HP LeftHand SAN Solutions

Support Document

Service Notes

VSA 7.0 for VMware ESX

Legal Notices

Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notices

Current Limitations in This Release

Installation and Upgrandes

Post-Install Qualification Window Doesn't Appear After Completing An Upgrade (6543)

Scenario

In rare cases when doing an upgrade, backout, or a patch that requires a reboot, on storage modules running any 6.x.x release, you may see a message in the Install Status Window indicating that the storage module is being rebooted, but nothing else happens.

If this is the case, you next see the following message:

> The entry in the upgrade install status window "01/11/2007 03:02:15 PM MST: It has been 5 minutes since the install server has responded. This install will time out after waiting 60 minutes with no response."

Workaround

1. Close the Console and then restart the Console.
2. Find the storage module on the network.
   The storage module may still be rebooting and may take several minutes more before it can be found.
3. When the storage module is discovered, log in to the management group, if appropriate, and verify that the storage module is running normally and is now on the upgraded version of the SAN/iQ software.
4. Complete any remaining installations.

Centralized Management Console Fails To Install On Linux (3177)

Scenario

When downloading the installer for the Console from the vendor's FTP site, the FTP program reports that the download completed successfully. However, when you run the installer, you receive an error message indicating that a Java error occurred and the installation cannot continue.

This occurs because some FTP programs may not download the complete installation package. You can verify that the download was complete by comparing the MD5 checksum of the file that was downloaded with the MD5 checksum that is published on the FTP site.

Workaround

Upgrade the FTP client you are using or use a different FTP client.

Console Loses Connection To Storage Module During Upgrade (7530)

Scenario

Sometimes the Console does not receive notification that the upgrade is complete.

During an upgrade or installation, the Console displays a message every five minutes saying that it has not received any messages from the storage module. After 60 minutes, the installation is aborted.

Workaround

Verify that the software version installed on the storage module is correct and that the system is healthy. Then perform any other upgrades.

Mixed Data Schemas May Occur (7394)

Scenario

If you start an upgrade from 6.6 to 7.0, and then move the storage module into a management group before the upgrade is finished, the result is a mix of data schemas.

Workaround

Do either of these tasks:

- Let the upgrade finish before moving the storage module into a management group.
- Move the storage module into a management group and then perform the upgrade.

Upgrade Post-Qualification May Grab Focus Every 20 Seconds (2754)

Scenario

During a software upgrade, the Console may come to the front of other windows open on the desktop and may grab focus as well.

Workaround

Click in a different window to re-establish focus elsewhere.

Upgrading Storage Modules And Management Groups May Take Some Time (4234)

Scenario

Upgrading a storage module to the current release may take 30 to 40 minutes depending upon the specific platform and configuration.

Even after the storage modules are upgraded, have rebooted, and have all been found on the network in the Console, the upgrade process may take up to another 10 minutes.

During the upgrade process, you may see messages such as “Waiting for MG1 to come up. The issue is - An NSM is down.” The storage module is not down. It is actually resynchronizing with the other storage modules in the management group.

Workaround

Wait for the resynchronization to complete.

SAN/iQ Software Upgrade On A Management Group with Mixed Version Storage Modules May Not Upgrade The Management Group Version (7698)

Scenario

In a management group with mixed SAN/iQ software versions, the management group database version does not get upgraded, For example, if the current management group version is 6.5, and some storage modules are upgraded to 70 6.6 and others are upgraded to 70 7.0, the management group version remains at 6.5.

Workaround

Upgrade all of the storage modules in the management group to the same 7.0 version, then the Console will upgrade the management group version. In the above example, upgrading all storage modules to 70 7.0 will correctly upgrade the management group version to 7.0.

Drive Status Alerts Appear When The Storage Module Is Rebooted (7700)

Scenario

When an storage module reboots, the system will display information about drive status in the Alert panel and send email alerts (if notification is configured) even though the status of the drives are unchanged from their state before the reboot.

Explanation

Normally Alerts are displayed and sent only when something in the system changed. In this case the previous drive status was not saved so that when the system reboots it cannot tell if the status has changed. To ensure the current status is known, the alerts are generated after reboot.

Finding Storage Modules on Network

Windows Fire Wall Prevents Storage Module Discovery In The Centralized Management Console (5855)

Scenario

After upgrading the SAN/iQ software, the Centralized Management Console fails to discover storage modules.

Workaround

1. Determine if Windows Fire Wall is running.
2. If Windows Fire Wall is running, disable it.

Storage Modules

When Removing A Storage Module From A Group, Status Returns State As Missing (7462)

Scenario

When swapping a storage module out of a cluster, it cannot be removed from the management group after a restripe is completed. The storage module is now in an unresponsive state and its status is: joining/leaving management group, storage server state missing.

Workaround

Reboot the storage module. When it comes up, it is marked as available, as expected.

Repair Storage Module Procedure Generates "Will Cause Restripe" Warning Message (7433)

Scenario

When adding the repaired module to the cluster during the Repair Storage Module procedure, a warning message is displayed that may give the impression that the entire cluster will be restriped. This is not the case. Only the repaired NSM will be restriped.

