# Oracle® SuperCluster M6-32

Owner's Guide: Administration

Part No.: E41537-05
November 2015

Please
Recycle

Pour plus d'informations sur l'engagement d'Oracle pour l'accessibilité à la documentation, visitez le site Web Oracle Accessibility Program, à l'adresse http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

**Accès au support électronique**

Les clients Oracle qui ont souscrit un contrat de support ont accès au support électronique via My Oracle Support. Pour plus d'informations, visitez le site http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info ou le site http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs si vous êtes malentendant.

# Contents

# Using This Documentation

This document describes how to monitor and administer Oracle SuperCluster M6-32.

- **Overview** – Describes how to monitor and administer Oracle SuperCluster M6-32
- **Audience** – Technicians, system administrators, and authorized service providers
- **Required knowledge** – Advanced experience in system administration
- "Product Documentation Library" on page viii
- "Feedback" on page viii

## Product Documentation Library

Documentation and resources for this product and related products are available on the system. Access the documentation by using a browser to view this directory on the first compute server installed in SuperCluster:

`/opt/oracle/node/doc/E41531_01`

## Feedback

Provide feedback about this documentation at:

`http://www.oracle.com/goto/docfeedback`

# Understanding SuperCluster Software

These topics describe how to identify the software version and some of the tools that are available on Oracle SuperCluster systems.



## ▼ Identify the Version of SuperCluster Software

Perform this procedure to determine the version of SuperCluster software.

1. **On the management network, log in to one of the SPARC servers.**
2. **Type.**

```
# svcprop -p configuration/build svc:/system/oes/id:default
```

3. **Determine the version of software based on the output.**
   - `ssc-1.x.x`
   - `ssc-2.x.x` (or later)

## SuperCluster Tools

Use one of these tables according to the version of SuperCluster software. Use the tables to see what tools are available.

To identify your version, see "Identify the Version of SuperCluster Software" on page 1.

**Note –** The tables do not provide a complete list of tools for SuperCluster. Instead, the tables list the tools are covered in subsequent sections and provides information on what is available based on specific version of SuperCluster software.

**TABLE:** SuperCluster v1.x Tools

| Tool | Description | Links |
|---|---|---|
| `ssctuner` | A set of scripts and configuration files that run on SuperCluster Oracle Solaris 10 and Oracle Solaris 11 global zones to monitor and tune various parameters. | "Tuning SuperCluster M6-32" on page 29 |
| `setcoremem` | This command is no longer available. Instead, use the `osc-setcoremem` tool to change how CPU and memory resources are allocated across domains. | "Configuring CPU and Memory Resources (osc-setcoremem)" on page 39 |

**TABLE:** SuperCluster v2.x Tools

| Tool | Description | Links |
|---|---|---|
| Oracle I/O Domain Creation tool | Enables you to create I/O domains on demand, assigning CPU, memory and I/O resources of your choice. | Refer to the *Oracle I/O Domain Administration Guide*. |
| `osc-setcoremem` | Enables you to change how CPU and memory resources are allocated across domains. The tool automatically assigns the appropriate amount of memory to each domain based on how you allocated CPU resources, ensuring optimal performance by minimizing NUMA effects. | "Configuring CPU and Memory Resources (osc-setcoremem)" on page 39 |
| `ssctuner` | A set of scripts and configuration files that run on SuperCluster Oracle Solaris 10 and Oracle Solaris 11 global zones to monitor and tune various parameters. | "Tuning SuperCluster M6-32" on page 29 |

# Controlling SuperCluster M6-32

These topics describe power on and power off information for SuperCluster M6-32.



## Cautions

The following cautions apply to SuperCluster M6-32.

**Caution –** Do not touch the parts of this product that use high-voltage power. Touching them might result in serious injury.

**Caution –** Keep the front and rear cabinet doors closed. Failure to do so might cause system failure or result in damage to hardware components.

**Caution –** Keep the top, front, and back of the cabinets clear to allow proper airflow and prevent overheating of components.

Use only the supplied hardware.

# ▼ Power On SuperCluster M6-32

Power on SuperCluster M6-32 in the reverse order of shutdown.

1. **Turn on both circuit breakers that provide power to the rack.**

   The switches power on, and the storage servers, base servers, and the ZFS storage appliance return to standby mode.

2. **Boot each ZFS storage appliance.**
3. **Boot each base server.**
4. **Boot each storage server.**

# Powering Off SuperCluster M6-32 Gracefully

Perform these tasks to shut down and power off SuperCluster M6-32 gracefully.

| Step | Description | Links |
|---|---|---|
| 1. | Shut down the cluster. | "Shut Down the Oracle Solaris Cluster" on page 5 |
| 2. | Shut down the Enterprise Controller if Ops Center is running. | "Shut Down the Enterprise Controller (Ops Center)" on page 5 |
| 3. | Shut down the database. | "Shut Down the Database" on page 5 |
| 4. | Shut down the storage servers. | "Shut Down the Storage Servers" on page 5 |
| 5. | Power off the storage servers. | "Power Off the Storage Servers" on page 6 |
| 6. | Shut down LDoms. | "Shut Down the LDoms" on page 6 |
| 7. | Shut down the OS running on the servers. | "Shut Down the OS on the Compute Servers" on page 7 |
| 8. | Shut down the storage appliance. | "Shut Down the ZFS Storage Appliance" on page 7 |
| 9. | Power off the switches and racks. | "Power Off the Switches and Racks" on page 7 |

## ▼ Shut Down the Oracle Solaris Cluster

- **Type.**

  # **/usr/cluster/bin/cluster shutdown -g 0 -y**

## ▼ Shut Down the Enterprise Controller (Ops Center)

- **If Ops Center is running, shut down the enterprise controller.**

  # **/opt/SUNWxvmoc/bin/ecadm stop**

  For HA environments use this command:

  # **/opt/SUNWxvmoc/bin/ecadm ha-stop-no-relocate**

## ▼ Shut Down the Database

- **Shut down the database using one of the methods described at this URL:**

  http://docs.oracle.com/cd/B28359_01/server.111/b28310/start003.htm

## ▼ Shut Down the Storage Servers

Perform this procedure for each storage server before you power them off. For more information on this task, refer to the Exadata documentation at:

http://wd0338.oracle.com/archive/cd_ns/E13877_01/doc/doc.112/e13874/maintenance.htm#CEGBHCJG

1. **Check for other offline disks.**

   CellCLI> **LIST GRIDDISK ATTRIBUTES name WHERE asmdeactivationoutcome != 'Yes'**

   If any grid disks are returned, then it is not safe to take the storage server offline, because proper Oracle ASM disk group redundancy will not be maintained. Taking a storage server offline when one or more grid disks are in this state causes Oracle ASM to dismount the affected disk group, which causes the databases to shut down abruptly.

2. **When a storage server is safe to take offline, inactivate all the grid disks.**

   ```
   CellCLI> ALTER GRIDDISK ALL INACTIVE
   ```

   This command completes after all disks are inactive and offline.

3. **Verify that all grid disks are inactive.**

   ```
   CellCLI> LIST GRIDDISK WHERE STATUS != 'inactive'
   ```

   If all grid disks are inactive, then you can shut down a storage server without affecting database availability.

4. **Shut down the cell.**

   See "Power Off the Storage Servers" on page 6.

## ▼ Power Off the Storage Servers

Perform the following procedure for each storage server.

Note the following when powering off storage servers:

- All database and Oracle Clusterware processes must be shut down prior to shutting down more than one storage server.
- Powering off one storage server does not affect running database processes or Oracle ASM.
- Powering off or restarting a storage server can impact database availability.

- **Shut down a storage server immediately.**

  ```
  # shutdown -h -y now
  ```

## ▼ Shut Down the LDoms

The base server configurations vary based on the configuration chosen during installation.

- If the compute server is running with one LDom, shut down the machine just as you would any other server by cleanly shutting down the OS.
- If the compute server is running two LDoms, shut down the guest domain first and then the primary (control) domain.
- If the compute server is running with three or more domains, identify the domains that are running off virtualized hardware and shut them down first before moving on to shutting down the guest domain and finally the primary (control) domain.

1. **Shut down, stop, and unbind each of the non-I/O domains.**

```
# ldm stop domainname
LDom domainname stopped
# ldm unbind-domain domainname
```

2. **Shut down, stop, and unbind any active I/O domains.**

```
# ldm stop activedomainname
LDom activedomainname stopped
# ldm unbind-domain activedomainname
```

3. **Halt the primary domain.**

   **# shutdown -i5 -g0 -y**

   Because no other domains are bound, the firmware automatically powers off the system.

## ▼ Shut Down the OS on the Compute Servers

- **Gracefully shut down the Oracle Solaris OS on each compute server:**

  **# init 0**

## ▼ Shut Down the ZFS Storage Appliance

- **Gracefully shut down the ZFS storage appliance, by logging in to the browser interface and clicking the power icon on the left side of the top pane.**

## ▼ Power Off the Switches and Racks

- **Power off the switches, and the entire rack, by turning off the circuit breakers.**

# ▼ Power Off SuperCluster M6-32 in an Emergency

If there is an emergency, such as earthquake or flood, an abnormal smell or smoke coming from SuperCluster M6-32, or a threat to human safety, you must power off SuperCluster M6-32 immediately. In that case, use one of the following ways to power off SuperCluster M6-32.

- **Power off SuperCluster M6-32 one of the following ways:**
  - Turn off power at the circuit breaker, or pull the emergency power-off switch in the computer room.
  - Turn off the site EPO switch to remove power from SuperCluster M6-32.
  - Turn off the two PDUs in the rack.

  After the emergency, contact Oracle Support Services to restore power to SuperCluster M6-32.

# Monitoring SuperCluster M6-32 (OCM)

These topics describe how Oracle Configuration Manager can be used to monitor SuperCluster M6-32 and where to obtain the latest OCM documentation.



## OCM Overview

OCM collects configuration information and uploads it to the Oracle repository. When the configuration information is uploaded daily, Oracle Support Services can analyze the data and provide better service. When a service request is logged, the configuration data is associated with the service request. The following are some of the benefits of OCM:

- Reduced time for problem resolution
- Proactive problem avoidance
- Improved access to best practices and the Oracle knowledge base
- Improved understanding of the customer's business needs
- Consistent responses and services

The OCM software is installed and configured in each `ORACLE_HOME` directory on a host. For clustered databases, only one instance is configured for OCM. A configuration script is run on every database on the host. The OCM collects and then sends the data to a centralized Oracle repository.

For more information, refer to the OCM documentation. See "Access OCM Documentation" on page 10:

# ▼ Access OCM Documentation

To access the latest OCM documentation, visit the OCM web page. The documentation describes how to install, administer, and use OCM.

- **In a browser, go to:**

  http://www.oracle.com/technetwork/documentation/ocm-092152.html

# Monitoring the System With ASR

These topics provide information about ASR and how to use it to monitor SuperCluster M6-32.

**Note –** Oracle personnel might have configured ASR during the installation of SuperCluster.



## ASR Overview

ASR automatically opens service requests when specific hardware faults occur. In many cases, Oracle Support Services can begin resolving the issue immediately, often before the system administrator is aware that a problem exists.

The telemetry data that is sent from the ASR Manager to Oracle is encrypted.

When ASR detects a fault, ASR sends an email message to your MOS email account for ASR, and to the technical contact for the activated asset, notifying them of the creation of the service request.

To enable this feature, ASR Manager software must be installed on a server (a server other than SuperCluster). The ASR Manager server must have connectivity to SuperCluster, and an outbound Internet connection using HTTPS or an HTTPS proxy. Certain SuperCluster components must be configured to send hardware fault telemetry to the ASR Manager server.

For SuperCluster systems, ASR uses these telemetry sources to detect fault events:

- **FMA** – Provides CPU and memory fault information from the host.
- **Oracle ILOM** – Provides fault information, power and environmental, and CPU and memory fault information from the service processor.
- **Exadata-detected Events (HALRT)** – Provides fault coverage for disks, flash, and PCI cards within Oracle SuperCluster.
- **ZFS storage appliance** – Provides fault events detected within the systems and disk arrays of the included Storage Appliance.
- **IB switch management module** – Provides fault coverage for power, memory, storage, and battery.

Consider this information when using ASR:

- ASR is applicable only for component faults. Not all component failures are covered, however components that are most likely to generate faults, such as disks, fans, and power supplies, are covered by ASR.
- ASR is not a replacement for other monitoring mechanisms, such as SMTP and SNMP alerts. ASR is a complementary mechanism that expedites and simplifies the delivery of replacement hardware. ASR should not be used for downtime events in high-priority systems. For high-priority events, contact Oracle Support Services directly.
- There are occasions when a service request might not be automatically filed. This can happen due to problems with the SNMP protocol or loss of connectivity to the ASR Manager. You must continue to monitor the systems for faults and call Oracle Support Services if you do not receive notice that a service request has been automatically filed.

# ASR Resources

For ASR details, refer to these resources:

| Resource | Links |
|---|---|
| The main ASR web page | http://www.oracle.com/asr |
| ASR documentation | http://www.oracle.com/technetwork/systems/asr/documentation |
| ASR download page | http://www.oracle.com/technetwork/systems/asr/downloads |
| Oracle Services Tools Bundle documentation | http://docs.oracle.com/cd/E35557_01 |
| A list of SuperCluster components that are qualified ASR assets. | http://docs.oracle.com/cd/E37710_01/doc.41/e37287/toc.htm |
| ASR fault coverage | http://docs.oracle.com/cd/E37710_01/doc.41/e55817/toc.htm |

# ASR Installation Overview

This is a high-level description of the tasks you perform to configure ASR on SuperCluster. For more details, refer to ASR documentation. See "ASR Resources" on page 12.

1. Create or verify your MOS account at http://support.oracle.com.

   Ensure that your MOS account is correctly set up:

   - Oracle Premier Support for Systems or Oracle/Sun Limited Warranty
   - Technical contact responsible for SuperCluster system
   - Valid shipping address for parts

2. Designate a standalone system to serve as the ASR Manager and install the ASR Manager software.

   The server must run either Oracle Solaris or Linux, and Java

   You must have superuser access to the ASR Manager system.

   To download ASR Manager software, go to:
   http://www.oracle.com/technetwork/systems/asr/downloads

   For installation instructions, refer to the ASR documentation that corresponds to the version of ASR Manager you plan to install. The documentation is available at:

   http://docs.oracle.com/cd/E37710_01/index.htm

3. Ensure that the ASR Manager server has connectivity to the Internet using HTTPS.

You might need to open certain ports to your datacenter. For more information, see the *Oracle ASR Security White paper*, located here:

`http://docs.oracle.com/cd/E37710_01/index.htm`

4. (Optional) Obtain these documents:
   - *Oracle SuperCluster M6-32 System Site Checklists*
   - *Oracle SuperCluster M6-32 System Configuration Worksheets*

   The information in these documents can provide helpful information when you configure SuperCluster for ASR.

5. Configure and activate SuperCluster ASR assets.

   Refer to ASR documentation and see "Configuring ASR on the Compute Servers (Oracle Solaris 11)" on page 21.

---

**Note –** An active ASR Manager must be installed and running before you configure ASR assets.

---

---

**Note –** To monitor Oracle Solaris 10 assets, you must install the latest STB bundle on SuperCluster. Refer to Doc ID `1153444.1` to download the latest STB bundle from MOS: `https://support.oracle.com`

---

6. Approve ASR assets in MOS.

   Follow the instructions in the ASR documentation.

# ▼ Configure ASR on the Compute Servers (Oracle ILOM)

---

**Note –** Do not attempt to copy and paste commands that span across multiple lines from this section. Manually type commands that span across multiple lines to ensure the commands are typed properly.

---

---

**Note –** You must have ASR Manager running in your environment before you can perform these tasks.

---

To configure the Oracle ILOM for compute servers, complete these steps on each compute server:

1. **Log in to Oracle ILOM on the first compute server.**
2. **Display the available rules:**

   # **show /SP/alertmgmt/rules**

   This lists the rules available, similar to the following:

   ```
   1
   2
   3
   ...
   15
   ```

3. **Pick one of the rules and type this command to determine if that rule is currently in use:**

   # **show /SP/alertmgmt/rules/***rule-number*

   For example:

   # **show /SP/alertmgmt/rules/1**

   - If you see output similar to this:

     ```
     Properties:
      type = snmptrap
      level = minor
      destination = 10.60.10.243
      destination_port = 0
      community_or_username = public
      snmp_version = 2c
      testrule = (Cannot show property)
     ```

     This rule is currently being used and should not be used for this exercise (the destination address shown is the IP address of the ASR Manager in this case).

     If you see output similar to the preceding example, pick another rule until you find one that is not in use.

   - If you see output similar to this:

     ```
     Properties:
     type = snmptrap
     level = disable
     destination = 0.0.0.0
     ```

```
destination_port = 0
community_or_username = public
snmp_version = 1
testrule = (Cannot show property)
```

this rule is currently unused and can be used for this exercise.

4. **Type this command using the unused rule:**

   # **set /SP/alertmgmt/rules/***unused-rule-number* **type=snmptrap level=minor destination=***IP-address-of-ASR-Manager* **destination_port=***port_number* **snmp_version=2c community_or_username=public**

5. **Log in to the ASR Manager server.**

6. **Activate Oracle ILOM for the compute server:**

   asr> **activate_asset -i** *ILOM-IP-address*

7. **Repeat these instructions on Oracle ILOM for all compute servers in your SuperCluster system.**

8. **Go to the next configuration task.**

   See "Configure SNMP Trap Destinations for Storage Servers" on page 16.

# ▼ Configure SNMP Trap Destinations for Storage Servers

**Note –** Do not attempt to copy and paste commands that span across multiple lines from this section. Manually type commands that span across multiple lines to ensure the commands are typed properly.

Perform this task on each storage server:

1. **Log in as** `celladmin` **on the storage server.**

2. **On the storage server, add SNMP trap destinations:**

   # **cellcli -e "alter cell snmpSubscriber=(host = '***ASR-Manager-name-or-IP-address***',port=162,community=public,type=asr)"**

   Single quotes are required around the *ASR-Manager-name-or-IP-address* entry. These are the element definitions for the command above:

- `host='ASR-Manager-name-or-IP-address'` – The ASR Manager hostname can be used when DNS is enabled for the site. If DNS is not running, the IP address is preferred, but the ASR Manager hostname can be used if the entry is added to the `/etc/hosts` file.
- `type=asr` – Specifies the ASR Manager as being a special type of SNMP subscriber.
- `community=public` – The required value of the community string. This value can be modified to be a different string based on your network requirements.
- `port=162` – The SNMP port. This port value is site dependant. It can be configured as a different port based on network requirements, or it might need to be changed for ASR to work correctly in a managed environment.

3. **Ensure that Oracle ILOM auto-activation occurred**

   This confirms that the network and Oracle ILOM are set up correctly.

   # **asr list_asset**

   Example:

```
IP_ADDRESS     HOST_NAME        SERIAL_NUMBER   ASR       PROTOCOL    SOURCE
----------     ---------        -------------   ---       --------    -----
10.60.40.105   ssc1cel01        1234FMM0CA      Enabled   SNMP        ILOM
10.60.40.106   ssc1cel02        1235FMM0CA      Enabled   SNMP        ILOM
10.60.40.107   ssc1cel03        1236FMM0CA      Enabled   SNMP        ILOM
10.60.40.117   ssc1cel01-ilom   1234FMM0CA      Enabled   SNMP,HTTP   EXADATA-SW
10.60.40.118   ssc1cel02-ilom   1235FMM0CA      Enabled   SNMP,HTTP   EXADATA-SW
10.60.40.119   ssc1cel03-ilom   1236FMM0CA      Enabled   SNMP,HTTP   EXADATA-SW
```

   - If all Oracle ILOMs for the storage servers are in the list, go to Step 5.
   - If Oracle ILOMs are not in the list, go to Step 4.

