Technical Report

Oracle VM and NetApp Storage Best Practices Guide

Oracle Alliance Engineering Team, NetApp
May 2011 | TR-3712

EXECUTIVE SUMMARY

NetApp's innovative technologies enable organizations to extract benefits from their virtual infrastructures by seamlessly integrating advanced virtualized storage alongside the virtual servers.

NetApp® provides industry-leading solutions in the areas of data protection, thin storage provisioning, data deduplication, file-based backups, instantaneous virtual machine (VM) backup and restores, and instantaneous VM cloning for testing, application development, and training purposes.

This technical report reviews the best practices for implementing an Oracle® VM (OVM) virtual infrastructure with NetApp storage systems. Scenarios corresponding to all three storage protocols—NFS, iSCSI, and FC—are covered, along with an Oracle VM–NetApp deployment case study.

TABLE OF CONTENTS

LIST OF FIGURES

1 INTRODUCTION

This technical report provides the best practices and a step-by-step guide for deploying Oracle VM virtualization software with shared storage solutions from NetApp.

1.1 SCOPE

This TR focuses on the deployment of Oracle VM virtualization software with NetApp storage (NFS, iSCSI, and FC SAN). It does not cover the comprehensive list of features available with Oracle VM virtualization software and it cannot be used as a replacement for Oracle VM user or administrative manuals.

1.2 INTENDED AUDIENCE AND ASSUMPTIONS

This document is for system and storage administrators, product management, and IT infrastructure managers, familiar with concepts of Oracle VM Server v2.2 and NetApp Data ONTAP® 7G.

2 OVERVIEW OF THE ORACLE VM AND NETAPP STORAGE SOLUTION

2.1 ADVANTAGES OF VIRTUALIZATION

Virtualization plays a key role in the ongoing transformation of organizations' IT departments. IT needs to respond rapidly to ever-changing business demands. To do so, an extremely flexible IT infrastructure is needed. This is where virtualization steps in.

Oracle's virtualization strategy is an extension of the Oracle Grid strategy. Oracle Grid seamlessly leverages the latest advancements in virtualization to further enhance the value of the grid infrastructure. It enables customers to dynamically share resources such as servers, storage, and network among different enterprise applications. In this way, virtualization is fast becoming the key enabler of dynamic IT.

With the advent of new generations of commodity servers supporting higher processor core counts and more powerful processors, servers deployed in enterprise data centers are well below optimal levels. Such underutilization represents a major opportunity for improvement. The combination of server and storage virtualization deployed in the Oracle Grid environment increases both server and storage utilization to maximize data center efficiency while lowering the cost of the infrastructure and operations.

Advanced server and storage virtualization significantly adds value to an existing Oracle Grid infrastructure by improving management efficiency and delivering higher service-level agreements (SLAs). This new generation of Oracle Grid will increasingly utilize advanced server virtualization and the latest storage virtualization technologies to:

- Increase server consolidation.
- Reduce power and space requirements.
- Reduce complexity.
- Achieve greater software isolation and platform uniformity.
- Provide better support for legacy applications.
- Simplify data management tasks.
- Provide capacity on demand.

2.2 ORACLE VM FOR ENTERPRISE READY VIRTUALIZATION

In November 2007, Oracle introduced Oracle VM, Xen architecture–based server virtualization software that fully supports both Oracle and non-Oracle applications. Oracle VM offers scalable, highly efficient, and low-cost server virtualization. Consisting of open source server software and an integrated browser-based management console, Oracle VM provides an easy-to-use GUI for creating and managing virtual server pools running on x86 and x86 64-based systems across an enterprise.

Users can create and manage virtual machines that exist on the same physical server but behave like independent physical servers. Each virtual machine created with Oracle VM has its own virtual CPUs, network interfaces, storage, and operating system. With Oracle VM, users have an easy-to-use browser-based tool for creating, cloning, sharing, configuring, booting, and migrating VMs.

Oracle VM is fully server-centric and designed for data center use. With Oracle VM, it is easy to load-balance to make sure that resources are fully utilized and to move live virtual machines from one physical server to another without reconfiguration or experiencing downtime. A number of versions of Linux® and Windows® are supported as guest operating systems on either 32-bit or 64-bit server platforms. Oracle VM uses native Linux device drivers so that users do not have to wait for the latest hardware to be supported by the virtualization solution layer.

Oracle VM supports two types of virtual machines:

- **Hardware virtualized**: The guest operating system does not need to be modified. It is available only on Intel® VT and AMD® SVM CPUs.
- **Paravirtualized**: The guest operating system is recompiled for the virtual environment for optimal performance.

Oracle VM consists of two components:

- **Oracle VM Manager**: This component provides a standard Application Development Framework (ADF) Web application to manage Oracle VM-based virtual machines. It also provides an API for Oracle VM Server.
- **Oracle VM Server**: This component provides a virtualization environment designed to provide a self-contained, secure, server-based platform for running virtualized guests. Oracle VM Agent is included for communication with Oracle VM Manager.

Figure 1) Oracle VM architecture.

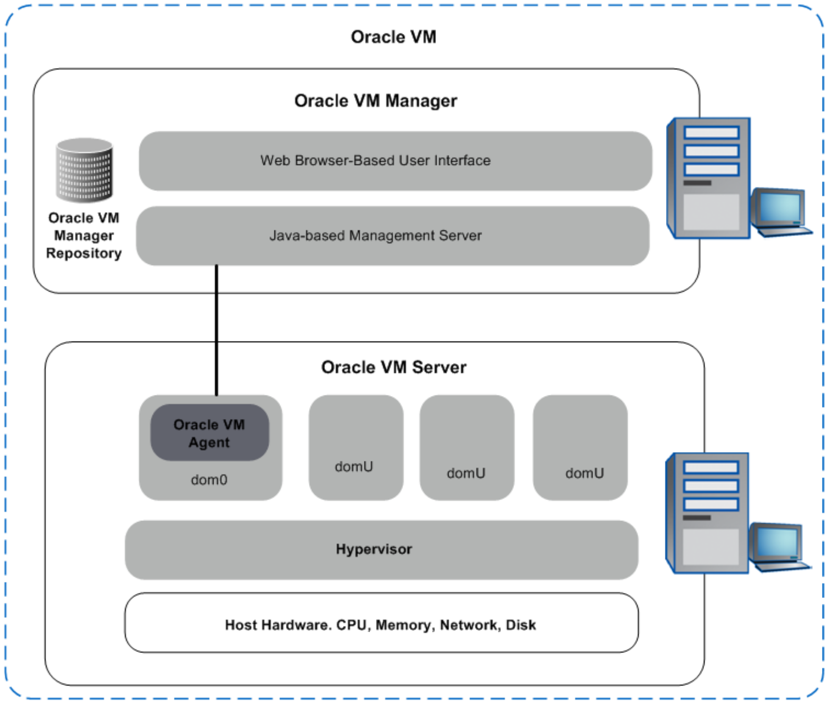

2.3 NETAPP STORAGE SOLUTIONS FOR ENTERPRISE READY VIRTUALIZATION

Server virtualization is only *one-half* of an Oracle VM infrastructure. Virtualized storage is required to complete an Oracle VM environment—in fact, the benefits of server virtualization are fully realized only when deployed with a shared storage solution.

Unified storage solutions from NetApp perfectly complement the manageability, utilization, and cost-saving benefits of Oracle VM. NetApp solutions enable powerful thin provisioning, simplified data management, and scalable and consistent I/O performance for all IT protocols across NAS (NFS) and SAN (Fibre Channel and iSCSI) in a single pool of storage. Key benefits and features are:

- Support for SAN (Fibre Channel and iSCSI) or NAS
- Nondisruptive scalability to hundreds of TB
- Easy installation, configuration, management, and maintenance
- Rapid backup and recovery with zero-penalty Snapshot™ copies
- Simple, cost-effective replication for disaster recovery
- Instant space-efficient data clones for provisioning and testing
- Dynamically expand and contract storage volumes as needed
- Data deduplication to reduce capacity requirements

NetApp storage solutions offer these powerful data management and data protection capabilities, which allow the Oracle VM to lower costs while meeting capacity, utilization, and performance requirements.

MultiStore® (a licensable feature) is another storage virtualization approach provided by NetApp that perfectly complements the server consolidation achieved through Oracle VM. MultiStore subdivides a NetApp physical storage system into multiple logical domains or virtual storage server partitions known as vFiler™ units, each with its own unique identity, management, and IP address. Diverse applications running on different virtual machines consolidated into the same physical server and common storage systems can be isolated and secured by creating separate vFiler units to store the application data.

MultiStore also enables vFiler units to transparently migrate to different physical systems without requiring reconfiguring client application servers and the mount points used for accessing data. For more details regarding NetApp MultiStore, see TR-3462, Storage Virtualization and DR Using vFiler.

Storage management and monitoring are very critical to the successful operations of a virtual infrastructure. NetApp Operations Manager suite of products can be used for the day-to-day activities on storage systems like discovering storage systems, monitoring the device health, the capacity utilization and performance characteristics of a storage system, configuring alerts and thresholds for event managements and so on. It also supports configuration of role-based access control (RBAC) for user login and role permissions. RBAC supported by NetApp Operations Manager allows administrators to manage groups of users by defining roles based on their specific responsibilities. For example, in a virtual infrastructure where an application (such as Oracle Database) administrator, virtual server (such as Oracle VM) administrator and the storage (NetApp) administrator may have to work in a tight synergy, role based control from the Operations Manager can dictate how they access the different storage resources——including NetApp FAS systems, aggregates, volumes, LUNs, protection policies, provisioning policies, vFiler templates, and so on. For details about NetApp Operations Manager, refer to http://www.netapp.com/us/products/management-software/operations-manager.html

2.4 ORACLE VM VALIDATED CONFIGURATION WITH NETAPP STORAGE

Oracle has extended its “Validated Configuration” program to include Oracle VM Server virtualization software. This program offers pre-tested, validated Oracle VM architectures—including software, hardware, storage, and network components—along with documented best practices. The program results in easier, faster, and lower-cost deployment of Oracle VM solutions in enterprises leading to improved performance, scalability, and reliability of Oracle VM solutions, with faster, lower-cost implementations.

NetApp collaborated with Oracle to come up with an Oracle VM validated configuration with NetApp storage. For more details, see Validated Config for Oracle VM with NetApp.

For more information on Oracle validated configurations, see http://www.oracle.com/technology/tech/linux/validated-configurations/index.html.

2.5 SHARED STORAGE OPTIONS FOR ORACLE VM

The guest operating systems running inside the virtual machines hosted on the OVM Server view the disk storage allocated while creating the virtual machine (for example, using the `virt-install` command) as a single virtual hard disk, analogous to a physical hard disk. This appears as “ad" and can be partitioned and managed in the guest exactly as if it were regular physical hardware.

Three types of shared storage options are available for an OVM Server configuration:

- Network attached storage using NFS
- iSCSI SAN
- Fibre Channel SAN

NetApp NFS shared storage gives unmatched flexibility to a virtual infrastructure deployed with Oracle VM Server. The files corresponding to the virtual disks of the virtual machines are thin provisioned by default and also deduplicated (if the deduplication license is enabled). This leads to very high utilization of storage as well as a drastic reduction in the total amount of storage required.

Oracle VM Server supports either a software-based iSCSI initiator or a supported iSCSI HBA. Similarly it supports Fibre Channel SANs using the supported FC HBA. The iSCSI or FC shared storage of OVM Server requires configuration of the Oracle Cluster File System (OCFS2) for use as a shared virtual disk for migration.

Note: Each virtual disk file belonging to the virtual machines has its own I/O queue to the NetApp storage system for all of these storage solutions.

Figure 2) Physical and logical storage configuration of NetApp FAS controllers.

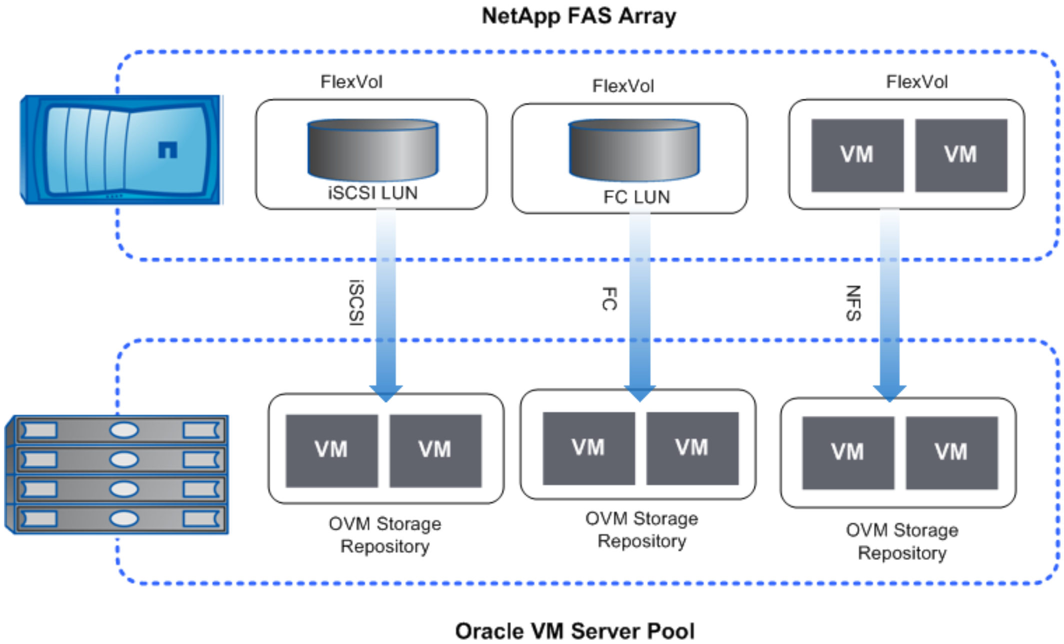

2.6 ORACLE VM SERVER REPOSITORIES

- Oracle VM uses the concept of storage repositories to define where Oracle VM resources may reside. Resources include guest virtual machines, virtual machines templates (guest seed images), ISO images, shared virtual disks, and so on. A storage repository is used for live migration of guests, high availability (HA), and local storage.
- **For Oracle VM Server v2.2:**
  - OVM repositories can be created and managed using the `/opt/ovs-agent-2.3/utils/repos.py` script.

```
Usage:
/opt/ovs-agent-2.3/utils/repos.py -n <storage>          : create an sr
/opt/ovs-agent-2.3/utils/repos.py -d <uuid>             : delete an sr
/opt/ovs-agent-2.3/utils/repos.py -r <uuid>             : set root sr
/opt/ovs-agent-2.3/utils/repos.py -l                    : list all sr
/opt/ovs-agent-2.3/utils/repos.py -i                    : init all sr
```

  - The /OVS directory is the location of the default storage repository created when Oracle VM Server is installed with the default partition layout.
  - The `/OVS` directory is the cluster root and is a symbolic link mounted to the `/var/ovs/mount/uuid` directory.
- **For Oracle VM Server v2.1.5 or below:**
  - Repositories can be added by using the `ovs-makerepo` script, and can be removed by using the `ovs-offlinerepo –d` script.
  - A list of the Oracle VM Server repositories is available at `/etc/ovs/repositories`.
  - A configuration file for mounting repositories is available at `/etc/ovs/repositories.options.`
- When a server pool is created in Oracle VM Manager, the storage repository is automatically mounted with the source type, for example, NFS, OCFS2, or ext3. The storage repository directory

structure (`running_pool, seed_pool, iso_pool, publish_pool, sharedDisk`) is also automatically created under the /OVS directory on the storage repository.

- Repositories are managed by Oracle VM Agent.
- To create OVM Server repositories for various shared storage options like NFS, iSCSI, and FC, see Section 4.1.
- Section 3.4 describes how virtual machines created on these repositories can be easily imported to the OVM Manager.

3 ORACLE VM AND NETAPP STORAGE BEST PRACTICES

3.1 NETAPP SHARED STORAGE BEST PRACTICES FOR HIGH AVAILABILITY AND RESILENCY OF ORACLE VM INFRASTRUCTURE

The most critical challenge that any server virtualization environment including Oracle VM faces is increased risk. As physical servers are converted to virtual machines and multiple virtual machines are consolidated onto a single physical server, the impact of the failure of even a single component of the consolidated platform can be catastrophic.

In an Oracle VM environment, the availability and performance of the shared storage infrastructure is critical; it is therefore vital to consider the required level of availability and performance when selecting and designing the storage solution for the virtualized server environment.

When focusing on storage availability, many levels of redundancy are available for deployments, including purchasing physical servers with multiple storage interconnects or HBAs, deploying redundant storage networking and network paths, and leveraging storage arrays with redundant controllers. A deployed storage design that meets all of these criteria can be considered to eliminate all single points of failure. The reality is that data protection requirements in a virtual infrastructure (like Oracle VM) are greater than those in a traditional physical server infrastructure. So data protection has to be the paramount feature of the shared storage solution.

NETAPP SHARED STORAGE SYSTEM CONFIGURATION AND SETUP

NetApp offers a comprehensive set of software and hardware solutions to address the most stringent requirements for availability and performance of large, scalable Oracle VM environments. The following sections provide a high-level overview of the NetApp components and features that should be considered when deploying Oracle VM Server virtualization on NetApp storage solutions.

