Intel® Storage System SSR316MJ2

Software Release Notes

Revision 1.3

November, 2006

Storage Systems Technical Marketing

Revision History

Date	Revision Number	Modifications
July, 2006	1.0	1st Release copy.
August, 2006	1.1	Added information regarding SP1 and BBU issues
November, 2006	1.2	Added 6.5 w/SP1 updates & issues, noted by *(6.5 w/SP1)* in the issue title.

Disclaimers

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION, OR SAMPLE.

Information in this document is provided in connection with Intel® products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications.

Intel retains the right to make changes to its test specifications at any time, without notice.

The hardware vendor remains solely responsible for the design, sale and functionality of its product, including any liability arising from product infringement or product warranty.

Table of Contents

1 Introduction 11

2 Summary of Issues Fixed in SAN Software 6.5 w/SP1.......... 11

3 Upgrading to 6.5 w/SP1 12

Platforms Supported for This Release.......... 12

3.1 Special Feature Key Upgrade Procedure for This Release.......... 12

3.1.1 Feature Key Overview 12

3.1.2 Prerequisites.......... 12

3.1.3 Download The Upgrade Components 13

3.2 Install the Storage System Software 6.5 w/SP1 Console 13

3.2.1 Best Practice.......... 13

3.2.2 Selecting The Type of Upgrade 14

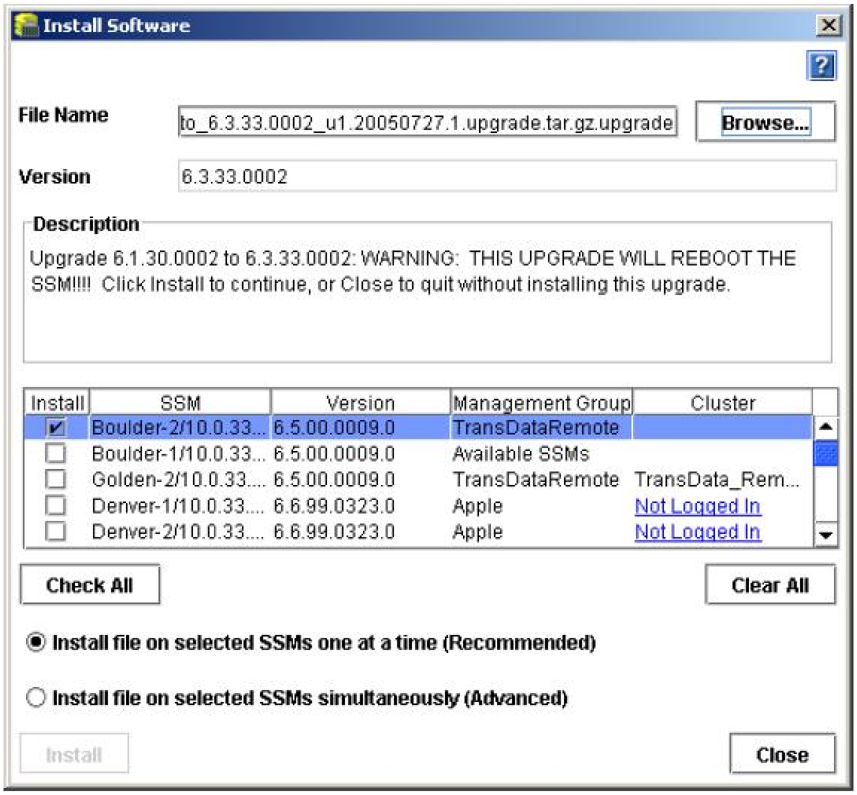

File Name: to_6.3.33.0002_u1.20050727.1.upgrade.tar.gz.upgrade

Version: 6.3.33.0002

Description

Upgrade 6.1.30.0002 to 6.3.33.0002: WARNING: THIS UPGRADE WILL REBOOT THE SSM!!!! Click Install to continue, or Close to quit without installing this upgrade.

Install	SSM	Version	Management Group	Cluster
☑	Boulder-2/10.0.33...	6.5.00.0009.0	TransDataRemote	
☐	Boulder-1/10.0.33...	6.5.00.0009.0	Available SSMs	
☐	Golden-2/10.0.33....	6.5.00.0009.0	TransDataRemote	TransData_Rem...
☐	Denver-1/10.0.33....	6.6.99.0323.0	Apple	Not Logged In
☐	Denver-2/10.0.33....	6.6.99.0323.0	Apple	Not Logged In

● Install file on selected SSMs one at a time (Recommended)

○ Install file on selected SSMs simultaneously (Advanced)

.... 14

3.2.3 Verify Management Group Version 14

4 Current SAN & SSC Software Limitations.......... 16

4.1 Storage System Console (SSC) 16

4.1.1 Storage System Console Fails To Install On Linux 16

4.2 Upgrades 17

4.2.1 Upgrade Post-Qualification May Grab Focus Every 20 Seconds 17

4.2.2 Storage System Module Running 6.5 w/SP1 Will Lose Connection with Lower, 6.3.43 and below, Version Modules (6.5 w/SP1) 17

4.2.3 Upgrading Storage Modules And Management Groups May Take Some Time (6.5 w/SP1) 17

4.3 Storage System Module 18

4.3.1 Disk Should Not Be Hot-swapped. 18

4.3.2 How To Correctly Identify A Faulty Power Supply 18

4.3.3 Rebooting The Storage Module While RAID is Rebuilding Causes Reboot to Take Up to 20 Minutes 18

4.3.4 Repair Storage Module Stalls When Attempting to Remove The Storage Module 19

4.4 RAID and Disk Management 19

4.4.1 Adding Disks to RAID In Degraded Mode Results In RAI Going Off 19

4.4.2 Why RAID May Go Off If A Foreign Drive Is Inserted Prior To Powering Up The Storage Module 19

4.4.3 Swapping One Or More Disks Across Controllers Causes Data Loss 19

4.4.4 Rebuilding RAID 5 Takes To Long When Minimum Setting is 1 20

4.4.5 When Replacing A Disk, If New Disk is Seated Improperly Disk Status Displays DMA Off With Yellow Exclamation Icon 20

4.4.6 Removing Drive from Storage Module Without First Removing Disk from RAID Requires Rebooting the Storage Module to Recover from Degraded Mode........................ 20

