

Oracle® Enterprise Manager Ops Center

Configure a Secondary Service Domain

12c Release 3 (12.3.0.0.0)

E60039-01

June 2015

This guide provides an end-to-end example for how to use Oracle Enterprise Manager Ops Center.

Introduction

This document focuses on using the latest features introduced in Oracle Enterprise Manager Ops Center 12.2 to create a highly available environment using a dual I/O configuration as well as the application of several best practices.

In earlier versions of Oracle Enterprise Manager Ops Center, it was only possible to create and manage the control domain which owned all the physical I/O devices and also provided virtual devices to the guest domains. It was impossible to configure, deploy, manage, or discover a secondary domain. An advantage of having a secondary service domain and creating a resilient environment is that maintenance on the control domain (such as OS upgrades, patching, and increasing memory) happens without affecting any running guest domains.

This example describes how to use Oracle Enterprise Manager Ops Center 12.2.1 to deploy:

- A control domain, also referred to as primary domain and denoted as `primary` wherever applicable in the Oracle Enterprise Manager Ops Center UI.
- A secondary service domain, also referred to as alternate I/O domain or redundant I/O domain. The secondary service domain is a root domain, that is a domain configured with a PCIe root complex with a configuration similar to the control domain.
- A high available (HA) guest domain, the guest domain will use services from both, the control domain and secondary service domain for improving resilient.

The scenario in this document uses an Oracle SPARC server and Oracle VM Server for SPARC to create the control domain, the secondary service domain, and the guest domain.

What You Will Need

The following resources are used for this example:

- A server hosting Oracle Enterprise Manager Ops Center, the Enterprise controller and the co-located Proxy controller.
- Oracle Enterprise Manager Ops Center 12.2.1 installed. Enterprise Controller (EC) with co-located Proxy Controller (PC), and an embedded database running on Solaris 11.1 SRU 19.6.

- Oracle Solaris 11 Software Update Library. This example uses Oracle Solaris 11.1 SRU 19.6 content
- Oracle SPARC T5-2 ILOM discovered and managed in Oracle Enterprise Manager Ops Center.
- Sun ZFS Storage 7120 appliance discovered in Oracle Enterprise Manager Ops Center.
- Fiber Channel (FC) LUNs available from the Sun ZFS Storage 7120 appliance.
- NFS share from the Sun ZFS Storage 7120 appliance.
- Free IP addresses allocated.
- VLAN tagged network created in Enterprise Manager Ops Center.
- Oracle SPARC T5-2 for provisioning Oracle VM for SPARC Server, version 3.1.1. The server requires a free PCIe bus for configuring the secondary service domain.
- Virtualization Admin, and Profile and Plan Admin role to perform the actions.

Preview of the Configuration for the Resilient Environment

An Oracle SPARC T5-2 Server is used to create the service domains and the HA guest domain. A summary of how the domains and guest will be configured is explained in the following section, followed by a table which includes the PCIe buses used, and then an image depicting the layout.

Due to lab constraints, I/O configuration options were limited; therefore, setting up a best practice environment in all areas was not possible. However, the document mentions areas in which best practices could be applied.

Configuration Summary

The primary and secondary domains will be configured similarly, with:

- Connections to a Public and an Admin network
- Link aggregation for the public Network
- Connections to an Sun 7120 Storage Appliance
- Oracle Solaris 11.1 SRU 19.6 operating system (OS)
- Oracle VM Server for SPARC 3.1.1
- 4 whole-core CPUs for the control domain, and 2 whole-cores CPUs for the secondary domain
- 16 GB of memory for the control domain and 4GB of memory for the secondary domain.

Each domain will have access to different PCIe buses for resilience.

For best practice, consider having multipathed LUNs from the storage to the primary and secondary domains.

The guest domain will have:

- Two network connections for the public network, one from each domain
- Two network connections for the admin network, one from each domain
- Redundant storage I/O paths provided by primary and secondary domains

- Oracle Solaris s11.1 SRU 19.6 operating system
- 4 core CPUs
- 4 GB of memory

Configuration Table

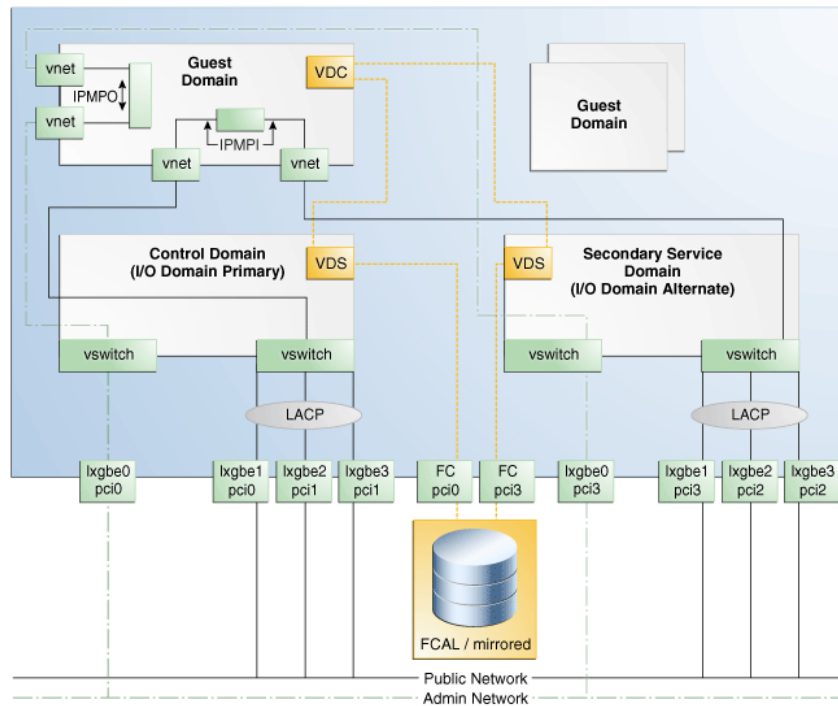
The following table shows how PCIe buses are used for each domain:

Domain	Configuration	Network	Storage
Control domain	Oracle Solaris 11 SRU 19.6 4 CPUs 16 GB of memory	4 Ethernet connections: ixgbe0 (pci0) connected to admin network. ixgbe1(pci0), ixgbe2(pci1), and ixgbe3(pci1) connected to the public network, all aggregated.	FC LUN (pci0) provided by a Sun ZFS Storage 7120 appliance
Secondary service domain	Oracle Solaris 11 SRU 19.6 2 CPUs 4 GB of memory	4 Ethernet connections: ixgbe0 (pci3) connected to admin network. ixgbe1(pci3), ixgbe2(pci2), and ixgbe3(pci2) connected to the public network, all aggregated.	FC LUN (pci3) provided by a Sun ZFS Storage 7120 storage appliance
Guest domain	Oracle Solaris 11 SRU 19.6 4 CPUs 4 GB of memory	4 virtual network connections: The admin network has two virtual networks (vnets) one from the control domain and one from the secondary domain which is IP Multipathed (IPMP0). The public network has two network connections, one from the control domain and one from the secondary domain which will be IP Multipathed (IPMP1).	Redundant I/O paths to virtual disk server in control domain and secondary domain.

Configuration Illustration

[Figure 1](#) illustrates how the Oracle SPARC server will be configured to allow for resilient logical domains.

Figure 1 Resilient Environment for Logical Domains



Steps for Creating the Environment

To following steps define the procedure to create a resilient logical domain environment, each step is discussed in detail:

1. [Configuring and Provisioning the Control Domain](#)
2. [Adding a PCI bus to the Control Domain](#)
3. [Aggregating Networks and Creating a Virtual Switch to the Control Domain](#)
4. [Configuring and Provisioning the Secondary Service Domain](#)
5. [Creating the HA Guest Domain](#)

Configuring and Provisioning the Control Domain

The first domain in Oracle VM Server for SPARC is always the control domain. The main purpose of this domain is to provide an environment where the Oracle VM Server for SPARC Manager runs. It also provides virtual console services as well as virtual services to guest domains.

The following steps define the procedure to successfully configure and provision Oracle VM Server for SPARC 3.1.1 version:

1. [Creating an OS Provisioning Profile](#)
2. [Creating an OS Configuration Profile](#)
3. [Creating the Deployment Plan and Provisioning the Control Domain](#)

Creating an OS Provisioning Profile

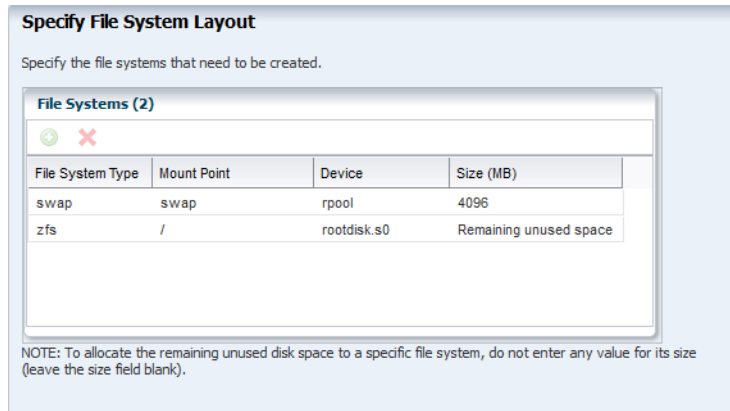
The OS provisioning profile (OSP) allows you to set the operating system parameters, including file system layout, time zones, passwords, and other parameters.

