

# **Dell Acceleration Appliance for Databases 2.0**

## **Field Replacement Unit Guide**



© 2009-2015 Dell Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. Dell™ and the Dell logo are trademarks of Dell Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

2015 - 06

Rev. A04

# Contents

About this guide .....	5
Typographical conventions .....	5
<b>1 Software installation .....</b>	<b>7</b>
<b>2 DAAD drive and fabric card components .....</b>	<b>9</b>
Configuration guidelines — Hardware replacement .....	9
Replacing a drive in RAID 10 (reliable performance) .....	10
Replacing a drive in RAID 0 (maximum performance) .....	11
Standalone mode .....	11
HA mode .....	12
Replacing a drive in RAID 5 (reliable capacity) .....	12
Replacing a drive in a direct access (JBOD) pool .....	12
HA mode .....	13
Replacing a Fibre Channel HBA .....	13
Standalone mode .....	13
HA mode .....	14
Replacing a cluster interconnect card (HA) .....	14
Replacing an InfiniBand/SRP card .....	15
Case 1: IB and cluster interconnect .....	15
Case 2: Both ports configured as IB .....	15
Replacing an iSCSI card .....	16
Standalone mode .....	16
HA mode .....	17
Replacing a split-function-port iSCSI card .....	17
Correct iSCSI target entry after a replacement .....	18
<b>3 Replacing the DAAD chassis .....</b>	<b>19</b>
<b>4 Contacting technical support .....</b>	<b>23</b>



## About this guide

This guide contains information about servicing and maintaining the Dell Acceleration Appliance for Databases (DAAD) software . This guide is intended for administrators responsible for server and storage systems. It is assumed that the reader is familiar with basic server administration.

## Typographical conventions

This document follows these conventions:

Convention	Usage	Examples
<b>NOTE:</b>	Important additional information or further explanation of a topic.	<b>NOTE:</b> A weekly backup is recommended.
<b>CAUTION!</b>	The task or operation might have serious consequences if conducted incorrectly or without appropriate safeguards. If you are not an expert in the use of this product, consult support for assistance.	<b>CAUTION!</b> Do not change configuration parameters.
<b>Bold</b>	A command or system input that you type, or text or a button you click on a graphical user interface (GUI).	Click <b>Help</b> for details about disaster recovery.
<i>Italic</i>	Italic font indicates any of the following: <ul style="list-style-type: none"> <li>• A term with a specific meaning in the context of this document.</li> <li>• Emphasis on specific information.</li> <li>• Reference to another document.</li> <li>• Variables in a syntax statement for which values are substituted.</li> </ul>	Detailed information about disaster recovery methods is available in the Administrator Guide.  <code>network:ping <i>hostname</i></code>
Courier	System output, file names or path names. Bold Courier for commands typed by user.	> Recovery in progress  <b>network:ping 10.1.100.14</b>
< > Angle Brackets	A required entry or variable parameter	installer-<version#>.run
Square [ ] Brackets	An optional entry or variable parameter.	<b>tar [zxvf] file.tgz</b>
Curly { } Brackets	A list of options separated by a the pipe symbol " " from which any one must be selected.	Click { <b>OK</b>   <b>Cancel</b> }.



## Software installation

**NOTE:** The Dell Acceleration Appliance for Databases software comes pre-installed. If you need to re-install the software, follow the steps in this section.


**NOTE:** To navigate through the installation screens, press **Tab** until your choice is highlighted, and then press **Enter**. Or, you can press **Alt+<highlighted letter>**, such as **Alt+H** for Help, or **Alt+N** for Next. Press the down arrow key to scroll through a list.

To install the Dell Acceleration Appliance for Databases software on your target server, follow the steps below:


- 1 From the customer support site, [dell.com/support/home](http://dell.com/support/home), download the .ISO image for the DAAD software you purchased.
- 2 Burn an installation DVD that contains the .ISO image. Or you can create a bootable USB or softmount the ISO through iDRAC.
- 3 Boot from the ISO image. The following screen is displayed:



