

An Oracle White Paper  
September 2012

# Best Practices for Building a Virtualized SPARC Computing Environment

---

<b>Introduction .....</b>	<b>2</b>
<b>Server Platform: Oracle SPARC T4 Server .....</b>	<b>2</b>
<b>Virtualization Technology: Oracle VM Server for SPARC .....</b>	<b>2</b>
<b>Datacenter Management: Enterprise Manager Ops Center 12c .....</b>	<b>3</b>
<b>Network Infrastructure: Sun Network 10GbE Switch 72p .....</b>	<b>3</b>
<b>Storage Infrastructure: Sun ZFS Storage 7320 Appliance .....</b>	<b>3</b>
<b>High Level Architecture .....</b>	<b>4</b>
<b>Virtualization Infrastructure .....</b>	<b>5</b>
<b>Virtualized Infrastructure .....</b>	<b>5</b>
<b>Design Best Practices .....</b>	<b>6</b>
<b>Project Planning .....</b>	<b>6</b>
<b>Planning the Virtualization Infrastructure .....</b>	<b>6</b>
<b>Configuring Storage Libraries .....</b>	<b>7</b>
<b>Configuring Networks .....</b>	<b>9</b>
<b>Planning Your Networks .....</b>	<b>13</b>
<b>Planning the Virtualized Infrastructure: Guests .....</b>	<b>14</b>
<b>Operational Best Practices .....</b>	<b>15</b>
<b>Enable Automatic Recovery for Server Pools .....</b>	<b>15</b>
<b>Deploy Guests .....</b>	<b>16</b>
<b>Redeploy .....</b>	<b>16</b>
<b>Backup and Recover .....</b>	<b>16</b>
<b>Live Migration .....</b>	<b>17</b>
<b>Load Balance the Server Pools .....</b>	<b>17</b>
<b>Conclusion.....</b>	<b>19</b>
<b>Glossary.....</b>	<b>20</b>
<b>References.....</b>	<b>21</b>
<b>Appendix A – Process Flow for Building a Server Pool .....</b>	<b>22</b>
<b>Preparation .....</b>	<b>22</b>
<b>Building the Server Pool .....</b>	<b>23</b>
<b>Deploying LDOM Guests.....</b>	<b>23</b>

## Introduction

This best practices guide provides a solution for a SPARC virtualized environment that hosts general computing workloads using the following technologies:

- **Server Platform:** Oracle SPARC T4 servers (T4-2 or T4-4)
- **Virtualization Technology:** Oracle VM Server for SPARC 2.2 or later
- **Datacenter Management:** Enterprise Manager Ops Center 12c
- **Network Infrastructure:** Sun Network 10GbE Switch 72p
- **Storage Infrastructure:** Sun ZFS Storage 7320 Appliance

The solution provides a scalable, performant implementation which can be effectively used as a foundation for a variety of applications, including standalone datacenters and cloud computing solutions. By following these best practices, you'll be able to produce an easy to manage environment that supports automatic high-availability failover and live migration of guest virtual machines.

This guide describes key design factors to consider during a deployment, and leverages best practices that have been developed through years of experience in datacenter and virtualization management.

### Server Platform: Oracle SPARC T4 Server

This best practices guide uses the Oracle SPARC T4-2 or T4-4 Server as its standard platform. The Oracle SPARC T4 family of servers is well-suited to virtualization management. The servers support built-in, no-cost virtualization technologies which simplify server consolidation, improve utilization, and reduce operational overhead. They run Oracle Solaris 10 and 11 with guaranteed binary compatibility and support for legacy applications.

The Oracle SPARC T4 servers are also an optimal server platform for high-end enterprise applications, server and application consolidation and Oracle database applications. They provide a hosting environment to applications which require high reliability, availability, security, and high single-thread performance.

### Virtualization Technology: Oracle VM Server for SPARC

This best practices guide uses Oracle VM Server for SPARC 2.2 or greater as the virtualization technology. The virtualization technology provides a good combination of flexibility and hard resource boundaries and most closely resembles the physical servers it replaces. Oracle VM Server for SPARC provides the following key benefits:

- Hard CPU and memory resource constraints
- Guest OS fault isolation

- Guest run various package/patch combinations of Oracle Solaris 10 and 11, when minimum requirements for hardware and Oracle VM Server for SPARC software are met.
- Live migration of guests among hosts in the server pool
- Automatic HA restart of guests on other hosts in the server pool after a physical server failure. This feature is provided by Enterprise Manager Ops Center.

### **Datacenter Management: Enterprise Manager Ops Center 12c**

Oracle Enterprise Manager Ops Center provides the management technology for this implementation to rapidly deploy and maintain a virtualized environment on Oracle VM server for SPARC technology.

Enterprise Manager Ops Center provides key hardware, operating system and virtualization management capabilities which enable administrators to provision and maintain enterprise clouds and traditional IT environments alike. Key capabilities for this implementation include the following:

- Intelligent management of the Oracle stack and Engineered Systems
- Repeatable and consistent deployment of operating systems and applications
- Rapid deployment of a virtualized and physical infrastructure
- Life cycle management and compliance reporting
- Integration with My Oracle Support (MOS) for call logging and tracking

### **Network Infrastructure: Sun Network 10GbE Switch 72p**

Oracle's Sun Network 10GbE Switch 72p Top of Rack (ToR) switches provide excellent throughput between physical server and storage nodes in this implementation. These switches provide high-density, low-latency switching and offer the port capacity to provide even greater bandwidth with link aggregation, all in a 1RU form factor.