4.4.7 No Warning If Remove and Re-Add Disk To RAID 0.. 21

4.5 Network Management.. 22

4.5.1 Configuring The SAN On A Private vs. Public Network (SP1).............................. 22

4.6 NTP.. 23

4.6.1 Setting Time On Storage Module Caused Volume To Become Unavailable......... 23

4.7 Management Groups 24

4.7.1 Restoring A Management Group Configuration Fails Because The Configuration Backup Does Not Reflect An IP Address Change Made To An Storage Module................. 24

4.8 Clusters.. 25

4.8.1 If Incorrect Virtual IP Information Is Entered, SSR316MJ2s Go Offline And Volumes Become Unavailable 25

4.8.2 VIP Is Not Being Enforced for iSCSI Load Balancing.. 25

4.9 Volumes & Snapshots 26

4.9.1 Snapshots Don't Autogrow When Restriping Exceeds Hard Threshold (SP1) 26

4.9.2 Snapshot Schedules Do Not Adjust For Daylight Savings Time (SP1) 26

4.9.3 Frequent Re-synchronizing Of Snapshots Prevents Other Operations (6.5 w/SP1)26

4.9.4 Creating A Volume With A Duplicate Name Does Not Give An Error (SP1).......... 27

4.9.5 Volumes Go Offline And The Microsoft* Initiator Cannot Find A Volume When Snapshot Hard Threshold Is Reached 27

4.9.6 Cannot Retain Fewer Scheduled Remote Snapshots Than Primary Snapshots... 27

4.9.7 Snapshot Schedules Do Not Adjust for Daylight Savings Time............................ 27

4.10 Remote IP Copy 28

4.10.1 Remote IP Copy From Multiple Management Groups To A Single Remote Management Group Causes Performance Drop In Remote Management Group 28

4.10.2 Remote Copy Schedules Failed To Recreate When Restoring A Remote Management Group. 28

4.10.3 No Error Message Received When Remote Copy From Management Groups With Different Storage System Software Version Is Started (6.5 w/SP1)................................... 28

4.11 iSCSI.. 30

4.11.1 iSCSI Disconnect/Reconnect May Take An Application Cluster Offline 30

4.11.2 Volume Migration Can Cause Application Cluster To Loose Disk(s).................... 30

4.11.3 iSCSI Closes All Shares After Reboot 30

4.11.4 An iSCSI Volume That Becomes Unavailable For Approximately 60 Seconds Or Longer May Cause Data Loss.. 31

4.11.5 When Mounting Existing iSCSI Volumes On Different Servers, Volumes May Be Assigned Duplicate Drive Letters Or No Drive Letters 31

4.11.6 Linux-iSCSI Initiator Cannot Reboot When Storage System Software Volume Is Unavailable.. 31

4.11.7 If Changing Permissions On An iSCSI Volume, Log On To A New Initiator Session To Complete The Changes 32

4.11.8 Microsoft* iSCSI Initiator Does Not Support Dynamic Disks................................ 32

4.11.9 iSCSI Volume Disappears From The iSCSI Initiator “Active Sessions” Window When Using Scheduled Snapshots... 32

4.11.10 Failover Works As Long As A Virtual IP Address Is Used With iSCSI Initiators from Microsoft*, Intel®, Solaris*, Qlogic*, Adaptec*, Novell*, HP* and IBM*............................... 33

4.11.11 Starting Or Stopping A Manager Causes iSCSI TO Disconnect And Reconnect.. 33

4.11.12 RedHat: Changing Authentication Type Causes Existing iSCSI Devices To Be Renamed.. 34

4.11.13 After Power Cycle, Load Balancing Does Not Distribute Requests Properly From A Microsoft Cluster (6.5 w/SP1).. 34

4.11.14 When Using Storage System Software DSM for MPIO, If User Logs Off A Target Session While Client Is Accessing A Volume, The Path Disappears From MS Initiator (6.5 w/SP1) …………………………………………………………………………………….35

4.11.15 iSCSI Load Balancing Does Not Properly Balance iSCSI Sessions When Running A Mix Of Servers With Storage System Software DSM for MPIO And Servers with iSCSI Load Balancing Enabled (6.5 w/SP1)... 35

4.11.16 iSCSI Load Balanced Connections Using Virtual IP (VIP) Are Not Re-assigned After Volume Is Migrated or Storage Module Removed From Cluster (6.5 w/SP1) 36

4.11.17 Failed Initiator Session Records Are Not Always Removed From Database 36

4.11.18 An Extra Microsoft iSCSI Session Is Created In The Console After Rebooting The Host 36

4.11.19 Microsoft iSCSI Initiator Stops With Error 37

4.12 Configuration Backup and Restore 38

4.12.1 If IP Address On SSR316MJ2 Is Changed Using the Configuration Interface, Some Processes Continue to Use The Old IP Address 38

4.12.2 Single Disk Errors Are Not Recovered In Clusters With SSMs Running Mixed SAN Software Versions 38

4.13 MSCS 39

4.13.1 Need Windows Patch If Running With VSS In A MCS Environment (SP1) 39

4.13.2 Rebooting SSR316MJ2 That Are Hosting More Than 2 Microsoft Cluster Nodes May Cause Failover Errors 39

4.13.3 MCS Cluster Failover While SSR316MJ2 Cluster Under Heavy Load Takes MCS Cluster Off-line 39

4.14 Dell Open Manage Secure Port Server 40

4.14.1 Unable To Install Or Load Console With Dell's Secure Port Server Service Started… 40

4.15 Reporting and SNMP 40

4.15.1 Fan Status Log Details Are Unclear 40

4.15.2 After Changing Host Name of A SSR316MJ2 in A Management Group, Alerts Use Old Host Name 40

4.15.3 “NVRAM Card = Corrupt” Alert Generated When The SSR316MJ2 Is Restarted after Being Shut Down for Some Time 40

4.15.4 Monitoring Variables That Have Been Removed Cannot Be Re-Added in SSM 100s and SSM 200s 41

4.15.5 Reconfiguring RAID Disables SNMP 41

4.16 Volume Lists and Authentication Groups 41

4.16.1 Volume Lists Must Contain Only Authentication Groups With Same Load Balancing Configurations (6.5 w/SP1) 41

4.16.2 Enable or Disable Load Balancing On An Authentication Group Requires Logging Off And Re-logging On To Volume (6.5 w/SP1) 41

4.17 Fiber Channel (SSR316MJ2 Only) 42

4.17.1 If RAID Fails, LUNs Still Show As Available To Fiber Channel Network 42

4.17.2 Windows Server 2003 SP1 Systems Cannot See Full Capacity Of Volume When Volumes Exceed 2TB 42

1 Introduction

The following Release Notes provide information about current limitations in this 6.5.01.0017 (SSR316MJ2) release of the Storage System SAN software with Service Pack 1 (SP1), and 6.5.01.0015 release of the Storage System Console (SSC) software.