4. **On the ASR Manager, activate the Oracle ILOMs of the storage servers:**

   # **asr activate_asset -i** *ILOM-IP-address*

   or

   # **asr activate_asset -h** *ILOM-hostname*

---

**Note –** If the last step fails, verify that port 6481 on the Oracle ILOM is open. If port 6481 is open and the step still fails, contact ASR Support.

---

5. **Activate the Exadata OS side of the ASR support:**

   # **asr activate_exadata -i** *host-management-IP-address* **-h** *host-management-hostname* **-l** *ILOM-IP-address*

   or

   # **asr activate_exadata -i** *host-management-IP-address* **-h** *host-management-hostname* **-n** *ILOM-hostname*

6. **Ensure that all storage servers are visible on the ASR Manager.**

   # **asr list_asset**

   You should see both the Oracle ILOM and the host referenced in the list, with the same serial number, as shown in this example:

```
IP_ADDRESS     HOST_NAME        SERIAL_NUMBER   ASR       PROTOCOL    SOURCE
----------     ---------        -------------   ---       --------    -----
10.60.40.105   ssc1cel01        1234FMM0CA      Enabled   SNMP        ILOM
10.60.40.106   ssc1cel02        1235FMM0CA      Enabled   SNMP        ILOM
10.60.40.107   ssc1cel03        1236FMM0CA      Enabled   SNMP        ILOM
10.60.40.117   ssc1cel01-ilom   1234FMM0CA      Enabled   SNMP,HTTP   EXADATA-SW
10.60.40.118   ssc1cel02-ilom   1235FMM0CA      Enabled   SNMP,HTTP   EXADATA-SW
10.60.40.119   ssc1cel03-ilom   1236FMM0CA      Enabled   SNMP,HTTP   EXADATA-SW
```

7. **On the storage server, validate the configuration:**

   # **cellcli -e "list cell attributes snmpsubscriber"**

8. **On the storage server, validate the SNMP configuration:**

   # **cellcli -e "alter cell validate snmp type=asr"**

   The MOS contact receives an email as confirmation.

9. **Repeat these instructions for every storage server in your SuperCluster.**

10. **Go to the next configuration task.**

    See "Configure ASR on the ZFS Storage Appliance" on page 18.

# ▼ Configure ASR on the ZFS Storage Appliance

To activate the storage appliance included in your SuperCluster system, complete these steps on each ZFS storage controller:

1. **In a web browser, type the IP address or host name you assigned to the host management port of either ZFS storage controller:**

   `https://`*storage-controller-ipaddress*`:215`

   or

   `https://`*storage-controller-hostname*`:215`

   The login screen appears.

2. **Type `root` into the `Username` field and the root password into this login screen, and press the `Enter` key.**

3. **Click the Configuration tab, and click SERVICES, and then on the left navigation pane, click Services to display the list of services.**

4. **Scroll down in the screen and click Phone Home.**

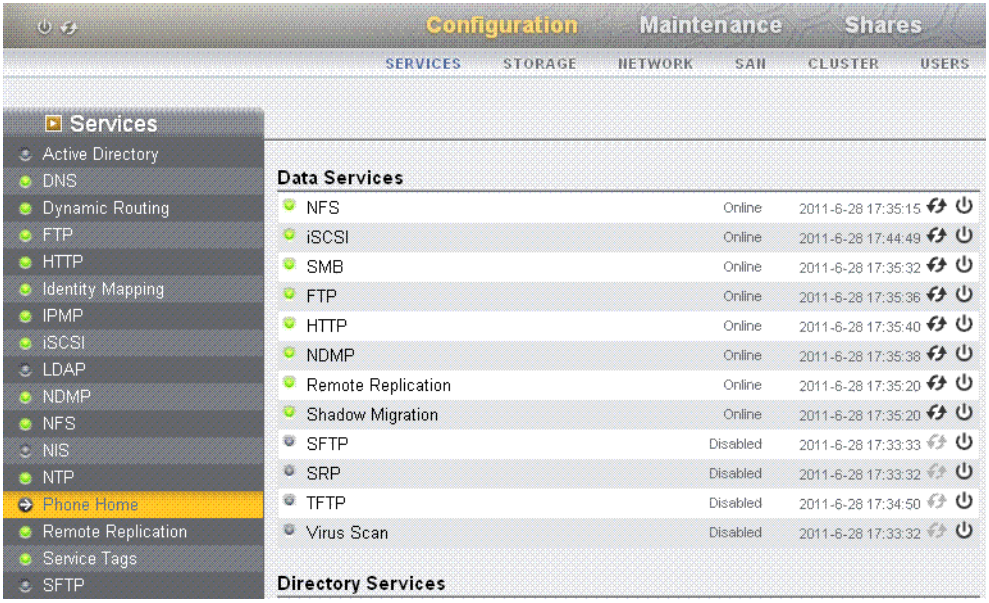


The Phone Home page is displayed:

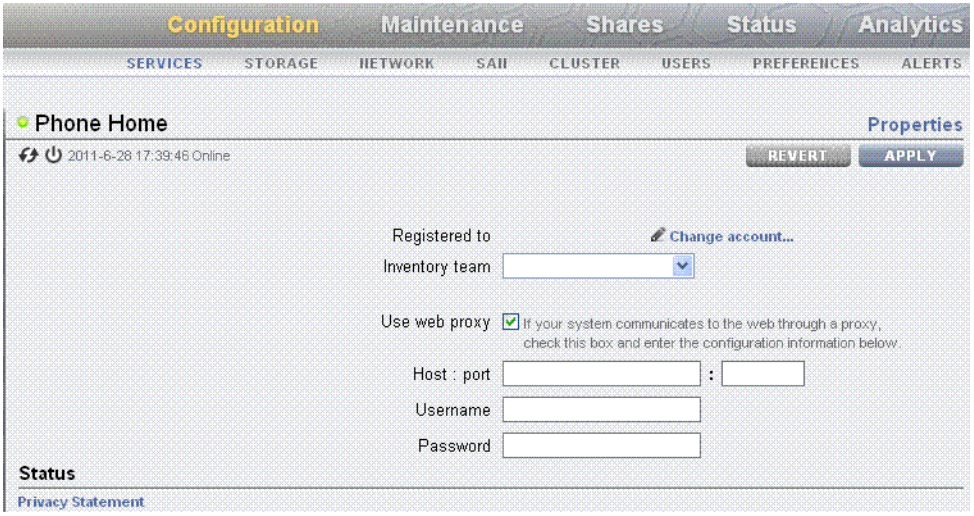


5. **If you are using a web proxy to connect to the Internet from the storage appliance, select the Use web proxy option, and provide this information:**
   - In the Host:port field, type the complete host name of your web proxy server and the port.
   - In the *Username* field, type your user name for the accessing the web proxy server.
   - In the *Password* field, type the password.
6. **Click the pencil icon in the registration section.**

   A Privacy Statement is displayed. Click OK, complete the section for My Oracle Support and password, and click OK.
7. **When the account is verified, select the Sun Inventory and Enable Phone Home options.**
8. **After typing the information, click APPLY.**
9. **When the Service Enable / Disable popup is presented, select the Enable option.**
10. **Repeat these instructions for every storage controller in your SuperCluster system**
11. **Go to the next configuration task.**

    See "Configure ASR on the Compute Servers (Oracle ILOM)" on page 14.

# Configuring ASR on the Compute Servers (Oracle Solaris 11)

**Note –** Do not attempt to copy and paste commands that span across multiple lines from this section. Manually type commands that span across multiple lines to ensure the commands are typed properly.

Oracle Solaris 11 includes the ability to send ASR fault events and telemetry to Oracle using xml over HTTP to the ASR Manager.

To enable this capability, perform the following tasks:

- "Enable the HTTP Receiver on the ASR Manager" on page 21
- "Enable HTTPS on ASR Manager (Optional)" on page 22
- "Register Compute Servers With Oracle Solaris 11 or Database Domains to ASR Manager" on page 23

## ▼ Enable the HTTP Receiver on the ASR Manager

Perform this procedure on the ASR Manager to enable the HTTP receiver for Oracle Solaris 11 ASR Assets.

1. **Log in to the ASR Manager system as superuser.**
2. **Verify the existing settings:**

   # **asr show_http_receiver**
3. **Enable the HTTP receiver:**

   # **asr enable_http_receiver -p** *port-number*

   where *port-number* is the port that you are designating for HTTP traffic.

**Note –** If you need to disable the HTTP receiver, run `asr disable_http_receiver`.

4. **Verify the updated configuration:**

   # **asr show_http_receiver**

5. **Verify the HTTP receiver is up and running.**

   In a browser, go to: `http://`*ASR-Manager-name*`:`*port-number*`/asr`

   A message displays indicating that the HTTP receiver is up and running.

## ▼ Enable HTTPS on ASR Manager (Optional)

If you need to use HTTPS for security purposes, you can set up HTTPS/SSL for the ASR Manager HTTP receiver.

---

**Note –** The detailed steps for enabling HTTPs/SSL for Jetty are documented at: http://docs.codehaus.org/display/JETTY/How+to+configure+SSL

---

1. **Once the SSL certificate from a trusted authority is loaded into keystore, add this SSL connector in** `/var/opt/SUNWsasm/configuration/jetty/jetty.xml` **below the** `<Call name="addConnector">` **sections:**

```
<Call name="addConnector">
  <Arg>
    <New class="org.mortbay.jetty.security.SslSocketConnector">
      <Set name="Port">443</Set>
      <Set name="maxIdleTime">30000</Set>
      <Set name="keystore">path-to-keystore</Set>
      <Set name="password">password</Set>
      <Set name="keyPassword"key-password</Set>
      <Set name="truststore">path-to-keystore</Set>
      <Set name="trustPassword"trust-password</Set>
    </New>
  </Arg>
</Call>
```

   Passwords above can be plain text or obfuscated as follows:

   `java -classpath lib/jetty-6.1.7.jar:lib/jetty-util-6.1.7.jar org.mortbay.jetty.security.Password` *plaintext-password*

   Then copy and paste the output line starting with `OBF:` (including the `OBF:` part) into this `jetty.xml` config file.

2. **Restart OASM.**

   - On a system running Oracle Solaris, type:

     # **`svcadm restart sasm`**

   - On a system running Oracle Linux, type:

```
# /opt/SUNWsasm/bin/sasm stop-instance
# /opt/SUNWsasm/bin/sasm start-instance
```

3. **Verify the SSL setup by accessing the following URL from a browser:**

   `https://`*ASR-Manager-name*`/asr`

## ▼ Register Compute Servers With Oracle Solaris 11 or Database Domains to ASR Manager

Use this procedure to register compute servers with Oracle Solaris 11 or Database Domains to the ASR Manager.

1. **Log in to the compute server as superuser.**
2. **Confirm that the** `asr-notify` **service is working:**

   ```
   # svcs asr-notify
   ```

   - If you see this message:

     ```
     svcs: Pattern 'asr-notify' doesn't match any instances
     ```

     then confirm that the `asr-notify` service is installed:

     ```
     # pkg list asr-notify
     ```

     If you see this message:

     ```
     pkg list: no packages matching 'asr-modify' installed
     ```

     then install the `asr-notify` service:

     ```
     # pkg install system/fault-management/asr-notify
     ```

     Enter the `svcs asr-notify` command again to confirm that the `asr-notify` service is working.

   - If you see this message:

```
# svcs asr-notify
STATE     STIME       FMRI
online    16:06:05    svc:/system/fm/asr-notify:default
```

   then the `asr-notify` service is installed and is working properly

3. **Register the ASR manager using this command:**

   # **asradm register -e http://***asr-manager-host***:***port-number***/asr**

   Example:

   # **asradm register -e http://asrmanager1.mycompany.com:8777/asr**

   Screens are displayed asking for your Oracle Support account name and password. After entering your Oracle Support account name and password, a notification is displayed, indicating that your registration is complete:

```
Enter Orcle SSO User Name:
Enter password:

Registration complete.
```

4. **Run this command:**

   # **asradm list**

   Example output:

```
PROPERTY VALUE
Status Successfully Registered with ASR manager
System Id system-identification-number
Asset Id asset-identification-number
User username
Endpoint URL http://asr-manager-host:port-number/asr
```

   The registration of the ASR Manager is complete.

5. **Repeat these instructions for all compute servers with Oracle Solaris 11 or Database Domains in your SuperCluster system.**

6. **When you have completed the configuration of all the SuperCluster assets, approve and verify contacts in MOS.**

   See "Approve and Verify ASR Asset Activation" on page 24 for those instructions.

# ▼ Approve and Verify ASR Asset Activation

Perform this task after completing these tasks:

- "Configure SNMP Trap Destinations for Storage Servers" on page 16

- "Configure ASR on the ZFS Storage Appliance" on page 18
- "Configure ASR on the Compute Servers (Oracle ILOM)" on page 14
- "Configuring ASR on the Compute Servers (Oracle Solaris 11)" on page 21

For more information on the process, see ASR MOS 5.3+ Activation Process (Doc ID 1329200.1).

---

**Note –** If a subscriber has not been set up, then the subsequent Auto Service Request activation fails.

---

1. **On the standalone system where ASR Manager is running, run this command to verify the status of your system assets:**

   `list_asset`

   The ASR assets in your SuperCluster system are listed, including compute servers, storage servers, and ZFS storage controllers.

2. **Log in to My Oracle Support (`https://support.oracle.com`).**

3. **In the My Oracle Support Dashboard, click the More... tab, then click Settings from the menu.**

4. **In the Settings pane on the left of the window, select Pending ASR Activations (located under the Administrative sub menu).**

   A complete list of all qualified ASR assets that are awaiting approval are displayed.

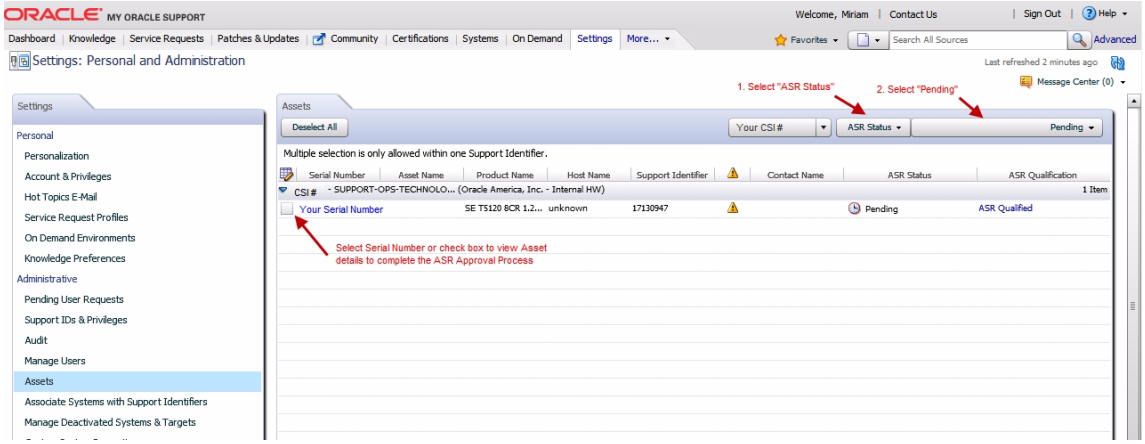


By default, all support identifiers that you are associated with are displayed. If this list of assets is long, you can limit the display to show only assets associated to one support identifier. You can also search for an asset's serial number.

**Tip –** For each component in the SuperCluster system, you should see two host names associated with each serial number. If you see only the Oracle ILOM host name, that means that you did not activate ASR for that component. If you see more than two host names associated with each serial number, you might need to request support. To do this, open a hardware SR with "Problem Category" set to "My - Auto Service Request (ASR) Installation and Configuration Issues."

5. **Click the asset's serial number.**

   If there is any missing asset information the ASR Activation window is displayed, prompting you to enter the missing information.


**Note –** ASR Host name is updated when an activation request is sent to Oracle from the ASR software on the asset. (For example, from the `asr activate_asset` command on the ASR Manager.)

Required fields for ASR asset activation are:

- Contact Name: You can only select a name associated with the support identifier. Click the drop-down menu to see the list of available names.

  A contact must have the Create SR privilege for the asset's support identifier.
- Street Address 1: Type the street address for the asset.

**Note –** By default, all support identifiers that you are associated with are displayed. If this list of assets is long, you can limit the display to show only assets associated to one support identifier. You can also search for an asset's serial number.

- Country: Select the asset's country location from the drop-down menu.
- ZIP/Postal Code: type the ZIP/postal code for the asset's location. If there is no postcode insert "-".

- Distribution Email List: Add email addresses that receive all ASR mail notifications. Separate multiple email addresses with a comma. For example:

  asr-notifications-1@mycompany.com, asr-notifications-2@mycompany.com

  ASR sends email to the Contact's email address and the Distribution Email List, if provided. This is a useful feature if your organization has a team that should be informed about Service Requests created by ASR.

**6. Click the Approve button to complete the ASR activation.**

---

**Note –** A system asset must be in an active ASR state in My Oracle Support for Service Request autocreate to work.

---

**7. To confirm that ASR can send information to the transport server, run:**

# **asradm send test** *email-address@company.com*

This command sends a test alert e-mail to the e-mail address.

# Tuning SuperCluster M6-32

These topics describe the utility (`ssctuner`) used to tune SuperCluster M6-32. For the latest information about `ssctuner`, see the `README` file installed with the utility.



### Related Information

- For more information about SMF services on the Oracle Solaris OS, see the *Oracle Solaris System Administration Guide: Common System Management Tasks* at:

  `http://docs.oracle.com/cd/E23824_01/html/821-1451/hbrunlevels-25516.html#scrolltoc`

## `ssctuner` Overview

The `ssctuner` utility is a small set of Perl and Korn shell scripts and configuration files that run on SuperCluster Oracle Solaris 10 and Oracle Solaris 11 global zones. By default, `ssctuner` is installed and enabled at installation time.

The utility runs in realtime as an SMF service to monitor and tune ndd parameters and various system configuration parameters including these files:

- `/etc/system`
- `/kernel/drv/sd.conf`
- `/kernel/drv/ssd.conf`
- `/etc/inet/ntp.conf`

The utility also periodically checks for the use of DISM or suboptimal NFS mount options.

By default, the utility runs every two hours and modifies parameters as needed.

The utility also checks every two minutes to see if there are any virtual disk devices that were in a degraded state and have come back online, and if so, clears that zpool.

---

**Note –** If you manually tune a parameter for which `ssctuner` requires a different value, `ssctuner` sets the value of that parameter back to what `ssctuner` requires and logs the changes at this interval check. If you must control one or more of the parameters `ssctuner` manages, consider turning off those specific components rather than disabling `ssctuner` completely. See "Change ssctuner Properties and Disable Features" on page 32.

---

**Note –** Oracle Solaris 11 must be at SRU 11.4 or later, or `ssd.conf/sd.conf` settings might cause panics.

---

**Note –** Do not set `ndd` parameters through another SMF service or `init` script. `ssctuner` must manage the `ndd` parameters.

