HP LeftHand SAN Solutions

Support Document

Service Notes
Service Notes for SANiQ 8.0

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

Upgrade Stops With (illegalArgumentException) Error (7035)

Scenario

During a software upgrade using the CMC, you may see the following message:

```
java.lang.illegalArgumentException:port out of range:-1
```

This can happen during any type of upgrade (patch, service pack, firmware etc.) and it is intermittent.

Workaround

1. 1. Click OK in the message window.
2. 2. Close the software upgrade window.
3. 3. Exit the CMC and restart it.

Once the storage node has been re-discovered, log in and retry the software upgrade. If the problem persists, call support.

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

Scenario

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

Even after the storage nodes are upgraded, have rebooted, and have all been found on the network in the CMC, 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 node is not down. It is actually resynchronizing with the other storage nodes in the management group.

Workaround

Wait for the resynchronization to complete.

Centralized Management Console

Configuration Summary in CMC Uses 8.0 Scalability Values for Release 7.x Management Groups (10132)

Scenario

When using the release 8.0 CMC and logging into a release7.x management group, the Configuration Summary displays the 7.x management group with release 8.0 scalability values. These values are not valid for release 7.x management groups.

Workaround

There is no workaround.

Centralized Management Console Fails To Install On Linux (3177)

Scenario

When downloading the installer for the CMC 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.

Storage Nodes

ID LED Status May Not Be Accurate in CMC [NSM 2120 G2] (10198)

Scenario

The ID LED helps you locate a storage node in a rack. With an NSM 2120 G2 storage node selected in the CMC, the ID LED status may not accurately reflect whether the ID LED is on or off.

If you turn the ID LED on using the CMC, perform some task on the storage node, then press the UID LED button on the front panel, the LED turns off. However, the status of the ID LED in the CMC will still show the LED as on.

Explanation

When you power on the storage node and the open the CMC, the default status for the ID LED is off. You can use the CMC to turn the ID LED on or off, and the CMC correctly reports the status.

If you press the UID LED button on the NSM 2120 G2 front panel, the CMC may not report the changed/correct status.

Workaround

- If you turned off the LED on the storage node front panel and the CMC shows the status as on, use the CMC to turn the LED off. This resyncs the LED and the reported status in the CMC. If you turn on the ID LED again using the CMC, the LED will turn on.
- If you cannot locate a storage node that shows the ID LED status as on in the CMC, turn the LED off, then on in the CMC. In this situation, the LED was actually off and using the CMC resynced the LED and the status in the CMC and turns on the LED.

Receive Error Message When Attempting To Log In To Storage Module [DL 320s] (9674)

Scenario

When using the CMC to log into a storage node you receive the error message “there are no more available sessions”.

Workaround

If you receive this error message, contact Support.

Drive Is Undetected By Storage Node [DL 320s, NSM 2120] (7834)

Scenario

If the storage node is booted or re-booted with an empty drive bay, a drive added later will not be detected, it will be listed in the Storage category Disk Setup tab as “Off or Removed”.

Solution

Reboot the storage server with the drive installed. After reboot, the drive will be detected and listed in the Storage category Disk Setup tab.

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

Scenario

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

Workaround

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

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

Scenario

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

This behavior is rarely encountered because a storage node 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 node to a management group.

Disks Are Not Hot-Swappable In The IBM x3650 (3533)

These models do not support hot swap disk drives.

When Replacing or Reseating A Power Supply, The CMC 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 CMC.
- If the AC power cord is plugged into the power supply during installation, the CMC 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 node.
2. Remove both power supply power cables.
3. Wait 10 seconds.
4. Plug in the power cables.

5 Power on the storage node.

Best Practice

Install patch 10006-00 on NSM 160s running pre-7.0 SAN/iQ software and 10006-07 on NSM 160s running SAN/iQ software version 7.0 and later.

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

Scenario

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

Explanation

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

After Changing RAID Configuration Flash Status Changes On NSM 160 and NSM 260 (5498)

Scenario

After changing the RAID configuration, flash status changes on platforms with boot devices, and you see flash status alerts.

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.

Best Practice

After reconfiguring RAID, wait for the RAID resync to complete.

Administrative Users and Groups

Users with Read-Only Password Access Can Change Passwords (9902)

Scenario

If a user is a member of a group that has read-only access to passwords but read-write or full access to management groups, the user will be able to change passwords.

Explanation

This is because the operation to change passwords is now at the management group level instead of individual node level and the user in the scenario has read-write access to management group.

Workaround

There is no workaround.

When Creating an Administrative Group with the SAN/iQ Command Line Interface (CLI), the Only Allowable Permission for Reports is Read-Only (10326)

Scenario

When using the SAN/iQ CLI to create an administrative group (createAdminGroup), if the reports permission is set to anything other than read-only (r), the following error message with be given:

```
RESPONSE
result          80001012
processingTime …
name            CliqInvalidParameter
description     Invalid parameter Permissions
```

Solution

Always use the read-only permission for reports. Example:

```
CLIQ>createadmingroup admingroupname=group_name login=123.45.678.900
username=admininistrator password=password permissions=rfrfr
```

Access to the On-Node SAN/iQ Command Line Interface (CLI) Can Be Compromised if a Management Group is Created with an Initial Administrative User Named "admin" (10292, 10310)

Scenarios

The release 8.0 default administrative user, with full permissions, for a storage node in the Available pool is named "admin." This default user provides you access to the on-node SAN/iQ CLI before putting the storage node into a management group.