RAID DATA PROTECTION

RAID-DP® is an advanced RAID technology that is provided as the default RAID level on all NetApp storage systems. RAID-DP protects against the simultaneous loss of two drives in a single RAID group. It is very economical to deploy; the overhead with default RAID groups is a mere 12.5%. This level of resiliency and storage efficiency makes data residing on RAID-DP safer than data residing on RAID 5 and more cost effective than RAID 10. NetApp recommends using RAID-DP on all RAID groups that store Oracle VM data.

For more information about NetApp's deduplication technology, see [TR-3505: NetApp Deduplication for FAS and V-Series Deployment and Implementation Guide](TR-3505).

Figure 3) RAID-DP.

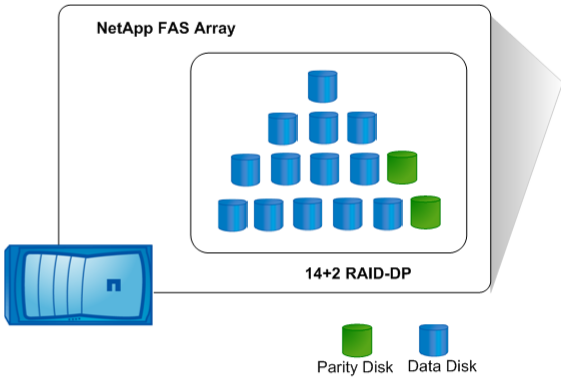

NetApp RAID-DP
– *Protects against double disk failures*
– *Provides high storage utilization (87.5% by default)*
– *No performance impact*
– *Least probability of data loss over a period of time*

	RAID 5	RAID 10	RAID DP
Cost	Low	**High**	**Low**
Performance	**Low**	High	**High**
Resiliency	**Low**	High	**High**

ACTIVE-ACTIVE NETAPP CONTROLLERS

NetApp clusters, referred to as active-active HA pairs, consist of two independent storage controllers that provide fault tolerance and high-availability storage for virtual environments. The cluster mechanism provides nondisruptive failover between controllers in the event of a controller failure. Redundant power supplies in each controller maintain constant power. Storage HBAs and Ethernet NICs are all configured redundantly within each controller. The failure of up to two disks in a single RAID group is accounted for by RAID-DP.

For more details, see:

www.netapp.com/us/products/platform-os/active-active.html

NetApp TR-3450: Active-Active Controller Overview and Best Practices Guidelines

Figure 4) High-availability cluster.

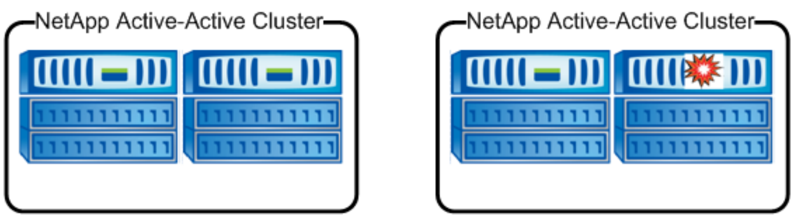

NetApp active-active provides high-availability from controller failures and up to two-disk failure (RAID-DP)

The NetApp HA cluster model can be enhanced by synchronously mirroring data at the RAID level using NetApp SyncMirror®. When SyncMirror is used with HA clustering, the cluster has the ability to survive the loss of complete RAID groups or shelves of disks on either side of the mirror.

Figure 5) NetApp SyncMirror.

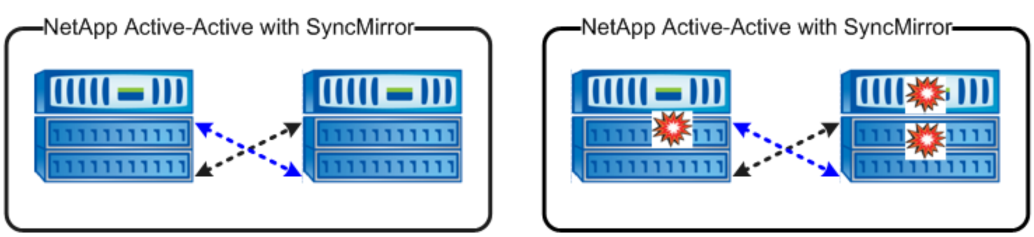

SyncMirror with NetApp active-active cluster provides high availability from complete RAID groups and disk shelves failure in addition to controllers.

MULTIPATH HA

Multipath HA storage configurations further enhance the resiliency and performance of active-active controller configurations. Although cluster failover software provides high availability by providing fault tolerance in the event of controller failure, storage-triggered events often result in unneeded failovers or prevent successful takeovers. Multipath HA storage enhances storage resiliency by reducing unnecessary takeover by a partner node due to a storage fault, thus improving overall system availability and promoting higher performance consistency. Multipath HA provides added protection against various storage faults, including HBA or port failure, controller-to-shelf cable failure, shelf module failure, dual inter-shelf cable failure, and secondary path failure. Multipath HA helps provide consistent performance in active-active configurations by providing larger aggregate storage loop bandwidth.

For more details, see TR-3437: Storage Subsystem Resiliency Guide.

REMOTE LAN MANAGEMENT CARD

The Remote LAN Management (RLM) card provides a secure out-of-band access to the storage controllers, which can be used regardless of the state of the controllers. The RLM offers a number of remote management capabilities for NetApp controllers including remote access, monitoring, troubleshooting, logging, and alerting features. The RLM also extends the AutoSupport capabilities of the NetApp controllers by sending alerts or "down-filer" notifications via an AutoSupport message when the controller goes down, regardless of whether the controller can send AutoSupport messages. These AutoSupport messages also provide proactive alerts to NetApp to help provide faster service.

For more details on RLM, see: http://now.netapp.com/NOW/download/tools/rlm_fw/info.shtml

Best Practices

NetApp recommends the following configuration options for best-in-class resiliency:

- Use RAID-DP, the NetApp high-performance implementation of RAID 6, for better data protection.
- Use multipath HA with active-active storage configurations to improve overall system availability as well as promote higher performance consistency.
- Use the default RAID group size (16) when creating aggregates.
- Allow Data ONTAP to select the disks automatically when creating aggregates or volumes.
- Use the latest storage controller, shelf, and disk firmware and the Data ONTAP general deployment release available from the NetApp Support (formerly NOW™) site.
- Maintain at least two hot spares for each type of disk drive in the storage system to take advantage of Maintenance Center (MC).
- Maintenance Center software is part of the NetApp suite of proactive, self-healing storage resiliency tools. MC provides configurable in-place disk drive diagnostics to determine the health of suspect disk drives. If Data ONTAP disk health monitoring determines that a disk drive has surpassed an error threshold, Rapid RAID Recovery is initiated to a hot spare. Afterward, the suspect disk can be placed into MC, where it undergoes a series of diagnostic tests. Consisting of Storage Health Monitor (SHM), NetApp Health Triggers, and NetApp Drive Self-Tests software, Maintenance Center promotes drive self-healing and preventive or corrective maintenance.
- Do not put SAN LUNs or user data into the root volume.
- Replicate data with SnapMirror® or SnapVault® for disaster recovery (DR) protection.
- Replicate to remote locations to increase data protection levels.
- Use an active-active storage controller configuration (clustered failover) to eliminate single points of failure (SPOFs).
- Deploy SyncMirror for the highest level of storage resiliency.

For more information on storage resiliency, see:

- [TR-3437: Storage Best Practices and Resiliency Guide]
- [TR-3450: Active-Active Controller Overview and Best Practices Guidelines]

3.2 NETAPP STORAGE NETWORKING BEST PRACTICES

Design a network infrastructure (FC or IP) so it has no single point of failure. A highly available solution includes having two or more FC or IP network switches, two or more HBAs or network interface cards (NICs) per host and two or more target ports or NICs per storage controller. In addition, if using Fibre Channel, two fabrics are required to have a truly redundant architecture.

Best Practice
For designing an FC or IP storage network infrastructure, refer to the FC/iSCSI Configuration Guide: http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/fc_iscsi_config_guide.pdf

NETAPP IP STORAGE NETWORKING

NetApp Virtual Interfaces

A virtual network interface (VIF) is a mechanism that supports aggregation of network interfaces into one logical interface unit. Once created, a VIF is indistinguishable from a physical network interface. VIFs are used to provide fault tolerance of the network connection and in some cases higher throughput to the storage device.

Multimode VIFs are compliant with IEEE 802.3ad. In a multimode VIF, all physical connections in the VIF are simultaneously active and can carry traffic. This mode requires all interfaces to be connected to a switch that supports trunking or aggregation over multiple port connections. The switch must be configured to understand that all of the port connections share a common MAC address and are part of a single logical interface.

Figure 6 is an example of a multimode VIF. Interfaces e0, e1, e2, and e3 are part of the MultiTrunk1 multimode VIF. All four interfaces in the MultiTrunk1 multimode VIF are active.

Figure 6) Multi-mode VIF.

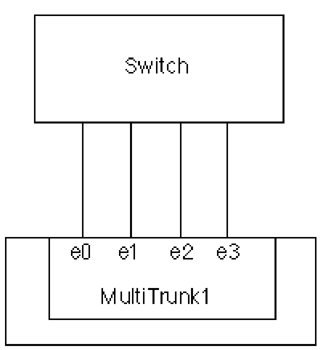

In a single-mode VIF, only one of the physical connections is active at a time. If the storage controller detects a fault in the active connection, a standby connection is activated. No configuration is necessary on the switch to use a single-mode VIF, and the physical interfaces that constitute the VIF do not have to connect to the same switch.

Note: IP load balancing is not supported on single-mode VIFs.

Figure 7 illustrates an example of a single-mode VIF. In this figure, e0 and e1 are part of the SingleTrunk1 single-mode VIF. If the active interface, e0, fails, the standby e1 interface takes over and maintains the connection to the switch.

Figure 7) Single-mode VIF.

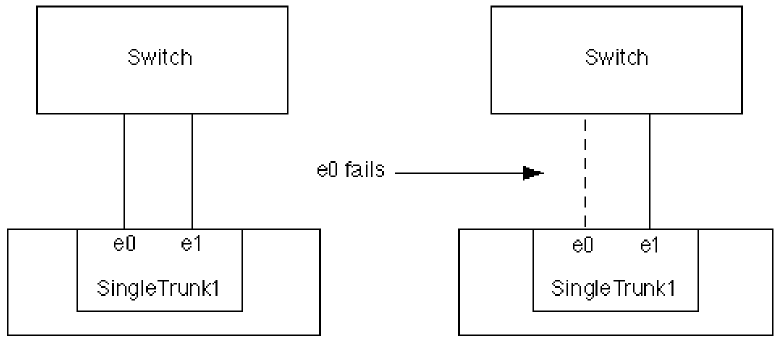

It is also possible to create second-level single or multimode VIFs. By using second-level VIFs, you can take advantage of both the link aggregation features of a multimode VIF and the failover capability of a single-mode VIF.

In the configuration shown in Figure 8, two multimode VIFs are created, each connected to a different switch. A single-mode VIF is then created composed of the two multimode VIFs. In normal operation, traffic flows over only one of the multimode VIFs, but, in the event of an interface or a switch failure, the storage controller moves the network traffic to the other multimode VIF. For more information on the different types of VIFs, see the Data ONTAP Network Management Guide available at the NetApp Support site.

Figure 8) Second-level single or multimode VIFs.

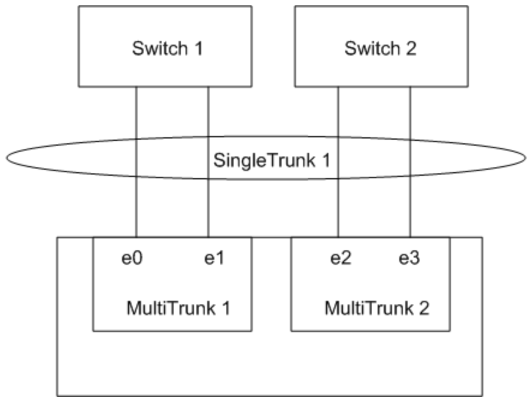

iSCSI Traffic Security

NetApp storage controllers also allow the restriction of the iSCSI protocol to specific interfaces and/or VLAN tags. These simple configuration settings have an enormous effect on the security and availability of IP-based host disks.

ETHERNET SWITCH CONNECTIVITY

An IP storage infrastructure provides the flexibility to connect to storage in different configurations, depending on the needs of the environment. A basic architecture can provide a single nonredundant link to a physical disk, suitable for storing ISO images, or various backups. A redundant architecture, suitable for most production environments, has multiple links, providing failover for switches and network interfaces. Link-aggregated and load-balanced environments make use of multiple switches and interfaces simultaneously to provide failover and additional overall throughput for the environment. Some Ethernet switch models support “stacking,” in which multiple switches are linked by a high-speed connection to allow greater bandwidth between individual switches. A subset of the stackable switch models supports “cross-stack EtherChannel” trunks, in which interfaces on different physical switches in the stack are combined into an 802.3ad EtherChannel trunk that spans the stack. The advantage of cross-stack EtherChannel trunks is that they can eliminate the need for additional passive links that are accessed only during failure scenarios in some configurations.

NETAPP FIBRE CHANNEL STORAGE NETWORKING

Best Practice
• NetApp recommends that the storage controllers have two or more target ports configured to two separate fabrics to make multiple paths available to the Oracle VM Servers. • Have at least two FC HBA ports available for storage connectivity paths on the Oracle VM Server.

FIBRE CHANNEL MULTIPATHING OPTION

NetApp clustered storage systems have an option known as cluster failover mode (cfmode) that defines how Fibre Channel ports behave during failover in an active-active configuration. Selecting the right cfmode is critical to having your LUNs accessible and optimizing your storage system's performance in the event of a failover. If you deploy storage solutions that provide storage for an Oracle VM environment, NetApp strongly recommends that the cfmode be set to single system image (SSI) because this provides LUNs accessibility across all storage ports. NetApp also strongly recommends using Data ONTAP version 7.3 or higher.

To verify the current cfmode using the NetApp console, complete the following steps:

Step	Action
1	Log in to the NetApp console via either SSH, telnet, or console connection.
2	Type: `fcp show cfmode`
3	If cfmode needs to be changed to SSI, type: `priv set advanced`
4	Type: **`fcp set cfmode`** *<mode type>* For more information about the different cfmodes available and the impact of changing a cfmode, see section 8 in the *Data ONTAP Block Management Guide*.

ORACLE VM IP NETWORK CONFIGURATION

Best Practices
• Bond multiple NICs in the OVM Server for the IP storage access path. • Use separate bonded NIC groups for IP storage access and server management.

Oracle VM includes the same native bonding module that is common across all Enterprise Linux 5.x distributions. The native bonding can be implemented in many fashions as indicated by the “mode” in the configuration file. Two of the common values of mode used are:

Mode 0 - balance-rr- Round-robin policy: Transmit packets in sequential order from the first available slave through the last. This mode provides load balancing and fault tolerance.

Mode 1 - active-backup - Active-backup policy: Only one slave in the bond is active. A different slave becomes active if and only if the active slave fails. The bond's MAC address is externally visible on only one port (network adapter) to avoid confusing the switch.

The active-backup policy (mode 1) is the preferred mode for Oracle VM.

STEPS FOR CREATING NETWORK BONDING IN OVM SERVER

Step	Description	Action
1	Disable Xen Script	In the `/etc/xen/xend-config.sxp` Oracle VM config file comment out the network startup script: `# (network-script network-bridges)`

2	Configure System Network Interface	Create a bond device - bond0 and enslave two NIC adapters.
		Create a bond0 device file under `/etc/sysconfig/network-scripts/` named `ifcfg-bond0`
		Ifcfg-bond0
		`DEVICE=bond0`
		`ONBOOT=yes`
		`USERCTL=no`
		`BRIDGE=xenbr0`

		Create a `ifcfg-xenbr0` file under `/etc/sysconfig/network-scripts/`
		Ifcfg-xenbr0
		`DEVICE=xenbr0`
		`TYPE=Bridge`
		`IPADDR=XX.XX.XX.XX`
		`NETMASK=XX.XX.XX.XX`
		`NETWORK=XX.XX.XX.XX`
		`BROADCAST=XX.XX.XX.XX`
		`ONBOOT=yes`

		Enslave devices eth0 and eth1 to the bond0 device.
		Ifcfg-eth0
		`DEVICE=eth0`
		`ONBOOT=yes`
		`MASTER=bond0`
		`SLAVE=yes`
		`USERCTL=no`
		
		Ifcfg-eth1
		`DEVICE=eth1`
		`ONBOOT=yes`
		`MASTER=bond0`
		`SLAVE=yes`
		`USERCTL=no`

Step	Description	Action

3	Bond Configuration in the System	In the `/etc /modprobe.conf` configuration file add the following lines:
		`alias bond0 bonding`
		`options bonding miimon=100 mode=1 primary=eth<n>`
		
		Where `eth<n>` can be replaced with either `eth1` or `eth0` depending on which adapter we want to use as the primary.
		