---

### Related Information

- "Monitor ssctuner Activity" on page 30
- "View Log Files" on page 31
- "Change ssctuner Properties and Disable Features" on page 32
- "Install ssctuner" on page 35
- "Enable ssctuner" on page 36

## ▼ Monitor `ssctuner` Activity

- **View `ssctuner` activity.**

```
# svcs -l ssctuner
```

### Related Information

- "ssctuner Overview" on page 29
- "View Log Files" on page 31
- "Change ssctuner Properties and Disable Features" on page 32
- "Install ssctuner" on page 35
- "Enable ssctuner" on page 36

## ▼ View Log Files

1. **View the `ssctuner` service log.**

   `ssctuner` writes messages to `syslog` and to the `ssctuner` service log. Those messages are tagged as `ssctuner` and might point to other file locations for more information.

```
# svcs -x ssctuner
svc:/site/application/sysadmin/ssctuner:default (ssctuner for Oracle
SuperCluster)
State: online since September 28, 2012 07:30:15 AM PDT
See: ssctuner(1)
See: /var/svc/log/site-application-sysadmin-ssctuner:default.log
Impact: None.

# more /var/svc/log/site-application-sysadmin-ssctuner\:default.log
[ Sep 28 07:30:00 Disabled. ]
[ Sep 28 07:30:00 Rereading configuration. ]
[ Sep 28 07:30:10 Enabled. ]
[ Sep 28 07:30:10 Executing start method ("/opt/oracle.supercluster/ssctuner.ksh
start"). ]
ssctuner local0.notice success: Saved rollback for : /etc/system
ssctuner local0.notice success: Saved ndd rollback.
ssctuner local0.notice success: Saved rollback for : /kernel/drv/sd.conf
ssctuner local0.notice success: enabled, version 0.99e. daemon PID= 14599
[ Sep 28 07:30:15 Method "start" exited with status 0. ]
ssctuner local0.notice success: daemon executing
ssctuner local0.notice success: Changes made to /etc/system
ssctuner local0.notice success: Changes made to /kernel/drv/sd.conf
```

2. **View** `ssctuner` **messages in** `/var/adm`**.**

```
# grep -i ssctuner /var/adm/messages
Sep 28 07:30:10 etc6cn04 ssctuner: [ID 702911 local0.notice] success: Saved
rollback for : /etc/system
Sep 28 07:30:10 etc6cn04 ssctuner: [ID 702911 local0.notice] success: Saved ndd
rollback.
Sep 28 07:30:10 etc6cn04 ssctuner: [ID 702911 local0.notice] success: Saved
rollback for : /kernel/drv/sd.conf
Sep 28 07:30:15 etc6cn04 ssctuner: [ID 702911 local0.notice] success: enabled,
version 0.99e. daemon PID= 14599
Sep 28 07:30:15 etc6cn04 ssctuner: [ID 702911 local0.notice] success: daemon
executing
Sep 28 07:30:15 etc6cn04 ssctuner: [ID 702911 local0.notice] success: Changes
made to /etc/system
Sep 28 07:30:15 etc6cn04 ssctuner: [ID 702911 local0.notice] success: Changes
made to /kernel/drv/sd.conf
```

### Related Information

- "ssctuner Overview" on page 29
- "Monitor ssctuner Activity" on page 30
- "Change ssctuner Properties and Disable Features" on page 32
- "Install ssctuner" on page 35
- "Enable ssctuner" on page 36

## ▼ Change `ssctuner` Properties and Disable Features

**Caution –** Do not perform this procedure without Oracle Support approval. Changing properties or disabling `ssctuner` features can have unpredictable consequences.

Changing certain `ssctuner` properties such as `EMAIL_ADDRESS` and disk or memory usage warning levels might be advantageous in some environments.

1. **List the** `ssctuner` **properties to identify the property you want to change.**

```
# svccfg -s ssctuner listprop 'ssctuner_vars/*'
ssctuner_vars/CRIT_THREADS_FIX          boolean     true
ssctuner_vars/CRIT_THREADS_NONEXA       boolean     false
ssctuner_vars/DISK_SPACE_CHECK          boolean     true
ssctuner_vars/DISK_USAGE_CRIT           integer     90
ssctuner_vars/DISK_USAGE_WARN           integer     85
ssctuner_vars/DISM_CHECK                boolean     true
ssctuner_vars/EMAIL_ADDRESS             astring     root@localhost
ssctuner_vars/EMAIL_MESSAGES            boolean     true
ssctuner_vars/FORCELOAD_VDC             boolean     false
ssctuner_vars/INTRD_DISABLE             boolean     true
ssctuner_vars/ISCSI_TUNE                boolean     true
ssctuner_vars/MAJOR_INTERVAL            integer     120
ssctuner_vars/MEM_USAGE_CRIT            integer     97
ssctuner_vars/MEM_USAGE_WARN            integer     94
ssctuner_vars/MINOR_INTERVAL            integer     2
ssctuner_vars/NDD_TUNE                  boolean     true
ssctuner_vars/NFS_CHECK                 boolean     true
ssctuner_vars/NFS_EXCLUDE               astring
ssctuner_vars/NFS_INCLUDE               astring
ssctuner_vars/NTPCONF_TUNE              boolean     true
ssctuner_vars/POWERADM_DISABLE          boolean     true
ssctuner_vars/SDCONF_TUNE               boolean     true
ssctuner_vars/SERD_THRESHOLD_TUNE       boolean     true
ssctuner_vars/SSDCONF_TUNE              boolean     true
ssctuner_vars/SYSLOG_DUP_SUPPRESS_HOURS integer     8
ssctuner_vars/SYSTEM_TUNE               boolean     true
ssctuner_vars/ZPOOL_FIX                 boolean     true
ssctuner_vars/ZPOOL_NAME_CUST           astring
```

2. **Use the** `svccfg` **command to change property settings.**

   These are examples of properties you might need to change:

   - Configure the system so that critical messages are sent to your email address.

```
~# svccfg -s ssctuner setprop ssctuner_vars/EMAIL_ADDRESS=
"my_name@mycorp.com"
```

   - Change the disk (/ and zone roots) usage warning level to 80%.

```
~# svccfg -s ssctuner setprop ssctuner_vars/DISK_USAGE_WARN=80
```

- Enable thread priority changing for non-exa Oracle DB domains:.

```
~# svccfg -s ssctuner setprop ssctuner_vars/CRIT_THREADS_NONEXA=
true
```

- Enable zpool check and repair of vdisk zpools that are not generated by the SuperCluster installer.

```
~# svccfg -s ssctuner setprop ssctuner_vars/ZPOOL_NAME_CUST=
my_vdisk_pool
```

- Exclude NFS mounts from warning mechanisms.

```
~# svccfg -s ssctuner setprop ssctuner_vars/NFS_EXCLUDE=
'mount_name_or_device'
```

- Include NFS mounts in warning mechanism (overrides exclude).

```
~# svccfg -s ssctuner setprop ssctuner_vars/NFS_INCLUDE=
'mount_name_or_device'
```

- Disable all NFS mount warnings (not recommended).

```
~# svccfg -s ssctuner setprop ssctuner_vars/NFS_CHECK=false
```

The `NFS_EXCLUDE`, `NFS_INCLUDE` and `ZPOOL_NAME_CUST` properties must be simple strings but you can use simple regular expressions.

If you need the flexibility of regular expressions, be extremely careful to double quote the expressions. Also verify that the `ssctuner` service comes back after restarting and no errors are in the SMF log file.

**3. Restart the SMF service for changes to take effect.**

```
# svcadm restart ssctuner
```

4. **Ensure that the `ssctuner` service is enabled and no error messages are reported.**

   If you changed a property using an incorrect syntax, the service does not come back. If this happens, identify the offending property that you must fix:

```
# grep -i parameter
/var/svc/log/site-application-sysadmin-ssctuner:default.log
```

   After making any corrections or changes, repeat Step 3.

### Related Information

- "ssctuner Overview" on page 29
- "Monitor ssctuner Activity" on page 30
- "View Log Files" on page 31
- "Install ssctuner" on page 35
- "Enable ssctuner" on page 36

## ▼ Install `ssctuner`

By default, `ssctuner` is installed and running. If for some reason `ssctuner` is not installed, use this procedure to install it.

1. **Install the `ssctuner` package.**

   Use the Oracle Solaris package command and package name based on the version of the OS.

   - **Oracle Solaris 10 OS**:

**Note –** You must be in the directory where the `ORCLssctuner` package tree resides.

```
# pkgadd -d ORCLssctuner
```

   - **Oracle Solaris 11 OS**:

**Note –** You must have the latest exa-family repository set as a publisher.

```
# pkg install ssctuner
```

2. **Verify the package installation.**
   - **Oracle Solaris 10 OS**:

```
# pkginfo ORCLssctuner
```

   - **Oracle Solaris 11 OS**:

```
# pkg info ssctuner
```

3. **Verify that the `ssctuner` service is automatically started after the package installation.**

```
# svcs ssctuner
```

   If the service does not transition to an online state after a minute or two, check the service log file. See "View Log Files" on page 31.

4. **Reboot the OS.**

   When `ssctuner` changes configuration files, you must reboot the OS for those changes to take effect.

### Related Information

- "ssctuner Overview" on page 29
- "Monitor ssctuner Activity" on page 30
- "View Log Files" on page 31
- "Change ssctuner Properties and Disable Features" on page 32
- "Enable ssctuner" on page 36

## ▼ Enable `ssctuner`

Usually `ssctuner` is running. If for some reason `ssctuner` is not running, use this procedure to enable it.

1. **Enable `ssctuner`.**

```
# svcadm enable ssctuner
```

2. **Verify that the `ssctuner` service started.**

```
# svcs ssctuner
```

   If the service does not transition to an online state after a minute or two, check the service log file. See "View Log Files" on page 31.

3. **Check the `/var/adm/messages` log file to see if `ssctuner` changed any configuration file settings.**

   See "View Log Files" on page 31.

   If configuration settings changed, you must reboot the OS for the changes to take effect. If settings did not change, you do not need to reboot the OS.

### Related Information

- "ssctuner Overview" on page 29
- "Monitor ssctuner Activity" on page 30
- "View Log Files" on page 31
- "Change ssctuner Properties and Disable Features" on page 32
- "Install ssctuner" on page 35

# Configuring CPU and Memory Resources (`osc-setcoremem`)

This section describes how to configure Oracle SuperCluster CPU and memory resources using `osc-setcoremem`.

Prior to the Oracle SuperCluster July 2015 quarterly update, you configured CPU and memory resources using the `setcoremem` tool (in some cases called `setcoremem-t4`, `setcoremem-t5`, or `setcoremem-m6` based on the SuperCluster model).

As of the Oracle SuperCluster July 2015 quarterly update, you use the `osc-setcoremem` tool, and the previous commands are no longer available.

The `osc-setcoremem` tool is supported on all SuperCluster systems that run SuperCluster 1.*x* software with the July 2015 quarterly update, and SuperCluster 2.*x* software.

Use these topics to change CPU and memory allocations for domains using the CPU/Memory tool called `osc-setcoremem`.

| Description | Links |
|---|---|
| Learn about the CPU/Memory tool. | "osc-setcoremem Overview" on page 40<br>"Minimum and Maximum Resources (Dedicated Domains)" on page 41 |
| Find out if SuperCluster resources can be modified using the CPU/Memory tool. | "Supported Domain Configurations" on page 43 |
| Plan CPU and memory allocations. | "Plan CPU and Memory Allocations" on page 44 |
| Identify domain configurations. | "Display the Current Domain Configuration (osc-setcoremem)" on page 47<br>"Display the Current Domain Configuration (osc-setcoremem)" on page 47<br>"Access osc-setcoremem Log Files" on page 65<br>"View the SP Configuration" on page 70 |

| Description | Links |
|---|---|
| Configure domain CPU and memory resources at the socket or core level. | "Change CPU/Memory Allocations (Socket Granularity)" on page 51<br>"Change CPU/Memory Allocations (Core Granularity)" on page 55 |
| Configure domain CPU and memory resources so that some resources are parked. | "Park Cores and Memory" on page 60 |
| Access information about previous executions of `osc-setcoremem`. | "Access osc-setcoremem Log Files" on page 65<br>"View the SP Configuration" on page 70 |
| Revert to or remove a previous CPU/memory configuration. | "Revert to a Previous CPU/Memory Configuration" on page 72 |

# `osc-setcoremem` Overview

SuperCluster compute server CPU and memory resources are initially allocated during installation as defined by your configuration. CPU sockets are assigned to domains in the same proportion as IB HCAs. Memory is assigned in the same proportions.

The `osc-setcoremem` tool enables you to migrate CPU cores and memory resources between dedicated domains, and from dedicated domains to CPU and memory repositories for the use of IO domains.

These points provide important information related to the use of `osc-setcoremem`:

- The final CPU and memory layout for a dedicated domain is optimized for locality to minimize accesses to non-local resources.
- The granularity of CPU and memory migration is 1 core and 16GB.
- Empty dedicated domains (domains with zero cores and zero memory) are not supported.
- The tool tracks resource allocation and ensures that the selections you make are valid. See "Minimum and Maximum Resources (Dedicated Domains)" on page 41.
- Affected dedicated domains must be rebooted after any change.

The tool enables you to change the CPU and memory allocations in one of two levels of granularity:

- **Socket granularity** – The tool automatically allocates each domain a minimum of one socket, then enables you to allocate remaining sockets to the domains. See "Change CPU/Memory Allocations (Socket Granularity)" on page 51.

- **Core granularity** – The tool automatically allocates each domain a minimum number of cores, then enables you to allocate additional cores in one-core increments. See "Change CPU/Memory Allocations (Core Granularity)" on page 55.

If you configure the CPU and memory resources so that some resources are not allocated to any domain, those unallocated resources are parked. Parked resources are placed in a logical CPU and memory repository and are available for I/O Domains. See "Park Cores and Memory" on page 60.

You can park resources from dedicated domains anytime, but you cannot move parked resources to dedicated domains once I/O Domains are created.

Also see "Supported Domain Configurations" on page 43.

### Related Information

- "Minimum and Maximum Resources (Dedicated Domains)" on page 41
- "Supported Domain Configurations" on page 43
- "Plan CPU and Memory Allocations" on page 44
- "Display the Current Domain Configuration (osc-setcoremem)" on page 47
- "Display the Current Domain Configuration (ldm)" on page 49
- "Change CPU/Memory Allocations (Socket Granularity)" on page 51
- "Change CPU/Memory Allocations (Core Granularity)" on page 55
- "Park Cores and Memory" on page 60

# Minimum and Maximum Resources (Dedicated Domains)

The tool tracks resource allocation and ensures that the selections you make are valid. This section describes how the minimum and maximum resources are determined.

This table summarizes the minimum resource requirements for dedicated domains on SuperCluster M6-32:

| Configuration | Minimum Resource Requirements |
|---|---|
| Dedicated Domain with 1 HCA | 2 cores / 32GB memory |

| Configuration | Minimum Resource Requirements |
|---|---|
| Dedicated Domain with 2 HCAs | 4 cores / 64GB memory |
| Dedicated Domain with 4 HCAs | 8 cores / 128GB memory |

The minimum amount of CPU resource that can be assigned to a dedicated domain is determined by the number of IB and 10GbE devices in the domain (2 cores are required per IB HCA)

The minimum amount of memory that can be assigned to a dedicated domain is determined as follows:

- The number of IB and 10GbE devices in the domain (2 16GB memory granules are required per IB HCA)
- The number of cores assigned to the domain (one 16GB granule in the same locality group is required per 4 additional cores)

The maximum amount of CPU resource that can be assigned to a dedicated domain is determined by the amount of resources available after taking these points into account:

- Resources already assigned to other dedicated domains
- Required minimal resource for dedicated domains with no resource yet assigned

The maximum amount of memory resources that can be assigned to a dedicated domain is determined by the amount of resources available after taking these points into account:

- Resources already assigned to other dedicated domains
- Required minimum resources for dedicated domains with no resource yet assigned
- The requirement that for each dedicated domain a memory granule footprint is placed in all locality groups with allocated cores

### Related Information


- "osc-setcoremem Overview" on page 40
- "Supported Domain Configurations" on page 43
- "Plan CPU and Memory Allocations" on page 44
- "Display the Current Domain Configuration (osc-setcoremem)" on page 47
- "Display the Current Domain Configuration (ldm)" on page 49
- "Change CPU/Memory Allocations (Socket Granularity)" on page 51
- "Change CPU/Memory Allocations (Core Granularity)" on page 55
- "Park Cores and Memory" on page 60

# Supported Domain Configurations

Use this table to identify your SuperCluster configuration, then review the supported resource allocation activities.

**Note –** A dedicated domain can be any Application or Database Domain that is not associated with I/O Domains. For more information about the different types of SuperCluster domains, refer to the section called Understanding the Software Configurations in *Oracle SuperCluster M6-32 Owner's Guide: Overview*.

| Domain Configuration | Supported Resource Allocation Activities | Links |
|---|---|---|
| All domains are dedicated domains | Plan how the CPU and memory resources are allocated to the domains. | "Plan CPU and Memory Allocations" on page 44 |
| | Reallocate all of the resources across domains at the socket or core level (a reboot is required if primary domain resources are changed). | "Change CPU/Memory Allocations (Socket Granularity)" on page 51<br>"Change CPU/Memory Allocations (Core Granularity)" on page 55 |
| | Remove (park) resources from dedicated domains for licensing purposes.<br>**Note -** Parked resources are not available for use by any domains. | "Park Cores and Memory" on page 60 |
| | Revert to a previous resource configuration. | "Revert to a Previous CPU/Memory Configuration" on page 72 |
| | Remove a CPU/memory configuration. | "Remove a CPU/Memory Configuration" on page 73 |

| Domain Configuration | Supported Resource Allocation Activities | Links |
|---|---|---|
| Mixed domains – some are dedicated, some are Root Domains | Activities you can only perform at initial installation, before any I/O Domains are created: | |
| | • Plan how the CPU and memory resources are allocated to the domains. | "Plan CPU and Memory Allocations" on page 44 |
| | • Reallocate all of the resources across domains at the socket or core level (a reboot is required if primary domain resources are changed). | "Change CPU/Memory Allocations (Socket Granularity)" on page 51<br>"Change CPU/Memory Allocations (Core Granularity)" on page 55 |
| | • Revert to a previous allocation configuration. | "Revert to a Previous CPU/Memory Configuration" on page 72 |
| | Activities you can perform anytime: | |
| | • Configure resources for I/O Domains. | Refer to the *I/O Domain Administration Guide*. |
| | • Move resources from dedicated domains so that the resources are available to I/O Domains. | "Park Cores and Memory" on page 60 |
| | • Move resources between dedicated domains. | "Change CPU/Memory Allocations (Socket Granularity)" on page 51<br>"Change CPU/Memory Allocations (Core Granularity)" on page 55 |
| | • Remove a CPU/memory configuration. | "Remove a CPU/Memory Configuration" on page 73 |

### Related Information

- "osc-setcoremem Overview" on page 40
- "Plan CPU and Memory Allocations" on page 44
- "Display the Current Domain Configuration (osc-setcoremem)" on page 47
- "Display the Current Domain Configuration (ldm)" on page 49
- "Change CPU/Memory Allocations (Socket Granularity)" on page 51
- "Change CPU/Memory Allocations (Core Granularity)" on page 55
- "Park Cores and Memory" on page 60

# ▼ Plan CPU and Memory Allocations

There are two main approaches to modifying resource allocations:

- **All resources allocated** – You move resources from domains to other domains, and ensure that all resources are allocated.
- **Some resources are unallocated** – You allocate less than the maximum available cores and memory for a compute node. Any unused cores are considered *parked* cores and are not counted for licensing purposes. However, parked cores are added to the logical CPU and memory repository. If you have Root Domains, you can later allocate the repository resources to I/O Domains. See "Park Cores and Memory" on page 60.

Depending on which command you use to view domain resources, you might need to convert socket, core, and VCPU values.

| | SuperCluster M6-32 | SuperCluster T5-8 |
|---|---|---|
| 1 socket = | 12 cores (96 VCPUs) | 16 cores (128 VCPUs) |
| 1 core = | 8 VCPUs | 8 VCPUs |

1. **Identify the resource configuration for each compute node.**

   See one of these procedures:

   - "Display the Current Domain Configuration (osc-setcoremem)" on page 47
   - "Display the Current Domain Configuration (ldm)" on page 49

   In this example, one compute node on a SuperCluster M6-32 has two dedicated domains and two Root Domains.

| Domain | Domain Type | Cores | Memory (GB) |
|---|---|---|---|
| `primary` | Dedicated | 18 | 1536 |
| `ssccn3-dom1` | Dedicated | 30 | 2560 |
| `ssccn3-dom2` | Root | n/a | n/a |
| `ssccn3-dom3` | Root | n/a | n/a |
| Unallocated Resources | | 45 | 4048 |

2. **Add the domain resources together to determine the total number of resources.**

   Calculating the total amount of CPU and memory resources gives you a starting point for determining your resource plan.