Workaround

Ignore the warning; just dismiss it.

Storage Module Takes Longer Than Normal To Come Up [NSM 150, NSM 160, NSM 260] (5994)

Scenario

You may sometimes notice a delay when adding storage module to a management group immediately following RAID reconfiguration.

This behavior is rarely encountered because a storage module has RAID configured at the factory.

Workaround

If you choose to reconfigure RAID from its factory settings to some other configuration, wait five minutes before adding the storage module to a management group.

Disks Are Not Hot-Swappable In The NSM 150, NSM 200, Or IBM x3650 (3533)

These models do not support hot swap disk drives.

When Replacing or Reseating A Power Supply, The Console May Report Improper Power Supply Status [NSM 160] (2997, 3532, 7060)

Scenarios

- Replacing a power supply may cause both power supplies to show "Missing" in the Console.
- If the AC power cord is plugged into the power supply during installation, the Console may report "missing" for one or both power supplies even though they are both installed and working properly.

Workaround

To restore proper status reporting, perform these steps:

1. Power down the storage module.
2. Remove both power supply power cables.
3. Wait 10 seconds.
4. Plug in the power cables.
5. Power on the storage module.

Rebooting The Storage Module While RAID Is Rebuilding Causes Reboot To Take Up To 20 Minutes [DL 380] (4048)

Scenario

If you reboot the storage module while RAID is rebuilding, the reboot can take up to 20 minutes to complete.

Cause

The lower the priority setting of the RAID, the longer it will take the reboot to complete.

Boot Flash1 Status Changes After Changing RAID Configuration On NSM 150 (5498)

Scenario

Changing the RAID configuration on an NSM 150 causes the Boot Flash1 status to change to Inactive for about 2 minutes. The status then changes to Updating, and then back to Normal.

Explanation

This status change is due to the system processing the RAID reconfiguration. If you use the factory default RAID configuration, you never see this alert. If you configure the RAID to a different RAID setting, you see the status changes one time.

RAID and Disk Management

Drive Shows As "Off or Removed" If The Storage Module Is Powered On With A Missing Drive [NSM 160, NSM 260, DL 380, DL 320s] (7331, 7480)

Scenario

When a storage module is rebooted or powered on with a drive missing from the drive bay, that drive will be reported as "Off or removed" by the Console. If that drive is then inserted into the drive bay, it will not be recognized by the Console.

Workaround

Reboot the storage module from the Console.

Disk Replacement [IBM x3650] (5968, 7543)

Warning: Incorrect disk replacement can corrupt the entire array. To avoid such corruption, be sure to follow the procedures below:

Replacing A Disk

1 Power off the original disk in the Console.

2 Remove the disk from the drive bay and insert the replacement disk.

3 Wait for the RAID status to show “rebuilding.”

4 Click the Power Disk On button.
Even if the drive appears to be on and everything appears normal, this enables drive monitoring functions for that drive.

Reseating A Disk

1 Power off the disk in the Console.

2 Power off the IBM x3650 in the Console.

3 Reseat the disk in the drive bay.

4 Manually power back on the IBM x3650.

5 Wait for the RAID status to show “rebuilding.”

6 Click the Power Disk On button.
Even if the drive appears to be on and everything appears normal, this enables drive monitoring functions for that drive.

Use A Different Disk for Disk Replacement

- If you remove a disk, you should replace it with a different disk. If you replace it with the same disk, the necessary RAID rebuild may not be initiated, even with a server reboot.

Delays with Disk Management and Disk Reporting

- When powering off a disk, there may be a lag before the status changes in the Console.
- When you replace a disk, there may be a long delay (up to 10 minutes) before the array starts rebuilding.
- In a cluster, the manager and/or storage module may temporarily go off-line when inserting a disk. The services should appear active again after a wait, probably not more than 2-3 minutes. There may be client access delays during that pause. Ensure that the client initiator timeouts are set as recommended for the SAN/iQ software.

Intermediate Disk Status Reporting

- When a disk is powered on or inserted in a drive, certain intermediate states may be reported. For example, if a drive is added to a degraded RAID 5 array, it may temporarily say Normal, before correctly changing to Degraded and then to Rebuilding.

Swapping One Or More Disks Across Controllers Causes Data Loss [NSM 260] (3342)

If the storage module powers up with one or more drives foreign to the configuration of a controller, data corruption occurs.

Scenario

The storage module is moved to a different physical location. Before the move, the storage module is powered down and all drives are removed. While replacing the drives back in the drive bays, one or more drives are accidentally inserted into slots handled by a different controller. When the storage module is powered up, data corruption occurs.

Workaround

Label the drives before removing them so that you can replace them in the correct bays.

When Replacing A Disk, If New Disk Is Seated Improperly, Disk Status Displays DMA Off With Yellow Exclamation Icon [NSM 150, NSM 200] (2848)

Scenario

A disk is replaced in an storage module. After the RAID rebuild is completed, the disk status displays DMA Off. This status occurs due to an improperly seated disk.

Workaround

Repeat the procedures for replacing the disk, paying careful attention to reseat the disk properly in the drive bay.

After the RAID rebuild is finished, the disk status should be correct.