2 Summary of Issues Fixed in SAN Software 6.5 w/SP1

In configurations using Microsoft Cluster Services (MSCS), changing hot spare status causes MSCS clusters to lose connections to volumes
iSCSI disconnect/reconnect may take an application cluster offline
Volume Migration can cause application cluster to lose disk(s)
When configuring Virtual IP, if subnet mask or default gateway are entered and no IP address is entered, operation does not finish and console hangs
Volumes go offline and the MS Initiator cannot find volume when snapshot hard threshold is reached
After upgrading to Storage System Software release 6.5 w/SP1, Cluster Utilization monitored variable is not installed on the storage module
Add ALB Network Load Balancing (balance-alb) to all storage module platforms that supports the feature. (316SSRMJ+ only)
After upgrading to Storage System Software release 6.5 w/SP1, changes to the BBU Capacity Test Overdue monitored variable are not installed on the storage module
Power supply voltage being reported incorrectly and thus causing false alerts
Messages were being logged internally and prematurely filling the log partition
Solaris write performance is slow
Upgrade fails if the storage module performs an NTP update during the upgrade process
Excessive network communication occurs when Remote Copy is copying sparsely filled snapshots
RAID Configuration Reference Documentation is not correct for some configurations
Setting or modifying the time by more than 30 seconds on a storage module can cause the storage server to become temporarily unavailable
Problem with network communication causes a restripe to occur more slowly than necessary
Configuration settings (NTP, SNMP, NIC Bonding) are not retained after a DOM replacement

3 Upgrading to 6.5 w/SP1

Platforms Supported for This Release

Upgrades to Release 6.5 w/SP1 are available for SSR316MJ2. Please use the following procedures to upgrade storage modules from 6.3.xx.xxxx to Release 6.5 w/SP1.

If you see the following type of message during the upgrade process, please call customer support.

```
UPGRDE: 6.5.xx.xxxx.main—Upgrade will now be aborted.

Thu May 19 22:34:56 GMT 2005: UPGRADE: 6.5.xx.xxxx.main aborted
```

3.1 Special Feature Key Upgrade Procedure for This Release

3.1.1 Feature Key Overview

Storage System Software 6.3 and above

With Storage System Software 6.3 and above, keyed features are enabled per SSM. Feature keys are required for the following add-on features and applications:

- Scalability Pak
- Configurable Snapshot Pak
- Remote Data Protection Pak

Customers may use these add-on features and applications without a feature key, but are limited to a 30-day trial period. After the 30-day trial period, if a feature key is not purchased, any volumes and snapshots associated with add-on features or applications will become unavailable until a feature key is purchased and applied.

3.1.2 Prerequisites

- If you are running a software version earlier than 6.3 on the SSR316MJ2, you must first upgrade to 6.3. You can then upgrade to Release 6.5 w/SP1
- If any iSCSI volumes are in use, stop any activity to those volumes and unmount or log off before beginning the upgrade. The storage module reboots as part of the upgrade process. Consequently, your volumes may go off-line, depending on your configuration.
- Ensure that you are using version 6.5.00.0034 of the Console before upgrading the Storage System Software on the storage modules.

3.1.3 Download The Upgrade Components

1. Download the upgrade components to a temporary location.
2. Download the latest Storage System Console (6.5.01.0015) from the SSR316MJ2 support site or from your IBL account.
3. Download the appropriate Storage System Software upgrade package(s) from the SSR316MJ2 support site or from your IBL account.

Platform	For Storage System Software Version	Use Upgrade File
SSR316MJ2	6.5 w/SP1	6.x_to_6.5.01.0017.20060915.SSR316MJ2.upgrade

3.2 Install the Storage System Software 6.5 w/SP1 Console

1. Install the Storage System Software 6.5.01.0015 Console and discover storage module(s) on the network.
2. Use the Console to install the Storage System Software 6.5.xx.xxxx upgrade. If you do not have a direct path to the 6.5 release, you may be required to first upgrade the software on the SSM(s) to a version that can then upgrade to 6.5.

3.2.1 Best Practice

- Virtual IP Addresses - If a Virtual IP (VIP) address is assigned to a storage module in a cluster, the VIP storage module needs to be upgraded last. The VIP storage module is shown in a field in the clusters detail tab.
  - First upgrade the non-VIP storage modules that are running managers one at a time.
  - Then upgrade the non-VIP non-manager storage modules.
  - Lastly, upgrade the VIP storage module.
- Remote Copy – If you are upgrading management groups with Remote Copy associations, you must upgrade the remote management group first. If you upgrade the primary group first, Remote Copy will stop working.

3.2.2 Selecting The Type of Upgrade

The Storage System Software Console supports two types of upgrades, as shown in the figure below.

- One-at-a-time (recommended) - this is the default and only method if the storage modules exist in a management group.
- Simultaneous (advanced) - this allows you to upgrade multiple storage modules at the same time if they are not in a management group. Use this for new storage modules and/or re-configured storage modules.

