Hard disk replacement

This section provides instructions and information about replacing the hard disks in the following servers:

- Hitachi NAS Platform, Model 3090
- Hitachi NAS Platform, Model 3080
- NoteIn the remainder of this document, all server models are referred to as a "NAS server."

Intended Audience

These instructions are intended for Hitachi Vantara field personnel, and appropriately trained authorized third-party service providers. To perform this procedure, you must be able to:

- Use a terminal emulator to access the HNAS server CLI and Bali console.
- Log in to NAS Manager.
- Migrate EVSs.
- Physically remove and replace fan assemblies and hard disks.

NoteYou may also be required to upgrade the firmware. See Requirements for hard disk replacement for information about the minimum required firmware version.

Downtime considerations for hard disk replacement

Downtime is required because hard disk replacement is not a hot-swap operation. Replacing a hard disk requires that you shut down the server, disconnect the power cables from the Power Supply Units (PSUs), physically replace server parts, and start the process of rebuilding the server's internal RAID subsystem.

- Standalone server

  The complete disk replacement process requires approximately 2.5 hours, and the server will be offline during this time. You could restore services in approximately 1.5 hours by restoring services before the second disk of the server's RAID subsystem has completed synchronizing.
  Caution**Early service restoration is not recommended.** If the second disk of the internal RAID subsystem has not completed synchronizing, and there is a disk failure, you may lose data. Do not restore services before the RAID subsystem has been completely rebuilt unless the customer understands, and agrees to, the risks involved in an early restoration of services.
- Cluster node

  The complete disk replacement process requires approximately 2.5 hours for each node, and the node will be

offline during this time. You can, however, replace a node's internal hard disks with minimal service interruption for the customer by migrating file serving EVSs between nodes. Migrating EVSs allows the cluster to continue to serve data in a degraded state. Using EVS migration, each EVS will be migrated twice, once away from the node, and then to return the EVS to the node after hard disk replacement.

Requirements for hard disk replacement

Before replacing the hard disks, ensure that you have:

- Completed a disk health check. This health check should be performed at least one week in advance of the planned disk replacement. See Step1 Performing an Internal Drive Health Check for more information.
- The following tools and equipment:
  - #2 Phillips screwdriver.
  - A laptop that can be used to connect to the server's serial port. This laptop must have an SSH (Secure Shell) client or terminal emulator installed. The SSH client or terminal emulator must support the UTF-8 (Unicode) character encoding. See Accessing Linux on the server and node for more information.
  - A null modem cable.
  - An Ethernet cable.
  - Replacement hard disks.
  - Minimum firmware revision of 7.0.2050.17E2:

    If the system firmware version is older than 7.0.2050.17E2, update it to the latest mandatory or recommended firmware level before beginning the hard disk replacement procedure. Refer to the *Server and Cluster Administration Guide* for more information on upgrading firmware.
- The password for the "manager," "supervisor," and "root" user accounts on the server with the hard disks to be replaced.
- A maintenance period as described in Downtime considerations for hard disk replacement.
- Access to the Linux operating system of the server/node. See Accessing Linux on the server and node for more information.

Overview of the Procedure

This section provides a high-level overview of the hard disk replacement process. See the sections referenced in each step for detailed instructions.
NoteApproximately one week before starting this disk replacement, perform the disk health check. See "Step 1: Performing an Internal Drive Health Check" on page 55 for more information.

The hard disk replacement process is as follows:

Procedure

1. Perform a health check: See Step1 Performing an Internal Drive Health Check for more information.

2. Gather and record IP address and disk status information about the server: See Step 2 Gathering information about the server or node.
3. Back up the server's configuration: See Step 3 Backing up the server configuration.
4. Physically locate the server: See Step 4 Locating the server.
5. For cluster nodes, save the preferred mapping, and migrate EVSs to a different node in the cluster: See Step 5 Save the preferred mapping and migrate EVSs (cluster node only).
6. Physically replace the first disk: See Step 6 Replacing a Server's Internal Hard Disk.
7. Synchronize the first new disk and the existing disk: See Step 7 Synchronizing server's new disk.
8. Physically replace the server's second hard disk: See Step 8 Replacing the server's second disk.
9. Synchronize the second new disk and the first new disk: See Step 9 Synchronizing the second new disk.
10. For cluster nodes, restore migrated EVSs to their preferred node: See Step 10 Restore EVSs (cluster node only).

    When performing parts of the disk replacement process, you must access the Linux operating system and/or the Bali console of the NAS server/node. Instructions on how to access these components are provided in Accessing Linux on the server and node

Accessing Linux on the server and node

To run some of the commands, you must access the Linux layer of the NAS server or node using one of two methods:

- The serial (console) port, located on the rear panel of the server. See Using the Serial (Console) Port for more information.
- SSH connection. See Using SSH for an Internal SMU or Using SSH for an External SMU,

Using the Serial (Console) Port

Use the terminal emulator and null modem cable to access the NAS server's Linux operating system.

Procedure

1. Configure the terminal emulator as follows:
    - Speed: 115200
    - Data bits: 8 bits
    - Parity: None
    - Stop bits: 1
    - Flow control: No flow control

NoteTo increase readability of text when connected, set your terminal emulator to display 132 columns.

2. Log in as '**root**.'
3. Connect to localhost using the SSC (server control) utility to run the Bali commands by entering the command:
   `ssc localhost`

Using SSH for an embedded SMU

These instructions apply if you have an embedded SMU. If you have an external SMU, see Using SSH for an External SMU.

Procedure

1. Use SSH to log in to the embedded SMU as 'manager.' Enter the following command:
   `ssh manager@`[*IP Address*]

   where [*IP Address*] is the IP address of the NAS server administrative service EVS.
2. Enter the password for the 'manager' user account.