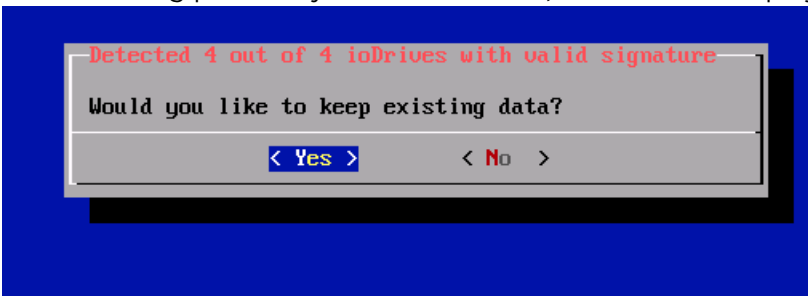
- 4 Select **Install ION Accelerator**.
- 5 If you have multiple disks in the appliance, select the disk where the Dell Acceleration Appliance for Databases software will be installed, such as `/dev/sda`.
- 6 Select **Yes**, when the following warning prompt is displayed:



A progress dialog box is displayed, first for loading the software and then for verifying the `/dev/sda` section on the boot disk).



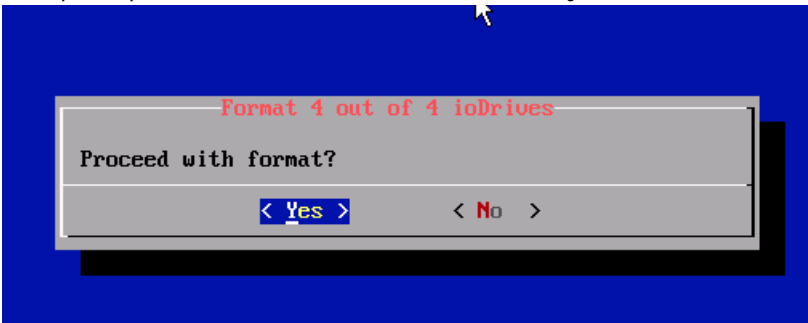
Detecting previously formatted drives, the software displays this prompt:



7 To keep the previous data on all the drives, select **Yes**.

Or

To remove the previous data from all the drives, select **No** . In this case, you are prompted to format all Fusion ioMemory devices in the DAAD system.



8 Select **Yes**.

The Dell Acceleration Appliance for Databases software installs. Complete the First Boot configuration of the new software using information in the *Dell Acceleration Appliance for Databases Configuration Guide*, and then, you can use the `config:restore` command to restore a previously saved configuration (if a configuration backup is available).



## DAAD drive and fabric card components

The basic procedure for replacing components in your Dell Acceleration Appliance for Databases is explained in this section. If you need instructions for the physical installation or removal of Fusion ioMemory devices, visit [dell.com/support/home](http://dell.com/support/home) and read the hardware guide that corresponds to the version of the Fusion ioMemory device you are replacing.

**CAUTION!** Improper removal or installation of a Fusion ioMemory device or network fabric adapter could void its warranty, as well as damage the device or cause personal injury.

**NOTE:** In standalone (non-HA) configurations, the Dell Acceleration Appliance for Databases system will be out of service during the replacement process.

The following scenarios are described:

- Replacing a Fusion ioMemory device in RAID 10 (Reliable Performance), RAID 0 (Maximum Performance), RAID 5 (Reliable Capacity), or JBOD (Direct Access).
- Replacing a fabric-connect adapter
- Replacing a cluster interconnect card

The instructions in this section may vary, depending on whether a single or multiple storage pools have been set up. Each variation is marked by *Single* storage pool or *Multiple* storage pools.

### Configuration guidelines — Hardware replacement

To ensure that the Dell Acceleration Appliance for Databases configuration is not inadvertently lost when hardware is replaced in a system:

- 1 Use the last known good configuration saved before component failure occurred (use `config --output`). It is a best practice to save and back up configurations on a regular basis.
- 2 If no saved configuration exists, save a new one before shutting down the Dell Acceleration Appliance for Databases to replace the failed part. Configuration information may be lost for failed cards if the appliance is restarted.