1. Select from the list which storage modules to upgrade.
2. Select the type of upgrade.
3. Click Install.

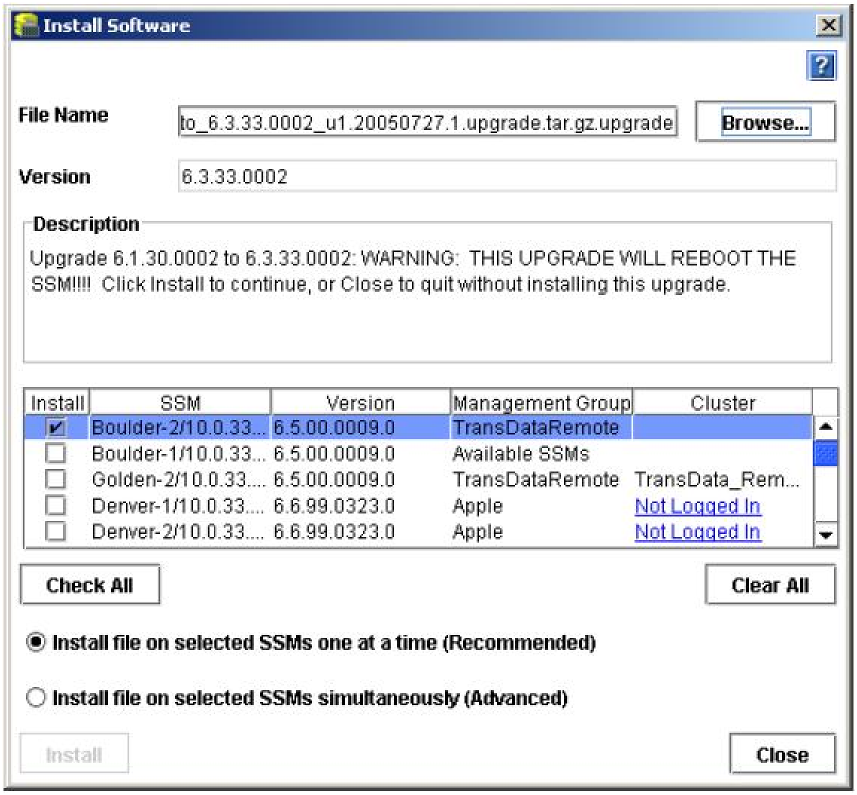

File Name to_6.3.33.0002_u1.20050727.1.upgrade.tar.gz.upgrade

Version 6.3.33.0002

Description

Upgrade 6.1.30.0002 to 6.3.33.0002: WARNING: THIS UPGRADE WILL REBOOT THE SSM!!!! Click Install to continue, or Close to quit without installing this upgrade.

Install	SSM	Version	Management Group	Cluster
☑	Boulder-2/10.0.33...	6.5.00.0009.0	TransDataRemote	
□	Boulder-1/10.0.33...	6.5.00.0009.0	Available SSMs	
□	Golden-2/10.0.33....	6.5.00.0009.0	TransDataRemote	TransData_Rem...
□	Denver-1/10.0.33....	6.6.99.0323.0	Apple	Not Logged In
□	Denver-2/10.0.33....	6.6.99.0323.0	Apple	Not Logged In

● Install file on selected SSMs one at a time (Recommended)

○ Install file on selected SSMs simultaneously (Advanced)

3.2.3 Verify Management Group Version

You must verify the management group version only when upgrading Release 6.3 to 6.5 w/SP1. If you are upgrading from Release 6.5 to 6.5 w/SP1, the management group version remains at 6.5.

After the upgrades are complete, the Console attempts to upgrade the management group version. The rules for this operations are as follows

- Upgrade management group version to 6.5 w/SP1 if all storage module serial numbers (eth0 MAC address) are in the list of known serial numbers included in the upgrade file
- Do not upgrade the management group version to 6.5 w/SP1 if any of the storage module serial numbers are unknown.

If you get the following message (or similar message), the management group version upgrade has not completed. Until the management group is upgraded to 6.5 w/SP1, you will not be able to take advantage of all the 6.5 w/SP1 features such as iSCSI load balancing.

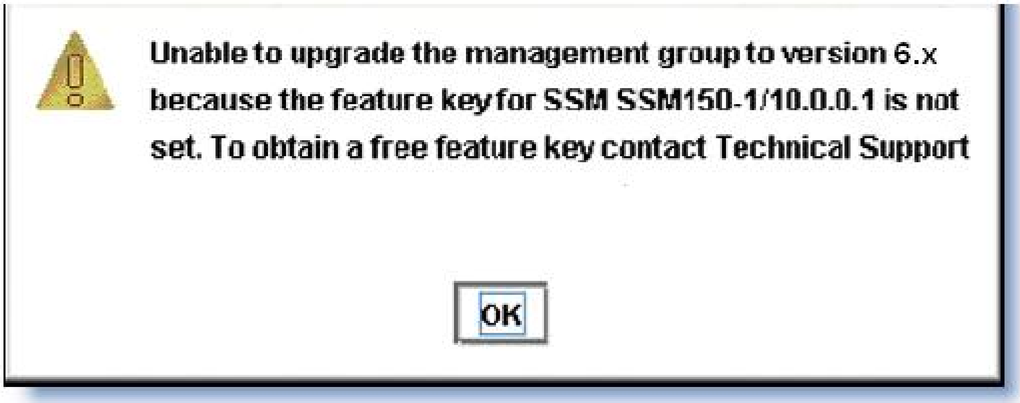

Unable to upgrade the management group to version 6.x because the feature key for SSM SSM150-1/10.0.0.1 is not set. To obtain a free feature key contact Technical Support

To receive Feature Key assistance, please contact you vendor.
After the management group version is upgraded, the upgrade is complete.

4 Current SAN & SSC Software Limitations

4.1 Storage System Console (SSC)

4.1.1 Storage System Console Fails To Install On Linux

Issue

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 certify 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.

4.2 Upgrades

4.2.1 Upgrade Post-Qualification May Grab Focus Every 20 Seconds

Issue

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

Workaround

None.

4.2.2 Storage System Module Running 6.5 w/SP1 Will Lose Connection with Lower, 6.3.43 and below, Version Modules (6.5 w/SP1)

Issue

SSM running 6.5 w/SP1 will not be able to communicate with modules running lower releases of the software, which may result in connection lose to volumes when working in clusters.

Workaround

Upgrade all modules to 6.5 w/SP1 by following the Best Practices outline in section 3.2.1 of the Release Notes.

4.2.3 Upgrading Storage Modules And Management Groups May Take Some Time (6.5 w/SP1)

Issue

Upgrading a storage module from version 6.3.xx to 6.5.xx takes from 15 to 45 minutes depending upon the specific platform and configuration.

Additionally, after the storage modules are upgraded, they have rebooted, and have all been found on the network in the Console, the management group health check may take up to another 10 minutes.

During the management group health check you may see messages such as “Waiting for MG1 to come up. The issue is – An SSM is down.” The module is not down. It is actually resyncing with the other modules in the management group.

Workaround

Upgrade all modules to 6.5 w/SP1 by following the Best Practices outline in section 3.2.1 of the Release Notes.

4.3 Storage System Module

4.3.1 Disk Should Not Be Hot-swapped.

Issue

After performing the hot-swap and systems reboot, user cannot login to the management group.