   This logs you into the Bali console.
3. Access the Linux prompt by exiting the Bali console. Enter the following command:
   `exit`

   or press the Ctrl+D keys.
4. Log in as the 'root' user. Enter the following command:
   `su -;` [*password*]

   where [*password*] is the password for the root user account.

Using SSH for an External SMU

These instructions apply if you have an external SMU. If you have an internal SMU, see Using SSH for an Internal SMU.

Procedure

1. SSH into the external SMU as manager. Enter the following command:

   `ssh manager@`[*IP Address*]

   where [*IP Address*] is the IP address of the NAS server/node.

   This logs you into the siconsole.

2. Select the system (the server or the cluster node) that has the hard disks to be replaced.
   This logs you into the Bali console.
3. Synchronous Disaster Recovery Cluster the cluster node IP addresses. Enter the following command:
   `ipaddr`
4. Record the cluster IP addresses.
5. Access the Linux prompt by exiting the Bali console. Enter the following command:
   `exit`

   or press the Ctrl+D keys.

   This logs you into the siconsole.
6. Quit to the SMU's Linux prompt. Enter the following command:
   `q`
7. Access cluster IP address using SSH and logging in as the 'supervisor' user. Enter the following command:
   `ssh supervisor@`[*Cluster_IP_Address*]

   where [*Cluster_IP_Address*] is the IP address of the NAS server/node.
8. Enter the password for the 'supervisor' user. By default, the password for the 'supervisor' user account is the "supervisor," but this may have been changed.
9. Log in as the 'root' user. Enter the following command:
   su -; [*password*]

   where [*password*] is the password for the root user account.

   You are now at the Linux prompt.

Step1: Performing an Internal Drive Health Check

The health check evaluates both internal disks to determine if there are any pending disk failures. Perform the health check twice:

- Approximately one week before hard disk replacement to allow time to resolve any errors before running the disk replacement procedure.
- When you start the hard disk replacement procedure to make sure the disks are ready for the replacement.

  The health check includes retrieving and evaluating the disk's SMART (Self-Monitoring, Analysis, and Reporting Technology) information and reviewing the server's internal RAID subsystem status.

  If you find errors on either of the two disks, note the disk and make sure that the disk with the errors is the first one to be replaced. If both disks have errors, contact technical support and escalate the errors based on the health

check output.

To run the health check:

Procedure

1. Log in to each node/server using the SSH process, which is described in Accessing Linux on the server and node.

2. Verify the mapping of physical disks to SCSI devices.
   To display the mapping between the physical drive and the dev/sdX name, there are symlinks displayed by the output from the `/ls -l /dev/disk/by-path` command.

   In the example below, the portion of the output that displays the mapping between the SATA port and the SCSI device number is underlined. This example shows the standard post boot situation, where SATA port 0 (Physical Drive A) is `/dev/sda` and port 2 (Physical Drive B) is `/dev/sdb`.

   ```
   mercury100:~$ ls -l /dev/disk/by-path
   total 0
   lrwxrwxrwx 1 root root  9 2011-06-27 12:17 pci-0000:00:1f.2-scsi-0:0:0:0 -> ../../sda
   lrwxrwxrwx 1 root root 10 2011-06-27 12:17 pci-0000:00:1f.2-scsi-0:0:0:0-part1 -> ../../sda1
   lrwxrwxrwx 1 root root 10 2011-06-27 12:17 pci-0000:00:1f.2-scsi-0:0:0:0-part2 -> ../../sda2
   lrwxrwxrwx 1 root root 10 2011-06-27 12:17 pci-0000:00:1f.2-scsi-0:0:0:0-part3 -> ../../sda3
   lrwxrwxrwx 1 root root 10 2011-06-27 12:17 pci-0000:00:1f.2-scsi-0:0:0:0-part5 -> ../../sda5
   lrwxrwxrwx 1 root root 10 2011-06-27 12:17 pci-0000:00:1f.2-scsi-0:0:0:0-part6 -> ../../sda6
   lrwxrwxrwx 1 root root  9 2011-06-27 12:17 pci-0000:00:1f.2-scsi-2:0:0:0 -> ../../sdb
   lrwxrwxrwx 1 root root 10 2011-06-27 12:17 pci-0000:00:1f.2-scsi-2:0:0:0-part1 -> ../../sdb1
   lrwxrwxrwx 1 root root 10 2011-06-27 12:17 pci-0000:00:1f.2-scsi-2:0:0:0-part2 -> ../../sdb2
   lrwxrwxrwx 1 root root 10 2011-06-27 12:17 pci-0000:00:1f.2-scsi-2:0:0:0-part3 -> ../../sdb3
   lrwxrwxrwx 1 root root 10 2011-06-27 12:17 pci-0000:00:1f.2-scsi-2:0:0:0-part5 -> ../../sdb5
   lrwxrwxrwx 1 root root 10 2011-06-27 12:17 pci-0000:00:1f.2-scsi-2:0:0:0-part6 -> ../../sdb6
   mercury100:~$
   ```

3. Retrieve the SMART data for each of the internal disks by entering the following commands:
   - For disk A: smartctl –a /dev/sda
   - For disk B: smartctl –a /dev/sdb

4. Review the Information section of the retrieved data to verify that the SMART support is available and enabled on both disks.