		For an OVM Server pool/cluster with `m` number of OVM Servers connected to the NetApp storage through two network switches (corresponding to `eth0` and `eth1`), on one-half of the servers (`m/2`) use `eth0` as the primary device and on the other half use `eth1` as the primary device.

3.3 NETAPP STORAGE PROVISIONING BEST PRACTICES FOR ORACLE VM SERVER

AGGREGATES

An aggregate is NetApp's virtualization layer, which abstracts physical disks from logical data sets that are referred to as flexible volumes. Aggregates are the means by which the total IOPs available to all of the physical disks are pooled as a resource. This design is well suited to meet the needs of an unpredictable and mixed workload.

Best Practice

NetApp recommends that whenever possible a small aggregate be used as the root aggregate. This aggregate stores the files required for running and providing GUI management tools for the FAS system.

The remaining storage should be placed into a small number of large aggregates. The overall disk I/O from Oracle VM environments is traditionally random by nature, so this storage design gives optimal performance because a large number of physical spindles are available to service I/O requests.

On smaller FAS arrays, it may not be practical to have more than a single aggregate, due to the restricted number of disk drives on the system. In these cases, it is acceptable to have only a single aggregate.

FLEXIBLE VOLUMES

Flexible volumes contain either LUNs (FC or iSCSI) or virtual disk files that are accessed by Oracle VM Servers.

Best Practice

NetApp recommends one-to-one alignment of the Oracle VM storage repository to flexible volumes.

This design offers an easy means to understand the Oracle VM storage repository layout when viewing the storage configuration from the FAS array. This mapping model also makes it easy to implement Snapshot backups and SnapMirror® replication policies at the Oracle VM storage repository level, because NetApp implements these storage-side features at the flexible volume level.

SNAPSHOT RESERVE

NetApp flexible volumes should be configured with the snap reserve set to 0 and the default Snapshot schedule disabled. NetApp Snapshot copies are covered in detail in section 3.6.

To set the volume options for Snapshot copies to the recommended setting, enter the following commands in the FAS system console:

```
snap sched <vol-name> 0 0 0
snap reserve <vol-name> 0
```

STORAGE PROVISIONING

Figure 9 depicts the high-level architecture of provisioning NetApp storage to Oracle VM Servers.

For an end-to-end example of provisioning NetApp storage to the Oracle VM Server for all three protocols: NFS, FC, and iSCSI, see section 4.1.

Figure 9) Storage provisioning.

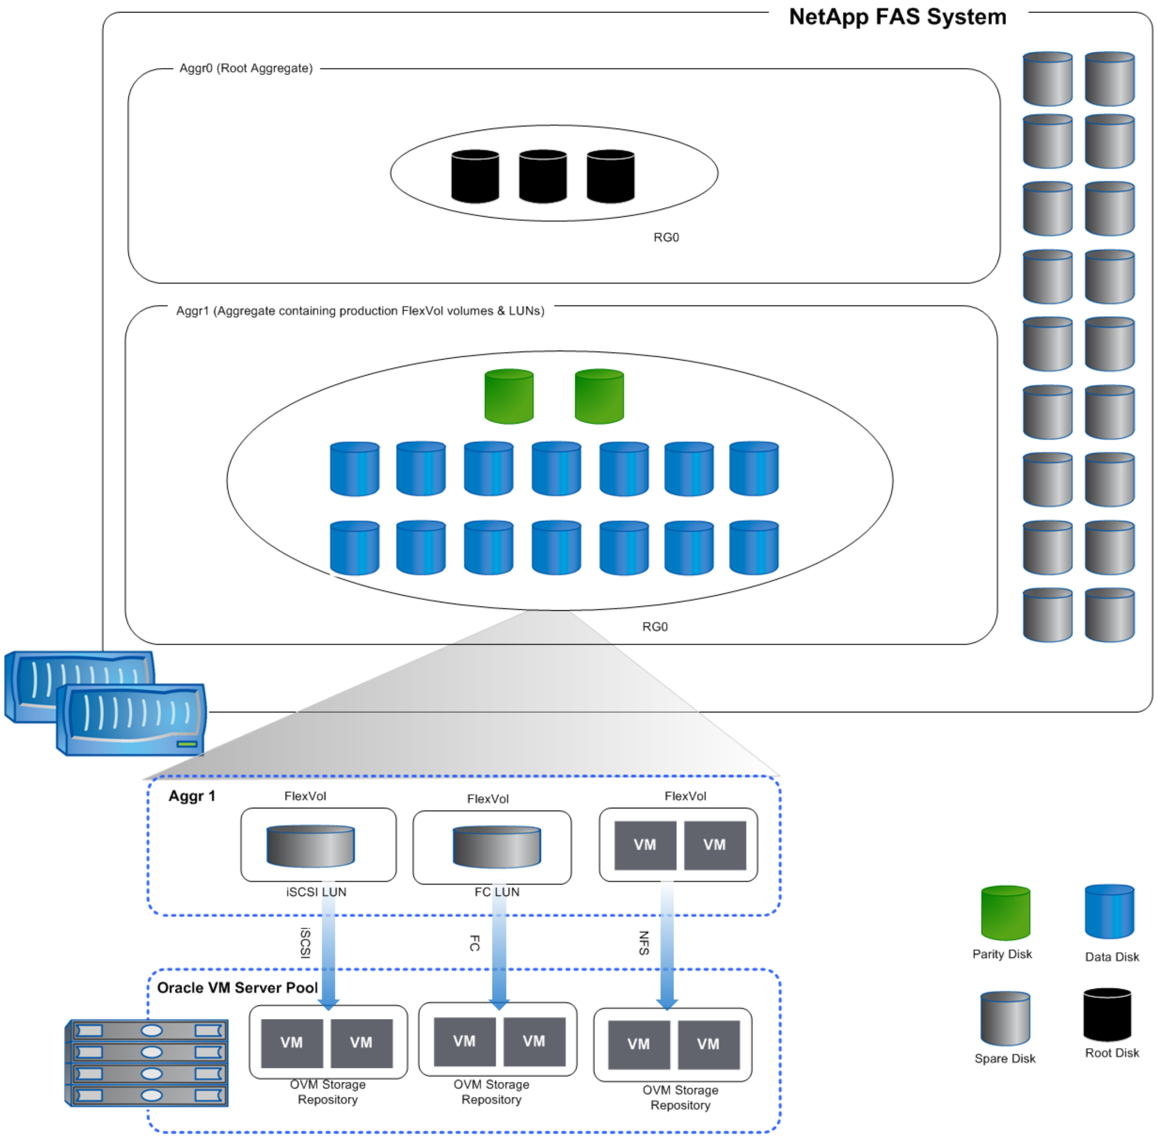

3.4 ORACLE VM VIRTUAL MACHINE STORAGE BEST PRACTICES—CREATE, IMPORT, LIVE MIGRATE

VM CREATION

Oracle VM virtual machines can be created on the storage repositories configured on the storage provisioned earlier.

Oracle VM provides various ways of creating virtual machines. Some common methods are:

- Using the `virt-install` command
- Using the Oracle VM templates

BEST PRACTICES

1. If using the `virt-install` command to create VMs, specify the storage repository created on the NetApp storage (NFS volume or FC/iSCSI LUN) as the path when prompted:

   `What would you like to use as the disk (path)?`

2. Oracle VM templates provide the most convenient and a faster way for deploying VMs. The Oracle VM templates can be simply copied to the following locations, and then expanded and used to create the desired VMs:
   - For OVM 2.1.5 or below, use one of the following:

   `/OVS/<repository uuid>/running_pool`

   or

   `/OVS/<repository uuid>/seed_pool`

   - For OVM 2.2, use one of the following:

   `/var/ovs/mount/<repo uuid>/running_pool`

   or

   `/var/ovs/mount/<repo uuid>/seed_pool`

3. Different ways to add additional storage to the VMs (that is, VMs created through Oracle VM templates):
   a. Storage provisioned from the virtual machine itself:
      i. NFS volumes from the NetApp storage can be directly mounted from the virtual machines.
      ii. iSCSI LUNs from the NetApp storage can be directly accessed from the virtual machines using the software iSCSI initiator.

Best Practice
NetApp SnapDrive® products can be used for easier management of options [i] and [ii] above.

   b. Storage provisioned from the OVM Server to the virtual machine:
      i. FC or iSCSI shared storage LUN mapped to the Oracle VM Server can be attached to the virtual machine by modifying the `vm.cfg` file.

      First format the storage device (LUN):

```
mkfs –t ext3 /dev/mapper/<mpathdevice>
```

Or

```
mkfs –t ocfs2 /dev/mapper/<mpathdevice>
```

Then add entries in the `vm.cfg` file:

```
disk = [ 'file://OVS//running_pool//OVM_EL4U8_X86_64_PVHVM_4GB
```

```
//System.img,hda,w', 'phy:/dev/mapper/<mpathdevice>,xvdc,w' ]
```

ii. An image file created on an existing OVM storage repository (NFS, FC, or iSCSI storage repository) can be attached to the virtual machine after modifying the `vm.cfg` file. For creating the image file, the **dd** command can be used.

First create an image file inside the existing OVM storage repository:

```
dd if=/dev/zero of=/OVS/running_pool/<vmname>/AppData.img
oflag=direct bs=”xxxx” count=”yyy”
```

Then add entries in the `vm.cfg` file:

```
disk = ['file://OVS/running_pool/OVM_EL5U3_X86_64_PVM_4GB//
System.img,xvda,w','file://OVS/running_pool/
OVM_EL5U3_X86_64_PVM_4GB/AppData.img,xvdb,w' ]
```

IMPORTING VMS TO THE ORACLE VM MANAGER

If the shared storage is mounted as an Oracle VM storage repository (`/var/ovs/mount/<Storage Repository UUID> or /OVS/<Storage Repository UUID)>`) or Oracle VM cluster root (`/OVS`) using the `repos.py` or `ovs-makerepo` command, all the VMs created inside the `running_pool` directory can be directly imported into the Oracle VM Manager.

Best Practice
If planning to manage the VMs through Oracle VM Manager, always mount the NetApp shared storage as the Oracle VM storage repository or Oracle VM cluster root using `/opt/ovs-agent-2.3/utils/repos.py` (for OVM v2.2) **or** the ovs-makerepo utility (for OVM v2.1.5 or earlier). All the VMs to be imported and managed by the Oracle VM Manager must reside inside the `running_pool` directory of the storage repository (`/var/ovs/mount/<Storage Repository UUID>/running_pool`’–for OVM v2.2 or `/OVS/<Storage Repository UUID>/running_pool`’– for OVM v2.1.5 or earlier) or cluster root (‘`/OVS/running_pool`’).

If the VM is not located inside the running_pool directory (but resides in a user-defined mounted point), the procedure of importing the VM to the Oracle VM Manager is not automatic—it needs the steps mentioned below.

4. Create a directory for the virtual machine (`vm1`) inside `/OVS/running_pool` or `/var/ovs/mount/<Storage Repository UUID>/running_pool` (for OVM v2.2) or `/OVS/<Storage Repository UUID>/running_pool` (for OVM v2.1.5 or earlier).
5. In the directory just created:
   c. Create a soft link for the virtual machine image file in the shared storage path.
   d. Create a config file called `vm.cfg`.

```
[root@ovmserver vm1]# pwd
/OVS/running_pool/vm1
[root@ovmserver vm1]# ll
total 1
lrwxrwxrwx 1 root root  39 Jun 18 14:07 system.img -> /OVS/netapp/nfsdatastore/vm1/system.img
-rw-r--r-- 1 root root 215 Jun 18 13:54 vm.cfg
[root@ovmserver vm1]# cat vm.cfg
# Automatically generated xen config file
name = "vm1"
memory = "1024"
disk = [ 'file:/OVS/running_pool/vm1/system.img,xvda,w', ]
bootloader="/usr/bin/pygrub"
vcpus=1
on_reboot   = 'restart'
on_crash    = 'restart'
[root@ovmserver vm1]#
```

6. In the OVM Manager, click Resources > Virtual Machine Images >Import >Select from Server Pool (Discover and Register).
7. From the drop-down list, select the Server Pool name and Virtual Machine Image (vm1) and click Confirm.

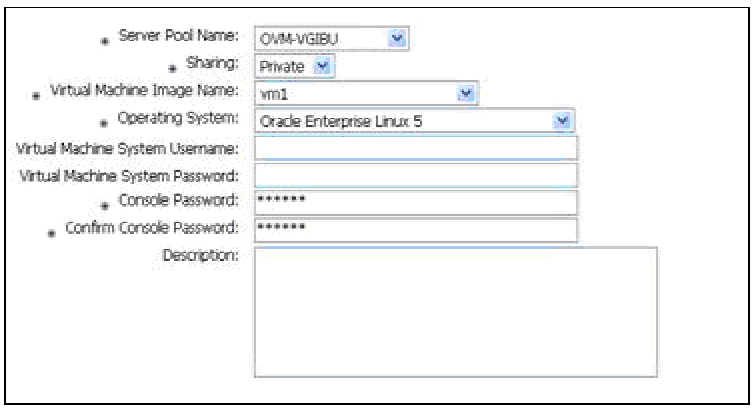

8. Now the virtual machine `vm1` can be managed and controlled through the Oracle VM Manager.

VM LIVE MIGRATION USING NETAPP SHARED STORAGE

Virtual machines residing on shared OVM Server repositories created on NetApp storage can be live migrated from one OVM Server to another using Oracle VM Manager or from the Oracle VM Server command line. Both the OVM Servers need to be inside the same OVM Server pool.

From OVM Manager:

Select the VM to be migrated, from More Action, select Live Migration from the drop-down list. Click Confirm.

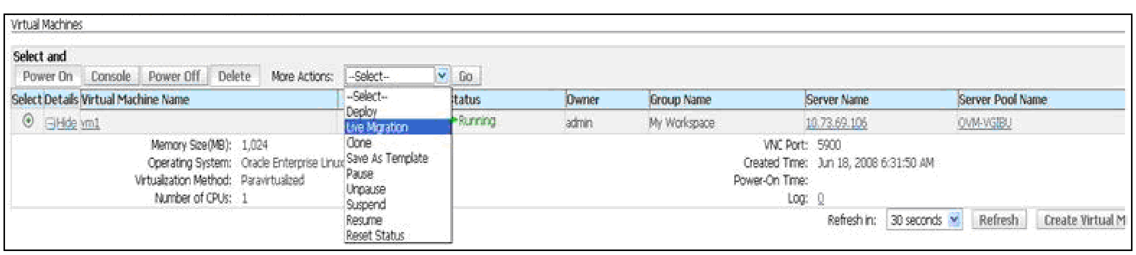

From Oracle VM Server command line:

```
xm migrate <vmname> <destination-OVM-Server> --live
```

```
[root@ovmserver2 /]#
[root@ovmserver2 /]# xm migrate vm1 ovmserver --live
[root@ovmserver2 /]#
```

3.5 INCREASING STORAGE UTILIZATION—DEDUPLICATION, THIN CLONING, THIN PROVISIONING

DEDUPLICATION

One of the most important features of Oracle VM is its ability to rapidly deploy new virtual machines from stored VM templates. A VM template includes a VM configuration file and one or more virtual disk files, which include an operating system, common applications, and patch files or system updates. Deploying from templates saves administrative time by copying the configuration and virtual disk files and registering this second copy as an independent VM. By design, this process introduces duplicate data for each new VM deployed.

Figure 10 shows an example of typical storage consumption in a normal Oracle VM deployment:

Figure 10) Traditional array.

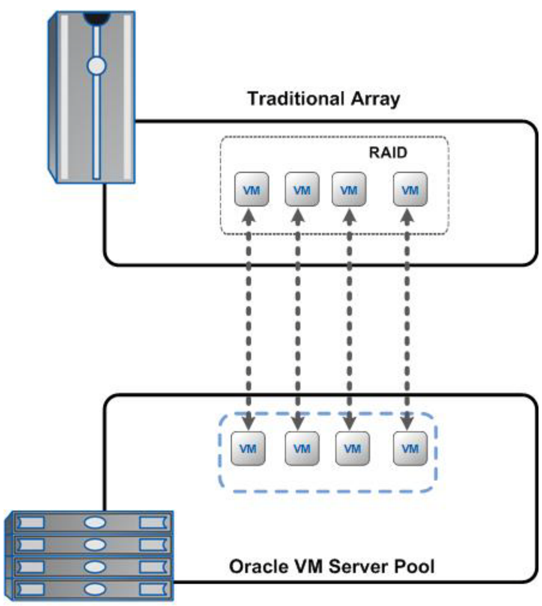

Deduplication technologies from NetApp assist Oracle VM deployments in eliminating duplicate data in their environment, enabling greater storage utilization on the production environment.

NetApp deduplication technology enables multiple virtual machines in an Oracle VM environment to share the same physical blocks on storage. Deduplication can be seamlessly introduced into a virtual infrastructure without having to make any changes to the Oracle VM administration, practices, or tasks. Figure 11 shows an example of the impact of deduplication on storage consumption in an Oracle VM deployment.

Figure 11) NetApp FAS array.