When you create a new 8.0 management group, and create a new administrative user, that initial administrative user has full permissions. If you name that user "admin," there are two scenarios that can prevent you from accessing the on-node SAN/iQ CLI:

- Scenario 1: An intermittent failure can occur that may prevent this newly created administrative user called "admin" from accessing the on-node SAN/iQ CLI.
- Scenario 2: This management group user named admin does have access to the on-node SAN/iQ CLI. However, if you create a second administrative user for that same management group with any name, the user named admin no longer has access to the on-node SAN/iQ CLI.
- Scenario 3: When upgrading an existing management group to 8.0, if the existing user is named admin, then after the upgrade to 8.0, that user will continue to log in via the CMC but will be unable to log in to the on-node SAN/iQ CLI.

Workarounds

- Scenario 1: Because of this potential issue, we recommend that you choose some name other than admin for your administrative user. If you do encounter this issue, contact Customer Support to re-enable access to the on-node SAN/iQ CLI.
- Scenario 2: If you want to have a management group administrative user named admin for onnode SAN/iQ CLI access, you must create this user last. If the user is already created, delete it and recreate it as the last full-permissions administrative user for that management group.

- Scenario 3: To work around the upgrade situation, do the following:

1. Using the CMC, create a new user, for example, joe, as part of the full-access group.
2. Log in to the CMC as User joe and delete the user admin. At this point, we recommend that you continue to use the User joe. However, you could also re-create the User admin as part of full-access group.

RAID and Disk Management

When Upgrading to Release 8.0, the Installation Aborts with an Error Message About OS Array Degraded (9814)

Scenario

Prior to an upgrade, if the storage node experienced any transient Read errors, they could cause the RAID array to perceive the disk drive as faulty, although the drive is good. Since these errors occurred on the disk partition that was part of the OS array, the array gets marked degraded. During a subsequent upgrade, the following error may appear in the install status screen:

```
"RAID status is not satisfactory for Storage Node xxx (xx.xx.xx.xx).
OS array md200 degraded and cannot be restored automatically:
err=3072"
```

Explanation

In the above scenario, since the RAID array with user data did not experience any errors, the RAID Status in the CMC continues to show Normal.

Workaround

There is no workaround. This error requires support intervention.

Confusing Error Messages When Attempting To Power Cycle A Drive That is Offline (9895)

Scenario

A drive goes offline with a status of Off or Removed and a safe to remove status of Yes. Using the CMC Disk Setup Tasks, you attempt to power off or power on disk. Either option returns a series of confusing error messages without indicating whether the power cycle worked.

Workaround

Ignore the error messages and use the prescribed drive replacement procedures for the platform. Find drive replacement procedures in the Storage chapter of the LeftHand SAN User Manual, in the section “Replacing a Disk.” This information is also in the Online Help available from the CMC.

Console Reports Drive Health Status As Faulty or Normal While The Drive Is Rebuilding [DL 320s, NSM 2120, NSM 2120 G2] (9169)

Scenario

While a drive is rebuilding the CMC may incorrectly show drive health as faulty or normal.

Explanation

The software may not be able to access drive health while that drive is rebuilding.

Workaround

To verify the correct drive health status use the HP Systems Insight Manager.

RAID Reconfiguration Needs To Complete Before Using Large Capacity Storage Nodes Like The NSM 2120, DL 320s and NSM 4150 (8935)

Explanation

If you reconfigure RAID on an NSM 2120, NSM4150, or DL 320s with 750GB SATA drives, you should wait for the RAID reconfiguration to completely finish before adding the storage module to a management group. Not letting the RAID initialization complete may cause a decrease in performance under heavy IO.

Solution

We recommend waiting 48 hours after reconfiguring RAID before using any large capacity storage node. Currently there is no indication given in the CMC when initialization has completed.

Workaround

There is no workaround.

Lower Capacity Disk Is Not Detected In A DL 320s, DL 380 or NSM 2120 (6651)

Scenario

If you replace one of the drives with a smaller capacity drive, the storage node won't recognize it. The disk status will be Off or Removed because there is not enough capacity on that drive to rebuild.

Solution

Replace the original drive with the same or higher capacity. Do not replace any drive with a lower capacity drive.

On A Storage Node With A Lagging System Time, RAID Status And Disk Status Is Shown As Normal When One Of The Disk Fails And Is Removed [Dell 2950, DL 380, DL 320s, IBM x3650, NSM 2060, NSM 2120, NSM 4150] (8133)

Scenario

If RAID and disk status are Normal and Active after removing a disk, and the status does not get updated to Degraded, Off, Missing or other appropriate status, after waiting for as much as 120 seconds, the system time on the storage node may not be correct. Some reasons that the system time is not synced include:

- The storage node is not configured for NTP
- The storage node cannot communicate with the NTP server
- The NTP server is having problems, etc.