Workaround

You must first power down the SSM or just the individual drive via the SSC, before removing and replacing a hard disk drive.

4.3.2 How To Correctly Identify A Faulty Power Supply

The Intel® Storage System SR212MJ2 ships with only one power supply; therefore, only one power supply is listed in the passive report. Status will be either normal or faulty.

Issue

In a system that has been upgraded to add a redundant module, if a power supply is not working properly, the storage console passive report it will report power supply status as faulty. The faulty module number will not be identified.

Workaround

To identify the storage module with a faulty power supply.

1. On the Module Information tab, click Set ID LED On.
2. The ID LED on the left front of the module illuminates a bright blue. Another ID LED is located on the back of the module on the right side under the empty slot.
3. Go to the back of the storage module and look at the two power supplies.
4. A green LED will be illuminated on the working power supply and an amber LED on the faulty power supply.
5. Replace the faulty power supply.

Note: To ensure redundancy, the two power cords must be connected to separate and independent power sources.

4.3.3 Rebooting The Storage Module While RAID is Rebuilding Causes Reboot to Take Up to 20 Minutes

Issue

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

Workaround

None

4.3.4 Repair Storage Module Stalls When Attempting to Remove The Storage Module

Workaround

1. Close the Storage System Console and reopen it. The storage module has moved from the cluster to the management group and the ghost storage module is in the cluster.
2. Remove the storage module from the management group.

4.4 RAID and Disk Management

4.4.1 Adding Disks to RAID In Degraded Mode Results In RAI Going Off

Issue

When adding disk capacity to a RAID configuration, adding a drive that is already part of the RAID configuration but is failed or offline will cause RAID to go off.

Workaround

Select only the newly added drives when adding capacity to any RAID configuration. If only the newly powered-on drives are added to RAID, the drives are added successfully.

4.4.2 Why RAID May Go Off If A Foreign Drive Is Inserted Prior To Powering Up The Storage Module

Issue

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

Drive replacement should ALWAYS be done using the Console. Select the drive to replace, click power-off, insert a new drive, click power-on, and then click add-to-RAID.

4.4.3 Swapping One Or More Disks Across Controllers Causes Data Loss

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

Issue

The storage module is moved 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

Labels should be added to drive carriers when first installed. If this has not been done, label the drives before removing them so that you can replace them in the correct bays.

4.4.4 Rebuilding RAID 5 Takes To Long When Minimum Setting is 1

Issue

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.

4.4.5 When Replacing A Disk, If New Disk is Seated Improperly Disk Status Displays DMA Off With Yellow Exclamation Icon

Issue

A disk is replaced in an Storage Module. After the RAID rebuild is complete, 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.

4.4.6 Removing Drive from Storage Module Without First Removing Disk from RAID Requires Rebooting the Storage Module to Recover from Degraded Mode

Issue

If a drive is removed 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.

4.4.7 No Warning If Remove and Re-Add Disk To RAID 0

Issue

Storage Module is configured with RAID 0. While the Storage Module is running, user manually removes 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 Issue 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.

4.5 Network Management

4.5.1 Configuring The SAN On A Private vs. Public Network

Issue

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.

Workaround

None.

4.6 NTP

4.6.1 Setting Time On Storage Module Caused Volume To Become Unavailable

Issue

If you change the time on the storage module that is the coordinating manager, the volumes attached to that cluster can become unavailable for a short period. This situation can occur if you change the time by more than about 15 seconds, or if you configure NTP. The time change causes communication between the managers to get out of sync.

Workaround

1. Change time by increments of 15 seconds or less.
2. After changing the time by a larger increment, wait for about 30 seconds and the system will come back.
3. Change the time when this temporary unavailability is not a problem.

Or

1. First change the time on non-coordinating manager.
2. Next, stop the coordinating manager, change its time, and start it back up.

4.7 Management Groups

4.7.1 Restoring A Management Group Configuration Fails Because The Configuration Backup Does Not Reflect An IP Address Change Made To An Storage Module

Issue

This Issue can occur when the IP address change is applied to an Storage Module, but does not get stored in the system before the management group configuration is backed up.

Workaround

If you do restore a management group configuration with an incorrect IP address, take the following steps:

1. Log in to the SSR316MJ2 with the incorrect IP address and change it back to the IP address stored in the configuration backup file.
2. Complete the management group restoration.
3. When the management group is successfully restored, change the IP address of the SSR316MJ2 to the desired address.

Best Practice to Prevent This Problem

1. In the Edit Configuration window, change the IP address of the SSR316MJ2.
2. In the Console Network View, select the management group.
3. Right-click and select Backup Configuration of Management Group.
4. In the Back up Configuration window that opens, scroll down to the section labeled SSR316MJ2s.
5. In the SSR316MJ2s section, verify that the Communication IP is the new IP address. If it is not the new IP address, then click OK to cancel out of the Backup Configuration window.

Wait for a few minutes and then repeat steps 3 through 5 . When the correct new IP address appears, select Back Up Configuration. The management group configuration is backed up with the correct IP address.

4.8 Clusters

4.8.1 If Incorrect Virtual IP Information Is Entered, SSR316MJ2s Go Offline And Volumes Become Unavailable

Issue

When configuring VIP for a cluster, entering incorrect information for any of the components (IP Address, Subnet Mask and Default Gateway) causes the SSR316MJ2s in that cluster to go down and any volumes associated with the cluster to become unavailable.

Workaround

The iSCSI VIP must be in the same subnet as all the SSR316MJ2s in the cluster.

1. Enter the correct information for the Virtual IP configuration.
2. Reboot the SSR316MJ2s in the cluster.

4.8.2 VIP Is Not Being Enforced for iSCSI Load Balancing

Issue

1. User can create an authentication group and enable load balancing without adding a VIP to the cluster.
2. User can remove a VIP form a cluster and leave load balancing enabled for authentication groups associated to volumes on that cluster.

.

4.9 Volumes & Snapshots

4.9.1 Snapshots Don't Autogrow When Restriping Exceeds Hard Threshold

Issue

Snapshots that are restriping reach the hard threshold on one snapshot. Autogrow does not take place and the entire restripe process is stopped.

Workaround

Use the Console to increase the hard threshold on the snapshot that has reached its hard threshold.

4.9.2 Snapshot Schedules Do Not Adjust For Daylight Savings Time

Issue

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 (under Standard Time) at 2:00 PM. Then, when the local time changes to Daylight Savings Time, the schedule starts running at 3:00 PM instead of 2:00 PM. This is happening 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. The Storage System Software does not include automatic adjustments 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 changes in the spring and autumn.