To create the OSP in Oracle Enterprise Manager Ops Center:

1. Select **Plan Management** section and expand **Profiles and Policies** in the Navigation pane.
2. Select **OS Provisioning** profile and click **Create Profile** in the Actions pane.
3. Enter the following details in the **Create OS Profile - OS Provisioning** wizard and then click **Next** to specify the OSP parameters.:
 - Name and description of the profile.
 - Select Oracle VM Server for SPARC as the Subtype.
4. Select the following parameters in the Specify OSP Parameters step and then click **Next** to specify the OS setup:
 - Select the Oracle VM Server for SPARC version along with the required Oracle Solaris 11 OS and SRU. In this example, select the version Solaris 11.1 SRU 19.6.0 (LDom 3.1.1.0).
 - Select solaris-small-server from the Systems Software Group.
5. Retain the default values for the OS setup parameters or edit the language, time zone, and NFS4 Domain values for your environment. Enter the root password and confirm the password. Click **Next** to specify the user account for Oracle Solaris 11 OS.

You can also choose to select the **Manual Net Boot** and **Save NVRAMRC** options. For more information about this, refer to the [Related Articles and Resources](#) section.

6. Root login is not enabled in Oracle Solaris 11 OS. Create a user account to SSH to the OS after provisioning. Provide a user name and password for the account.
Click **Next** to specify whether you want to use iSCSI disks for provisioning Oracle VM Server for SPARC.
7. Do not select the option to use iSCSI disk as this scenario does not involve the use of iSCSI disk for provisioning Oracle VM Server for SPARC.
Click **Next** to specify the file system layout.
8. Retain the default values for the root (/) and swap file systems. You have the options to change the swap size and add more ZFS file systems.



Click **Next** to specify the name service.

9. If you have a naming service in place, select the appropriate one and provide the setup details. In this example, select **None** for the naming service.

If you have any naming service in your setup, refer to the help in the wizard or the [Related Articles and Resources](#) section for information about specifying the naming services.

Click **Next** to view the summary of the parameters selected for the profile.

10. Review the parameters selected for the profile and click **Finish** to create the OS provisioning profile.

Creating an OS Configuration Profile

The OS configuration profile (OSC) is used to configure the resources for the control domain, like CPU cores, as well as detach unused buses, set network options, and other options.

To create the OSC in Oracle Enterprise Manager Ops Center:

1. Select the **Plan Management** section in the Navigation pane and expand **Profiles and Policies**.
2. Select OS Configuration and click **Create Profile** in the Actions pane.
3. Enter the following details in the **Create Profile - OS Configuration** wizard:
 - Name and description of the profile.
 - Select Oracle VM Server for SPARC as the Subtype.

Click **Next** to specify the control domain parameters.

4. Specify the resources that you want to assign to the control domain. The remaining resources are available for the logical domains. The following are new additions to the 12.2 version:
 - Allocate CPU resources as Whole Cores. All threads are allocated from the CPU to the control domain.
 - Virtual Disk Server name can be altered from the default.
 - Detach unused buses, PCIe buses that are not required for use by the control domain are unassigned and can be used for other domains.

- Enable SR-IOV, permits virtual functions to be shared from the PCIe card if the card permits.

In this scenario, select the following configuration parameters for the control domain:

- **Oracle VM Server Version:** Select the Solaris 11.1 SRU 19.6.0 (LDom 3.1.1.0) version to be installed.
- **CPU Model:** Use Whole-core to allocate the CPU resource in cores.
- **CPU Cores:** Enter four CPU core to be allocated to the control domain.
- **Memory:** Enter 16 GB as the memory required for control domain.
- **Virtual Console Port Range:** Enter the range between 5000 to 6000.

Do not provide any values for **Requested Crypto Units** and **Max CPU Cores**. Retain or edit the default name of the virtual disk server. Select the option **Detach Unused Buses**.

Click **Next** to specify the OS management details.

5. Select the option **Enable Multiplexed I/O** so that you can associate block storage libraries such as SAN and iSCSI for storage with the control domain.

Click **Next** to specify the networking details.

6. Select **None** as the networking option for Oracle VM Server for SPARC.

Click **Next** to specify the networking details for Oracle VM Server for SPARC.

7. Select the network interface to use for OS provisioning. Select the **Controller** that hosts the network interfaces and the corresponding network interface. The **Address Allocation** is Use Static IP by default and cannot be modified.

Select **Auto** to create a virtual switch automatically for the network connection to the control domain. The virtual switch is created in the default format. For example, the network *192.0.2.0/24*, the virtual switch is created as *192.0.2.0_24*.

Note: As [Figure 1](#) shows, the networks of the control domain are aggregated. However, as the option **Detach Unused Buses** is selected, this would mean `pci1` would not be owned by the control domain, so aggregation as required (for using `pci0` and `pci1`) is not possible. Aggregation will happen later when `pci1` is re-added and owned by the control domain.

Click **Next** to view the summary of the parameters selected for OS configuration.

8. Review the parameters and click **Finish** to create the OS configuration profile.

Creating the Deployment Plan and Provisioning the Control Domain

The Deployment Plan combines the OSP and OSC to provision and configure the control domain.

To create the deployment plan in Oracle Enterprise Manager Ops Center:

1. Select **Plan Management** section in the Navigation pane.
2. Expand **Deployment Plans** and select **Provision OS**.

3. Click **Create Deployment Plan** in the Actions pane.
4. Enter the following information for the deployment plan:
 - Name and description for the plan.
 - Select **Stop at Failure** as the Failure Policy.
 - Select the corresponding OS provisioning and configuration profile created for the control domain.
5. Click **Save** to create the deployment plan.

Once the Deployment plan has been created, the next task is to provision the control domain. To do this selects the ILOM server from the assets menu and select **Install Server** from the Actions panel. A job will be created, similar to the one below after completing the steps:

Task	Target of the task	Result	Elapsed Time
▲ t5-2-tvp540	t5-2-tvp540	Flow execution is in progress	00 hr, 00 min, 08 sec
▲ t5-2-tvp540			-
▲ Provision OS			-
✓ Setup Install Client	10.163.212.79	Setup Clients task is successfully completed. (32519)	Less than a second
✓ Configure Virtualizat	10.163.212.79	Configured the Oracle VM Server for SPARC as per the values set in Oracle VM Server for SPARC profile.	Less than a second
▬ Configure Install Ser	10.163.212.79	Created Auto Install client t5-2-tvp540-c' having macaddress 00:10:E0:2A:3C:E4.	00 hr, 00 min, 06 sec
▬ Clear ZoneController	t5-2-tvp540		-
▬ Reset to Factory De	t5-2-tvp540-		-
▬ Network Install	t5-2-tvp540-		-
▬ Monitor OS Installat	10.163.212.79		-
▬ Cleanup Client	10.163.212.79		-
▲ Configure OS			-
▬ Setup Configure Cle	10.163.212.79		-
▲ Download Agent Imc			-
▬ Download Image PC			-
▬ Configure Target Sei	10.163.212.79		-
▬ Monitor OS Configur	10.163.212.79		-
▬ Verify Agent Install	10.163.212.79		-
▬ Cleanup OSC Client	10.163.212.79		-

Once the Provisioning job has completed successfully, check if the parameters have been set. This can be seen from the UI by selecting the Summary tab.

The screenshot displays the Oracle VM Server Summary page for the instance t5-2-tvp540-c. The page is divided into several sections:

- Name:** t5-2-tvp540-c
- Description:** Oracle VM Server for SPARC
- CPU Info:** Oracle SPARC T5, 2 socket(s), 32 core(s), 256 thread(s)
- Available CPU Threads:** 240 out of 256
- Available CPU Cores:** 30 out of 32
- Available Memory (RAM):** 251 GB out of 255 GB
- Running Time:** 1 day(s), 15:51 (HH:MM)
- Oracle VM Server Status:** Healthy
- Reachable:** Yes
- Server Pool:** -
- Oracle VM Server Version:** 3.1.1.0.5
- Tags:** legacy.tags, idom virtualizationcontroller
- Control Domain Specification (highlighted):**
 - CPU Model:** Whole-Core
 - CPU Cores:** 2
 - Memory Size (MB):** 4096
- Max CPU Cores:** -
- CPU Utilization:** 0%

Since the option **Detach unused buses** was selected in the OSC, the control domain only owns the *pci0* bus. You can verify it in the UI by selecting the I/O Resources tab, or in the CLI by logging in to the control domain and using the *ldm ls-io* command. From the CLI the output will look similar to the following:

```
root@t5-2-tvp540-c:/var/tmp# ldm ls-io
NAME                                     TYPE  BUS      DOMAIN  STATUS
----                                     -
pci_1                                    BUS   pci_1
pci_0                                    BUS   pci_0    primary
pci_3                                    BUS   pci_3
pci_2                                    BUS   pci_2
/SYS/MB/PCIE5                           PCIE  pci_1    UNK
/SYS/MB/PCIE6                           PCIE  pci_1    UNK
/SYS/MB/PCIE7                           PCIE  pci_1    UNK
/SYS/MB/PCIE1                           PCIE  pci_0    primary OCC
/SYS/MB/SASHBA0                          PCIE  pci_0    primary OCC
/SYS/MB/NET0                             PCIE  pci_0    primary OCC
/SYS/MB/PCIE8                           PCIE  pci_3    UNK
/SYS/MB/SASHBA1                          PCIE  pci_3    UNK
/SYS/MB/NET2                             PCIE  pci_3    UNK
/SYS/MB/PCIE2                           PCIE  pci_2    UNK
/SYS/MB/PCIE3                           PCIE  pci_2    UNK
/SYS/MB/PCIE4                           PCIE  pci_2    UNK
/SYS/MB/NET0/IOVNET.PF0                 PF    pci_0    primary
/SYS/MB/NET0/IOVNET.PF1                 PF    pci_0    primary
root@t5-2-tvp540-c:/var/tmp#
```

As shown above, the control domain now only owns *pci_0*.

Adding a PCI bus to the Control Domain

As seen in the output generated by the command *ldm ls-io*, *pci1* needs to be added to the control domain, which will provide another network interface and allow network aggregation across the *pci0* and *pci1* buses.