   While identifying resources, keep these points in mind:

   - **Root Domain resources** – Are a small amount of resources that are reserved for the exclusive use of Root Domains. Do not factor these resources into your plan.

- **Unallocated resources** – These resources are placed in the logical CPU and memory repositories when Root Domains are created, or by leaving some resources unallocated when you use the `osc-setcoremem` command.

In this example, the resources for the dedicated domains and the unallocated resources are summed to provide total resources. The Root Domain resources are not included in total resources.

| Domain | Domain Type | Cores | Memory (GB) |
|---|---|---|---|
| `primary` | Dedicated | 18 | 1536 |
| `ssccn3-dom1` | Dedicated | 30 | 2560 |
| `ssccn3-dom2` | Root | n/a | n/a |
| `ssccn3-dom3` | Root | n/a | n/a |
| Unallocated Resources | | 45 | 4048 |
| Total Resources | | 93 | 8144 |

3. **Based on your site requirements, and the type and number of domains on SuperCluster, decide how to allocate CPU and memory for each domain.**

   In this example, 12 cores and 1 TB memory are parked from the primary domain, and 18 cores and 1536 GB memory are parked from the `ssccn3-dom1` domain.

   The total resources for before and after columns should match. This check ensures that all resources are accounted for in your plan.

| Domain | Domain Type | Cores Before | Cores After | Memory Before (GB) | Memory After (GB) |
|---|---|---|---|---|---|
| `primary` | Dedicated | 18 | 6 | 1536 | 512 |
| `ssccn3-dom1` | Dedicated | 30 | 12 | 2560 | 1024 |
| `ssccn3-dom2` | Root | n/a | n/a | n/a | n/a |
| `ssccn3-dom3` | Root | n/a | n/a | n/a | n/a |
| Unallocated Resources | | 45 | | 4048 | |
| Total Resources | | 93 | 93 | 8144 | 8144 |

4. **Consider your next action:**
   - Change resource allocations at the socket granularity level.

     See "Change CPU/Memory Allocations (Socket Granularity)" on page 51
   - Change resource allocations at the core granularity level.

     . See "Change CPU/Memory Allocations (Core Granularity)" on page 55

- Increase unallocated resources.

  See "Park Cores and Memory" on page 60

### Related Information

- "osc-setcoremem Overview" on page 40
- "Supported Domain Configurations" on page 43
- "Display the Current Domain Configuration (osc-setcoremem)" on page 47
- "Display the Current Domain Configuration (ldm)" on page 49
- "Change CPU/Memory Allocations (Socket Granularity)" on page 51
- "Change CPU/Memory Allocations (Core Granularity)" on page 55
- "Park Cores and Memory" on page 60

# ▼ Display the Current Domain Configuration (`osc-setcoremem`)

This procedure describes how to display a compute node domain configuration using the `osc-setcoremem` command.

---

**Note –** Alternatively, you can use `ldm` commands to get similar information See "Display the Current Domain Configuration (ldm)" on page 49.

---

1. **Log in as superuser on the compute node's control domain.**
2. **Use the `osc-setcoremem` command to view domains and resources.**

---

**Note –** If you don't want to continue to use the `osc-setcoremem` command to change resource allocations, enter CTL-C at the first prompt.

---

Example:

```
# /opt/oracle.supercluster/bin/osc-setcoremem

                              osc-setcoremem
                    v2.0  built on Aug 27 2015 23:09:35


 Current Configuration: SuperCluster Fully-Populated M6-32 Base
```

```
 +------------------------------+-------+--------+-----------+--- MINIMUM ----+
 | DOMAIN                       | CORES | MEM GB |   TYPE    | CORES | MEM GB |

+------------------------------+-------+--------+-----------+-------+--------+
 | primary                      |    24 |   2048 | Dedicated |     2 |     32 |
 | ssccn3-dom1                  |    24 |   2048 | Dedicated |     2 |     32 |
 | ssccn3-dom2                  |     1 |     16 |   Root    |     1 |     16 |
 | ssccn3-dom3                  |     2 |     32 |   Root    |     2 |     32 |

+------------------------------+-------+--------+-----------+-------+--------+
 | unallocated or parked        |    45 |   4048 |    --     |    -- |   --   |

+------------------------------+-------+--------+-----------+-------+--------+

 [Note] Following domains will be skipped in this session.

 Root Domains
 ------------
 ssccn3-dom2
 ssccn3-dom3


 CPU allocation preference:

        1. Socket level
        2. Core level

 In case of Socket level granularity, proportional memory capacity is
  automatically selected for you.

 Choose Socket or Core level [S or C] <CTL-C>
```

### Related Information


- "osc-setcoremem Overview" on page 40
- "Supported Domain Configurations" on page 43
- "Plan CPU and Memory Allocations" on page 44
- "Display the Current Domain Configuration (ldm)" on page 49
- "Change CPU/Memory Allocations (Socket Granularity)" on page 51
- "Change CPU/Memory Allocations (Core Granularity)" on page 55
- "Park Cores and Memory" on page 60

# ▼ Display the Current Domain Configuration (`ldm`)

This procedure describes how to display a compute node domain configuration using a series of `ldm` commands.

---

**Note –** Alternatively, you can use the `osc-setcoremem` command to get similar information See "Display the Current Domain Configuration (osc-setcoremem)" on page 47.

---

1. **Log in as `root` on the compute node's control domain.**

2. **Identify which domains are Root Domains:**

   Root Domains are identified by `IOV` in the `STATUS` column.

   In this example, `ssccn3-dom2` and `ssccn3-dom3` are Root Domains. The other domains are dedicated domains.

```
# ldm list-io | grep BUS
NAME                                      TYPE   BUS      DOMAIN     STATUS
pci_32                                    BUS    pci_32   primary
pci_33                                    BUS    pci_33   primary
pci_34                                    BUS    pci_34   primary
pci_35                                    BUS    pci_35   primary
pci_36                                    BUS    pci_36   ssccn3-dom2  IOV
pci_37                                    BUS    pci_37   ssccn3-dom2  IOV
pci_38                                    BUS    pci_38   ssccn3-dom2  IOV
pci_39                                    BUS    pci_39   ssccn3-dom2  IOV
pci_40                                    BUS    pci_40   ssccn3-dom1
pci_41                                    BUS    pci_41   ssccn3-dom1
pci_42                                    BUS    pci_42   ssccn3-dom1
pci_43                                    BUS    pci_43   ssccn3-dom1
pci_44                                    BUS    pci_44   ssccn3-dom3  IOV
pci_45                                    BUS    pci_45   ssccn3-dom3  IOV
pci_46                                    BUS    pci_46   ssccn3-dom3  IOV
pci_47                                    BUS    pci_47   ssccn3-dom3  IOV
```

3. **View domains and resource allocation information.**

   In this example, `ssccn3-dom2` and `ssccn3-dom3` are Root Domains (from Step 2). The resources listed for Root Domains only represent the resources that are reserved for the Root Domain itself. Parked resources are not displayed.

```
# ldm list
NAME              STATE      FLAGS   CONS    VCPU  MEMORY   UTIL  NORM  UPTIME
primary           active     -n-cv-  UART    192   2095872M 0.1%  0.1%  12h 28m
ssccn3-dom1       active     -n----  5001    192   2T       0.1%  0.1%  12h 25m
ssccn3-dom2       active     -n----  5002    8     16G      0.1%  0.1%  2d 23h 34m
ssccn3-dom3       active     -n--v-  5003    16    32G      0.1%  0.1%  2d 23h 34m
```

4. **View the amount of parked resources.**

   In this example, the first command line reports the number of cores in the logical CPU repository. The second command line reports the amount of memory in the memory repository.

```
# ldm list-devices -p core | grep cid | wc -l
      45


# ldm list-devices memory
MEMORY
    PA                   SIZE
    0x100000000000       1008G
    0x180000000000       1T
    0x300000000000       1008G
    0x380000000000       1008G
```

### Related Information

- "osc-setcoremem Overview" on page 40
- "Supported Domain Configurations" on page 43
- "Plan CPU and Memory Allocations" on page 44
- "Display the Current Domain Configuration (osc-setcoremem)" on page 47
- "Change CPU/Memory Allocations (Socket Granularity)" on page 51
- "Change CPU/Memory Allocations (Core Granularity)" on page 55
- "Park Cores and Memory" on page 60

# ▼ Change CPU/Memory Allocations (Socket Granularity)

Perform this procedure on each compute node to change its CPU and memory resource allocation at the socket granularity level.

---

**Note –** To find out if you can perform this procedure, see "Supported Domain Configurations" on page 43.

---

The tool makes these changes:

- Automatically detects Root Domains.
- Calculates the minimum and maximum resources for all domains, and only enables you to select valid quantities.
- Modifies domain resources according to the choices you make.
- Automatically assigns memory capacity in the same proportion to CPU resources.
- (If needed) Stops nonprimary domains.
- (If needed) Reboots the primary domain with new resources.
- (If needed) Brings up nonprimary domains with new resources.

In this example, one socket and 1 TB memory are removed from the primary domain and allocated to `ssccn3-dom1`.

This table shows the allocation plan (see "Plan CPU and Memory Allocations" on page 44).

| Domain | Domain Type | Sockets Before | Sockets After | Memory Before (GB) | Memory After (GB) |
|---|---|---|---|---|---|
| primary | Dedicated | 2 | 1 | 2048 | 1024 |
| ssccn3-dom1 | Dedicated | 2 | 3 | 2048 | 3072 |
| ssccn3-dom2 | Root | n/a | n/a | n/a | n/a |
| ssccn3-dom3 | Root | n/a | n/a | n/a | n/a |
| Unallocated resources | | 45 | 45 | 4048 | 4048 |
| Total resources | | 49 | 39 | 4144 | 4144 |

1. **Log in as superuser on the compute node's control domain.**
2. **Ensure that applications are shut down and that there is no production activity.**
3. **Activate any inactive domains using the** `ldm bind` **command.**
   The tool does not continue if any inactive domains are present.
4. **Run** `osc-setcoremem` **to reconfigure the resources.**
   Respond when prompted. Press Enter to select the default value.

```
# /opt/oracle.supercluster/bin/osc-setcoremem

                          osc-setcoremem
                    v2.0  built on Aug 27 2015 23:09:35


 Current Configuration: SuperCluster Fully-Populated M6-32 Base

 +------------------------------+-------+--------+-----------+--- MINIMUM ----+
 | DOMAIN                       | CORES | MEM GB |   TYPE    | CORES | MEM GB |
 +------------------------------+-------+--------+-----------+-------+--------+
 | primary                      |    24 |   2048 | Dedicated |     2 |     32 |
 | ssccn3-dom1                  |    24 |   2048 | Dedicated |     2 |     32 |
 | ssccn3-dom2                  |     1 |     16 |   Root    |     1 |     16 |
 | ssccn3-dom3                  |     2 |     32 |   Root    |     2 |     32 |
 +------------------------------+-------+--------+-----------+-------+--------+
 | unallocated or parked        |    45 |   4048 |    --     |    -- |   --   |
 +------------------------------+-------+--------+-----------+-------+--------+

 [Note] Following domains will be skipped in this session.
```

```
Root Domains
------------
ssccn3-dom2
ssccn3-dom3


CPU allocation preference:

       1. Socket level
       2. Core level

In case of Socket level granularity, proportional memory capacity is
 automatically selected for you.

Choose Socket or Core level [S or C] s


Step 1 of 1: Socket Count

primary      : specify socket count [min: 1, max: 3. default: 1] : 1
               you chose [1] socket for primary domain

ssccn3-dom1  : specify socket count [min: 1, max: 3. default: 1] : 3
               you chose [3] sockets for ssccn3-dom1 domain


Configuration In Progress After Socket Count Selection:

+----------------------------------+----------+----------+-----------+
| DOMAIN                           |  SOCKETS |   MEM GB |   TYPE    |
+----------------------------------+----------+----------+-----------+
| primary                          |     1    |     1024 | Dedicated |
| ssccn3-dom1                      |     3    |     3072 | Dedicated |
| *ssccn3-dom2                     |   0.083  |       16 |   Root    |
| *ssccn3-dom3                     |   0.167  |       32 |   Root    |
+----------------------------------+----------+----------+-----------+
| unallocated or parked            |   3.750  |     4048 |    --     |
+----------------------------------+----------+----------+-----------+


Following domains will be stopped and restarted:

       ssccn3-dom1

This configuration requires rebooting the control domain.
Do you want to proceed? Y/N : Y

+- IMPORTANT NOTE: -+
```

```
 |  After the reboot, osc-setcorememem attempts to complete CPU, memory
re-configuration. |
 |  Please check syslog and the state of all domains before using the system.        |
 |  eg.,  dmesg | grep osc-setcorememem ; ldm list | grep -v active ; date |
 +- -+

 All activity is being recorded in log file:

/opt/oracle.supercluster/osc-setcorememem/log/osc-setcorememem_activity_08-28-2015_15:
31:27.log

 Please wait while osc-setcorememem is setting up the new CPU, memory configuration.
 It may take a while. Be patient and do not interrupt.

 0%    10    20    30    40    50    60    70    80    90   100%
 |-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|

 [Info] Domain ssccn3-dom1 is taking too long to stop. Waiting ..
 [Info] Domain ssccn3-dom1 is taking too long to stop. Still waiting for the domain
to stop ..
 *=====*=====*=====*=====*=====*=====*=====*=====*=====*=====*

Broadcast Message from root (pts/1) on etc5mdbadm0301 Fri Aug 28 15:36:45...
THE SYSTEM etc5mdbadm0301 IS BEING SHUT DOWN NOW ! ! !
Log off now or risk your files being damaged

               Task complete with no errors.

#
```

5. **Check the system log and the status of all logical domains to ensure that they are in active state before proceeding with the regular activity.**

   Example:

```
# dmesg | grep osc-setcorememem
Aug 28 15:43:46 etc5mdbadm0301 root[2074]: [ID 702911 user.alert]
osc-setcorememem: core, memory re-configuration complete. system can be used for
regular work.
```

6. **Verify the new resource allocation.**

   You can verify the resource allocation and check for possible `osc-setcorememem` errors in several ways:

   - "Display the Current Domain Configuration (osc-setcorememem)" on page 47
   - "Display the Current Domain Configuration (ldm)" on page 49
   - "Access osc-setcorememem Log Files" on page 65

7. **Repeat this procedure if you need to change resource allocations on another compute node.**

### Related Information

- "osc-setcorememem Overview" on page 40
- "Supported Domain Configurations" on page 43
- "Plan CPU and Memory Allocations" on page 44
- "Display the Current Domain Configuration (osc-setcorememem)" on page 47
- "Display the Current Domain Configuration (ldm)" on page 49
- "Change CPU/Memory Allocations (Core Granularity)" on page 55
- "Park Cores and Memory" on page 60

# ▼ Change CPU/Memory Allocations (Core Granularity)

Perform this procedure on each compute node to change its CPU and memory resource allocation at the core level.

---

**Note –** To find out if you can perform this procedure, see "Supported Domain Configurations" on page 43.

---

The tool makes these changes:

- Automatically detects Root Domains.
- Calculates the minimum and maximum resources for all domains, and only enables you to select valid quantities.
- Presents viable memory capacities for you to select, based on your core allocations.
- Modifies domain resources according to the choices you make.
- (If needed) Stops nonprimary domains.
- (If needed) Reboots the primary domain with new resources.
- (If needed) Brings up nonprimary domains with new resources.

In this example, 6 cores and 512 GB memory are moved from dedicated domain `ssccn3-dom1` to another dedicated domain, primary.

This table shows the allocation plan (see "Plan CPU and Memory Allocations" on page 44).

| Domain | Domain Type | Cores Before | Cores After | Memory Before (GB) | Memory After (GB) |
|---|---|---|---|---|---|
| primary | Dedicated | 12 | 18 | 1024 | 1536 |
| ssccn3-dom1 | Dedicated | 36 | 30 | 3072 | 2560 |
| ssccn3-dom2 | Root | N/A | N/A | N/A | N/A |
| ssccn3-dom3 | Root | N/A | N/A | N/A | N/A |
| Unallocated | | 45 | 45 | 4048 | 4048 |
| Total resources | | 93 | 93 | 8144 | 8144 |

1. **Log in as superuser on the compute node's control domain.**
2. **Ensure that all applications are shut down and that there is no production activity running.**
3. **Activate any inactive domains using the `ldm bind` command.**

   The tool does not continue if any inactive domains are present.
4. **Run `osc-setcoremem` to reconfigure the resources.**

   Respond when prompted. Press Enter to select the default value.

```
# /opt/oracle.supercluster/bin/osc-setcoremem

                          osc-setcoremem
                  v2.0  built on Aug 27 2015 23:09:35


 Current Configuration: SuperCluster Fully-Populated M6-32 Base

 +---------------------------------+-------+--------+----------+--- MINIMUM
----+
 | DOMAIN                          | CORES | MEM GB |   TYPE    | CORES | MEM GB |

+---------------------------------+-------+--------+----------+-------+--------
+
 | primary                         |    12 |   1024 | Dedicated |     2 |     32 |
 | ssccn3-dom1                     |    36 |   3072 | Dedicated |     2 |     32 |
 | ssccn3-dom2                     |     1 |     16 |   Root    |     1 |     16 |
 | ssccn3-dom3                     |     2 |     32 |   Root    |     2 |     32 |

+---------------------------------+-------+--------+----------+-------+--------
+
```

```
 | unallocated or parked            |    45 |   4048 |    --     |    -- |   --   |
 +----------------------------------+-------+--------+-----------+-------+--------+

 [Note] Following domains will be skipped in this session.

 Root Domains
 ------------
 ssccn3-dom2
 ssccn3-dom3


 CPU allocation preference:

        1. Socket level
        2. Core level

 In case of Socket level granularity, proportional memory capacity is
  automatically selected for you.

 Choose Socket or Core level [S or C] C


 Step 1 of 2: Core Count

 primary      : specify number of cores [min: 2, max: 46. default: 12] : 18
                you chose [18] cores for primary domain

 ssccn3-dom1  : specify number of cores [min: 2, max: 30. default: 2] : 30
                you chose [30] cores for ssccn3-dom1 domain


 Configuration In Progress After Core Count Selection:

 +----------------------------------+-------+--------+-----------+--- MINIMUM
----+
 | DOMAIN                           | CORES | MEM GB |   TYPE    | CORES | MEM GB |
 +----------------------------------+-------+--------+-----------+-------+--------+
 | primary                          |    18 |   1024 | Dedicated |     2 |     96 |
 | ssccn3-dom1                      |    30 |   3072 | Dedicated |     2 |    128 |
 | *ssccn3-dom2                     |     1 |     16 |   Root    |     1 |     16 |
 | *ssccn3-dom3                     |     2 |     32 |   Root    |     2 |     32 |

+----------------------------------+-------+--------+-----------+-------+--------
+
 | unallocated or parked            |    45 |   4048 |    --     |    -- |   --   |
 +----------------------------------+-------+--------+-----------+-------+--------+


 Step 2 of 2: Memory Capacity
```

```
        (must be 16 GB aligned)

 primary: specify memory capacity in GB [min: 96, max: 2016. default: 2016]: 1536
               you chose [1536 GB] memory for primary domain

 ssccn3-dom1: specify memory capacity in GB [min: 128, max: 2560. default: 2560]:
2560
               you chose [2560 GB] memory for ssccn3-dom1 domain


 Configuration In progress After Memory Capacity Selection:

 +---------------------------------+-------+--------+-----------+--- MINIMUM ----+
 | DOMAIN                          | CORES | MEM GB |   TYPE    | CORES | MEM GB |
 +---------------------------------+-------+--------+-----------+-------+--------+
 | primary                         |    18 |   1536 | Dedicated |     2 |     96 |
 | ssccn3-dom1                     |    30 |   2560 | Dedicated |     2 |    128 |
 | *ssccn3-dom2                    |     1 |     16 |   Root    |     1 |     16 |
 | *ssccn3-dom3                    |     2 |     32 |   Root    |     2 |     32 |
 +---------------------------------+-------+--------+-----------+-------+--------+
 | unallocated or parked           |    45 |   4048 |    --     |    -- |   --   |
 +---------------------------------+-------+--------+-----------+-------+--------+


 Following domains will be stopped and restarted:

       ssccn3-dom1
This configuration requires rebooting the control domain.
 Do you want to proceed? Y/N : y

 IMPORTANT NOTE:
 +-                                                                              -+
 |  After the reboot, osc-setcoremem attempts to complete CPU, memory
re-configuration. |
 |  Please check syslog and the state of all domains before using the system.
|
 |  eg.,  dmesg | grep osc-setcoremem ; ldm list | grep -v active ; date            |
 +-                                                                              -+

 All activity is being recorded in log file:

/opt/oracle.supercluster/osc-setcoremem/log/osc-setcoremem_activity_08-28-2015_15
:59:31.log

 Please wait while osc-setcoremem is setting up the new CPU, memory configuration.
 It may take a while. Be patient and do not interrupt.

 0%    10    20    30    40    50    60    70    80    90   100%
 |-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
```

```
 *=====*=====*=====*=====*=====*=====*=====*=====*=====*=====*

Broadcast Message from root (pts/1) on etc5mdbadm0301 Fri Aug 28 16:03:13...
THE SYSTEM etc5mdbadm0301 IS BEING SHUT DOWN NOW ! ! !
Log off now or risk your files being damaged

                Task complete with no errors.

#
```