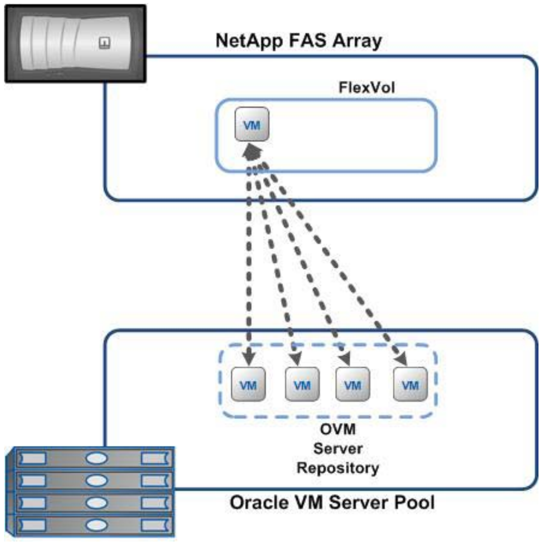

Deduplication is enabled on a volume; the amount of data deduplication realized is based on the commonality of the data stored in a deduplication-enabled volume.

Best Practice
To leverage the largest storage savings when using deduplication, NetApp recommends grouping similar operating systems and similar applications into the same deduplication-enabled volumes.

DEDUPLICATION CONSIDERATIONS WITH ISCSI AND FC LUNS

Storage savings are apparent if deduplication is enabled while provisioning LUNs. However, the default behavior of a LUN is to reserve an amount of storage equal to the provisioned LUN. This design means that although the storage array reduces the total amount of capacity consumed, any gains made with deduplication are, for the most part, unrecognizable, because the space reserved for LUNs is not reduced.

To benefit from the storage savings of deduplication with LUNs, a LUN must be thin provisioned. NetApp thin provisioning is covered later in this section.

In addition, although deduplication reduces the amount of consumed storage, this benefit is not seen directly by the Oracle VM administrative team, because their view of the storage is at a LUN level, and LUNs always represent their provisioned capacity, whether they are traditional or thin provisioned.

DEDUPLICATION CONSIDERATIONS WITH NFS

Unlike with LUNs, when deduplication is enabled with NFS, the storage savings are both immediately available and also recognizable by the Oracle VM administrative team. No special considerations are required for its usage.

Best Practice
For NetApp deduplication best practices, including scheduling and performance considerations see [TR-3505: NetApp Deduplication for FAS and V-Series Deployment and Implementation Guide].

For a step-by-step procedure for applying NetApp deduplication to Oracle VM Server repositories, refer to section [4.2].

NETAPP THIN PROVISIONING

Oracle VM provides an excellent means to increase the hardware utilization of physical servers. By increasing hardware utilization, the amount of hardware in a data center can be reduced, thus lowering the cost of data center operations. In a typical Oracle VM environment, the process of migrating physical servers to virtual machines does not reduce the amount of data stored or the amount of storage provisioned. By default, server virtualization does not have any impact on improving storage utilization (and, in many cases, it may have the opposite effect).

In traditional storage provisioning, the storage is already allocated and assigned to a server, or, in the case of an Oracle VM, a virtual machine. It is also common practice for server administrators to overprovision storage to avoid running out of storage and avoid the associated application downtime when expanding the provisioned storage. Although no system can be run at 100% storage utilization, there are methods of storage virtualization that allow administrators to oversubscribe storage in the same manner as with server resources (such as CPU, memory, networking, and so on). This form of storage virtualization is referred to as thin provisioning.

Thin provisioning provides storage on demand; traditional provisioning preallocates storage. The value of thin-provisioned storage is that storage is treated as a shared resource pool and is consumed only as each individual VM requires it. This sharing increases the total utilization rate of storage by eliminating the unused but provisioned areas of storage that are associated with traditional storage. The drawback of thin provisioning and oversubscribing storage is that (without the addition of physical storage) if every VM requires its maximum possible storage at the same time, there will not be enough storage to satisfy the requests.

It is important to note that the benefits of NetApp thin provisioning can be realized across all kinds of shared storage repositories in an Oracle VM environment, this is, NFS, iSCSI, or FC.

Best Practices

1. If using NFS storage, NetApp flexible volumes are thin-provisioned by default. No extra configuration steps are necessary.

   However, when using iSCSI or FC storage, make sure that the Space Reserved checkbox in the LUN wizard is not selected.
2. When enabling NetApp thin provisioning, also configure the storage management policies on the volumes that contain the thin-provisioned LUNs. These policies aid in providing the thin-provisioned LUNs with storage capacity as they require it.

   The important policies are—automatic sizing of a volume, automatic Snapshot copy deletion and LUN fractional reserve.

Volume Autosize is a policy-based space-management feature in Data ONTAP that allows a volume to grow in defined increments up to a predefined limit if the volume is nearly full. For Oracle VM environments, NetApp recommends setting this value to ON. Doing so requires setting the maximum volume and increment size options. To enable these options, do the following:

Step	Action
1	Log in to the NetApp console.
2	Set the volume autosize policy by typing: `vol autosize <vol-name> [-m <size>[k|m|g|t]] [-i <size>[k|m|g|t]] on`

Snapshot Auto Delete is a policy-based space-management feature that automatically deletes the oldest Snapshot copies on a volume when that volume is nearly full. For Oracle VM environments, NetApp recommends setting this value to delete Snapshot copies at 5% of available space. In addition, you should set the volume option to have the system attempt to grow the volume before deleting Snapshot copies. To enable these options, do the following:

Step	Action
1	Log in to the NetApp console.
2	Set the Snapshot copy autodelete policy by typing: **`snap autodelete`** *`<vol-name>`* **`commitment try trigger volume target_free_space 5 delete_order oldest_first`**
3	Set the volume autodelete policy by typing: **`vol options`** *`<vol-name>`* **`try_first volume_grow`**

LUN Fractional Reserve is a policy that is required when you use NetApp Snapshot copies on volumes that contain Oracle VM LUNs. This policy defines the amount of additional space reserved to guarantee LUN writes if a volume becomes 100% full. For Oracle VM environments in which `Volume Auto Size` and `Snapshot Auto Delete` are used, NetApp recommends setting this value to 0%. Otherwise, leave this setting at its default of 100%. To enable this option, do the following:

Step	Action
1	Log in to the NetApp console.
2	Set the volume Snapshot fractional reserve by typing: **`vol options`** *`<vol-name>`* **`fractional_reserve 0`**

NETAPP VIRTUAL CLONING—VOLUME, LUN, AND FILE LEVEL CLONING

Virtual cloning technologies from NetApp can be used for rapidly provisioning zero-cost Oracle VM virtual machine clones.

There can be different flavors of the virtual cloning technology from NetApp: volume-level cloning (or FlexClone®), LUN-level cloning, and file-level cloning.

Best Practices
1. Depending on the requirement and necessity, decide the design of the cloning methodology (volume, LUN, or file level) and other NetApp technologies (like deduplication) to be applied. There can be many possibilities for achieving the same end result. 2. File-level cloning can be used only if the cloned VMs need to reside on NFS storage repositories. 3. While using file-level cloning, make sure that the source/golden VM from where the clones are to be created does not share any blocks with others.

Figures 15 through 16 depict some common ways of using NetApp thin-cloning technologies with Oracle VM Server for an NFS storage repository. For step-by-step process, see section 4.3.

Figure 12) Thin cloning deployment with Oracle VM server.

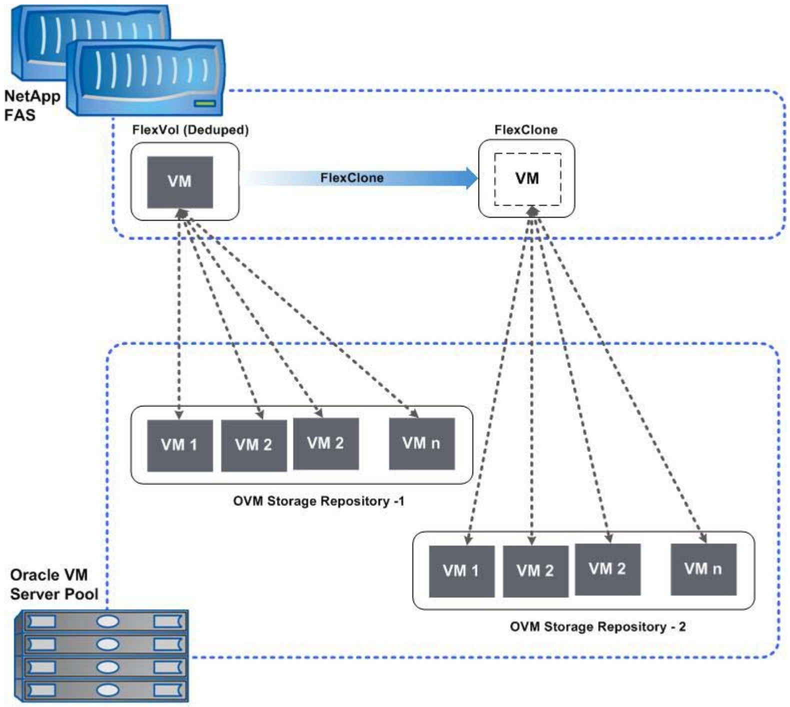

Figure 13) Thin cloning deployment with Oracle VM server.

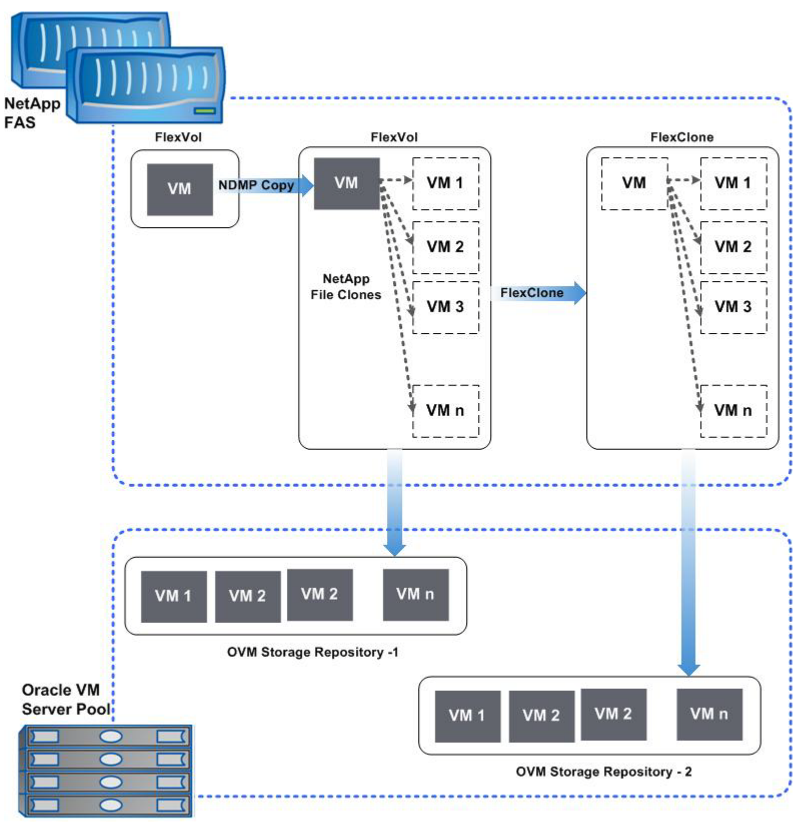

3.6 DISK-BASED SNAPSHOT BACKUPS FOR ORACLE VMS USING NETAPP SNAPSHOT AND SNAPRESTORE

NetApp Snapshot technology can be used to back up and restore the virtual machines, their virtual disks, and Oracle VM Server repositories residing on NetApp shared storage. It can be accessed using any protocol: NFS, iSCSI, or FC.

Note that the Snapshot backup of the OVM server repositories, virtual machine, and their virtual disks residing on NetApp storage (NFS, iSCSI, or FC) will be crash consistent.

NetApp SnapManager along with SnapDrive can be used to take application-consistent backup of the applications running inside the virtual machines. NetApp provides SnapManager products for several enterprise applications and databases including Oracle, MS Exchange, MS SQL, SAP, and so on. For more information about the SnapManager product suite, see http://www.netapp.com/us/products/management-software/.

The following table describes the procedure to take crash-consistent backup of an Oracle VM server repository.

Step	Action
1	Consider that an NFS volume from a NetApp FAS system has been mounted as a repository in the OVM Server. `/opt/ovs-agent-2.3/utils/repos.py  -l` `[ * ] 3e4fcb05-7a0c-4808-a322-f4e205af24b4 => 10.73.68.199:`**`/vol/OVM_NFS`**
2	Create a Snapshot copy of the volume in the NetApp FAS system using NetApp System Manager, FilerView®, or the CLI. The Snapshot copy used in this example is named "`OVMNFSSNAPSHOT.`" To recover the virtual machines on this volume, you can mount this Snapshot copy and recover the individual VM images. Additionally, the Snapshot technology can be seamlessly integrated with the NetApp SnapMirror solution for disaster recovery solutions. 

RESTORING VIRTUAL MACHINE IMAGES FILES OR VIRTUAL DISKS USING SNAP RESTORE

NetApp's `snap restore` command can be used either from NetApp System Manager, FilerView, or the CLI for restoring any virtual disks or images of a virtual machine.

```
snap restore –f –t file –s <Snapshot copy name> /vol/<volume name>/running_pool/<VM directory>/<Virtual Disk file>
```

For example:

```
snap restore –f –t file –s OVMNFSSNAPSHOT
/vol/OVM_NFS/running_pool/OVM_EL5U2_X86_64_ORACLE11G_PVM1/System.img
```

SINGLE FILE RESTORE USING UFS EXPLORER

UFS Explorer is a licensed utility that can be used to browse the contents of the virtual disk (like `System.img`). Any lost file inside the virtual disk can then be copied for using with UFS Explorer.

Figure 14 shows the contents of a `System.img` file (root file system of a DomU) using UFS Explorer.

Figure 14) Contents of the system.img file.

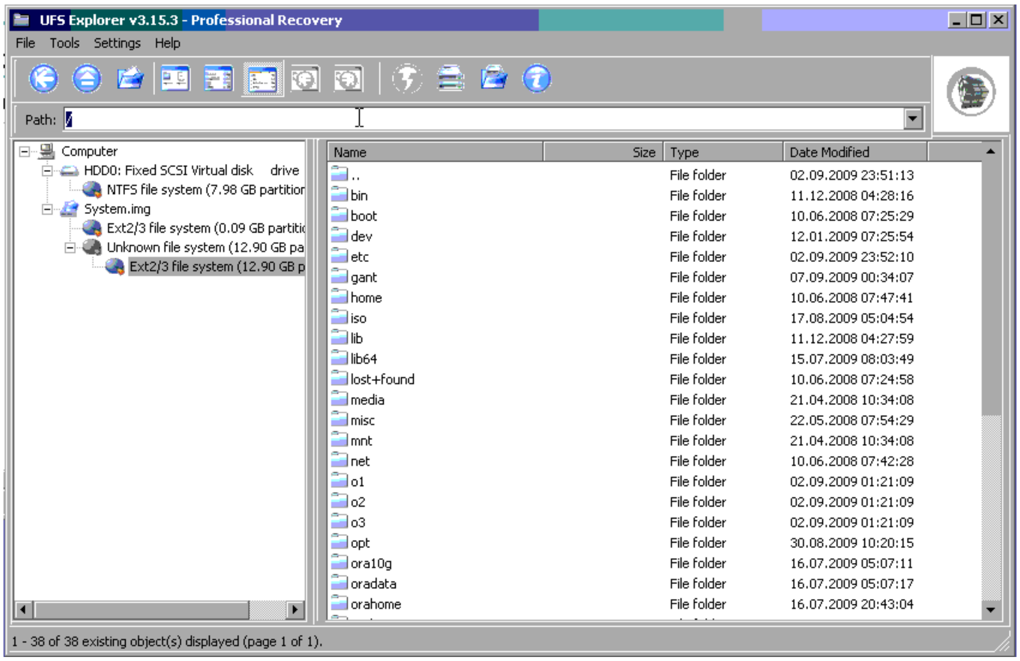

3.7 DISASTER RECOVERY OF AN OVM VIRTUAL INFRASTRUCTURE USING NETAPP SNAPMIRROR

For disaster recovery of an entire Oracle VM infrastructure hosted on NetApp storage, NetApp SnapMirror can be used.

For more information on NetApp SnapMirror, refer to the NetApp online backup and recovery guide at: http://now.netapp.com/NOW/knowledge/docs/ontap/rel732/pdfs/ontap/onlinebk.pdf.

Figure 15 shows a typical NetApp SnapMirror deployment with Oracle VM.

Figure 15) SnapMirror deployment with Oracle VM.