Currently, the network devices on the control domain are *ixgbe2* (*net0*) and *ixgbe3* (*net1*). You can verify it from the CLI by using the *dladm show-phys* command.

```
root@t5-2-tvp540-c:/var/tmp# dladm show-phys
LINK      MEDIA      STATE      SPEED  DUPLEX  DEVICE
net0      Ethernet  up         100    full    ixgbe2
net1      Ethernet  unknown    0      unknown ixgbe3
net10     Ethernet  up         100    full    vsw0
root@t5-2-tvp540-c:/var/tmp#
```

The device numbers might seem strange at first, as seeing *ixgbe0* and *ixgbe1* would be more logical. The reason for the numbering is because of the way Oracle Solaris maps instance names to physical devices in */etc/path_to_inst*. Oracle Solaris keeps a history about all devices and how they are mapped.

PCIe buses cannot be dynamically added to the control domain. Therefore, the control domain is set in delayed reconfiguration mode. In this mode, operation changes take effect after the next reboot.

Certain actions, usually one-time actions, cannot be performed from the UI, such as adding a PCIe bus. These actions need to be performed from the CLI by using *ldm* commands on the control domain. A new feature in Ops Center 12.1 and above allows using CLI or BUI in parallel to perform administration tasks.

To add a new PCIe bus to the control domain:

1. Initiate the delayed reconfiguration state in the control domain by using the `ldm start-reconf` command.

```
root@t5-2-tvp540-c:/var/tmp# ldm start-reconf
No LDom specified

Usage:
    ldm start-reconf <ldom>

root@t5-2-tvp540-c:/var/tmp# ldm start-reconf primary
Initiating a delayed reconfiguration operation on the primary domain.
All configuration changes for other domains are disabled until the primary
domain reboots, at which time the new configuration for the primary domain
will also take effect.
```

2. Add the PCIe bus to the control domain by using the `ldm add-io pci_1 primary` command, in this scenario the PCIe bus name `pci1` is added.

```
root@t5-2-tvp540-c:/var/tmp# ldm add-io pci_1 primary
-----
Notice: The primary domain is in the process of a delayed reconfiguration.
Any changes made to the primary domain will only take effect after it reboots.
-----
root@t5-2-tvp540-c:/var/tmp# █
```

3. Reboot the primary domain for changes to take effect.

After the reboot, the control domain is now the owner of `pci1` and `pci0`, as shown below.

```
root@t5-2-tvp540-c:/var/tmp# ldm ls-io
NAME                                     TYPE  BUS      DOMAIN  STATUS
----
pci_1                                   BUS   pci_1    primary
pci_0                                   BUS   pci_0    primary
pci_3                                   BUS   pci_3
pci_2                                   BUS   pci_2
/SYS/MB/PCIE5                           PCIE  pci_1    primary UNK
/SYS/MB/PCIE6                           PCIE  pci_1    primary UNK
/SYS/MB/PCIE7                           PCIE  pci_1    primary UNK
/SYS/MB/PCIE1                           PCIE  pci_0    primary OCC
/SYS/MB/SASHBA0                         PCIE  pci_0    primary OCC
/SYS/MB/NET0                            PCIE  pci_0    primary OCC
/SYS/MB/PCIE8                           PCIE  pci_3    UNK
/SYS/MB/SASHBA1                         PCIE  pci_3    UNK
/SYS/MB/NET2                            PCIE  pci_3    UNK
/SYS/MB/PCIE2                           PCIE  pci_2    UNK
/SYS/MB/PCIE3                           PCIE  pci_2    UNK
/SYS/MB/PCIE4                           PCIE  pci_2    UNK
/SYS/MB/NET0/IOVNET.PF0                 PF    pci_0    primary
/SYS/MB/NET0/IOVNET.PF1                 PF    pci_0    primary
root@t5-2-tvp540-c:/var/tmp# █
```

The network devices available on the control domain are now `xgbe2` (`net0`), `ixgbe3` (`net1`), `xgbe0` (`net4`), and `xgbe1` (`net5`):

```
root@t5-2-tvp540-c:/etc/default# dladm show-phys
LINK  MEDIA  STATE  SPEED  DUPLEX  DEVICE
net8   Ethernet  up     10     full    usbecm2
net0   Ethernet  up     100    full    ixgbe2
net4   Ethernet  unknown 0     unknown ixgbe0
net1   Ethernet  unknown 0     unknown ixgbe3
net5   Ethernet  unknown 0     unknown ixgbe1
net10  Ethernet  up     100    full    vsw0
root@t5-2-tvp540-c:/etc/default# █
```

This yields the four network interfaces that are required. In the next step, three of the interfaces for the public network will be aggregated.

Aggregating Networks and Creating a Virtual Switch to the Control Domain

Link aggregation increases bandwidth, provides degradation as failure occurs, and increases availability. In addition, it provides network redundancy by load-balancing traffic across all available links. If a link fails, the system automatically load-balances traffic across all remaining links as specified when creating the policy.

For redundancy, network links are used from the pci0 and pci1 buses. These are aggregated for better throughput and added redundancy.


The following steps define the procedure to create the aggregated links and the virtual switch:

1. [Creating Aggregated Links](#)
2. [Creating a Virtual Switch](#)

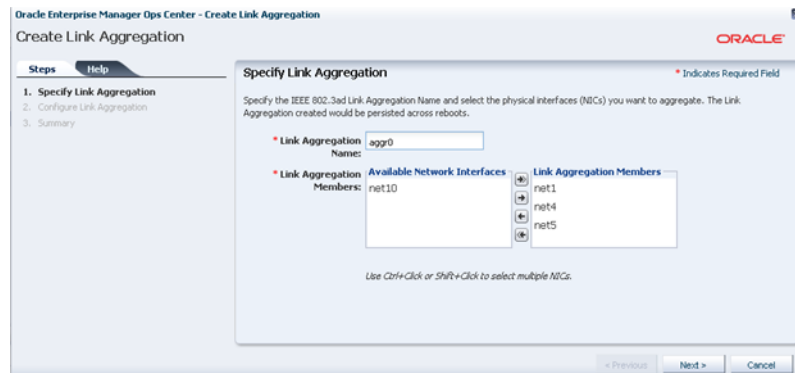
Creating Aggregated Links

To create aggregated link from the UI:

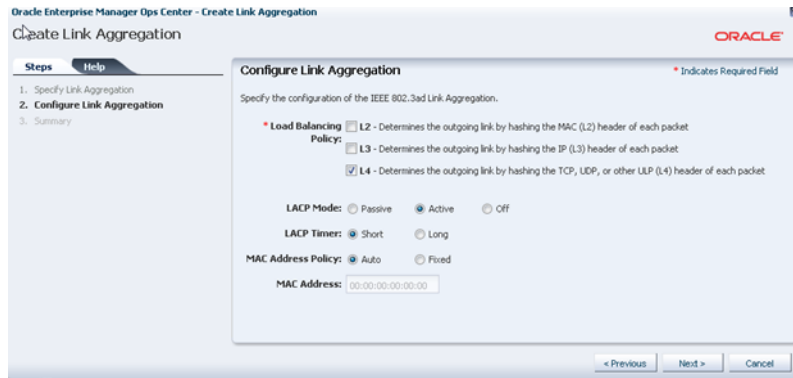
1. Expand Assets in the Navigation pane and select the OS of the control domain.
2. Select the Networks tab in the center pane, then the Link Aggregations sub tab, and click the **Create Link Aggregation** icon.



3. Specify a name for the link aggregation and move the network interfaces to the link aggregation members list, as shown below.



4. Configure the aggregation links, as shown in the image below.



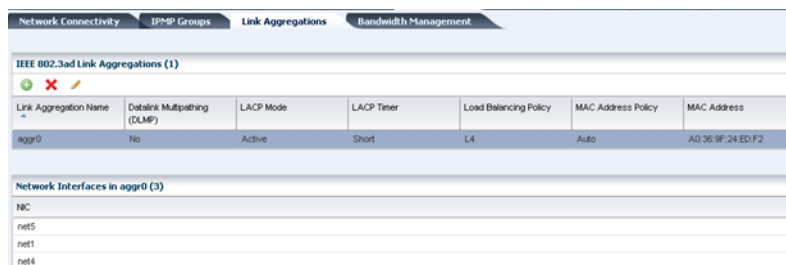
5. Review the summary and click **Finish**.

You can use the UI or the command line to confirm that the correct changes have taken place.

From the command line, login to the control domain and run the following command:

```
root@t5-2-tyvp540-c:/etc/default# dladm show-link
LINK          CLASS      MTU      STATE    OVER
net8          phys      1500    up       --
net0          phys      1500    up       --
net4          phys      1500    up       --
net1          phys      1500    up       --
net5          phys      1500    up       --
net10         phys      1500    up       --
aggr0         aggr      1500    up       net5 net1 net4
root@t5-2-tyvp540-c:/etc/default#
```

As shown in the output above, an aggregated link called `aggr0` has been created, consisting of `net 1`, `net 4`, and `net 5`. This information can also be seen from the UI, as depicted in the image below:




Creating a Virtual Switch

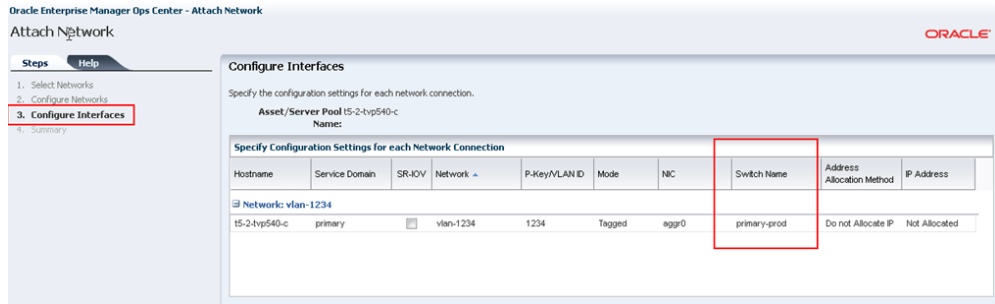
At this stage, a virtual switch (`vsw`) for the aggregation link `aggr0` needs to be created for the control domain. This virtual switch connects the guest domain to the public and admin networks.

