardware/Software/Firmware

- All E-Series Storage Arrays

Problem or Restriction

iSCSI network port settings get changed (overwritten to unfavorable defaults) after the RHEV hypervisor host reboots. Hence, it couldn't activate the iSCSI sessions at the bootup. Therefore the attached storage domains were not visible to RHEV manager. Bugzilla 1207377 has been filed with RedHat.

Workaround

Once the Hypervisor host completes the reboot, change the network port settings as desired (i.e., ONBOOT=yes BOOTPROTO=static), and then re-activate the iSCSI sessions through RHEV manager.

Reference Number

200703664

LUNs are not visible to Solaris 10 hosts with 16GB HBA directly connected to the storage using 16GB speed

Severity: 2 Functional

Operating System

Solaris 10 (All versions)

Hardware/Software/Firmware

- HBA: All supported 16GB HBAs (QLogic and Emulex)
- Direct Attached Configuration
- All E-Series Storage Arrays

Problem or Restriction

LUNs are not visible to Solaris 10 host in configurations where the host configured with 16GB HBAs (Qlogic, Emluex) are directly connected to 16GB targets. The default connection option/topology configured in the HBA driver fails because the Fibre Channel specification for FC 16 does not support loop mode and Solaris 10 does not support point to point connections.

Workaround

Change the connection options/topology to loop mode only - this forces the adapter to drop down to 8Gb FC which does support loop connection.

- For the "qlc" driver edit the file `/kernel/drv/qlc.conf` file to change the value of "connection-options" to be 0 (loop only) and reboot the server.
- For the "emlxs" driver edit the file `/kernel/drv/emlxs.conf` to change the value of "topology" to be 4 (attempt loop mode only) and reboot the server.

Reference Number

200713100

SMagent and SMdevices not functional on RHEL7.1

Severity: 2 Functional**Operating System**

RedHat Enterprise Linux RHEL 7.1 base kernel 3.10.0-229. sg driver part of `sg3_utils-1.37-5.el7.x86_64`.

Hardware/Software/Firmware

- All E-Series Storage Arrays
- NetApp E-Series SMdevices
- NetApp E-Series SMagent

Problem or Restriction

RHEL7.1 introduced a change which caused the sg drive to not load on system start. This causes the SMagent and SMdevices to not function because they require sg devices.

Workaround

Create a script that loads the sg module on boot.

From the RHEL7.1 release notes:

- **sg3_utils component, BZ#1186462**
Due to the conversion of the `iprutils` package to use `systemd` instead of legacy init scripts, the sg driver is no longer loaded during system boot. Consequently, if the sg driver is not loaded, the `/dev/sg*` devices will not be present.

To work around this issue, manually issue `modprobe sg` or add it to an init script. Once the sg driver is loaded, the `/dev/sg*` devices will be present and the sg driver may be used to access SCSI devices.

Reference Number

200713750

Windows SAN Boot installation failure

Severity: 2 Functional**Operating System**

Windows

Hardware/Software/Firmware

- All E-Series Storage Arrays
- Emulex LPe16002 HBA with Firmware/BIOS version 10.2.470.14
- Dell 12G or newer server (R620, R720, R820, R730...)
- Direct Attach topology
- SAN Boot installation

Problem or Restriction

When trying to select the SAN disk to install on during the Windows installation, Windows reported the following error and will not allow the installation to continue:

"Windows cannot be installed to this disk. This computer's hardware may not support booting to this disk. Ensure that the disk's controller is enabled in the computer's BIOS menu."

Workaround

None.

Reference Number

200736662

Solaris 10 may encounter I/O errors when array controller reboots

Severity: 2 Functional

Operating System

Solaris 10 (All updates) both Sparc and x86

Hardware/Software/Firmware

- All E-Series Storage Arrays

Problem or Restriction

During asymmetric access state transition, Solaris hosts may experience SCSI error when controllers reboot. The following message is logged when this error occurs:

```
Jun 25 03:42:11 ictm-ovm2 scsi: [ID 107833 kern.warning] WARNING: /
scsi_vhci/ssd@g600a0980005a1802000084715579aff1 (ssd52):
```

```
Jun 25 03:42:11 ictm-ovm2 Error for Command: write(10) Error Level: Fatal
Jun 25 03:42:11 ictm-ovm2 scsi: [ID 107833 kern.notice] Requested Block:
58808 Error Block: 58808
```

```
Jun 25 03:42:11 ictm-ovm2 scsi: [ID 107833 kern.notice] Vendor: NETAPP
Serial Number:
```

```
Jun 25 03:42:11 ictm-ovm2 scsi: [ID 107833 kern.notice] Sense Key:
Not_Ready
```

```
Jun 25 03:42:11 ictm-ovm2 scsi: [ID 107833 kern.notice] ASC: 0x4 (LUN not
accessible, asymmetric access state transition), ASCQ:0xa,FRU:0x0
```

SR 10965264821 has been submitted to Oracle.