5. **Verify the new resource allocation.**

   You can verify the resource allocation and check for possible `osc-setcoremem` errors in several ways:

   - "Display the Current Domain Configuration (osc-setcoremem)" on page 47
   - "Display the Current Domain Configuration (osc-setcoremem)" on page 47
   - "Access osc-setcoremem Log Files" on page 65

   Example:

```
# dmesg | grep osc-setcoremem
Aug 28 16:08:56 etc5mdbadm0301 root[1913]: [ID 702911 user.alert]
osc-setcoremem: core, memory re-configuration complete. system can be used for
regular work.


# ldm list
NAME              STATE      FLAGS   CONS    VCPU  MEMORY   UTIL  NORM  UPTIME
primary           active     -n-cv-  UART    144   1572096M 0.1%  0.1%  5m
ssccn3-dom1       active     -n----  5001    240   2620928M 1.3%  1.3%  2m
ssccn3-dom2       active     -n----  5002    8     16G      0.1%  0.1%  3d 16m
ssccn3-dom3       active     -n--v-  5003    16    32G      0.1%  0.1%  3d 16m
```

6. **Repeat this procedure if you need to change resource allocations on another compute node.**

### Related Information

- "osc-setcoremem Overview" on page 40
- "Supported Domain Configurations" on page 43
- "Plan CPU and Memory Allocations" on page 44
- "Display the Current Domain Configuration (osc-setcoremem)" on page 47
- "Display the Current Domain Configuration (ldm)" on page 49
- "Change CPU/Memory Allocations (Socket Granularity)" on page 51
- "Park Cores and Memory" on page 60

# ▼ Park Cores and Memory

Perform this procedure on each compute node to move CPU and memory resources from dedicated domains into logical CPU and memory repositories, making the resources available for I/O Domains.

If you are parking cores and memory, plan carefully. Once you park resources and create I/O Domains you cannot move resources back to dedicated domains.

---

**Note –** To find out if you can perform this procedure, see "Supported Domain Configurations" on page 43.

---

In this example, 12 cores and 1 TB memory are parked from the primary domain, and 18 cores and 1536 GB memory are parked from the `ssccn3-dom1` domain.

This table shows the allocation plan (see "Plan CPU and Memory Allocations" on page 44).

| Domain | Domain Type | Cores Before | Cores After | Memory Before (GB) | Memory After (GB) |
|---|---|---|---|---|---|
| `primary` | Dedicated | 18 | 6 | 1536 | 512 |
| `ssccn3-dom1` | Dedicated | 30 | 12 | 2560 | 1024 |
| `ssccn3-dom2` | Root | n/a | n/a | n/a | n/a |
| `ssccn3-dom3` | Root | n/a | n/a | n/a | n/a |
| Unallocated Resources | | 45 | | 4048 | |
| Total Resources | | 93 | 93 | 8144 | 8144 |

1. **Log in as superuser on the compute node's control domain.**
2. **Ensure that all applications are shut down and that there is no production activity running.**
3. **Activate any inactive domains using the `ldm bind` command.**

   The tool does not continue if any inactive domains are present.

4. **Run** `osc-setcoremem` **to change resource allocations.**

   In this example, some resources are left unallocated which parks them.

   Respond when prompted. Press Enter to select the default value.

```
# /opt/oracle.supercluster/bin/osc-setcorememm

                              osc-setcoremem
                        v2.0  built on Aug 27 2015 23:09:35


 Current Configuration: SuperCluster Fully-Populated M6-32 Base

 +---------------------------------+-------+--------+-----------+--- MINIMUM ----+
 | DOMAIN                          | CORES | MEM GB |   TYPE    | CORES | MEM GB |

+---------------------------------+-------+--------+-----------+-------+--------+
 | primary                         |    18 |   1536 | Dedicated |     2 |     32 |
 | ssccn3-dom1                     |    30 |   2560 | Dedicated |     2 |     32 |
 | ssccn3-dom2                     |     1 |     16 |   Root    |     1 |     16 |
 | ssccn3-dom3                     |     2 |     32 |   Root    |     2 |     32 |

+---------------------------------+-------+--------+-----------+-------+--------+
 | unallocated or parked           |    45 |   4048 |    --     |    -- |   --   |

+---------------------------------+-------+--------+-----------+-------+--------+

 [Note] Following domains will be skipped in this session.

 Root Domains
 ------------
 ssccn3-dom2
 ssccn3-dom3


 CPU allocation preference:

        1. Socket level
        2. Core level

 In case of Socket level granularity, proportional memory capacity is
  automatically selected for you.

 Choose Socket or Core level [S or C] c


 Step 1 of 2: Core Count
```

```
 primary      : specify number of cores [min: 2, max: 46. default: 18] : 6
                you chose [6] cores for primary domain

 ssccn3-dom1  : specify number of cores [min: 2, max: 42. default: 30] : 12
                you chose [12] cores for ssccn3-dom1 domain


 Configuration In Progress After Core Count Selection:

 +--------------------------------+-------+--------+-----------+--- MINIMUM ----+
 | DOMAIN                         | CORES | MEM GB |   TYPE    | CORES | MEM GB |

+--------------------------------+-------+--------+-----------+-------+--------+
 | primary                        |     6 |   1536 | Dedicated |     2 |     32 |
 | ssccn3-dom1                    |    12 |   2560 | Dedicated |     2 |     64 |
 | *ssccn3-dom2                   |     1 |     16 |   Root    |     1 |     16 |
 | *ssccn3-dom3                   |     2 |     32 |   Root    |     2 |     32 |

+--------------------------------+-------+--------+-----------+-------+--------+
 | unallocated or parked          |    75 |   4048 |    --     |    -- |   --   |

+--------------------------------+-------+--------+-----------+-------+--------+


 Step 2 of 2: Memory Capacity
        (must be 16 GB aligned)

 primary: specify memory capacity in GB [min: 32, max: 2048. default: 2048] : 512
                you chose [512 GB] memory for primary domain

 ssccn3-dom1:specify memory capacity in GB [min: 64, max: 2048. default: 2048] :
1024
                you chose [1024 GB] memory for ssccn3-dom1 domain


 Configuration In progress After Memory Capacity Selection:

 +--------------------------------+-------+--------+-----------+--- MINIMUM ----+
 | DOMAIN                         | CORES | MEM GB |   TYPE    | CORES | MEM GB |

+--------------------------------+-------+--------+-----------+-------+--------+
 | primary                        |     6 |    512 | Dedicated |     2 |     32 |
 | ssccn3-dom1                    |    12 |   1024 | Dedicated |     2 |     64 |
 | *ssccn3-dom2                   |     1 |     16 |   Root    |     1 |     16 |
 | *ssccn3-dom3                   |     2 |     32 |   Root    |     2 |     32 |

+--------------------------------+-------+--------+-----------+-------+--------+
 | unallocated or parked          |    75 |   6608 |    --     |    -- |   --   |
 +--------------------------------+-------+--------+-----------+-------+--------+
```

```
 Following domains will be stopped and restarted:

        ssccn3-dom1

 This configuration requires rebooting the control domain.
 Do you want to proceed? Y/N : y

 IMPORTANT NOTE:
 +-                                                                              -+
 |  After the reboot, osc-setcoremem attempts to complete CPU, memory
re-configuration. |
 |  Please check syslog and the state of all domains before using the system.
|
 |  eg.,  dmesg | grep osc-setcoremem ; ldm list | grep -v active ; date            |
 +-                                                                              -+

 All activity is being recorded in log file:

/opt/oracle.supercluster/osc-setcoremem/log/osc-setcoremem_activity_08-28-2015_16
:18:57.log
Please wait while osc-setcoremem is setting up the new CPU, memory configuration.
 It may take a while. Be patient and do not interrupt.

 0%    10    20    30    40    50    60    70    80    90   100%
 |-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
 *=====*=====*=====*=====*=====*=====*=====*=====*=====*=====*

Broadcast Message from root (pts/1) on etc5mdbadm0301 Fri Aug 28 16:22:07...
THE SYSTEM etc5mdbadm0301 IS BEING SHUT DOWN NOW ! ! !
Log off now or risk your files being damaged

                Task complete with no errors.

#
```

5. **If the tool indicated that a reboot was needed, after the system reboots, log in as `root` on the compute node's control domain.**

6. **Verify the new resource allocation.**

   You can verify the resource allocation and check for possible `osc-setcoremem` errors in several ways:

   - "Display the Current Domain Configuration (osc-setcoremem)" on page 47
   - "Display the Current Domain Configuration (osc-setcoremem)" on page 47
   - "Access osc-setcoremem Log Files" on page 65

7. **Check the log file to ensure that all reconfiguration steps were successful.**

```
# cd /opt/oracle.supercluster/osc-setcoremem/log
# ls (identify the name of the log file)
# tail -17 osc-setcoremem_activity_08-28-2015_16\:18\:57.log

 ::Post-reboot activity::

 Please wait while osc-setcoremem is setting up the new CPU, memory
configuration.
 It may take a while. Be patient and do not interrupt.


 Executing ldm commands ..

 0%    10    20    30    40    50    60    70    80    90   100%
 |-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
 *=====*=====*=====*=====*=====*=====*=====*=====*=====*=====*

                Task complete with no errors.
               This concludes socket/core, memory reconfiguration.
                You can continue using the system.
```

8. **Verify the new resource allocation.**

   You can verify the resource allocation and check for possible `osc-setcoremem` errors in several ways:

   - "Display the Current Domain Configuration (osc-setcoremem)" on page 47
   - "Display the Current Domain Configuration (osc-setcoremem)" on page 47
   - "Access osc-setcoremem Log Files" on page 65

   Example:

```
# dmesg | grep osc-setcoremem
Aug 28 16:27:50 etc5mdbadm0301 root[1926]: [ID 702911 user.alert]
osc-setcoremem: core, memory re-configuration complete. system can be used for
regular work.

# ldm list
NAME              STATE      FLAGS   CONS    VCPU  MEMORY   UTIL  NORM  UPTIME
primary           active     -n-cv-  UART    48    523008M  0.4%  0.4%  6m
ssccn3-dom1       active     -n----  5001    96    1T       0.2%  0.2%  3m
ssccn3-dom2       active     -n----  5002    8     16G      0.1%  0.1%  3d 36m
ssccn3-dom3       active     -n--v-  5003    16    32G      0.1%  0.1%  3d 36m
```

9. **Verify the parked cores.**

   See "Display the Current Domain Configuration (ldm)" on page 49:

```
# ldm list-devices -p core | grep cid | wc -l
      75
```

10. **Verify the parked memory.**

    See "Display the Current Domain Configuration (ldm)" on page 49:

```
# ldm list-devices memory
MEMORY
    PA                   SIZE
    0x3c00000000         768G
    0x84000000000        768G
    0x100000000000       1008G
    0x180000000000       1T
    0x208000000000       512G
    0x288000000000       512G
    0x300000000000       1008G
    0x380000000000       1008G
```

11. **Repeat this procedure if you need to change resource allocations on the other compute node.**

### Related Information

- "Supported Domain Configurations" on page 43
- "Plan CPU and Memory Allocations" on page 44
- "Display the Current Domain Configuration (osc-setcoremem)" on page 47
- "Display the Current Domain Configuration (ldm)" on page 49
- "Access osc-setcoremem Log Files" on page 65

## ▼ Access `osc-setcoremem` Log Files

The `osc-setcoremem` command creates a timestamped log file for each session.

1. **Log in as superuser on the compute node's control domain.**

2. **Change directories to the log file directory and list the contents to get the name of the log file.**

```
# cd /opt/oracle.supercluster/osc-setcoremem/log
# ls
```

3. **Use a text reader of your choice to view the contents of a log file.**

```
# more log_file_name
```

Example:

```
# cat osc-setcoremem_activity_08-28-2015_15\:59\:31.log

# ./osc-setcoremem

                            osc-setcoremem
                    v2.0  built on Aug 27 2015 23:09:35


 Current Configuration: SuperCluster Fully-Populated M6-32 Base

 +----------------------------------+-------+--------+-----------+--- MINIMUM
----+
 | DOMAIN                          | CORES | MEM GB |   TYPE    | CORES | MEM GB |

+----------------------------------+-------+--------+-----------+-------+-----
---+
 | primary                          |    12 |   1024 | Dedicated |     2 |     32 |
 | ssccn3-dom1                      |    36 |   3072 | Dedicated |     2 |     32 |
 | ssccn3-dom2                      |     1 |     16 |   Root    |     1 |     16 |
 | ssccn3-dom3                      |     2 |     32 |   Root    |     2 |     32 |

+----------------------------------+-------+--------+-----------+-------+-----
---+
 | unallocated or parked            |    45 |   4048 |    --     |    -- |   --   |

+----------------------------------+-------+--------+-----------+-------+-----
---+

 [Note] Following domains will be skipped in this session.

 Root Domains
 ------------
 ssccn3-dom2
 ssccn3-dom3
```

```
 CPU allocation preference:

        1. Socket level
        2. Core level

 In case of Socket level granularity, proportional memory capacity is
  automatically selected for you.

 Choose Socket or Core level [S or C]
 user input: 'C'


 Step 1 of 2: Core Count

 primary      : specify number of cores [min: 2, max: 46. default: 12] :
 user input (desired cores): '18'              you chose [18] cores for primary
domain

 ssccn3-dom1  : specify number of cores [min: 2, max: 30. default: 2] :
 user input (desired cores): '30'           you chose [30] cores for ssccn3-dom1
domain


 Configuration In Progress After Core Count Selection:

 +----------------------------------+-------+--------+-----------+--- MINIMUM
----+
 | DOMAIN                          | CORES | MEM GB |   TYPE    | CORES | MEM GB |

+----------------------------------+-------+--------+-----------+-------+-----
---+
 | primary                         |    18 |   1024 | Dedicated |     2 |     96 |
 | ssccn3-dom1                     |    30 |   3072 | Dedicated |     2 |    128 |
 | *ssccn3-dom2                    |     1 |     16 |   Root    |     1 |     16 |
 | *ssccn3-dom3                    |     2 |     32 |   Root    |     2 |     32 |

+----------------------------------+-------+--------+-----------+-------+-----
---+
 | unallocated or parked           |    45 |   4048 |    --     |    -- |   --   |

+----------------------------------+-------+--------+-----------+-------+-----
---+


 Step 2 of 2: Memory Capacity
        (must be 16 GB aligned)
```

```
 primary     : specify memory capacity in GB [min: 96, max: 2016. default: 2016] :
 user input (desired memory): '1536' GB         you chose [1536 GB] memory for
primary domain

 ssccn3-dom1  : specify memory capacity in GB [min: 128, max: 2560. default:
2560] :
 user input (desired memory): '' GB             you chose [2560 GB] memory for
ssccn3-dom1 domain


 Configuration In progress After Memory Capacity Selection:

 +---------------------------------+-------+--------+-----------+--- MINIMUM
----+
 | DOMAIN                          | CORES | MEM GB |   TYPE    | CORES | MEM GB |

+---------------------------------+-------+--------+-----------+-------+-----
---+
 | primary                         |    18 |   1536 | Dedicated |     2 |     96 |
 | ssccn3-dom1                     |    30 |   2560 | Dedicated |     2 |    128 |
 | *ssccn3-dom2                    |     1 |     16 |   Root    |     1 |     16 |
 | *ssccn3-dom3                    |     2 |     32 |   Root    |     2 |     32 |

+---------------------------------+-------+--------+-----------+-------+-----
---+
 | unallocated or parked           |    45 |   4048 |    --     |    -- |   --   |

+---------------------------------+-------+--------+-----------+-------+-----
---+


 Following domains will be stopped and restarted:

        ssccn3-dom1

 This configuration requires rebooting the control domain.
 Do you want to proceed? Y/N :
 user input: 'y'

 IMPORTANT NOTE:
 +-                                                                          -+
 |  After the reboot, osc-setcoremem attempts to complete CPU, memory
re-configuration. |
 | Please check syslog and the state of all domains before using the system.
|
 |  eg.,  dmesg | grep osc-setcoremem ; ldm list | grep -v active ; date
|
 +-                                                                          -+
```

```
 Please wait while osc-setcorememem is setting up the new CPU, memory
configuration.
 It may take a while. Be patient and do not interrupt.


 Executing ldm commands ..

 0%    10    20    30    40    50    60    70    80    90   100%
 |-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
 *=====*=====*=====*=====*=====*=====*=====*=====*=====*=====*

               Task complete with no errors.


 ::Post-reboot activity::

 Please wait while osc-setcorememem is setting up the new CPU, memory
configuration.
 It may take a while. Be patient and do not interrupt.


 Executing ldm commands ..

 0%    10    20    30    40    50    60    70    80    90   100%
 |-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
 *=====*=====*=====*=====*=====*=====*=====*=====*=====*=====*

               Task complete with no errors.
               This concludes socket/core, memory reconfiguration.
               You can continue using the system.
```

### Related Information


- "osc-setcorememem Overview" on page 40
- "Supported Domain Configurations" on page 43
- "Plan CPU and Memory Allocations" on page 44
- "Display the Current Domain Configuration (osc-setcorememem)" on page 47
- "Display the Current Domain Configuration (ldm)" on page 49
- "Change CPU/Memory Allocations (Socket Granularity)" on page 51
- "Change CPU/Memory Allocations (Core Granularity)" on page 55
- "Park Cores and Memory" on page 60

# ▼ View the SP Configuration

When you reallocate resources using the `osc-setcoremem` command, `osc-setcoremem` saves the new configuration to the service processor (SP) in this format:

`CM_`*dom1_dom2_dom3_..._TimeStamp*

where:

- `CM_` – indicates a core/memory configuration that was created sometime after the initial installation.
- *domx* is expressed with this nomenclature:
  - *x*`C` or *x*`S` – CPU resources in number (*x*) of cores (`C`) or sockets (`S`)
  - *x*`G` or *x*`T` – Memory resources in number (*x*) of gigabytes (`G`) or number of terabytes (`T`)
- *TimeStamp* – in the format `MMDDYYYYHHMM`

This file name example . . .

`CM_2S1T_1S512G_3S1536G_082020151354`

. . . represents a configuration created on August 20, 2015 at 13:54 and has three domains with these resources:

- 2-sockets, 1-TB memory
- 1-socket, 512 GB memory
- 3-sockets, 1536 GB memory

To see more details about the resource allocations, you can use the SP configuration timestamp to locate and view the corresponding `osc-setcoremem` log file.