   In the sample output from the smartctl command below, the portion of the information that indicates SMART support is underlined:

   ```
   === START OF INFORMATION SECTION ===
   Device Model:     ST9250610NS
   Serial Number:    9XE00JL3
   Firmware Version: SN01
   User Capacity:    250,059,350,016 bytes
   Device is:        Not in smartctl database [for details use: -P showall]
   ATA Version is:   8
   ATA Standard is:  ATA-8-ACS revision 4
   ```

```
Local Time is:    Thu Mar  3 12:48:44 2011 PST
SMART support is: Available - device has SMART capability.
SMART support is: Enabled
```

5. Scroll past the Read SMART Data section, which looks similar to the following example.

```
=== START OF READ SMART DATA SECTION ===
SMART overall-health self-assessment test result: PASSED

General SMART Values:
Offline data collection status:  (0x82) Offline data collection activity
                                        was completed without error.
                                        Auto Offline Data Collection: Enabled.
Self-test execution status:      (   0) The previous self-test routine completed
                                        without error or no self-test has been run.
Total time to complete Offline
data collection:                 ( 634) seconds.
Offline data collection
capabilities:                    (0x7b) SMART execute Offline immediate.
                                        Auto Offline data collection on/off
                                        support.
                                        Suspend Offline collection upon new
                                        command.
                                        Offline surface scan supported.
                                        Self-test supported.
                                        Conveyance Self-test supported.
                                        Selective Self-test supported.
SMART capabilities:              (0x0003) Saves SMART data before entering
                                        power-saving mode.
                                        Supports SMART auto save timer.
Error logging capability:        (0x01) Error logging supported.
                                        General Purpose Logging supported.
Short self-test routine
recommended polling time:        (   1) minutes.
Extended self-test routine
recommended polling time:        (  49) minutes.
Conveyance self-test routine
recommended polling time:        (   2) minutes.
SCT capabilities:              (0x10bd) SCT Status supported.
                                        SCT Feature Control supported.
                                        SCT Data Table supported.
```

6. Review the SMART Attributes Data section of the retrieved data to verify that there are no “Current_Pending_Sector” or “Offline_Uncorrectable” events on either drive.

   In the sample output from the smartctl command below, the portion of the information that indicates “Current_Pending_Sector” or “Offline_Uncorrectable” events is underlined:

```
SMART Attributes Data Structure revision number: 10
Vendor Specific SMART Attributes with Thresholds:
ID# ATTRIBUTE_NAME          FLAG     VALUE WORST THRESH TYPE      UPDATED  WHEN_FAILED
RAW_VALUE
  1 Raw_Read_Error_Rate     0x000f   080   064   044    Pre-fail  Always   -   102792136
```

```
  3 Spin_Up_Time            0x0003   096   096   000    Pre-fail  Always       -       0
  4 Start_Stop_Count        0x0032   100   100   020    Old_age   Always       -       13
  5 Reallocated_Sector_Ct   0x0033   100   100   036    Pre-fail  Always       -       0
  7 Seek_Error_Rate         0x000f   065   060   030    Pre-fail  Always       - 3326385
  9 Power_On_Hours          0x0032   100   100   000    Old_age   Always       -       156
 10 Spin_Retry_Count        0x0013   100   100   097    Pre-fail  Always       -       0
 12 Power_Cycle_Count       0x0032   100   100   020    Old_age   Always       -       13
184 Unknown_Attribute       0x0032   100   100   099    Old_age   Always       -       0
187 Reported_Uncorrect      0x0032   100   100   000    Old_age   Always       -       0
188 Unknown_Attribute       0x0032   100   100   000    Old_age   Always       -       0
189 High_Fly_Writes         0x003a   100   100   000    Old_age   Always       -       0
190 Airflow_Temperature_Cel 0x0022   074   048   045    Old_age   Always       - 26 (Lifetime Min/Max 25/
27)
191 G-Sense_Error_Rate      0x0032   100   100   000    Old_age   Always       -       0
192 Power-Off_Retract_Count 0x0032   100   100   000    Old_age   Always       -       12
193 Load_Cycle_Count        0x0032   100   100   000    Old_age   Always       -       13
194 Temperature_Celsius     0x0022   026   052   000    Old_age   Always       -       26 (0 20 0 0)
195 Hardware_ECC_Recovered  0x001a   116   100   000    Old_age   Always       -       102792136
197 Current_Pending_Sector  0x0012   100   100   000    Old_age   Always       -       0
198 Offline_Uncorrectable   0x0010   100   100   000    Old_age   Offline      -       0
199 UDMA_CRC_Error_Count    0x003e   200   200   000    Old_age   Always       -       0
```

If the RAW_VALUE for "Current_Pending_Sector" or "Offline_Uncorrectable" events are more than zero, this indicates that those events have been detected, and that the drive may be failing.

7. Check the SMART Error log for any events.
   In the sample output from the `smartctl` command below, the portion of the information that indicates SMART Error Log events is underlined:

```
SMART Error Log Version: 1
No Errors Logged
```

8. Validate all self test short and extended tests have passed.
   In the sample output from the `smartctl` command, the portion of the information that indicates SMART Self-test log events is underlined:

```
SMART Self-test log structure revision number 1
Num  Test_Description    Status                  Remaining  LifeTime(hours)  LBA_of_first_error
# 1  Short offline       Completed without error       00%       143         -
# 2  Short offline       Completed without error       00%       119         -
# 3  Short offline       Completed without error       00%        94         -
# 4  Short offline       Completed without error       00%        70         -
# 5  Extended offline    Completed without error       00%        46         -
# 6  Short offline       Completed without error       00%        21
```

If you find that one disk has no errors, but the other disk does have errors, replace the disk **with** errors first. If you find errors on both disks, contact technical support and provide them with the `smartctl` output.