Monitoring RAID status and disk status is based on time, so if the system date/time is set backwards unexpectedly, the next polling cycle to obtain RAID and disk status will be in the future and until that time arrives, the status remains at the existing setting.

Workaround

Verify the current system time.

1. Select the management group in the navigation window.
2. Select the Time tab.
3. Check the management group time listed at the top.
4. Click the Time Tasks menu at the bottom of the tab and select Refresh All.

If the time lags by a significant amount, either

- Configure NTP, or if it is already configured,
- Verify the communication to that NTP server or try another NTP server and observe if the time remains synced.

Reconfiguring RAID On An IBM x3650 And Rebooting The Storage Node May Return RAID To Rebuilding State (5986)

Scenario

If an IBM x3650 has RAID rebuilding and if, to save time, you select Reconfigure RAID from the RAID Setup Tasks menu on the Storage > RAID Setup tab and choose the same RAID type as before, the RAID status will go to Normal. However, if that storage node is then rebooted, the RAID status changes from Normal to Rebuilding.

Explanation

Reconfiguring RAID on an IBM x3650 that was already rebuilding the RAID array causes the RAID controller to perform an un-necessary rebuild. This is normal controller behavior. No data is lost or corrupted.

Solution

Wait for the rebuild to complete.

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 CMC.
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 CMC.
2. Power off the IBM x3650 in the CMC.
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 CMC.
- 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 node 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.

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

Scenario

The storage node is moved to a different physical location. Before the move, the storage node 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 node is powered up, data corruption occurs.

Workaround

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

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 CMC, 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 nodes with higher capacity disks is not supported.

Workaround

1. Using the CMC, 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.

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 CMC 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 Node [NSM 260] (3341)

Scenario

If the storage node 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 node from coming online. Replacing the original drive may not result in RAID going to normal.

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

Workaround

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

Contact Support to determine if data for this storage node 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 node is powered down while there are data in the RAID cache. If the storage node stays powered-off long enough (more than 72 hours), data in the cache will be corrupted. When the storage node powers back up, the cache corruption is detected, and an alert is posted indicating the cache is corrupt. The storage node will not be allowed to come online in order to prevent corruption within the cluster. A “storage node down” alert will also be posted. Please note that data on the storage node 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.

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 node. The drive comes online and begins rebuilding.
- Replace the drive and rebuild the array.

Network Management

Unclear Warning Message When Configuring Management Group and the SAN/iQ Interface is Disabled (10086)

Scenario

You successfully upgrade storage nodes in the Available pool from 7.1 to 8.0. Next you create a management group using those nodes. When you assign a VIP in the Create Management Group wizard, and click Next, an error message opens stating that the storage nodes' IP addresses are not reachable by the VIP.

Explanation

If the SAN/iQ interface is disabled, then the VIP cannot reach the storage node.

Workaround

Enable the SAN/iQ interface on each storage node and try again. Find the SAN/iQ interface using the following steps.

1. Navigate to the storage node TCP/IP configuration category and select the Communication tab.
2. Select the manager IP address and open the Communication Tasks menu.
3. Choose Select SAN/iQ Interface from the menu.

Editing DNS Suffixes in CMC Adds New Entry Instead of Changing Existing Entry (10260)

Scenario

When attempting to edit a DNS suffix in the TCP/IP Network configuration category, the edited domain name is added to the list of DNS suffixes in the Edit window, instead of changing the original.

Workaround

Delete existing entry and add new entry.

Unable to Create A Bond With Incorrect Message That NIC Flow Control Settings Are Different (9968)

Scenario

When attempting to bond two NICs, the bond fails with the message that flow control settings are different on the NICs, although the CMC shows that flow control is off on both the physical interfaces.

Workaround

1. For the NSM 160 and NSM 260, apply Patch 10033-00.
2. Using the CMC, manually set the flow control off on both NICs, even if the settings are already off.
3. Now create the bond.

NIC Description is Displayed as Unknown for VSA and FOM-ESX (9728)

Scenario

When you look at the TCP/IP configuration category, TCP/IP tab, for a VSA or the Failover Manager for ESX, the Description is listed as “unknown.”

Explanation

This information is obtained directly from the VMware Ethernet driver. We do not modify non-proprietary drivers.

ALB And 802.3ad Bond On The Storage Node May Show A NIC Failure If Network Switch Autonegotiate Is Disabled [DL 320s, DL 380 or NSM 2120] (7855)

Scenario

If you configure a bonded network interface using ALB or 802.3ad, it may show one NIC as failed if the network switch ports connected to the NICs are not set to auto negotiate.

Solution

- Update the firmware on the switch to the latest level available from the switch manufacturer.
- Change the switch ports configuration to 'auto negotiate'.