1. **Log in as superuser on the compute node's control domain.**
2. **Display the SP configuration.**

   Examples:

   - Output indicating no custom CPU/memory configurations:

     The file called `V_B4_4_1_20150804141204` is the initial resource configuration file that was created when the system was installed.

```
# ldm list-config
factory-default
V_B4_4_1_20150825155356 [next poweron]
```

- Output indicating three additional CPU/memory configurations:

```
# ldm list-config
factory-default
V_B4_4_1_20150825155356
CM_3S3T_1S1T_082820151531
CM_30C2560G_18C1536G_082820151559
CM_1S1T_6C512G_082820151618 [current]
```

3. **View the corresponding log file.**

```
# more
/opt/oracle.supercluster/osc-setcoremem/log/osc-setcoremem_activity_08-28-2015
_16\:18*.log
```

### Related Information


- "osc-setcoremem Overview" on page 40
- "Supported Domain Configurations" on page 43
- "Plan CPU and Memory Allocations" on page 44
- "Display the Current Domain Configuration (osc-setcoremem)" on page 47
- "Display the Current Domain Configuration (ldm)" on page 49
- "Change CPU/Memory Allocations (Socket Granularity)" on page 51
- "Change CPU/Memory Allocations (Core Granularity)" on page 55
- "Park Cores and Memory" on page 60

# ▼ Revert to a Previous CPU/Memory Configuration

Use this procedure to revert a compute node to a previous CPU/Memory configuration. You must perform this procedure on each member in a cluster. The tool does not automatically propagate changes to every cluster member.

---

**Note –** To find out if you can perform this procedure, see "Supported Domain Configurations" on page 43.

---

1. **Log in as superuser on the compute node's control domain.**
2. **List previous configurations.**

---

**Note –** You can also view previous configurations in the log files. See "Access osc-setcoremem Log Files" on page 65.

---

```
# ldm list-config
factory-default
V_B4_4_1_20150825155356
CM_3S3T_1S1T_082820151531
CM_30C2560G_18C1536G_082820151559
CM_1S1T_6C512G_082820151618 [current]
```

For details about SP configuration files see "View the SP Configuration" on page 70.

3. **Revert to a previous configuration.**

```
# ldm set-config CM_30C2560G_18C1536G_082820151559
```

4. **Halt all domains, then halt the primary domain.**
5. **Restart the system from the service processor.**

```
# #.

-> cd /SP
-> stop /SYS
Are you sure you want to stop /SYS (y/n) ? y
```

```
Stopping /SYS

-> start /SYS
Are you sure you want to start /SYS (y/n) ? y
Starting /SYS
```

6. **Boot all domains and zones.**

### Related Information

- "osc-setcoremem Overview" on page 40
- "Supported Domain Configurations" on page 43
- "Plan CPU and Memory Allocations" on page 44
- "Display the Current Domain Configuration (osc-setcoremem)" on page 47
- "Display the Current Domain Configuration (ldm)" on page 49
- "Change CPU/Memory Allocations (Socket Granularity)" on page 51
- "Change CPU/Memory Allocations (Core Granularity)" on page 55
- "Park Cores and Memory" on page 60

# ▼ Remove a CPU/Memory Configuration

The compute node's service processor has a limited amount of memory. If you are unable to create a new configuration because the service processor ran out of memory, delete unused configurations using this procedure.

1. **List all current configurations.**

```
# ldm list-config
factory-default
V_B4_4_1_20150825155356
CM_3S3T_1S1T_082820151531
CM_30C2560G_18C1536G_082820151559
CM_1S1T_6C512G_082820151618 [current]
```

2. **Determine which configurations are safe to remove.**

   It is safe to remove any configuration that contains the string `CM_` or `_ML`, as long as it is not marked `[current]` or `[next poweron]`.

3. **Remove a configuration.**

   Example:

```
# ldm remove-spconfig CM_3S3T_1S1T_082820151531
```

### Related Information

- "osc-setcoremem Overview" on page 40
- "Supported Domain Configurations" on page 43
- "Plan CPU and Memory Allocations" on page 44
- "Display the Current Domain Configuration (osc-setcoremem)" on page 47
- "Display the Current Domain Configuration (ldm)" on page 49
- "Change CPU/Memory Allocations (Socket Granularity)" on page 51
- "Change CPU/Memory Allocations (Core Granularity)" on page 55
- "Park Cores and Memory" on page 60

# Obtaining the EM Exadata Plug-in

Starting with Oracle SuperCluster 1.1, you can monitor all Exadata-related software and hardware components in the cluster using the Oracle Enterprise Manager Exadata 12.1.0.3 Plug-in only in the supported configuration described in these topics.



## ▼ Confirm System Requirements

Only SuperCluster systems with software version 1.1 (or later) with Database Domain on Control LDom-only environments are supported. Earlier versions of SuperCluster systems can be made compatible if you update to the October 2012 QMU release.

---

**Note –** With the Oracle SuperCluster software version 2.x, the `compmon` command name changed to `osc-compmon`. Use the new name if SuperCluster is installed with the Oracle SuperCluster v2.x release bundle. See "Identify the Version of SuperCluster Software" on page 1.

---

- **Confirm you have a version of the `compmon pkg` installed on SuperCluster M6-32 using either `pkg info compmon` or `pkg list compmon` commands.**

  You must have the following minimum version of `compmon` installed:

  ```
  pkg://exa-family/system/platform/exadata/compmon@0.5.11,5.11-0
  .1.0.11:20120726T024158Z
  ```

# Known Issues With the EM Exadata Plug-in

- The prerequisite check script `exadataDiscoveryPreCheck.pl` that is bundled in the EM Exadata plug-in 12.1.0.3 does not support the `catalog.xml` file.

  Download the latest `exadataDiscoveryPreCheck.pl` file from MOS as described in the "Discovery Precheck Script" section of the *Oracle Enterprise Manager Exadata Management Getting Started Guide* at:

  `docs.oracle.com/cd/E24628_01/doc.121/e27442/title.htm`

- If multiple database clusters share the same storage server, in one Enterprise Manager management server environment, you can discover and monitor the first DB machine target and all its components. However, for additional DB machine targets sharing the same storage server, the Oracle Storage Server Grid system and the Oracle Database Storage Server System have no storage server members because they are already monitored.
- If the `perfquery` command installed on SuperCluster M6-32 is version 1.5.8 or later, you might encounter a bug (ID 15919339) where most columns in the HCA Port Errors metric in the host targets for the compute nodes are blank. Any errors occurring on the HCA ports, are not reported in Enterprise Manager.

  To check your version, run the following command:

  ```
  perfquery -V
  ```

# Configuring the Exalogic Software

These topics describe how to use Exalogic software on Oracle SuperCluster M6-32.



## Exalogic Software Overview

Oracle EECS includes performance optimizations for SuperCluster M6-32 to improve input/output, thread management, and request handling efficiency.

Additional optimizations include reduced buffer copies, which result in more efficient input/output. Finally, session replication performance and CPU utilization is improved through lazy deserialization, which avoids performing extra work on every session update that is only necessary when a server fails.

WebLogic Server clusters can be configured with cluster-wide optimizations that further improve server-to-server communication. The first optimization enables multiple replication channels, which improve network throughput among WebLogic Server cluster nodes. The second cluster optimization enables IB support for Sockets Direct Protocol, which reduces CPU utilization as network traffic bypasses the TCP stack.

# ▼ Prepare to Configure the Exalogic Software

1. **Configure the environment, including database, storage, and network.**

   Refer to Chapter 3, "Network, Storage, and Database Preconfiguration" of the *Oracle Exalogic Enterprise Deployment Guide*, at:
   http://docs.oracle.com/cd/E18476_01/doc.220/e18479/toc.htm

2. **Configure your Oracle Exalogic Domain.**

   Refer to Chapter 5, "Configuration Oracle Fusion Middleware" in the *Oracle Exalogic Enterprise Deployment Guide*, at:
   http://docs.oracle.com/cd/E18476_01/doc.220/e18479/toc.htm

# ▼ Enable Domain-Level Enhancements

1. **Log in to the Oracle WebLogic Server Administration Console.**
2. **Select Domainname in the left navigation pane.**

   The Settings for Domainname screen is displayed.

3. **Click the General tab.**
4. **In your domain home page, select Enable Exalogic Optimizations, and click Save.**
5. **Activate changes.**
6. **Stop and start your domain.**

   The Enable Exalogic Optimizations setting collectively enables all of the individual features described in this table. The Startup Option indicates how to independently enable and disable each feature.

| Feature | Options | Description |
|---|---|---|
| Scattered Reads | Description | Increased efficiency during I/O in environments with high network throughput |
| | Startup Option | `-Dweblogic.ScatteredReadsEnabled=true/false` |
| | MBean | `KernelMBean.setScatteredReadsEnabled` |
| Gathered Writes | Description | Increased efficiency during I/O in environments with high network throughput |
| | Startup Option | `-Dweblogic.GatheredWritesEnabled=true/false` |
| | MBean | `KernelMBean.setGatheredWritesEnabled` |
| Lazy Deserialization | Description | Increased efficiency with session replication |
| | Startup Option | `-Dweblogic.replication.enableLazyDeserialization=true/false` |
| | MBean | `ClusterMBean.setSessionLazyDeserializationEnabled` |

**Note –** After enabling the optimizations, you might see the following message:

```
java.io.IOException: Broken pipe.
```

You might see the same message when storage failover occurs. In either case, you can ignore the error message.

# ▼ Enable Cluster-Level Session Replication Enhancements

You can enable session replication enhancements for managed servers in a WebLogic cluster to which you will deploy a web application at a later time.

**Note –** If you are using Coherence*web, these session replication enhancements do not apply. Skip these steps if you use the `dizzyworld.ear` application as described in Chapter 8, "Deploying a Sample Web Application to and Oracle WebLogic Cluster" in the *Oracle Fusion Middleware Exalogic Enterprise Deployment Guide* at: http://docs.oracle.com/cd/E18476_01/doc.220/e18479/deploy.htm

To enable session replication enhancements for `Dept1_Cluster1`, complete the following steps:

1. **Ensure that managed servers in the `Dept1_Cluster1` cluster are up and running.**

   Refer to Section 5.16 "Starting Managed Servers on ComputeNode1 and ComputeNode2" of the *Oracle® Fusion Middleware Exalogic Enterprise Deployment Guide* at:
   `http://docs.oracle.com/cd/E18476_01/doc.220/e18479/create_domain.htm#BABEGAFB`

2. **Set replication ports for a managed server, such as WLS1.**

   a. **Under Domain Structure, click Environment and Servers.**

      The Summary of Servers page is displayed.

   b. **Click WLS1 on the list of servers.**

      The Settings for WLS1 is displayed.

   c. **Click the Cluster tab.**

   d. **In the Replication Ports field, enter a range of ports for configuring multiple replication channels.**

      For example, replication channels for managed servers in `Dept_1_Cluster1` can listen on ports starting from 7005 to 7015. To specify this range of ports, enter 7005-7015.

3. **Create a custom network channel for each managed server in the cluster (for example, WLS1.**

   a. **Log in to the Oracle WebLogic Server Administration Console.**

   b. **If you have not already done so, click Lock & Edit in the Change Center.**

   c. **In the left pane of the Console, expand Environment and select Servers.**

      The Summary of Servers page is displayed.

   d. **In the Servers table, click WLS1 Managed Server instance.**

   e. **Select Protocols and then Channels.**

   f. **Click New.**

   g. **Enter ReplicationChannel as the name of the new network channel and select t3 as the protocol, then click Next.**

   h. **Enter the following information:**

      - **Listen address:** `10.0.0.1`

**Note –** This is the floating IP assigned to WLS1.

- **Listen port:** `7005`

i. **Click Next, and in the Network Channel Properties page, select Enabled and Outbound Enabled.**

j. **Click Finish.**

k. **Under the Network Channels table, select ReplicationChannel, which is the network channel you created for the WLS1 Managed Server.**

l. **Expand Advanced, and select Enable SDP Protocol.**

m. **Click Save.**

n. **To activate these changes, click Activate Changes in the Change Center of the Administration Console.**

o. **Repeat the preceding steps to create a network channel each for the remaining managed servers in the `Dept1_Cluster1` cluster. Enter the required properties, as described in this table.**

| Managed Servers in `Dept1_Cluster1` | Name | Protocol | Listen Address | Listen Port | Additional Channel Ports |
|---|---|---|---|---|---|
| WLS2 | ReplicationChannel | t3 | `10.0.0.2` | `7005` | `7006 to 7014` |
| WLS3 | ReplicationChannel | t3 | `10.0.0.3` | `7005` | `7006 to 7014` |
| WLS4 | ReplicationChannel | t3 | `10.0.0.4` | `7005` | `7006 to 7014` |
| WLS5 | ReplicationChannel | t3 | `10.0.0.5` | `7005` | `7006 to 7014` |
| WLS6 | ReplicationChannel | t3 | `10.0.0.6` | `7005` | `7006 to 7014` |
| WLS7 | ReplicationChannel | t3 | `10.0.0.7` | `7005` | `7006 to 7014` |
| WLS8 | ReplicationChannel | t3 | `10.0.0.8` | `7005` | `7006 to 7014` |

4. **After creating the network channel for each of the managed servers in your cluster, click Environment→ Clusters.**

   The Summary of Clusters page is displayed.

5. **Click `Dept1_Cluster1`. This is the example cluster to which you will deploy a web application at a later time.**

   The Settings for Dept1_Cluster1 page is displayed.

6. **Click the Replication tab.**

7. **In the Replication Channel field, ensure that ReplicationChannel is set as the name of the channel to be used for replication traffic.**
8. **In the Advanced section, select the Enable One Way RMI for Replication option, and click Save.**
9. **Activate changes, and restart the managed servers.**
10. **Manually add the system property** `-Djava.net.preferIPv4Stack=true` **to the** `startWebLogic.sh` **script, which is located in the** `bin` **directory of base_domain, using a text editor as follows:**
    a. **Locate the following line in the** `startWebLogic.sh` **script:**

    ```
    . ${DOMAIN_HOME}/bin/setDomainEnv.sh $*
    ```

    b. **Add the following property immediately after the preceding entry:**

    ```
    JAVA_OPTIONS="${JAVA_OPTIONS} -Djava.net.preferIPv4Stack=
    true"
    ```

    c. **Save and close the file.**
11. **Restart all managed servers.**
    a. **In the administration console, click Environment → Servers.**

    The Summary of Servers page is displayed.

    b. **Select a managed server, such as WLS1, by clicking WLS1.**

    The Settings for WLS1 page is displayed.

    c. **Click the Control tab. Select WLS1 in the Server Status table. Click Start.**
    d. **Repeat these steps for each of the managed servers in the WebLogic cluster.**
12. **Verify that multiple listening ports were opened in one of these ways:**
    - Type the `netstat -na` command.
    - Check the managed server logs.

# Configuring Grid Link Data Source for `Dept1_Cluster1`

You must create a grid link data source for JDBC connectivity between Oracle WebLogic Server and a service targeted to a RAC cluster. The grid link data source uses the ONS to adaptively respond to state changes in an Oracle RAC instance.

These topics describe the grid link data source and how to create it:



# Fast Connection Failover

A grid link data source uses fast connection failover to:

- Provide rapid failure detection.
- Abort and remove invalid connections from the connection pool.
- Perform graceful shutdown for planned and unplanned Oracle RAC node outages. The data source allows in-progress transactions to complete before closing connections. New requests are load balanced to an active Oracle RAC node.
- Adapt to changes in topology, such as adding a new node.
- Distribute runtime work requests to all active Oracle RAC instances.

Refer to "Fast Connection Failover" in the *Oracle Database JDBC Developer's Guide and Reference* at:
`http://docs.oracle.com/cd/B19306_01/java.102/b14355/fstconfo.htm`.

# Runtime Connection Load Balancing

Runtime connection load balancing allows WebLogic Server to:

- Adjust the distribution of work based on back end node capacities such as CPU, availability, and response time.
- React to changes in RAC topology.
- Manage pooled connections for high performance and scalability.

If FAN is not enabled, grid link data sources use a round-robin load balancing algorithm to allocate connections to RAC nodes.

## XA Affinity

XA Affinity for global transactions ensures that all the data base operations for a global transaction performed on a RAC cluster are directed to the same RAC instance. The first connection request for an XA transaction is load balanced using RCLB and is assigned an Affinity context. All subsequent connection requests are routed to the same RAC instance using the Affinity context of the first connection.

## SCAN Addresses

SCAN addresses can be used to specify the host and port for both the TNS listener and the ONS listener in the WebLogic console. A grid link data source containing SCAN addresses does not need to change if you add or remove RAC nodes. Contact your network administrator for appropriately configured SCAN URLs for your environment. For more information, refer to: http://www.oracle.com/technetwork/database/clustering/overview/scan-129069.pdf.

## Secure Communication With Oracle Wallet

Oracle Wallet allows you to configure secure communication with the ONS listener.

## ▼ Create a Grid Link Data Source on `Dept1_Cluster1`

Create a grid link data source for each of the Oracle database instances during the process of setting up the multidata source, both for these data sources and the global leasing multidata source.

1. **Prepare to create the data source.**
   a. Ensure that this is a non-**xa** data source.
   b. **Target these data sources to the `Dept1_Cluster1` cluster.**
   c. **Set the data sources connection pool initial capacity to 0.**
      i. **In the Oracle WebLogic Server Administration Console, select Services, JDBC, and then Datasources.**
      ii. **In the Datasources screen, click the Datasource Name, then click the Connection Pool tab, and enter 0 in the Initial capacity field.**

**d. Ensure that an ONS daemon is running on your database servers at all times. Start the ONS daemon on a database server by running the `onsctl` command:**

```
start
```

2. **Log in to the Oracle WebLogic Server Administration Console.**
3. **If you have not already done so, in the Change Center of the Administration Console, click Lock & Edit.**
4. **In the Domain Structure tree, expand Services, then select Data Sources.**
5. **On the Summary of Data Sources page, click New and select GridLink Data Source.**

   The Create a New JDBC GridLink Data Source page is displayed.

6. **Enter the following information, then click Next.**
   - Logical name for the datasource in the Name field. For example, gridlink.
   - Name for JNDI. For example, jdbc/gridlink.
7. **In the Transaction Options page, deselect Supports Global Transactions, and click Next.**
8. **Select Enter individual listener information and click Next.**
9. **Enter the following connection properties:**
   - **Service Name**: Enter the name of the RAC service in the Service Name field. For example, enter `myService` in Service Name.

---

**Note –** The Oracle RAC Service name is defined on the database, and it is not a fixed name.

---

   - **Host Name**: Enter the DNS name or IP address of the server that hosts the database. For an Oracle GridLink service-instance connection, this must be the same for each data source in a given multi data source.
   - **Port**: Enter the port on which the database server listens for connections requests.
   - **Database User Name**: Enter the database user name. For example, myDataBase.
   - **Password**: Enter the password. For example, myPassword1.

     Confirm password and click Next.

