Oracle® SuperCluster M8 and SuperCluster M7 Administration Guide

Part No: E58635-11
June 2020

ORACLE

Oracle SuperCluster M8 and SuperCluster M7 Administration Guide

Part No: E58635-11

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info or visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs if you are hearing impaired.

Référence: E58635-11

Restrictions de licence/Avis d'exclusion de responsabilité en cas de dommage indirect et/ou consécutif

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf stipulation expresse de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, accorder de licence, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Exonération de garantie

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Avis sur la limitation des droits

Si ce logiciel, ou la documentation qui l'accompagne, est livré sous licence au Gouvernement des Etats-Unis, ou à quiconque qui aurait souscrit la licence de ce logiciel pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software" or "commercial computer software documentation" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

Avis sur les applications dangereuses

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer un risque de dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour des applications dangereuses.

Marques

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

Intel et Intel Inside sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. AMD, Epyc, et le logo AMD sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée de The Open Group.

Avis d'exclusion de responsabilité concernant les services, produits et contenu tiers

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers, sauf mention contraire stipulée dans un contrat entre vous et Oracle. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation, sauf mention contraire stipulée dans un contrat entre vous et Oracle.

Date de publication et mention de la version préliminaire de Disponibilité Générale ("Pre-GA")

Version préliminaire de Disponibilité Générale ("Pre-GA") : 15.01.2020

Avis sur la version préliminaire de Disponibilité Générale ("Pre-GA") de la documentation

Si ce document est fourni dans la Version préliminaire de Disponibilité Générale ("Pre-GA") à caractère public ou privé :

Cette documentation est fournie dans la Version préliminaire de Disponibilité Générale ("Pre-GA") et uniquement à des fins de démonstration et d'usage à titre préliminaire de la version finale. Celle-ci n'est pas toujours spécifique du matériel informatique sur lequel vous utilisez ce logiciel. Oracle Corporation et ses affiliés déclinent expressément toute responsabilité ou garantie expresse quant au contenu de cette documentation. Oracle Corporation et ses affiliés ne sauraient en aucun cas être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'utilisation de cette documentation.

Mention sur les informations confidentielles Oracle

INFORMATIONS CONFIDENTIELLES ORACLE. Destinées uniquement à un usage autorisé. Ne pas distribuer à des tiers.

Avis sur la reconnaissance du revenu

Si ce document est fourni dans la Version préliminaire de Disponibilité Générale ("Pre-GA") à caractère privé :

Les informations contenues dans ce document sont fournies à titre informatif uniquement et doivent être prises en compte en votre qualité de membre du customer advisory board ou conformément à votre contrat d'essai de Version préliminaire de Disponibilité Générale ("Pre-GA") uniquement. Ce document ne constitue en aucun cas un engagement à fournir des composants, du code ou des fonctionnalités et ne doit pas être retenu comme base d'une quelconque décision d'achat. Le développement, la commercialisation et la mise à disposition des fonctions ou fonctionnalités décrites restent à la seule discrétion d'Oracle.

Accessibilité de la documentation

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

Accès aux services de support Oracle

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

- **Overview** – Describes how to monitor and administer Oracle SuperCluster M8 and Oracle SuperCluster M7 engineered systems using tools and utilities provided with the SuperCluster software.
- **Audience** – Technicians, system administrators, and authorized service providers
- **Required knowledge** – Advanced experience administering computer systems.

Product Documentation Library

Documentation and resources for this product and related products are available at http://docs.oracle.com/cd/E58626_01/.

Feedback

Provide feedback about this documentation at http://www.oracle.com/goto/docfeedback.

Understanding SuperCluster Software

These topics describe administration resources.

Related Information

Administration Resources

Use this table to identify the task you want to perform and where to find information about the task.

Topic	Description	Links
Perform administrative tasks using the Oracle Solaris OS.	Oracle Solaris documentation includes information about getting started, booting the OS, administering networks, managing users, creating virtual environments, and setting up an application development environment.	`http://docs.oracle.com/en/operating-systems`
Perform administrative tasks on the Oracle Database.	Oracle Database documentation provides information about getting started, managing users, application development, using Enterprise Manager plug-ins.	`http://docs.oracle.com/en/database`
Perform administrative tasks on the ZFS storage appliance.	Oracle ZFS Storage ZS3-ES documentation includes an administration guide, analytics guide, and release notes.	`http://www.oracle.com/goto/zs3-es/docs`

Topic	Description	Links
Administer logical domains.	Oracle VM Server for SPARC documentation includes administration, security, and best practices information.	`https://docs.oracle.com/en/virtualization/oracle-vm-server-sparc/`
Administer I/O domains.	This document describes how to administer Oracle SuperCluster I/O Domains using the Oracle I/O Domain Creation tool.	Refer to the *Oracle I/O Domain Administration Guide*.
Perform administrative tasks through Oracle ILOM.	Oracle ILOM documentation includes information on getting started, administration, monitoring and diagnostics, and configuring Oracle ILOM with SNMP and IPMI.	`http://www.oracle.com/goto/ilom/docs`
Administer Enterprise Manager Ops Center.	Oracle Enterprise Manager Ops Center documentation includes getting started, installation, and administration information.	`http://docs.oracle.com/en/enterprise-manager/?tab=2`
Administer cluster software.	Oracle Solaris Cluster documentation includes getting started, installation, administration, and security information.	`http://docs.oracle.com/en/operating-systems`
Identify the version of SuperCluster software.	In this guide.	“Identify the SuperCluster Software Version” on page 17
Identify utilities that are unique to SuperCluster.	In this guide.	“SuperCluster Tools” on page 17
Create a Database listener in the IB network.	In this guide.	“Create a DB Listener on the IB Network” on page 18
Power on the system.	In this guide.	“Starting and Shutting Down SuperCluster” on page 27
Shut down or power off the system.	In this guide.	“Starting and Shutting Down SuperCluster” on page 27
Configure Oracle Engineered Systems Hardware Manager.	In this guide.	“Configuring Oracle Engineered Systems Hardware Manager” on page 37
Monitor the system.	In this guide.	“Monitoring SuperCluster Systems (OCM)” on page 43
Administer the tuning scripts.	In this guide.	“Tuning SuperCluster Systems (`ssctuner`)” on page 65
Configure CPU and memory resources.	In this guide.	“Configuring CPU and Memory Resources (`osc-setcoremem`)” on page 77
Obtain the EM Exadata plug-in.	In this guide.	“Obtaining the EM Exadata Plug-in” on page 127
Administer Exalogic software.	In this guide.	“Configuring the Exalogic Software” on page 129

Related Information

- “Identify the SuperCluster Software Version” on page 17
- “SuperCluster Tools” on page 17
- “Understanding Clustering Software” on page 23

▼ Identify the SuperCluster Software Version

If you want to identify the specific version of SuperCluster software, perform this procedure.

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

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

In the output, the numbers appended to `ssc` represent the software version.

Related Information

- "Administration Resources" on page 15
- "SuperCluster Tools" on page 17
- "Understanding Clustering Software" on page 23

SuperCluster Tools

In addition to the full compliment of Oracle Solaris OS and Oracle Database software features, SuperCluster provides additional tools and utilities that help you manage the system.

The Python programming language is used to create SuperCluster-specific utilities such as the SuperCluster Virtual Assistant. The SuperCluster 2020 Q2 Quarterly Patch Update (SuperCluster software version 3.0.0.1875 or later) provides utilities that are based on Python 3. Prior to the 2020 Q2 update, the utilities were developed using Python 2. To determine your software version, see "Identify the SuperCluster Software Version" on page 17.

This table lists the additional tools and utilities.

Tool	Description	Links
SuperCluster Virtual Assistant (Formerly call the 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	"Configuring CPU and Memory Resources (`osc-setcoremem`)" on page 77

Tool	Description	Links
	allocated CPU resources, ensuring optimal performance by minimizing NUMA effects.	
`ssctuner`	Monitors and tunes various parameters through a set of scripts and configuration files that run on SuperCluster Oracle Solaris 10 and Oracle Solaris 11 global zones.	"Tuning SuperCluster Systems (`ssctuner`)" on page 65

Related Information

- "Administration Resources" on page 15
- "Identify the SuperCluster Software Version" on page 17
- "Understanding Clustering Software" on page 23

▼ Create a DB Listener on the IB Network

1. **Edit the `/etc/hosts` file on each Database Domain in the cluster to add the virtual IP addresses to be used for the IB network.**
   For this example, these addresses are added to the /etc/hosts file on both nodes. Addresses starting with x.y are VIPs on the 10 GbE client access network. Those starting with a.b. are on the storage IB.
   Ensure that the IP addresses that you specify are not already in use.

   x.y`.132.103` *node1*`-vip`
   `a.b.30.65` *node1*`-vipIB`
   x.y`.132.104` *node2*`-vip`
   `a.b.30.66` *node2*`-vipIB`

2. **As superuser on node 1, register the IB subnet in the Grid by running these commands from the `bin` directory of the Grid home (the usual Grid home is `/u01/app/11.2.x.y/grid/bin`). Also register the two VIPs:**

```
root@node1# srvctl add network -k 2 -S 192.0.2.0/255.255.252.0/stor_ipmp0 -w static -v
root@node1# srvctl add vip -n node1 -A node1-vipIB/255.255.252.0/stor_ipmp0 -k 2
root@node1# srvctl add vip -n node2 -A node2-vipIB/255.255.252.0/stor_ipmp0 -k 2
root@node1# srvctl start vip -i node1-vipIB
root@node1# srvctl start vip -i node2-vipIB
```

3. **As the grid user on node 1, create a `LISTENER_IP` by running these commands from the Grid home `bin` directory (if there is no grid account use the oracle account):**

Note - If you want to enable SDP in addition to TCP, replace: `grid@node1$` **`srvctl add listener -l LISTENER_IB -k 2 -p 1522`** with: `grid@node1$` **`srvctl add listener -l LISTENER_IB -k 2 -p TCP:1522,/SDP:1522`**

```
grid@node1$ srvctl add listener -l LISTENER_IB -k 2 -p 1522
grid@node1$ srvctl start listener -l LISTENER_IB
grid@node1$ srvctl status listener -l LISTENER_IB
  Listener LISTENER_IB is enabled
  Listener LISTENER_IB is running on node(s): node1,node2
```

4. As oracle one node 1, check that the listener is properly registered by running these commands from the Grid home `bin` directory:

```
oracle@node1$ lsnrctl status LISTENER_IB

   LSNRCTL for Solaris: Version 11.2.0.3.0 - Production on 30-SEP-2018 15:49:2
  Copyright (c) 1991, 2018, Oracle.  All rights reserved.
  Connecting to (DESCRIPTION=(ADDRESS=(PROTOCOL=IPC)(KEY=LISTENER_IB)))
  STATUS of the LISTENER
  ------------------------
  Alias                     LISTENER_IB
  Version                   TNSLSNR for Solaris: Version 11.2.0.3.0 - Production
  Start Date                30-SEP-2018 15:48:38
  Uptime                    0 days 0 hr. 0 min. 50 sec
  Trace Level               off
  Security                  ON: Local OS Authentication
  SNMP                      OFF
  Listener Parameter File   /u01/app/11.2.x.y/grid/network/admin/listener.ora
  Listener Log File         /u01/app/11.2.x.y/grid/log/diag/tnslsnr/node1/listener_ib/
alert/log.xml
  Listening Endpoints Summary...
  (DESCRIPTION=(ADDRESS=(PROTOCOL=ipc)(KEY=LISTENER_IB)))
  (DESCRIPTION=(ADDRESS=(PROTOCOL=sdp)(HOST=a.b.30.65)(PORT=1522)))
  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=a.b.30.65)(PORT=1522)))
  The listener supports no services
  The command completed successfully
```

5. As oracle on node 1, edit `tnsnames.ora` file located in the Grid home `network/admin` directory.

Note - If the IB listener is to be used only with a specific standalone database, then modify the `tnsnames.ora` from the Oracle home instead.

---

a. Get the database name DBNAME from the Grid home bin directory.

This value is assigned to SERVICE_NAME in the tnsnames.ora file.

```
oracle@node1$ srvctl config database
  DBNAME
```

b. Use vi with the :set list option to view and remove invisible characters.

Invisible characters such as tabs (^l) and end of line ($) can create problems during the next step and should be removed.

```
## BEGIN
DBNAME_IB =
(DESCRIPTION =
        (LOAD_BALANCE=on)
        (ADDRESS = (PROTOCOL = TCP)(HOST = node1-vipIB)(PORT = 1522))
        (ADDRESS = (PROTOCOL = TCP)(HOST = node2-vipIB)(PORT = 1522))
        (CONNECT_DATA =
                (SERVER = DEDICATED)
                (SERVICE_NAME = DBNAME)
        )
)

LISTENER_IBREMOTE =
(DESCRIPTION =
        (ADDRESS_LIST =
                (ADDRESS = (PROTOCOL = TCP)(HOST =  node2-vipIB)(PORT = 1522))
        )
)

LISTENER_IBLOCAL =
(DESCRIPTION =
        (ADDRESS_LIST =
                (ADDRESS = (PROTOCOL = TCP)(HOST =  node1-vipIB)(PORT = 1522))
                (ADDRESS = (PROTOCOL = SDP)(HOST =  node1-vipIB)(PORT = 1522))
        )
)

LISTENER_IPLOCAL =
(DESCRIPTION =
        (ADDRESS_LIST =
                (ADDRESS = (PROTOCOL = TCP)(HOST = node1-vip)(PORT = 1521))
        )
)

LISTENER_IPREMOTE =
(DESCRIPTION =
        (ADDRESS_LIST =
                (ADDRESS = (PROTOCOL = TCP)(HOST = node2-vip)(PORT = 1521))
```

```
            )
    )
    ## END
```

6. **As oracle on node 2, add these lines to the `tnsnames.ora` file in the Grid home network/admin directory:**

```
DBNAME_IB =
(DESCRIPTION =
        (LOAD_BALANCE=on)
        (ADDRESS = (PROTOCOL = TCP)(HOST = node1-vipIB)(PORT = 1522))
        (ADDRESS = (PROTOCOL = TCP)(HOST = node2-vipIB)(PORT = 1522))
        (CONNECT_DATA =
                (SERVER = DEDICATED)
                (SERVICE_NAME = DBNAME)
        )
)

LISTENER_IBREMOTE =
(DESCRIPTION =
        (ADDRESS_LIST =
                (ADDRESS = (PROTOCOL = TCP)(HOST =  node1-vipIB)(PORT = 1522))
        )
)

LISTENER_IBLOCAL =
(DESCRIPTION =
        (ADDRESS_LIST =
                (ADDRESS = (PROTOCOL = TCP)(HOST =  node2-vipIB)(PORT = 1522))
                (ADDRESS = (PROTOCOL = SDP)(HOST =  node2-vipIB)(PORT = 1522))
        )
)

LISTENER_IPLOCAL =
(DESCRIPTION =
        (ADDRESS_LIST =
                (ADDRESS = (PROTOCOL = TCP)(HOST = node2-vip)(PORT = 1521))
        )
)

LISTENER_IPREMOTE =
(DESCRIPTION =
        (ADDRESS_LIST =
                (ADDRESS = (PROTOCOL = TCP)(HOST = node1-vip)(PORT = 1521))
        )
)
## END
```

7. **As oracle, register the new listener at the database level.**

   Make sure to set `ORACLE_SID` to the proper value before running sqlplus.

   Note - If the `LOCAL_LISTENER` and `REMOTE_LISTENER` parameters are defined in the database, they should not appear in the `LISTENER_NETWORKS` parameter or cross registration occurs and connections will be redirected across networks. In this example, if `LOCAL_LISTENER` and `REMOTE_LISTENER` are defined respectively to `LISTENER_IPLOCAL` and `LISTERNER_IPREMOTE`, then the SQL statements should be as follows.