Workaround

Retry the failed I/O

Reference Number

200751051

RHEL6.7 (IB-iSER) xfs filesystem I/O timeout may occur when storage array is reset

Severity: 2 Functional

Operating System

- RedHat Enterprise Linux (RHEL) 6.7 (Kernel 2.6.32-573.el6.x86_64)
- Software Stack: rdma-6.7_3.15-4.el6.noarch
- filesystem type: xfs

Hardware/Software/Firmware

- All E-Series Storage Arrays
- Infiniband HICs using iSER protocol

Problem or Restriction

xfs type filesystems may get IO timeout over infiniband iSER protocol when the storage array is reset. Bugzilla 1236141 has been filed with RedHat.

Workaround

For planned controller reboots by the user, halt any input output operations during a controller reset.

For unplanned controller reboots, wait until the controllers become optimal, make sure the paths are back to optimal state, and retry the operation.

Reference Number

200751433

RHEL6.6 disk offlined during cable pull

Severity: 2 Functional

Operating System

RedHat Enterprise Linux Server 6.6 (Kernel 2.6.32-504.el6.x86_64)

Hardware/Software/Firmware

- All E-Series Storage Arrays
- Cisco UCS El Cap MR3 - 2.2(4) configuration
 - o Cisco UCS B200 M3 blade server
 - o Cisco UCS VIC 1340 adapter with driver version 1.6.0.17, and FW version 2.2(4)
 - o Cisco UCS 6248UP fabric interconnect with FW version 2.2(4)

- Cisco Nexus5548 switch

Problem or Restriction

RedHat Enterprise Linux Server 6.6 blade server residing in the Cisco UCS system received the error: "Medium access timeout failure. offlining disk!" when one of the uplink cables (between fabric interconnect and Cisco switch) is pulled or fails. In a Cisco UCS environment the fnic driver allows more outstanding io than fnic can process. Cisco filed the RedHat Bugzilla #1024066 to update the fnic driver to get max io throttle count from UCSM to configure maximum outstanding IO supported by fnic and push that value to scsi mid-layer.

Workaround

1. You can online the disks manually

```
[root@host]#echo running > /sys/block/sd<x>/device/state
```

2. Verify the changes

```
[root@host]#cat /sys/block/sdbm/device/state
```

3. If using multipath, restart the multipath daemon

```
[root@host]#service multipathd restart
```

Another option is a complete server powercycle.

Reference Number

200766549

Oracle Linux 6.6 disk offlined during cable pull

Severity: 2 Functional

Operating System

Oracle Enterprise Linux Server 6.6 UEK 3 Update5 (Kernel 3.8.13.68.el6uek.x86_64)

Hardware/Software/Firmware

- All E-Series Storage Arrays
- Cisco UCS El Cap MR3 - 2.2(4) configuration
 - o Cisco UCS B200 M3 blade server
 - o Cisco UCS VIC 1340 adapter with driver version 1.6.0.18, and FW version 2.2(4)
 - o Cisco UCS 6248UP fabric interconnect with FW version 2.2(4)
- Cisco Nexus5548 switch

Problem or Restriction

Oracle Enterprise Linux Server 6.6 UEK 3 U5 blade server residing in the in the Cisco UCS system received the error: "Medium access timeout failure. offlining disk!" when one of the uplink cables (between fabric interconnect and Cisco switch) is pulled or fails. In a Cisco UCS environment the fnic driver allows more outstanding io than fnic can process. Cisco filed the Oracle Bugzilla #15124 to track this issue.

Workaround

1. Can online the disks manually

```
[root@host]#echo running > /sys/block/sd<x>/device/state
```

2. Verify the changes

```
[root@host]#cat /sys/block/sdbm/device/state
```

3. If using multipath, restart the multipath daemon

```
[root@host]#service multipathd restart
```

Another option is a complete server powercycle.

Reference Number

200771810

iSCSI sessions fail to restart after a OEL6.7 UEK3 U5 host reboot

Severity: 2 Functional**Operating System**

- Oracle Enterprise Linux Server 6.7 UEK3 - iscsi-initiator-utils-6.2.0.873-14.0.1.el6.x86_64
- Oracle Enterprise Linux Server 6.7 - iscsi-initiator-utils-6.2.0.873-14.0.1.el6.x86_64

Hardware/Software/Firmware

- Software iSCSI initiator

Problem or Restriction

This problem is seen when both the regular network and NetworkManager services are running, where the iSCSI network is using the regular network service and the public network is using the NetworkManager service. During system boot up, if the iSCSI sessions are started right after the regular network service, but right before starting the NetworkManager service, all the iSCSI sessions are shut down. The iSCSI sessions fail to restart afterward.