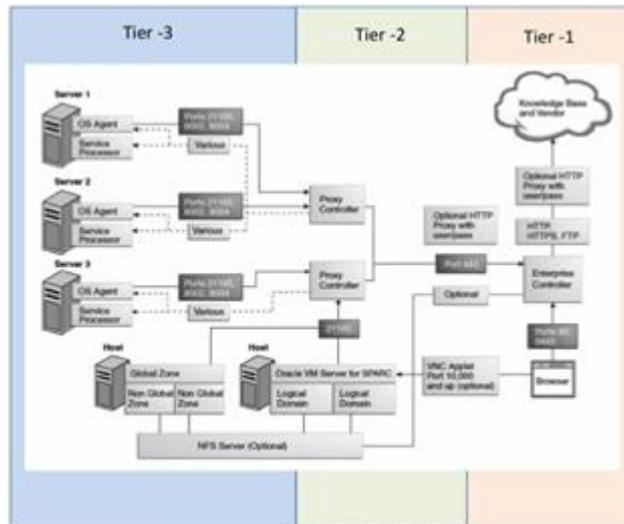
### **Storage Infrastructure: Sun ZFS Storage 7320 Appliance**

Oracle's Sun ZFS Storage 7320 or 7420 Appliance provides performant, highly-available storage for this implementation. Dual storage heads provide clustered access to a ZFS storage pool configured with the desired mix of redundancy and capacity. Storage from the pool can be a combination of filesystem and block storage, which can be accessed through data protocols including NFS, FC-SAN and iSCSI among others. When the Sun ZFS Storage Appliance is managed by Oracle Enterprise Manager Ops Center, it is possible to create, resize, and delete filesystems and block storage LUNs. Additionally, when using iSCSI, it is possible for Oracle Enterprise Manager Ops Center to create LUNs dynamically, as needed, to provide storage for virtualized servers.

## High Level Architecture

In an Oracle VM server for SPARC environment managed by Enterprise Manager Ops Center, the environment consists of three functional areas: Management and Deployment Infrastructure, Virtualization Infrastructure, and the Virtualized Infrastructure.

Enterprise Manager Ops Center provides the Management and Deployment Infrastructure, which consists of a three-tier architecture.



- Tier 1 - Enterprise Controller (EC) – The EC is the central point of control for the environment and provides the following functions:
  - Web-based management interface for the environment
  - Stores information in an Oracle database that can be either local to the EC or remote
  - Utilizes Software Libraries for caching patches, packages, and OS deployment media. This file system storage can be either local or through an NFS mount.
- Tier 2 - Proxy Controller (PC) – The PC provides a point of presence in various places in the network and provides the following functions:
  - A local cache point for data transferred to and from a managed asset
  - Provides an SSL-encrypted, single network port connection to the EC
  - Communicates on all network protocols to deploy and manage assets.
- Tier 3 - Agent controller (AC) – An agent can be deployed to any managed Operating System (OS) instance and provides the following:
  - Monitoring telemetry
  - Metric data collection
  - Configuration information
  - Patch analysis
  - Remote execution of Enterprise Manager Ops Center jobs
  - Note: Enterprise Manager Ops Center provides an agentless option for OS monitoring/management.

## Virtualization Infrastructure

The Oracle SPARC T4 server provides the virtualization capability, the hypervisor, that supports the Oracle VM SPARC virtualization software. While the T4 server comes in three specifications, the T4-2 and T4-4 are normally used for the virtualization infrastructure layer.



**SPARC T4-2**



**SPARC T4-4**

Within each T4 server, the first domain is the control domain (CDOM), which provides the management capability for the platform.

## Virtualized Infrastructure

With each control domain (CDOM) are virtualized guest domains, referred to as logical domains (LDOM) or, simply, guests. Operating system instances and applications run in these guests, supported by the hypervisor's virtual I/O devices.

## Design Best Practices

### Project Planning

The most critical part of a virtualization project is planning. It is a common misconception that since it is a virtualized environment, it all just happens dynamically and on demand without the need to plan. In physical server environments, planning and infrastructure provisioning are done with each installation. In a virtualized environment, the same planning and infrastructure provisioning needs to be done before the first virtualized host is deployed. While it is possible to dynamically add storage, networks and servers to the virtualized environment, these additions must be planned as part of the initial design and provisioned where physical resources allow to avoid delays and cost. The more planning done at the commencement of the project, the fewer issues are encountered during the life of the environment. This level of planning leads to smooth, simple, point-and-click provisioning of the virtualized environment.

Although the Management and Deployment Infrastructure is customized to suit the scale and complexity of each solution's requirements, a typical deployment consists of the following:

- An EC running Oracle Solaris 11 with 32 GB memory, 4(8)<sup>†</sup> x 2.8 GHz CPU with attached Enterprise class disk storage
- One or more pairs of HA PCs running Oracle Solaris 11 with 16 GB memory 4(8)<sup>†</sup> x 2.8 GHz CPU with attached Enterprise class disk storage
- A remote database server to support an Oracle 11G R2 database

The design considerations for the the Management and Deployment Infrastructure do not form part of this best practices guide, but can be determined from the following documents:

- [Oracle Enterprise Manager Ops Center Documentation Library](#)
- [Enterprise Manager Ops Center 12c Release 1 Deployment Considerations](#)

### Planning the Virtualization Infrastructure

Although it is possible for Enterprise Manager Ops Center to discover and manage existing CDOMs, to enable the full functionality that Enterprise Manager Ops Center provides to the Virtualization Infrastructure layer, use Enterprise Manager Ops Center to discover the hardware and to deploy the CDOM on top of bare metal hardware. It is recommended that any server that is deployed as a CDOM be either provisioned directly into or placed in a server pool before any guests are deployed to it.

---

<sup>†</sup> T4 architecture chipsets allocate resources on a 8 vCPU boundary

### High Availability

While there are a number of different ways to provide high availability within an Oracle VM server for SPARC environment, Enterprise Manager Ops Center provides this capability by leveraging the live migration capability of Oracle VM server for SPARC, promoting the use of server pools and shared storage libraries, and adding an automatic restart capability for guests on a failed hardware platform.