   On Node 1:

   ```
   SQL> alter system set listener_networks='((NAME=network2)
   (LOCAL_LISTENER=LISTENER_IBLOCAL)(REMOTE_LISTENER=LISTENER_IBREMOTE))';
   ```

   On Node 2:

   ```
   SQL> alter system set listener_networks='((NAME=network2)
   (LOCAL_LISTENER=LISTENER_IBLOCAL)(REMOTE_LISTENER=LISTENER_IBREMOTE))';
   ```

   Register the listener on both nodes, as shown:

   ```
   oracle@node1$ export ORACLE_SID=DBNAME1
   oracle@node1$ sqlplus / as sysdba
   SQL> 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;


   oracle@node2$ export ORACLE_SID=DBNAME2
   oracle@node2$ sqlplus / as sysdba
   SQL> 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. **As oracle on node 1, restart the `LISTENER_IB` and check its status:**

   ```
   oracle@node1$ srvctl stop listener -l LISTENER_IB
   oracle@node1$ srvctl start listener -l LISTENER_IB
   oracle@node1$ export TNS_ADMIN=/u01/app/11.2.x.y/grid/network/admin
   oracle@node1$ lsnrctl status LISTENER_IB

   LSNRCTL for Solaris: Version 11.2.0.3.0 - Production on 04-SEP-2018 11:53:20
   Copyright (c) 1991, 2018, Oracle.  All rights reserved.
   ```

```
Connecting to (DESCRIPTION=(ADDRESS=(PROTOCOL=IPC)(KEY=LISTENER_IB)))
STATUS of the LISTENER
------------------------
Alias                     LISTENER_IB
Version                   TNSLSNR for Solaris: Version 11.2.0.3.0 - Production
Start Date                04-SEP-2018 11:52:32
Uptime                    0 days 0 hr. 0 min. 47 sec
Trace Level               off
Security                  ON: Local OS Authentication
SNMP                      OFF
Listener Parameter File   /u01/app/11.2.0.3/grid/network/admin/listener.ora
Listener Log File         /u01/app/11.2.0.3/grid/log/diag/tnslsnr/rmb-zpr-db-int1/
listener_ib/alert/log.xml
Listening Endpoints Summary...
  (DESCRIPTION=(ADDRESS=(PROTOCOL=ipc)(KEY=LISTENER_IB)))
  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=a.b.30.86)(PORT=1522)))
Services Summary...
Service "DBNAME" has 2 instance(s).
  Instance "DBNAME1", status READY, has 1 handler(s) for this service...
  Instance "DBNAME2", status READY, has 1 handler(s) for this service...
The command completed successfully
```

Related Information

- “Configuring the Exalogic Software” on page 129

Understanding Clustering Software

Clustering software is used on multiple interconnected servers so that they appear to end users and applications as one server. For SuperCluster systems, clustering software clusters certain logical domains on the compute nodes together with the same type of domain on other compute node. The benefits of clustering software include:

- Reduce or eliminate system downtime because of software or hardware failure.
- Ensure availability of data and applications to end users, regardless of the kind of failure that would normally take down a single-server system.
- Increase application throughput by enabling services to scale to additional processors by adding nodes to the cluster and balancing the load.
- Provide enhanced availability of the system by enabling you to perform maintenance without shutting down the entire cluster.

SuperCluster M8 and SuperCluster M7 use the following clustering software:

- "Cluster Software for Database Domains" on page 24
- "Cluster Software for Oracle Solaris Application Domains" on page 24

Cluster Software for Database Domains

Oracle Database 11g Real Application Clusters (Oracle RAC) enables the clustering of the Oracle Database on the Database Domain. Oracle RAC uses Oracle Clusterware for the infrastructure to cluster the Database Domain on the compute nodes.

Oracle Clusterware is a portable cluster management solution that is integrated with the Oracle database. Oracle Clusterware is also a required component for using Oracle RAC. Oracle Clusterware enables you to create a clustered pool of storage to be used by any combination of single-instance and Oracle RAC databases.

Single-instance Oracle databases have a one-to-one relationship between the Oracle database and the instance. Oracle RAC environments, however, have a one-to-many relationship between the database and instances. In Oracle RAC environments, the cluster database instances access one database. The combined processing power of the multiple servers can provide greater throughput and scalability than is available from a single server. Oracle RAC is the Oracle Database option that provides a single system image for multiple servers to access one Oracle database.

Oracle RAC technology provides high availability and scalability for all application types. The Oracle RAC infrastructure is also a key component for implementing the Oracle enterprise grid computing architecture. Having multiple instances access a single database prevents the server from being a single point of failure. Applications that you deploy on Oracle RAC databases can operate without code changes.

Related Information

- "Cluster Software for Oracle Solaris Application Domains" on page 24

Cluster Software for Oracle Solaris Application Domains

The Oracle Solaris Cluster software is an optional clustering tool used for the Oracle Solaris Application Domains. On SuperCluster M8 and SuperCluster M7, the Oracle Solaris Cluster software is used to cluster the Application Domains.

Related Information

- “Cluster Software for Database Domains” on page 24

Starting and Shutting Down SuperCluster

These topics describe how to power on, shut down, and power down SuperCluster M8 and SuperCluster M7.

Cautions

These cautions apply to SuperCluster M8 and SuperCluster M7.

Caution - Do not touch the parts of this product that use high-voltage power. Touching these parts 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

1. **Turn on both circuit breakers in each rack.**
   Power is applied to the components in the rack. The storage servers, compute servers, and the ZFS storage appliance enter standby mode.

2. **Boot each ZFS storage appliance.**
3. **Boot each compute node.**
4. **Boot each storage server.**

Related Information

- “Shutting Down SuperCluster Gracefully” on page 28
- “Power Off SuperCluster in an Emergency” on page 35

Shutting Down SuperCluster Gracefully

To shut down but not power off SuperCluster, only perform the shut down procedures in the order listed.

If you plan to power off SuperCluster, perform all the tasks in the order listed.

Note - Depending on the configuration of your SuperCluster, some of the procedures might not apply to your situation.

Step	Description	Links
1.	Prepare users and the system for a shutdown.	“Prepare Users and the System for a Shutdown” on page 28
2.	Obtain system configuration information needed for shutdown tasks.	“Obtain Information Needed for Shutdown Tasks” on page 29
3.	Shut down zones.	“Shut Down Zones” on page 30
4.	Shut down guest domains.	“Shut Down the Guest Domains” on page 31
5.	Shut down the primary domains.	“Shut Down the Primary Domains” on page 33
6.	Shut down the storage servers.	“(If Needed) Shut Down the Storage Servers” on page 34
7.	Power off the ZFS storage appliance.	“(If Needed) Power Off the ZFS Storage Appliance” on page 35
8.	Power off the racks.	“(If Needed) Power Off the Racks” on page 35

▼ Prepare Users and the System for a Shutdown

1. **For a complete system shutdown, ensure that users are notified and not using the system.**

2. **Follow applicable procedures and documentation to shut down any running applications and databases.**

Related Information

- "Obtain Information Needed for Shutdown Tasks" on page 29
- "Shutting Down SuperCluster Gracefully" on page 28

▼ Obtain Information Needed for Shutdown Tasks

The information gathered in this procedure is used in later procedures.

1. **Identify the domain configuration.**

   a. **Login to Oracle ILOM on the active SP for a PDomain.**

   b. **Start the host console.**
   In this example, replace *x* with the number of the host.

```
-> start /HOSTx/console
Are you sure you want to start /HOSTx/console [y/n] y
Serial console started. To stop, type #.
```

   c. **Login to the primary domain as root.**

```
sccn01 console login: root
Password:
Last login: Thur Feb 9 16:52:10 form /dev/console
Oracle Corporation      SunOS 511     11.2      Apr 2016
```

   d. **List the domains for this PDomain.**

```
# ldm list
NAME          STATE      FLAGS   CONS    VCPU  MEMORY   UTIL  NORM  UPTIME
primary       active     -n-cv-  UART    192   1047296M 0.1%  0.1%  21d 12h 55m
ssccn1-dom1   active     -n----  5001    192   1T       0.0%  0.0%  21d 5h 58m
ssccn1-dom2   active     -n----  5002    192   1T       0.0%  0.0%  21d 5h 49m
ssccn1-dom3   active     -n----  5003    192   1T       0.1%  0.1%  21d 5h 17m
```

   e. **From the output, make note of this domain information:**

   - Each domain name and its CONS number (for example, `ssccn1-dom1 5001`, and so on).

- Identify the primary domain (for example, `primary`).
- Identify the end (service) domain. It is the domain with the highest number at the end of its name (for example, `ssccn1-dom3`).
- Identify the middle domains (all domains between the primary and end domain. For example, `ssccn1-dom1` and `ssccn1-dom2`).

f. Identify which domains are dedicated domains and which domains are I/O Domains.

In this example, replace *NAME* with a valid domain name from the **`ldm list`** command.

Syntax:

```
# ldm ls-variable NAME | grep oes_ldom_type
```

Examples:

```
# ldm ls-variable ssccn1-dom1 | grep oes_ldom_type
oes_ldom_type=dedicated

# ldm ls-variable ssccn1-dom2 | grep oes_ldom_type
oes_ldom_type=io
```

g. Obtain a list of running zones.

Perform this command in each domain to determine if the domain includes zones.

```
# zoneadm list
global
orlm6db01z2
orlm6db01z1
orlm6db01z3
orlm6db01_T
```

2. Repeat this procedure for all PDomains.

Related Information

- “Shut Down Zones” on page 30
- “Shutting Down SuperCluster Gracefully” on page 28

▼ Shut Down Zones

If your system is configured with zones, use this procedure to shut them down.

1. **For the domain that includes zones, login as root.**

2. **List running zones.**

   For example:

   ```
   # zoneadm list
   global
   ABC_zone
   ```

3. **Shut down a zone.**

   For example:

   ```
   # zoneadm -z ABC_zone shutdown
   ```

4. **Repeat this procedure to shut down all zones in all PDomains.**

Related Information

- "Shut Down the Guest Domains" on page 31
- "Shutting Down SuperCluster Gracefully" on page 28

▼ Shut Down the Guest Domains

Note - Domain configurations vary based on the configuration chosen during installation, and based on any I/O Domains created.

The order in which you shut down domains is important and varies according to the type and number of domains. Shut down domains in this order:

1. I/O Domains (if present)

2. Middle domains (dedicated guest domains that are neither the primary (control) domain nor the end (service) domain

3. End (service) domain (domain named `ssccnX-dom_`*highest_number*)

4. Primary domain (described in "Shut Down the Primary Domains" on page 33)

If the PDomain is running with only one domain, shut down the system just as you would any other server by cleanly shutting down the OS (see "Shut Down the Primary Domains" on page 33).

To identify the number and types of domains on your system, perform "Obtain Information Needed for Shutdown Tasks" on page 29.

1. **Login to the primary domain:**

   a. **Login to Oracle ILOM on the active SP for a PDomain.**

   b. **Start the host console.**

   In this example, replace *x* with the number of the host.

```
-> start /HOSTx/console
Are you sure you want to start /HOSTx/console [y/n] y
Serial console started. To stop, type #.
```

   c. **Login to the primary domain as root.**

```
sccn01 console login: root
Password:
Last login: Thur Feb 9 16:52:10 form /dev/console
Oracle Corporation      SunOS 511     11.2      Apr 2016
```

2. **List the running domains.**

```
# 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
```

3. **Stop all active I/O domains (if present).**

```
# ldm stop domain_name
LDom domain_name stopped
```

4. **Connect to a running dedicated guest domain (not the primary or end domain).**

This example uses `telnet` and the CONS port number obtained in "Obtain Information Needed for Shutdown Tasks" on page 29.

```
# telnet 0 CONS_Port_NO
Trying 0.0.0.0...
Connecting to 0.
Escape character is '^]'
Connecting to console "ssccn3-dom1" in group "ssccn3-dom1" . . . .
Press ~? for control options  ..
```

```
ssccn3-dom1 console login: root
Password:
Feb 14 04:34:52 ssccn3-dom1 login: ROOT LOGIN /dev/console
Last login: Wed Feb  5 07:44:41 on console
Oracle Corporation      SunOS 5.10      Generic Patch   January 2005
#
```

5. **Shut down the dedicated guest domain (not the primary or end domain).**

Note - The `init 0` command is the preferred command to shut down a domain rather than other commands such as `shutdown`.

```
# init 0
```

6. **Repeat Step 3 and Step 4 to shut down any additional middle guest domains.**

7. **Connect to and shut down the end (service) domain.**

```
# telnet 0 CONS_Port_NO
Trying 0.0.0.0...
Connecting to 0.
Escape character is '^]'
Connecting to console "ssccn3-dom1" in group "ssccn3-dom1" . . . .
Press ~? for control options  ..

ssccn3-dom1 console login: root
Password:
Feb 14 04:34:52 ssccn3-dom1 login: ROOT LOGIN /dev/console
Last login: Wed Feb  5 07:44:41 on console
Oracle Corporation      SunOS 5.10      Generic Patch   January 2005
# init 0
```

8. **Repeat the procedure to shut down guest domains on all PDomains.**

Related Information

- “Shut Down the Primary Domains” on page 33
- “Shutting Down SuperCluster Gracefully” on page 28

▼ Shut Down the Primary Domains

1. **Login to the primary domain:**

a. **Login to Oracle ILOM on the active SP for a PDomain.**

b. **Start the host console.**

In this example, replace *x* with the number of the host.

```
-> start /HOSTx/console
Are you sure you want to start /HOSTx/console [y/n] y
Serial console started. To stop, type #.
```

c. **Login to the primary domain as root.**

```
sccn01 console login: root
Password:
Last login: Thur Feb 9 16:52:10 form /dev/console
Oracle Corporation      SunOS 511     11.2      Apr 2016
```

2. **Shut down the primary domain using the `init` command.**

```
# init 0
```

3. **Repeat this procedure for all PDomains.**

4. **Consider your next action.**

Once all the primary domains are shut down, SuperCluster is fully shut down. Depending on the purpose for shutting the system down, you might want to power off the system using these procedures:

- “(If Needed) Shut Down the Storage Servers” on page 34
- “(If Needed) Power Off the ZFS Storage Appliance” on page 35
- “(If Needed) Power Off the Racks” on page 35

▼ (If Needed) Shut Down the Storage Servers

1. **Login to a storage server.**

2. **Shut the storage server down.**