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 CMC 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.

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 accepted 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.

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

Scenario

You may see storage 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.

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

Best Practice

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

Reports, Logs, Status, SNMP and Performance Monitoring

The Monitored Variable for Volume Threshold Change Is No Longer Supported (10092)

Explanation

The monitored variable for Volume Threshold Change appears in the Alert Setup tab as a monitored variable. In fact, it is not supported in release 8.0.

Unable to Set Threshold Action on CPU Utilization Variable Under Alert Setup (10262)

Scenario

If you want to set notifications for the CPU Utilization variable, and you use the Set Threshold Actions menu choice from the Alert Setup Tasks menu, the settings you configure are not saved. Consequently, you will not receive the notifications you expect.

Explanation

This is because the CPU Utilization variable does not have the CMC alert enabled by default.

Workaround

Set the notifications you want by using the Edit Monitored Variable function from the Alert Setup Tasks. Step 2 of that process allows you to configure notifications for the variable threshold changes.

Network Utilization Values Are Not Accurate and May Exceed 100% (10165)

Scenario

On certain platforms, when monitoring network utilization using the Performance Dashboard, occasionally the values are not accurate and they may exceed 100%.

Workaround

Use either one of these methods.

- In the Performance Monitoring Table, use the data in the Average column, and ignore the data in the columns Value, Minimum, and Maximum.
- Change sampling interval from the default value of 5 seconds to 10 seconds. This setting will not persist after closing the CMC, so you must change it every time you open the CMC and want to view network utilization.

"Volume Threshold Change" Monitoring Variable Not Applicable to 8.0 and Later (10092)

Scenario

On the Alert Setup tab in the CMC, the Volume Threshold Change variable still appears. This variable is not applicable to 8.0 and later releases.

Workaround

Disregard the variable.

Some Log Files Show “localhost” (9594)

Scenario

After imaging a storage node, then rebooting the storage node, some logs show “localhost” instead of the network name.

Explanation

The syslog function comes up and starts logging events before the network has initialized, therefore syslog does not know the network name. Once the network initializes, all logs show the network name.

Workaround

There is no workaround.

Performance Monitor Sometimes Pauses When Storage Nodes Reboot (10065)

Scenario

Using the Performance Monitor node, and monitoring the performance of cluster I/O, reboot one of the storage nodes in the cluster. The rebooted storage node may cause a log out of the management group, which will pause the monitoring. If the storage node does not cause a log out from the management group, the reboot may still cause the performance monitor to pause the monitoring.

Workaround

Navigate to the Performance Monitor window and click Resume Monitoring on the toolbar.

Hardware Information Report Incorrectly Labels the NSM 4150 Disk Enclosure (8214)

Scenario

In the Hardware Information tab of an NSM 4150, the Sensors section lists temperatures for the “disk shelf.”

The *LeftHand SAN User Manual* refers to this component of the NSM 4150 as the “disk enclosure.”

In the NSM 4150, the Power Supply Status for the Disk Enclosure Is Not Clear in the Hardware Information Log (7999)

Scenario

If you connect only one power supply of the disk enclosure to a power source, the status of the unconnected power supply is reported as faulty.

On the Hardware Information Report, if you see this under Power Supplies:

Number 3 faulty

or

Number 4 faulty

This most likely means that the power supply is not connected to a power source. The Hardware Information Report should list the disconnected power supply as offline.

Workaround

Check that the power supply is not connected to a power source. If it is not, ignore the faulty report or connect the power supply.

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 CMC and manually change the BBU Capacity Test monitoring variable frequency to four weeks. Select a NSM 160 storage node > Alerts >Alert Setup> Edit Monitored Variable. Change the Schedule Week field to Every Four Weeks.

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.

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

Workaround

Call Support.

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

Workaround

Reboot the storage node.

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.

Management Groups and Managers

CMC Falsely Indicates Loss Of Quorum (6555)

Scenario

The CMC suggests there is no quorum to be found when the manager it is connected to is not in the quorum.

Workaround

Log out of the management group, and log back in. If there is quorum, the CMC should log into a manager which is correctly reporting a quorum.

Powering on Failover Manager and Its Clone At Same Time Causes the Original Failover Manager To Go Into "Manager Starting" Status (7970)

Warning: You should not clone a Failover Manager or VSA after either on is in a management group. You must only clone a VSA while it is in the Available Nodes pool.

Scenario

Start up the original Failover Manager and the clone, and then log into the management group. While both nodes are shown in the group, the original Failover Manager status is reported as 'Manager starting' and the clone node status is 'Normal.'

Workaround

Power off the cloned Failover Manager.

CMC Alerts May Be Received When Failover Manager Is Started Or Added To A Management Group (9747, 9705, 9699)

Scenario

There are two scenarios for this issue:

- When adding a Failover Manager to a management group you may receive an alert that the storage server status = no quorum. The alert is a benign warning and is caused by the Failover Manager initializing. The alert can be ignored, and it stops occurring once the Failover Manager is initialized.
- When stopping or starting a regular manager in a management group with a Failover Manager running, the Failover Manager may alert that the CPU utilization is 100%. These alerts are temporary and should cease once the regular manager has completed its operation.