### Server Pools

A server pool is a group of virtualization servers that have the same CPU architecture/CPU speed, have common access to the same defined networks and access to the same shared storage libraries. The server pool provides a migration boundary and a framework for Enterprise Manager Ops Center's HA LDOM restart capability.

Server pools must be provisioned with sufficient CPU/memory capacity to meet your organization's N+1 or N+2 requirements and to allow for capacity to restart guests from a failed physical server.

Oracle VM Server for SPARC software does not allow for the over-commitment of memory and CPU for active guests. Enterprise Manager Ops Center does give you the ability to define additional guests in excess of the physical resources and these guests are stored in a shutdown state on shared storage. These guests can be started anytime there are sufficient free CPU and memory resources available.

### Configuring Storage Libraries

It is also necessary to plan the storage for the virtualized environment. A storage library must be accessible to all of the virtualization hosts in order to allow proper guest placement by load balancing policy and migration and failover between the members of the server pool. Guests' root disks must be on shared LUNs using either Fibre Channel (FC) or iSCSI SAN storage. In addition, Enterprise Manager Ops Center requires the use of a small NFS shared library to store guest metadata, which also must be accessible to all the members of the server pool. It is recommended that this metadata share be shared from an enterprise class NFS filer such as the Oracle Sun ZFS Storage Appliance. Each storage library must have a separate metadata share or, at a minimum, a separate metadata share for each server pool.

A block-based storage library is backed by either Fibre Channel disks or iSCSI disks in the form of Logical Units (LUNs) and can be considered either static or dynamic.

- Static block based storage
  - Existing LUNs (previously carved out of either a Fibre Channel or iSCSI array) are mapped to the library for use by Enterprise Manager Ops Center. Enterprise Manager Ops Center cannot modify, delete, create, or clone these LUNs. These functions must be done through the storage device's user interface outside Enterprise Manager Ops Center control.
- Dynamic block based storage

- Enterprise Manager Ops Center can create, resize, delete, and clone LUNs. A dynamic block storage library can be created if an Oracle's Sun ZFS Storage Appliance is discovered and managed by Enterprise Manager Ops Center.

This best practices guide discusses considerations for both iSCSI and FC SAN setup and use, but it is recommended to use dynamic storage libraries based on Oracle's Sun ZFS Storage Appliance wherever possible.

#### **Dynamic Storage Library Configuration**

This best practices guide describes how to create a storage library for guest metadata using iSCSI LUNs from Oracle's Sun ZFS Storage Appliance. In addition, the Sun ZFS Storage Appliance offers a unique capability in Enterprise Manager Ops Center to create and resize LUNs as needed.

#### **Sun ZFS Storage Appliance Setup**

Details of setting up and configuring a Sun ZFS Storage Appliance can be found in the product documentation available in [Oracle Unified Storage Systems Documentation](#).

For the best performance and reliability, use dual storage controllers with clustering enabled and configure IP multipathing and jumbo frames on the network devices. Enterprise Manager Ops Center requires the presence of Storage Connect workflows in the Sun ZFS Storage Appliance. Those workflows are part of the appliance kit from version ak/SUNW,iwashi@2011.04.24.1.1,1-1.11 and beyond. It is also necessary to set up an iSCSI Target and iSCSI Target Group within the Sun ZFS Storage Appliance and then include these names in the Enterprise Manager Ops Center's discovery profile. More details on the procedures to set these up can be found in the guide [Oracle® Enterprise Manager Ops Center Discovering a Sun ZFS Storage Appliance and Configuring Storage Libraries](#).

#### **Enterprise Manager Ops Center Discovery of Sun ZFS Storage Appliance**

After the Sun ZFS Storage Appliance is configured, use Enterprise Manager Ops Center to discover it. The detailed steps are provided in the Enterprise Manager Ops Center document referenced in the previous paragraph. Ensure that the discovery profile has the iSCSI target checked and that the Plugin Specific Information field contains the iSCSI Target Group and iSCSI Target values that were set up on the storage appliance. It is also possible to select the NFS target in the discovery profile, thereby enabling creation, management and use of NFS shares from within Enterprise Manager Ops Center. No plugin-specific information is required for the NFS target. Execute the Add Assets action using this discovery profile and when job completes, a new dynamic library is available in the Dynamic Block Storage section of the Libraries navigation accordion. This library allows creation or deletion of LUNs, shows existing LUNs, and allows them to be resized. As described later in this guide, it is also possible to use this dynamic storage library as the source for LDOM guest disks, in which case the LUN is created at the time the LDOM guest is created.

#### **Static Block (FC) Storage Library Configuration**

SAN libraries can be configured out of either Fibre Channel or iSCSI LUNs. LUNs must have the following configuration characteristics:

- Be already created and shared from the storage array
- Be accessible to all CDOMs in a server pool with which the library is associated
- Be MPXIO style devices (EMC Powerpath and VXDMP are not supported)

LUNs should have the following configuration characteristics as a best practice:

- Be presented from redundant array storage. Mirroring/RAID is provided by the array.
- Access to the array is redundant. Using redundant storage array controllers and redundant HBAs to provide access to the array is recommended.
- For iSCSI LUNs, it is recommended that the network be configured with jumbo frames
- Both the FC/network connection and the backend array are required to be sized to meet the throughput and IOPS of the combined virtualized guests in the server pool.

#### **Storage Library Tuning**

The default alignment of storage/virtual disk server (VDS)/file system is more than ample to provide the required performance for virtualized environments. If storage performance is seen as an issue, the default blocking factor of 512 Bytes of the VDS layer can be modified to match that of the underlying storage.