```
# shutdown -h now
```

3. **Repeat this procedure for each storage server.**

(If Needed) Power Off the ZFS Storage Appliance

- **Gracefully shut down each ZFS storage appliance controllers using one of these methods:**
  - Enter the appliance hostname or IP address into a browser address field, log into the BUI, then click the power icon.
  - SSH into the appliance and enter the **`maintenance system poweroff`** command.
  - Use a pen or non-conducting pointed object to press and release the Power button on the front panel.

Related Information

- “(If Needed) Power Off the Racks” on page 35
- “Shutting Down SuperCluster Gracefully” on page 28

(If Needed) Power Off the Racks

- **Power off the remaining components by turning off the rack circuit breakers.**

Note - There are two circuit breakers in each rack.

Related Information

- “Power Off SuperCluster in an Emergency” on page 35
- “Shutting Down SuperCluster Gracefully” on page 28

Power Off SuperCluster in an Emergency

If there is an emergency, such as earthquake or flood, an abnormal smell or smoke, or a threat to human safety, you must power off SuperCluster immediately.

- **Use one of these methods to power off SuperCluster M8 or SuperCluster M7 in an emergency.**
  - **Turn off power at the circuit breaker, or pull the emergency power-off switch in the computer room.**

- **Turn off the site EPO switch.**
- **Turn off the two PDUs in the rack.**

After the emergency, contact Oracle Support Services to restore power to the system.

Related Information

- "Cautions" on page 27
- "Power On SuperCluster" on page 27
- "Shutting Down SuperCluster Gracefully" on page 28

Configuring Oracle Engineered Systems Hardware Manager

Oracle Engineered Systems Hardware Manager is a BUI-based rack-level hardware management tool intended for use by Oracle Service personnel.

You must configure and keep Oracle Engineered Systems Hardware Manager up-to-date to ensure that Oracle Service personnel can use the tool to manage SuperCluster components.

These topics are covered in this section:

Related Information

- "Monitoring SuperCluster Systems (OCM)" on page 43
- "Tuning SuperCluster Systems (`ssctuner`)" on page 65

Engineered Systems Hardware Manager Overview

Oracle Engineered Systems Hardware Manager is a BUI-based rack-level hardware management tool used by Oracle Service personnel. The tool provides service personnel with these capabilities:

- Hardware dashboard/health view
- Hardware inventory summary
- Hardware component details
- Ability to launch other Oracle SuperCluster tools
- Easy access to Oracle ILOM and other hardware administration interfaces

- Problem summary across all components
- Rack level problem history, and the ability to manually clear hardware faults and warnings
- Automatic and manual collection of support file bundles
- Manual delivery of support file bundles to My Oracle Support (MOS)

You must maintain two accounts on Oracle Engineered Systems Hardware Manager.

- **`admin`** – Use this account to configure Oracle Engineered Systems Hardware Manager.
- **`service`** – An account used by Oracle Service personnel to manage SuperCluster M8 and SuperCluster M7 systems.

Oracle Engineered Systems Hardware Manager must be configured with the passwords for these component accounts:

Component	User Account
All Oracle ILOMs	`root`
Exadata storage servers OS	`root`
ZFS storage controllers OS	`root`
IB switches	`root`
Ethernet management switch	`admin`
PDUs	`admin`

Note - Oracle Engineered Systems Hardware Manager only requires passwords for components listed in the table. The tool does not need to know the passwords for any logical domains or zones.

Related Information

- “Configure Oracle Engineered Systems Hardware Manager” on page 38
- “Update Component Passwords in Oracle Engineered Systems Hardware Manager” on page 40

▼ Configure Oracle Engineered Systems Hardware Manager

Perform this procedure to prepare Oracle Engineered Systems Hardware Manager for use by Oracle Service personnel, or any time you want to change the tool's passwords.

1. **Obtain a copy of your installation summary.**

2. **In a browser, enter this address:**

   `https://`*address_of_master_control_domain*`:8001`

   To find the address of your master control domain, refer to your installation summary.

   The Oracle Engineered Systems Hardware Manager login screen is displayed.

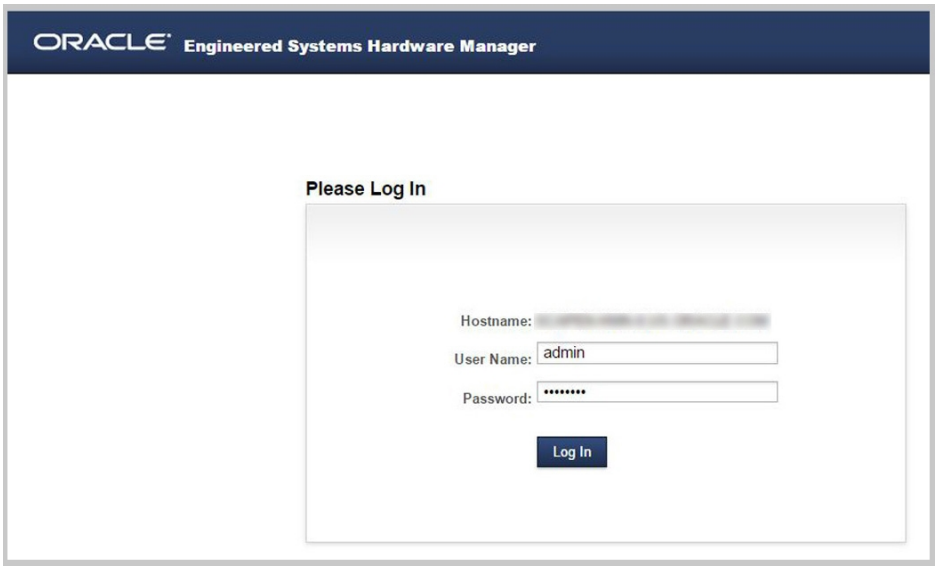

3. **Log in to the `admin` account.**

   The user name is `admin`. the password is provided in your installation summary.

   Tip - To access online help, click the More details link in the upper right corner of the page.

4. **Configure the password policy for these accounts so that it matches your corporate policy.**

   Refer to Setup > Authentication > Password Policy.

5. **Change the default passwords for the admin and service accounts.**

   Refer to Setup > Authentication > User Management.

6. **Update component passwords.**

Refer to Setup > Authentication > Component Access.

This step is mandatory so that Oracle Service personnel can use the tool to manage SuperCluster components. For details on which component passwords are required see "Engineered Systems Hardware Manager Overview" on page 37.

7. **(Optional) Configure Oracle Engineered Systems Manager to use your own certificates instead of the site- and instance-specific certificates that the tool generates.**

   Refer to Setup > Certificates.

8. **(Optional) Change network ports if they conflict with your environment.**

   For example, if an application running on Oracle SuperCluster uses the same port that Oracle Engineered Systems Hardware Manager uses (ports 8001 to 8004), you or Oracle Service can configure Oracle Engineered Systems Hardware Manager to use a different port.

   Refer to Setup > Network.

Related Information

- "Engineered Systems Hardware Manager Overview" on page 37
- "Update Component Passwords in Oracle Engineered Systems Hardware Manager" on page 40

▼ Update Component Passwords in Oracle Engineered Systems Hardware Manager

You must perform this procedure whenever a component password is changed on the component. Keeping Oracle Engineered Systems Hardware Manager up to date ensures that Oracle Service personnel can use the tool to manage SuperCluster components.

For details on which component passwords are required see "Engineered Systems Hardware Manager Overview" on page 37.

1. **Access Oracle Engineered Systems Hardware Manager.**

   In a browser, enter this address:

   `https://`*IP_address_of_master_control_domain*`:8001`

   To find the IP address of your master control domain, refer to your installation summary.

   The Oracle Engineered Systems Hardware Manager login screen is displayed.

2. **Log in to the `admin` account.**

   The user name is `admin`. the password was set in "Configure Oracle Engineered Systems Hardware Manager" on page 38.

3. **Update component passwords.**

   Refer to Setup > Authentication > Component Access.

Related Information

- "Engineered Systems Hardware Manager Overview" on page 37
- "Configure Oracle Engineered Systems Hardware Manager" on page 38

Monitoring SuperCluster Systems (OCM)

These topics describe how to monitor SuperCluster M8 and SuperCluster M7 with Oracle Configuration Manager.

Related Information

- "Tuning SuperCluster Systems (`ssctuner`)" on page 65
- "Configuring CPU and Memory Resources (`osc-setcoremem`)" on page 77

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 44:

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:**

  https://www.oracle.com/technical-resources/documentation/ocm.html

Monitoring SuperCluster Systems (ASR)

These topics describe how to configure ASR and use it to monitor SuperCluster M8 and SuperCluster M7.

ASR Overview

Auto Service Request (ASR) is designed to automatically open service requests when specific SuperCluster hardware faults occur. To enable this feature, the SuperCluster components must be configured to send hardware fault telemetry to the ASR Manager software. The ASR Manager must be installed on a server that has connectivity to the SuperCluster, and an outbound Internet connection using HTTPS or an HTTPS proxy.

When a hardware problem is detected, the ASR Manager submits a service request to Oracle Support Services. In many cases, Oracle Support Services can begin work on resolving the issue before the database/system administrator is even aware the problem exists.

Prior to using ASR, set up the following:

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

An e-mail message is sent to both the My Oracle Support (MOS) e-mail account for Auto Service Request and the technical contact for the activated asset, notifying them of the creation of the service request.

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

Consider the following information when using ASR:

- ASR is applicable only for component faults. Not all component failures are covered, though the most common components, such as disk, fan, and power supplies are covered. For more information, refer to this web page: `https://www.oracle.com/support/premier/auto-service-request.html`

  Click the `Documentation` link on this page, then refer to the "ASR Fault Coverage Information" section at the bottom of the page.
- ASR is not a replacement for other monitoring mechanisms, such as SMTP, and SNMP alerts, within the customer data center. It 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 may not be automatically filed. This can happen because of the unreliable nature of the SNMP protocol, or loss of connectivity to the ASR Manager. Oracle recommends that customers continue to monitor their systems for faults and call Oracle Support Services if they do not receive notice that a service request has been automatically filed.

Tip - Refer to the Oracle Auto Service Request web page at `https://www.oracle.com/support/premier/auto-service-request.html` for more information on ASR.

Preparing to Configure ASR

Confirm your environment is supported and prepared before installing and configuring ASR on the SuperCluster M8 and SuperCluster M7:

- "Prepare the ASR Environment" on page 47
- "SAR Manager Software Requirements" on page 47
- "ASR Software Requirements" on page 48

▼ Prepare the ASR Environment

Before You Begin Before installing ASR, ensure the following conditions are met:

1. **Create a My Oracle Support (MOS) account at `http://support.oracle.com`.**

   Ensure the following are correctly set up:

   - Oracle Premier Support for Systems or Oracle/Sun Limited Warranty
   - Technical contact responsible for SuperCluster M7
   - Valid shipping address for SuperCluster M7 parts

2. **Identify and designate a system to serve as ASR Manager.**

   For more information, see:

   `https://www.oracle.com/support/premier/auto-service-request.html`

   Click `additional details`, then click `Hardware and Network Configuration Recommendations`.

3. **Identify and verify ASR assets.**

4. **Ensure connectivity to the Internet using HTTPS.**

   You may 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`

   Click on the `Oracle ASR user documentation` link.

5. **Provide the necessary information in the following documents:**

   - *SuperCluster M8 and SuperCluster M7 Site Checklists*
   - *SuperCluster M8 and SuperCluster M7 Configuration Worksheets*

SAR Manager Software Requirements

You need root access to install the software to set up the ASR Manager.

- ASR Manager, version 3.6 or higher
- Oracle Services Tool Bundle (STB) for Oracle Solaris only

ASR Software Requirements

You need root access to install the software to set up the ASR Manager.

- Operating System: Oracle Linux 5.3 and later, or Oracle Solaris 10 Update 10 (10u10) plus patches, and later
- Java Version: at least JRE/JDK 6.2.28
- Database Server: Exadata Software 11.2.3.1 and higher
- Storage Server: Exadata Storage Server Software 11.2.0.3 DB with Bundle Patch 1, and higher
- Sun ZFS Storage 7320 Storage Controller: Firmware version AK2011.1.0 and higher

Installing ASR Manager Components

Perform these procedures to install the ASR Manager components on the external system designated as the ASR Manager. You may use a pre-existing ASR Manager, as long as it conforms to the requirements listed in "Prepare the ASR Environment" on page 47.

▼ Install the OASM Package

1. **Verify that you have version 1.3.1 or later (if needed, download OASM).**

   As root:

   - Oracle Solaris: `pkginfo -l SUNWsasm`
   - Oracle Linux: `rpm -q SUNWsasm`

2. **Install the OASM package.**

   As root:

   - Oracle Solaris: `pkgadd -d SUNWsasm.`*version-number*`.pkg`
   - Oracle Linux: `rpm -i SUNWsasm.`*version-number*`.rpm`

Install STB (Oracle Solaris Only)

1. **If needed, download Services Tools Bundle from:**

   https://www.oracle.com/support/premier/auto-service-request.html

   and click on the `Download` link.

2. **Untar the STB bundle and run the installation script (`install_stb.sh`).**

   As part of the installation, select:

   - Type I for "install"
   - Type Y to replace existing SNEEP packages
   - Type Y to replace existing Service Tags packages

   Note - See Doc ID 1153444.1 to download the latest STB bundle from My Oracle Support (log in required): `https://support.oracle.com`.

3. **Confirm that SNEEP is installed correctly:**

   ```
   sneep -a
   ```

4. **Verify that Service Tags is reporting your system attributes correctly:**

   ```
   stclient -E
   ```

   If the serial number does not display, then register the serial number manually:

   `sneep -s` *serial-number*

Install the Oracle ASR Package

1. **Download and unzip the ASR package.**

   As root:

   - Oracle Solaris: `pkgadd -d SUNWswasr.`*version-number*`.pkg`
   - Oracle Linux: `rpm -i SUNWswasr.`*version-number*`.rpm`

2. **Add the `asr` command to the `PATH` (update to the root's `.profile`, `.cshrc`, `.kshrc` or `.bashrc` as needed):**

   ```
   PATH=$PATH:/opt/SUNWswasr/bin
   ```

```
export PATH
```

▼ Register ASR Manager

Before You Begin When registering ASR Manager, type your MOS single sign on information and any proxy servers, if needed.

1. **As root on the ASR Manager system, type:**

   ```
   # asr register
   ```

2. **Type "1" or "alternate URL for Managed OPS use only":**

   ```
   1) transport.oracle.com
   ```

3. **If you are using a proxy server to access the Internet, type the proxy server information.**

   For example:

   ```
   Proxy server name: ? <proxy server name>
   Proxy port number: ? <proxy port number>
   Proxy authentication; if authentication is not required, enter -.
   Proxy user: <proxy user name>
   Proxy password: <proxy password>
   If this is an NTLM type proxy, enter the information below.
   Otherwise, enter -
   NTLM Domain: [?] <NTLM domain name>
   Enter the host the NTLM authentication request is originating
   from. Usually this is the hostname of the SASM server.
   NTLM Host: [?] <NTLM host name>
   NTLM support is not bundled with SASM but may be added now.

   1) Download jCIFS from http://jcifs.samba.org/
   2) Extract contents and locate the jcifs-*.jar file
   3) Enter full path to this file below

   jCIFS jar file: [?] <full path of jCIFS jar file>
   Note: These properties are stored in the
   /var/opt/SUNWsasm/configuration/config.ini file. You can update
   these properties if needed and then restart SASM.
   ```

4. **When prompted, type your My Oracle Support (MOS) user name and password. ASR validates the login. Once validated, the registration is complete. Note: Passwords are not stored.**

Your MOS e-mail address receives output from notification of ASR problems and Service Request (SR) generation.

Verify ASR Manager

1. **On the ASR Manager, verify that you have the correct version of ASR Manager:**

   ```
   # asr show_rules_version
   ```

   You should see that the version is 3.6 or later.

2. **Check the registration status:**

   ```
   # asr show_reg_status
   ```

3. **Test the connection to ensure that ASR can send information to the transport server:**

   ```
   # asr test_connection
   ```

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.

Complete the following steps 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)”
   ```

   Note that single quotes are required around the *ASR-Manager-name-or-IP-address* entry. Following 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 enable 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 – Shows 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 customer network requirements.

- port=162 – The SNMP port. This port value is customer dependant. It can be configured as a different port based on network requirements, or it may need to be changed for ASR to work correctly in a managed environment.

3. **Validate if Oracle ILOM auto-activation occurred (if the network and Oracle ILOM are set up correctly):**

   # asr list_asset

   For 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. **Validate 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 the following example output:

```
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 system.**

10. **When you have completed these instructions for every storage server in your SuperCluster system, approve and verify contacts to the storage servers on MOS.**

    See "Approve and Verify ASR Activation for SuperCluster systems Assets" on page 61 for those instructions.

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

▼ Configure ASR on the Storage Appliance

To activate the storage appliance included in SuperCluster M8 and SuperCluster M7, complete these steps on each storage controller:

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

   `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, as shown in the following figure.**

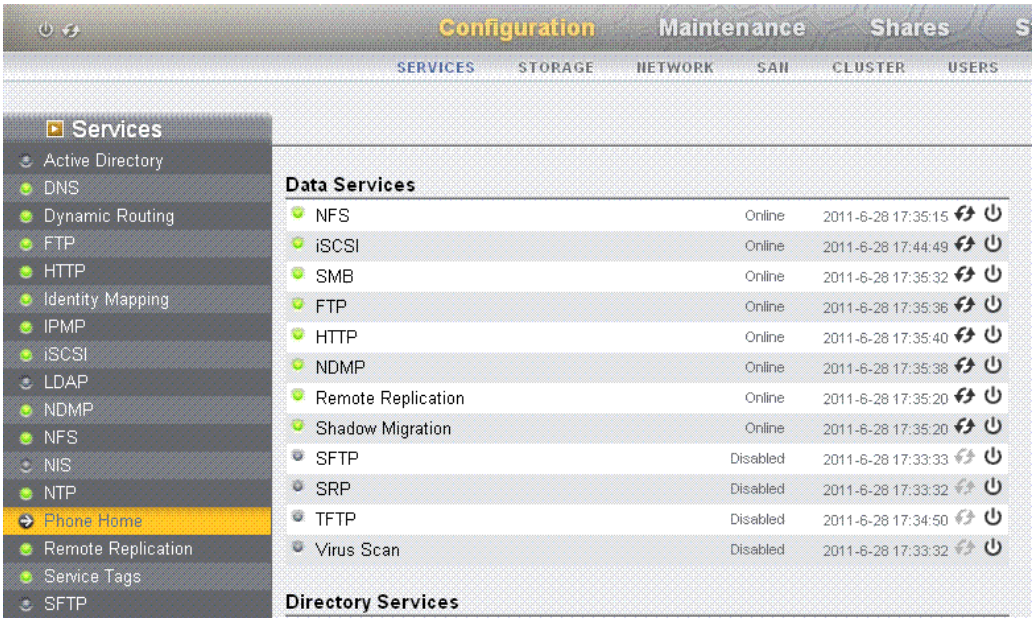

When you click Phone Home, the Phone Home Properties page is displayed, as shown in the following figure.

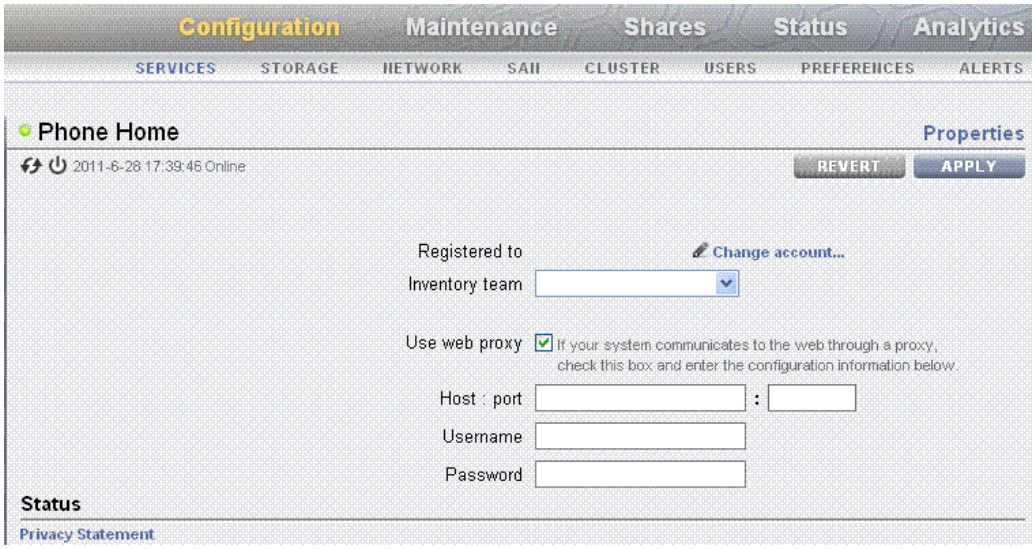

5. **If you are using a web proxy to connect to the Internet from the storage appliance, select the Use web proxy option, and type the following information:**
   - In the *Host:port* field, type the complete host name of your web proxy server and the port.
   - In the *Username* field, type your user 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 pop-up is presented, select the Enable option.**
10. **Repeat these instructions for every storage controller in your SuperCluster system.**

11. **When you have completed these instructions for every storage controller in your SuperCluster M7, approve and verify contacts to the Sun ZFS Storage appliances on MOS.**
    See “Approve and Verify ASR Activation for SuperCluster systems Assets” on page 61 for those instructions.
    For more information on the process, see ASR MOS 5.3+ Activation Process (Doc ID 1329200.1).

▼ 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.

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

1. **Log in to the base server Oracle ILOM.**

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 the following command to determine if that rule is currently being used:**
   # show /SP/alertmgmt/rules/*rule-number*
   For example:
   # show /SP/alertmgmt/rules/1

   - If you see output similar to the following:
     ```
     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 would be the IP address of the ASR Manager in this case). If you see output similar to the preceding example, pick another rule and type the `show /SP/alertmgmt/rules/`*rule-number* command again, this time using another rule in the list.

- If you see output similar to the following:

```
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* **snmp_version=2c community_or_username=public**

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

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

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

7. **Repeat these instructions on Oracle ILOM for all base servers in your SuperCluster systems.**

8. **When you have completed these instructions for all base servers in SuperCluster systems, approve and verify contacts to the base servers on MOS.**

See “Approve and Verify ASR Activation for SuperCluster systems Assets” on page 61 for those instructions.