9. Perform the RAID subsystem health check to review the current status of the RAID subsystem synchronization.
   Enter the following command:
   `cat /proc/mdstat outout`

```
Group5-node1:~# cat /proc/mdstat
Personalities : [raid1]
md1 : active raid1 sda6[0] sdb6[1] <-- Shows disk and partition (volume) status
      55841792 blocks [2/2] [UU] <-- [UU] = Up/Up and [U_] = Up/Down
      bitmap: 1/1 pages [4KB], 65536KB chunk

md0 : active raid1 sda5[0] sdb5[1]
      7823552 blocks [2/2] [UU]
      bitmap: 1/1 pages [4KB], 65536KB chunk

md2 : active raid1 sda3[0] sdb3[1]
      7823552 blocks [2/2] [UU]
      bitmap: 0/1 pages [0KB], 65536KB chunk

unused devices: <none>
Group5-node1:~#
```

Step 2: Gathering information about the server or node

Before shutting down the server/node to replace disks, you must gather and record information about the related IP addresses and check the status and synchronization of the devices. To obtain this information:

Procedure

1. Log in to the Bali console. See Accessing Linux on the server and node.
2. Select the server or node that has the disks you want to replace.

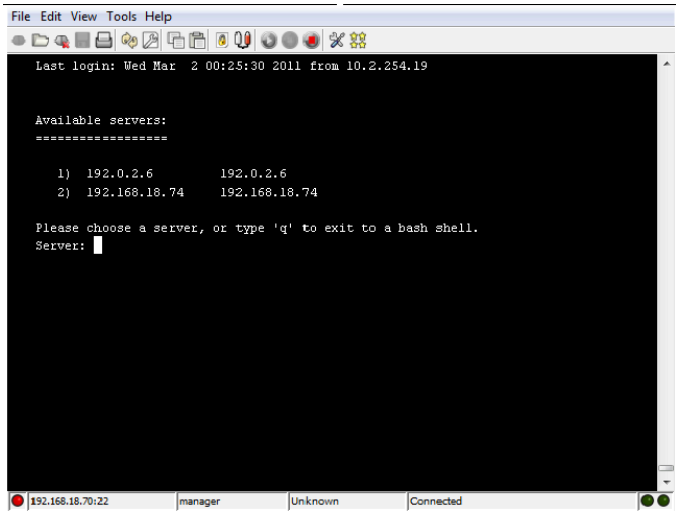

3. Record the IP Address of the system you choose.
4. Run the `evs list` command.
   - For a single-node cluster or a standalone server, record the administrative services EVS IP address.
   - For a multi-node cluster, record all cluster node IP addresses.

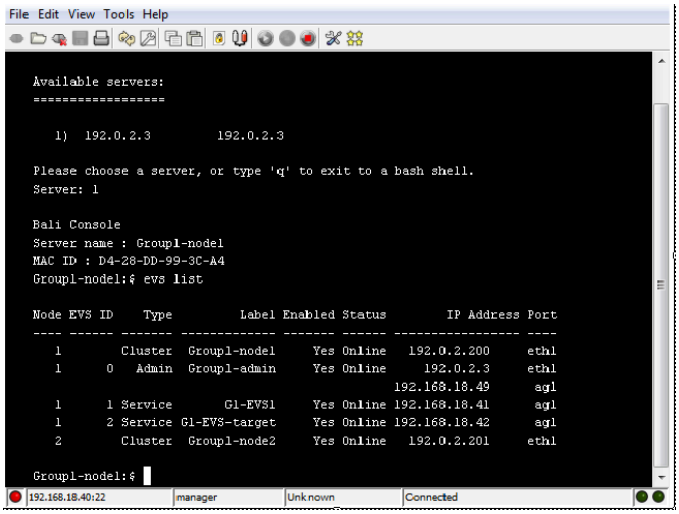

5. Run the `chassis-drive-status` command

6. Review the values in the **Status** and **% Rebuild** columns for each device.
   The response to the command should be similar to the following:

```
Device  Status       % Used  Size (4k blks) Used (4k blks) % Rebuild
------  ------------ ------  -------------- -------------- ------------
0       Good         32      3846436        1266962        Synchronized
1       Good         3       12302144       463572         Synchronized
2       Good         0       0              0              Synchronized
Success
```

   For each device, the **Status** should be “Good” and the **%Rebuild** should be “Synchronized.”

   - If the values are correct, repeat the health check, as described in Step1 Performing an Internal Drive Health Check.
   - If the values are not correct, run the `trouble chassis-drive` command. If the command response displays “No faults found,” repeat the health check, as described in Step1 Performing an Internal Drive Health Check. If the command response displays issues, resolve them if possible, or contact technical support for assistance.

Step 3: Backing up the server configuration

Backing up the server’s configuration for an internal or external SMU saves the server’s configuration, including the SI configuration. When backing up a server with an internal SMU, the configuration backup also includes a ZIP file of the SMU configuration.

Procedure

1. Connect your laptop to the management Ethernet switch using an Ethernet cable.

2. Log in to NAS Manager.

3. Navigate to Home Server Settings Configuration Backup & Restore.

4. Click backup to save the configuration file to your laptop.

5. Verify that the backup file is complete and make sure the file size is not 0 bytes

Step 4: Locating the server

Before shutting down the server/node to replace disks, you must physically locate the server.

Procedure

1. Run the led-identify-node X command.
   where *X* is the number of cluster node (the pnode-id) to identify.

   The result of this command is that the server's fault and power LEDs (located on the left side of the server's rear panel) flash simultaneously.

   

2. Physically locate the server that has the disks to be replaced. After you have identified the server, press any key to stop the LEDs from flashing.

Step 5: Save the preferred mapping and migrate EVSs (cluster node only)

If replacing the hard disks in a standalone server, skip this step. If replacing the hard disks in a cluster node, before shutting down the node to replace disks, migrate the EVSs to another node. You can migrate an individual EVS to a different node within the same cluster, or you can migrate all EVSs to another server or another cluster.

The current mapping of EVSs to cluster nodes can be preserved, and the saved map is called a preferred mapping.

Saving the current EVS-to-cluster configuration as the preferred mapping helps when restoring EVSs to cluster nodes. For example, if a failed cluster node is being restored, the preferred mapping can be used to restore the original cluster configuration.