#### **Configuring Networks**

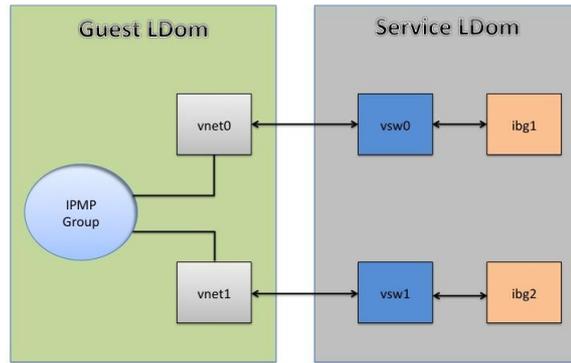
This is an area that seems to cause a great deal of confusion, probably because it requires the Unix admin/architect to cross into an area that has been the reserve of the networking specialist. While Enterprise Manager Ops Center allows you to configure the networks that are required for your environments, it does impose some constraints.

#### **Network Throughput and Redundancy Options**

Enterprise Manager Ops Center can define, deploy and manage single networks and multiple networks as either aggregations or IPMP groups. It is important to consider the purpose of each network attached to a LDOM or CDOM to determine if redundancy or additional throughput is required.

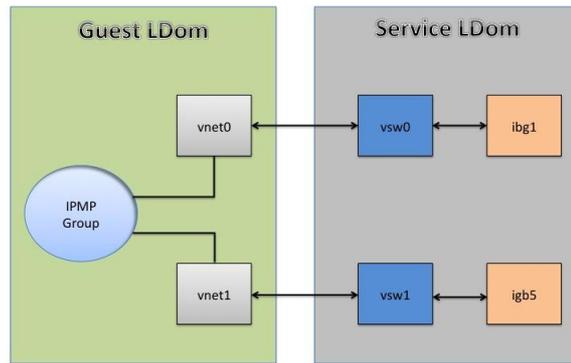
For networks that require enhanced capabilities:

- To gain redundancy against link failure, two physical NICs are presented through two vSwitches and the resulting vnets are in an IPMP group for each guest (including the control domain if it requires access).



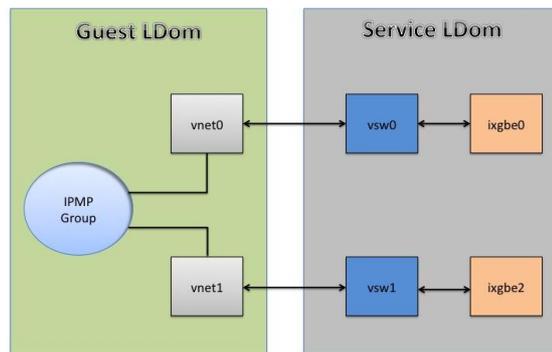
Note: Virtual interface names will differ between S10 and S11 guests

- To gain resilience against card failure, two physical NICs are built across onboard and optional PCIe NIC cards or two separate PCIe NIC cards.



Note: Virtual interface names will differ between S10 and S11 guests

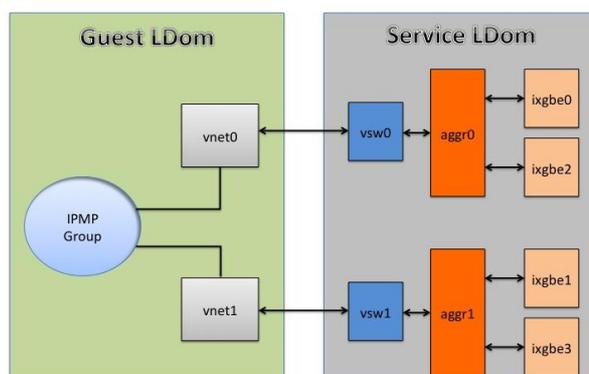
- To increase throughput, the 1 GBE interfaces can be replaced with 10 GBE interfaces.



Note: Virtual interface names will differ between S10 and S11 guests

- If even greater throughput is required, the physical NICs can be aggregated to form a new NIC that the vSwitch is then built upon. From each vSwitch, a vnet is presented into a guest (including the control domain if it requires access) that is then included in an IPM group to maintain redundancy and resilience.

Note: Aggregation requires additional functionality that must be supported by core switch infrastructure.



Note: Virtual interface names will differ between S10 and S11 guests

While other configurations of LDOM networking are possible using routing within the control domain, they are not recommended.

### Virtualization Infrastructure and Guest Networks

Each customer environment has vastly differing network environment requirements and it is impossible to represent all possible configurations. This best practices guide gives examples of common network requirements. Many customers can reduce complexity by consolidating network functions to suit their network environment.

#### Network Definitions

- Out of Band (OOB) Management –For ILOM interfaces
- In Band Management (MGMT) – For deployment and agent/agentless management
- NAS – For connectivity to NFS/iSCSI storage for vdisks
- Guest Initiated Storage (GIS) - For application NFS/iSCSI storage initiated by the virtualized guests
- Application Data Network (DATA) – For application data traffic. There can be more than one DATA network if application traffic separation is required
- Backup Network (BKUP) – For backup software to provide file-based recovery

In a best practices environment, these networks are provided on a combination of multiple physical NICs and VLAN tagging is used to allow multiple networks to co-exist on these physical NICs. Creating these dedicated single-purpose networks increases security and manageability (that is, bandwidth control) but also increases complexity. In a complex virtualized networking environment, all network interfaces must be set up as VLAN trunks to allow only the required VLANs.

Use the default VLAN to accommodate interfaces that do not natively support VLAN tagging. The OOB and MGMT networks do not natively support VLAN tagging. The OOB network is attached to the physical LOM on the hardware and does not have VLAN tag support. The MGMT network, while capable of VLAN tag support, cannot function in this mode during Operating System Provisioning (OSP). It is

recommended that both these networks be integrated into a VLAN infrastructure by setting the core switch port as a trunk, but specifying a default VLAN in which the untagged traffic is placed.