After Reboot, Lower Capacity Disk Status Is Shown As On And Secured In An IBM x3650 That Has Higher Capacity Disks (6740)

Scenario

You insert a lower capacity disk in an IBM x3650 with higher capacity disks and reboot it. In the Console, the physical drive status appears as Active, and RAID status appears as Degraded. You will not be able to power off the lower capacity disk to replace it with the higher capacity one.

Note

Adding lower capacity disks to storage modules with higher capacity disks is not supported.

Workaround

1. Using the Console, power off the IBM x3650.
2. Replace lower capacity disk with a new, higher capacity disk.
3. Power on the IBM x3650.
   When the IBM x3650 comes up, the RAID status appears as Rebuilding and the physical drive status appears as Active.

In An NSM 200 With RAID 10, 3ware Controller May Not Bring New Disk Online Without A Reboot (7629)

Scenario

If you replace a disk, there is a chance that the replacement disk will not be recognized by the 3ware controller until you reboot the NSM 200.

Workaround

Reboot the NSM 200.

When Powering Off A Mirrored Disk And RAID Is Rebuilding, The Mirrored Disk Is Not Powered Off [NSM 150] (7368)

Scenario

If you try to power off the mirrored disk when RAID is rebuilding, it is not powered off. However, no message appears to inform you that your request for power off has been denied.

The Disk Setup panel indicates that the drive is still active, confirming that the disk has not been powered down.

Workaround

Wait for the RAID to rebuild and then power off the disk.

Changing The RAID Rebuild Rate Does Not Retain The New Setting [IBM x3650] (5780)

Scenario

If you try to change the RAID Rebuild Rate, the slider returns to the default setting of High. This setting of High affects other activities on the IBM x3650. For example, if the system is rebooted, the storage server takes a long time to start. The long start time means that the icon will continue blinking red in the Console and, even after the system is up and the storage server started, the unit's performance is affected until the RAID rebuild completes.

Workaround

There is no workaround.

Explanation

This inability to change the RAID Rebuild Rate is due to a limitation in the IBM controller firmware.

Why RAID May Go Off If A Foreign Drive Is Inserted Prior To Powering Up The Storage Module [NSM 260] (3341)

Scenario

If the storage module powers up with a drive that does not belong to the RAID configuration, data corruption may occur causing RAID to go off and preventing the storage module from coming online. Replacing the original drive may not result in RAID going to normal.

Data may be lost on this storage module in this case.

Workaround

Never replace a drive when the storage module is off. Replace a drive while the system is still operational and your are working from the Console.

“Contact support at **www.hp.com/go/support** to determine if data...” for this storage module must be rebuilt or restored.

What To Do When A Cache Corruption Alert Is Received [NSM 260] (3321)

Scenario

Cache corruption can occur if the storage module is powered down while there are data in the RAID cache. If the storage module stays powered-off long enough (more than 72 hours), data in the cache will be corrupted. When the storage module powers back up, the cache corruption is detected, and an alert is posted indicating the cache is corrupt. The storage module will not be allowed to come online in order to prevent corruption within the cluster. A “storage module down” alert will also be posted. Please note that data on the storage module has been lost in this case and must be rebuilt from the cluster, assuming replication was configured.

Workaround

To resolve the issue, please “Contact support at **www.hp.com/go/support**".

Rebuilding RAID 5 Takes Too Long When Minimum Setting Is 1 [NSM 200] (2763)

The default setting for the minimum RAID rebuild rate is 1. This setting may cause RAID 5 rebuild to take too long.

Workaround

Increase the minimum rebuild rate to a value of 10 or greater. The following guidelines describe the effects of the RAID rebuild rates.

- Setting the rate high is good for rebuilding RAID quickly and protecting data; however, it will slow down user access.
- Setting the rate low maintains user access to data during the rebuild.

Removing Drive From Storage Module Without First Removing Disk From RAID Requires Rebooting The Storage Module To Recover From Degraded Mode [NSM 150] (707)

See the section, Managing Disks, in the *LeftHand SAN User Manual* or the online help, for instructions about removing and replacing drives.

Scenario

If you remove a drive without first removing it from RAID in the Console, RAID becomes degraded and the storage module becomes inaccessible.

Workaround

1. Re-insert the drive.
2. Reboot the module.
3. Add the disk to RAID. RAID will start rebuilding after the drive is powered on.

No Warning If You Remove and Re-Add Disk To RAID 0 [NSM 150, NSM 200] (2302)

Scenario

Storage module is configured with RAID 0. While the storage module is running, you manually remove any disk from the storage module. On the Disk Setup window, the disk status is “Off or missing.” On the RAID Setup window, RAID status is Normal.

This scenario occurs when the disk is removed while there is no activity to the volume. As soon as any activity to that volume occurs, such as a client attempting to read or write data, then the volume becomes unavailable.

Single Drive Error [NSM 160, NSM 260] (6502)

Scenario

A drive may become unavailable, causing the RAID status to go Degraded or Off, depending on the RAID configuration.

Workarounds

The following three options should be tried, in order. If one does not fix the problem, try the next one.