To create a virtual switch:

1. Expand Assets in the Navigation pane and select the control domain.
2. Select the Networks tab in the center pane and click **Attach Network**.
3. Select the public network, in this example is `vlan-1234`.



4. Configure the aggregation links, as shown in the image below.



5. Review the summary and click **Finish**.

You can use the command line to view the newly created virtual switch for the aggregated links (aggr0) by using the `ldm ls-services` command.

```

root@t5-2-tvp540-ci-# ldm ls-services
VSC
NAME          LDOM          PORT-RANGE
primary-vec0  primary       5000-6000

VSW
NAME          LDOM          MAC          NET-DEV    ID  DEVICE  LINKPROP  DEFAULT-VLAN-ID  PVID  VID
10.163.208.0_20  primary      00:14:4f:fb:2e:c5  net0      0   switch0  phys-state  1                1
primary-prod    primary      00:14:4f:2e:32:6e  aggr0     1   switch1  phys-state  1                1234

VDS
NAME          LDOM          VOLUME      OPTIONS    MPGROUP    DEVICE
primary-vds0  primary

root@t5-2-tvp540-ci-#

```

Configuring and Provisioning the Secondary Service Domain

As mentioned before the secondary service domain has various names, such as alternate I/O domain or redundant I/O Domain.

By design, the secondary domain will have some physical I/O devices assigned, which may be a PCIe bus root complex, a PCI device, or a SR-IOV (Single Root I/O Virtualization) virtual function.

Within Oracle Enterprise Manager Ops Center, when creating a secondary domain, the terms physical I/O domain and root domain are used. The physical I/O domain maps PCI device end points, and the root domain maps PCIe buses, which also has an option to create SR-IOV functions.

In this example, a Root Domain is created by assigning PCIe buses. As seen, the Oracle SPARC T5-2 server has two free PCIe buses that have not been assigned to the primary Domain (pci2 and pci3). These PCIe buses are used to create the secondary domain (root domain).

The following steps define the procedure to successfully configure a secondary service domain:

1. [Creating the Logical Domain Profile](#)
2. [Applying the Logical Domain Profile](#)

3. [Provisioning the Secondary Domain OS](#)
4. [Creating Aggregation Link and Creating Virtual Switches](#)

Creating the Logical Domain Profile

To create the secondary service domain in Oracle Enterprise Manager Ops Center, start by creating a Logical Domain Profile.

To create the Logical Domain profile:

1. Select **Plan Management** section and expand **Profiles and Policies** in the Navigation pane.
2. Select **Logical Domain** profile and click **Create Profile** in the Actions pane.
3. Enter the following details in the **Create Profile - Logical Domain** wizard and then click **Next**.
 - Name and description of the profile.
 - Retain the option to create a deployment plan for this profile.
 - Select Root Domain as the Subtype.
4. Provide a name for the secondary service domain and click **Next**.
5. Enter the configuration details for the secondary service domain and click **Next**. In this example, you use Whole-Core CPU model with 2 CPU cores and 4 GB of memory.
6. Specify the number of PCIe buses to be assigned to the secondary I/O and click **Next**. In this example, you assign two buses.
7. Select a library to store the metadata and the secondary service domain's storage, then click **Next**. In this example, select the local filesystem library and enter the disk size as 20 GB.

The metadata is on the local disk, `file://guest`. This is fine for the secondary domain as it will not be migrated. Only the logical domains and their guests will get migrated; this makes it mandatory to have the metadata on a shared storage.
8. Skip the step to specify networks and click **Next**. This step is optional for this example.
9. Review the parameters selected for the profile and click **Finish** to create the profile.

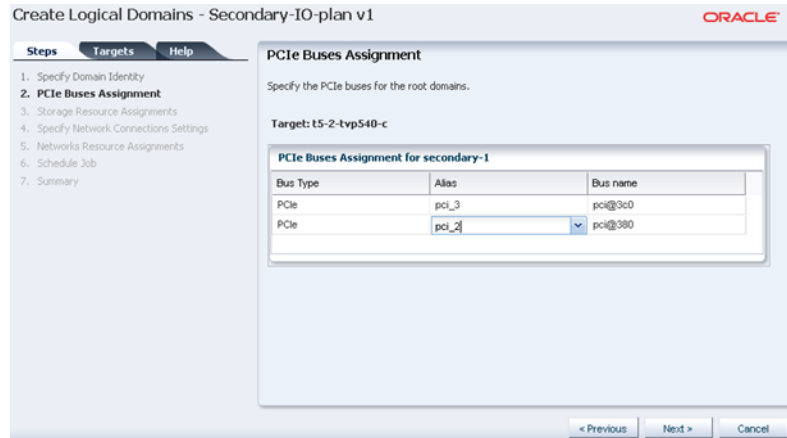
Applying the Logical Domain Profile

To apply the logical domain profile in Oracle Enterprise Manager Ops Center:

1. Select **Plan Management** section in the Navigation pane and expand **Deployment Plans**.
2. Expand **Create Logical Domains** and select the newly create plan from the list.
3. Click **Apply Deployment Plan** in the Actions pane.
4. In the **Select Targets** window, select the control domain.

Click **Add to Target List** to move the selected target Oracle VM Server to the Target List, then click **Next**.
5. Specify the name for the secondary service domain, then click **Next**.

- Select the PCIe buses that will be assigned to the secondary domain and then click **Next**. In this example, you assign `pci1_3` and `pci_2`.



- Retain or edit the default name of the virtual disk server (vds), then click **Next**.
- In the profile to create root domain, the network connections details were not provided. Therefore, the Network Connection Settings step and Network Resource Assignments steps are empty. For this example, skip both steps by clicking **Next**.
- Schedule the job to run now and click **Next**.
- Review the properties and click **Apply** to apply the deployment plan to create secondary domain.

Once the job completes successfully, you can use the UI or the command line to confirm that the changes.

From the command line, login to the control domain and run the following commands:

```
root@t5-2-tvp540-c:~# ldm ls
NAME          STATE   FLAGS  CONS  VCPU  MEMORY  UTIL  NORM  UPTIME
primary      active -n-cv-  UART   16    4G      0.4%  0.4%  4d 18h 40m
secondary    active -t--v- 5000   16    4G      6.2%  6.2%  2m
root@t5-2-tvp540-c:~#
```

```
root@t5-2-tvp540-c:~# ldm ls-io
NAME          TYPE    BUS      DOMAIN  STATUS
-----
pci_1         BUS     pci_1    primary
pci_0         BUS     pci_0    primary
pci_3         BUS     pci_3    secondaryIOV
pci_2         BUS     pci_2    secondaryIOV
/SYS/MB/PCIE5 PCIE    pci_1    primary  OCC
/SYS/MB/PCIE6 PCIE    pci_1    primary  EMP
/SYS/MB/PCIE7 PCIE    pci_1    primary  EMP
/SYS/MB/PCIE1 PCIE    pci_0    primary  OCC
/SYS/MB/SASHBA0 PCIE    pci_0    primary  OCC
/SYS/MB/NET0   PCIE    pci_0    primary  OCC
/SYS/MB/PCIE8 PCIE    pci_3    secondaryUNK
/SYS/MB/SASHBA1 PCIE    pci_3    secondaryUNK
/SYS/MB/NET2   PCIE    pci_3    secondaryUNK
/SYS/MB/PCIE2 PCIE    pci_2    secondaryUNK
/SYS/MB/PCIE3 PCIE    pci_2    secondaryUNK
/SYS/MB/PCIE4 PCIE    pci_2    secondaryUNK
/SYS/MB/PCIE5/IOVNET.PF0 PF       pci_1    primary
/SYS/MB/PCIE5/IOVNET.PF1 PF       pci_1    primary
/SYS/MB/NET0/IOVNET.PF0 PF       pci_0    primary
/SYS/MB/NET0/IOVNET.PF1 PF       pci_0    primary
root@t5-2-tvp540-c:~#
```

From the UI, confirm that the buses have been assigned correctly in the I/O Resources tab of the secondary domain:

Alias	Bus name	Type
pci_2	pci@300	BUS
pci_3	pci@300	BUS

Alias	Device Name	Root Domain	PCIe Bus	PCIe Slot Status	Domain
Ethernet Device (2)					
/SYSMBNET2	pci@300pci@1pci...	secondary	pci_3	-	secondary
/SYSMBPCE2	pci@300pci@1pci...	secondary	pci_2	-	secondary
Unrecognized Device Type (4)					
/SYSMBPCE3	pci@300pci@1pci...	secondary	pci_2	-	secondary
/SYSMBPCE4	pci@300pci@1pci...	secondary	pci_2	-	secondary
/SYSMBPCE8	pci@300pci@1pci...	secondary	pci_3	-	secondary
/SYSMBASAHBA1	pci@300pci@1pci...	secondary	pci_3	-	secondary

Provisioning the Secondary Domain OS

To provision the operating system to the secondary service domain, you need to create an OS provisioning (OSP) profile, create an OS configuration (OSC) profile, create a deployment plan for the OSP and OSC profiles, and apply the deployment plan to the secondary service domain.