Workaround

Move all the network ports to either regular network or NetworkManager and shutdown the unused network service.

-Or -

Manually restart the iSCSI connection session, by running `iscsiadm -m node -L all`.

Another option is a complete server powercycle.

Reference Number

200774506

Data Assurance - A transient data loss may occasionally

occur with 2 failed drives in a RAID 6 volume group that is still formatting

Severity: 1 Major

Operating System

All Operating Systems

Hardware/Software/Firmware

- E2700
- E5500
- EF550
- E5600
- EF560

Problem or Restriction

With Data Assurance enabled, if 2 drives fail in a RAID 6 volume group while it is still undergoing IAF (Immediate Availability Format) and host reads arrive to blocks that have not yet been formatted, a transient data loss condition may occur.

Workaround

Wait for the IAF to complete, which will fix the issue by regenerating the Data Assurance tags or reissue the read command.

Reference Number

200784401

Linux devices offlined with multiple ISLs in a Cisco UCS config

Severity: 2 Functional

Operating System

- RedHat Enterprise Linux RHEL 6.x, RHEL 7.x
- Oracle Linux OL 6.x and OL 7.x with base kernel and UEK
- SuSE Linux Enterprise Server SLES 11 SPx, SLES 12

Hardware/Software/Firmware

- All E-Series Storage Arrays
- Cisco UCS El Cap MR3 and MR4 - 2.2(4) and 2.2(5)
- (Virtual Interface Card) Fabric Interconnect in Switching Mode to an upstream switch

Problem or Restriction

A handful of SCSI devices are offlined due to a "Medium Access Timeout" generated by the Cisco

VIC (Virtual Interface Card) fnic driver whenever multiple connections between a Fabric Interconnect (FI) running in switching mode and an upstream switch exist in a UCS Environment running 2.2(4) or 2.2(5) and one of these links is failed for any length of time.

Workaround

Operate with only a single ISL link between each switch and Fabric Interconnect (FI) for each host. The recommended setup has two FI's running to separate switches so redundancy won't be lost. In the event of a failure, traffic is just sent down the remaining FI instead.

Reference Number

200788610

ATTO Multipath Director fails to claim LUNs after Windows 10 upgrade

Severity: 2 Functional**Operating System**

- Windows 10
- ATTO Multipath Director version 1.70 (first version to support Windows 10)

Hardware/Software/Firmware

- All E-Series Storage Arrays

Problem or Restriction

When performing an in-place upgrade to Windows 10 client, the ATTO Multipath Director failover driver will no longer claim LUNs until it is uninstalled and re-installed. Running I/O without a failover driver presents the potential for data corruption, since two instances of each LUN will appear to the Windows client.

Workaround

Multipath Director should be uninstalled and re-installed as soon as possible after an in place upgrade.

Reference Number

200792000

Standard features and premium features

You should be familiar with the standard and premium feature offerings available with SANtricity Storage Manager 11.20 for E-Series storage systems and the EF540 and EF550 flash arrays.

Standard features available with SANtricity Storage Manager 11.20

A number of standard features are automatically available with SANtricity Storage Manager 11.20 for your NetApp E-Series storage system or EF540 and EF550 flash array without additional purchase. No keys are required to activate these standard features on your system.

The following features are standard in SANtricity version 11.20:

- Snapshot (enhanced)
- Volume Copy
- Synchronous Mirroring (not available for EF550 or EF560 flash arrays)
- Asynchronous Mirroring
- RAID 6 Volumes
- SANshare Storage Partitioning
- Data Assurance (DA)
- Solid State Drive (SSD) Support
- SSD Cache (not available for EF540, EF550 or EF560 flash arrays)
- High Performance Tier
- Thin Provisioning
- Dynamic Disk Pools

See the *Concepts Guide for SANtricity Storage Manager* for information about these standard features, as well as the premium features that are available for your E-Series system.

Note: Standard features do not display in the SANtricity Array Management Window (AMW) Premium Features and Feature Pack Information dialog.

Note: The SANtricity Storage Manager 11.20 online help does not distinguish between standard and premium features. The online help describes many SANtricity features as premium features, but the features listed above are in fact standard for E-Series and EF-series storage systems (except as noted).

Availability of premium features for E-Series storage systems

The Drive Security feature requires a premium feature key on E-series and EF-series systems.

The Synchronous Mirroring feature requires a premium feature key on the EF550 and EF560 flash arrays.

The following features are not available for the EF550, EF540 and EF560 flash arrays.

- Snapshot (legacy)
- SSD Cache

Copyright information

Copyright © 1994–2016 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL and other names are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277