- Reseat the drive using the instructions in the User Manual or the Online Help. If the drive does not start rebuilding, and the drive status shows Inactive in the Disk Setup tab, select the drive and click Add to RAID.
- Reboot the storage module. The drive comes online and begins rebuilding.
- Replace the drive and rebuild the array.

Network Management

Unable To Set Frame Size Or Flow Control On The VSA (8070)

Explanation

There are options in both the Centralized Management Console and the Configuration Interface to modify the frame size and flow control network parameters. These options are not currently supported by the VMware guest OS network driver. Any changes made to these variables using either interface will be accpted but no change to the physical network connection will be made. Any changes required for performance or redundancy of the network interface should be made in the ESX configuration using the VMware Virtual Infrastructure Client Interface.

Flow Control Behavior Is Erratic On Bonded NICs (7575)

Scenario

The flow control on the NIC bond does not consistently affect the flow control on either NIC interface. Modifying the flow control on the NIC bond produces unpredictable results.

When creating a NIC bond on a storage module, flow control for the NIC bond is not representative of each NIC. A NIC bond may indicate that flow control for one NIC is “enabled” and that for the second NIC, flow control is “disabled.”

Workaround

Set flow control using the following guidelines:

- Do not change flow control settings after the bond is created.
- Flow control setting on a disabled physical NIC interface cannot be changed.
- Flow control setting on the NIC bond is meaningless. Ignore the flow control setting for a bonded interface.

Caution: Changing flow control settings after creating a bond results in unpredictable flow control settings.

Caution: Creating a bond with two individual NICs that have different flow control settings results in unpredictable flow control.

Procedures for Enabling Flow Control on a NIC or NIC Bond:

To set flow control on a physical NIC interface

1. Enable the NIC by configuring the IP address.
2. Set flow control On for that NIC.

To set up a NIC bond with flow control enabled:

1. Enable each NIC by configuring the IP address.
2. Set flow control On for each NIC.
3. Create the bond.

To set up a NIC bond with flow control disabled:

1. Enable each NIC by configuring the IP address.
2. Set flow control Off for each NIC.
3. Create the bond.

Table 1 lists the expected flow control settings for various types of NIC bonds. The flow control setting should remain the same after you create any bond. If you check the NIC bond and find that the settings are not the same, delete the bond and reset the flow control settings to ensure that they are the same.

Table 1 Expected Flow Control Settings for Bond Types

	Before Creating Bond		After Creating Bond	
	Port 1	Port 2	Port 1	Port 2
Active - Passive	Enabled	Enabled	Enabled	Enabled
	Disabled	Disabled	Disabled	Disabled
Link Aggregation Dynamic Mode 802.3ad	Enabled	Enabled	Enabled	Enabled
	Disabled	Disabled	Disabled	Disabled
Adaptive Load Balancing	Enabled	Enabled	Enabled	Enabled
	Disabled	Disabled	Disabled	Disabled

Configuring The SAN On A Private versus Public Network (3836)

Best Practice

The recommended best practice is to isolate the SAN, including Console traffic, on a separate network. If the SAN must run on a public network, use a VPN to secure data and Console traffic.

Storage Traffic Is Using The Wrong Network Interface Card (NIC) (5168)

Scenario

You may see SAN/iQ traffic on NICs other than the designated one.

Explanation

This is unavoidable when two or more NICs are assigned IP addresses in the same subnet. It can occur in any configuration where hosts are configured with multiple NICs.

Workaround

Assign “public” adapters, intended for servicing users, to a subnet distinct from storage adapters.

Time On The VSA Is Out Of Sync With The Time On The ESX Server (8101)

Scenario

The customer will experience a noticeable time difference between the actual time and the time displayed on the Console for the Virtual SAN Appliance (VSA).

Solution

Using the VMware VI Client, configure ESX to sync the system clock with NTP (See ESX configuration documentation). The VSA’s time is ultimately controlled by the physical systems’ hardware clock. ESX controls the relative hardware clock for each guest operating system. If the ESX server has the incorrect time, the guest operating system will also display the incorrect time.

Reports, Logs and SNMP

In The DL 380, The Cache Battery Status Is Not Clear In The Hardware Information Log (5387)

Scenario

If you remove only the battery from the controller card, the battery status is reported as Faulty.

On the Hardware Information Report, if you see this:

Battery 1 Status faulty

This means that the battery on the controller is missing, although the controller card itself may be present.

Workaround

Replace the BBU.

Battery Capacity Test Timing Changed [NSM 160] (7040)

Scenario

If you upgrade an NSM 160 from release 6.6.x to 7.0, the battery capacity test runs every week instead of once every four weeks.

Workaround

After an upgrade, use the Console and manually change the BBU Capacity Test monitoring variable frequency to four weeks. Select a NSM 160 storage module > Alerts >Alert Setup> Edit Monitored Variable. Change the Schedule Week field to Every Four Weeks.

"NVRAM Card = Corrupt" Alert May Be Generated When The Storage Module Is Restarted After Being Shut Down For Some Time (4362)

Workaround

Call Support.

Management Groups

New User Unable To Change Original Password (7396)

Scenario

A new user logs in for the first time and wants to change his or her administrator-assigned password. The new user tries the click path starting at a management group > Administration> Edit User > Change Password and Confirmation Password.