Workaround

Since these are temporary alerts, there is no workaround.

Manager List On Storage Nodes Not Updated With Failover Manager's IP Address (8930)

Scenario

When a Failover Manager is added into a Management Group, its IP address is not added to the manager list of any storage node.

Workaround

To correct this:

1. Login to the storage node
2. Select the “TCP/IP Network” node in the tree view (on the left side of the CMC)
3. Select the “Communication” tab on the right side of the CMC
4. Select “Update Communication List” from the “Communications Tasks” menu
5. Repeat this for each storage node in the management group

Volumes and Snapshots

Convert Temporary Space Allows Creating a Volume with No Name (10076)

Scenario

Use Convert Temp Space command to create a new volume from a snapshot's temporary space. Do not enter anything in the Volume Name field of the Convert Temp Space dialog. The new volume is created without a name.

Workaround

1. Take a snapshot of the no-name volume.
2. Create a SmartClone volume from the snapshot.
3. Delete the no-name volume.
4. [optional] To move the snapshot data to the new named volume, delete the no-name snapshot that was created as part of the SmartClone process.

Cannot Mount SmartClone Volume Created From A Snapshot That Has Temporary Space in Cluster with Less Than 4 MB Space (9943)

Scenario

If you create an additional SmartClone volume from a snapshot that is, or has been, mounted and has temporary space, you cannot mount the additional SmartClone volume. This is a rare occurrence that may occur when more than one SmartClone volume is created from a clone point.

Workaround

Free up 4 MB of space in the cluster. One way to do this is to delete unused temporary space from other snapshots. See the chapter Provisioning Storage in the LeftHand SAN User Manual for more information on managing space in the cluster.

Snapshot Delete May Fail Due to Cluster Being Out of Space (9944)

Scenario

In certain cases when a cluster is close to full, and you try to delete a snapshot manually, or a scheduled snapshot is being deleted, the delete may fail. For a manual deletion you see an error message. For the scheduled snapshot that fails to be deleted, it remains on the cluster.

Workaround

Free up 4 MB of space in the cluster. One way to do this is to delete unused temporary space from other snapshots. See the chapter Provisioning Storage in the LeftHand SAN User Manual for more information on managing space in the cluster.

Volume Utilization Alerts May Not Be Accurate (8052)

Scenario

In release 7.0, if you enable the Volume Utilization alerts and a volume is written to within 90 or 95% of its size, and snapshots exist on that volume, the alerts are sent for the snapshots instead of the volume itself.

In release 8.0, the Volume Utilization variable no longer exists.

Workaround

In release 8.0, use your operating system's disk space monitoring tools for volumes and use the Cluster Utilization variable in the CMC to monitor space in the cluster.

Pre-7.0 Volumes With Autogrow Enabled Use More Space Than Is Required After Upgrading To SAN/iQ Software Version 7.0 (7644)

Scenario

When upgrading from a pre-7.0 SAN/iQ software version, like 6.6 or 6.6 SP1, to version 7.0, snapshots of volumes with autogrow enabled take up more space than is required. This is because, after upgrading to 7.0, the pre-7.0 volumes configured with the autogrow setting of "Auto" have a default autogrow value of 512 MB, whereas the autogrow value for volumes newly created with 7.0 is 128 MB.

Workaround

We do not recommend using the command line interface (java commandline.CommandLine) to override the software's default autogrow values. However, if you are in a space constraint in the cluster, the autogrow value of pre-7.0 volumes can be set using the command line interface. The autogrow value should be set to 128MB.

java commandline.CommandLine <admin name> <admin password> <manager ip> volume_autogrow_set <volume name> 128

To verify or query the current autogrow value, run java commandline.CommandLine <admin name> <admin password> <manager ip> volume_autogrow_get <volume name>

In A DL 320s Or NSM 2120 With RAID Rebuild Priority Set To High, Volume Becomes Unavailable During RAID Rebuild (7554)

Scenario

When the RAID is rebuilding on a DL 320s or NSM 2120 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 node. This allows the storage node to complete the rebuild quickly.

In A Cluster With A Virtual IP Address, Cannot Mount Volume Using Storage Node 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 node's physical IP address. The volume is detected, but you are unable to log in to it using the storage node IP.

Workaround

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

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 nodes 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.

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

Remote Copy Displays Incorrect Error Status While Primary Volume is Restriping (9938)

Explanation

When the status column of the remote snapshot shows Error, that may be because the remote copy activity is paused or delayed due to volume restriping. However, this delay is not an error state, and the remote copy activity resumes without intervention.

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

When Upgrading to Release 8.0, Servers Created with CHAP Required Do Not Appear in Performance Dashboard and Show Up as N/A in the iSCSI Sessions Tab (10435)

Scenario

During the upgrade from 7.x to release 8.0, authentication groups that were created with CHAP required, or modified to use CHAP, are changed to Servers. However, those Servers do not appear in the Performance Dashboard and they show up as N/A in the iSCSI Sessions tab.