- 3 To save the configuration on a remote Unix system, enter:

```
config:backup --output-scp [user]:[password]<hostname>:<backup-filename>  
The backup file will be saved in the home (or /root) directory.
```

- 4 To restore a previously saved configuration from a remote Unix system, enter:

```
sftp --host <hostname> --user <username> --password <password> get  
<path_to_filename>  
config:restore --input-file <filename>
```

- 5 To save the configuration to a Windows system, enter:

```
config:backup --output-share (domain/user[:pass]@host/share[/dest])  
This will prompt for password after issuing the commands for a CIFS share. For  
example:
```

```
config:backup --output-share mydomain.com/testuser@10.10.10.10/docs/doc.xml
```

- 6 To restore a previously saved configuration from a Windows host, enter:

```
config:restore --input-share (domain/user[:pass]@host/share[/dest])  
This will prompt for password after issuing the commands for a CIFS share. For  
example:
```

```
config:restore --input-share mydomain.com/testuser@10.10.10.10/docs/doc.xml
```

## Replacing a drive in RAID 10 (reliable performance)

**NOTE:** There is no data loss for either standalone or HA mode when replacing Fusion ioMemory devices. For standalone, the system will be unavailable during Fusion ioMemory device replacement. For HA, the system will be available for only one node while the Fusion ioMemory device replacement is underway.

- 1 Using the GUI or CLI, shut down the node (standalone or a node in the cluster) with the failed Fusion ioMemory device by entering `system:shutdown --node <node to service>`

You can also use the Shut Down Host link in the GUI (**Configuration > Hosts > (host name) > Shut Down Host**) to accomplish the same thing.

- 2 Remove the Fusion ioMemory device in question and replace it with a good one.
- 3 Restart the system, which will recognize the newly added Fusion ioMemory device.
- 4 At the console prompt, select **Yes** to format the Fusion ioMemory device.

When the system starts, the RAID will be in a degraded state, and the replacement Fusion ioMemory device may have the same name as the replaced drive (such as fioa).

- 5 At the console enter the CLI command `raid:update -a <name of replacement Fusion ioMemory device>` to add the new Fusion ioMemory device to the degraded RAID.

The RAID status will change from degraded to rebuilding. After the rebuild is completed, the RAID state will go back to normal.

## Replacing a drive in RAID 0 (maximum performance)

This procedure can be performed on systems in standalone or HA modes.

### Standalone mode

**CAUTION!** There will be data loss in non-HA mode for an Fusion ioMemory device failure in a RAID 0 configuration.

- 1 At the console enter the CLI command `config:config --output-file backup.xml` to capture the configuration so it can be restored after Fusion ioMemory device replacement. This makes and keeps a copy to the local system.  
Or, enter `config:backup` to make a copy of the configuration to the remote system. See the *Dell Acceleration Appliance for Databases CLI Reference* for instructions on copying the configuration to the remote system.
- 2 Shut down the node with the failed Fusion ioMemory device by entering `system:shutdown`.  
For HA systems, enter `system:shutdown --node <node to service>`
- 3 Add the new Fusion ioMemory device(s), or remove the Fusion ioMemory device in question and replace it with a good one.
- 4 Restart the system, which will recognize the swapped Fusion ioMemory device.  
Single pool: Maintenance mode is not invoked, so you must enter `system:maintenance on`  
For HA systems, update the prompt status by entering `node:local` and then entering `system:maintenance on`  
Multiple pools: The system starts normally but is missing the pool for the failed Fusion ioMemory device. The RAID 0 configuration will be changed, and data will be lost only for this RAID 0 pool.
- 5 During startup, at the console screen, select **Yes** to format the Fusion ioMemory device when prompted.
- 6 Log in to the CLI.
- 7 Single pool: Turn on maintenance mode by entering `system: maintenance on`  
For HA systems, update the prompt status by entering `node:local` and then entering `system:maintenance on`
- 8 Locate the failed RAID by entering `raids -dt`

- 9 Delete the failed RAID by entering `raid:delete <md#>`
- 10 Restore the configuration.  
For example, if the backup file were *backup.xml* on a Windows (CIFS) share named `\\host1\share\DAAD` the command would be  

```
config:restore --input-share <domain/user>@host1/share/DAAD/backup.xml
```

This will rebuild the configuration on this node.
- 11 Restart the system by entering `system:restart`  
For HA systems, enter `system:restart --node <node to service>`

Single pool: The system will come back online with all configuration intact but with a new RAID, pools, volumes and LUNs.