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

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 use the `asr enable_http_receiver` command on the ASR Manager. Select a port for the HTTP receiver that is appropriate for your network environment and does not conflict with other network services.

Perform the following tasks:

- “Enable the HTTP Receiver on the ASR Manager” on page 58
- “Register Compute Servers With Oracle Solaris 11 or Database Domains to ASR Manager” on page 59

▼ Enable the HTTP Receiver on the ASR Manager

Follow 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 `root`.**
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.

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

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

1. **Log in to the base server as `root`.**

2. **Confirm that the `asr-notify` service is working:**

   ```
   # svcs asr-notify
   ```

   - If you see the following 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 the following 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 the following message, the `asr-notify` service is installed and is working properly:

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

3. **To register the ASR manager, run:**

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

   For example:

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

   You should see screens asking for your Oracle Support account name and password. After entering your Oracle Support account name and password, you should see a notification, saying that your registration is complete:

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

Registration complete.
```

4. **Run the following command:**

   # asradm list

   The screen output should be similar to the following:

```
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
```

   Upon successful results of the above commands, the registration of the ASR Manager is complete.

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

6. **After you have completed these instructions for both base servers in your SuperCluster system, approve and verify contacts to the base servers on MOS.**

   See “Approve and Verify ASR Activation for SuperCluster systems Assets” on page 61 for those instructions.

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

Approve and Verify ASR Activation for SuperCluster systems Assets

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

   `list_asset`

   This command lists ASR assets in your SuperCluster system, including compute servers, storage servers, and 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.

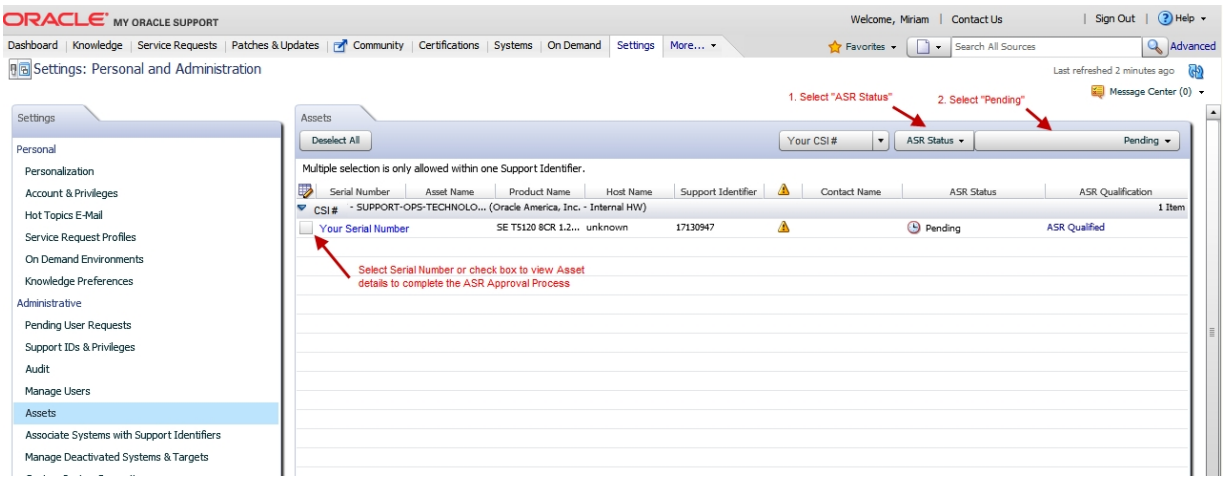

Note - By default, all support identifiers that you are associated with are displayed. If this list of assets is long, you can limit the display to show only assets associated to one support identifier. You can also search for an asset's serial number.

Note - For each SuperCluster component, 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 help for ASR. 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 any missing information about the asset is required, the information pop-up indicates the needed information. The ASR Activation window appears and look like the following figure.

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@example.com, asr-notifications-2@example.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 in order for Service Request autocreate to work.

---

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

   # **`asradm send test`** *email-address*@*example.com*

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

Tuning SuperCluster Systems (ssctuner)

These topics describe the utility (ssctuner) used to tune SuperCluster M8 and SuperCluster M7. For the latest information about ssctuner, refer to the README file installed with the utility.

Related Information

- For more information about SMF services on the Oracle Solaris OS, refer to 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 when SuperCluster is installed.

The utility runs in real time 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 69.

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

There is an ssctuner SMF variable called ssctuner_vars/COMPLIANCE_RUN that you set to an appropriate benchmark and then restart ssctuner to configure a compliance assessment. By default, this variable is set to none. For security purposes, you must enable this feature, see "Configure ssctuner to Run compliance(1M) Benchmarks" on page 71.

Related Information

- "Monitor ssctuner Activity" on page 67
- "View Log Files" on page 67
- "Configure the EMAIL_ADDRESS Property" on page 68
- "Change ssctuner Properties and Disable Features" on page 69
- "Configure ssctuner to Run compliance(1M) Benchmarks" on page 71
- "Monitor and View the Compliance Benchmark" on page 72
- "Install ssctuner" on page 74
- "Enable ssctuner" on page 75

Monitor ssctuner Activity

- **View ssctuner activity.**

  ```
  # svcs -l ssctuner
  ```

  ### Related Information

  - “ssctuner Overview” on page 65
  - “View Log Files” on page 67
  - “Configure the EMAIL_ADDRESS Property” on page 68
  - “Change ssctuner Properties and Disable Features” on page 69
  - “Configure ssctuner to Run compliance(1M) Benchmarks” on page 71
  - “Monitor and View the Compliance Benchmark” on page 72
  - “Install ssctuner” on page 74
  - “Enable ssctuner” on page 75

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(l)
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 65
- “Monitor ssctuner Activity” on page 67
- “Configure the EMAIL_ADDRESS Property” on page 68
- “Change ssctuner Properties and Disable Features” on page 69
- “Configure ssctuner to Run compliance(1M) Benchmarks” on page 71
- “Monitor and View the Compliance Benchmark” on page 72
- “Install ssctuner” on page 74
- “Enable ssctuner” on page 75

▼ Configure the `EMAIL_ADDRESS` Property

You must configure the `EMAIL_ADDRESS` property so that `ssctuner` messages are emailed to the appropriate person, even when the message is not logged into the system.

1. **Configure the `ssctuner` so that critical messages are sent to your email address.**

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

2. **If you plan to change any other ssctuner properties, do so before you perform the remaining steps in this task.**

   See “Change ssctuner Properties and Disable Features” on page 69.

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 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 65
- "Monitor `ssctuner` Activity" on page 67
- "View Log Files" on page 67
- "Configure `ssctuner` to Run `compliance(1M)` Benchmarks" on page 71
- "Monitor and View the Compliance Benchmark" on page 72
- "Install `ssctuner`" on page 74
- "Enable `ssctuner`" on page 75

▼ 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 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:

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

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

   - **For cases when you run Oracle Databases in Application Domains, setting this property to `true` enables `ssctuner` to change thread priorities as it does in Database Domains. By default, the value is `false`.**

     ```
     ~# 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 that there are 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 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 65
- "Monitor ssctuner Activity" on page 67
- "View Log Files" on page 67
- "Configure the EMAIL_ADDRESS Property" on page 68
- "Configure ssctuner to Run compliance(1M) Benchmarks" on page 71
- "Monitor and View the Compliance Benchmark" on page 72
- "Install ssctuner" on page 74
- "Enable ssctuner" on page 75

▼ Configure ssctuner to Run compliance(1M) Benchmarks

Use this procedure to configure ssctuner to run compliance benchmarks.

The assessment begins within 12 minutes and then to reruns after each time the node is rebooted.

By default, this variable is set to `none`., but you must enable this feature.

1. **Identify available benchmarks.**

   In this example, two benchmarks are available, `pci-dss` and `solaris`.

   ```
   # compliance list -b
   pci-dss solaris
   ```

2. **Set the ssctuner SMF variable to the chosen benchmark.**

   This example uses the `solaris` benchmark, which runs the recommended profile.

   ```
   # svccfg -s ssctuner setprop ssctuner_vars/COMPLIANCE_RUN=solaris
   # svcadm restart ssctuner
   ```

3. **Verify that the compliance run is scheduled by viewing the SMF log file.**

   Note - Compliance runs are staggered to prevent DOS attacks on the ZFS storage appliance.

   ```
   # grep compliance /var/svc/log/site-application-sysadmin-ssctuner\:default.log
   [ Nov 16 11:47:54 notice: Performing compliance run after delay of 519 seconds... ]
   ```

▼ Monitor and View the Compliance Benchmark

1. **(Optional) View the SMF log as the benchmark runs:**

   ```
   # tail -f /var/svc/log/site-application-sysadmin-ssctuner\:default.log
   root@etc28zadm0101:~# tail -f /var/svc/log/site-application-sysadmin-ssctuner
   \:default.log
   [ Nov 16 11:47:55 CURRENT STATUS version=1.3.8  crit issue count=1  disabled feature
    count=0 ]
   [ Nov 16 11:47:55 CURRENT ISSUES : Please change ssctuner email address from
    root@localhost ]
   [ Nov 16 11:47:55 notice: Checking Oracle log writer and LMS thread priority. ]
   [ Nov 16 11:47:56 notice: completed initialization. ]
   [ Nov 16 11:47:56 Method "start" exited with status 0.]
   [ Nov 16 11:49:55 notice: Checking Oracle log writer and LMS thread priority. ]
   [ Nov 16 11:51:55 notice: Checking Oracle log writer and LMS thread priority. ]
   [ Nov 16 11:53:55 notice: Checking Oracle log writer and LMS thread priority. ]
   [ Nov 16 11:55:55 notice: Checking Oracle log writer and LMS thread priority. ]
     Assessment will be named 'solaris.Baseline.2015-11-16,11:56'
         Package integrity is verified
         OSC-54005
         pass
   ```

```
The OS version is current
OSC-53005
pass

Package signature checking is globally activated
OSC-53505
pass
```

2. **(Optional) Determine if the assessment is complete.**

   When you see `Compliance assessment completed`, either from Step 1 or by using this `grep` command, continue to the next step.

```
# grep -i compliance /var/svc/log/site-application-sysadmin-ssctuner\:default.log
        [ Nov 16 11:47:54 notice: Performing compliance run after delay of 519
 seconds... ]
        [ Nov 16 11:57:47 notice: Compliance assessment completed.]
```

3. **List assessments.**

```
# compliance list -a
solaris.Baseline.2015-11-16,11:56
```

4. **Obtain an assessment html report.**

Note - New assessments run (staggered in time) after each time the domain reboots.