Click OK.

A message appears: Unable To Edit The User.

Workaround

Most likely, the new user belongs to a group that has only Read-Modify permissions. A trustworthy administrator must assign a password to users having this level of permissions. Read-Modify users cannot change their own passwords.

Clusters

Cannot Create A Cluster Using VSA And Any Storage Module Running SAN/iQ Software Release 6.6 (7874)

Explanation

SAN/iQ software release 7.0 is the first release that supports mixing RAID levels in a cluster. The VSA runs virtual RAID which is new in the 7.0 release. Therefore, the VSA cannot be added to any cluster with storage modules running release 6.6 or earlier.

Volumes and Snapshots

Failover/Failback wizard does not work correctly when more than one schedule exists on volume (7696)

Scenario

When a Primary volume has multiple Remote copies to different volumes, the Failover/Failback wizard does not work correctly when initiated from the Primary volume.

Workaround

Initiate the Failover/Failback wizard from the remote volume you want to failover to. This will cause appropriate behavior to occur.

The Volume Edit Dialog Uses Existing Snapshot's Autogrow Value (7694)

Scenario

We do not recommend using the command line interface (java commandline.CommandLine) to override the SAN/iQ software autogrow values. However, if you are in a space constraint and must use the command line interface to change the autogrow value on a volume with snapshots, and you later edit the volume using Edit Volume in the Console, the autogrow value will reset to its original value.

Workaround

Wait to edit the volume using the Console until all the snapshots created prior to changing the volume's autogrow value are deleted by the snapshot schedule.

If you anticipate an immediate need for editing the volume, delete all snapshots and if the cluster space constraint is still there, only then change the volume's autogrow value.

In A DL 320s With RAID Rebuild Priority Set To High, Volume Becomes Unavailable During RAID Rebuild (7554)

Scenario

When the RAID is rebuilding on a DL 320s for which the Rebuild Priority has been set to High, and the volume is under heavy load, it is possible that the system may have difficulty keeping up with I/O and may lose the iSCSI connection.

Workaround

If the volume goes offline while RAID is rebuilding, do either of these workarounds:

- Move the RAID rebuilding priority to low. This lengthens the time that the array is rebuilding, but allows heavy I/O volume to continue.
- Reduce the load on the storage module. This allows the storage module to complete the rebuild quickly.

Converting A Volume From Remote To Primary With Thin Provisioning Makes The Volume Un-editable (7545)

Scenario

After promoting a remote volume to primary, the length of the volume is set to 0. If you attempt to edit this primary volume at any time (even months later) after you have converted it, you receive the following error message:

Volume “x” cannot be modified because the initial quota must be greater than zero and less than its length.

Workaround

1. Convert the volume back to a remote volume.
2. Convert it to a primary, but with full provisioning.
3. Edit the volume and make it thin provisioned.

MS Cluster Failovers When Migrating A Large Number Of Volumes Concurrently (7485)

Scenario

You may experience delayed write failures and cluster failovers. The client servers become unresponsive.

When migrating volumes from one cluster to another, multiple disk groups fail. The iSCSI session is overloaded and write failures and cluster failovers by accessing the iSCSI initiator UI and Cluster Admin UI during the process.

The application log is noted with a hungapp error.

Workaround

When migrating volumes, stagger the migrations. Do not start all at the same time.

Volumes With Degraded Replication Are Not Apparent (7363)

Scenario

There are two scenarios for this issue:

- You change a volume from 0-way replication to 2-way replication, and the software displays a message at the start of this transition that the volume is now 2-way replicated. However, during the transition period from 0-way to 2-way replication, if you lose a storage module or perform an upgrade, the changed volume goes offline. This scenario often transgresses to the next one.
- A storage module is down, and 2-way replicated volumes are really not fully replicated anymore. The Console shows these volumes as Normal. There is no indication that the volume has a lower reliability level than the replication goal that is set on the volume.

Workaround

Bring the storage module back online and check the replication level, changing it if necessary. Replication proceeds from where it left off when the storage module went down.

In A Cluster With A Virtual IP Address, Cannot Mount Volume Using Storage Module IP As A Discovery Address (7369)

Scenario

If a cluster has a virtual IP address, and that IP address is not used for discovery in the iSCSI initiator, you cannot mount a volume from that cluster using the storage module's physical IP address. The volume is detected, but you are unable to log in to it using the storage module IP.

Workaround

Use the virtual IP address of the cluster to log in.

Console Has Memory Problem After Deleting Many Remote Snapshots (7547)

Scenario

After deleting as few as 50 remote copy snapshots, memory usage increases to about 100MB from a more usual 40MB or so. The memory increase can be seen by selecting Help > About from the menu bar and then clicking System at the bottom of the Help About window.

In this instance, overtaxed memory causes the Console to become unresponsive.

Workaround

Exit the Console and restart it.

Snapshot Schedules Do Not Adjust For Daylight Savings Time (4383, 4913)

Scenario

When snapshot schedules are created under Standard Time, the schedules continue to execute at the originally scheduled Standard Time, even though the storage modules are operating under Daylight Savings Time.