Procedure

1. Connect your laptop to the customer's network.
2. Using a browser, go to http://[*SMU_IP_Address*]/
   where [*SMU_IP_Address*] is the IP address of the SMU (System Management Unit) managing the cluster
3. Log into NAS Manager as the "manager" user.
4. Navigate to Home Server Settings EVS Migration to display the EVS Migration page.
   Note If the SMU is currently managing a cluster and at least one other cluster or standalone server, the following page appears:

   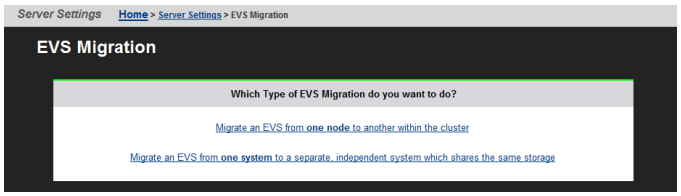

   If this page does appear, click Migrate an EVS from one node to another within the cluster to display the main EVS Migration page.

   If the SMU is managing one cluster and no standalone servers, the main EVS Migration page appears:

   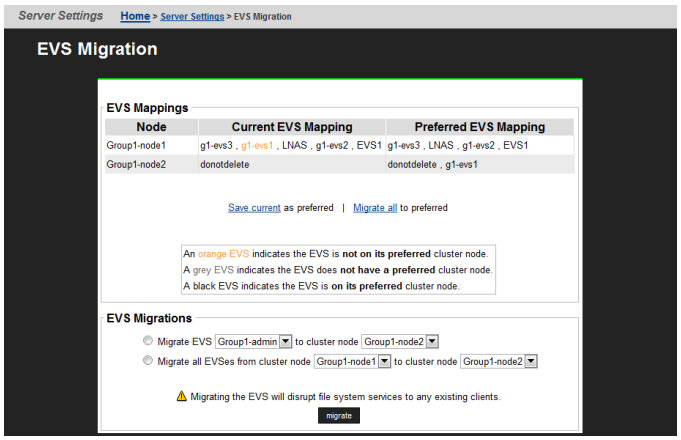

5. Migrate the EVSs between the cluster nodes until the preferred mapping has been defined. The current mapping is displayed in the Current EVS Mappings column of the EVS Mappings section of the page.
6. Save the current EVS-to-cluster node mapping by clicking Save current as preferred in the EVS Mappings section.
7. Migrate EVSs as required:
   - To migrate all EVSs between cluster nodes:

1. Select Migrate all EVS from cluster node ___ to cluster node ___.
2. From the first drop-down list, select the cluster node from which to migrate all EVS.
3. From the second drop-down list, select the cluster node to which the EVSs will be migrated.
4. Click Migrate.

- To migrate a single EVS to a cluster node:

1. Select Migrate EVS ____ to cluster node ___.
2. From the first drop-down list, select the cluster node to migrate.
3. From the second drop-down list, select the cluster node to which the EVS will be migrated.
4. Click Migrate.

Step 6: Replacing a Server's Internal Hard Disk

Because physically replacing hard disks is not a hot-swap operation, you must shut down the server and disconnect the power cables from the PSUs before beginning physical replacement.

Procedure

1. Shut down the server.
   Using NAS Manager, navigate to the Server Settings page, and:
   - For a cluster node, navigate to Home Restart, Reboot or Shutdown Server Shutdown.

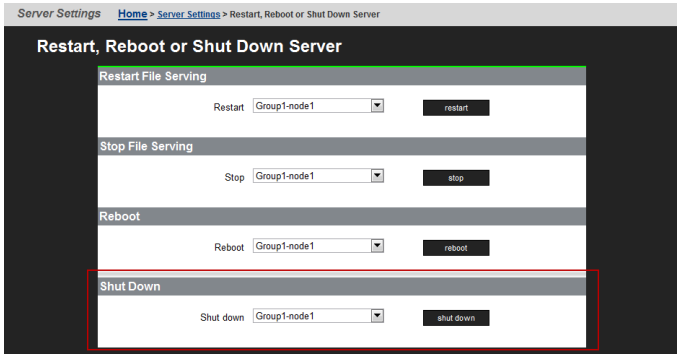

   - For a standalone server, navigate to Home Restart, Reboot or Shutdown Server Shutdown.
   - Using the CLI, shut down the server using the following command:

     shutdown –-powerdown –-ship -f

2. Wait for the status LEDs on the rear panel of the server to stop flashing, which may take up to five (5) minutes.
   If the LEDs do not stop flashing after five minutes, make sure the Linux operating system has shut down by looking at your terminal emulator program. If Linux has not shut down, enter the shutdown now command.

3. Remove the power cables from the PSUs.
4. Remove the fascia. See Bezel removal for details.
5. Remove the fan.

   Typically, hard disk “B” is replaced before hard disk “A.” Hard disk “B” is behind fan assembly number 2 (the center fan), Hard disk “A” is behind fan assembly number 1 (the left fan).
   CautionAfter one hard disk is replaced, you must restart the server and resynchronize its internal RAID subsystem before replacing the second hard disk. See Step 7 Synchronizing server’s new disk for more information.
6. Disconnect the fan power connector by pressing down on the connector’s retention latch and gently pulling the connector apart.