**Tip –** For more information, refer to the *Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help*.

The console automatically generates the complete JDBC URL. For example:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=
(PROTOCOL=TCP)(HOST=left)(PORT=1234))(ADDRESS=(PROTOCOL=
TCP)(HOST=right)(PORT=1234))(ADDRESS=(PROTOCOL=TCP)(HOST=
center)(PORT=1234)))(CONNECT_DATA=(SERVICE_NAME=myService)))
```

10. **On the Test GridLink Database Connection page, review the connection parameters and click Test All Listeners.**

    Oracle WebLogic attempts to create a connection from the administration server to the database. Results from the connection test are displayed at the top of the page. If the test is unsuccessful, you should correct any configuration errors and retry the test.

    Click Next.

11. **In the ONS Client Configuration page, do the following:**

    a. **Select Fan Enabled to subscribe to and process FAN events.**

    b. **In the ONS host and port fields, enter a comma-separated list of ONS daemon listen addresses and ports for receiving ONS-based FAN events.**

       You can use SCAN addresses to access FAN notifications.

    c. **Click Next.**

12. **On the Test ONS Client Configuration page, review the connection parameters and click Test All ONS Nodes.**

    Click Next.

13. **In the Select Targets page, select Dept1_Cluster1 as the target and All Servers in the cluster.**

14. **Click Finish.**

15. **Click Activate Changes.**

16. **Configure SDP-enabled JDBC drivers for the cluster.**

    For instructions, see "Configuring SDP-Enabled JDBC Drivers for Dept1_Cluster1" on page 87.

# Configuring SDP-Enabled JDBC Drivers for `Dept1_Cluster1`

These topics describe how you must configure SDP-enabled JDBC drivers for the `Dept1_Cluster1` cluster.



## ▼ Configure the Database to Support IB

- **Before enabling SDP support for JDBC, configure the database to support IB.**

  Refer to the Configuring SDP Protocol Support for Infinband Network Communication to the Database Server section in the *Oracle Database Net Services Administrator's Guide*, located at:

  `http://download.oracle.com/docs/cd/B28359_01/network.111/b28316/performance.htm#i1008413`

  Ensure that you set the protocol to SDP.

## ▼ Enable SDP Support for JDBC

1. **Create the grid link data sources for the JDBC connectivity on `ComputeNode1` and `ComputeNode2`.**

   Refer to Section 7.6 "Configuring Grid Link Data Source for Dept1_Cluster1" of the *Oracle® Fusion Middleware Exalogic Enterprise Deployment Guide* at: `http://docs.oracle.com/cd/E18476_01/doc.220/e18479/optimization.htm#BABHEDI`.

   The console automatically generates the complete JDBC URL, as shown in the following example:

   ```
   jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=
   192.x.x.x)(PORT=1522))(CONNECT_DATA=(SERVICE_NAME=myservice)))
   ```

2. **In the JDBC URL, replace TCP protocol with SDP protocol.**

   For example:

   ```
   jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=sdp)(HOST=
   192.x.x.x)(PORT=1522))(CONNECT_DATA=(SERVICE_NAME=myservice)))
   ```

3. **Manually add the system property** `-Djava.net.preferIPv4Stack=true` **to the** `startWebLogic.sh` **script.**

   The script is located in the `bin` directory of base_domain.

   Use a text editor as follows:

   a. **Locate the following line in the** `startWebLogic.sh` **script:**

      ```
      . ${DOMAIN_HOME}/bin/setDomainEnv.sh $*
      ```

   b. **Add the following property immediately after the preceding entry:**

      ```
      JAVA_OPTIONS="${JAVA_OPTIONS} -Djava.net.preferIPv4Stack=
      true -Doracle.net.SDP=true"
      ```

   c. **Save and close the file.**

4. **Restart the managed server.**

   a. **In the administration console, click Environment → Servers. The Summary of Servers page is displayed.**

   b. **Select a managed server, such as WLS1, by clicking WLS1. The Settings for WLS1 page is displayed.**

   c. **Click the Control tab. Select WLS1 in the Server Status table. Click Start.**

## ▼ Monitor SDP Sockets

You can monitor SDP sockets by running the `netstat` command on the Application Domains running Oracle Solaris 11 that contain EECS in SuperCluster M6-32. Run the `netstat` command on these Application Domains running Oracle Solaris 11 and on the Database Domains, to monitor SDP traffic between the Application Domains running Oracle Solaris 11 and the Database Domains.

1. **Log in to the operating system as** `root`**.**

2. **Type the following command.**

   # **netstat -f sdp -s 1**

   This command displays the status of all SDP sockets (established or not), as in the following example output:

```
SDP     sdpActiveOpens       = 66357    sdpCurrEstab         =   748
        sdpPrFails           =     0    sdpRejects           =     0
        sdpOutSegs           =39985638793
        sdpInDataBytes       =9450383834191
        sdpOutDataBytes      =6228930927986

SDP     sdpActiveOpens       =     0    sdpCurrEstab         =     0
        sdpPrFails           =     0    sdpRejects           =     0
        sdpInSegs            = 14547
        sdpOutSegs           = 14525
        sdpInDataBytes       =3537194
        sdpOutDataBytes      =2470907
```

# ▼ Create an SDP Listener on the IB Network

Oracle RAC 11*g* Release 2 supports client connections across multiple networks. This release also provides load balancing and failover of client connections within the network that they are connecting. To add a listener for the EECS connections coming in on the IB network, first add a network resource for the IB network with Virtual IP addresses.

---

**Note –** This example lists two Database Domains. If you have more than two Database Domains in your SuperCluster M6-32, you must repeat Database Domain-specific lines for each Database Domain in the cluster.

---

1. **Edit the** `/etc/hosts` **file on each Database Domain in the cluster to add the virtual IP addresses you will use for the IB network.**

   Ensure that these IP addresses are not used.

   The following is an example:

```
# Added for Listener over IB
192.168.10.21 ssc01db01-ibvip.mycompany.com ssc01db01-ibvip
192.168.10.22 ssc01db02-ibvip.mycompany.com ssc01db02-ibvip
```

2. **On one of the Database Domains, as the root user, create a network resource for the IB network, as in the following example:**

   ```
   # /u01/app/grid/product/11.2.0.2/bin/srvctl add network -k 2
   -S 192.168.10.0/255.255.255.0/bondib0
   ```

3. **Validate that the network was added correctly, by running one of the following commands:**

   ```
   # /u01/app/grid/product/11.2.0.2/bin/crsctl stat res -t | grep
   net

   ora.net1.network

   ora.net2.network -- Output indicating new Network resource
   ```

   or

   ```
   # /u01/app/grid/product/11.2.0.2/bin/srvctl config network -k
   2

   Network exists: 2/192.168.10.0/255.255.255.0/bondib0, type
   static -- Output indicating Network resource on the
   192.168.10.0 subnet
   ```

4. **Add the Virtual IP addresses on the network created in Step 2, for each node in the cluster.**

   ```
   srvctl add vip -n ssc01db01 -A
   ssc01db01-ibvip/255.255.255.0/bondib0 -k 2

   srvctl add vip -n ssc01db02 -A
   ssc01db02-ibvip/255.255.255.0/bondib0 -k 2
   ```

5. **As the "*oracle*" user (who owns the grid infrastructure home), add a listener which listens on the VIP addresses created in Step 3.**

   ```
   srvctl add listener -l LISTENER_IB -k 2 -p TCP:1522,/SDP:1522
   ```

6. **For each database that will accept connections from the middle tier, modify the `listener_networks init` parameter to allow load balancing and failover across multiple networks (Ethernet and IB).**

   You can either enter the full tnsnames syntax in the initialization parameter or create entries in `tnsnames.ora` in the `$ORACLE_HOME/network/admin` directory. The `tnsnames.ora` entries must exist in the `GRID_HOME`.

   The following example first updates `tnsnames.ora`. Complete this step on each Database Domain in the cluster with the correct IP addresses for that Database Domain. `LISTENER_IBREMOTE` should list all other Database Domains that are in the cluster. `DBM_IB` should list all Database Domains in the cluster.

**Note –** The TNSNAMES entry is only read by the database instance on startup. If you modify the entry that is referred to by any init.ora parameter (LISTENER_NETWORKS), you must restart the instance or issue an ALTER SYSTEM SET LISTENER_NETWORKS command for the modifications to take affect by the instance.

```
(DESCRIPTION =
DBM =
(ADDRESS = (PROTOCOL = TCP)(HOST = ssc01-scan)(PORT = 1521))
(CONNECT_DATA =
(SERVER = DEDICATED)

(SERVICE_NAME = dbm)
))

DBM_IB =
(DESCRIPTION =
(LOAD_BALANCE=on)
(ADDRESS = (PROTOCOL = TCP)(HOST = ssc01db01-ibvip)(PORT = 1522))
(ADDRESS = (PROTOCOL = TCP)(HOST = ssc01db02-ibvip)(PORT = 1522))
(CONNECT_DATA =
(SERVER = DEDICATED)
(SERVICE_NAME = dbm)
))

LISTENER_IBREMOTE =
(DESCRIPTION =
(ADDRESS_LIST =
(ADDRESS = (PROTOCOL = TCP)(HOST =
ssc01db02-ibvip.mycompany.com)(PORT = 1522))
))

LISTENER_IBLOCAL =
(DESCRIPTION =
(ADDRESS_LIST =
(ADDRESS = (PROTOCOL = TCP)(HOST =
ssc01db01-ibvip.mycompany.com)(PORT = 1522))
(ADDRESS = (PROTOCOL = SDP)(HOST =
ssc01db01-ibvip.mycompany.com)(PORT = 1522))
))

LISTENER_IPLOCAL =
(DESCRIPTION =
(ADDRESS_LIST =
(ADDRESS = (PROTOCOL = TCP)(HOST = ssc0101-vip.mycompany.com)(PORT
= 1521))
))
```

```
LISTENER_IPREMOTE =
(DESCRIPTION =
(ADDRESS_LIST =
(ADDRESS = (PROTOCOL = TCP)(HOST = ssc01-scan.mycompany.com)(PORT
= 1521))
))
```

7. **Modify the** `listener_networks` **init parameter.**

   Connect to the database instance as `sysdba`.

   ```
   SQLPLUS> alter system set listener_networks='((NAME=network2)
   (LOCAL_LISTENER=LISTENER_IBLOCAL)(REMOTE_LISTENER=
   LISTENER_IBREMOTE))', '((NAME=network1)(LOCAL_LISTENER=
   LISTENER_IPLOCAL)(REMOTE_LISTENER=LISTENER_IPREMOTE))' scope=
   both;
   ```

8. **Stop and start** `LISTENER_IB` **for the modification in** Step 7.

   ```
   srvctl stop listener -l LISTENER_IB
   srvctl start listener -l LISTENER_IB
   ```

# Administering Oracle Solaris 11 Boot Environments

When the Oracle Solaris OS is first installed on SuperCluster M6-32, a boot environment is created. You can use the `beadm(1M)` utility to create and administer additional boot environments on your SuperCluster M6-32.

After your SuperCluster M6-32 is installed, create a backup of the original boot environment. If needed, you can then boot to the backup of the original boot environment.

For more information about Oracle Solaris 11 boot environments, refer to:

`http://docs.oracle.com/cd/E23824_01/html/E21801/toc.html`

These topics describe how to manage the Oracle Solaris 11 boot environments.



## Advantages to Maintaining Multiple Boot Environments

Multiple boot environments reduce risk when updating or changing software because system administrators can create backup boot environments before making any updates to the SuperCluster M6-32. If needed, they have the option of booting a backup boot environment.

The following examples show how having more than one Oracle Solaris boot environment and managing them with the `beadm` utility can be useful.

- You can maintain more than one boot environment on your SuperCluster M6-32 and perform various updates on each of them as needed. For example, you can clone a boot environment by using the `beadm create` command. The clone you create is a bootable copy of the original. Then, you can install, test, and update different software packages on the original boot environment and on its clone.

  Although only one boot environment can be active at a time, you can mount an inactive boot environment by using the `beadm mount` command. Then, you could use the `pkg` command with the alternate root (`-R`) option to install or update specific packages on that environment.

- If you are modifying a boot environment, you can take a snapshot of that environment at any stage during modifications by using the `beadm create` command. For example, if you are doing monthly upgrades to your boot environment, you can capture snapshots for each monthly upgrade. See "Create a Snapshot of a Boot Environment" on page 97.

For more information about the advantages of multiple Oracle Solaris 11 boot environments, go to:

`http://docs.oracle.com/cd/E23824_01/html/E21801/snap3.html#scrolltoc`

# ▼ Create a Boot Environment

If you want to create a backup of an existing boot environment, for example, prior to modifying the original boot environment, you can use the `beadm` command to create and mount a new boot environment that is a clone of your active boot environment. This clone is listed as an alternate boot environment in the boot menu for the compute server.

**1. Log in to the target compute server.**

```
localsys% ssh systemname -l root
Password:
Last login: Wed Nov 13 20:27:29 2011 from dhcp-vpn-r
Oracle Corporation SunOS 5.11 solaris April 2011
root@sup46:~#
```

**2. Manage ZFS boot environments with `beadm`.**

```
root@sup46:~# beadm list

BE       Active    Mountpoint    Space    Policy     Created
-------------------------------------------------------------
solaris    NR          /         2.17G    static   2011-07-13
12:01
```

---

**Note –** In the `Active` column, the first letter indicates the boot environment current status and the second letter indicates the status at next reboot. In the preceding example, `N` indicates the current (or Now) boot environment, while the `R` indicates which boot environment will be active at next reboot.

---

**3. Create a new ZFS boot environment based on the current environment.**

```
root@sup46:~# beadm create solaris_backup
root@sup46:~# beadm list

      BE       Active  Mountpoint Space  Policy     Created
-------------------------------------------------------------------
solaris         NR        /       2.17G  static   2011-07-13 12:01
solaris_backup  -         -       35.0K  static   2011-07-17 21:01
```

**4. Change to the next boot environment.**

```
root@sup46:~# beadm activate solaris_backup
root@sup46:~# beadm list

      BE       Active  Mountpoint Space  Policy     Created
-------------------------------------------------------------------
solaris_backup  R         -       2.17G  static   2011-07-17 21:01
solaris         N        /       1.86G  static   2011-07-13 12:01
```

**5. Reboot to the new boot environment.**

```
root@sup46:~# reboot
Connection to systemname closed by remote host.
Connection to systemname closed.
localsys% ssh systemname -l root
Password:
Last login: Thu Jul 14 14:37:34 2011 from dhcp-vpn-
Oracle Corporation SunOS 5.11 solaris April 2011

root@sup46:~# beadm list

      BE       Active  Mountpoint Space  Policy     Created
------------------------------------------------------------------
solaris_backup  NR        -       2.19G  static   2011-07-17 21:01
solaris          -        /       4.12G  static   2011-07-13 12:01
```

# ▼ Mount to a Different Build Environment

- **Mount to a different build environment and unmount the other build environment.**

```
root@sup46:~# beadm mount s_backup /mnt
root@sup46:~# df -k /mnt
Filesystem           1024-blocks Used     Available  Capacity  Mounted on
rpool1/ROOT/s_backup 286949376   2195449  232785749  1%        /mnt

root@sup46:~# df -k /
Filesystem           1024-blocks Used     Available  Capacity  Mounted on
rpool1/ROOT/s_backup 286949376   2214203  232785749  1%        /

root@sup46:~# ls /mnt
bin etc lib opt rpool1 system wwss
boot export media pkg sbin tmp
cdrom home micro platform scde usr
dev import mnt proc share var
devices java net re shared workspace
doe kernel nfs4 root src ws
root@sup46:~#
```

```
root@sup46:~# beadm umount solaris
root@sup46:~#
```

## ▼ Reboot to the Original Boot Environment

- **Type.**

```
root@sup46:~# beadm activate solaris
root@sup46:~# reboot
Connection to systemname closed by remote host.
Connection to systemname closed.
localsys%
ssh systemname -l root
Password: Last login: Thu Jul 14 14:37:34 2011 from dhcp-vpn-
Oracle Corporation SunOS 5.11 solaris April 2011
root@sup46:~#
```

## ▼ Create a Snapshot of a Boot Environment

You can take a snapshot of a boot environment, for backup or tracking purposes, at any stage during modifications by using the `beadm create` command.

- **Type.**

  # **beadm create** *BeName@snapshotNamedescription*

  where *BeName* is the name of an existing boot environment that you want to make a snapshot from. Enter a custom *snapshotdescription* to identify the date or purpose of the snapshot.

  Although a snapshot is not bootable, you can create a boot environment based on that snapshot by using the `-e` option in the `beadm create` command. Then you can use the `beadm activate` command to specify that this boot environment becomes the default boot environment on the next reboot.

  You can use the `beadm list -s` command to view the available snapshots for a boot environment.

# ▼ Remove Unwanted Boot Environments

- **Type.**

```
root@sup46:~# beadm list

    BE       Active  Mountpoint Space  Policy     Created
------------------------------------------------------------------
solaris_backup  -         -      13.25G  static   2011-07-17 21:19
solaris        NR         -       4.12G  static   2011-07-13 12:01

root@sup46:~# beadm destroy solaris_backup
Are you sure you want to destroy solaris_backup? This action cannot
be undone(y/[n]): y
root@sup46:~# beadm list

   BE       Active  Mountpoint Space  Policy     Created
------------------------------------------------------------------
solaris       NR        /      4.12G  static   2011-07-13 12:01

root@sup46:~#
```

# Administering DISM

These topics describe how to use dynamic intimate shared memory:



## DISM Restrictions

DISM is not supported for use on SuperCluster M6-32 Oracle Solaris environments in instances other than the ASM instance. The use of DISM on SuperCluster M6-32 outside of the ASM instance can lead to several different issues ranging from excessive swap usage (even when memory is available) to kernel panics to performance problems. It has been determined that the ASM instance is typically such a small memory footprint that it should not cause an issue.

This behavior typically occurs on instances created after installation, because Solaris 11 uses ASMM by default. To prevent this DISM issue when creating Oracle Solaris 11 instances, disable DISM. For more information see: "Disable DISM" on page 100.

To decide if DISM is appropriate for your environment, and for more information about using DISM with an Oracle database, refer to the Oracle white paper *Dynamic SGA Tuning of Oracle Database on Oracle Solaris with DISM* at:

`http://www.oracle.com/technetwork/articles/systems-hardware-architecture/using-dynamic-intimate-memory-sparc-168402.pdf`

# ▼ Disable DISM

DISM is not supported for use on SuperCluster M6-32 Oracle Solaris environments in instances other than the Oracle ASM instance. For more information, see "DISM Restrictions" on page 99.

---

**Note –** Do not disable the use of ASMM within the database, which is a very useful and desirable feature to reduce DBA management of the database.

---

- **Disable the use of DISM by the database on Oracle Solaris in one of two ways:**
  - **Unset the** `SGA_MAX_SIZE / MEMORY_MAX_TARGET / MEMORY_TARGET` **parameters, or**
  - **Ensure** `SGA_MAX_SIZE` **is set to the same value as** `SGA_TARGET` **parameter or equal to the sum of all SGA components in the instance**

  For example, to set a `64 G` SGA:

```
alter system set SGA_TARGET=64G scope=spfile;
alter system set SGA_MAX_SIZE=64G scope=spfile;
alter system set MEMORY_MAX_TARGET=0 scope=spfile;
alter system set MEMORY_TARGET=0 scope=spfile;
```

# Administering Storage Servers

The storage servers are highly optimized for use with the Oracle DB, and employ a massively parallel architecture and Exadata Smart Flash Cache to dramatically accelerate Oracle DB processing and speed I/O operations. For more information, refer to the "Storage Servers" section of the *SuperCluster M6-32 Owner's Guide: Overview*.

For general maintenance information, refer to the storage server documentation, located in the following directory on the storage servers:

```
/opt/oracle/cell/doc
```

These topics describe maintenance relevant to storage servers in SuperCluster M6-32.


- "Monitor Write-through Caching Mode" on page 101
- "Shut Down or Reboot a Storage Server" on page 103
- "Drop a Storage Server" on page 106


### Related Information

- *Oracle Exadata Storage Server Software User's Guide* for additional information about the Oracle ASM disk repair timer

# ▼ Monitor Write-through Caching Mode

The disk controller on each storage server periodically performs a discharge and charge of the controller battery. During the operation, the write cache policy changes from write-back caching to write-through caching. Write-through cache mode is slower than write-back cache mode. However, write-back cache mode has a risk of data loss if the storage server loses power or fails. For storage server releases earlier than release 11.2.1.3, the operation occurs every month. For Oracle Exadata Storage Server Software release 11.2.1.3 and later, the operation occurs every three months, for example, at 01:00 on the 17th day of January, April, July, and October.

1. **Change the start time for when the learn cycle occurs, by typing a command similar to the following.**

   ```
   CellCLI> ALTER CELL bbuLearnCycleTime=
   "2011-01-22T02:00:00-08:00"
   ```

   The time reverts to the default learn cycle time after the cycle completes.

2. **View the time for the next learn cycle.**

   ```
   CellCLI> LIST CELL ATTRIBUTES bbuLearnCycleTime
   ```

   The storage server generates an informational alert about the status of the caching mode for logical drives on the cell, for example:

   ```
   HDD disk controller battery on disk controller at adapter 0 is
   going into a learn cycle. This is a normal maintenance
   activity that occurs quarterly and runs for approximately 1 to
   12 hours. The disk controller cache might go into WriteThrough
   caching mode during the learn cycle. Disk write throughput
   might be temporarily lower during this time. The message is
   informational only, no action is required.
   ```

3. **View the status of the battery.**

   ```
   # /opt/MegaRAID/MegaCli/MegaCli64 -AdpBbuCmd -GetBbuStatus -a0
   ```

   The following is example output.