Creating an OS Provisioning Profile

1. Select **Plan Management** section and expand **Profiles and Policies** in the Navigation pane.
2. Select **OS Provisioning** profile and click **Create Profile** in the Actions pane.
3. Enter the following details in the **Create OS Profile - OS Provisioning** wizard and then click **Next** to specify the OSP parameters.:
 - Name and description of the profile.
 - Select Logical Domain as the Subtype.
4. Complete the remaining steps according to your requirements. In the Summary screen click **Finish** to create the OSP profile.

Oracle Enterprise Manager Ops Center - Create Profile - OS Provisioning

Create Profile - OS Provisioning

Steps: 1. Identify Profile, 2. Specify OSP Parameters, 3. Specify OS Setup, 4. Specify User Account, 5. Specify iSCSI Disk Usage, 6. Specify File System Layout, 7. Specify Naming Services, 8. Summary

Summary

OS Image: Oracle Solaris 11.1 sparc (SRU 19.6.0) (A0)

Software Group: pkg://solaris/group/system/solaris-small-server

Language: English (7-bit ASCII)

Time Zone: GMT

Terminal Type: vt100

Console Serial Port: ttya

Console Baud Rate: 9600

NFS4 Domain: oracle.com

Manual Net Boot:

Solaris 11 Update Profile:

Username: jack

Full Name:

Use iSCSI Disk:

File System Type	Mount Point	Device	Size (MB)
swap	swap	rpool	4096
zfs	/	rootdisk.s0	Remaining unused space

Name Services: DNS

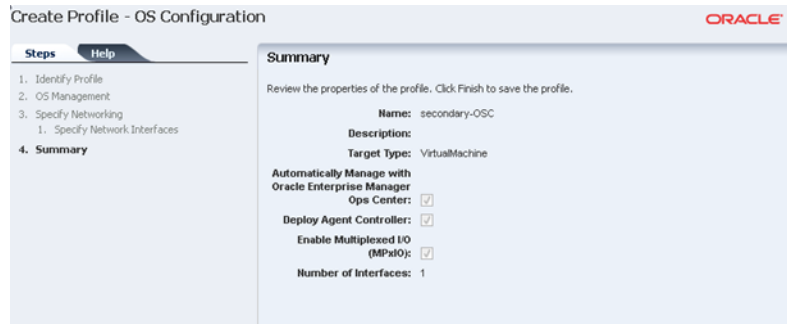
< Previous Finish Cancel

Creating an OS Configuration Profile

To create the OSC in Oracle Enterprise Manager Ops Center:

1. Select the **Plan Management** section in the Navigation pane and expand **Profiles and Policies**.
2. Select OS Configuration and click **Create Profile** in the Actions pane.
3. Enter the following details in the **Create Profile - OS Configuration** wizard:
 - Name and description of the profile.
 - Select logical domain as the Subtype.

Click **Next** to specify the logical domain parameters.
4. Complete the remaining steps according to your requirements. In the Summary screen click **Finish** to create the OSC profile.

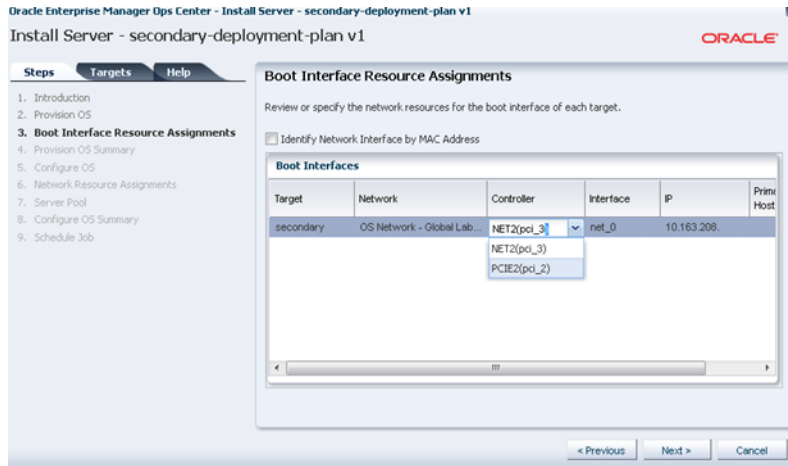


Creating the Deployment Plan and Provisioning the Secondary Domain OS

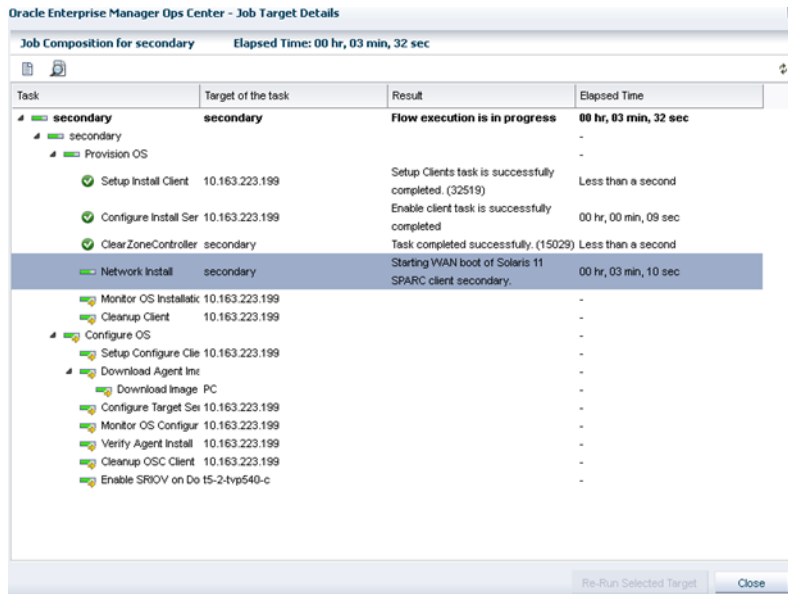
To create the deployment plan in Oracle Enterprise Manager Ops Center:

1. Select **Plan Management** section in the Navigation pane.
2. Expand **Deployment Plans** and select **Provision OS**.
3. Click **Create Deployment Plan** in the Actions pane.
4. Enter the following information for the deployment plan:
 - Name and description for the plan.
 - Select **Stop at Failure** as the Failure Policy.
 - Select the corresponding OS provisioning and configuration profile created for the secondary service domain.
5. Click **Save** to create the deployment plan.

Once the Deployment plan has been created, the next task is to apply the deployment plan to the secondary service domain.
6. Expand Assets in the Navigation pane and select the control domain.
7. Click **Create Logical Domains** in the Actions pane.
8. Select the recently created Deployment Plan and complete the Provision OS step.
9. In the Boot Interface Resource Assignment step, make sure that the correct boot interface is selected, as shown below.



10. Complete all other steps as required and then, click **Finish**. A job, similar to the one depicted below, will run.



It is possible and recommended to view how the install is progressing. To do this, log on the control domain and view the `/var/log/vntsd/secondary/console-log` file as shown below.

```

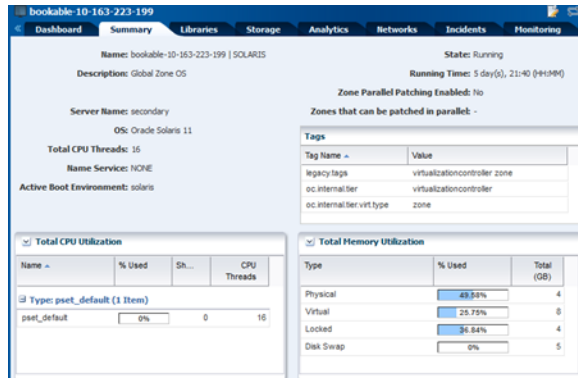
root@t5-2-tp540-c:/var/log/vntsd/secondary# tail -f console-log

Boot device: /pci@3c0/pci@1/pci@0/pci@1/network@0 File and args: - install
100 Mbps full duplex Link up
<time unavailable> wanboot info: WAN boot messages->console
<time unavailable> wanboot info: configuring /pci@3c0/pci@1/pci@0/pci@1/network@0

100 Mbps full duplex Link up
<time unavailable> wanboot progress: wanbootfs: Read 368 of 368 kB (100%)
<time unavailable> wanboot info: wanbootfs: Download complete
Tue Jun 3 11:33:51 wanboot progress: miniroot: Read 106168 of 252623 kB (42%)

```

Once the job has completed successfully, confirm the set-up is correct in the UI or the command line.



Creating Aggregation Link and Creating Virtual Switches


As with the control domain, the public network for the secondary service domain are aggregated links in order to provide resilience and better throughput.

Creating aggregated links in the secondary domain is a similar process to the one used to create those links in the control domain virtual switch (vsw). However, another virtual switch for the admin network needs to be created.

Creating Aggregated Links

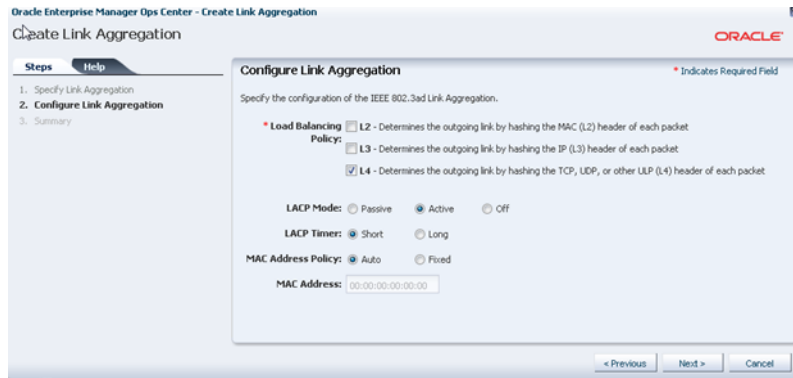
To create aggregated link from the UI:

1. Expand Assets in the Navigation pane and select the OS of the secondary domain.
2. Select the Networks tab in the center pane, then the Link Aggregations sub tab, and click the **Create Link Aggregation** icon.