Multiple pools: The system will recreate the missing pools, volumes and LUNs. The T10 IDs for recreated volumes will be different, but the names for the RAID, pool and volumes will be the same.

## HA mode

Follow the steps outlined in Standalone Mode, except:

- In Step 1, use `config:config --cluster --output-file backup.xml`  
There will be no data loss in HA mode for an Fusion ioMemory device failure on one node of a cluster in a RAID 0 configuration. Volumes will resync with the peer node automatically after the node joins the cluster.
- Before step 11 (`system:restart`), wait for all the volumes to be resynced. To verify, enter `volumes -dt` to show all volume status as Connected.
- After the last step, enter `setup resetios` on both nodes.

## Replacing a drive in RAID 5 (reliable capacity)

The steps for replacing an Fusion ioMemory device in a RAID 5 configuration are the same as with a RAID 10 configuration. For details, see [Replacing a drive in RAID 10 \(reliable performance\)](#) on page 10.

## Replacing a drive in a direct access (JBOD) pool

This procedure can be performed on systems in standalone or HA modes.

Standalone Mode

- 1 At the console enter the CLI command `config:config --output-file backup.xml` to capture the configuration so it can be restored after Fusion ioMemory device replacement. This makes and keeps a copy to the local system.  
Or, enter `config:backup` to make a copy of the configuration to the remote system. See *Dell Acceleration Appliance for Databases CLI Reference* for instructions on copying the configuration to the remote system.

- 2 Shut down the node with the failed Fusion ioMemory device by entering `system:shutdown`  
For HA systems, enter `system:shutdown --node <node to service>`
- 3 Remove the Fusion ioMemory device in question and replace it with a good one.
- 4 Restart the system, which will recognize the newly added Fusion ioMemory device.
- 5 At the console prompt, select **Yes** to format the Fusion ioMemory device.  
The system starts normally, but with the missing JBOD. The JBOD configuration will be changed, and data will be lost only for this JBOD pool.
- 6 Log in to the CLI.
- 7 Restore the configuration by entering:  
`config:restore --input-file backup.xml`

This will rebuild the lost JBOD configuration on this node, including corresponding pools, volumes, and LUNs. The T10 IDs for volumes will be different, but the names for the RAID, pool and volume will be the same.

## HA mode

Follow the steps outlined in Standalone Mode, except:

- In Step 1, use: `config:config --cluster --output-file backup.xml`

Volumes will resync with the peer node automatically after configuration is restored. There will be no data loss in HA mode for a Fusion ioMemory device failure on one node of a cluster in the Direct Access Profile (JBOD) configuration. In HA, the system will not enter maintenance mode, as it has to service other JBODs in the cluster.

## Replacing a Fibre Channel HBA

This procedure can be performed on systems in standalone or HA modes.

**NOTE:** If the Fibre Channel HBA is connected to a fabric switch, ensure the new HBA is zoned correctly after installation.

### Standalone mode

- 1 At the console enter the CLI command `config:config --output-file backup.xml` to capture the configuration so it can be restored after a Fibre Channel HBA replacement. This makes and keeps a copy to the local system.  
Or, enter `config:backup` to make a copy of the configuration to the remote system. See *Dell Acceleration Appliance for Databases CLI Reference* for instructions on copying the configuration to the remote system.
- 2 Shut down the node with the failed HBA by entering `system:shutdown`
- 3 Remove the HBA device in question and replace it with a good one.

- 4 Restart the system.  
The system will come up in maintenance mode. In this case, only LUNs are lost.
- 5 Update the prompt status by entering `node:local`
- 6 Restore the configuration on the node.  
For example, if the backup file were *backup.xml* on a Windows (CIFS) share named `\\host1\share\DAAD` the command would be  

```
config:restore --input-share <domain/user>@host1/share/DAAD/backup.xml
```
- 7 Restart the system by entering  
`system:restart`

## HA mode

- 1 At the console enter the CLI command `config:config --cluster --output-file backup.xml` to capture the configuration so it can be restored after FC HBA replacement. This makes and keeps a copy to the local system.  
Or, enter `config:backup` to make a copy of the configuration to the remote system. See *Dell Acceleration Appliance for Databases CLI Reference* for instructions on copying the configuration to the remote system.
- 2 Shut down the node with the failed HBA by entering `system:shutdown --node <node to service>`
- 3 Remove the HBA device in question and replace it with a good one.
- 4 Restart the system, which will come up in maintenance mode. In this case, only LUNs are lost.
- 5 Restore the configuration on the node by entering:  

```
config:restore --input-file backup.xml
```
- 6 Update the prompt status by entering `node:local`
- 7 Wait for all the volumes to be resynced by entering `volumes -dt`  
This command should show all volume status as Connected.
- 8 Restart the node by entering  
`system:restart.--node <node to service>`

There will be no data loss in HA mode for an HBA failure.