4.9.3 Frequent Re-synchronizing Of Snapshots Prevents Other Operations (6.5 w/SP1)

Issue

Re-synchronizing snapshots occupies system resources and may prevent other management group operations.

Best Practice

Create snapshot schedules with a frequency greater than one hour.

4.9.4 Creating A Volume With A Duplicate Name Does Not Give An Error

Issue

You create a volume in a cluster named Volume_0. You then create a second volume in that cluster named Volume_0. No error message is generated, the second volume does not appear, but some settings on the first volume are changed to those entered for the second volume.

Workaround

Do not use the same name for a volume more than once in a cluster.

4.9.5 Volumes Go Offline And The Microsoft* Initiator Cannot Find A Volume When Snapshot Hard Threshold Is Reached

Issue

Volume configured without Auto Grow reaches the hard threshold at the time a snapshot is taken. The top-level volume does not show a problem while the snapshot flashes red and the Microsoft* initiator cannot find the volume.

Workaround

Increase the hard threshold on the snapshot by as small an amount as possible. For example, if the hard threshold shows as 100 G, increase that amount to 100.01 G.

Solution

Configure auto grow on the volume before creating snapshots.

4.9.6 Cannot Retain Fewer Scheduled Remote Snapshots Than Primary Snapshots

Issue

Configure a remote snapshot schedule with a retention policy of 7 primary snapshots and 4 remote snapshots. The San retains 7 primary and 7 remote snapshots.

You cannot retain fewer remote snapshots than primary snapshots. However, you can retain more remote snapshots than primary snapshots

Workaround

None at this time.

4.9.7 Snapshot Schedules Do Not Adjust for Daylight Savings Time

Issue

Snapshot schedules that are created under standard time continue to execute at the originally scheduled Standard Time, even when time changes to daylight savings time.

Workaround

Manually change the time on the snapshot schedule.

4.10 Remote IP Copy

4.10.1 Remote IP Copy From Multiple Management Groups To A Single Remote Management Group Causes Performance Drop In Remote Management Group

Issue

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

Workaround

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

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

4.10.2 Remote Copy Schedules Failed To Recreate When Restoring A Remote Management Group.

Issue

Two management groups, A and B, had a remote copy schedule copying from A to B. Both management group configurations had been backed up.

1. Management group A goes down and is restored from the configuration backup.
2. Later, management group B goes down and is restored from backup but now the remote schedule is lost.

Workaround

After restoring management group A, back up both management group configurations again. Otherwise you must manually re-enter the remote copy schedule.

4.10.3 No Error Message Received When Remote Copy From Management Groups With Different Storage System Software Version Is Started (6.5 w/SP1)

Issue

Using management group ABC with version 6.5 and management group XYZ with version 6.3, using the Console, you attempt a remote copy operation from ABC to XYZ and the Console shows that the remote copy never starts and shows 0 progress.

Expectation

Console should not start the remote copy operation and display an error message that indicates this copy operation cannot be completed because of a version mismatch.

Fix

Upgrade the storage modules in the down-level management group. Once all storage modules are upgraded, the Console will auto-upgrade the management group version and the remote copy operation can be restarted.

To locate the management group version, click once on the management group name and select the Register tab.

Note: When Remote Copy is involved, please remember the following:

1. Storage System Software Remote Copy support copying between groups that have the same management group version, For instance, a group running Release 6.5 w/SP1 can copy to/from a group also running Release 6.5 w/SP1.
2. Storage System Software Remote Copy supports copying to up-level groups. For instance, a group running Release 6.3 can copy to a group running Release 6.5 w/SP1.
3. Storage System Software Remote Copy does not support copying to down-level groups. For instance, a group running Release 6.5 w/SP1 cannot copy to a group running Release 6.3.

4.11 iSCSI

4.11.1 iSCSI Disconnect/Reconnect May Take An Application Cluster Offline

Issue

Various maintenance activities can cause an iSCSI disconnect and reconnect on a client. Each time such a disconnect/reconnect occurs, it can take down an application cluster. Example activities that can cause an iSCSI disconnect/reconnect include the following:

- Starting or stopping a manager in a management group
- Swapping a hot spare in to or out of a cluster
- Migrating a volume between clusters

Workaround

Log off the volume from the client before performing maintenance activities.

4.11.2 Volume Migration Can Cause Application Cluster To Loose Disk(s)

Issue

Volume migration can cause an iSCSI disconnect and reconnect on a client. This disconnect/reconnect may interrupt client access to volume.

Workaround

- Log off the volume from the client before starting the volume migration in the Console.
- When the volume migration has started in the Console, log back in to the volume from the client.

4.11.3 iSCSI Closes All Shares After Reboot

Issue

If your iSCSI volumes are used by automatically-started Windows services (e.g., 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 Microsoft* support article 870964 on the Microsoft support web site.

Also, see the section entitled “Running automatic start services on iSCSI disks” in the Microsoft* iSCSI Initiator Users Guide for more details.

4.11.4 An iSCSI Volume That Becomes Unavailable For Approximately 60 Seconds Or Longer May Cause Data Loss

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

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

Solution

Change the Windows Registry setting for the default Maximum Request Hold Time to a very large (infinite) value.

Important: Back up your registry before making any changes.

1. Run regedit.exe.
2. Navigate to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\4D36E97B-E325-11CE-BFC1-08002BE10318\####\Parameters (where #### is the index of the Microsoft* iSCSI initiator in the set of SCSI and RAID Controllers).
3. Double-click MaxRequestHoldTime in the right-hand pane. The Edit DMWord Value window opens.
4. Change the Base to decimal.
5. Enter a value of 600.
6. Click OK.
7. Save your changes and exit the Registry.
8. Reboot the system.

4.11.5 When Mounting Existing iSCSI Volumes On Different Servers, Volumes May Be Assigned Duplicate Drive Letters Or No Drive Letters

Issue

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.

4.11.6 Linux-iSCSI Initiator Cannot Reboot When Storage System Software Volume Is Unavailable

The iSCSI Device Manager hangs when network problems prevent it from communicating with an SSR316MJ2. Because the default timeout for the Linux-iSCSI initiator is infinite, the initiator cannot reboot when it is unable to access the iSCSI volume on the SSR316MJ2.*Workaround*

Restore full network connectivity between iSCSI initiators and SSR316MJ2s. If this is not possible, disconnect the SSR316MJ2 that the initiator can't communicate with from the network. Disconnecting will cause the managers to tell the client that it should stop attempting to contact that SSR316MJ2.