3. Specify a name for the link aggregation and move the network interfaces to the link aggregation members list, as shown below.

4. Configure the aggregation links, as shown in the image below.



5. Review the summary and click **Finish**.

You can use the UI or the command line to confirm that the correct changes have taken place.


From the command line, login to the secondary domain and run the following command:

```

root@bookable-10-163-223-199:~# dladm show-aggr
LINK      MODE POLICY  ADDRPOLICY  LACPACTIVITY LACPTIMER
aggr0     trunk L4    auto        active      short
root@bookable-10-163-223-199:~# dladm show-link
LINK      CLASS  MTU    STATE  OVER
net0      phys   1500   up     --
net2      phys   1500   up     --
net1      phys   1500   up     --
net3      phys   1500   up     --
aggr0     aggr   1500   up     net2 net1 net3
root@bookable-10-163-223-199:~#

```

As shown in the output above, an aggregated link called aggr0 has been created, consisting of net 1, net 2, and net 3. This information can also be seen from the UI, as shown below:



Creating a Virtual Switch

The command line output shown below shows that the primary domain has two virtual switches; one from the admin network and another from the public network. The secondary domain, on the other hand, has no virtual switches.

```

root@t5-2-tvp540-c:/var/log/vntsd/secondary# ldm ls -o network primary
NAME
primary
MAC
00:10:e0:2a:3c:e4
VSM
NAME      MAC          NET-DEV  ID  DEVICE  LINKPROP  DEFAULT-VLAN-ID  PVID  VID
10.163.208.0_20  00:14:4f:fb:2e:c5  net0    0   switch00  phys-state  1             1
primary-prod  00:14:4f:fe:32:6e  aggr0   1   switch01  phys-state  1             1234
root@t5-2-tvp540-c:/var/log/vntsd/secondary# ldm ls -o network secondary
NAME
secondary
MAC
00:14:4f:ff:a0:02
root@t5-2-tvp540-c:/var/log/vntsd/secondary#

```

To create a virtual switch for the admin network:

1. Expand Assets in the Navigation pane and select the control domain that was just created.
2. Select the Networks tab in the center pane and click **Attach Network**.
3. Select the admin network, in this example is OS Network.

Network Name	VLAN ID	Service Domain	Switch name	Link	Over	IP Address
OS Network - Global Labs	-	primary	10.163.208.0_20	net0	-	10.163.212.79
OS Network - Global Labs	-	secondary	-	net0	-	-
Vlan-1234	1234	primary	primary-prod	vnic17.20412588	aggr0	0.0.0.0
-	-	primary	primary-prod	aggr0	-	-

4. Configure the aggregation links, as shown in the image below.
 - Ensure the secondary service domain is selected.
 - Select the correct NIC, in this example is net0.
 - Enter a name for the Switch Name, in this example secondary-mgmt.
 - No Network IP address is allocated.

Attach Network

Steps: 1. Select Networks, 2. Configure Networks, 3. Configure Interfaces, 4. Summary

Configure Interfaces

Specify the configuration settings for each network connection.

Asset/Server Pool: t5-2-tvp540-c

Hostname	Service Domain	SRIOV	Network	P.Key/VLAN ID	Mode	NIC	Switch Name	Address Allocation Method	IP Address
OS Network - Global Labs	UK								
t5-2-tvp540-c	secondary	<input type="checkbox"/>	OS Network - Glo...		Untagged	net0	secondary-mgmt	Do not Allocate IP	Not Allocated

5. Review the summary and click **Finish**.

You can use the command line to view the newly created virtual switch.

```

root@t5-2-tvp540-c:/var/log/vntsd/secondary# ldm ls -o network secondary
NAME
secondary
MAC
00:14:4f:ff:a0:02
VSM
NAME      MAC          NET-DEV  ID  DEVICE  LINKPROP  DEFAULT-VLAN-ID  PVID  VID
secondary-mgmt  00:14:4f:fb:68:d8  net0    0   switch00  phys-state  1             1
root@t5-2-tvp540-c:/var/log/vntsd/secondary#

```

6. Repeat the steps to create another virtual switch for the aggregated link on the secondary service domain. Ensure that:

- The secondary service domain is selected.
- Select the correct NIC, in this example is `aggr0`.
- Enter a name for the Switch Name, in this example `secondary-prod`.
- No Network IP address is allocated.

You can verify both virtual switches were created from the command line or UI. In the control domain, it is possible to see the networks created for the control domain and secondary domain.

```

root@t5-2-tvp540-ci/var/log/vntsd/secondary# idm ls -o network secondary
NAME
secondary
MAC
00:14:4f:ff:a0:02
VSW
NAME      MAC      NET-DEV  ID  DEVICE  LINKPROP  DEFAULT-VLAN-ID  PVID  VID  MTU  MODE  INTER-VNET-LINK
secondary-rgmt  00:14:4f:f8:68:d8  net0  0  switch80  phys-state  1  1  1500  on
secondary-prod  00:14:4f:fb:c3:67  aggr0  1  switch81  phys-state  1  1  1234  1500  on

root@t5-2-tvp540-ci/var/log/vntsd/secondary# idm ls -o network primary
NAME
primary
MAC
00:10:e0:2a:3c:e4
VSW
NAME      MAC      NET-DEV  ID  DEVICE  LINKPROP  DEFAULT-VLAN-ID  PVID  VID  MTU  MODE  INTER-VNET-LINK
10.163.208.0_20  00:14:4f:fb:2e:c5  net0  0  switch80  phys-state  1  1  1500  on
primary-prod  00:14:4f:fe:32:6e  aggr0  1  switch81  phys-state  1  1  1234  1500  on

```

Creating the HA Guest Domain

In Oracle Enterprise Manager Ops Center, it is possible to create two types of Guest Domains: Guest Domain and HA Guest Domain. The required criteria for creating a HA Guest Domain is:

- Two network connections for each network
- Redundant storage access to the virtual disks

Since the goal of this paper is to help create a resilient environment, a HA Guest domain is created. The following steps define the procedure:

1. [Creating the HA Guest Domain Profile](#)
2. [Applying the HA Guest Domain Profile](#)
3. [Provisioning the HA Guest Domain OS](#)

It is also possible to combine multiple profiles mentioned above in a complex plan, thus enabling just one job to create the Guest HA container and install the OS.

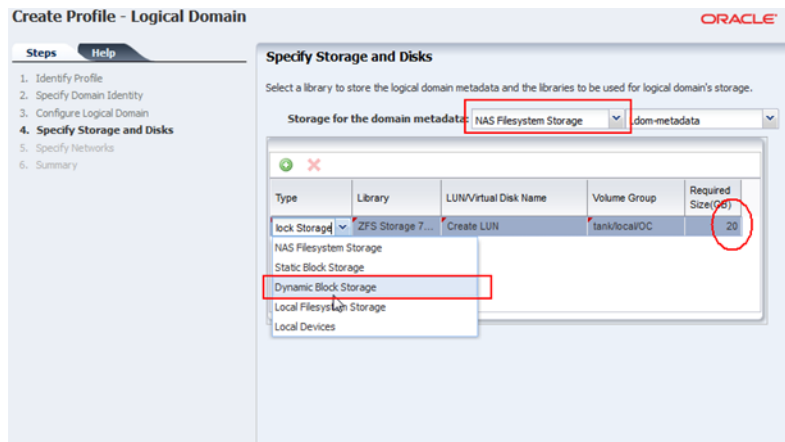
Creating the HA Guest Domain Profile

To create the HA guest domain:

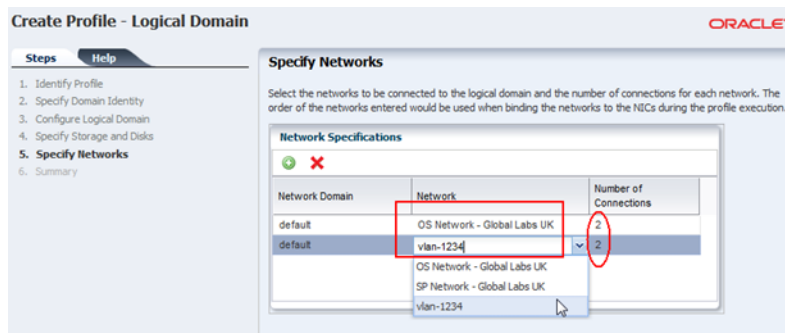
1. Select **Plan Management** section and expand **Profiles and Policies** in the Navigation pane.
2. Select **Logical Domain** profile and click **Create Profile** in the Actions pane.
3. Enter the following details in the **Create Profile - Logical Domain** wizard and then click **Next**.
 - Name and description of the profile.
 - Retain the option to create a deployment plan for this profile.

- Select HA Guest Domain as the Subtype.
- 4. Provide a name for the HA guest domain and click **Next**.
- 5. Enter the configuration details for the HA guest domain and then click **Next**. In this example, you use Virtual CPU model with 4 CPU threads and 4 GB of memory.
- 6. Select a library to store the metadata and the HA guest domain's storage, then click **Next**.