Workaround

Contact Support for help with this issue.

Favorite/Persistent Targets Not Reconnected When Server Is Rebooted (9987)

Scenario

Volumes or targets are connected to servers as “persistent” or “favorite” targets in the iSCSI initiator. If you log in to a target and specify an IP address to use for persistent connections, the persistent target is not always reconnected after the server reboots.

Workaround

This is due to a problem in the Microsoft initiator version 2.07. There is no workaround at this time.

Connecting Multiple Servers with Different iSCSI Load Balancing Settings to the Same Volume Causes Server-Volume Connectivity Problems (9514)

Scenario

Connecting multiple iSCSI servers with different load balancing settings to the same SAN volume causes problems with server-volume connectivity. A warning message displays only once the first time you initially create the condition or modify the configuration to cause a server load balancing mismatch. If you do not correct the mismatch, and set up all the servers with the same load balancing configuration, then when a server in this scenario reboots, it may not be able to reconnect to the volume.

Solution

Set the load balancing option to be the same on all servers accessing a volume. If any one server does not support load balancing, turn off the load balancing option for all servers for that volume.

Load balancing support is based on the iSCSI initiator in use on the server. To see a list of supported initiators, open the New Server or Edit Server window, and click the link next to the load balancing option entitled “Information on compliant initiators.”

1-way CHAP Does Not Work With IBM AIX (9422)

Scenario

Volumes associated with a server configured for 1-way CHAP cannot be mounted on IBM AIX.

Workaround

Do not use 1-way CHAP with IBM AIX.

Site Network Failure In Multi-Site Configuration Takes Cluster Resources Offline (8982)

Scenario

In a campus SAN with multiple subnets (with a Virtual IP on each subnet) and both VIPs in the discover list, persistently bound volumes will fail if one of the sites fails. The iSCSI connections do not fail over to the other site that is up.

Workaround

This is due to a problem in the Microsoft initiator. Though this was observed with version 2.06, we have no confirmation that earlier versions would behave differently. There is no workaround for this – short of not using persistent binding. If this happens then administrator intervention is required to bring disk volumes back on line. High availability applications with Microsoft clusters are probably the biggest impact target of this. LeftHand Networks has opened a case with Microsoft to investigate this problem.

Campus SAN Site Failovers Do Not Work With Virtual IP Load Balancing Enabled(8499)

Scenario

In a campus SAN with multiple subnets and multiple Virtual IPs, VIP load balancing will result in loss of access to volumes in the event of a site failure.

Workaround

Use Microsoft iSCSI Software Initiator Version 2.07.

LUN Resets Taking Too Much Time For Applications On Windows 2008 (8590)

Scenario

If a storage node fails when there is I/O pending on an iSCSI connection to that storage node there is a chance that the MS iSCSI initiator will become confused about the state of its connection. This is due to the following process:

1. The initiator does not detect that the connection to that storage node has failed.
2. The initiator issues a LUN reset to cancel the pending I/O so that the initiator can reissue the I/O.
3. The LUN reset fails because it's sent to a failed target.
4. There is a race condition in the initiator between it finally detecting the connection is failed and the LUN reset failing.

If the initiator determines that the LUN reset failed before it determines that the connection has failed, it might issue device failure status to applications, thereby causing the applications to fail. If this failure happens then administrator intervention is required to restart the applications.

Microsoft clusters are particularly sensitive to this problem. LeftHand Networks has opened a case with Microsoft to investigate this problem.

Workaround

The preventative workaround is to manually add/modify the following setting in the registry:

HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Disk\TimeOutValue

This should be a DWORD and should be added/modified to have the value 30 (decimal) instead of the default of 10.

See the Microsoft KnowledgeBase article for more details. **http://support.microsoft.com/kb/954088**

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:

  http://www.lefthandnetworks.com/searchcontrol.aspx?t=Best%20Practices%20for%20Enabling%20Microsoft%20Windows&sd1=ASC&sort1=title

- 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.

See the LeftHand Networks document at:

http://www.lefthandnetworks.com/searchcontrol.aspx?t=Best%20Practices%20for%20Enabling%20Microsoft%20Windows&sd1=ASC&sort1=title

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 node. 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 node.

Workaround

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

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 CMC.

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

Workaround

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

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

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 CMC After Rebooting The Host (5023)

Scenario

An extra iSCSI session is created in the CMC 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 CMC, 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 node 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 CMC Authentication Group as the QLogic Target Secret.

For 2-way CHAP, first use the Initiator Secret from the CMC Authentication Group as the QLogic Target Secret. Next, add the Target Secret from CMC 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 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 Node Configuration Backup and Restore

Storage Node 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 node configuration file (Unit-1). Unit-1 becomes unavailable and you restore the backed up configuration of Unit-1 to a second storage node 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 CMC 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 nodes from single backup file causes an IP address conflict.