```
# compliance report -a solaris.Baseline.2015-11-16,11:56
/var/share/compliance/assessments/solaris.Baseline.2015-11-16,11:56/report.html
```

Related Information

- "`ssctuner` Overview" on page 65
- "Monitor `ssctuner` Activity" on page 67
- "View Log Files" on page 67
- "Configure the `EMAIL_ADDRESS` Property" on page 68
- "Change `ssctuner` Properties and Disable Features" on page 69
- "Install `ssctuner`" on page 74

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 using the Oracle Solaris 11 pkg command.**

   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 67.

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 65
- "Monitor ssctuner Activity" on page 67
- "View Log Files" on page 67
- "Configure the EMAIL_ADDRESS Property" on page 68
- "Change ssctuner Properties and Disable Features" on page 69
- "Configure ssctuner to Run compliance(1M) Benchmarks" on page 71

- "Monitor and View the Compliance Benchmark" on page 72
- "Enable ssctuner" on page 75

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 67.

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

See "View Log Files" on page 67.

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 65
- "Monitor ssctuner Activity" on page 67
- "View Log Files" on page 67
- "Configure the EMAIL_ADDRESS Property" on page 68
- "Change ssctuner Properties and Disable Features" on page 69
- "Configure ssctuner to Run compliance(1M) Benchmarks" on page 71
- "Monitor and View the Compliance Benchmark" on page 72
- "Install ssctuner" on page 74

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

These sections describe how to configure Oracle SuperCluster CPU and memory resources using the `osc-setcoremem` command. You can also use the `osc-setcoremem` simulator to run simulations to determine the best way to configure CPU and memory resources.

Use these topics to change CPU and memory allocations for I/O Domains.

Description	Links
Learn about the CPU/Memory tool.	"`osc-setcoremem` Overview" on page 78 "Minimum and Maximum Resources (Dedicated Domains)" on page 80
Determine if SuperCluster resources can be modified using the CPU/Memory tool.	"Supported Domain Configurations" on page 81
Plan CPU and memory allocations.	"Plan CPU and Memory Allocations" on page 83
Identify domain configurations.	"Display the Current Domain Configuration (`osc-setcoremem`) Command" on page 86 "Display the Current Domain Configuration (`ldm`)" on page 87
Display command line help for the non-interactive options.	"Display `osc-setcoremem` Help" on page 88
Configure domain CPU and memory resources at the socket or core level.	"Guidelines for Modifying the Base Configuration File" on page 117 "Change CPU/Memory Allocations (Socket Granularity)" on page 89 "Change CPU/Memory Allocations (Core Granularity)" on page 94 "Change CPU/Memory Allocations Non-Interactively" on page 98
Configure domain CPU and memory resources so that some resources are parked.	"Park Cores and Memory" on page 102
View the SP configuration, revert to a previous configuration, or remove a configuration.	"View the SP Configuration" on page 107

Description	Links
	"Revert to a Previous CPU/Memory Configuration" on page 109 "Remove a CPU/Memory Configuration" on page 110
Reset the CPU and memory configuration to the factory default configuration.	"Reset the CPU/Memory Configuration" on page 111
Access information about previous executions of `osc-setcoremem`.	"Access `osc-setcoremem` Log Files" on page 113
Run simulations to test different configurations.	"Run a Simulation" on page 119 "Example: Simulating Changes on a Non-SuperCluster Node" on page 122

`osc-setcoremem` Overview

SuperCluster compute server CPU and memory resources are initially allocated during installation. This base configuration of CPU and memory resources are assigned to I/O Domains in the same proportion as IB HCAs. Memory is assigned in the same proportions. If your computing needs change, you can modify the base configuration later.

You can use the `osc-setcoremem` command to migrate CPU cores and memory resources between dedicated domains, and from dedicated domains to CPU and memory repositories for the use of I/O domains.

Starting with the 2.5 release, `osc-setcoremem` makes an attempt to eliminate domain reboots whenever possible. The assessment is transparent. However, if the tool determines that any domain must be restarted as part of a domain reconfiguration session, those domains are highlighted.

You can save time by using the `osc-setcoremem` simulator to view different possible configurations before you make the changes. The `osc-setcoremem` simulator uses the same production binary as the `osc-setcoremem` command. The `osc-setcoremem` simulator disregards the physical hardware it is running on as long as the target system contains SPARC processors that run the Oracle Solaris OS 11 or later. For instructions on using the simulator, see "Run a Simulation" on page 119.

Use the simulator to test these situations:

- Move from a base configuration to a desired configuration.
- Move from your current custom configuration to a desired configuration.

Consider this information when you use the `osc-setcoremem` command:

- 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 80.
- Affected dedicated domains must be rebooted after any change.

The `osc-setcoremem` command enables you to change the CPU and memory allocations at these levels of granularity:

- **Socket granularity** – The command 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 89.
- **Core granularity** – The command 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 94.

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 102.

Prior to SuperCluster software version 3.0, if you parked resources from dedicated domains, you could not move parked resources to back to dedicated domains after I/O Domains were created. As of version 3.0, you can move parked unallocated resources back to dedicated domains.

The `osc-setcoremem` command optimizes for locality. For example, in SuperCluster M7, each socket is populated with 32 cores and 512 GB memory. These resources form a locality group, because IO devices, cores, and memory are all local to the same system board. To minimize the impact of NUMA, `osc-setcoremem` assigns cores and memory from the same socket or locality group. In doing so, the tool enforces a minimum of one memory granule (16 GB memory) for every 4 cores assigned (rounded up) from the same socket. For instance, if a domain was assigned 14 cores from a socket, a minimum of 4 memory granules (14/4=3.5 rounded up to 4) or (4 * 16 GB) = 64 GB memory is required to be allocated from the same socket the 14 cores were allocated from.

When 32 or fewer cores are assigned to a dedicated domain, to optimize performance, those cores are all assigned from a single locality group. Because `osc-setcoremem` ensures that resources assigned from any locality group must include both cores and memory, memory is assigned from the same locality group as the cores, and is therefore limited to a maximum of 512 GB memory (or less than 512 GB, if fewer than 32 cores are assigned, because some memory must accompany the cores that are assigned to a different domain).

Dedicated domain memory is not limited to 512 GB, though. More than 512 GB of memory can be assigned to a dedicated domain by adding more than 32 cores to that domain. Because the cores, out of necessity, span multiple locality groups, memory can be drawn from multiple locality groups as well. So a domain with 64 cores can include up to 1024 GB of memory, for example, and a domain with 96 cores up to 1536 GB of memory.

Also see “Supported Domain Configurations” on page 81.

Related Information

- “Planning CPU and Memory Changes” on page 80

Planning CPU and Memory Changes

These topics describe how minimum and maximum resources are determined, which domain configurations are supported, and how to change your CPU and memory allocations.

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 M8 and SuperCluster M7.

Configuration	Minimum Resource Requirements
Dedicated Domain with 1 HCA	2 cores / 32GB memory
Dedicated Domain with 2 HCAs	4 cores / 64GB memory
Dedicated Domain with 4 HCAs	8 cores / 128GB memory

The minimum amount of CPU resources 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

- "Supported Domain Configurations" on page 81
- "Plan CPU and Memory Allocations" on page 83
- "Display the Current Domain Configuration (`osc-setcoremem`) Command" on page 86

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 *Understanding the Software Configurations* in the *Oracle SuperCluster M8 and SuperCluster M7 Overview Guide*.

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 83

Domain Configuration	Supported Resource Allocation Activities	Links
	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 89 "Change CPU/Memory Allocations (Core Granularity)" on page 94
	Remove (park) resources from dedicated domains for licensing purposes. **Note -** Parked resources are not available for use by any domains.	"Park Cores and Memory" on page 102
	Revert to a previous resource configuration.	"Revert to a Previous CPU/Memory Configuration" on page 109
	Remove a CPU or memory configuration.	"Remove a CPU/Memory Configuration" on page 110
Mixed domains – some are dedicated, some are Root Domains	(Only applicable to SuperCluster software version 2.6 or earlier) 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 83
	■ 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 89 "Change CPU/Memory Allocations (Core Granularity)" on page 94
	■ Revert to a previous allocation configuration.	"Revert to a Previous CPU/Memory Configuration" on page 109
	Activities you can perform anytime:	
	■ Configure resources for I/O Domains.	Refer to the *Oracle 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 102
	■ Move resources between dedicated domains.	"Change CPU/Memory Allocations (Socket Granularity)" on page 89 "Change CPU/Memory Allocations (Core Granularity)" on page 94
	■ Remove a CPU/memory configuration.	"Remove a CPU/Memory Configuration" on page 110

Related Information

- "Plan CPU and Memory Allocations" on page 83
- "Display the Current Domain Configuration (`osc-setcoremem`) Command" on page 86
- "Change CPU/Memory Allocations (Socket Granularity)" on page 89

- “Display `osc-setcoremem` Help” on page 88

▼ Plan CPU and Memory Allocations

There are two main approaches to modifying resource allocations:

- **All resources allocated** – Move resources from domains to other domains, and ensure that all resources are allocated.
- **Some resources are unallocated** – 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 102.

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

	SuperCluster M8	SuperCluster M7	SuperCluster M6-32	SuperCluster T5-8
1 socket =	32 cores (256 VCPUs)	32 cores (256 VCPUs)	12 cores (96 VCPUs)	16 cores (128 VCPUs)
1 core =	8 VCPUs	8 VCPUs	8 VCPUs	8 VCPUs

1. Identify the resource configuration for each compute node.

Follow the steps in one of these procedures:

- “Display the Current Domain Configuration (`osc-setcoremem`) Command” on page 86
- “Display the Current Domain Configuration (`ldm`)” on page 87

Note - Examples are based on SuperCluster M6-32, however, the same concepts apply to SuperCluster M7 and SuperCluster M8.

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

Domain	Domain Type	Cores	Memory (GB)
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 89.
   - Change resource allocations at the core granularity level. See "Change CPU/Memory Allocations (Core Granularity)" on page 94.
   - Increase unallocated resources. See "Park Cores and Memory" on page 102.

Related Information

- "Display the Current Domain Configuration (osc-setcoremem) Command" on page 86
- "Display the Current Domain Configuration (ldm)" on page 87
- "Change CPU/Memory Allocations (Socket Granularity)" on page 89
- "Change CPU/Memory Allocations (Core Granularity)" on page 94
- "Run a Simulation" on page 119

Using the osc-setcoremem Command

You can use the osc-setcoremem command to migrate CPU cores and memory resources between dedicated domains, and from dedicated domains to CPU and memory repositories for the use of I/O domains. Use the osc-setcoremem command to display and change CPU and memory allocations. You can also revert or reset to a previous configuration, or remove a configuration. You can also use the osc-setcoremem simulator to test different configurations.

Display the Current Domain Configuration (osc-setcoremem) Command

This procedure describes how to display a compute node's 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 87.

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

2. **View the existing domains and resources.**

   Use the osc-setcoremem -list command to view domains and resources.

   For example:

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

                    osc-setcoremem (non-interactive)
                  v2.3  built on May 25 2017 09:36:31

 Current Configuration: SuperCluster Half-Populated M7 Extended

 +-----+---------------------------+---- CURRENT ---+-----------+--- MIN REQD ---+
 |  ID | DOMAIN                    | CORES | MEM GB |   TYPE    | CORES | MEM GB |
 +-----+---------------------------+-------+--------+-----------+-------+--------+
 |   1 | primary                   |    32 |    240 | Dedicated |     2 |     32 |
 |   2 | ssccn2-dom1               |    32 |    240 | Dedicated |     2 |     32 |
 |   3 | ssccn2-dom3               |    32 |    240 | Dedicated |     2 |     32 |
 |
 +-----+---------------------------+-------+--------+-----------+-------+--------+
```

Related Information

- "Display the Current Domain Configuration (ldm)" on page 87
- "Guidelines for Modifying the Base Configuration File" on page 117
- "Change CPU/Memory Allocations (Core Granularity)" on page 94
- "Change CPU/Memory Allocations (Socket Granularity)" on page 89

▼ 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) Command" on page 86.

---

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

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

   Root Domains can be identified by matching the value of domain variable oes_domain_type with root. In this example, ssccn1-dom3 is a Root Domain.

```
# ldm ls-variable oes_domain_type primary
oes_domain_type=db

# ldm ls-variable oes_domain_type ssccn1-dom1
oes_domain_type=db

..

# ldm ls-variable oes_domain_type ssccn1-dom3
oes_domain_type=root
```

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

- “Display osc-setcoremem Help” on page 88
- “Display the Current Domain Configuration (osc-setcoremem) Command” on page 86
- “Guidelines for Modifying the Base Configuration File” on page 117
- “Using the osc-setcoremem Simulator” on page 115

▼ Display osc-setcoremem Help

1. **Log in as superuser on the compute node's control domain.**
2. **View the help for the osc-setcoremem command and the osc-setcoremem simulator.**

   For example:

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

                                  osc-setcoremem
                            v2.4  built on Feb 27 2017 18:23:43
 NON-INTERACTIVE ARGUMENTS
        -l, -list  .. list eligible domains

        -R, -reset [-y|-n] .. return the system to the original state of socket/core/memory allocations

        -type socket|core -res #sockets|#cores/memcap[:#sockets|#cores/memcap]* [-y|n]
                 "min|MIN" and "max|MAX" strings can be substituted in place of #sockets, #cores and
                         memory capacity to specify the minimum and maximum possible resource assignments
                         respectively.
                 "current|CURRENT" string can be used to accept the current config as is. In the
                         worst case, tool may use a value that is as close as possible to the current
                         value.
                 specify memory capacity in gigabytes aligned to 16 GB
          ** optional argument -y auto-confirms the desired change. likewise -n auto-rejects the changes
                 helps while testing).
```

```
SIMULATOR ARGUMENTS
        -gencfgtmpl|-g .. generate base configuration on a live SuperCluster node that can readily be
                 used along with -config|-c option
        -gencfgtmpl|-g -dc <domain-count> .. generate an editable base configuration file template
                 Base configuration is the core count & memory that each of the domains were
                 assigned during the initial setup.

         -platform|-p <T4|T5|M6|M7|M8> -config|-c <config-file-path> [ <<non-interactive arguments>> ]
                     .. establish a new base configuration
                 The file containing base configuration can be auto/hand generated, or based on a
                 template created by using "-gencfgtmpl|-g [ -dc <count> ]" options.

        << No option >> [ <<non-interactive arguments>> ] .. continue with the current state in
                     simulator, if exist, on a non-SuperCluster node
                 Set shell variable SSC_SCM_SIMULATE to any value on a live SuperCluster node to
                     enable this behavior
```

Related Information

- “Guidelines for Modifying the Base Configuration File” on page 117
- “Change CPU/Memory Allocations (Socket Granularity)” on page 89
- “Change CPU/Memory Allocations (Core Granularity)” on page 94
- “Display the Current Domain Configuration (`osc-setcoremem`) Command” on page 86
- “Using the `osc-setcoremem` Simulator” on page 115

▼ Change CPU/Memory Allocations (Socket Granularity)

Use the `osc-setcoremem` command to change how CPU and memory resources are allocated at the socket granularity level. Perform this procedure on each compute node.

You can use the `osc-setcoremem` simulator to view different allocations and select the best configuration for you. For instructions on using the `osc-setcoremem` simulator, see “Run a Simulation” on page 119. Examples are based on SuperCluster T5. However, the same concepts apply to other supported systems.

Note - To find out if you can perform this procedure, see “Supported Domain Configurations” on page 81.

The `osc-setcoremem` command 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 83).

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 the `osc-setcoremem` command to reconfigure the resources.**

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

```
# /opt/oracle.supercluster/bin/osc-setcoremem
                    osc-setcoremem (non-interactive)
                  v2.3  built on May 25 2017 09:36:31

 Current Configuration: SuperCluster Full-Rack T5-8

 +-----+----------------------------+---- CURRENT ---+-----------+--- MIN REQD ---+
```

```
|  ID | DOMAIN                       | CORES | MEM GB |   TYPE    | CORES | MEM GB |
+-----+------------------------------+-------+--------+-----------+-------+--------+
|   1 | primary                      |    32 |    512 | Dedicated |     8 |     64 |
|   2 | ssccn2-dom1                  |    16 |    256 | Dedicated |     4 |     32 |
|   3 | ssccn2-dom3                  |    16 |    256 | Dedicated |     4 |     32 |
|   4 | ssccn2-dom4                  |    16 |    256 | Dedicated |     4 |     32 |
+-----+------------------------------+-------+--------+-----------+-------+--------+

root@etc30zadm0201:/var/tmp/v24# ./osc-setcoremem

                          osc-setcoremem
                   v2.3  built on May 25 2017 09:36:31
 Current Configuration: SuperCluster Full-Rack T5-8

 +---------------------------------+-------+--------+-----------+--- MINIMUM ----+
 | DOMAIN                          | CORES | MEM GB |   TYPE    | CORES | MEM GB |
 +---------------------------------+-------+--------+-----------+-------+--------+
 | primary                         |    32 |    512 | Dedicated |     8 |     64 |
 | ssccn2-dom1                     |    16 |    256 | Dedicated |     4 |     32 |
 | ssccn2-dom2                     |     1 |     16 |   Root    |     1 |     16 |
 | ssccn2-dom3                     |    16 |    256 | Dedicated |     4 |     32 |
 | ssccn2-dom4                     |    16 |    256 | Dedicated |     4 |     32 |
 | ssccn2-dom5                     |     4 |     64 |   Root    |     4 |     64 |
 | *ssccn2-io-opcssc42             |     2 |     32 |    IO     |     1 |     16 |
 | *ssccn2-io-opcssc46             |     2 |     32 |    IO     |     1 |     16 |
 | *ssccn2-io-opcssc46-app         |     2 |     32 |    IO     |     1 |     16 |
 | *ssccn2-io-opcssc48             |     2 |     32 |    IO     |     1 |     16 |
 | *ssccn2-io-opcssc48-app         |     2 |     32 |    IO     |     1 |     16 |
 | *ssccn2-io-opcssc50             |     2 |     32 |    IO     |     1 |     16 |
 | *ssccn2-io-opcssc50-app         |     2 |     32 |    IO     |     1 |     16 |
 +---------------------------------+-------+--------+-----------+-------+--------+
 | Parked Resources (Approx)       |    29 |    240 |    --     |    -- |   --   |
 | Memory in Use by _sys_ Pool     |    -- |   1.25 |    --     |    -- |   --   |
 +---------------------------------+-------+--------+-----------+-------+--------+

 [ INFO ] following domains will be ignored in this session.

 Root Domains             IO Domains
 ------------             ----------
 ssccn2-dom2              ssccn2-io-opcssc42
 ssccn2-dom5              ssccn2-io-opcssc46
                          ssccn2-io-opcssc46-app
                          ssccn2-io-opcssc48
                          ssccn2-io-opcssc48-app
                          ssccn2-io-opcssc50
                          ssccn2-io-opcssc50-app

 CPU Granularity Preference:
```

```
       1. Socket
       2. Core

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

Choose Socket or Core [S or C] s

 Step 1 of 1: Socket Count
 primary      : specify socket count [min: 1, max: 2. default: 1] :
                you chose [1] socket for primary domain

 ssccn2-dom1  : specify socket count [min: 1, max: 2. default: 1] : 1
                you chose [1] socket for ssccn2-dom1 domain

 ssccn2-dom3  : specify socket count [min: 1, max: 2. default: 1] : 1
                you chose [1] socket for ssccn2-dom3 domain

 ssccn2-dom4  : specify socket count [min: 1, max: 2. default: 1]:
                you chose [1] socket for ssccn2-dom4 domain

 Configuration In Progress After Socket Count Selection:

 +---------------------------------+----------+----------+-----------+
 | DOMAIN                          |  SOCKETS |   MEM GB |   TYPE    |
 +---------------------------------+----------+----------+-----------+
 | primary                         |     1    |   228.00 | Dedicated |
 | ssccn2-dom1                     |     1    |   228.00 | Dedicated |
 | ssccn2-dom3                     |     1    |   228.00 | Dedicated |
 | ssccn2-dom4                     |     1    |   228.00 | Dedicated |
 | *ssccn2-dom2                    |   0.062  |       16 |   Root    |
 | *ssccn2-dom5                    |   0.250  |       64 |   Root    |
 | *ssccn2-io-opcssc42             |   0.125  |       32 |    IO     |
 | *ssccn2-io-opcssc46             |   0.125  |       32 |    IO     |
 | *ssccn2-io-opcssc46-app         |   0.125  |       32 |    IO     |
 | *ssccn2-io-opcssc48             |   0.125  |       32 |    IO     |
 | *ssccn2-io-opcssc48-app         |   0.125  |       32 |    IO     |
 | *ssccn2-io-opcssc50             |   0.125  |       32 |    IO     |
 | *ssccn2-io-opcssc50-app         |   0.125  |       32 |    IO     |
 +---------------------------------+----------+----------+-----------+
 | Parked Resources (Approx)       |   2.812  |      608 |    --     |
 | Memory in Use by _sys_ Pool     |      --  |     1.25 |    --     |
 +---------------------------------+----------+----------+-----------+

 [ INFO ] Equivalent non-interactive input string for the record:
         /opt/oracle.supercluster/bin/osc-setcorememm -type socket -res 1:1:1:1

 Following domains will be restarted:
```

```
        ssccn2-dom1
        ssccn2-dom3
        ssccn2-dom4

 This configuration change 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_05-25-2015_09:56:31.log
7
 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 ..
 *=====*=====*=====*=====*=====*=====*=====*=====*=====*=====*

                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.**

   For 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`) Command” on page 86
   - “Display the Current Domain Configuration (`ldm`)” on page 87

- “Access `osc-setcoremem` Log Files” on page 113

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

Related Information

- “Change CPU/Memory Allocations (Core Granularity)” on page 94
- “Change CPU/Memory Allocations Non-Interactively” on page 98
- “Display the Current Domain Configuration (`osc-setcoremem`) Command” on page 86
- “Using the `osc-setcoremem` Simulator” on page 115

▼ 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 81.