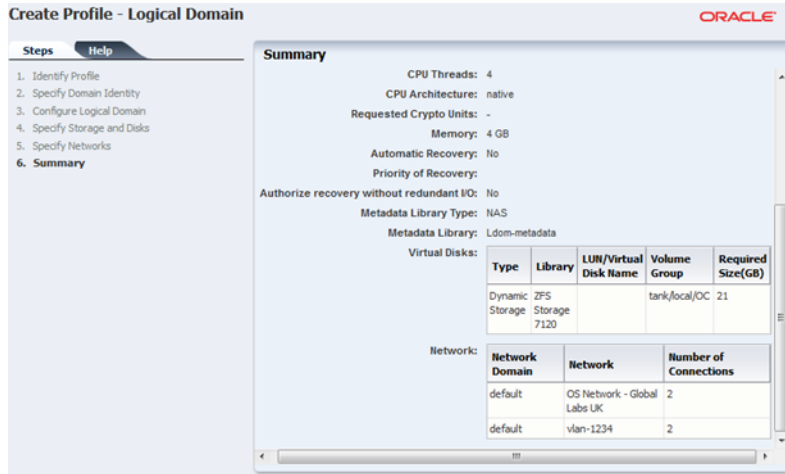
In this example, select the NAS Filesystem Storage as storage type to store the metadata and select the predefined NFS share. Select the Dynamic Block Storage from the menu, which is a Sun ZFS Storage 7120 appliance, to be the storage for the guest domain and specify the size of 20 GB.



- 7. Specify the networks to add to the HA Guest by adding two network connections for each of the networks (admin and public).



- 8. Review the parameters selected for the profile and click **Finish** to create the profile.

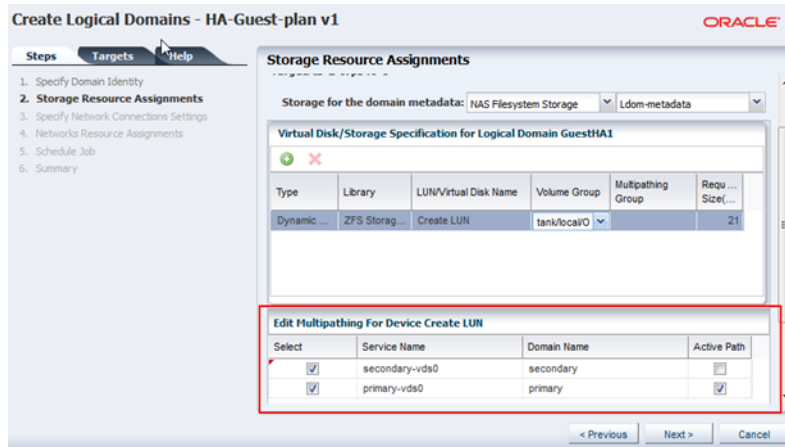


Applying the HA Guest Domain Profile

To apply the logical domain profile in Oracle Enterprise Manager Ops Center:

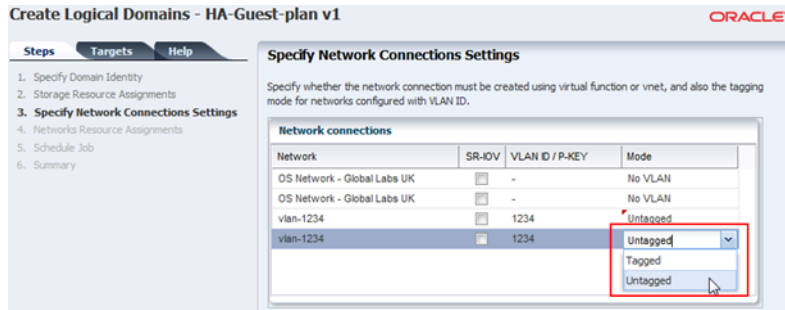
1. Expand Assets and select the control domain.
2. Click **Create Logical Domains** in the Actions pane.
3. Select the HA Guest profile that was just created and click **Apply Plan**.
4. Retain the default name and click **Next**.
5. Retain the values for metadata and virtual disks, select which devices to multi-path, and chose one for the active path.

Since the objective is to create a redundant guest domain, provide a LUN from both the primary and secondary domains. One path must be chosen as the active path.

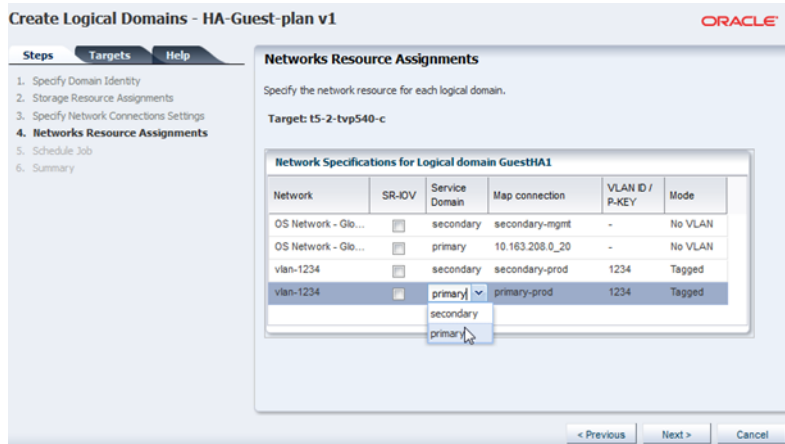


Click **Next**.

6. Choose Untagged as the mode for the public network (vlan network) so that the guest OS will see the network as untagged.



- Specify the network from each domain. One connection will be for the control domain and the other for the secondary domain, to ensure that there is redundancy.



- Review the Summary and click **Finish** to complete the HA guest domain installation.

After the job completes successfully, HA guest domain is created. The guest domain appears in the Navigation pane listed under the corresponding control domain.

Provisioning the HA Guest Domain OS

Installing the OS requires creating the OS Provisioning Profile (OSP) and OS Configuration Profile (OSC), then combining the Profiles in a Deployment Plan to apply it to the HA guest domain.

Creating an OS Provisioning Profile

- Select **Plan Management** section and expand **Profiles and Policies** in the Navigation pane.
- Select **OS Provisioning** profile and click **Create Profile** in the Actions pane.
- Enter the following details in the **Create OS Profile - OS Provisioning** wizard and then click **Next** to specify the OSP parameters.:
 - Name and description of the profile.
 - Select Logical Domain as the Subtype.
- Complete the remaining steps according to your requirements. In the Summary screen click **Finish** to create the OSP profile.

Creating an OS Configuration Profile

Within this profile you use IPMP for the networks to achieve the goal of having resilient networks.

To create the OSC in Oracle Enterprise Manager Ops Center:

1. Select the **Plan Management** section in the Navigation pane and expand **Profiles and Policies**.
2. Select OS Configuration and click **Create Profile** in the Actions pane.
3. Enter the following details in the **Create Profile - OS Configuration** wizard:
 - Name and description of the profile.
 - Select logical domain as the Subtype.

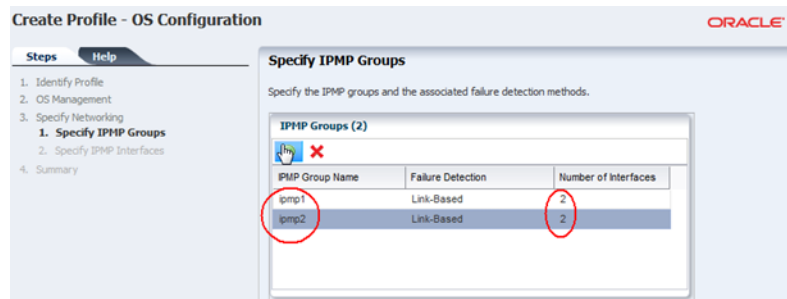
Click **Next**.

4. Select to manage the OS automatically and deploy the Agent Controller to manage the asset. Select the option Enable Multiplexed I/O so that you can associate block storage libraries such as FC and iSCSI for storage with the OS.

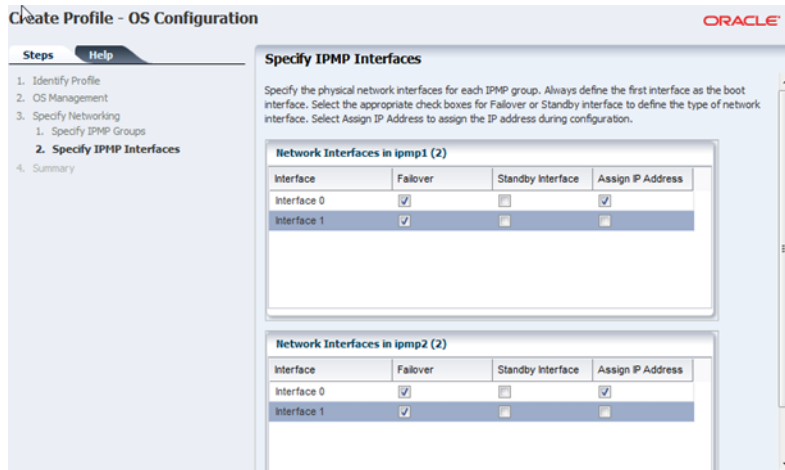
Click **Next** to specify the networking details.

5. Select Use IPMP as the networking option, then click **Next** to specify the networking details.
6. Create two IPMP groups and specify two interfaces for each, then click **Next**.

Because resilience is required, two IPMP groups are created. One will be used for the admin network and the other from the public network. Each network should have two connections, one from the control domain and another from the secondary service domain.