Note: The “Default VLAN” of 1 is unacceptable for this purpose.

In this configuration, Enterprise Manager Ops Center sees the OOB network as an untagged network and, depending on whether the MGMT network was connected to the PC as a tagged or untagged interface, as a tagged or untagged network. Enterprise Manager Ops Center ignores the tag on the MGMT network during OS Provisioning and leaves the deployed hosts with its MGMT network as untagged.

It is possible to configure more than one OOB and MGMT network if you require further separation.

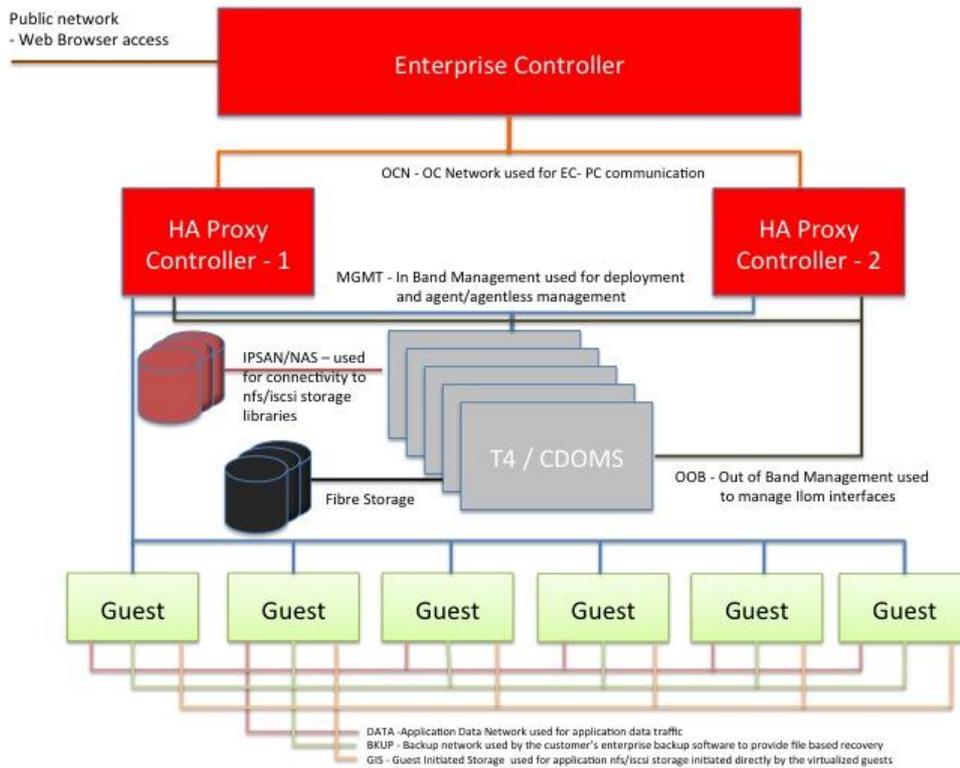
All CDOMs within a sever pool must be on the same MGMT network to communicate among themselves. The use of private VLANs (PVLAN) to isolate the CDOMs causes migration of guests to fail and might even cause MAC address conflicts and should be avoided.

The following table shows what networks are attached to a server pool:

Network	Connected to	Tagged	# Physical connections	IPMP
OOB	ILOM	No	1	No
MGMT	CDOM/Guests	No	1	No
NAS	CDOM	Yes	2	Yes
DATA	Guests	Yes	2	Yes
GIS	Guests	Yes	2	Yes
Backup	Guests	Yes	1(2)	No (optional)

Note: It is common to have more than one application DATA or GIS network.

An example of this type of configuration is seen below, where much of the network separation is being provided by VLAN tagging of the different networks.



## Planning Your Networks

It is important to identify the network requirements and to plan the network configuration as part of the initial deployment because:

- You can pre-define and in most cases pre-provision your physical cabling requirements so that these operations do not affect rapid rollout of your virtual guests.
- You can specify sufficient physical capacity of the network interfaces, that is, the number and throughput of physical NICs. In most organizations installing additional hardware and changing network configuration of hardware have long and fixed timelines for implementation and often require outages on the underlying virtualization servers, resulting in repeatedly rebalancing the guests at each upgrade.
- Enterprise Manager Ops Center requires a different level (higher) of privilege to define and add a network to the virtualization infrastructure than to define a guest network and to provision its operating system. These are two separate tasks and provisioning infrastructure should not be considered part of provisioning a virtualized guest.

It is easy and a zero cost operation to define networks and allocate the IP address ranges, VLANs, and MTU sizes as part of the initial installation so that they are available when needed. It is recommended to always define and allocate more IP address space than you think you need to allow for growth.

## Planning the Virtualized Infrastructure: Guests

While LDOMs can be configured in different ways, Enterprise Manager Ops Center deploys and maintains a LDOM in a manner that supports live migration and virtualized IO and networking.

### CPU performance

For optimum performance, LDOMs must be configured with a number of vCPUs that is a multiple of eight to ensure the best alignment to the underlying hardware. This constraint is true for the T4 CPU chip. Other T-series systems using different generation chips have differing alignment requirements.

### Cryptographic units

The SPARC T4 processor has implemented the inbuilt cryptographic framework through user-level instructions and the cyphers are implemented within the appropriate pipeline itself rather than as a co-processor, as it was done in previous T series chips. It is recommended when configuring the LDOM that the setting for “Requested Crypto Units:” be left blank. This allows your guest to take advantage of the T4’s cryptographic framework and not inhibit the live migration capability of the guest.

### Sizing

It is recommended that a minimum size of a guest be 8 VCPU and 8 GB of memory. Although it is possible to build a guest with a single vCPU and a single GB of memory, these guests have less than optimal performance and it is not recommended. If this is the size of the guest environment you want, consider using Oracle Solaris Zones, which can provide finer grain control without the overhead of a full Oracle Solaris kernel for each guest instance.