For example, if a schedule is configured, under Standard Time, to run at 2:00 PM, then the schedule initially runs at 2:00 P.M. Standard Time. When the local time changes to Daylight Savings Time, the schedule starts running at 3:00 PM instead of 2:00 PM.

This happens because the schedule is operating as if Daylight Savings Time doesn't exist; so the schedule continues to execute at 2:00 PM Standard Time.

Explanation

The SAN/iQ software does not include automatic adjustment for Daylight Savings Time.

Workaround

If you want snapshot schedules to operate at the same relative time all year, you must manually edit the schedules when the time change in spring and changes back in autumn.

"NVRAM Card = Corrupt" Alert Generated After RAID 0 Disk Replacement [NSM 160, NSM 200] (4359)

Workaround

Reboot the storage module.

Hardware Information Report Does Not Report CPU Temperature [IBM x3650] (5703)

Reading the hardware report, the status of the CPU temperature is "not available." This is due to a limitation in the IBM Baseboard Management Controller (BMC) firmware.

Volume Not Added To Volume List Appears In iSCSI Initiator (4215)

Scenario

You create a cluster and configure the cluster to use iSNS. You then create a volume but do not add the volume to a volume list. The volume appears as a target in the iSCSI initiator. However, if you attempt to log on to this target, you receive an Authorization Failure message. This is a function of iSNS discovery.

Workaround

If you need to log on to the volume, add it to a volume list and create an authentication group, as described in the user documentation.

Remote Copy

False Alerts May Be Generated While Remote Copy Is In Progress. (7681)

Scenario

During Remote Copy, alerts in the Management Console or via SNMP or email alerts (if notification is configured) may give the impression that the Remote Copy is not working properly.

Explanation

Ignore the “Status=incomplete” alert, this is a temporary state while the copy is in progress. Wait until the Remote Copy is complete and you receive a “complete alert”.

Remote Copy From Multiple Management Groups To A Single Remote Management Group Causes Performance Drop In Remote Management Group (3499)

Scenario

A remote management group experiences a performance drop if too much bandwidth is used for transfer of Remote Copy data.

Workaround

To designate enough bandwidth for I/O to the management group, reduce the local bandwidth used for Remote Copy.

1. Log in to the remote management group.
2. On the Edit Remote Bandwidth dialog window, reduce the local bandwidth setting.

iSCSI

Two-Way CHAP Can Be Done Using One-Way Chap Password (7370)

Scenario

For One-way CHAP, you have one password and use Outgoing Authentication.

For two-way CHAP, you have two different passwords, one for Incoming Authentication and one for Outgoing Authentication.

Sometimes, you are able to mount a volume with two-way CHAP only using one password.

Workaround

Use the single password for two-way CHAP until this issue is understood more fully.

Adaptec HBA Unable To See Target (2348)

Workaround

Do not use MS iSCSI initiator with the Adaptec HBA.

iSCSI Closes All Shares After Reboot (3367)

If your iSCSI volumes are used by automatically-started Windows services, for example, File Sharing, you must use the Microsoft Initiator's “Bind Volumes” operation to make sure that those volumes are available before the services that require them are started.

Workaround

- See the LeftHand Networks document at this URL: **https://www.lefthandnetworks.com/member_area/dl_file.php?fid=1037**
- Also, see the section entitled “Running automatic start services on iSCSI disks” in the Microsoft iSCSI Initiator Users Guide for more details.

An iSCSI Volume That Becomes Unavailable For Approximately 60 Seconds Or Longer May Cause Data Loss (3396, 3298, 573)

Scenario

The Windows Registry has a default maximum hold time setting of 60 seconds before a Windows system terminates a connection to an iSCSI device that is unavailable.

This means that an iSCSI volume that becomes unavailable for longer than 60 seconds may cause delayed write failures and potential data loss.

Workaround

Change the Windows Registry settings for the default Maximum Request Hold Time to 600 (decimal) value.

Important: Back up your registry.

“Refer to the HP document – Best Practices for Enabling Microsoft Windows with SANIQ”

When Mounting Existing iSCSI Volumes On Different Servers, Volumes May Be Assigned Duplicate Drive Letters Or No Drive Letters (469, 541)

Scenario

An iSCSI volume that was mounted on a server and assigned a drive letter is logged off from Server 1. It is then mounted on Server 2. Sometimes, it picks up a drive letter that is already in use on Server 2. Sometimes, it is not assigned a drive letter. The volume then becomes inaccessible.

Workaround

Open the Windows Disk Management console and assign a new drive letter to the volume. The volume should then appear in the directory structure.

Linux-iSCSI Initiator Cannot Reboot When SAN/iQ Volume is Unavailable (3346)

Scenario

The iSCSI Device Manager hangs when network problems prevent it from communicating with a storage module. Because the default time-out for Linux-iSCSI initiator is infinite, the initiator cannot reboot when it is unable to access the iSCSI volume on the storage module.

Workaround

Restore full network connectivity between iSCSI initiators and storage modules. If this is not possible, you also can disconnect from the network the storage module that the initiator can't communicate with. Disconnecting causes the managers to tell the client that it should stop attempts to contact that storage module.