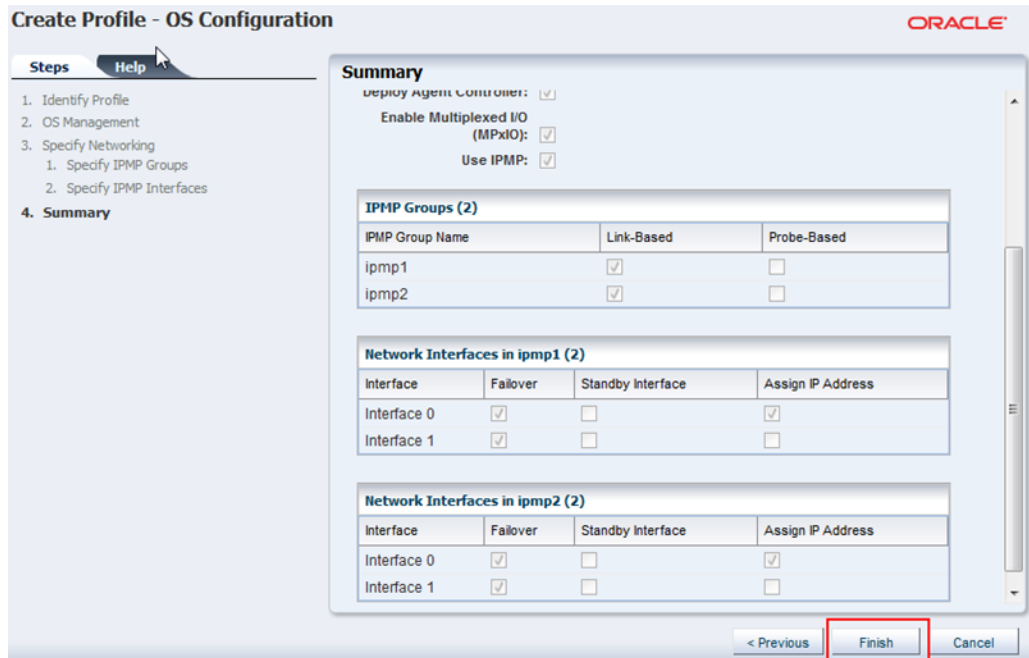


7. Specify the physical network interfaces for each IPMP groups as show in the image below, then click **Next**.



This configuration states how to set up the IPMP groups and how the failover should be initiated, such as Link Based or Probe-Based.

8. Review the summary information, and then click **Finish**.



Creating the Deployment Plan and Provisioning the Control Domain

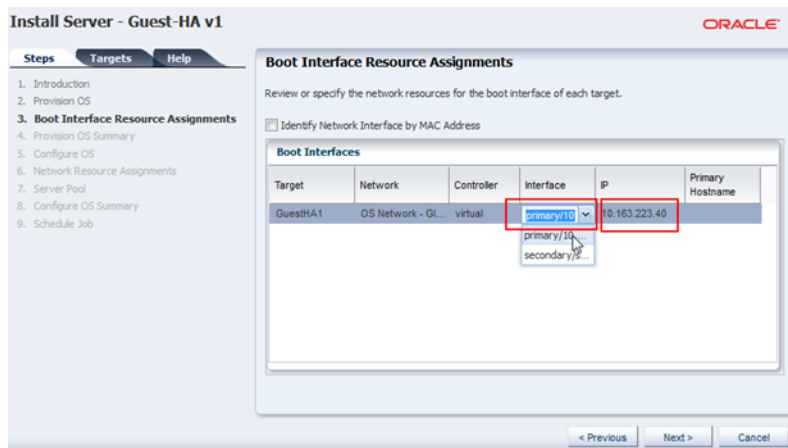
To create the deployment plan in Oracle Enterprise Manager Ops Center:

1. Select **Plan Management** section in the Navigation pane.
2. Expand **Deployment Plans** and select **Provision OS**.
3. Click **Create Deployment Plan** in the Actions pane.

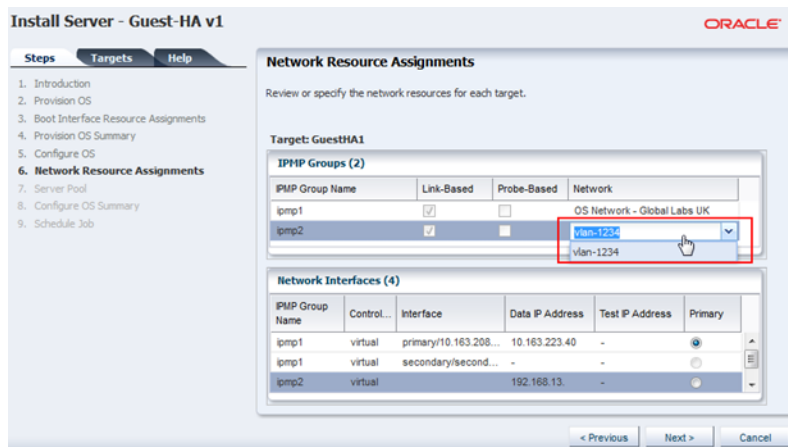
The **Create a Deployment Plan** window is displayed.

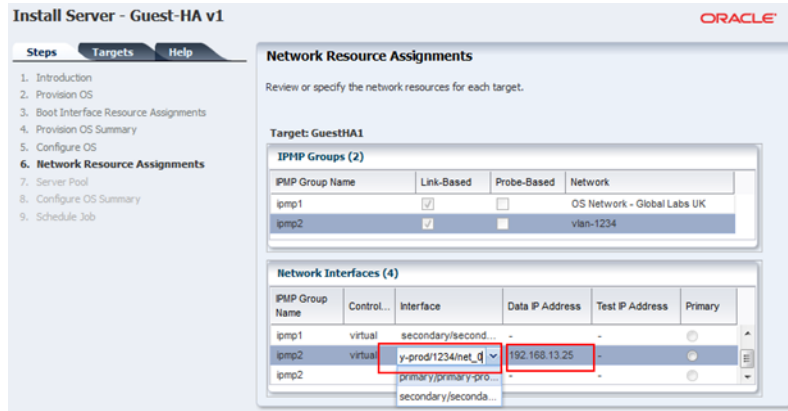
4. Enter the following information for the deployment plan:
 - Name and description for the plan.
 - Select **Stop at Failure** as the Failure Policy.

- Select the corresponding OSP and OSC profiles created for the HA guest domain.
5. Click **Save** to create the deployment plan.
Once the Deployment plan has been created, the next task is to apply the deployment plan to the HA guest domain.
 6. Expand Assets in the Navigation pane and select the HA guest domain.
 7. Click **Install Server** in the Actions pane.
 8. Select the recently created Deployment Plan for the HA guest domain, click **Apply Plan** and complete the first steps of Install Server wizard.
 9. In the Boot Interface Resource Assignment step, select the primary interface to install the OS and assign an IP address, then click **Next** and complete the next steps.



10. In the Network Resource Assignments step, select the correct network for each of the IPMP groups, then select the correct Network Interfaces for the IPMP groups, and assign an IP address. Then click **Next**.





11. Complete the remaining steps, schedule the job to run now, and click **Apply** to start a job to install the OS on the HA guest domain immediately.

It is possible and recommended to view how the install is progressing. To do this, log on the control domain and view the `/var/log/vntsd/ha-guest/console-log` file as shown below.

```

13:41:28 Download: 33022/50926 items 379.9/528.1MB 71% complete (4.1M/s)
13:41:33 Download: 35542/50926 items 399.2/528.1MB 75% complete (3.6M/s)
13:41:38 Download: 37065/50926 items 412.3/528.1MB 78% complete (3.2M/s)
13:41:43 Download: 39193/50926 items 435.7/528.1MB 82% complete (3.7M/s)
13:41:48 Download: 40727/50926 items 476.4/528.1MB 90% complete (6.4M/s)
13:41:53 Download: 42495/50926 items 499.0/528.1MB 94% complete (6.3M/s)
13:41:58 Download: 46558/50926 items 503.6/528.1MB 95% complete (2.7M/s)
13:42:03 Download: 48568/50926 items 519.3/528.1MB 98% complete (2.0M/s)
13:42:07 Download: Completed 528.14 MB in 114.85 seconds (4.6M/s)
13:42:15 Actions: 1/70640 actions (Installing new actions)
13:42:20 Actions: 15044/70640 actions (Installing new actions)
13:42:25 Actions: 17221/70640 actions (Installing new actions)
13:42:30 Actions: 19202/70640 actions (Installing new actions)
13:42:35 Actions: 21737/70640 actions (Installing new actions)
13:42:40 Actions: 23492/70640 actions (Installing new actions)
13:42:45 Actions: 25701/70640 actions (Installing new actions)

```

Once the job has completed successfully, confirm the setup is correct in the UI or the command line.

What's Next?

This document demonstrated how to improve the resilience of a guest domain by creating a secondary service domain. This was demonstrated by configuring an Oracle SPARC T5-2 server with a control domain and a secondary domain with redundant I/O access to and from a guest domain as configured within Oracle Enterprise Manager Ops Center. This configuration was tested by first halting the control domain and monitoring all guest operations. This test showed no loss of service in the guest domain. The test was repeated by halting the secondary domain, and again, there were no interruptions on the guest.

To take resilience to the next step, Server Pools should be introduced. These can perform automatic recovery of guests in the event of server failure. See the following documents for information about server pools:

- *Oracle Enterprise Manager Ops Center Virtualization Guide*
- *Oracle Enterprise Manager Ops Center Creating a Server Pool for Oracle VM Server for SPARC*

Related Articles and Resources

The Oracle Enterprise Manager Ops Center 12c Release 3 documentation is available at http://docs.oracle.com/cd/E59957_01/index.htm.

See the following documentation for more detailed information:

- *Oracle Enterprise Manager Ops Center Creating a Server Pool for Oracle VM Server for SPARC*
- *Oracle Enterprise Manager Ops Center Virtualize Reference* for information about LDom creation
- *Oracle Enterprise Manager Ops Center Operations Reference* for information about OS provisioning
- Oracle VM Server for SPARC documentation at <http://www.oracle.com/technetwork/documentation/vm-sparc-194287.html>

For current discussions, see the product blog at <https://blogs.oracle.com/opscenter>.

See the Deploy How To library at http://docs.oracle.com/cd/E59957_01/nav/deploy.htm and the Operate How To library at http://docs.oracle.com/cd/E59957_01/nav/operate.htm for deployment and operational examples.

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.

Oracle Enterprise Manager Ops Center Configure a Secondary Service Domain, 12c Release 3 (12.3.0.0.0)
E60039-01

Copyright © 2007, 2015, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