## Replacing a cluster interconnect card (HA)

**NOTE:** Ensure that the replacement card is a supported, Dell-branded cluster interconnect card.

- 1 Shut down the node with the failed Cluster Interconnect card by entering the CLI command `system:shutdown --node <node to service>`.
- 2 Remove the interconnect card in question and replace it with a good one.

- 3 Restart the system, which will automatically detect the new interconnect card and configure the IPs for the new card from the old configuration.

## Replacing an InfiniBand/SRP card

### Case 1: IB and cluster interconnect

In this example, the first port is configured as IB and the second port is configured as cluster interconnect.

- 1 At the console enter the CLI command `config:config --output-file backup.xml` to capture the configuration so it can be restored after InfiniBand/HCA replacement. This creates and saves a copy to the local system.  
Or, enter `config:backup` to make a copy of the configuration to the remote system. See *Dell Acceleration Appliance for Databases CLI Reference* for instructions on copying the configuration to the remote system.
- 2 Shut down the node with the failed HCA by entering `system:shutdown --node <node to service>`
- 3 Remove the HCA device in question and replace it with a good one.
- 4 Restart the system, which will come up in maintenance mode.
- 5 Update the prompt status by entering `node:local`
- 6 Enter `config:restore` to restore the configuration and LUNs on the node.
- 7 Restart the node by entering `system:restart.--node <node to service>`

### Case 2: Both ports configured as IB

In this example, both ports are configured as IB.

- 1 At the console enter the CLI command `config:config --output-file backup.xml` to capture the configuration and restore it after InfiniBand/HCA replacement. This makes and keeps a copy to the local system.  
Or, enter `config:backup` to make a copy of the configuration to the remote system. See *Dell Acceleration Appliance for Databases CLI Reference* for instructions on copying the configuration to the remote system.
- 2 Shut down the node with the failed HCA by entering `system:shutdown --node <node to service>`
- 3 Remove the HCA device in question and replace it with a good one.
- 4 Restart the system. Note that the system will not come up in maintenance mode.
- 5 Enter `config:restore` to restore the configuration and LUNs on the node.
- 6 Restart the node by entering `system:restart.--node <node to service>`

## Replacing an iSCSI card

**NOTE:** If you are using a split-port configuration with your iSCSI device, skip to [Replacing a split-function-port iSCSI card](#) on page 17.

### Standalone mode

- 1 At the console enter the CLI command `config:config --output-file backup.xml` to capture the configuration so it can be restored after iSCSI card replacement. This makes and keeps a copy to the local system.  
Or, enter `config:backup` to make a copy of the configuration to the remote system. See *Dell Acceleration Appliance for Databases CLI Reference* for instructions on copying the configuration to the remote system.
- 2 Enter `ports -dt` and write down the *Eth* numbers of the iSCSI ports being replaced.
- 3 Shut down the node with the failed iSCSI card by entering `system:shutdown`  
For HA systems, enter `system:shutdown --node <node to service>`
- 4 Remove the iSCSI card in question and replace it with a good one.
- 5 Restart the system, which will come up in maintenance mode.
- 6 For HA systems, update the prompt status by entering `node:local`

**NOTE:** Only LUNs are lost from the replaced iSCSI card.

- 7 Log in to the CLI and enter `setup lan`. This will display a Setup screen.
- 8 Remove the iSCSI card entries for the old card and configure the new card with previously used IP addresses.
- 9 Change the *Eth* number of this iSCSI port to the old iSCSI *Eth* number (as recorded in step 2).
- 10 In the Setup screen, select the port, then go to **Edit > Hardware > <device name>**.
- 11 Change the Device Name back to the original name (before the replacement was done).
- 12 Select **OK** and press **Enter**.
- 13 Restore the configuration on the node.  
For example, if the backup file were *backup.xml* on a Windows (CIFS) share named `\\host1\share\DAAD` the command would be  
`config:restore --input-share <domain/user>@host1/share/DAAD/backup.xml`
- 14 Restart the system by entering `system:restart`  
For HA systems, enter `system:restart --node <node to service>`

The system comes up with the configuration restored.