Note - Examples are based on SuperCluster M6-32, however, the same concepts apply to SuperCluster M7 and SuperCluster M8.

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 83).

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`) Command” on page 86
- “Display the Current Domain Configuration (`osc-setcoremem`) Command” on page 86
- “Access `osc-setcoremem` Log Files” on page 113

For 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

- "Change CPU/Memory Allocations Non-Interactively" on page 98
- "Park Cores and Memory" on page 102
- "Display the Current Domain Configuration (`osc-setcoremem`) Command" on page 86
- "Using the `osc-setcoremem` Simulator" on page 115

▼ Change CPU/Memory Allocations Non-Interactively

By default, `osc-setcoremem` runs in interactive mode, prompting you for input as it runs. Starting with SuperCluster software version 2.3, you can provide all of the required input using a combination options on a single command line.

You can use this procedure to change resource allocations at the core or socket granularity level.

Note - To find out if you can perform this procedure, see "Supported Domain Configurations" on page 81.

Non-Interactive syntax:

`osc-setcoremem -type socket|core -res` [`min`|`max`|*x_sockets*|*x_cores*/*x_memgb*]:[...]: [`min`|`max`|*x_sockets*|*x_cores*/*x_memgb*] [`-y`|`-n`] [`-h`|`-help`]

Option	Arguments	Description
`-type`	Specifies the type of granularity used for resource allocations. The choices are:	

Option	Arguments	Description
	`socket`	Specifies socket granularity.
	`core`	Specifies core granularity.
`-res`	Is the resource option that Is followed by the resource assignments for each of the eligible domains with each resource assignment separated by a `:` (colon). The order of resource assignments matches the order shown in the output from `osc-setcoremem -list` (see "Display the Current Domain Configuration (`osc-setcoremem`) Command" on page 86). You can specify symbolic values like `min` or `max`, and the tool automatically calculates the resource allocation values, or specify an explicit integer. You can mix the symbolic values and integer values. The choices are:	
	`min` or `MIN`	Specifies the minimum possible resource assignment.
	`max` or `MAX`	Specifies the maximum possible resource assignment.
	`current` or `CURRENT`	Accepts the current value as is. If it is not possible to retain the current value, the tool will use a value that is as close as possible to the current value.
	x_sockets or *x_cores*/*x_memgb*	Specify integer values for the resource assignments. Use one of these values: ■ *x_sockets* – For socket granularity, an integer that specifies the number of sockets. The tool automatically assigns proportional memory. ■ *x_cores*/*x_memgb* – For core granularity, The first integer specifies the number of cores, the second integer specifies the amount of memory in gigabytes (in multiples of 16 GB). For example: `12/256` indicates 12 cores and 256 GB of memory. **Note -** For *x_sockets*, *x_cores*, and *x_memgb*, you can specify `min` and `max` instead of a specific integer.
`-y` or `-n`	An optional option, that determines if you are prompted to confirm changes or not. The choices are: ■ `-n` – The tool displays the new resource allocations and you are prompted to confirm the changes. You can use this option to conduct a dry run without making actual changes by cancelling at the confirmation. ■ `-y` – As soon as you press Return, the tool immediately proceeds to make the resource allocations. You are not prompted for confirmation. If omitted, `-n` is assumed.	
`-h` or `-help`	Lists all of the non-interactive options and arguments that are listed in this table.	

These examples show the variety of arguments you can specify for the `-res` option:

```
# osc-setcoremem -type socket -res min:max:min

# osc-setcoremem -type socket -res 2:max:min -y

# osc-setcoremem -type socket -res 1:2:1:2 -n

# osc-setcoremem -type socket -res 2:current:min -y

# osc-setcoremem -type socket -res 1:min:2:max -y
```

```
# osc-setcoremem -type core -res min/min:min/max:min/min

# osc-setcoremem -type core -res 12/512:6/64:18/128 -y

# osc-setcoremem -type core -res min/current:max/32:current/current

# osc-setcoremem -type core -res 4/96:min/32:16/max -n

# osc-setcoremem -type core -res min/min:min/32:max/512:8/48 -y
```

1. **Log in as superuser on the compute node's control domain.**
2. **For any domains that might require restarting after the reconfiguration, 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. **List the current resource allocations.**

   For example:

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

                    osc-setcoremem (batch mode)
                 v2.3  built on Sep 26 2016 15:15:32

 Current Configuration: SuperCluster Half-Populated M7 Extended
 +-----+-----------------------------+---- CURRENT ---+-----------+--- MIN REQD ---+
 |  ID | DOMAIN                      | CORES | MEM GB |   TYPE    | CORES | MEM GB |
 +-----+-----------------------------+-------+--------+-----------+-------+--------+
 |   1 | primary                     |    32 |    240 | Dedicated |     2 |     32 |
 |   2 | ssccn1-dom1                 |    32 |    240 | Dedicated |     2 |     32 |
 |   3 | ssccn1-dom2                 |    32 |    240 | Dedicated |     2 |     32 |
 +-----+-----------------------------+-------+--------+-----------+-------+--------+
```

5. **Run `osc-setcoremem` non-interactively to reconfigure the resources.**

   For example:

```
# /opt/oracle.supercluster/bin/osc-setcoremem -type core -res 16/128:min/min:max/max

                    osc-setcoremem (batch mode)
                 v2.3  built on Sep 26 2016 15:15:32

 Current Configuration: SuperCluster Half-Populated M7 Extended
```

```
+----------------------------------+-------+--------+-----------+--- MINIMUM ----+
| DOMAIN                           | CORES | MEM GB |   TYPE    | CORES | MEM GB |
+----------------------------------+-------+--------+-----------+-------+--------+
| primary                          |    32 |    240 | Dedicated |     2 |     32 |
| ssccn1-dom1                      |    32 |    240 | Dedicated |     2 |     32 |
| ssccn1-dom2                      |    32 |    240 | Dedicated |     2 |     32 |
| ssccn1-dom3                      |     2 |     32 |   Root    |     2 |     32 |
| *ssccn1-io-etc4m7-ioappadm0101   |     2 |     32 |    IO     |     1 |     16 |
| *ssccn1-io-etc4m7zioadm0101      |     4 |  62.50 |    IO     |     1 |     16 |
+----------------------------------+-------+--------+-----------+-------+--------+
| Not in Use (Parked)              |    24 |    114 |    --     |    -- |   --   |
+----------------------------------+-------+--------+-----------+-------+--------+

[ INFO ] following domains will be ignored in this session.

Root Domains              IO Domains
------------              ----------
ssccn1-dom3               ssccn1-io-etc4m7-ioappadm0101
                          ssccn1-io-etc4m7zioadm0101

Configuration In Progress After Memory Capacity Selection:
+----------------------------------+-------+--------+-----------+--- MINIMUM ----+
| DOMAIN                           | CORES | MEM GB |   TYPE    | CORES | MEM GB |
+----------------------------------+-------+--------+-----------+-------+--------+
| primary                          |    16 |    128 | Dedicated |     2 |     64 |
| ssccn1-dom1                      |     2 |     32 | Dedicated |     2 |     32 |
| ssccn1-dom2                      |    78 |    560 | Dedicated |     2 |    320 |
| *ssccn1-dom3                     |     2 |     32 |   Root    |     2 |     32 |
| *ssccn1-io-etc4m7-ioappadm0101   |     2 |     32 |    IO     |     1 |     16 |
| *ssccn1-io-etc4m7zioadm0101      |     4 |  62.50 |    IO     |     1 |     16 |
+----------------------------------+-------+--------+-----------+-------+--------+
| Not in Use (Parked)              |    24 |    114 |    --     |    -- |   --   |
+----------------------------------+-------+--------+-----------+-------+--------+

Following domains will be restarted:
       ssccn1-dom2
       ssccn1-dom1

This configuration change requires rebooting the Control Domain.
Do you want to proceed? Y/N : y
```

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`) Command” on page 86
- “Access `osc-setcoremem` Log Files” on page 113

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

Related Information

- “Park Cores and Memory” on page 102
- “View the SP Configuration” on page 107
- “Revert to a Previous CPU/Memory Configuration” on page 109
- “Display the Current Domain Configuration (`osc-setcoremem`) Command” on page 86

▼ 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 might not be able to move resources back to dedicated domains.

Prior to SuperCluster software version 3.0, if you parked resources from dedicated domains, you could not move parked resources to back to dedicated domains after I/O Domains were created. As of version 3.0, you can move parked unallocated resources back to dedicated domains.

Note - To find out if you can perform this procedure, see “Supported Domain Configurations” on page 81.

Note - Examples are based on SuperCluster M6-32, however, the same concepts apply to SuperCluster M7 and SuperCluster M8.

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 83).

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

Domain	Domain Type	Cores Before	Cores After	Memory Before (GB)	Memory After (GB)
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-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                        |    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-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_16:18:57.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%
 |-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
 *=====*=====*=====*=====*=====*=====*=====*=====*=====*=====*

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`) Command” on page 86
   - “Access `osc-setcoremem` Log Files” on page 113

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`) Command" on page 86
   - "Access `osc-setcoremem` Log Files" on page 113

   For 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 87:

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

10. **Verify the parked memory.**

    See "Display the Current Domain Configuration (`ldm`)" on page 87:

```
# 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

- "View the SP Configuration" on page 107
- "Revert to a Previous CPU/Memory Configuration" on page 109
- "Display the Current Domain Configuration (`osc-setcoremem`) Command" on page 86
- "Using the `osc-setcoremem` Simulator" on page 115

▼ 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 time stamp 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

- “Revert to a Previous CPU/Memory Configuration” on page 109
- “Remove a CPU/Memory Configuration” on page 110
- “Reset the CPU/Memory Configuration” on page 111
- “Using the `osc-setcoremem` Simulator” on page 115

▼ 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 81.

---

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 113.

```
# 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 107.

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

- “Remove a CPU/Memory Configuration” on page 110
- “Reset the CPU/Memory Configuration” on page 111
- “Using the `osc-setcoremem` Simulator” on page 115

▼ 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.**

   For example:

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

Related Information

- “Reset the CPU/Memory Configuration” on page 111
- “Display the Current Domain Configuration (`osc-setcoremem`) Command” on page 86
- “Using the `osc-setcoremem` Simulator” on page 115

Reset the CPU/Memory Configuration

Use this procedure to return the system to the original state of socket/core/memory allocations without affecting other changes that were made in each of the domains since the system was set up initially. This option does not reset the SP configuration.

The syntax for this `osc-setcoremem` feature was introduced in SuperCluster version 2.3 software:

`osc-setcoremem -R|reset [-y|-n]`

The command supports these options:

- `-R` or `reset` – Returns the CPU and memory configuration to the original configuration.
- `-y` or `-n` – (Optional) If `-y` is specified, you are not prompted to confirm the operation, and he reset takes effect immediately. If `-n` is specified, the tool shows you the output as if the reset operation ran, but the CPU and memory allocations are not changed.

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

2. **Reset the CPU and memory configuration.**

   For example:

```
# ./osc-setcoremem -reset

 Current Configuration: SuperCluster Half-Populated M7 Extended

 +---------------------------------+-------+--------+-----------+--- MINIMUM ----+
 | DOMAIN                          | CORES | MEM GB |   TYPE    | CORES | MEM GB |
 +---------------------------------+-------+--------+-----------+-------+--------+
 | primary                         |    32 |    240 | Dedicated |     2 |     32 |
 | ssccn1-dom1                     |    32 |    240 | Dedicated |     2 |     32 |
 | ssccn1-dom2                     |    32 |    240 | Dedicated |     2 |     32 |
 | ssccn1-dom3                     |     2 |     32 |   Root    |     2 |     32 |
 | *ssccn1-io-etc4m7-ioappadm0101  |     2 |     32 |    IO     |     1 |     16 |
 | *ssccn1-io-etc4m7zioadm0101     |     4 |  62.50 |    IO     |     1 |     16 |
 +---------------------------------+-------+--------+-----------+-------+--------+
 | Not in Use (Parked)             |    24 |    114 |    --     |    -- |   --   |
 +---------------------------------+-------+--------+-----------+-------+--------+

 [ INFO ] following domains will be ignored in this session.

 Root Domains             IO Domains
 ------------             ----------
 ssccn1-dom3              ssccn1-io-etc4m7-ioappadm0101
```

```
                          ssccn1-io-etc4m7zioadm0101

CPU Granularity Preference:

       1. Socket
       2. Core

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

Choose Socket or Core [S or C] S

Step 1 of 1: Socket Count

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

ssccn1-dom1  : specify socket count [min: 1, max: 1. default: 1] : 1
               you chose [1] socket for ssccn1-dom1 domain

ssccn1-dom2  : specify socket count [min: 1, max: 1. default: 1] : 1
               you chose [1] socket for ssccn1-dom2 domain


Configuration In Progress After Socket Count Selection:

+----------------------------------+----------+----------+-----------+
| DOMAIN                           |  SOCKETS |   MEM GB |   TYPE    |
+----------------------------------+----------+----------+-----------+
| primary                          |     1    |      240 | Dedicated |
| ssccn1-dom1                      |     1    |      240 | Dedicated |
| ssccn1-dom2                      |     1    |      240 | Dedicated |
| *ssccn1-dom3                     |   0.062  |       32 |   Root    |
| *ssccn1-io-etc4m7-ioappadm0101   |   0.062  |       32 |    IO     |
| *ssccn1-io-etc4m7zioadm0101      |   0.125  |    62.50 |    IO     |
+----------------------------------+----------+----------+-----------+
| Not in Use (Parked)              |   0.750  |      114 |    --     |
+----------------------------------+----------+----------+-----------+

       Nothing to reconfigure.
       You can continue using the system.
```

Related Information

- “Using the `osc-setcoremem` Simulator” on page 115
- “Display the Current Domain Configuration (`osc-setcoremem`) Command” on page 86

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*

   For 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-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.

::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.
```

Related Information

- “Using the osc-setcoremem Simulator” on page 115
- “Display the Current Domain Configuration (osc-setcoremem) Command” on page 86

Using the osc-setcoremem Simulator

This osc-setcoremem feature was introduced in SuperCluster software version 2.4.

The osc-setcoremem simulator runs simulations to view possible configurations to determine the best way to configure CPU and memory resources. The simulation runs only on the hardware and does not permanently change the CPU or memory resources.

You can run the simulator with any of these configurations:

- No configuration – You must generate a base configuration file using the -p and -c options or the -g option. For instructions, see "Run a Simulation" on page 119.
- A base configuration – You can return to the base configuration any time with the -R option. For instructions, see "Reset the CPU/Memory Configuration" on page 111.
- A random configuration – You might have updated the base configuration or some other configuration. For example, you reconfigured the domains in a simulated environment.

For more information on command line options, see "Display osc-setcoremem Help" on page 88.

Use the osc-setcoremem simulator for these types of scenarios:

- Move from a base configuration to a desired configuration – Run the osc-setcoremem simulator once to view the desired configuration in a virtual environment. When you are ready to make the changes, run the osc-setcoremem command. See "Using the osc-setcoremem Command" on page 85.
- Move from a current configuration to a new configuration – You must run the osc-setcoremem simulator at least two times. The first time simulates the move from a base configuration to the current configuration. The second time simulates the move from the current configuration to the desired configuration. After viewing the changes in the virtual environment, run the osc-setcoremem command to make the changes.

  For example, if you have a modified configuration on a live SuperCluster system, and you plan to move to a different custom configuration, use the osc-setcoremem simulator to see what is possible. The simulator cannot start from a random configuration, so you must first run the simulator to start with the base configuration and provide changes to emulate your current custom domain configuration. Then run the simulator again and start with the virtual configuration that matches your current configuration and proceed with different simulations. After viewing the changes in the virtual environment, run the osc-setcoremem command to make the changes.