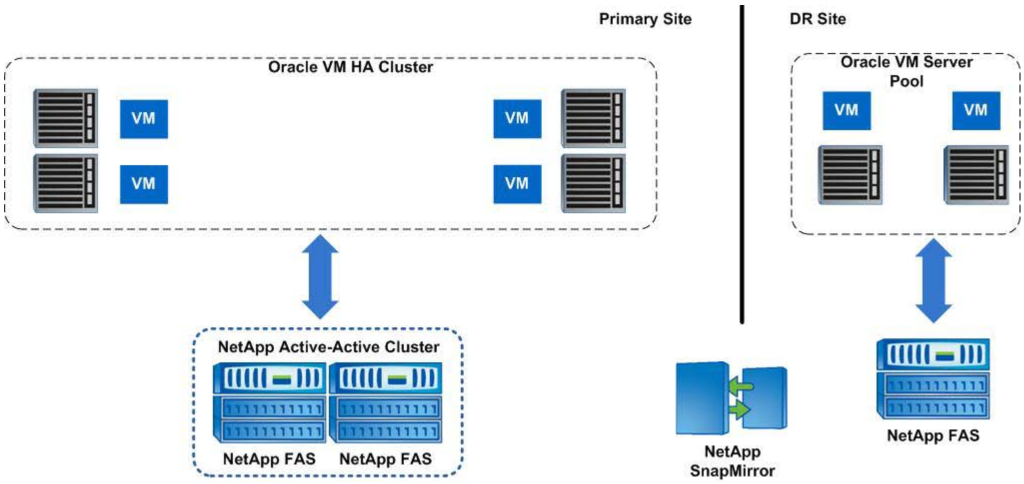

Best Practices
• NetApp SnapMirror Async best practices: [TR-3446: SnapMirror Async Overview and Best Practices Guide]   • NetApp SnapMirror Sync and Semi-Sync best practices: [TR-3326: SnapMirror Sync and SnapMirror Semi-Sync Overview and Design Considerations]

3.8 DESIGNING A HIGHLY AVAILABLE OVM VIRTUAL INFRASTRUCTURE USING OVM HA AND METROCLUSTER

Using Oracle VM HA and NetApp MetroCluster in conjunction can lead to an end-to-end highly available virtual infrastructure.

For more details on NetApp MetroCluster refer to the NetApp Active-Active Configuration Guide at: www.now.netapp.com/NOW/knowledge/docs/ontap/rel732/pdfs/ontap/aaconfig.pdf.

Figure 16 shows a typical NetApp MetroCluster deployment with Oracle VM HA.

Figure 16) Stretch MetroCluster.

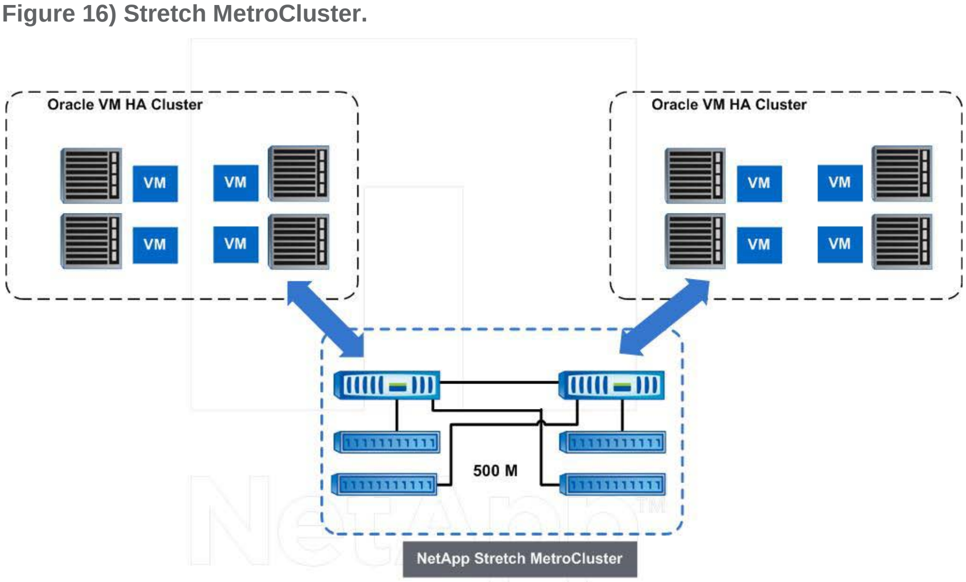

Figure 17) Fabric MetroCluster.

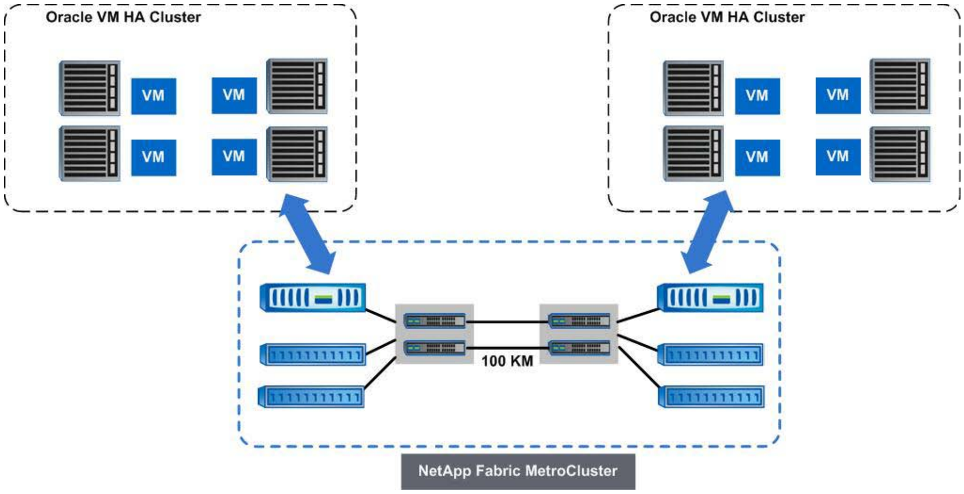

Best Practices
NetApp MetroCluster best practices: http://media.netapp.com/documents/tr-3548.pdf

4 BEST PRACTICES IMPLEMENTATION EXAMPLES

4.1 STORAGE PROVISIONING

CONFIGURING AN AGGREGATE IN NETAPP FAS SYSTEM

1. From NetApp System Manager, start the Create Aggregate Wizard.

2. Choose dual parity as the RAID type and make disk selection automatic.

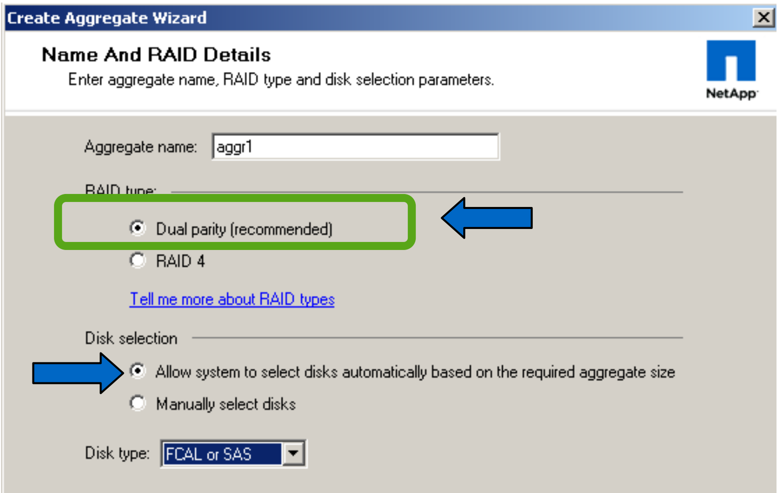

3. Choose the size of the aggregate depending on the number of disks.

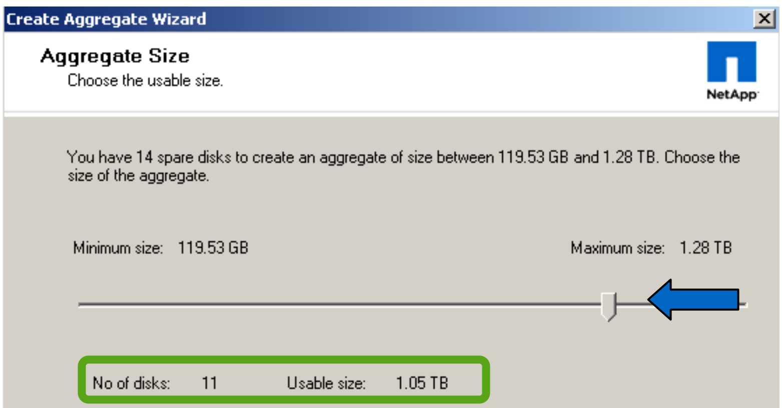

4. Commit the selection to complete the aggregate creation process.

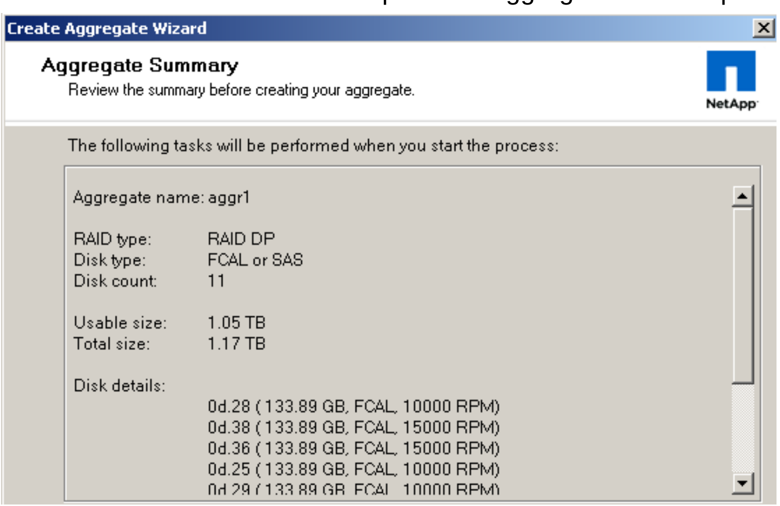

CONFIGURING NFS STORAGE ON ORACLE VM SERVER

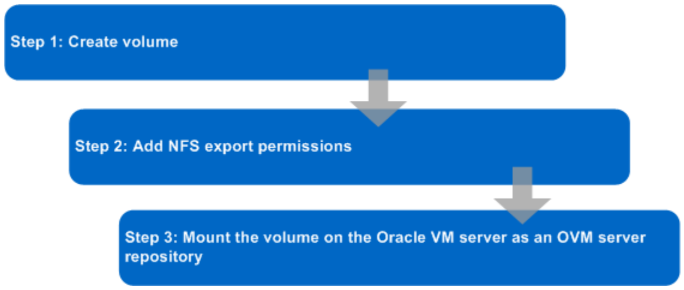

Step 1: Create volume

1. From NetApp System Manager, create the volume for the NFS storage repository.

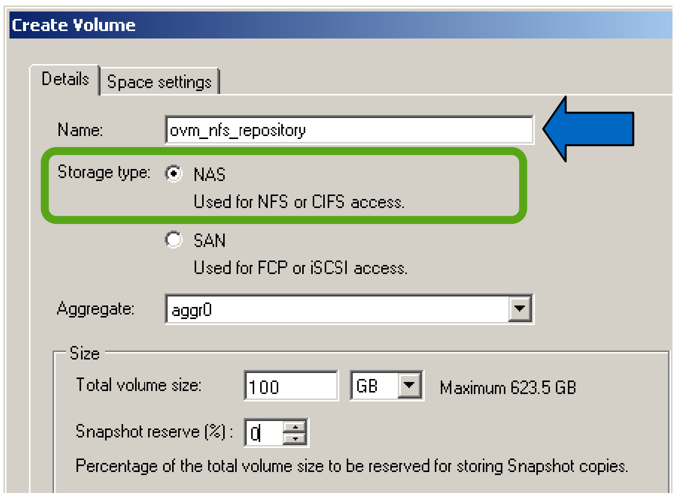

2. To thin-provision the volume and manage space at the aggregate level, in the Space settings tab, set the Guarantee to None.

   Once the volume is created, its properties can be further modified, like the Volume Autogrow settings.

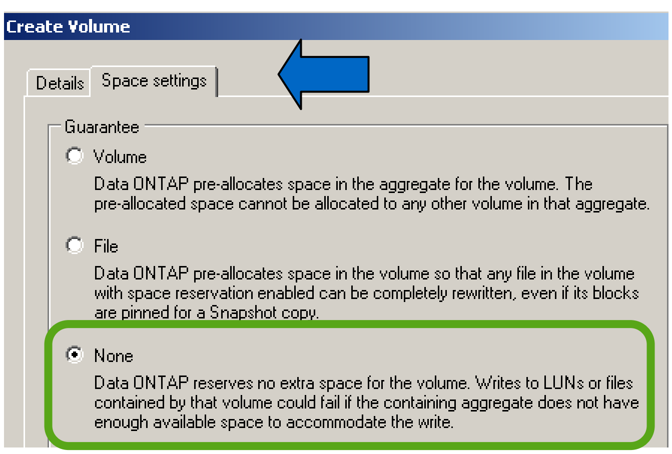

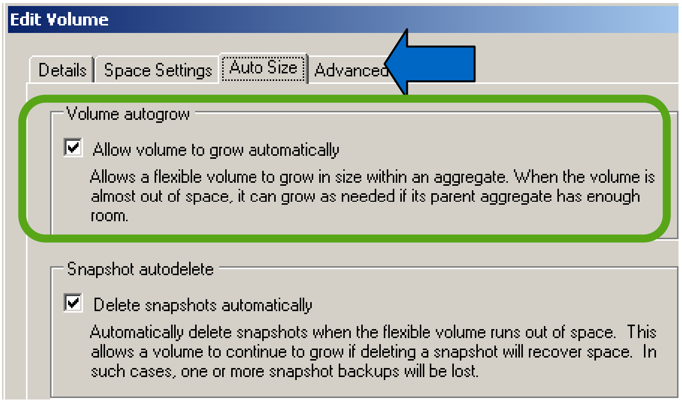

Step 2: Add NFS export permission

1. Click Shares/Exports under Shared Folders in NetApp System Manager.

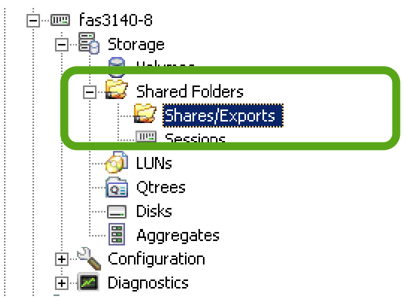

2. Select the volume on the right-hand pane and edit its settings.

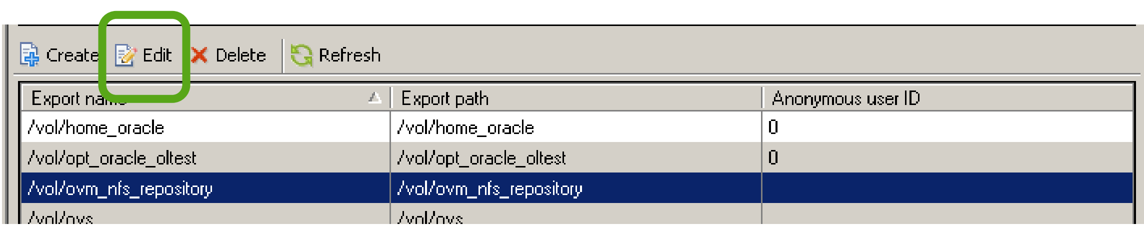

3. Add root access to the Oracle VM Server IP address and set the security setting.

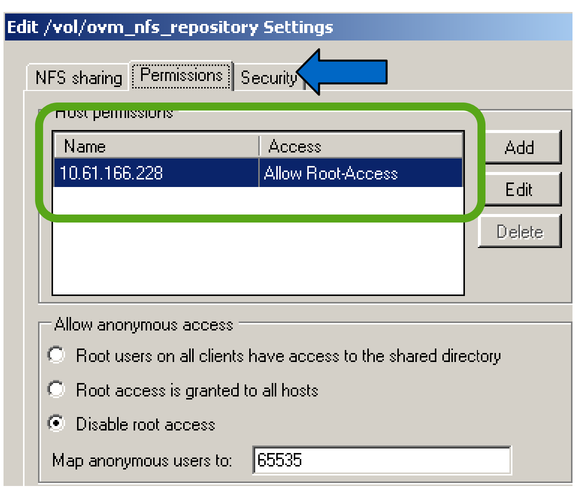

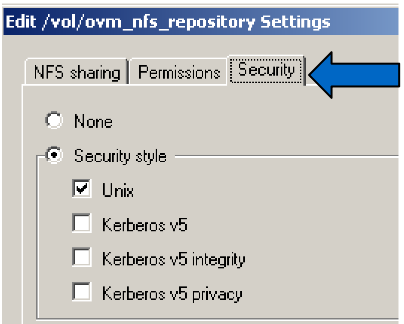

Step 3: Mount the volume on the Oracle VM Server

1. Mount the volume on the Oracle VM Server using either `/opt/ovs-agent-2.3/utils/repos.py` (for OVM v2.2) or the `ovs-makerepo` utility (for OVM v2.1.5 or earlier).

   The volume now can be used as an Oracle VM Server repository.