Workaround

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

Workaround

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

1. On the Post-install qualification window, click Cancel All Installs.
2. Either search for the storage node 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 Nodes 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 nodes in clusters where all storage nodes are upgraded to version 6.3 or later. If a cluster has one or more storage nodes running an earlier version of the software, than the recovery functionality will not work.

Command Line Interface (CLI)

Windows 2008 User Access Control Modification Is Required to Run Certain SAN/iQ Command Line Interface (CLI) Commands (10145)

Scenario

With Windows 2008 and Windows Vista, there is a new security feature called User Access Control (UAC). UAC enables additional security confirmation checks when accessing portions of the operating system that are considered core.

Any Windows 2008 server configured with the default security settings can not run the following SAN/iQ CLI commands until the solution below is applied.

getLocalVolumes

provisionVolume

connectVolume

disconnectLocalVolume

getScsInfo

removeVolume

createKey

Solution

Enable the SAN/iQ CLI to run as an administrator. To do this:

1. Navigate to the SAN/iQ CLI executable, using File Explorer: C:\Program Files\Common Files\LeftHand Networks\SIQSP.
2. Right click on the SAN/iQ CLI executable, cliq.exe, and select properties.
3. Select the compatibility tab and select “Run this program as an administrator.”
4. Run the SAN/iQ CLI to execute the commands listed above. Executing these commands should succeed.

Changing the Virtual IP (VIP) Address Using the SAN/iQ Command Line Interface (CLI) Causes the VIP to Be Disabled If the Optional Parameter, Usevip, Is Not Specified and Set To 1 (10174)

Scenario

Using the SAN/iQ CLI modifyCluster command to change the VIP address successfully changes the VIP to the address specified. However, the "Use this virtual IP for this cluster" becomes disabled if the optional parameter, useVIP, is not explicitly set to 1.

In the CMC, the affected cluster flashes a warning, and the Details tab is red. The cluster status is set to:

VIP error; at least one enabled VIP is required.

Workaround

When using the modifycluster command to change the Virtual IP address, include the modifycluster parameter useVip and set it to 1 (useVip=1).

If the command is executed without the useVip parameter set to 1, the Virtual IP can be re-enabled through the CMC, or by using the SAN/iQ CLI and executing the modifycluster command with the useVip parameter set to 1 (useVip=1).

The SAN/iQ Command Line Interface (CLI) and the SAN/iQ VSS Provider Share Common Software Components (8365)

Scenario

If both the SAN/iQ CLI and the SAN/iQ VSS Provider are installed, uninstalling one program will not remove the common install directory for these two software components: C:\Program Files\Common Files\LeftHand Networks\SIQSP.

Explanation

The SAN/iQ CLI and the SAN/iQ VSS Provider share common software components that, when installed, exist in the location: C:\Program Files\Common Files\LeftHand Networks\SIQSP. To remove this directory and files, both the SAN/iQ CLI and SAN/iQ VSS Provider must be uninstalled.

Executing a SAN/iQ Command Line Interface (CLI) Command to Query the Status of the Storage Node Manager Returns Whether a Manager is Configured, Not Whether a Manager is Running (9280)

Explanation

Any SAN/iQ CLI command that returns the managerRunning status is really returning status about whether a manager is configured for that storage node, not the operating status of the manager.

Workaround

Use the CMC to determine the actual operating status of the manager.

Using the SAN/iQ Command Line Interface (CLI) provisionVolume Command in a Single Node Management Group Requires That the Optional Replication Parameter Be Used and Set to 1 (10013)

Scenario

When using the SAN/iQ CLI provisionVolume command in a single node management group, if the replication parameter is not explicitly entered and set to 1, the following error message is returned and the command fails:

RESPONSE

result 80001010

processingTime …

name CliqOperationFailed

description CreateVolume failed: CIM_ERR_FAILED: A general error occurred that is not covered by a more specific error code: “Create Volume Operation Failed: Volume 'v' cannot be created because the requested replication level is incorrect. It must be less than or equal to the number of storage nodes in the cluster.”

Note: the error message is indicating that CreateVolume failed, even though this is a provision volume command. There is an intermediate create volume step in the provision volume command and that is where the failure is occurring.

Solution

Specify the optional replication parameter and set it to 1 in this situation. The optional minreplication parameter must also be explicitly stated for this command to succeed.

For example:

CLIQ>provisionvolume volumename=v4 clustername=c size=2GB login=172.31.146.138 username=user password=password replication=1 minreplication=1

A New User Cannot Be Added to the Default Full_administrator Group Using the SAN/iQ Command Line Interface (CLI) modifyAdminGroup Command (10014)

Scenario

Attempting to add a new user to the default full_administrator group using the SAN/iQ CLI modify – AdminGroup command results in the following error message and the command fails:

RESPONSE

result 80001026

processingTime …

name CliqDefaultAdmin

description You cannot delete, modify permissions, or remove the last user from the default administration group

Workaround

There are two possible workarounds:

1. Using the CMC, create and add a new user to the full_adminstrator group
2. Using the SAN/iQ CLI, create a new group with full administrative rights and add the new user to this new group.