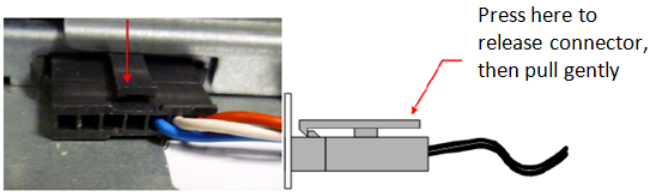

7. Remove the upper and lower fan retention brackets.

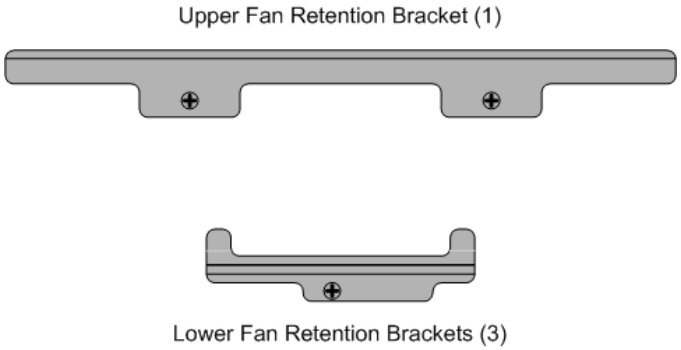

- When replacing hard disk B, remove the upper fan retention bracket and the lower fan retention bracket under fan assembly 2 (the center fan assembly).
- When replacing hard disk A, remove the upper fan retention bracket and the lower fan assembly bracket under fan assembly 1 (the left fan assembly).

8. Remove the fan assembly covering the disk you want to replace.

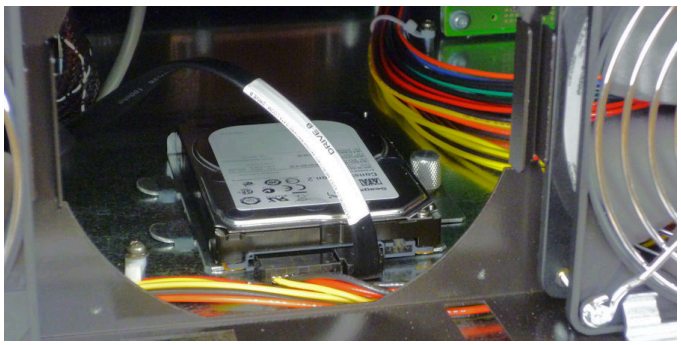

When replacing hard disk B, remove fan assembly 2 (the center fan assembly). Hard disk B should now be visible.

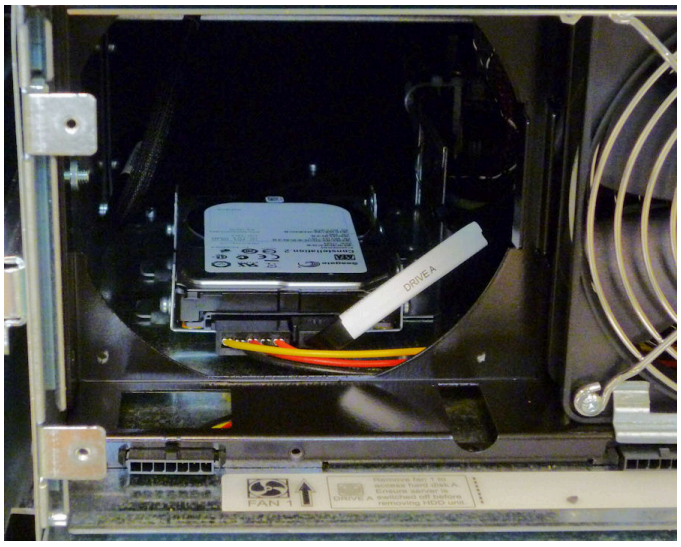

The hard disk is in a carrier (bracket) held to the bottom of the chassis by a thumbscrew on the right side and tabs that fit into slots on the chassis floor on the left side.

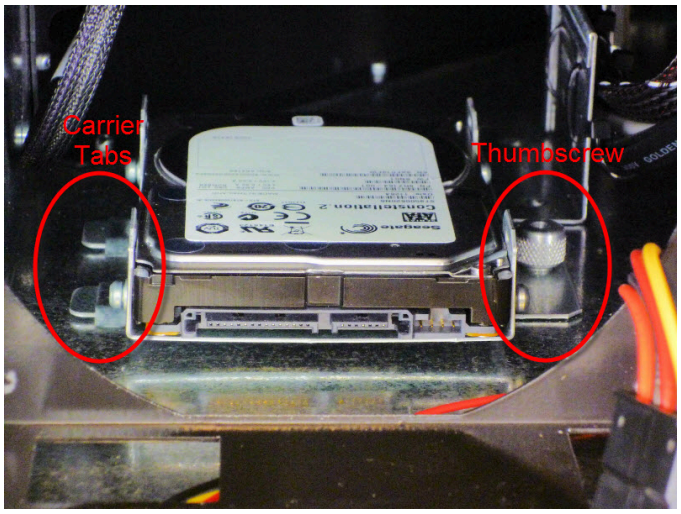

NoteThe carrier used for replacement hard disks may be different than the carrier holding the old hard disks. The new carriers fit into the same place and in the same way as the older carriers.

- Old carrier: the hard disk is mounted through tabs on the sides of the carrier.

- New carrier: the hard disk is mounted through the bottom plate of the carrier.

9. Disconnect the power and SATA cables from the hard disk.

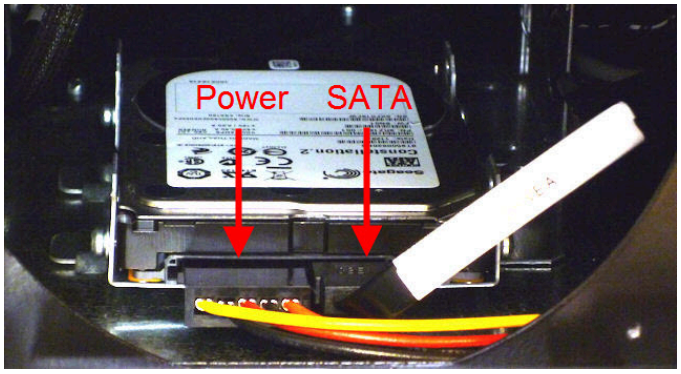

10. Loosen the thumbscrew on the right side of the hard disk carrier. Note that the thumbscrew cannot be removed from the carrier.