### OS Versions

To gain full functionality of the underlying Oracle VM for SPARC software and to support the LDOM guest operating systems, the installed operating system must be Oracle Solaris 10 9/10 OS, Oracle Solaris 11 OS or greater.

### OS/Application Storage Separation

The existing best practices around OS and application storage separation that existed on physical servers still apply to LDOM guests. Wherever possible, an application’s storage should be separated from the OS root file systems. This separate storage can be provided either by presenting a separate vdisk or by mapping iSCSI or NFS storage on the GIS network.

## Operational Best Practices

### Enable Automatic Recovery for Server Pools

For a step-by-step process flow of building a server pool with control domains and multiple networks, see Appendix A. For a detailed how-to on creating server pools, refer to [Oracle® Enterprise Manager Ops Center Creating Server Pool for Oracle VM Server for SPARC Guide 12c Release 1 \(12.1.0.0.0\) Part Number E27353-01](#)

When you create a server pool, select the Automatic Recovery option.

The screenshot shows the 'Specify Policies' step of the 'Create Server Pool' wizard. The 'Automatic Recovery' checkbox is checked, indicating that power will be off a failed server when capabilities are available before automatic recovery of its logical domains.

**Specify Policies**

**Placement Policy:**  Place guest on Oracle VM Server with lowest relative load  
 Place guest on Oracle VM Server with lowest allocated CPU and memory  
 Place guest on Oracle VM Server minimizing overall power consumption  
 A Server is over-utilized when the following values are exceeded:  
 CPU Utilization exceeds:  % for:  minutes

**Auto-Balancing Policy:**  Do not auto-balance the server pool  
 Automatic balancing of the server pool  
 Note: The date and time are in the Enterprise Controller's time zone.  
 Every:  Weeks on a  at

**Approval Policy:**  Approval not required, automatically move the guests  
 Approval required, send notification

**Automatic Recovery:**  Power off a failed server when the capabilities are available before the automatic recovery of its logical domains

For a sever pool that does not have the Automatic Recovery policy, enabled it using the “Edit Attributes” action.

The screenshot shows the 'Edit Attributes' page for the 'OVM\_SPARC\_POOL'. The 'Automatic Recovery' checkbox is checked, indicating that power will be off a failed server from Service Processor, given capabilities, before automatic recovery of attached logical domains.

**OVM\_SPARC\_POOL**

**Summary** | Libraries | Networks | Fabrics | Incidents | Monitoring | Charts

Name:   
 Type: Oracle VM-SPARC Server Pool  
 Network Domain: gmo-net-domain

**Placement Policy:**  Place guest on Oracle VM Server with lowest relative load  
 Place guest on Oracle VM Server with lowest allocated CPU and memory  
 Place guest on Oracle VM Server minimizing overall power consumption  
 A Server is over-utilized when the following values are exceeded:  
 CPU Utilization exceeds:  % for  minutes

**Auto Balancing Policy:**  Do not automatically balance the server pool  
 Remind me, according to the schedule, to manually balance the server pool  
 Automatically balance the server pool according to the schedule

**Migration Networks:**

**Automatic Recovery:**  Power off a failed server from Service Processor, given capabilities, before automatic recovery of attached logical domains.

## Deploy Guests

Use an LDOM profile that supports live migration, virtualized IO, and networking as described in *Planning the Virtualized Infrastructure: Guests*. The LDOM profile contains all of the network interfaces; the OS Provisioning profile contains only the primary provisioning network and its details. After provisioning, use the BUI to attach and create IPMP groups for any required networks.

While the Automatic HA capability is enabled by default in the LDOM profile, this capability can also be enabled/disabled for an individual LDOM by selecting the “Enable Automatic Recovery” or “Disable Automatic Recovery” action.



Procedures for manual recovery of a failed guest can be found in the guide [Oracle® Enterprise Manager Ops Center Recovering Logical Domains from a Failed Server Release 12.1.1.0.0 Part Number E36058-01](#)

## Redeploy

A conceptual change is required for incident management in virtualized environments. In a physical environment, we repair, recover, restore, and finally redeploy if all else fails. In a virtual environment, the virtualization infrastructure is disposable. When a virtual host (CDOM/GDOM) fails, its guests are reallocated to other nodes within the server pool. The quickest way to regain access to the failed server’s capacity is to redeploy it instead of diagnosing software issues and repairing or restoring them to a previous state. In static environments, restoration creates a good representation of the current state of the environment. In a dynamic virtualized environment, especially one where the virtualization has HA capabilities, restoration has risks because it can restore to a state that does not reflect the current environment.

This concept can be taken even further than the virtualized infrastructure when the service (application) is stored on Guest Initiated Storage (GIS) and is not tied to virtualized OS. For most organizations this journey is a gradual process and starts with one physical server equals one virtualized server.

## Backup and Recover

### Enterprise Controller Backup and Recovery

The Enterprise Controller is backed up and restored using the “ecadm backup” and “ecadm restore” commands. This first command backs up the configuration files and local database. For recovery to a different machine or the same machine after a rebuild, the machine must first be prepared by installing the OS and Enterprise Manager Ops Center software. This procedure is covered in the manual [Oracle® Enterprise Manager Ops Center Administration Guide 12c Release 1 \(12.1.1.0.0\) Part Number E25143-04](#).

In addition, it is recommended that standard Enterprise backup software be used to provide file-based recovery for a single file(s).

#### **Database Backup and Recovery**

The co-located database is backed up/recovered as part of the ecadm backup/recovery process. Remote databases must be backed up using the standard database backup/recovery methods for your organization.