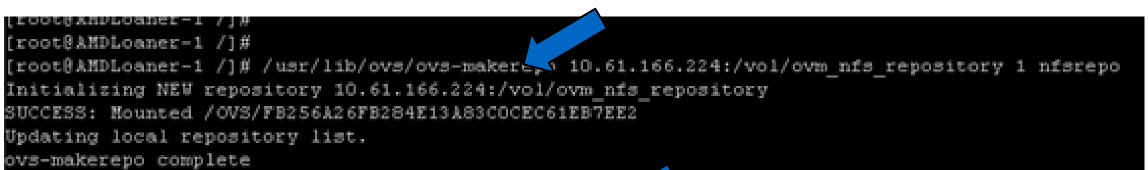

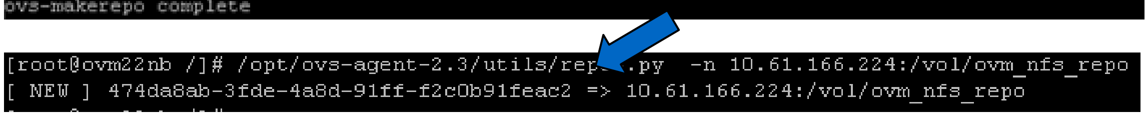

2. Although not a best practice, to mount the volume manually on the Oracle VM server, add the following line to the `/etc/fstab` file:

```
 10.61.166.224:/vol/ovm_nfs_repository    /OVS   nfs
rw,vers=3,rsize=65536,wsize=65536,hard,proto=tcp,timeo=600 0 0
```

CONFIGURING FC SHARED STORAGE ON ORACLE VM SERVER

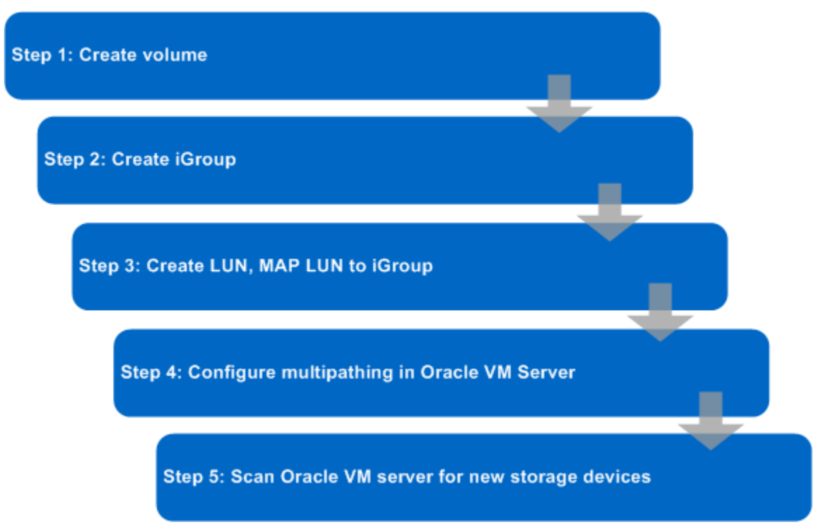

Step 1: Create volume

1. From NetApp System Manager, create the volume that will contain the LUN.

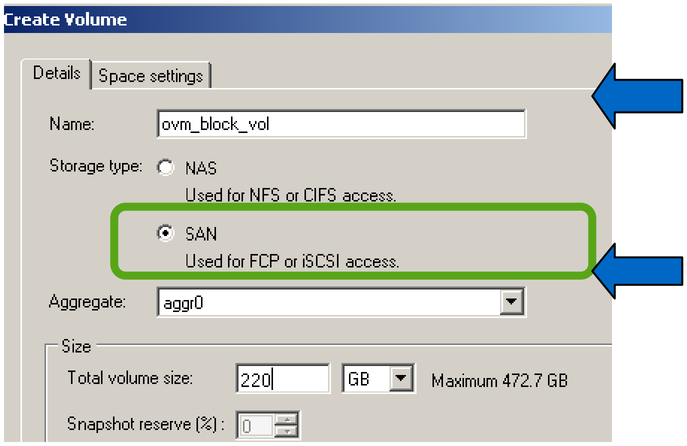

2. To thin-provision the volume and manage space at the aggregate level, in the Space settings tab set the Guarantee to None.

   Once the volume is created, its properties can be further modified, like the Volume Autogrow settings.

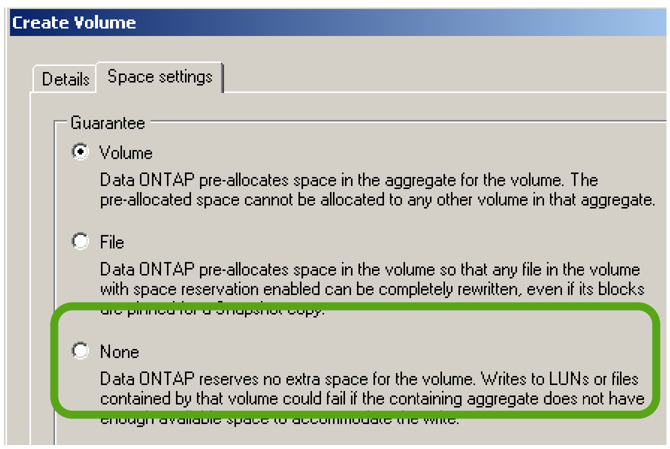

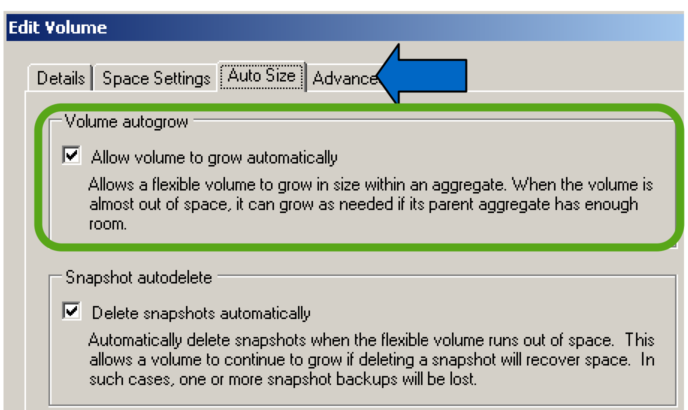

Step 2: Create iGroup

1. From the Oracle VM Server, determine the WWPN numbers of the host HBAs.

   If the Oracle VM Server is configured in a server pool and the storage need to shared among all the nodes of the cluster, WWPN numbers of all the nodes need to be collected so that they can be put into the same iGroup.

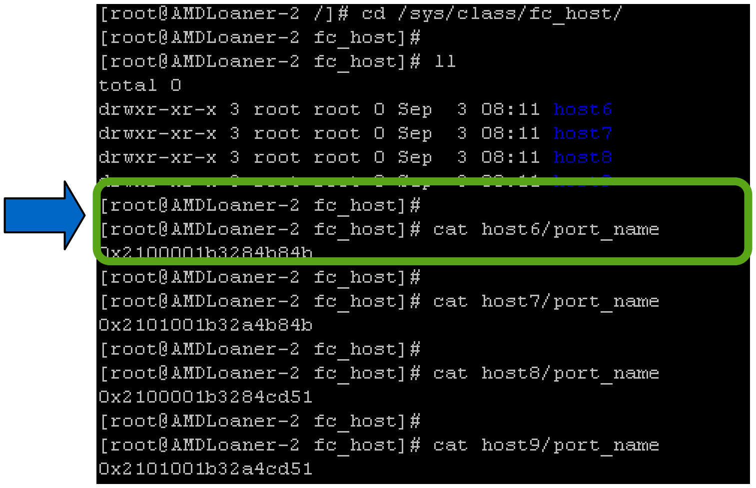

2. Create the iGroups corresponding to these WWPN numbers from NetApp System Manager.

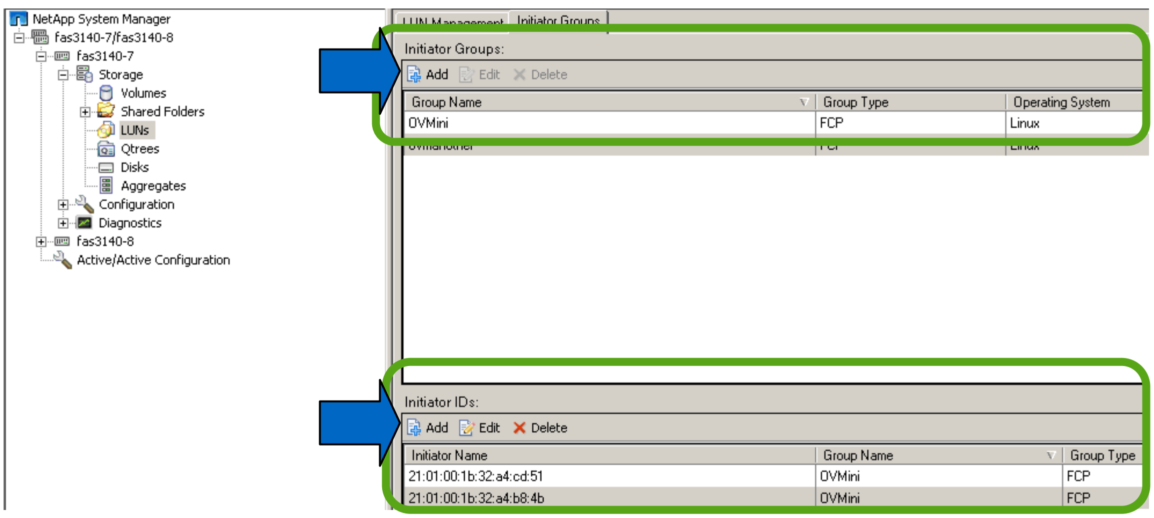

Step 3: Create LUN, and map LUN to iGroup

1. From the LUN Wizard of NetApp System Manager, create the LUN inside the volume created in Step 1 and map it to the iGroup created in Step 2.

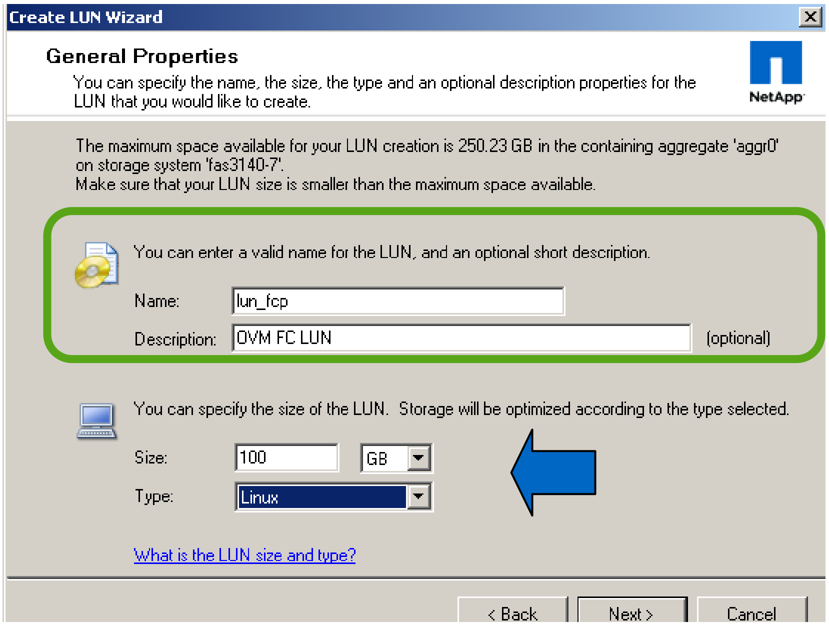

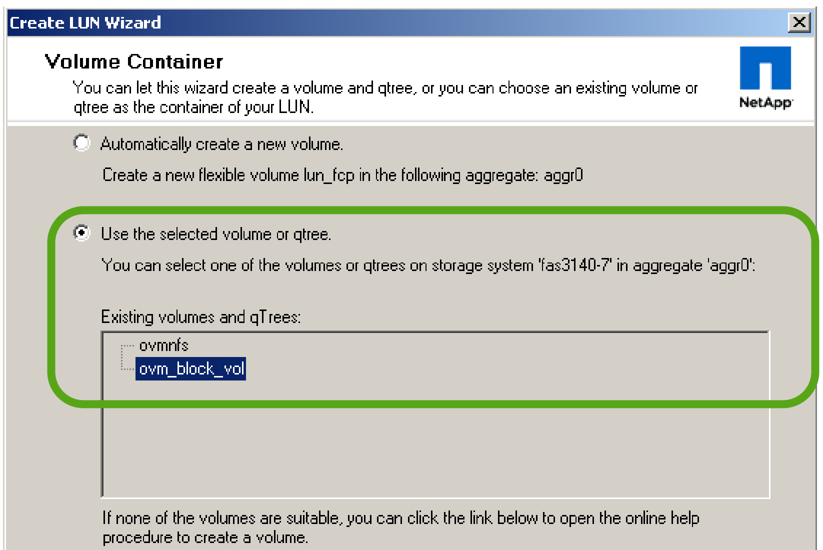

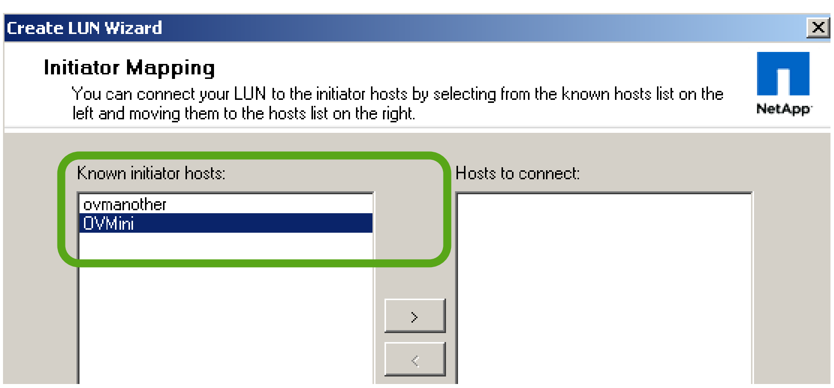

2. Complete the wizard to finish the LUN creation process.

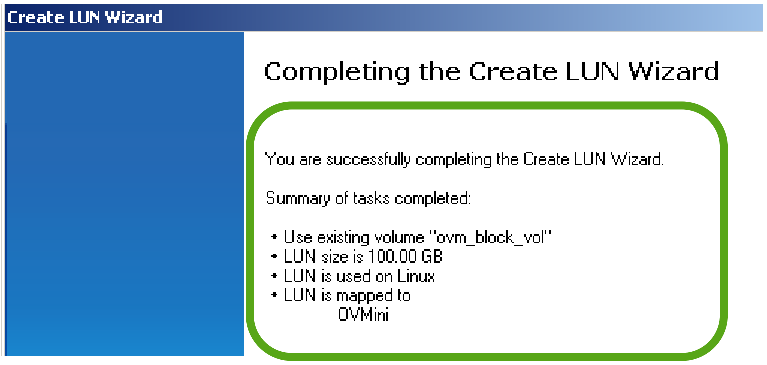

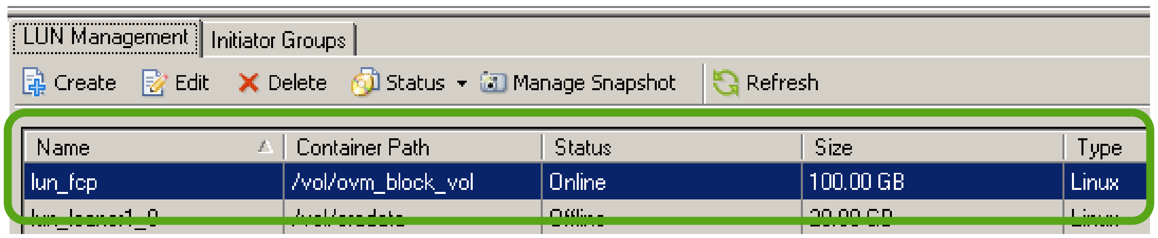

3. If using Thin Provisioning, clear the Space Reserved checkbox in the LUN properties.

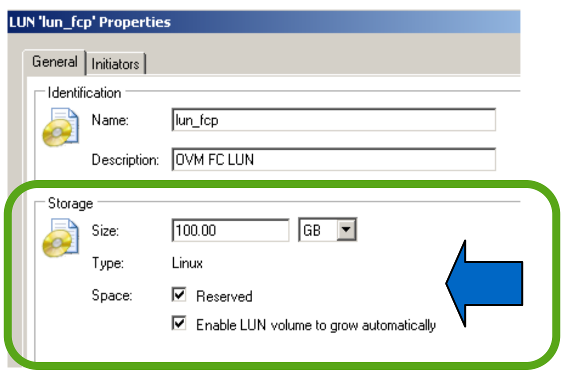

Step 4: Configure multipathing in the Oracle VM Server

1. Use the native multipathing—DM-MP (Device Mapper Multipath)—support provided by the Oracle VM Server to configure multipathing.