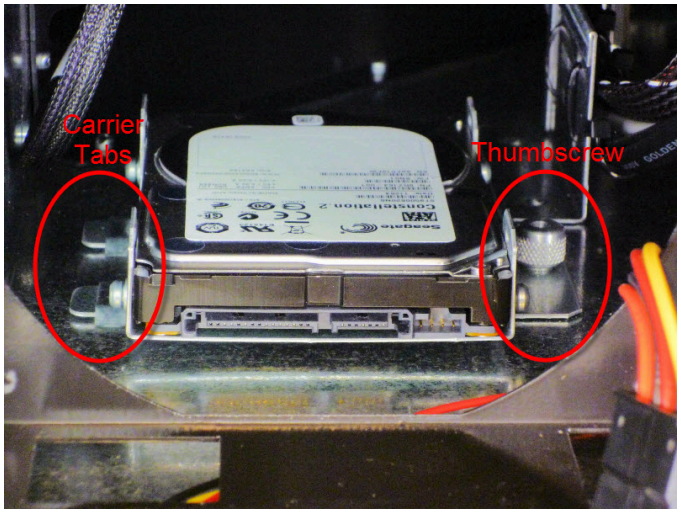

11. Gently lift the right side of the hard disk carrier and slide it to the right to disengage the tabs on the left side of the carrier.

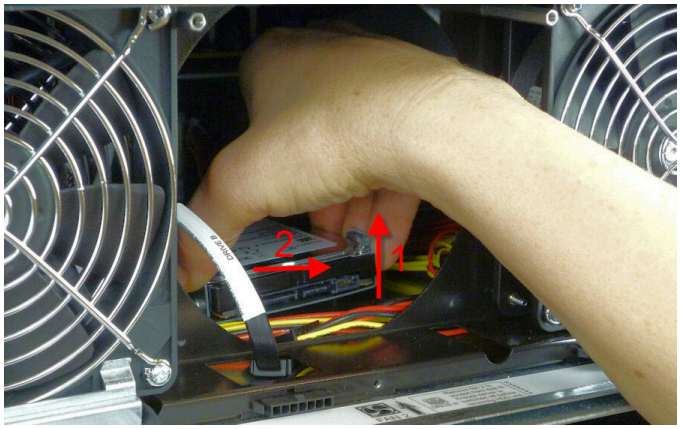

12. Once the disk carrier is completely disengaged from the chassis, remove it from the server, label it appropriately (for example, “server X, disk A”), and store it in a safe location.

13. To install the replacement hard disk, lift the right side of the carrier until you can insert the tabs on the left side of

the disk carrier into the slots on the floor of the server chassis.

14. Move the carrier to the left until the ends of the tabs are visible and the thumbscrew is aligned to fit down onto the threaded stud.

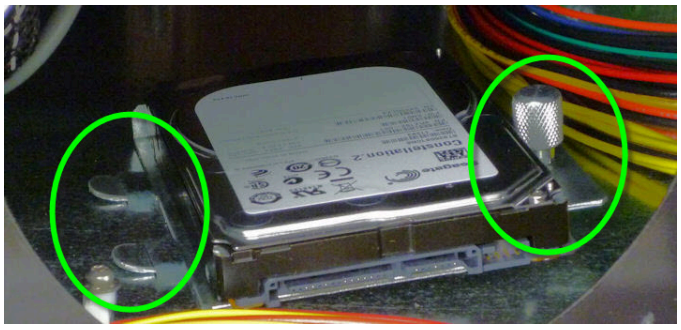

15. Tighten the thumbscrew to secure the disk carrier. Do not over tighten the thumbscrew.

16. Connect the power and SATA cables to the replacement hard disk.

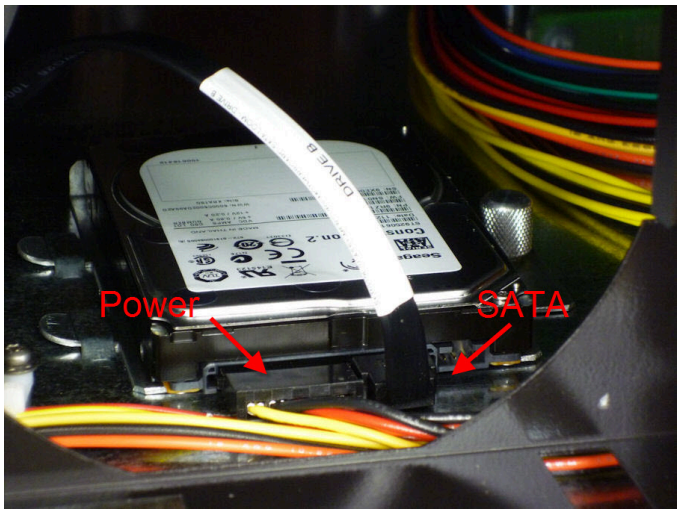

17. Reinstall the fan in the mounting slot, with the cable routed through the chassis cut-out.

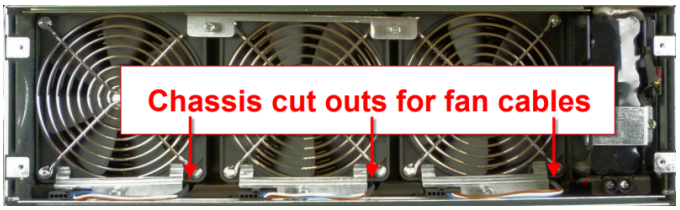

18. Reinstall the fan retention brackets. Do not over tighten the screws.
19. Reconnect the fan cable.
20. If you replaced only the first hard disk, continue with the next step. If you have replaced both disks, reinstall the fascia.
21. Reconnect the power cables to the PSUs.
    When the server starts, the LEDs on the front of the server flash quickly, indicating that the server is starting up.