If a SAN/iQ Manager is in the Process of Being Started or Stopped, Executing a SAN/iQ Command Line Interface (CLI) Command to Query the Status of the Manager May Return an Incorrect Value (10078)

Scenario

Any SAN/iQ CLI command that returns the managerRunning status may return an incorrect status if a manager is in the process of being started or stopped. This is most likely to occur when running a script that stops or starts a manager and immediately follows with a command to return the manager status.

Commands that return the manager status (managerRunning) include: getGroupInfo, GetNsminfo, modifyGroup, deleteGroup, createGroup.

Explanation

managerRunning status is really returning whether a manager is configured for that storage node, not the actual status of whether the manager is running on the storage node or not. A manager is configured if it is started on a storage node, and not configured if it is stopped. See the Service Note for 9280 for more details.

Workaround

To query the manager status after stopping or starting a manager via a script, add a time delay before issuing a command to check the manager status. The duration of the delay depends on many factors such as the size of management group, load on the system, etc. However, a delay of 30 seconds should be sufficient.

Microsoft Does Not Support the Creation of a VShadow Copy of a FAT32 Drive (9866)

Scenario

MS Windows does not support using the SAN/iQ VSS Provider or the SAN/iQ Command Line Interface (CLI) vssSnapshot command to create a VShadow copy of an MS Windows FAT32 mounted volume.

Workaround

An NTFS vshadow copy of an NTFS drive is supported. Unmount the FAT32 volume and remount the volume as a supported NTFS drive.

Service Console (Health Check)

Removing Directories from C:\Documents and Settings\Administrator\Local Settings\Temp Can Cause the Service Console Utility to No Longer Function (10121)

Scenario

The Service Console creates a directory in the C:\Documents and Settings\Administrator\Local Settings\Temp directory and requires that directory to be present in order to function. This directory is typically named with a single number.

If this directory is removed, the Service Console will not run and there will be an Application Event error message containing the following description:

> "The following information is part of the event: LeftHand Networks Health Check, unable to create temporary directory: 203"

Solution

Uninstall and reinstall the Service Console utility.

The Service Console May Leave Large files in the C:\Documents and Settings\Administrator\Local Settings\Temp If It Is Unable to Transfer Windows Logs (10103)

Scenario

There are situations not yet categorized that can cause the Service Console to fail transferring the logs that have been gathered. If this failure occurs, the log files are copied to C:\Documents and Settings\Administrator\Local Settings\Temp in directories, named lhn-sysinfo-xxxx.dir. These files can become very large.

Solution

The lhn-sysinfo-xxxx.dir log directories can be removed manually, or automatically by using the disk cleaning manager. To run the disk cleaning manager, execute “cleanmgr” at a command prompt, and select Temp files.

If these log directories are present, we recommend that you contact LeftHand Networks Customer Support to determine if the Service Console is correctly transferring logs.

Installing the Service Console and Returning to the Schedule and Email Contact Page Converts the Scheduled Time to an Invalid Format (8735)

Scenario

When installing the Service Console, after completing the Schedule and Email contact panel, if you return to the Schedule and Email contact panel, the Schedule Time is converted to an invalid format. Clicking Next from this panel without correcting the invalid format gives the following error:

> Invalid time format. Time must be formatted like HH:MM AM or HH:MM PM.

Workaround

Before clicking Next to continue to the next installation panel, re-enter the Schedule Time in the requested format - HH:MM AM or HH:MM PM.

Unexpected Results Can Occur When Removing or Modifying the LeftHand Service Console (9916)

Scenarios

With the LeftHand Networks Service Console Utility installed, execute the installer and select the Modify option. A command window opens with a warning message “WARNING: Are you sure you want to remove the task “LhnHealthChkV7” (Y/N)?

Explanation

The Y or N response is mainly to confirm if the scheduled task should be removed or not. We recommend that you answer Y to the question and the installer will proceed with the remaining tasks of the Remove or Modify operation.

Explanation of Service Console (Health Check) Errors That Appear in the Microsoft Windows Application Log (10012)

Scenario

The Service Console (HealthCheck) does not include a event log source. When it fails to gather logs from a storage node or if there is an issue with sending the logs to ftp://ftp.lefthandnetworks.com, the error message logged in the application event log shows –

Event Type:	Error
Event Source:	Application
Event Category:	None
Event ID:	0
	…

Description:

The description for Event ID (0) in Source (Application) cannot be found. The local computer may not have the necessary registry information or message DLL files to display messages from a remote computer. You may be able to use the /AUXSOURCE= flag to retrieve this description; see Help and Support for details. The following information is part of the event: LeftHand Networks Health Check, failed to get logs from [#.#.#.#].

Solution

Ensure that all storage nodes are online using the CMC. Ensure that the ftp.lefthandnetworks.com site can be accessed from the client server. If a storage node IP address has changed or a storage node has been removed from the management group, please modify the Service Console to update the list of IP addresses to be monitored.