   First check if the DM-MP is installed by entering:

   `rpm -q device-mapper`

   Then check if the DM-MP services are running:

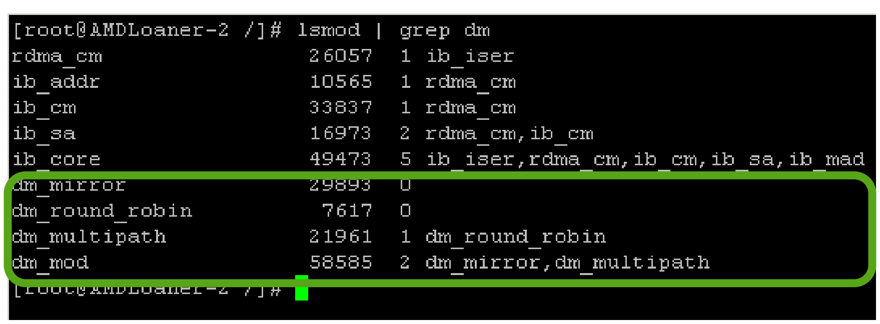

2. Install the NetApp Linux Host Utilities kit from:
   http://now.netapp.com/NOW/download/software/sanhost_linux/5.0/
3. Create or modify the `/etc/multipath.conf` file.

   A sample `/etc/multipath.con'` is shown below:

```
defaults
{
       user_friendly_names yes
       max_fds max
       queue_without_daemon no
}


blacklist
{
       wwid <DevId>
       devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
       devnode "^hd[a-z]"
       devnode "^cciss!c[0-9]d[0-9]*[p[0-9]*]"
}
devices
{
       device
       {
       vendor "NETAPP"
       product "LUN"
    getuid_callout "/sbin/scsi_id -g -u -s /block/%n"
    prio_callout "/sbin/mpath_prio_ontap /dev/%n"
       features "1 queue_if_no_path"
       hardware_handler "0"
       path_grouping_policy group_by_prio
       failback immediate
       rr_weight uniform
       rr_min_io 128
       path_checker directio
       flush_on_last_del yes
       }
       }
```

   The `<DevID>` refers to WWID of any SCSI device (not from NetApp) that is installed on the OVM Server; for example, the local SCSI drive: `/dev/sda`.

   The multipath devices need to have the same device identifier and device path on each Oracle VM Server in the OVM server pool. So if the `user_friendly_names` parameter is set to `yes` in the `/etc/multipath.conf` file, NetApp recommends using the `multipaths` section within `/etc/multipath.conf` to specify aliases corresponding to SCSI ID of each multipath device. This step can be executed after getting the SCSI ID of the devices from step 3. This will make sure that all multipath devices have a consistent name across all the nodes of the Oracle VM server pool.

```
multipaths {
          multipath {
                 wwid <SCSI ID of the multipath device 1>
                 alias <user friendly name>
          }
          multipath {
                 wwid < SCSI ID of the multipath device 2>
                 alias <user friendly name>
          }
}
```

4. Start the multipath service:

```
/etc/init.d/multipathd start
```

5. Configure the multipath service:

```
multipath
```

6. Add the multipath service to the boot sequence:

```
chkconfig --add multipathd
chkconfig multipathd on
```

7. Verify the multipath configuration:

```
multipath -v3 -d –ll
/etc/init.d/multipathd status
```

8. As shown in Step 5 below, once the mapped shared storage is scanned on the OVM Server, the multipaths can be viewed using either the `multipath –ll` or the `sanlun` command.

Step 5: Scan the Oracle VM Server for new storage

1. Rescan the Oracle VM Server to detect the newly mapped LUN.

```
[root@AMDLoaner-2 /]# cd /sys/class/scsi_host/
[root@AMDLoaner-2 scsi_host]#
[root@AMDLoaner-2 scsi_host]# ls
host0  host10  host2  host4  host6  host8
host1  host11  host3  host5  host7  host9
[root@AMDLoaner-2 scsi_host]#
[root@AMDLoaner-2 scsi_host]# for i in `seq 0 11` ; do echo "- - -" > /sys/class
/scsi_host/host${i}/scan; done
[root@AMDLoaner-2 scsi_host]#
[root@AMDLoaner-2 scsi_host]#
```

2. View the newly mapped LUN and the corresponding multipath device.

   The **sanlun** utility that comes with the NetApp Linux Host Utility Kit can display the information in a very user-friendly manner. Similar information is also displayed by the `multipath –ll` command.

```
sanlun lun show
```

```
[root@AMDL[illegible]aner-2 /]# sanlun lun show
controller          lun-pathname          device filena[illegible]e  adapter  protocol          lun size           lun state
  fas3140-[illegible]:  /vol/ovm_block_vol/lun_fcp  /dev/sdb       host6    FCP          100g (107374182400)   GOOD
  fas3140-[illegible]:  /vol/ovm_block_vol/lun_fcp  /dev/sdd       host7    FCP          100g (107374182400)   GOOD
  fas3140-[illegible]:  /vol/ovm_block_vol/lun_fcp  /dev/sdg       host8    FCP          100g (107374182400)   GOOD
  fas3140-[illegible]:  /vol/ovm_block_vol/lun_fcp  /dev/sdh       host9    FCP          100g (107374182400)   GOOD
  fas3140-[illegible]:  /vol/ovm_block_vol/lun_fcp  /dev/sdi       host9    FCP          100g (107374182400)   GOOD
  fas3140-[illegible]:  /vol/ovm_block_vol/lun_fcp  /dev/sdc       host6    FCP          100g (107374182400)   GOOD
  fas3140-7[illegible]  /vol/ovm_block_vol/lun_fcp  /dev/sdf       host8    FCP          100g (107374182400)   GOOD
  fas3140-7:  /vol/ovm_block_vol/lun_fcp  /dev/sde       host7    FCP          100g (107374182400)   GOOD
[root@AMDLoaner-2 /]#
```

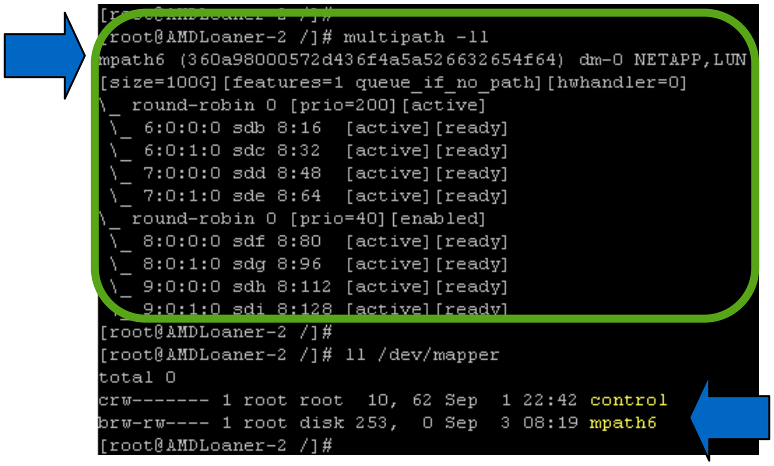

```
[root@AMDLoaner-2 /]#
[root@AMDLoaner-2 /]# multipath -ll
mpath6 (360a98000572d436f4a5a526632654f64) dm-0 NETAPP,LUN
[size=100G][features=1 queue_if_no_path][hwhandler=0]
\_ round-robin 0 [prio=200][active]
 \_ 6:0:0:0 sdb 8:16  [active][ready]
 \_ 6:0:1:0 sdc 8:32  [active][ready]
 \_ 7:0:0:0 sdd 8:48  [active][ready]
 \_ 7:0:1:0 sde 8:64  [active][ready]
\_ round-robin 0 [prio=40][enabled]
 \_ 8:0:0:0 sdf 8:80  [active][ready]
 \_ 8:0:1:0 sdg 8:96  [active][ready]
 \_ 9:0:0:0 sdh 8:112 [active][ready]
 \_ 9:0:1:0 sdi 8:128 [active][ready]
[root@AMDLoaner-2 /]#
[root@AMDLoaner-2 /]# ll /dev/mapper
total 0
crw------- 1 root root  10, 62 Sep  1 22:42 control
brw-rw---- 1 root disk 253,  0 Sep  3 08:19 mpath6
[root@AMDLoaner-2 /]#
```

3. The **sanlun** utility can even display detailed information regarding multipathing, including the physical port-specific information that corresponds with each of the multipaths.

```
sanlun lun show -p
```

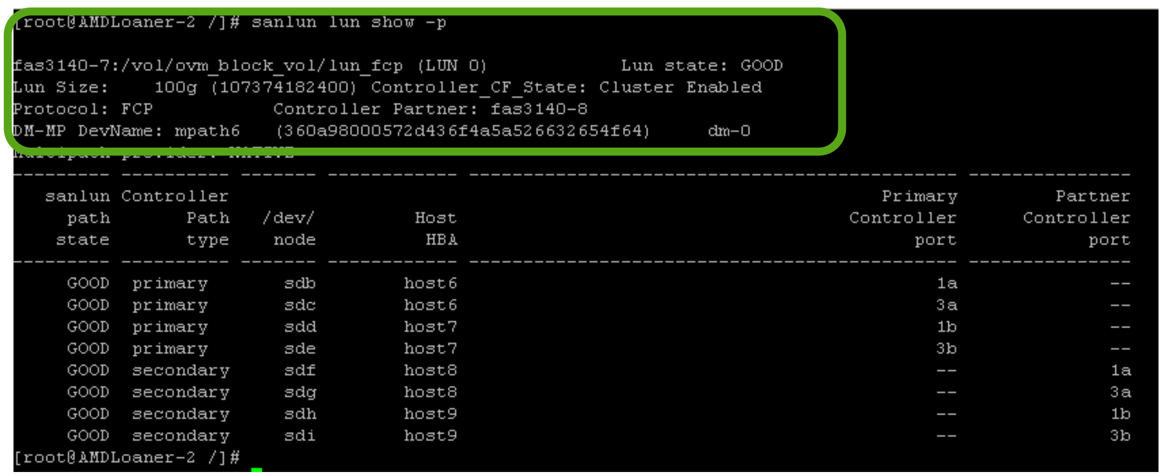

```
[root@AMDLoaner-2 /]# sanlun lun show -p

fas3140-7:/vol/ovm_block_vol/lun_fcp (LUN 0)          Lun state: GOOD
Lun Size:    100g (107374182400) Controller_CF_State: Cluster Enabled
Protocol: FCP            Controller Partner: fas3140-8
DM-MP DevName: mpath6   (360a98000572d436f4a5a526632654f64)      dm-0
Multipath provider: NATIVE
--------- ---------- ------- ------------ ------------------------------------------ ---------------
   sanlun Controller                                                           Primary         Partner
     path       Path   /dev/         Host                                   Controller      Controller
    state       type    node          HBA                                         port            port
--------- ---------- ------- ------------ ------------------------------------------ ---------------
     GOOD  primary       sdb        host6                                           1a              --
     GOOD  primary       sdc        host6                                           3a              --
     GOOD  primary       sdd        host7                                           1b              --
     GOOD  primary       sde        host7                                           3b              --
     GOOD  secondary     sdf        host8                                           --              1a
     GOOD  secondary     sdg        host8                                           --              3a
     GOOD  secondary     sdh        host9                                           --              1b
     GOOD  secondary     sdi        host9                                           --              3b
[root@AMDLoaner-2 /]#
```

4. The multipath device can be used either as an Oracle VM shared storage repository (OCFS2) or as a standalone storage repository (ext3).
5. While being using as a shared storage repository, OCFS2 first needs to be configured for the Oracle VM Server nodes in the server pool. A sample `/etc/ocfs2/cluster.conf` file that needs to present in each node of the server pool for the OCFS2 configuration may look like this:

```
node:
        ip_port = 7777
        ip_address = 10.73.69.112
        number = 0
        name = ovmserver1
        cluster = ocfs2

node:
        ip_port = 7777
        ip_address = 10.73.69.201
        number = 1
        name = ovmserver2
        cluster = ocfs2

cluster:
        node_count = 2
        name = ocfs2
```

To configure and start the OCFS2 cluster service:

```
service o2cb status
service o2cb load
service o2cb online
service o2cb start
service o2cb configure
```

NetApp recommends the following values:

```
O2CB_HEARTBEAT_THRESHOLD=81, O2CB_IDLE_TIMEOUT_MS=160000, O2CB_KEEPALIVE_DELAY_MS=4000,
O2CB_RECONNECT_DELAY_MS=4000
```

Now the multipath device can be formatted and mounted as an Oracle VM Server respository:

```
mkfs –t ocfs2 /dev/mapper/<mapthdevice>
```

For OVM v2.2:

```
/opt/ovs-agent-2.3/utils/repos.py –n /dev/mapper/<mapthdevice>
```

For OVM v2.1.5 and earlier:

```
/usr/lib/ovs/ovs-makerepo /dev/mapper/<mapthdevice> 1 <description>
```

6. Oracle VM 2.2 also supports the creation of a shared storage repository of a raw multipath device (or partition) without using OCFS2.

```
/opt/ovs-agent-2.3/utils/repos.py –n /dev/mapper/<mapthdevice>
```

7. Otherwise, the multipath device can be formatted and used in one of the following ways:

- As a standalone Oracle VM Server repository:

```
mkfs –t ext3 /dev/mapper/<mapthdevice>
```

  – (OVM v2.2)

```
/opt/ovs-agent-2.3/utils/repos.py –n /dev/mapper/<mapthdevice>
```

  – (OVM 2.1.5 or earlier)

```
/usr/lib/ovs/ovs-makerepo /dev/mapper/<mapthdevice> 1 <description>
```

- As a device attached to the virtual machines through the virtual machine config file (vm.cfg)

```
mkfs –t ext3 /dev/mapper/<mpathdevice>
```

Or

```
mkfs –t ocfs2 /dev/mapper/<mpathdevice>
```

Sample entries in the `vm.cfg` file:

```
disk = [ 'file://OVS/running_pool/OVM_EL4U8_X86_64_PVM_4GB//System.img,xvda,w' ,
'phy:/dev/mapper/<mpathdevice>,xvdc,w' ]
```

CONFIGURING ISCSI SHARED STORAGE ON ORACLE VM SERVER

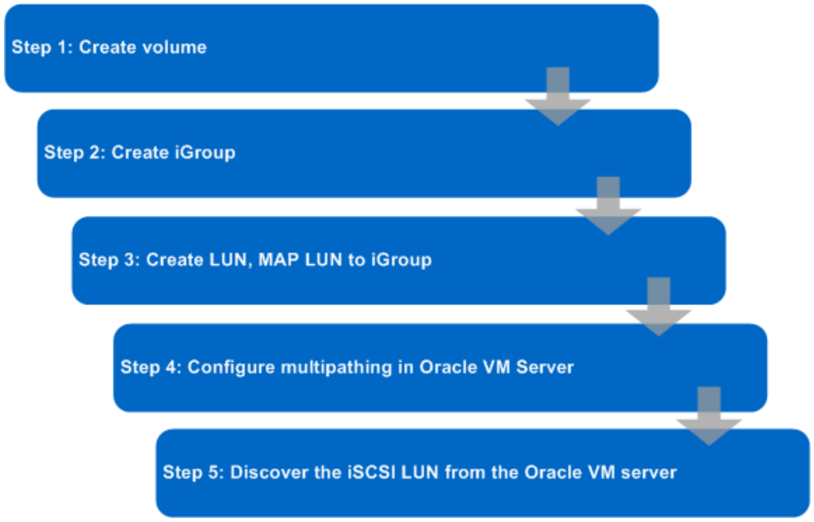

Step 1: Create Volume

Same as Step 1 in section Configuring FC Shared Storage on Oracle VM Server.

Step 2: Create iGroup

1. Check the status of the iSCSI service on the Oracle VM Server. If it is not running or is not installed, install it and start the iSCSI service.

   Make a note of the iSCSI initiator ID (IQN number).

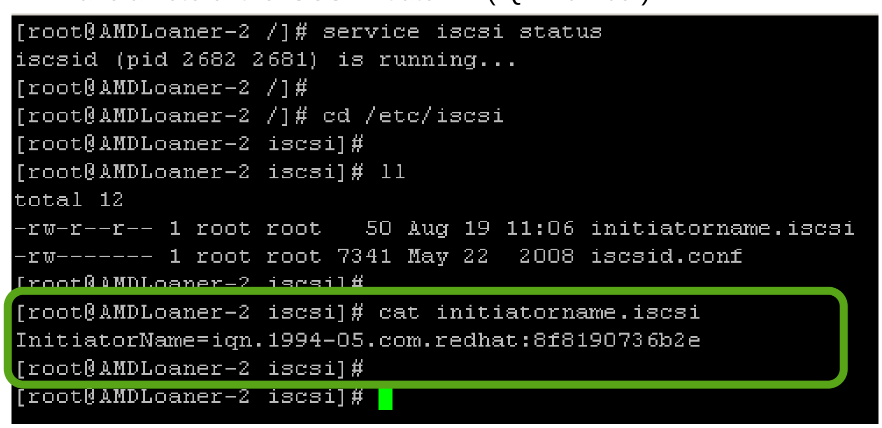

```
[root@AMDLoaner-2 /]# service iscsi status
iscsid (pid 2682 2681) is running...
[root@AMDLoaner-2 /]#
[root@AMDLoaner-2 /]# cd /etc/iscsi
[root@AMDLoaner-2 iscsi]#
[root@AMDLoaner-2 iscsi]# ll
total 12
-rw-r--r-- 1 root root   50 Aug 19 11:06 initiatorname.iscsi
-rw------- 1 root root 7341 May 22  2008 iscsid.conf
[root@AMDLoaner-2 iscsi]#
[root@AMDLoaner-2 iscsi]# cat initiatorname.iscsi
InitiatorName=iqn.1994-05.com.redhat:8f8190736b2e
[root@AMDLoaner-2 iscsi]#
[root@AMDLoaner-2 iscsi]#
```

   If the Oracle VM Server is in a server pool and the iSCSI storage needs to be configured as a shared Oracle VM repository, collect the IQN number of the nodes of the cluster so that they can be put inside the same iGroup.