#### **Proxy Controller Backup and Recovery**

Standard file-based backups of the operating system and PC software using Enterprise backup software allow recovery from a partial or total failure. If a HA Proxy Controller is available on the subnet, assets can fail over to the alternate Proxy Controller and the faulted PC be re-deployed and the assets rebalanced.

#### **CDOMs Backup and Recovery**

Although it is possible to backup a CDOM with file-based Enterprise backup software, a CDOM should be considered as a disposable asset and redeployed if the system becomes unrecoverable. Restoring out of date state information to a CDOM can cause issues.

#### **Guest Backup and Recovery**

File-based backup and recovery is done using the same Enterprise backup software that you currently use for physical hosts. Snapshotting of the underlying vDisk at storage array level provides an alternative backup solution but does not provide file-based recovery.

#### **Live Migration**

To take advantage of Enterprise Manager Ops Center's live migration capability, the server pool must be configured with the following:

- Oracle VM for SPARC 2.1 or higher (2.2 is recommended)
- The SPARC T4 server has a minimum firmware level of 8.1.1 (8.2.0f or higher is recommended)
- All networks and storage libraries are physically presented to all servers in the server pool as well as being defined in Enterprise Manager Ops Center
- All CDOMs have network connectivity to all other CDOMs in the server pool.
- The agent on the CDOMs is must be online.

If you have chosen to deploy zones into your LDOM guests, they migrate with the LDOM guest, but inherit the guest's capability for live migration.

#### **Load Balance the Server Pools**

Enterprise Manager Ops Center offers three placement policies for load balancing within a server and the option to balance on a recurring schedule or manually. In addition, you can balance a server pool manually using the "Migrate Guest" action.

**Placement Policy:**  Place guest on Oracle VM Server with lowest relative load  
 Place guest on Oracle VM Server with lowest allocated CPU and memory  
 Place guest on Oracle VM Server minimizing overall power consumption

A Server is over-utilized when the following values are exceeded:  
CPU Utilization exceeds:  % for  minutes

**Auto Balancing Policy:**  Do not automatically balance the server pool  
 Remind me, according to the schedule, to manually balance the server pool  
 Automatically balance the server pool according to the schedule

Because the placement policy also determines where new guests are placed, when you deploy a guest, you do not choose the physical server where the guest is created. In addition, when a guest is migrated to another host within the server pool, either as a load-balancing event or as a failure event, the HA capability restarts the guest.

## Conclusion

Oracle hardware and software are engineered to work together and this best practices document has shown how to plan for and implement a robust, scalable and highly performant virtualized computing environment using the full Oracle stack. From the Oracle SPARC T4 servers, Sun 10GbE switches, and Sun ZFS Storage Appliances through to Oracle VM Server for SPARC and Oracle Solaris 11, Oracle Enterprise Manager Ops Center is the single point for monitoring and managing the virtualized infrastructure as well as the physical infrastructure.

Enterprise Manager Ops Center enhances the underlying capabilities of Oracle VM Server for SPARC by removing the complexity of configuring LDOMs manually and increases high availability by the use of server pools and automatic guest recovery.

Although no two customer environments are the same, following the guidelines in this best practice guide allow you to deploy and manage a successful virtualized SPARC computing environment with the least amount of effort. Additional details can also be found in the [Oracle Optimized Solution for Enterprise Cloud Infrastructure on SPARC](#).

## Glossary

Acronym	Term	Description
ILOM	Integrated Lights out Management	The network attached service processor that allows power control, booting, configuration and alerting.
CDOM	Control domain	First domain in an Oracle VM Server for SPARC installation that provides management services for guest domains, also called the primary domain.
LDOM	Logical Domain	A guest domain where applications execute.
GDOM	Global Domain	A Solaris 10 or 11 operating system that supports zone deployment, also called a global zone.
NZ	Native Zone	A non-global zone that is the same version of OS as the parent Global Zone.
BZ	Branded Zone	A non-global zone that is running an older version of OS than the parent Global Zone.
EC	Enterprise Controller	Enterprise Manager Ops Center tier responsible for summarizing information across the lower tiers and displaying it to users via the browser user interface.
PC	Proxy Controller	Enterprise Manager Ops Center tier responsible for communication between Enterprise Controller tier and Agent Controller, responsible for provisioning and other services.
AC	Agent Controller	Enterprise Manager Ops Center tier that runs on the target OS instance responsible for managing and monitoring the OS, accepting jobs from Proxy Controller and sending state information to the Proxy Controller.
Pool	Server Pool	Shared network, storage, and compute resources attached a group of CDOMs or GDOMs. A server pool forms a migration boundary for virtualized guests that are part of that pool.
vSwitch	Virtual Switch	A virtual switch implemented in the Oracle VM Server for SPARC hypervisor layer to allow multiple LDOM guests to access a network through the same physical NIC.
NIC	Network interface Card	A physical network interface.
vNET		A virtualized network interface presented from the underlying hypervisor to the LDOM guest OS.
vnic	Virtual Network interface Card	A virtualized network interface under Solaris 11