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

Issue

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.

Solution

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

4.11.8 Microsoft* iSCSI Initiator Does Not Support Dynamic Disks

Issue

The Microsoft iSCSI initiator software does not support dynamic disks.

Workaround

Do not create dynamic disks to be used with the Microsoft iSCSI initiator.

4.11.9 iSCSI Volume Disappears From The iSCSI Initiator "Active Sessions" Window When Using Scheduled Snapshots

Issue

If you are using scheduled snapshots with an iSCSI volume, and the snapshot hard threshold is set to less than the volume hard threshold, the iSCSI volume disappears from the initiator Active Sessions window when the snapshot hard threshold is exceeded.

To recover from this situation:

1. In the Console, edit the snapshot schedule to increase the hard threshold.
2. Re-log in to the volume in the iSCSI initiator.

Workarounds

1. In the snapshot schedule, set the snapshot hard threshold to the same value as the volume thresholds, or
2. Use the auto_grow scripting feature to configure automatic threshold increases for the volume hard thresholds.

4.11.10 Failover Works As Long As A Virtual IP Address Is Used With iSCSI Initiators from Microsoft*, Intel®, Solaris*, Qlogic*, Adaptec*, Novell*, HP* and IBM*

Issue

To take advantage of the Storage System Software failover functionality in the iSCSI initiators from the listed companies, use a Virtual IP address when configuring clusters in the Console.

Workaround

Adaptec* Initiator

To ensure iSCSI volume availability in case of failover, it is recommended that the Session Recovery Timeout be set to 600 seconds. This is done using the Adaptec* iConfig utility.Qlogic*Initiator

To ensure iSCSI volume availability in case of failover, the following initiator configuration parameters must be set via the “Config Parameters” button on the Target Settings tab:

1. Default Timeout: 600 seconds
2. Connection Keep Alive Timeout: 600 seconds

Microsoft* Initiator

To ensure iSCSI volume availability in case of failover, the following registry key must be set:

```
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-
11CE-BFC1-08002BE10318}\0000\Parameters]
```

MaxRequestHoldTime: 600 seconds

Note: The “0000” before “Parameters” in the registry path might vary. It could be 0001, 0002, etc. Search for MaxRequestHoldTime to find the key.

4.11.11 Starting Or Stopping A Manager Causes iSCSI TO Disconnect And Reconnect

Issue

Starting or stopping a manager in a management group causes an iSCSI disconnect/reconnect that is visible in the Windows Event Log. Normally, the reconnect happens so quickly that there is no interruption in access to volumes in that cluster.

Workaround

If a disconnect occurs, reboot the client.

4.11.12 RedHat: Changing Authentication Type Causes Existing iSCSI Devices To Be Renamed

Issue

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)