This section covers the following topics:

- "Example: Simulating Changes on a Non-SuperCluster Node" on page 122

Guidelines for Modifying the Base Configuration File

Use the following guidelines when you are manually editing the base configuration file, or when you are editing the base configuration template that was generated by the `osc-setcoremem` command.

- Each row in the file contains information that pertains to a single eligible domain.

  DEDICATED and ROOT domains are the only domain types that are eligible. Any domain that is not a ROOT domain is treated as a DEDICATED domain. I/O Domains are not included in the base configuration file.
- There are no optional values. Each of the columns must contain a value. Integer fields are expected to have a positive non-zero value.
- On each line, separate each column with one or more white spaces or preferably with a tab.
- Prepend each insignificant line with a # symbol. These lines will be treated as comments.
- Core count should be evenly divisible by the total core count per processor that was specified with the `-platform` option. For example, 24 is a valid core count. Sixteen or 20 is not a valid core count when an M6 configuration is being simulated, because each M6 processor has 12 cores.
- Memory capacity should be a multiple of 16. The minimum memory capacity is 16 GB. For example, 1024 GB is a valid memory capacity but 1000 GB is not.
- Limit domain names to a maximum 48 characters. Domain names that are longer that 48 characters are automatically truncated to 48 characters during simulation.
- Acceptable inputs for ROOT & DOMAIN and SERVICE DOMAIN columns include YES, NO, Y, N, yes, no, y, and n.
- You can choose any file name for this edited base configuration file.
- To minimize validation errors or unexpected results, simulate an actual supported configuration rather than a fictional configuration.
- You can perform simulations as any user (root or non-root) on both SuperCluster and non-SuperCluster nodes.

The format of the base configuration file is shown below.

```
#DOMAIN        ROOT       SERVICE      SOCKET     CORE      MEMORY     HCA
#NAME          DOMAIN     DOMAIN       COUNT      COUNT     GB         COUNT
#-------------------------------------------------------------------------------
#<STRING>      YES|NO     YES|NO       <INT>      <INT>     <INT>      <INT>
```

Related Information

- “Prepare the Base Configuration File” on page 118
- “Run a Simulation” on page 119
- “Example: Simulating Changes on a Non-SuperCluster Node” on page 122
- “Using the `osc-setcoremem` Simulator” on page 115

▼ Prepare the Base Configuration File

The `osc-setcoremem` simulator requires a base configuration file as input to run the simulator for the first time, and each subsequent time you want to simulate a different configuration. You must provide the SuperCluster base configuration file that you want to simulate. This is required no matter where the simulation tests will be run—on a SuperCluster compute node or a non-SuperCluster SPARC system.

You can use a ready-to-use base configuration file or edit a base configuration file template. Choose one of the methods described in this procedure to prepare the base configuration file.

1. **Log in as superuser on the compute node's Control Domain.**

2. **Generate a ready-to-use base configuration file on a live SuperCluster environment that you plan to simulate in other places.**

   Ensure that you completed Step 1 before you generate the base configuration file.

```
% ./osc-setcoremem -g
[OK] simulator config file generated
location: /var/tmp/oscsetcorememcfg.txt

% cat /var/tmp/oscsetcorememcfg.txt
#DOMAIN         ROOT        SERVICE     SOCKET     CORE      MEMORY       HCA
#NAME           DOMAIN      DOMAIN      COUNT      COUNT     GB           COUNT
#-------------------------------------------------------------------------------
primary         NO          YES         2          24        2048         1
ssccn-dom1      NO          NO          2          24        2048         1
ssccn-dom2      NO          NO          2          24        2048         1
ssccn-dom3      YES         YES         2          2         32           1
```

   Use this configuration file as input to the simulator without any editing.

3. **Alternatively, you can generate a base configuration file template and manually populate the domain resource details of the target environment that you want to simulate.**

The template file can be generated on any SPARC system (not necessarily on a SuperCluster compute node). You can perform this step as any user on both SuperCluster and non-SuperCluster systems.

```
% ./osc-setcorememm -g -dc 4
[OK] simulator config file generated
location: /var/tmp/oscsetcorememcfg.txt

% cat /var/tmp/oscsetcorememcfg.txt
#DOMAIN        ROOT        SERVICE      SOCKET      CORE        MEMORY        HCA
#NAME          DOMAIN      DOMAIN       COUNT       COUNT       GB            COUNT
#-------------------------------------------------------------------------------
primary        YES|NO      YES|NO       <COUNT>     <COUNT>     <CAPACITY>    1|2
ssccn-dom1     YES|NO      YES|NO       <COUNT>     <COUNT>     <CAPACITY>    1|2
ssccn-dom2     YES|NO      YES|NO       <COUNT>     <COUNT>     <CAPACITY>    1|2
ssccn-dom3     YES|NO      YES|NO       <COUNT>     <COUNT>     <CAPACITY>    1|2
```

Edit the template file using the information in "Guidelines for Modifying the Base Configuration File" on page 117.

4. **Copy the base configuration file to a location that is accessible to the `osc-setcorememm` simulator.**

Tip - For syntax on all of the simulator options, see "Display `osc-setcorememm` Help" on page 88. At any time, you can reset the configuration to the original base configuration.

▼ Run a Simulation

1. **Prepare the base configuration file.**

   Follow the steps in "Prepare the Base Configuration File" on page 118 for instructions.

2. **Import the base configuration file from Step 1 and start the simulation.**

   Specify the SuperCluster platform to simulate, and the path to the base configuration file. You can also specify non-interactive options.

   For example:

```
% ./osc-setcorememm -p M6 -c ./oscsetcorememcfg-m6.txt
```

```
                    osc-setcoremem simulator
                 v2.4  built on Feb 24 2017 14:48:42

Current Configuration: SuperCluster M6

+---------------------------------+-------+--------+-----------+--- MINIMUM ----+
| DOMAIN                          | CORES | MEM GB |   TYPE    | CORES | MEM GB |
+---------------------------------+-------+--------+-----------+-------+--------+
| primary                         |    24 |   2048 | Dedicated |     2 |     32 |
| ssccn2-dom1                     |    24 |   2048 | Dedicated |     2 |     32 |
| ssccn2-dom2                     |    24 |   2048 | Dedicated |     2 |     32 |
| ssccn2-dom3                     |     2 |     32 |   Root    |     2 |     32 |
+---------------------------------+-------+--------+-----------+-------+--------+
| Parked Resources (Approx)       |    22 |   2016 |    --     |    -- |   --   |
+---------------------------------+-------+--------+-----------+-------+--------+

[ INFO ] following domains will be ignored in this session.

Root Domains
------------
ssccn2-dom3

CPU Granularity Preference:
       1. Socket
       2. Core
In case of Socket granularity, proportional memory capacity is automatically selected
for you.

Choose Socket or Core [S or C] c
...
<output omitted>
...
```

Caution - Use the `-p` and `-c` options only to remove the existing configuration and start a new base configuration.

3. **Using the modified configuration from Step 1, continue the simulation by running the `osc-setcoremem` binary file.**

   You can also specify non-interactive options.

   a. **To run this simulation on a SuperCluster node, set a shell variable called `SSC_SCM_SIMULATE` to any value and start the simulator.**

      For example:

      ```
      % SSC_SCM_SIMULATE=1 ./osc-setcoremem
      ```

As the simulator runs, enter changes to the cores and memory. To see an example of the entire file, see "Example: Simulating Changes on a Non-SuperCluster Node" on page 122.

Tip - At any time, you can continue a simulation on a SuperCluster node from where you left off by simply setting the `SSC_SCM_SIMULATE` variable. You can run the `osc-setcoremem` in normal execution mode in between simulations.

The absence of the `SSC_SCM_SIMULATE` variable triggers normal (non-simulated) `osc-setcoremem` execution.

b. To run this simulation on a non-SuperCluster node, do not set the shell variable when you start the simulator.

For example:

```
% ./osc-setcoremem
```

As the simulator runs, enter changes to the cores and memory.

4. View the modified domain configuration.

```
% SSC_SCM_SIMULATE=1 ./osc-setcoremem -list

                osc-setcoremem simulator (non-interactive)
                   v2.4  built on Feb 27 2017 18:23:43
 Current Configuration: SuperCluster M7
 +-----+---------------------------+---- CURRENT ---+-----------+--- MIN REQD ---+
 |  ID | DOMAIN                    | CORES | MEM GB |   TYPE    | CORES | MEM GB |
 +-----+---------------------------+-------+--------+-----------+-------+--------+
 |   1 | primary                   |     8 |    128 | Dedicated |     2 |     32 |
 |   2 | ssccn1-dom1               |    16 |    256 | Dedicated |     2 |     32 |
 |   3 | ssccn1-dom2               |    32 |    480 | Dedicated |     2 |     32 |
 +-----+---------------------------+-------+--------+-----------+-------+--------+
```

5. Repeat Step 2 and Step 3 to simulate different configurations and determine which configuration you want to implement.

To purge the current modified configuration file and start over with the original base configuration file, use the `-r` option to reset your changes. For instructions, see "Display `osc-setcoremem` Help" on page 88.

Related Information

- "Example: Simulating Changes on a Non-SuperCluster Node" on page 122

- "Using the osc-setcoremem Simulator" on page 115

Example: Simulating Changes on a Non-SuperCluster Node

This example shows how to generate an editable base configuration file template and run the simulator on a non-SuperCluster SPARC node. The template was edited to reflect a SuperCluster M6 configuration with four domains, with each domain running with two M6 processors and two TB memory.

It is not necessary to use the SSC_SCM_SIMULATE variable in this case. The shell variable is used only when running the simulator on a SuperCluster node.

```
% ./osc-setcoremem -g -dc 4
[OK] simulator config file generated
location: /var/tmp/oscsetcorememcfg.txt

% cat /var/tmp/oscsetcorememcfg.txt
#DOMAIN      ROOT       SERVICE     SOCKET     CORE       MEMORY        HCA
# NAME       DOMAIN     DOMAIN      COUNT      COUNT      GB            COUNT
#-------------------------------------------------------------------------------
primary      YES|NO     YES|NO      <COUNT>    <COUNT>    <CAPACITY>    1|2
ssccn-dom1   YES|NO     YES|NO      <COUNT>    <COUNT>    <CAPACITY>    1|2
ssccn-dom2   YES|NO     YES|NO      <COUNT>    <COUNT>    <CAPACITY>    1|2
ssccn-dom3   YES|NO     YES|NO      <COUNT>    <COUNT>    <CAPACITY>    1|2

% ./osc-setcoremem -p m6 -c /var/tmp/oscsetcorememcfg.txt

                    osc-setcoremem simulator
                v2.4  built on Feb 24 2017 14:48:42

 Current Configuration: SuperCluster M6
 +---------------------------------+-------+--------+-----------+--- MINIMUM ----+
 | DOMAIN                          | CORES | MEM GB |   TYPE    | CORES | MEM GB |
 +---------------------------------+-------+--------+-----------+-------+--------+
 | primary                         |    24 |   2048 | Dedicated |     2 |     32 |
 | ssccn2-dom1                     |    24 |   2048 | Dedicated |     2 |     32 |
 | ssccn2-dom2                     |    24 |   2048 | Dedicated |     2 |     32 |
 | ssccn2-dom3                     |     2 |     32 |   Root    |     2 |     32 |
 +---------------------------------+-------+--------+-----------+-------+--------+
 | Parked Resources (Approx)       |    22 |   2016 |    --     |    -- |   --   |
 +---------------------------------+-------+--------+-----------+-------+--------+
```

```
[ INFO ] following domains will be ignored in this session.

Root Domains
------------
ssccn2-dom3

CPU Granularity Preference:
       1. Socket
       2. Core

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

Choose Socket or Core [S or C] c

Step 1 of 2: Core Count

primary      : specify number of cores [min: 2, max: 68. default: 24] : 36
               you chose [36] cores for primary domain

ssccn2-dom1  : specify number of cores [min: 2, max: 34. default: 24] : 30
               you chose [30] cores for ssccn2-dom1 domain

ssccn2-dom2  : specify number of cores [min: 2, max: 6. default: 6] :
               you chose [6] cores for ssccn2-dom2 domain

Configuration In Progress After Core Count Selection:
+---------------------------------+-------+--------+-----------+--- MINIMUM ----+
| DOMAIN                          | CORES | MEM GB |   TYPE    | CORES | MEM GB |
+---------------------------------+-------+--------+-----------+-------+--------+
| primary                         |    36 |   2048 | Dedicated |     2 |    160 |
| ssccn2-dom1                     |    30 |   2048 | Dedicated |     2 |    128 |
| ssccn2-dom2                     |     6 |   2048 | Dedicated |     2 |     32 |
| *ssccn2-dom3                    |     2 |     32 |   Root    |     2 |     32 |
+---------------------------------+-------+--------+-----------+-------+--------+
| Parked Resources (Approx)       |    22 |   2016 |    --     |    -- |   --   |
+---------------------------------+-------+--------+-----------+-------+--------+

Step 2 of 2: Memory Capacity
       (must be 16 GB aligned)
primary      : specify memory capacity in GB [min: 160, max: 4032. default: 2048] :
3072
               you chose [3072 GB] memory for primary domain

ssccn2-dom1  : specify memory capacity in GB [min: 128, max: 3040. default: 2048] :
1536
               you chose [1536 GB] memory for ssccn2-dom1 domain
```

```
ssccn2-dom2  : specify memory capacity in GB [min: 32, max: 992. default: 992] : 512
               you chose [512 GB] memory for ssccn2-dom2 domain

Configuration In Progress After Memory Capacity Selection:
+---------------------------------+-------+--------+-----------+--- MINIMUM ----+
| DOMAIN                          | CORES | MEM GB |   TYPE    | CORES | MEM GB |
+---------------------------------+-------+--------+-----------+-------+--------+
| primary                         |    36 |   3072 | Dedicated |     2 |    160 |
| ssccn2-dom1                     |    30 |   1536 | Dedicated |     2 |    128 |
| ssccn2-dom2                     |     6 |    512 | Dedicated |     2 |     32 |
| *ssccn2-dom3                    |     2 |     32 |   Root    |     2 |     32 |
+---------------------------------+-------+--------+-----------+-------+--------+
| Parked Resources (Approx)       |    22 |   3040 |    --     |    -- |   --   |
+---------------------------------+-------+--------+-----------+-------+--------+

[ INFO ] Equivalent non-interactive input string for the record:
        osc-setcoremem -type core -res 36/3072:30/1536:6/512

        DOMAIN SUMMARY
| Domain name      : ssccn2-dom2
| Total core count : 6
| Total memory     : 512.00 GB

        Resource Distribution among Locality Groups

        LGRP   CORES   MEMGRANS
         1       0         0
         2       0         0
         3       0         0
         4       0         0
         5       3        16
         6       3        16
         7       0         0
         8       0         0

        Home Locality Groups
                => 5
                => 6

        ____________________

| Domain name      : ssccn2-dom1
| Total core count : 30
| Total memory     : 1536.00 GB

        Resource Distribution among Locality Groups

        LGRP   CORES   MEMGRANS
         1       0         0
```

```
         2       0         0
         3      12        47
         4      12        47
         5       3         1
         6       3         1
         7       0         0
         8       0         0

        Home Locality Groups
                => 3
                => 4

        --------------------

| Domain name      : primary
| Total core count : 36
| Total memory     : 3072.00 GB

        Resource Distribution among Locality Groups
        LGRP   CORES   MEMGRANS
         1      12        64
         2      12        64
         3       0         0
         4       0         0
         5       6        32
         6       6        32
         7       0         0
         8       0         0

        Home Locality Groups
                => 1
                => 2
        DOMAIN REBOOT SUMMARY

The following domains would have rebooted on a live system:
  ssccn2-dom2
  ssccn2-dom1
  primary

       POSSIBLE NEXT STEP

Continue the simulation with updated configuration
        eg., <path>/osc-setcorememem [<option(s)>]

                - OR -

Start with an existing or brand new base configuration
        eg., <path>/osc-setcorememem  -p [T4|T5|M6|M7|M8] -c <path_to_config_file>
```

To continue the simulation with the updated configuration, run:

```
% ./osc-setcoremem
```

Related Information

- “Access `osc-setcoremem` Log Files” on page 113
- “Using the `osc-setcoremem` Simulator” on page 115
- “Display `osc-setcoremem` Help” on page 88
- “Obtaining the EM Exadata Plug-in” on page 127

Obtaining the EM Exadata Plug-in

You can monitor all Exadata-related software and hardware components in the cluster using the Oracle Enterprise Manager Exadata plug-in the supported configuration described in these topics.

- “Confirm System Requirements” on page 127
- “Known Issues With the EM Exadata Plug-in” on page 127

Related Information

- “Configuring the Exalogic Software” on page 129
- “Administering Oracle Solaris 11 Boot Environments” on page 147

▼ Confirm System Requirements

With the Oracle SuperCluster software version 2.*x*, the `compmon` command name is `osc-compmon`.