## HA mode

- 1 Follow steps 1 through 13 in the Standalone Mode instructions, ensuring that both nodes are online.
- 2 Enter `maintenance off`
- 3 Wait for all the volumes to be resynced by entering `volumes -dt`  
This should show all volume status as Connected.
- 4 Restart the node by entering  
`system:restart.--node <node to service>`

After the node configuration is restored, the node will join cluster, and the cluster will come back to the normal, full service state. There will be no data loss in HA mode for an iSCSI card failure.

## Replacing a split-function-port iSCSI card

A split-function-port iSCSI device has one port configured for cluster interconnect and the other port configured for iSCSI traffic.

- 1 At the console enter the CLI command `config:config --output-file backup.xml` to capture the configuration so it can be restored after iSCSI card replacement. This makes and keeps a copy to the local system.  
Or, enter `config:backup` to make a copy of the configuration to the remote system. See *Dell Acceleration Appliance for Databases CLI Reference* for instructions on copying the configuration to the remote system.
- 2 Enter `ports -dt` and write down the *Eth* numbers of the iSCSI and cluster ports being replaced.
- 3 Enter `system:shutdown --node <node to service>` to shut down the node with the failed iSCSI card.
- 4 Remove the iSCSI card in question and replace it with a good one.
- 5 Restart the system by entering `system:restart --node <node to service>`  
The system will come up in maintenance mode.
- 6 Update the prompt status by entering `node:local`

**NOTE:** Only LUNs are lost from the replaced iSCSI card.

- 7 Log in to the CLI and enter `setup lan` to set the iSCSI port IP addresses. This will display a Network Configuration screen.  
New *Eth* device names will be seen from Network Configuration screen instead of the old *Eth* names mentioned in step 2.
- 8 Find the iSCSI port that was replaced and identify it using the MAC address.

- 9 Change the *Eth* number of this iSCSI port to the old iSCSI *Eth* number (as recorded in step 2).

**NOTE:** Only the new iSCSI *Eth* port must be changed; no change is needed for the *Eth* cluster interconnect.

- 10 In Network Configuration screen, select the port, then go to **Edit > Hardware > Device name**.
- 11 Change the Device Name back to the original name before the replacement was done.
- 12 Select **OK** and press **Enter**.
- 13 Enter `system:restart --node <node to service>` to restart and bring up the system with the configuration restored.

## Correct iSCSI target entry after a replacement

After an iSCSI device is replaced, creating a new target entry for the device in the Network Configuration screen will cause the old entry to go stale — it cannot be deleted, as it still has references to existing LUNs.

Instead of creating a new target entry in Network Configuration screen for the new hardware, you should reuse the old entry:

- 1 Delete the old hardware device from the Network Configuration screen.
- 2 Set the device name to be the same as the existing target under the Hardware tab.

## Replacing the DAAD chassis

**NOTE:** For a DAAD configured for HA mode, you will need installation media that matches the installed cluster.

To replace a DAAD chassis:

- 1 Create a backup of the current configuration by running `config:backup`. If the chassis is damaged to the point where a new backup cannot be made, use the last successful backup.

For example, the following commands create an XML backup that can be used later for a restore:

```
config:backup --host <hostname> --directory <path> --user <username>
--password <password>
```

- 2 Store the configuration file at a location that is not on the machine you are replacing.
- 3 Obtain the domain name from the backup and make a note of it.

When doing a chassis replacement, you need to match the domain name that was saved in the XML backup with the DNS hostname from the first boot Hostname/DNS screen. If this is not done, the `config:restore` command will fail.

**NOTE:** NOTE: Changing the DNS hostname is not currently supported after the First Boot process, so setting the name correctly must be done during First Boot.