## References

Description	Link
Enterprise Manager Ops Center Documentation	<a href="http://docs.oracle.com/cd/E27363_01/index.htm">http://docs.oracle.com/cd/E27363_01/index.htm</a>
Enterprise Manager Ops Center Deployment Considerations	<a href="http://www.oracle.com/technetwork/oem/ops-center/wp-emoc12c-deployment-1715835.pdf">http://www.oracle.com/technetwork/oem/ops-center/wp-emoc12c-deployment-1715835.pdf</a>
Oracle Unified Storage Systems Documentation	<a href="http://www.oracle.com/technetwork/documentation/oracle-unified-ss-193371.html?ssSourceSiteId=ocomen">http://www.oracle.com/technetwork/documentation/oracle-unified-ss-193371.html?ssSourceSiteId=ocomen</a>
Oracle® Enterprise Manager Ops Center Discovering a Sun ZFS Storage Appliance and Configuring Storage Libraries	<a href="http://docs.oracle.com/cd/E27363_01/doc.121/e27328/toc.htm">http://docs.oracle.com/cd/E27363_01/doc.121/e27328/toc.htm</a>
Oracle® Enterprise Manager Ops Center Creating Server Pool for Oracle VM Server for SPARC Guide	<a href="http://docs.oracle.com/cd/E27363_01/doc.121/e27328/toc.htm">http://docs.oracle.com/cd/E27363_01/doc.121/e27328/toc.htm</a>
Oracle® Enterprise Manager Ops Center Recovering Logical Domains from a Failed Server	<a href="http://docs.oracle.com/cd/E27363_01/doc.121/e36058/toc.htm">http://docs.oracle.com/cd/E27363_01/doc.121/e36058/toc.htm</a>
Oracle® Enterprise Manager Ops Center Administration Guide	<a href="http://docs.oracle.com/cd/E27363_01/doc.121/e25143/backup_and_recovery.htm">http://docs.oracle.com/cd/E27363_01/doc.121/e25143/backup_and_recovery.htm</a>
Oracle® Enterprise Manager Ops Center Creating and Managing Network Domains	<a href="http://docs.oracle.com/cd/E27363_01/doc.121/e27331/toc.htm">http://docs.oracle.com/cd/E27363_01/doc.121/e27331/toc.htm</a>
Oracle® Enterprise Manager Ops Center Configuring and Managing File System Storage Libraries Guide	<a href="http://docs.oracle.com/cd/E27363_01/doc.121/e27329/toc.htm">http://docs.oracle.com/cd/E27363_01/doc.121/e27329/toc.htm</a>
Oracle® Enterprise Manager Ops Center Configuring and Deploying Oracle VM Server for SPARC 12c	<a href="http://docs.oracle.com/cd/E27363_01/doc.121/e27349/toc.htm">http://docs.oracle.com/cd/E27363_01/doc.121/e27349/toc.htm</a>
Oracle® Enterprise Manager Ops Center Configuring and Installing Logical Domains	<a href="http://docs.oracle.com/cd/E27363_01/doc.121/e27355/toc.htm">http://docs.oracle.com/cd/E27363_01/doc.121/e27355/toc.htm</a>

## Appendix A – Process Flow for Building a Server Pool

### Preparation

1. Confirm all required networks are defined
  - Name – Make the name short and sensible
  - Network IP address – The first address of the subnet
  - Subnet Netmask – Specify in CDR format
  - Subnet default route – Points to the router off this subnet
  - VLAN ID (if required)
  - Jumbo frame enabled (if required)
  - If you plan to use this server pool as part of a vDC, create a network domain for the server pool(s) and assign all required networks to that domain. You can't use the Default network domain. For instructions on creating a network domain see [Oracle® Enterprise Manager Ops Center Creating and Managing Network Domains 12c Release 1 \(12.1.1.0.0\) Part Number E27331-01](#)
2. Confirm required Storage Libraries have been configured
  - [Oracle® Enterprise Manager Ops Center Configuring and Managing File System Storage Libraries Guide 12c Release 1 \(12.1.0.0.0\) Part Number E27329-01](#)
3. Define the OS Provisioning profile for the CDOM
  - Type: Oracle VM server for SPARC
  - Allocate vCPU as a multiple of 8 (T4)
  - Crypto units are not required (T4)
  - Select Solaris 11 operating system and remove the requirement for a later upgrade
  - For instructions on creating Oracle VM server for SPARC profile see [Oracle® Enterprise Manager Ops Center Configuring and Deploying Oracle VM Server for SPARC 12c Release 1 \(12.1.0.0.0\) Part Number E27349-01](#)
4. Define the Logical Domain profile for the guest
  - Allocate vCPU as a multiple of 8 (T4)
  - Crypto units are not required (T4)
  - Select “Automatic Recovery” and a “Recovery Priority”
  - Attach all networks required for the guest
  - For instructions on creating LDOM profile see [Oracle® Enterprise Manager Ops Center Configuring and Installing Logical Domains 12c Release 1 \(12.1.0.0.0\) Part Number E27355-01](#)
5. Define the OS Provisioning profile for the guest
  - Type: Solaris SPARC
6. Create a “Configure and Install Logical Domains” deployment plan that includes the Logical Domain profile and the OS Provisioning profile.

### Building the Server Pool

1. Provision the first CDOM, including only the OSP network
2. Create the server pool using the CDOM, Initial OSP network, and Storage library.
3. Provision the remaining CDOMs, either adding them to the pool at provisioning time or after provisioning.
4. Add additional required networks one at a time. Some networks are added twice to different physical interfaces to support IPMP groups.
5. Compare the vSwitch naming from the first and second CDOM. If they are different, remove the first CDOM from the server pool and then re-add it.
6. Use the wizard or an Operational plan to create an IPMP group of the networks where required on each CDOM.
7. If jumbo frames are required on any of the networks, reboot the CDOMs to activate jumbo frames.

### Deploying LDOM Guests

1. From the Assets menu, select “Server Pools” from the pull-down list and then select the new server pool
2. From the Actions pane, select “Create Logic Domains” and then select the “Configure and Install Logical Domains” plan you created.
3. Answer the questions in the wizard to deploy the guest.
4. Use the wizard or an operational plan to add any additional network configuration required.



Best Practices for Building a Virtualized SPARC  
Computing Environment  
September, 2012

Oracle Corporation  
World Headquarters  
500 Oracle Parkway  
Redwood Shores, CA 94065  
U.S.A.

Worldwide Inquiries:  
Phone: +1.650.506.7000  
Fax: +1.650.506.7200

oracle.com



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2012, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

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

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. 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. UNIX is a registered trademark licensed through X/Open Company, Ltd. 1010

**Hardware and Software, Engineered to Work Together**