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

Related Information

- “Known Issues With the EM Exadata Plug-in” on page 127
- “Check the perfquery Version” on page 128

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:

`http://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 your SuperCluster system is version 1.5.8 or later (see "Check the perfquery Version" on page 128), 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.

Related Information

- "Confirm System Requirements" on page 127
- "Check the perfquery Version" on page 128

▼ Check the perfquery Version

- **Type:**

```
perfquery -V
```

Related Information

- "Confirm System Requirements" on page 127
- "Known Issues With the EM Exadata Plug-in" on page 127

Configuring the Exalogic Software

These topics describe how to use Exalogic software on SuperCluster M8 and SuperCluster M7.

Related Information

Exalogic Software Overview

Oracle EECS includes performance optimizations for SuperCluster systems to improve I/O, and thread management, and request handling efficiency.

Additional optimizations include reduced buffer copies, which result in more efficient I/O. 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.

Related Information

- “Prepare to Configure the Exalogic Software” on page 130
- “Enable Domain-Level Enhancements” on page 130
- “Enable Cluster-Level Session Replication Enhancements” on page 132
- “Configuring Grid Link Data Source for `Dept1_Cluster1`” on page 135
- “Configuring SDP-Enabled JDBC Drivers for `Dept1_Cluster1`” on page 141
- “Create an SDP Listener on the IB Network” on page 143

▼ Prepare to Configure the Exalogic Software

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

   Refer to the instructions in 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 the instructions in the *Oracle Exalogic Enterprise Deployment Guide* at `http://docs.oracle.com/cd/E18476_01/doc.220/e18479/toc.htm`.

Related Information

- “Exalogic Software Overview” on page 129
- “Enable Domain-Level Enhancements” on page 130
- “Enable Cluster-Level Session Replication Enhancements” on page 132
- “Configuring Grid Link Data Source for `Dept1_Cluster1`” on page 135
- “Configuring SDP-Enabled JDBC Drivers for `Dept1_Cluster1`” on page 141
- “Create an SDP Listener on the IB Network” on page 143

▼ Enable Domain-Level Enhancements

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

2. **Select Domain name in the left navigation pane.**

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 this 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.

Related Information

- “Exalogic Software Overview” on page 129
- “Prepare to Configure the Exalogic Software” on page 130
- “Enable Cluster-Level Session Replication Enhancements” on page 132
- “Configuring Grid Link Data Source for `Dept1_Cluster1`” on page 135
- “Configuring SDP-Enabled JDBC Drivers for `Dept1_Cluster1`” on page 141
- “Create an SDP Listener on the IB Network” on page 143

▼ 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.**

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

   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.**

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 IP address 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 Step 3 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.**

5. **Click `Dept1_Cluster1`.**

   This is the example cluster where you will deploy a web application at a later time.

6. **Click the Replication tab.**

7. **In the Replication Channel field, ensure that Replication Channel 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.**

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

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

d. **Repeat Step 11 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.**

Related Information

- "Exalogic Software Overview" on page 129
- "Prepare to Configure the Exalogic Software" on page 130
- "Enable Domain-Level Enhancements" on page 130
- "Configuring Grid Link Data Source for Dept1_Cluster1" on page 135
- "Configuring SDP-Enabled JDBC Drivers for Dept1_Cluster1" on page 141
- "Create an SDP Listener on the IB Network" on page 143

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 enables 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.

Related Information

- "Runtime Connection Load Balancing" on page 136
- "XA Affinity" on page 137
- "SCAN Addresses" on page 137
- "Secure Communication With Oracle Wallet" on page 138
- "Create a Grid Link Data Source on Dept1_Cluster1" on page 138

Runtime Connection Load Balancing

Runtime connection load balancing enables 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.

Related Information

- "Fast Connection Failover" on page 136

XA Affinity

XA Affinity for global transactions ensures that all the database 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.

Related Information

- "Fast Connection Failover" on page 136
- "Runtime Connection Load Balancing" on page 136
- "SCAN Addresses" on page 137
- "Secure Communication With Oracle Wallet" on page 138
- "Create a Grid Link Data Source on Dept1_Cluster1" on page 138

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.

Related Information

- "Fast Connection Failover" on page 136
- "Runtime Connection Load Balancing" on page 136
- "XA Affinity" on page 137
- "Secure Communication With Oracle Wallet" on page 138

- "Create a Grid Link Data Source on Dept1_Cluster1" on page 138

Secure Communication With Oracle Wallet

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

Related Information

- "Fast Connection Failover" on page 136
- "Runtime Connection Load Balancing" on page 136
- "XA Affinity" on page 137
- "SCAN Addresses" on page 137
- "Create a Grid Link Data Source on Dept1_Cluster1" on page 138

▼ 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.**
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 the 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, set up FAN notifications.**

    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 141.

Related Information

- "Fast Connection Failover" on page 136
- "Runtime Connection Load Balancing" on page 136

- "XA Affinity" on page 137
- "SCAN Addresses" on page 137
- "Secure Communication With Oracle Wallet" on page 138

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" on page 141
- "Enable SDP Support for JDBC" on page 141
- "Monitor SDP Sockets" on page 143

▼ 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 Infiniband 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.

Related Information

- "Enable SDP Support for JDBC" on page 141
- "Monitor SDP Sockets" on page 143

▼ 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. For 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.**

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

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

Related Information

- “Configure the Database to Support IB” on page 141
- “Monitor SDP Sockets” on page 143

Monitor SDP Sockets

You can monitor SDP sockets by running the `netstat` command on the Application Domains running Oracle Solaris 11 that contain EECS. 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. **Display the status of all SDP sockets.**

   # **`netstat -f sdp -s l`**

   This command displays the status of all SDP sockets (established or not), For example:

```
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
```

Related Information

- "Configure the Database to Support IB" on page 141
- "Enable SDP Support for JDBC" on page 141

Create an SDP Listener on the IB Network

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.

   For example:

```
# Added for Listener over IB
```

```
192.168.28.21 ssc01db01-ibvip.example.com ssc01db01-ibvip

192.168.28.22 ssc01db02-ibvip.example.com ssc01db02-ibvip
```

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

```
# /u01/app/grid/product/11.2.0.2/bin/srvctl add network -k 2 -S
192.168.28.0/255.255.252.0/stor_ipmp0
```

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: 2/192.168.28.0/255.255.252.0/stor_ipmp0, type static -- Output
indicating Network resource on the 192.168.28.0 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/255.255.252.0/stor_ipmp0 -k 2

srvctl add vip -n ssc01db02 -A ssc01db02-ibvip/255.255.252.0/stor_ipmp0 -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 enable load balancing and failover across multiple networks (Ethernet and IB).**

   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`.

   This 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 enter an `ALTER SYSTEM SET LISTENER_NETWORKS` command for the modifications to take affect by the instance.

```
EXA1SB1 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = etc7m7-exa1-scan)(PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = exa1sb1)
    )
  )


EXA1SB1_IB =
  (DESCRIPTION =
    (LOAD_BALANCE=on)
    (ADDRESS = (PROTOCOL = TCP)(HOST = etc7m7-c0p1-d1-ibvip)(PORT = 1522))
    (ADDRESS = (PROTOCOL = TCP)(HOST = etc7m7-c1p1-d1-ibvip)(PORT = 1522))
    (ADDRESS = (PROTOCOL = SDP)(HOST = etc7m7-c0p1-d1-ibvip)(PORT = 1522))
    (ADDRESS = (PROTOCOL = SDP)(HOST = etc7m7-c1p1-d1-ibvip)(PORT = 1522))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = exa1sb1)
    )
  )


LISTENER_IPLOCAL =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)(HOST = etc7m7-c0p1-d1-vip)(PORT = 1521))
    )
  )


LISTENER_IPREMOTE =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)(HOST = etc7m7-exa1-scan)(PORT = 1521))
    )
  )

LISTENER_IBLOCAL =
  (DESCRIPTION =
```

```
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)(HOST = etc7m7-c0p1-d1-ibvip)(PORT = 1522))
      (ADDRESS = (PROTOCOL = SDP)(HOST = etc7m7-c0p1-d1-ibvip)(PORT = 1522))
    )
  )

LISTENER_IBREMOTE =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP)(HOST = etc7m7-c1p1-d1-ibvip)(PORT = 1522))
      (ADDRESS = (PROTOCOL = SDP)(HOST = etc7m7-c1p1-d1-ibvip)(PORT = 1522))
    )
  )
```

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
   ```

Related Information

- “Exalogic Software Overview” on page 129
- “Prepare to Configure the Exalogic Software” on page 130
- “Enable Domain-Level Enhancements” on page 130
- “Enable Cluster-Level Session Replication Enhancements” on page 132
- “Configuring Grid Link Data Source for `Dept1_Cluster1`” on page 135
- “Configuring SDP-Enabled JDBC Drivers for `Dept1_Cluster1`” on page 141

Administering Oracle Solaris 11 Boot Environments

When the Oracle Solaris OS is first installed on SuperCluster M8 and SuperCluster M7, a boot environment is created. You can use the `beadm(1M)` utility to create and administer additional boot environments on your SuperCluster system.

After your SuperCluster system 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.

Related Information

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 system. If needed, administrators have the option of booting a backup boot environment.

These 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 system 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 can 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 151.

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`

Related Information

- "Create a Boot Environment" on page 148
- "Mount to a Different Build Environment" on page 150
- "Reboot to the Original Boot Environment" on page 151
- "Create a Snapshot of a Boot Environment" on page 151
- "Remove Unwanted Boot Environments" on page 152

▼ 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 the 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
```

Related Information

- “Advantages to Maintaining Multiple Boot Environments” on page 147
- “Mount to a Different Build Environment” on page 150
- “Reboot to the Original Boot Environment” on page 151
- “Create a Snapshot of a Boot Environment” on page 151
- “Remove Unwanted Boot Environments” on page 152

▼ 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:~#
```

Related Information

- “Advantages to Maintaining Multiple Boot Environments” on page 147
- “Create a Boot Environment” on page 148
- “Reboot to the Original Boot Environment” on page 151
- “Create a Snapshot of a Boot Environment” on page 151
- “Remove Unwanted Boot Environments” on page 152

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:~#
```

Related Information

- “Advantages to Maintaining Multiple Boot Environments” on page 147
- “Create a Boot Environment” on page 148
- “Mount to a Different Build Environment” on page 150
- “Create a Snapshot of a Boot Environment” on page 151
- “Remove Unwanted Boot Environments” on page 152

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. Type 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.

Related Information

- “Advantages to Maintaining Multiple Boot Environments” on page 147
- “Create a Boot Environment” on page 148
- “Mount to a Different Build Environment” on page 150
- “Reboot to the Original Boot Environment” on page 151
- “Remove Unwanted Boot Environments” on page 152

▼ 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:~#
```

Related Information

- “Advantages to Maintaining Multiple Boot Environments” on page 147
- “Create a Boot Environment” on page 148
- “Mount to a Different Build Environment” on page 150
- “Reboot to the Original Boot Environment” on page 151

- “Create a Snapshot of a Boot Environment” on page 151

Administering DISM

These topics describe how to use DISM:

- “DISM Restrictions” on page 155
- “Disable DISM” on page 156

Related Information

- “Administering Storage Servers” on page 157
- “Understanding SuperCluster Software” on page 15

DISM Restrictions

DISM is not supported for use on SuperCluster M8 and SuperCluster M7 Oracle Solaris environments in instances other than the ASM instance. The use of DISM on the system 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. 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 ASM by default. To prevent this DISM issue when creating Oracle Solaris 11 instances, disable DISM. For more information see “Disable DISM” on page 156.

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`

Related Information

- “Disable DISM” on page 156

Disable DISM

DISM is not supported for use on SuperCluster M8 and SuperCluster M7 Oracle Solaris environments in instances other than the Oracle ASM instance. For more information, see "DISM Restrictions" on page 155.

Note - Do not disable the use of ASM 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**
  - **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;
```

Related Information

- "DISM Restrictions" on page 155

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 “Understanding Storage Servers” section of the *Oracle SuperCluster M8 and SuperCluster M7 Overview Guide*.

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 systems.

- “Monitor Write-Through Caching Mode” on page 157
- “Shut Down or Reboot a Storage Server” on page 159
- “Drop a Storage Server” on page 161

Related Information

- *Oracle Exadata Storage Server Software User's Guide* for additional information about the Oracle ASM disk repair timer
- “Understanding SuperCluster Software” on page 15

▼ 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
   ```

   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
```

Related Information

- “Shut Down or Reboot a Storage Server” on page 159
- “Drop a Storage Server” on page 161

▼ Shut Down or Reboot a Storage Server

When performing maintenance on storage servers, it might 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 enable 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"
```

Related Information

- "Monitor Write-Through Caching Mode" on page 157
- "Drop a Storage Server" on page 161

▼ 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 159 for additional information.

Related Information

- “Monitor Write-Through Caching Mode” on page 157
- “Shut Down or Reboot a Storage Server” on page 159

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 — The compute server CPU and memory resources that are initially allocated during a SuperCluster installation.

C

CFM — Cubic feet per minute.

Cisco Catalyst Ethernet switch and Cisco Nexus Ethernet switch — Provides the SuperCluster management network. Referred to in this documentation using the shortened name “Ethernet management switch.” See also Ethernet management switch.

CMIOU — CPU, memory, and I/O unit. Each CMIOU contains 1 CMP, 16 DIMM slots, and 1 I/O hub chip. Each CMIOU also hosts an eUSB device. Only SPARC M7 CMIOUs can be used in SuperCluster M7 and only SPARC M8 CMIOUs can be used in SuperCluster M8.

COD Capacity on Demand.

compute server Shortened name for the SPARC server (SPARC M7-8 or SPARC M8-8), a major component of SuperCluster system. See also SPARC M7-8 server and SPARC M8-8.

D

Database Domain The domain that contains the SuperCluster database.

DB Oracle Database.

DCM Domain configuration management. The reconfiguration of boards in PDomains for Enterprise-class systems. 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 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.

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 and Cisco Nexus Ethernet switch.

eUSB Embedded USB. A flash-based drive designed specifically to be used as a boot device. An eUSB does not provide storage for applications or customer data.

expansion rack Shortened name for optional Oracle Exadata Storage Expansion Racks (up to 17) that can be added to a SuperCluster system. See also Oracle Exadata Storage Expansion Rack.

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.

G

GB Gigabyte. 1 gigabyte = 1024 megabytes.

GbE Gigabit Ethernet.

GNS Grid Naming Service.

H

HCA Host channel adapter.

HDD Hard disk drive. In Oracle Solaris OS output, HDD can refer to hard disk drives or SSDs.

I

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.

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.

IPMI Intelligent Platform Management Interface.

IPMP IP network multipathing.

iSCSI Internet Small Computer System Interface.

K

KVMS Keyboard video mouse storage.

L

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.

leaf switch Two of the IB switches are configured as leaf switches, the third is configured as a spine switch. See also IB switch.

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.

OCM Oracle Configuration Manager.

ONS Oracle Notification Service.

Oracle ASM Oracle Automatic Storage Management. A volume manager and a file system that supports Oracle databases.

Oracle Exadata Storage Expansion Rack Optional expansion racks that can be added to SuperCluster 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 M7 and Oracle SuperCluster M8 Name of the SuperCluster engineered systems. Referred to in this documentation using the shortened name "SuperCluster M7" and SuperCluster M8". See also SuperCluster M7.

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 and Oracle ZFS ZS5-ES storage appliance Provides SuperCluster systems 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 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 into one or more logical groups (see also PDomain).

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.

SER MGT The serial management port on an SP. See also SP.

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 systems, SPs are located on the compute servers, storage servers, ZFS storage appliance controllers, and IB switches. See also Oracle ILOM.

SPARC M7-8 server and SPARC M8-8 A major component of SuperCluster M7 and SuperCluster M8 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 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 CMIOUs, memory controllers, and DIMMs. 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 server Storage servers in SuperCluster systems.

Sun Datacenter InfiniBand Switch 36 Interconnects SuperCluster 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 M7 Shortened name for Oracle SuperCluster M7. See also Oracle SuperCluster M7 and Oracle SuperCluster M8.

SuperCluster M8 Shortened name for Oracle SuperCluster M8. See also Oracle SuperCluster M7 and Oracle SuperCluster M8.

T

TCP Transmission Control Protocol.

TNS Transparent Network Substrate.

TPM Trusted platform module.

U

UPS Interruptible 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 disk shelf A component of the ZFS storage appliance that contains the storage. The ZFS disk shelf is controlled by the ZFS storage controllers. See also ZFS storage appliance and ZFS storage controller.

ZFS storage appliance Shortened name for Oracle ZFS Storage ZS3-ES and ZS5-ES storage appliances. See also Oracle ZFS ZS3-ES and Oracle ZFS ZS5-ES storage appliance.

ZFS storage controller Servers in the Oracle ZFS storage appliance that manage the storage appliance. See also ZFS storage appliance.

Index

D

E

G

I

L

M

O