Step 7: Synchronizing server's new disk

After replacing a hard disk, the new disk in the server's internal RAID subsystem must be synchronized with the older disk.

Procedure

1. Wait until the LEDs on the front of the server slow to indicate normal activity.
2. Use a serial cable connected to the serial (console) port of the server to access the Bali console.
3. Once you have successfully logged in, select the server or node that has the disks you want to synchronize.
4. Run the chassis-drive-status command, and look at the values in the **Status** and **% Rebuild** columns for each

device.

- The values in the **Status** column should be “Invalid.”
- The**% Rebuild** column should not display any values.

5. Run the script /opt/raid-monitor/bin/recover-replaced-drive.sh. This script partitions the replacement disk appropriately, updates the server’s internal RAID configuration, and initiates rebuilding the replaced disk.
   The RAID system rebuilds the disk as a background operation, which takes approximately 50 minutes to complete. Events are logged as the RAID partitions rebuild and become fully fault tolerant.
6. Monitor the rebuilding process by running the chassis-drive-status command, and check the values in the **Status** column for each device. The values in the **Status** column should be:
   - “Good” for synchronized volumes.
   - “Rebuilding” for the volume currently being synchronized.
   - “Degraded” for any volume(s) that have not yet started the synchronization process.
7. Once the rebuild process has successfully completed, run the trouble chassis-drive command.

   If the command response displays issues, resolve them if possible, or contact technical support for assistance.

   If the command response displays “No faults found,” continue the disk replacement process by replacing the second hard disk.
8. Shut down the server.

Step 8: Replacing the server’s second disk

Once the server’s first hard disk has been replaced and synchronized, replace the second disk. Refer to Step 6 Replacing a Server’s Internal Hard Disk for the steps required to replace the server’s second hard disk.

Step 9: Synchronizing the second new disk

Once the server’s second hard disk has been replaced, synchronize the server’s second hard disk to restore the integrity of the server’s internal RAID subsystem. Refer to Step 7 Synchronizing server’s new disk for the steps required to synchronize the server’s second hard disk.

Once the second hard disk is synchronized, log out by entering the exit command or pressing the Ctrl+D keys.

Step 10: Restore EVSs (cluster node only)

If replacing the hard disks in a standalone server, skip this step. If replacing the hard disks in a cluster node, return each of the EVSs to its preferred node (the node with the replaced disks).

The preferred mapping of EVSs to cluster nodes should have been saved in Step 5 Save the preferred mapping and

migrate EVSs (cluster node only). To return each EVSs to its preferred node using the preferred mapping:

Procedure

1. Connect your laptop to the customer's network.
2. Using a browser, go to http://[*SMU_IP_Address*]/
   where [*SMU_IP_Address*] is the IP address of the SMU (System Management Unit) managing the cluster
3. Log into NAS Manager as the "manager" user.
4. Navigate to Home Server Settings EVS Migration to display the EVS Migration page.
   Note If the SMU is currently managing a cluster and at least one other cluster or standalone server, the following page appears:

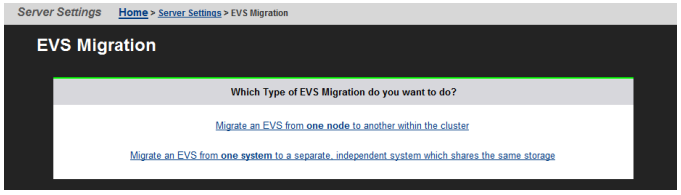

   If this page does appear, click Migrate an EVS from one node to another within the cluster to display the main EVS Migration page.

   If the SMU is managing one cluster and no standalone servers, the main EVS Migration page appears:

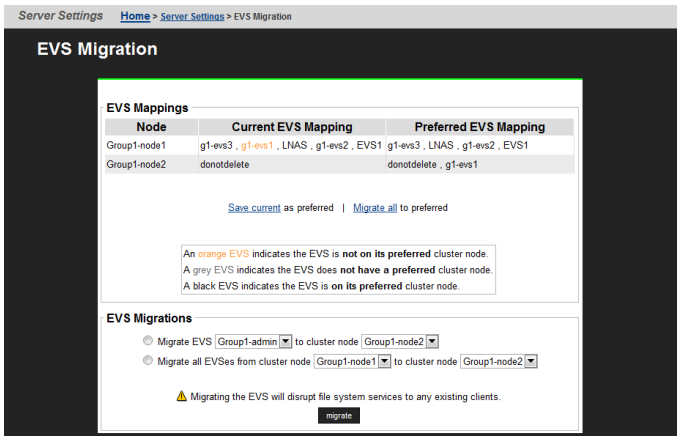

5. To return all EVSs to their preferred nodes:
   - If the preferred mapping was saved in Step 5 Save the preferred mapping and migrate EVSs (cluster node only), click Migrate all to preferred in the EVS Mappings section.
   - If the preferred mapping was not saved, migrate EVSs as required:
6. Migrate EVSs as required:
   - To migrate all EVSs between cluster nodes:

   1. Select Migrate all EVS from cluster node ___ to cluster node ___.

2. From the first drop-down list, select the cluster node from which to migrate all EVS.
3. From the second drop-down list, select the cluster node to which the EVSs will be migrated.
4. Click Migrate.

◦ To migrate a single EVS to a cluster node:

1. Select Migrate EVS ____ to cluster node ___.
2. From the first drop-down list, select the cluster node to migrate.
3. From the second drop-down list, select the cluster node to which the EVS will be migrated.
4. Click Migrate.