```
BBU status for Adapter: 0

BatteryType: iBBU08
Voltage: 3721 mV
Current: 541 mA
Temperature: 43 C

BBU Firmware Status:

Charging Status : Charging
Voltage : OK
Temperature : OK
Learn Cycle Requested : No
Learn Cycle Active : No
Learn Cycle Status : OK
Learn Cycle Timeout : No
I2c Errors Detected : No
Battery Pack Missing : No
Battery Replacement required : No
Remaining Capacity Low : Yes
Periodic Learn Required : No
Transparent Learn : No
```

```
Battery state:

GasGuageStatus:
Fully Discharged : No
Fully Charged : No
Discharging : No
Initialized : No
Remaining Time Alarm : Yes
Remaining Capacity Alarm: No
Discharge Terminated : No
Over Temperature : No
Charging Terminated : No
Over Charged : No

Relative State of Charge: 7 %
Charger System State: 1
Charger System Ctrl: 0
Charging current: 541 mA
Absolute State of Charge: 0%

Max Error: 0 %
Exit Code: 0x00
```

# ▼ Shut Down or Reboot a Storage Server

When performing maintenance on storage servers, it may be necessary to power down or reboot the cell. If a storage server is to be shut down when one or more databases are running, then verify that taking a storage server offline does not impact Oracle ASM disk group and database availability. The ability to take a storage server offline without affecting database availability depends on two items:

- Level of Oracle ASM redundancy used on the affected disk groups
- Current status of disks in other storage servers that have mirror copies of data on the storage server to be taken offline

**1. Check if there are other offline disks.**

```
CellCLI> LIST GRIDDISK ATTRIBUTES name WHERE
asmdeactivationoutcome != 'Yes'
```

If any grid disks are returned, then it is not safe to take a storage server offline because proper Oracle ASM disk group redundancy will not be maintained. Taking a storage server offline when one or more grid disks are in this state causes Oracle ASM to dismount the affected disk group, causing the databases to shut down abruptly.

2. **When the storage server is safe to take offline, inactivate all the grid disks.**

   ```
   CellCLI> ALTER GRIDDISK ALL INACTIVE
   ```

   The preceding command completes once all disks are inactive and offline.

3. **Verify that all grid disks are inactive to allow safe shut down of the storage server.**

   ```
   LIST GRIDDISK WHERE STATUS != 'inactive'
   ```

   If all grid disks are inactive, then you can shut down the storage server without affecting database availability.

4. **Shut down the cell.**

5. **After performing the maintenance, start the cell.**

   The cell services start automatically.

6. **Bring all grid disks online.**

   ```
   CellCLI> ALTER GRIDDISK ALL ACTIVE
   ```

   When the grid disks become active, Oracle ASM automatically synchronizes the grid disks to bring them back into the disk group.

7. **Verify that all grid disks have been successfully put online.**

   ```
   CellCLI> LIST GRIDDISK ATTRIBUTES name, asmmodestatus
   ```

   Wait until `asmmodestatus` is `ONLINE` or `UNUSED` for all grid disks. For example:

```
DATA_CD_00_dm01cel01            ONLINE
DATA_CD_01_dm01cel01            SYNCING
DATA_CD_02_dm01cel01            OFFLINE
DATA_CD_02_dm02cel01            OFFLINE
DATA_CD_02_dm03cel01            OFFLINE
DATA_CD_02_dm04cel01            OFFLINE
DATA_CD_02_dm05cel01            OFFLINE
DATA_CD_02_dm06cel01            OFFLINE
DATA_CD_02_dm07cel01            OFFLINE
DATA_CD_02_dm08cel01            OFFLINE
DATA_CD_02_dm09cel01            OFFLINE
DATA_CD_02_dm10cel01            OFFLINE
DATA_CD_02_dm11cel01            OFFLINE
```

   Oracle ASM synchronization is complete only when all grid disks show `asmmodestatus=ONLINE` or `asmmodestatus=UNUSED`. Before taking another storage server offline, Oracle ASM synchronization must complete on the restarted storage server. If synchronization is not complete, the check performed on another storage server fails. For example:

```
CellCLI> list griddisk attributes name where asmdeactivationoutcome != 'Yes'
 DATA_CD_00_dm01cel02    "Cannot de-activate due to other offline disks in the
diskgroup"
DATA_CD_01_dm01cel02    "Cannot de-activate due to other offline disks in the
diskgroup"
DATA_CD_02_dm01cel02    "Cannot de-activate due to other offline disks in the
diskgroup"
DATA_CD_03_dm01cel02    "Cannot de-activate due to other offline disks in the
diskgroup"
DATA_CD_04_dm01cel02    "Cannot de-activate due to other offline disks in the
diskgroup"
DATA_CD_05_dm01cel02    "Cannot de-activate due to other offline disks in the
diskgroup"
DATA_CD_06_dm01cel02    "Cannot de-activate due to other offline disks in the
diskgroup"
DATA_CD_07_dm01cel02    "Cannot de-activate due to other offline disks in the
diskgroup"
DATA_CD_08_dm01cel02    "Cannot de-activate due to other offline disks in the
diskgroup"
```

```
DATA_CD_09_dm01cel02    "Cannot de-activate due to other offline disks in the
diskgroup"
DATA_CD_10_dm01cel02    "Cannot de-activate due to other offline disks in the
diskgroup"
DATA_CD_11_dm01cel02    "Cannot de-activate due to other offline disks in the
diskgroup"
```

## ▼ Drop a Storage Server

1. **From Oracle ASM, drop the Oracle ASM disks on the physical disk.**

   **ALTER DISKGROUP** *diskgroup-name* **DROP DISK** *asm-disk-name*

   To ensure the correct redundancy level in Oracle ASM, wait for the rebalance to complete before proceeding.

2. **Remove the IP address entry from the** `cellip.ora` **file on each database server that accesses the storage server.**

3. **From the storage server, drop the grid disks, cell disks, and cell on the physical disk.**

   **DROP CELLDISK** *celldisk-on-this-lun* **FORCE**

4. **Shut down all services on the storage server.**

5. **Power down the cell.**

   See "Shut Down or Reboot a Storage Server" on page 103 for additional information.

# Glossary

## A

**Application Domain** A domain that runs Oracle Solaris and client applications.

**ASMM** Automatic shared memory management.

**ASR** Auto Service Request. A feature of Oracle or Sun hardware that automatically opens service requests when specific hardware faults occur. ASR is integrated with MOS and requires a support agreement. See also *MOS*.

## B

**base configuration PDomain** A SuperCluster M6-32 configuration comprised of two or four PDomains, where one DCU is associated with each PDomain. A base configuration PDomain can reside on a single compute server or can be split across two compute servers. See also *compute server*, *DCU*, *extended configuration PDomain*, and *PDomain*.

## C

**CFM** Cubic feet per minute.

**Cisco Catalyst Ethernet switch** Provides the SuperCluster M6-32 management network. Referred to in this documentation using the shortened name "Ethernet management switch." See also *Ethernet management switch*.

**CMP** Chip multiprocessing. Each CMU contains 2 CMP processors. The compute server can contain a maximum of 32 CMPs.

**CMU** CPU memory unit. Each CMU in the compute server contains two CMPs and two sets of DIMM slots.

**COD** Capacity on Demand.

**compute server** Shortened name for the SPARC M6-32 server, a major component of SuperCluster M6-32. See also *SPARC M6-32 server*.

# D

**Database Domain** The domain that contains the SuperCluster M6-32 database.

**DB** Oracle Database.

**DCM** Domain configuration management. The reconfiguration of boards in PDomains for Enterprise-class systems. See also *PDomain*.

**DCU** Domain configurable unit. The smallest building block for PDomains. Each DCU in the compute server contains two or four CMUs and one IOU. See also *PDomain*.

**dedicated domain** A SuperCluster LDom category that includes the domains configured at installation time as either a Database Domain or an Application Domain (running the Oracle Solaris 10 or Oracle Solaris 11 OS). Dedicated domains have direct access to the 10GbE NICs and IB HCAs (and Fibre Channel cards, if present). See also *Database Domain* and *Application Domain*.

**DHCP** Dynamic Host Configuration Protocol. Software that automatically assigns IP addresses to clients on a TCP/IP network. See also *TCP*.

**DIMM** Dual in-line memory module.

**DISM** Dynamic intimate shared memory.

# E

**EECS** Oracle Exalogic Elastic Cloud software.

**EMS** Express module SAS. Each EMS contains two 10GBASE-T network connections and provides access to four hard drives on the compute server.

**EPO switch** Emergency power-off switch.

**ESD** Electrostatic discharge.

**Ethernet management switch** Shortened name for the Cisco Catalyst Ethernet switch. See also *Cisco Catalyst Ethernet switch*.

**expansion rack** Shortened name for optional Oracle Exadata Storage Expansion Racks (up to 17) that can be added to SuperCluster M6-32. See also *Oracle Exadata Storage Expansion Rack*.

**extended configuration PDomain** A SuperCluster M6-32 configuration comprised of two PDomains, where two DCUs are associated with each PDomain. An extended configuration PDomain can reside on a single compute server or can be split across two compute servers. See also, *base configuration PDomain*, *compute server*, *DCU*, and *PDomain*.

# F

**FAN** Fast application notification event.

**FCoE** Fibre Channel over Ethernet.

**FM** Fan module.

**FMA** Fault management architecture. A feature of Oracle Solaris servers that includes error handlers, structured error telemetry, automated diagnostic software, response agents, and messaging.

**FRU** Field-replaceable unit.

**fully-populated DCU configuration** A configuration where each DCU in the compute servers contains four CMUs. See also *DCU* and *half-populated DCU configuration*.

# G

**GB** Gigabyte. 1 gigabyte = 1024 megabytes.

**GbE** Gigabit Ethernet.

**GNS** Grid Naming Service.

# H

**half-populated DCU configuration** A configuration where each DCU in the compute servers contains two CMUs. See also *DCU* and *fully-populated DCU configuration*.

**HCA** Host channel adapter.

**HDD** Hard disk drive. In Oracle Solaris OS output, HDD can refer to hard disk drives or SSDs.

# I

**IB** InfiniBand.

**IB switch** Shortened name for the Sun Datacenter InfiniBand Switch 36. See also *leaf switch*, *spine switch*, and *Sun Datacenter InfiniBand Switch 36*.

**ILOM** See *Oracle ILOM*.

**I/O Domain** If you have Root Domains, you create I/O Domains with your choice of resources at the time of your choosing. The I/O Domain Creation tool enables you to assign resources to I/O Domains from the CPU and memory repositories, and from virtual functions hosted by Root Domains. When you create an I/O Domain, you assign it as a Database Domain or Application Domain running the Oracle Solaris 11 OS. See also *Root Domain*.

**IOU** I/O unit. The compute server contains up to 4 IOUs, one for each DCU. Each IOU supports up to 16 PCIe slots, 8 10GBASE-T ports on 4 EMS modules, and 8 drives.

**IPMI** Intelligent Platform Management Interface.

**IPMP** IP network multipathing.

**iSCSI** Internet Small Computer System Interface.

# K

**KVMS** Keyboard video mouse storage.

# L

**leaf switch** Two of the IB switches are configured as leaf switches, the third is configured as a spine switch. See also *IB switch*.

**LDom** Logical domain. A virtual machine comprising a discrete logical grouping of resources that has its own operating system and identity within a single computer system. LDoms are created using Oracle VM Server for SPARC software. See also *Oracle VM Server for SPARC*.

# M

**MIB** Management information base.

**MOS** My Oracle Support.

# N

**NET MGT** The network management port on an SP. See also *SP*.

**NIC** Network interface card.

**NUMA** Nonuniform memory access.

# O

**OBP** OpenBoot PROM. Firmware on SPARC servers that enables the server to load platform-independent drivers directly from devices, and provides an interface through which you can boot the compute server and run low-level diagnostics.

**Oracle ASM** Oracle Automatic Storage Management. A volume manager and a file system that supports Oracle databases.

**OCM** Oracle Configuration Manager.

**ONS** Oracle Notification Service.

**Oracle Exadata Storage Expansion Rack** — Optional expansion racks (in full, half, or quarter configurations) that can be added to SuperCluster M6-32 systems that require additional storage. Referred to in this documentation using the shortened name "expansion rack." See also *expansion rack*.

**Oracle ILOM** — Oracle Integrated Lights Out Manager. Software on the SP that enables you to manage a server independently from the operating system. See also *SP*.

**Oracle Solaris OS** — Oracle Solaris operating system.

**Oracle SuperCluster** — Refers to all Oracle SuperCluster models.

**Oracle SuperCluster M6-32** — Full name of this SuperCluster model. Referred to in this documentation using the shortened name "SuperCluster M6-32." See also *SuperCluster M6-32*.

**Oracle SuperCluster M6-32 storage rack** — Full name of first storage rack that contains the storage servers, ZFS storage appliance, IB switches, and Ethernet management switch. Referred to in this documentation using the shortened name "storage rack." See also *storage rack*.

**Oracle VM Server for SPARC** — SPARC server virtualization and partitioning technology. See also *LDom*.

**Oracle VTS** — Oracle Validation Test Suite. An application, preinstalled with Oracle Solaris, that exercises the system, provides hardware validation, and identifies possible faulty components.

**Oracle XA** — Oracle's implementation of the X/Open distributed transaction processing XA interface that is included in Oracle DB software.

**Oracle ZFS ZS3-ES storage appliance** — Located in the storage rack, it provides SuperCluster M6-32 with shared storage capabilities. Referred to in this documentation using the shortened name "ZFS storage appliance." See also *ZFS storage appliance*.

**OS** — Operating system.

# P

**parked resources** — CPU and memory resources that are set aside in the CPU and memory repositories. You assign parked resources to I/O Domains with the I/O Domain Creation tool.

**PCIe** — Peripheral Component Interconnect Express.

**PDomain** — Physical domain. Each PDomain on the compute server is an independently configurable and bootable entity with full hardware domain isolation for fault isolation and security purposes. See also *compute server*, *DCU*, and *SSB*.

**PDomain-SPP** The lead SPP of a PDomain. The PDomain-SPP on the compute server manages tasks and provides rKVMS service for that PDomain. See also *PDomain*.

**PDU** Power distribution unit.

**PF** Physical function. Functions provided by physical I/O devices, such as the IB HCAs, 10GbE NICs, and any Fibre Channel cards installed in the PCIe slots. Logical devices, or virtual functions (VFs), are created from PFs, with each PF hosting 32 VFs.

**POST** Power-on self-test. A diagnostic that runs when the compute server is powered on.

**PS** Power supply.

**PSDB** Power system distribution board.

**PSH** Predictive self healing. An Oracle Solaris OS technology that continuously monitors the health of the compute server and works with Oracle ILOM to take a faulty component offline if needed.

# Q

**QMU** Quarterly maintenance update.

**QSFP** Quad small form-factor, pluggable. A transceiver specification for 10GbE technology.

# R

**RAC** Real Application Cluster.

**RCLB** Runtime connection load balancing.

**rKVMS** Remote keyboard video mouse and storage.

**root complex** CMP circuitry that provides the base to a PCIe I/O fabric. Each PCIe I/O fabric consists of the PCIe switches, PCIe slots, and leaf devices associated with the root complex.

**Root Domain** A logical domain that is configured at installation time. Root Domains are required if you plan to configure I/O Domains. Root Domains host PFs from which I/O Domains derive VFs. The majority of Root Domain CPU and memory resources are parked for later use by I/O Domains.

# S

**SAS** Serial attached SCSI.

**SATA** Serial advance technology attachment.

**scalability** The ability to increase (or scale up) processing power in a compute server by combining the server's physical configurable hardware (see also *DCU*) into one or more logical groups (see also *PDomain*).

**SER MGT** The serial management port on an SP. See also *SP*.

**SCAN** Single Client Access Name. A feature used in RAC environments that provides a single name for clients to access any Oracle Database running in a cluster. See also *RAC*.

**SDP** Session Description Protocol.

**SFP and SFP+** Small form-factor pluggable standard. SFP+ is a specification for a transceiver for 10GbE technology.

**SGA** System global area.

**SMF** Service Management Facility.

**SNEEP** Serial number in EEPROM.

**SNMP** Simple Management Network Protocol.

**SP** Service processor. A processor, separate from the host, that monitors and manages the host no matter what state the host is in. The SP runs Oracle ILOM, which provides remote lights out management. In SuperCluster M6-32, SPs are located on the compute servers, storage servers, ZFS storage appliance controllers, and IB switches. See also *Oracle ILOM*.

**SPARC M6-32 server** A major component of SuperCluster M6-32 that provides the main compute resources. Referred to in this documentation using the shortened name "compute server." See also *compute server*.

**spine switch** One of the SuperCluster M6-32 IB switches that is configured as a spine switch. See also *IB switch* and *leaf switch*.

**SPP** Service processor proxy. One SPP in the compute server is assigned to manage each PDomain. SPPs monitor environmental sensors and manage the CMUs, memory controllers, and DIMMs within the DCU. See also *PDomain-SPP*.

**SR-IOV Domain** Single-Root I/O Virtualization Domain -- A SuperCluster logical domain category that includes Root Domains and I/O Domains. This category of domains support single-root I/O virtualization. See also *I/O Domain* and *Root Domain*.

**SSB** Scalability switch board in the compute server.

**SSD** Solid state drive.

**STB** Oracle Services Tool Bundle.

**storage rack** Shortened name for the Oracle SuperCluster M6-32 storage rack that contains the storage servers. See also *Oracle SuperCluster M6-32 storage rack*.

**storage server** Storage servers in SuperCluster M6-32.

**Sun Datacenter InfiniBand Switch 36** Interconnects SuperCluster M6-32 components on a private network. Referred to in this documentation using the shortened name "IB switch." See also *IB switch*, *leaf switch*, and *spine switch*.

**SuperCluster M6-32** Shortened name for Oracle SuperCluster M6-32. See also *Oracle SuperCluster M6-32*.

## T

**TCP** Transmission Control Protocol.

**TNS** Transparent Network Substrate.

**TPM** Trusted platform module.

## U

**UPS** Uninterruptible power supply.

# V

**VAC** Voltage alternating current.

**VF** Virtual function. Logical I/O devices that are created from PFs, with each PF hosting 32 VFs.

**VIP** Virtual IP.

**VLAN** Virtual local area network.

**VNET** Virtual network.

# W

**WWN** World Wide Name.

# X

**XA** See *Oracle XA*.

# Z

**ZFS** A file system with added volume management capabilities. ZFS is the default file system in Oracle Solaris 11.

**ZFS storage appliance** Shortened name for Oracle ZFS Storage ZS3-ES storage appliance. See also *Oracle ZFS ZS3-ES storage appliance*.

**ZFS storage controller** Servers in the Oracle ZFS ZS3-ES storage appliance that manage the storage appliance. See also *ZFS storage appliance*.

# Index