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

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
   ```

Then change to use CHAP or whatever authentication you want to test next. Then restart things in the reverse order:

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

4.11.13 After Power Cycle, Load Balancing Does Not Distribute Requests Properly From A Microsoft Cluster (6.5 w/SP1)

Issue

A storage module is powered off and then powered on, and another storage module in the Storage system Software 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.

Fix

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 Storage System Software iSCSI disks.

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

4.11.14 When Using Storage System Software DSM for MPIO, If User Logs Off A Target Session While Client Is Accessing A Volume, The Path Disappears From MS Initiator (6.5 w/SP1)

Issue

In the MS iSCSI initiator, when you open the Details window from the Targets tab and click on the Connections button on the Sessions tab, the number of iSCSI sessions displayed for a given volume is fewer than expected.

Explanation

User should not try to manually log off additional iSCSI sessions that are generated by the LeftHand MPIO software, especially while clients are accessing the target volumes. If you do a manual logout, the session will usually be restored by the software within a minute or two. However, there have been a few cases where a replacement session does not appear. This may result in degraded performance for the volume.

Workaround

1. Quiesce client activity to the volume.
2. Completely log off the volume and reconnect.

4.11.15 iSCSI Load Balancing Does Not Properly Balance iSCSI Sessions When Running A Mix Of Servers With Storage System Software DSM for MPIO And Servers with iSCSI Load Balancing Enabled (6.5 w/SP1)

Issue

A mixture of servers with Storage System Software with DSM for MPIO (Server Group-1) and servers with iSCSI Load Balancing enabled (Server Group-2) are accessing volumes in a storage cluster, iSCSI sessions from Server Group-2 are not properly load-balanced. For example, in a cluster with Server Group -1 accessing 3 volumes and Server Group-2 accessing 3 other volumes, the 3 sessions for Server Group-2 all end up on a single storage module. This storage module can become a performance bottleneck.

Workaround

If practical, the problem can be avoided by partitioning the Management Group into storage clusters such that a given cluster is accessed by only DSM hosts or only non-DSM hosts.

This problem is not platform-specific.

4.11.16 iSCSI Load Balanced Connections Using Virtual IP (VIP) Are Not Re-assigned After Volume Is Migrated or Storage Module Removed From Cluster (6.5 w/SP1)

Issue

After migrating a volume or removing a storage module from a cluster, load-balanced iSCSI sessions are not re-assigned. While iSCSI connectivity is maintained throughout these operations, performance may not be optimal.

Workaround

1. Quiesce client activity to the volume.
2. Completely log off the volume and reconnect.

4.11.17 Failed Initiator Session Records Are Not Always Removed From Database

Issue

When a host disconnects ungracefully from a Storage System Software volume (e.g. the network fails, the host hardware is rebooted), the Console shows the iSCSI session for that host as Failed. If the host does not reestablish the connection within a day, the session is supposed to be considered dead and removed from the Console's iSCSI session displays. However, the failed sessions are usually not removed, resulting in lists of failed sessions displaying in the Console.

Workaround

None, However, as long as your volume is connected and accessible, and shows a connected session in the Console, you can ignore the failed session.

4.11.18 An Extra Microsoft iSCSI Session Is Created In The Console After Rebooting The Host

Issue

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 host-volume 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.

4.11.19 Microsoft iSCSI Initiator Stops With Error

Explanation

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.

4.12 Configuration Backup and Restore

4.12.1 If IP Address On SSR316MJ2 Is Changed Using the Configuration Interface, Some Processes Continue to Use The Old IP Address

Issue

An SSR316MJ2 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 Storage System 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 SSR316MJ2 with the new IP address.
2. In the SSR316MJ2 Configuration Interface, navigate to the TCP/IP Network category.

On the Communication tab, click Update to synchronize the IP addresses of all managers.

4.12.2 Single Disk Errors Are Not Recovered In Clusters With SSMs Running Mixed SAN Software Versions

Issue

Release 6.3 contains functionality to recover from any single disk unrecoverable data error. This recovery functionality only works on SSR316MJ2's in clusters where all SSR316MJ2's are upgraded to version 6.3. If a cluster has one or more SSR316MJ2's running an earlier version of the software, than the recovery functionality will not work.

Workaround

Upgrade all SSM's to release 6.3 SAN software.

Fix

None.

4.13 MSCS

4.13.1 Need Windows Patch If Running With VSS In A MCS Environment

Issue

In an active-passive Microsoft Cluster environment running on Windows 2003 SP1, the event viewer incorrectly reports messages like transaction log could not be flushed, data corruption could occur, delayed write failures etc. The application continues to run without any interruption, which confirms that the messages are incorrect.

Workaround

If you install VSS version 433 in a Microsoft Clustered Environment, you should also apply this Microsoft VSS patch: http://support.microsoft.com/?kbid=891957

4.13.2 Rebooting SSR316MJ2 That Are Hosting More Than 2 Microsoft Cluster Nodes May Cause Failover Errors

Issues

If you have more than 3-node clusters and your SSR316MJ2s have experienced multiple reboots and you see the following messages in the event logs, you need to first ensure that the storage modules are staying up, and then reboot the clustered hosts to recover from this situation.

Log files for NTFS show “delayed write failures”

Log files for ftdisk show “failing to flush the transaction log in the system event viewer”

Workaround

Upgrade to Release 6.6 when that becomes available.

4.13.3 MCS Cluster Failover While SSR316MJ2 Cluster Under Heavy Load Takes MCS Cluster Off-line

Issue

If an MCS cluster failover occurs while the SSR316MJ2 cluster is under very heavy load, the MCS cluster does not come back online until the load on the SSR316MJ2 cluster decreases.

Workaround

Increase the “pending timeout” of each of the disk resources on the MCS cluster to the same as the “maxrequestholdtime” of 600.

Do the following on each “physical disk” resource that is actually an iSCSI disk on the SSR316MJ2.

1. Right-click on the disk in the MCS cluster administrator.
2. Select Properties > Advanced tab.
3. Change the “pending timeout: seconds” from 180 to whatever you used as a “maxrequestholdtime” for iSCSI in the registry.

4.14 Dell Open Manage Secure Port Server

4.14.1 Unable To Install Or Load Console With Dell's Secure Port Server Service Started

Issue

Using Microsoft Windows* on a Dell* Server with the Dell* Open Manage Secure Port Server service, the user cannot properly install the Console or start the Console.

Workaround

Stop the Dell* Open Manage Secure Port Server service when installing or running the Console.

4.15 Reporting and SNMP

4.15.1 Fan Status Log Details Are Unclear

Issue

A SSR316MJ2 is in available mode with at least one fan not working. Fan status log details do not specify which fan is faulty.

Fix

Additional information about the fan status is in the Alerts.log file. This information includes the fan numbers of the faulty fans.

4.15.2 After Changing Host Name of A SSR316MJ2 in A Management Group, Alerts Use Old Host Name

Issue

Create a management group. At a later time, change the host name of the SSR316MJ2s in the management group. Alerts continue to use the old host names of the SSR316MJ2s, instead of the new names.

Workaround

Remove the SSR316MJ2 from the cluster and management group and re-add it. Note: Removing and re-adding storage modules to the management group will cause two restripes – one when the SSR316MJ2 is removed from the management group and the second when it is added back.

4.15.3 “NVRAM Card = Corrupt” Alert Generated When The SSR316MJ2 Is Restarted after Being Shut Down for Some Time

Workaround

Contact your vendor.

4.15.4 Monitoring Variables That Have Been Removed Cannot Be Re-Added in SSM 100s and SSM 200s

Issue

In the Active window of the Reporting category, a user removes a monitoring variable from the Monitored Variables list. If the user then tries to re-add the variable to the list using the add function, the following error occurs: 'Trigger value is null'.

Workaround

1. Log out of the SSM and log back in.
2. Navigate to the Active window of the Reporting category.
3. Try adding the variable again.

4.15.5 Reconfiguring RAID Disables SNMP

Workaround

Re-enable SNMP in the Console

4.16 Volume Lists and Authentication Groups

4.16.1 Volume Lists Must Contain Only Authentication Groups With Same Load Balancing Configurations (6.5 w/SP1)

Issue

If a volume list contains one authentication group with load balancing and one authentication group without load balancing, it may not be possible for both of two clustered iSCSI clients to connect to the volume at the same time.

Fix

Only add authentication groups with the same load balancing configurations to a volume list.

4.16.2 Enable or Disable Load Balancing On An Authentication Group Requires Logging Off And Re-logging On To Volume (6.5 w/SP1)

Issue

The user name changes the Enabling Load Balancing configuration of an authentication group. Afterwards, some iSCSI clients may not be able to reconnect to volumes because of the change

Workaround

Log off all iSCSI connections and log back on to reset the connections properly.

4.17 Fiber Channel (SSR316MJ2 Only)

4.17.1 If RAID Fails, LUNs Still Show As Available To Fiber Channel Network

Issue

If RAID fails, the SSR316MJ2 LUNs will appear available to the Fiber Channel SAN. The Client does not know that the LUNs are not available. This can cause a client to hang on startup, or its disk management window to hang.

Workaround

Pull the Fiber Channel link from the SSR316MJ2 with failing RAID until the client either boots or completes its disk management tasks.

4.17.2 Windows Server 2003 SP1 Systems Cannot See Full Capacity Of Volume When Volumes Exceed 2TB

Issue

The Storage System San software can configure volume larger than 2 TB. However, when operating in Windows Server 2003 SP1 with the FC protocol you will only see part of the volume, the portion greater than 2 TB. For example, if you create a 2.5 TB volume, only 0.5 TB can be seen in FC HBA utility. This issue only applies to FC SSR316MJ2s.

Workaround

Do not establish volumes larger than 2TB.