If Changing Permissions On An iSCSI Volume, Log On To A New Initiator Session To Complete The Changes (3326)

Scenario

An iSCSI volume is mounted as a read/write volume and is in use.

You change the access permissions to read-only for the authentication group in the Console.

The permissions have not changed for the clients that are accessing the volume. They are still able to write to the volume.

Workaround

To complete the process of changing permissions, log off the current initiator session for that volume and log on to a new session.

Red Hat: Changing Authentication Type Causes Existing iSCSI Devices To Be Renamed (3668)

Scenario

You configured an authentication group for iSCSI access. You then changed the access configuration, either to require CHAP or to remove or change CHAP requirements. After the change, the existing iSCSI devices are renamed and cannot be remounted.

Workaround

To change the authentication type of any volume (LVM or otherwise), follow these steps:

1. Unmount volumes and stop iSCSI services.

   ```
   # /etc/init.d/iscsi stop
   ```

2. Make appropriate changes to the authentication group (i.e. change from iqn to CHAP).
3. Make appropriate changes to the initiator (i.e. settings in /etc/iscsi.conf).
4. Start iSCSI services and remount volumes.

For LVM volume groups, the following steps are recommended since the system allows iSCSI services to be stopped even though iscsi_sfnet driver is still in use by the volume group.

To change authentication type of volumes being used in a volume group, follow this procedure:

1. Unmount volume/volume group.

   ```
   # umount /iSCSI
   ```

2. Deactivate the volume group.

   ```
   # vgchange -a n vgiSCSI
   ```

3. Stop iSCSI services.

   ```
   # /etc/init.d/iscsi stop
   ```

4. Use the change to use CHAP or whatever authentication you want to test next.
5. Restart things in the reverse order:

```
# /etc/init.d/iscsi start
# vgchange -a y vgiSCSI
# mount /dev/vgiSCSI/lvol0 /iSCSI
```

After Power Cycle, Load Balancing Does Not Distribute Requests Properly From A Microsoft Cluster (3993)

Scenario

A storage module is powered off and then powered on, and another storage module in the SAN/iQ cluster handles all the connections to the volumes connected to that cluster. When the storage module is powered on again, load balancing does not redirect I/O to that storage module.

Workaround

1. Take one of the MS Cluster groups offline.
2. Disconnect the iSCSI connections on both storage modules.
3. Reconnect the targets on both storage modules.
4. Bring the MS Cluster group back online.
5. Repeat steps 1 through 4 for all MS Cluster groups that host SAN/iQ iSCSI disks.

Load balancing will again distribute I/O requests across all storage modules.

2-way CHAP Does Not Work With Solaris 10 (4292)

Scenario

Volumes associated with an authentication group configured for 2-way CHAP cannot be mounted on Solaris 10.

Workaround

Use 1-way CHAP or no CHAP with Solaris 10.

An Extra Microsoft iSCSI Session Is Created In The Console After Rebooting The Host (5023)

Scenario

An extra iSCSI session is created in the Console after rebooting the host for the volume which is mounted with “Automatically restore this connection when the system boots” selected.

Explanation

This is a Microsoft issue in which different session IDs (iSCSI ISIDs) are used for the same hostvolume pair, depending on how the session was established. After an ungraceful host shutdown, you might see duplicate iSCSI sessions in the Console, one with a Status of Failed and one a Status of Connected.

Workaround

Log off the automatically logged on persistent session and manually log back on to get rid of the spurious session.

Microsoft iSCSI Initiator Stops With Error (5552)

Scenario

In rare cases, the Microsoft iSCSI Initiator version 2.02 and 2.03 may stop after a storage module reboots.

Workaround

Manually restart the Microsoft iSCSI Initiator Service.

Using 1-Way CHAP To Mount Volume In QLogic HBA Fails To Detect Volume (5289)

Scenario

Using the Centralized Management Console, configure an Authentication Group with a CHAP name, target secret, and initiator secret. After adding the volume list, you then attempt to mount a volume in the QLogic HBA using the target secret and initiator secret you set in the Authentication Group. The volume is not detected.

Workaround

For 1-way CHAP, use the Initiator Secret from the Console Authentication Group as the QLogic Target Secret.

For 2-way CHAP, first use the Initiator Secret from the Console Authentication Group as the QLogic Target Secret. Next, add the Target Secret from Console Authentication Group as the QLogic Initiator Secret.

Using QLogic HBA And Solaris 10, I/O Can Only Be Done On One Volume (5269)

Explanation

The QLogic HBA is not supported with Solaris 10 and the HP LeftHand SAN.

Workaround

Use the Sun Solaris native iSCSI initiator.

SuSE 9 and SuSE Linux iSCSI: Version 4.0.1-88.26 Initiator Reports Incorrect Driver State (5444)

Workaround

Use the iSCSI initiator provided with the SLES 9 distribution.