To find the domain name in the XML backup, examine the first few lines of the XML text. Under the `<domain>` section, the second line will be similar to `<name>MYNAME</name>` (the domain name is `MYNAME` in this example).

```
<domain>
<created>2013-02-27T12:18:43.030-08:00</created>
<name>MYNAME</name>
```

This name, identical in case, must be entered when configuring the replacement chassis on the "Hostname/DNS" screen.

- 4 For standalone mode, from each initiator, stop any data traffic, and logout from all the iSCSI sessions connecting to the system that needs to be replaced.

- 5 Prepare the replacement chassis.  
Unbox it, put it in a location suitable for the exchange of hardware, and install rack mount hardware if required. If replacement HBAs or NICs are being used, note any relevant GUIDs or MAC addresses.
- 6 Power down the chassis to be replaced by running `system:shutdown -n <nodename>`  
If this does not succeed, use the power button to power off the chassis.
- 7 Remove power from the chassis to be replaced.
- 8 Remove all cables from the chassis to be replaced.
- 9 Label each cable with the slot/port from which it has been removed.
- 10 Transfer necessary hardware from the original chassis to the new chassis.  
At a minimum, the Fusion ioMemory devices must be transferred. Retain the same slot locations for all transferred hardware. Note that hard drives, HBAs, cluster interconnect cards, and other NICs may be transferred as well.
- 11 Install the replacement chassis. If the server needs to be installed in a rack, do so now.
- 12 Replace all previously labeled cables in the correct locations.
- 13 Connect the console access.
- 14 Insert the power cords.
- 15 Power up the new chassis. If new hard drives are being used, make the installation media available (insert the DVD in the drive).

**NOTE:** It is extremely important that the Initial installation questions be answered correctly.

- 16 Select **Yes** when the following prompt appears:  
`Destroying ALL data on /dev/sda, continue?`
- 17 Select **Yes** when the following prompt appears (otherwise, the existing data will be erased):  
`Detected ioDrives for ION Data Accelerator  
Would you like to keep existing data?`  
  
Or, if all the Fusion ioMemory devices in the node have been replaced, select **Yes** to format the Fusion ioMemory when prompted.
- 18 Complete the First Boot screens (see About DAAD first boot in the *Dell Acceleration Appliance for Databases Configuration Guide*.  
The Hostname/DNS screen must be configured with the host name obtained in Step 3. The network interfaces and cluster interfaces must also be configured.

- 19 If you are using new hard drives or a fresh installation, restore the configuration from the saved XML file. For example:

```
sftp --host 10.60.33.123 --user <username> --password <password>  
get chassisreplacement.xml
```

If you have replaced Fibre Channel HBAs, you also need to run this command:

```
config:restore --input-file chassis-replacement.xml
```

- 20 Run `system:restart -n <nodename>` after successfully restoring the configuration.

- 21 In HA configurations, run `setup resetios` on both the nodes.

- 22 If you are using the existing hard drives in the new chassis and have replaced the cluster interconnect HBAs or NICs in the new chassis, without doing a fresh OS installation, then perform the following steps:

**NOTE:** Ensure the cluster is enabled. If the cluster is not enabled, the `config:restore` command in step b will fail.

- a Delete the entries for obsolete NICs and configure new NICs by running the `setup lan` command at the console.

- b Enter `config:restore --input-file chassis-replacement.xml`

- 23 For standalone mode, discover new target names and create new sessions in each initiator.

**NOTE:** If you have transferred all the components in the system except the built-in management NIC, run `setup lan`, delete the old entry with the `ETH0` device name, and edit the new entry for management port `ETH0` with the IP address that was previously assigned. Then restart the system.

**NOTE:** If you have transferred all the components in the system including boot drives, the boot drives may appear as foreign drives. If this happens, type **f** to go into the BIOS and accept them. Then restart the system



## Contacting technical support

Dell Acceleration Appliance for Databases drivers, utilities, and related documentation are available at:

[dell.com/support/home](http://dell.com/support/home)

Dell provides several online and telephone-based support and service options. Availability varies by country and product, and some services may not be available in your area. To get help with your Fusion ioMemory devices, contact your Dell Technical Service representative or access the Dell Support website.

Choose the method of contacting Dell that is convenient for you.

**NOTE:** The safety information that shipped with your system provides important safety and regulatory information. Warranty information may be included within this document or as a separate document.