Storage Module Configuration Backup and Restore

Storage Module Post-Install Qualification Of Restored Module Stalls If Restored Module Has Different IP Address Than That Of Original Module (939)

Scenario

Back up a storage module configuration file (Unit-1). Unit-1 becomes unavailable and you restore the backed up configuration of Unit-1 to a second storage module on the network (Unit-2). Unit-2 has a different IP address than the unavailable Unit-1. As part of the post-install qualification, the Console searches for the newly configured Unit-2 on the network. However, it is searching for the original IP address of Unit-2 instead of the IP address that

was saved in the Unit-1 configuration back-up file. That search never completes because the IP address on Unit-2 has changed and is now the IP address of Unit-1.

Note: Restoring multiple storage modules from single backup file causes an IP address conflict.

Workaround

Before restoring a backed-up storage module configuration file, make certain that the new storage module is configured with the IP address of the original storage module.

Workaround

If the backed up configuration has been restored and the post-install qualification process can't complete because it cannot find the storage module on the network, do the following:

1. On the Post-install qualification window, click Cancel All Installs.
2. Either search for the storage module on the network using the correct IP address or search with Find by Subnet and Mask.

Single Disk Errors Are Not Recovered In Clusters With Storage Modules Running Mixed Software Versions (1819)

Versions 6.3 and later contain functionality to recover from any single disk unrecoverable data error. This recovery functionality only works on storage modules in clusters where all storage modules are upgraded to version 6.3 or later. If a cluster has one or more storage modules running an earlier version of the software, than the recovery functionality will not work.

If IP Address On Storage Module Is Changed Using the Configuration Interface, Some Processes Continue to Use The Old IP Address (1711)

Scenario

A storage module in a management group has an IP address assigned. That IP address is changed using the Configuration Interface instead of using the Console. The new IP address is not universally updated in the SAN/iQ software and some functions continue to use the old IP address.

Workaround

To finish updating the IP address using the Console:

1. Log in to the storage module with the new IP address.
2. On the storage module, navigate to the TCP/IP Network category.
3. On the Communication tab, select Communications Tasks and click Update Communications List.
   This synchronizes the IP addresses of all managers.

MSCS

MSCS Cluster Failover While SAN/iQ Cluster Under Heavy Load Takes MSCS Cluster Offline (1784)

If an MSCS cluster failover occurs while the SAN/iQ cluster is under very heavy load, the MSCS cluster does not come back online until the load on the SAN/iQ cluster decreases.

Workaround

Increase the pending timeout of each of the disk resources on the MSCS cluster to match the maxrequestholdtime value, for example, 600.

Do the following on each physical disk resource that is actually an iSCSI disk on the storage module.

1. Right-click on the disk in the MSCS cluster administrator.
2. Select Properties > Advanced tab.
3. Change the Pending Timeout: Seconds Field to match the maxrequestholdtime value for iSCSI in the registry.

Dell Open Manage Secure Port Server

Unable To Install Or Load Console With Dell's Secure Port Server Service Started (909)

Scenario

Using Windows on a Dell Server with the Dell OpenManage Secure Port Server service, you cannot properly install the Console or start the Console.

Workaround

Stop the Dell OpenManage Secure Port Server service when installing or running the Console.

Novell

Netware Server Stops Responding If The Storage Module That Is Hosting The VIP Becomes Unresponsive (4008)

Scenario

Novell customers may see problems during iSCSI session recovery. In a SAN/iQ cluster, when the storage module that is hosting the VIP becomes unavailable, one of the following three things could happen:

1. The iSCSI session recovers within 30 seconds.
2. The iSCSI session takes a very long time to recover (hours).
3. The iSCSI session never recovers.

Workaround

In the case of (a), no action is required.

In the case of (b) or (c), the customer should explicitly disconnect and then reconnect to the iSCSI session using Netware server configuration tools (e.g. NoRM).

Red Hat Enterprise Linux

On A RHEL Cluster With A Volume In Use, A Network Outage Longer Than 45 Seconds Results In The Volume Not Automatically Remounting [NSM 150] (6545)

Workaround

1. Deactivate the volume that was being used when the node failed on all other nodes in the cluster:

   Example:

   [root@rac8] # umount /mnt/home1

   [root@rac8] # vgchange -an home1

2. Restart the cluster services on the failed node.
3. Reactivate the VolumeGroup on the other RHCS nodes.

   Example:

   [root@rac8] # vgchange -ay home1

   [root@rac8] # mount /mnt/home1

NOTE: If the network outage does not last long enough to trigger the fence action by the RHCS lock server, then the volume comes back online automatically, assuming that the iSCSI process has not timed out itself.

Virtual Managers

Cannot Start Virtual Manager (7367)

Scenario

In a two-node management group, there are two managers running, and a virtual manager is enabled. The manager the Console used at log in has stopped for some reason. The Console correctly reports that the manager disconnected.

You log back in to the management group to start a virtual manager. Now, the Console cyclically logs into a storage module where the manager is no longer running. The Start Virtual Manager menu item for the storage module is not displayed because the global database is not available to the Console.

There is no way to start the virtual manager to recover quorum.

Workaround

1. Log out of all the storage modules.
2. Log into the storage module that has a manager running.
3. Log out of the management group.
4. Log into the management group.
   Now, the Console should try to get the global database from the storage module you are logged into and that has a manager running.
5. Start the virtual manager.