2. Create an iGroup on the storage using NetApp System Manager and assign the IQN number(s) noted above to the iGroup.

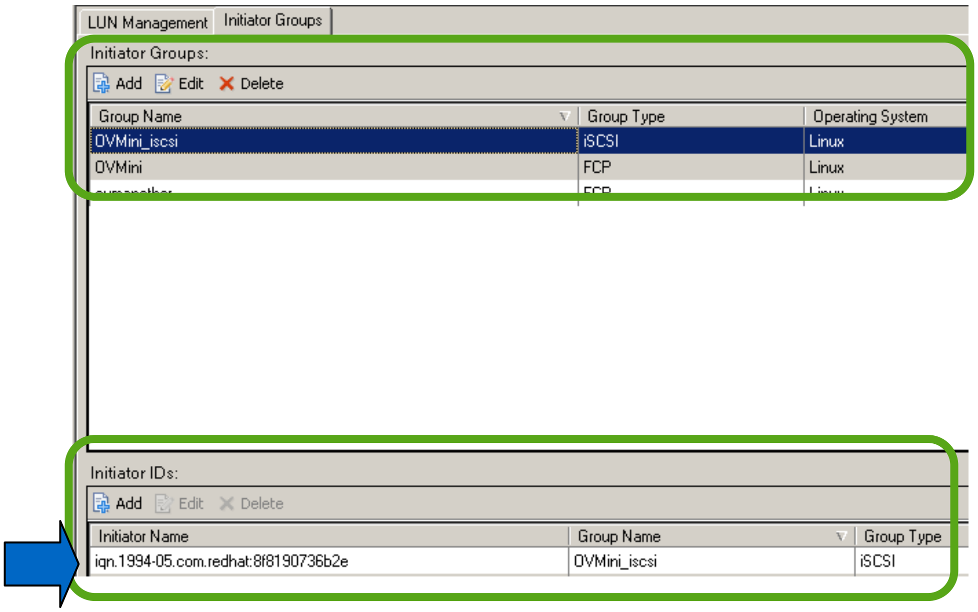

Step 3: Create LUN, map iGroup

1. From the LUN Wizard of NetApp System Manager, create the LUN inside the column created in [Step 1](#) and then map that LUN to the iGroup created in [Step 2](#).

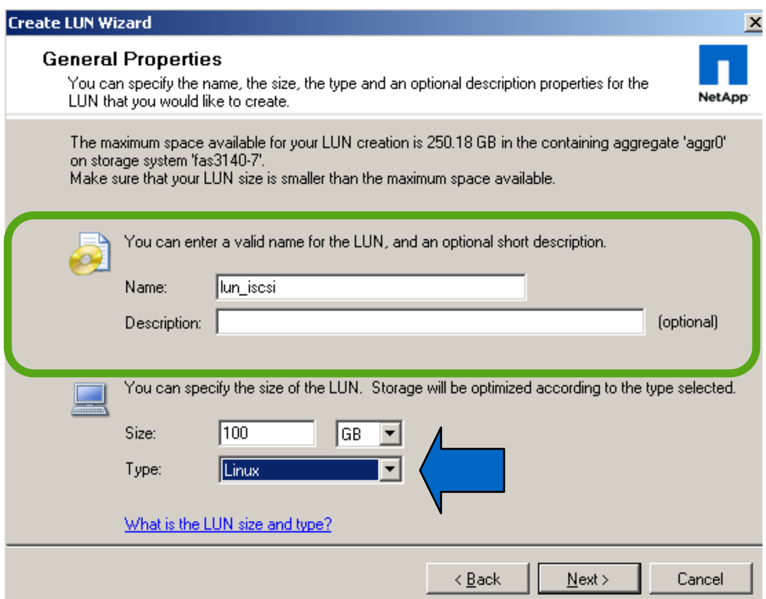

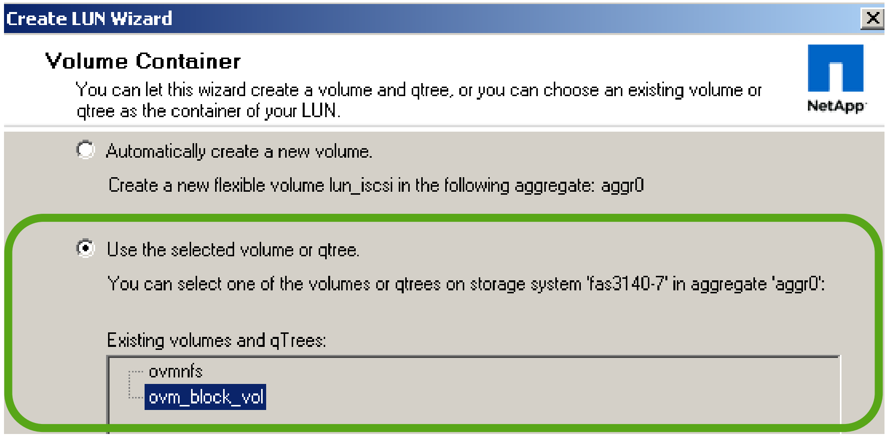

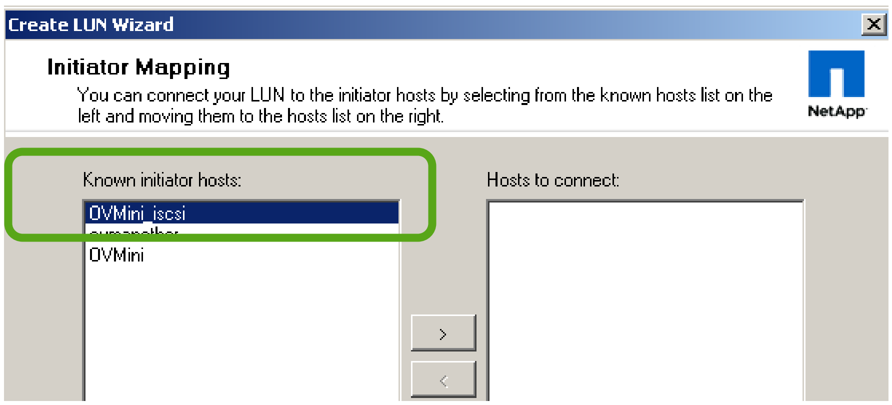

2. Complete the LUN Wizard to finish the LUN creation process.

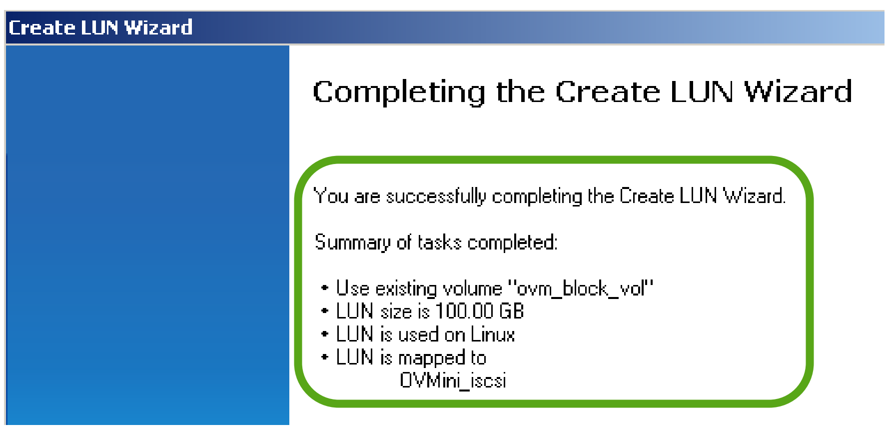

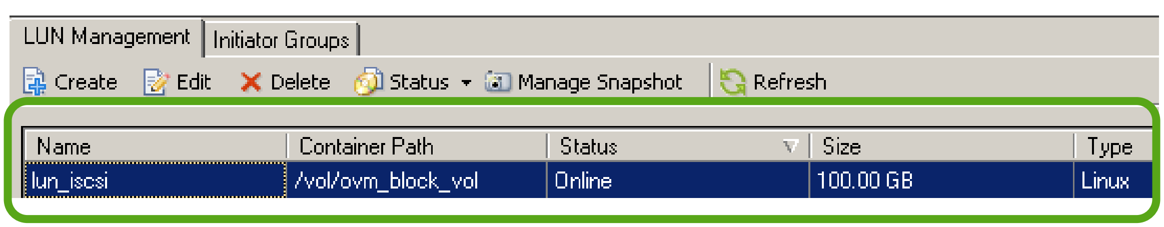

3. If using thin provisioning, clear the Space Reserved checkbox in the LUN properties.

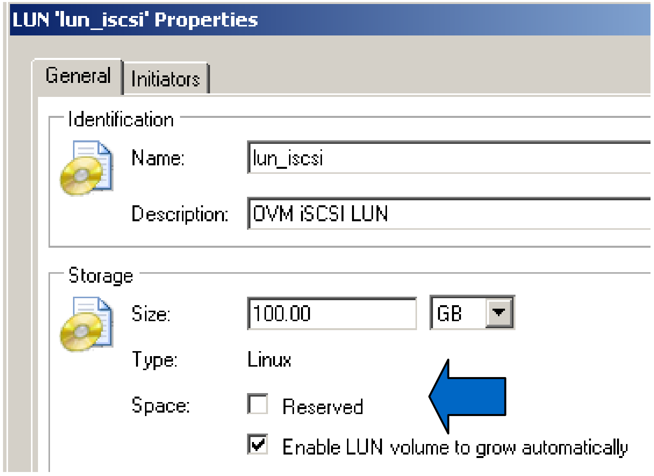

Step 4: Configure multipathing on the Oracle VM Server

Refer to the section on multipathing in the FC shared storage section.

Step 5: Discover the iSCSI LUN from the Oracle VM Server

1. Discover the iSCSI LUN from the Oracle VM Server.

```
[root@AMDLoaner-2 /]# iscsiadm -m discovery -t sendtargets -p 10.61.166.223
10.61.166.223:3260,1001 iqn.1992-08.com.netapp:sn.151697881
[root@AMDLoaner-2 /]#
[root@AMDLoaner-2 /]# iscsiadm -m node
10.61.166.223:3260,1001 iqn.1992-08.com.netapp:sn.151697881
[root@AMDLoaner-2 /]#
```

```
[root@AMDLoaner-2 /]# service iscsi restart
Stopping iSCSI daemon: /etc/init.d/iscsi: line 33: 16942 Killed
 /etc/init.d/iscsid stop
iscsid dead but pid file exists                              [  OK  ]
Turning off network shutdown. Starting iSCSI daemon:         [  OK  ]
                                                             [  OK  ]
Setting up iSCSI targets: Logging in to [iface: default, target: iqn.1992-08.com
.netapp:sn.151697881, portal: 10.61.166.223,3260]
Login to [iface: default, target: iqn.1992-08.com.netapp:sn.151697881, portal: 1
0.61.166.223,3260]: successful
                                                             [  OK  ]
[root@AMDLoaner-2 /]#
```

2. View the newly mapped LUN. The sanlun utility provided by NetApp Host Utilities displays the LUN information in a user-friendly manner.

```
fas3140-7:/vol/ovm_block_vol/lun_iscsi (LUN 0)          Lun state: GOOD
Lun Size:    100g (107374182400) Controller_CF_State: Cluster Enabled
Protocol: iSCSI          Controller Partner: fas3140-8
DM-MP DevName: mpath7    (360a98000572d436f4a5a5266354d4549)     dm-1
Multipath-provider: NATIVE
--------- ---------- ------- ------------ ---------------------------------------------- ---------------
   sanlun Controller                                                          Primary         Partner
     path       Path   /dev/         Host                                  Controller      Controller
    state       type    node          HBA                                        port            port
--------- ---------- ------- ------------ ---------------------------------------------- ---------------
     GOOD  iscsi         sdj       host12                                 10.61.166.223              --
[root@AMDLoaner-2 /]#
```

```
[root@AMDLoaner-2 /]# sanlun lun show
controller:           lun-pathname            device filename  adapter  protocol             lun size            lun state
```

3. Once the iSCSI LUN is visible to the Oracle VM host, it can be configured either as a shared storage repository (OCFS2), a standalone storage (ext3) repository, or a virtual disk attached to virtual machines exactly in the same way mentioned for the FC LUN in the previous section.

4.2 NETAPP DEDUPLICATION IN AN ORACLE VM ENVIRONMENT

1. Five identical VMs reside inside the `/OVS/running_pool` directory.

```
[root@AMDLoaner-2 running_pool]# pwd
/OVS/running_pool
[root@AMDLoaner-2 running_pool]#
[root@AMDLoaner-2 running_pool]# ll
total 20
drwxr-xr-x 2 root root 4096 Jan 23  2009 OVM_EL5U2_X86_64_PVM_4GB_VM_1
drwxr-xr-x 2 root root 4096 Jan 23  2009 OVM_EL5U2_X86_64_PVM_4GB_VM_2
drwxr-xr-x 2 root root 4096 Jan 23  2009 OVM_EL5U2_X86_64_PVM_4GB_VM_3
drwxr-xr-x 2 root root 4096 Jan 23  2009 OVM_EL5U2_X86_64_PVM_4GB_VM_4
drwxr-xr-x 2 root root 4096 Jan 23  2009 OVM_EL5U2_X86_64_PVM_4GB_VM_5
[root@AMDLoaner-2 running_pool]#
```

2. The /OVS cluster root has been created on a NetApp shared storage volume. Check the space consumed from the OVM Server.

```
10.61.166.224:/vol/OVM_NFS
                       40G    15G   26G  36% /OVS
```

3. The space consumed can also be observed from the NetApp System Manager or FilerView or NetApp FAS system command line.

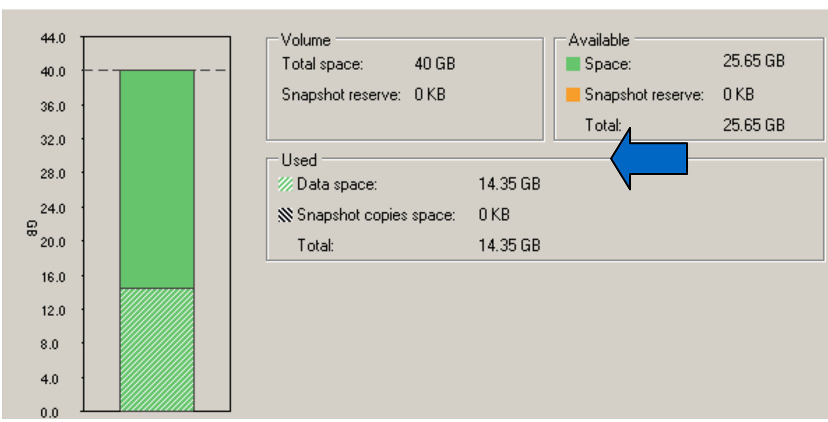

4. Enable deduplication on the volume from the NetApp FAS CLI using the sis on command or from NetApp System Manager.

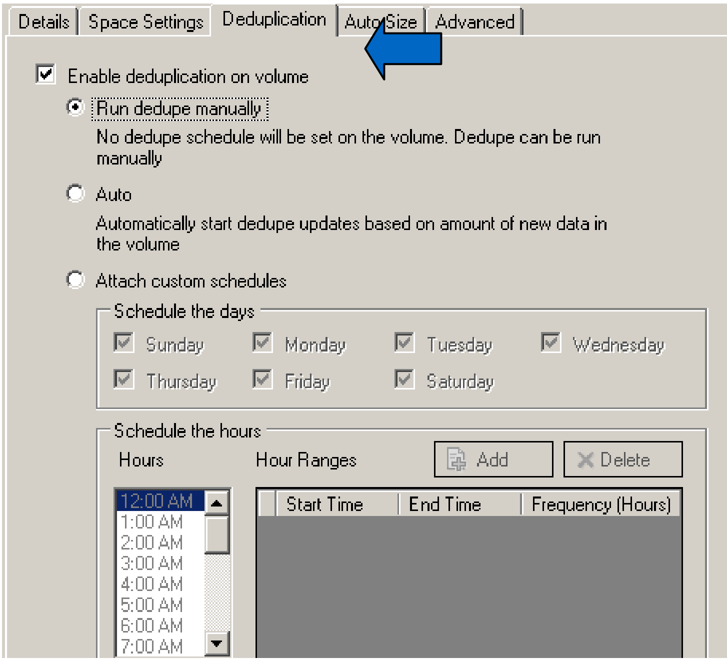

5. Once enabled on the volume, the deduplication process can be started manually or automatically or can be scheduled either from the NetApp FAS CLI using the `sis start` command or from NetApp System Manager.

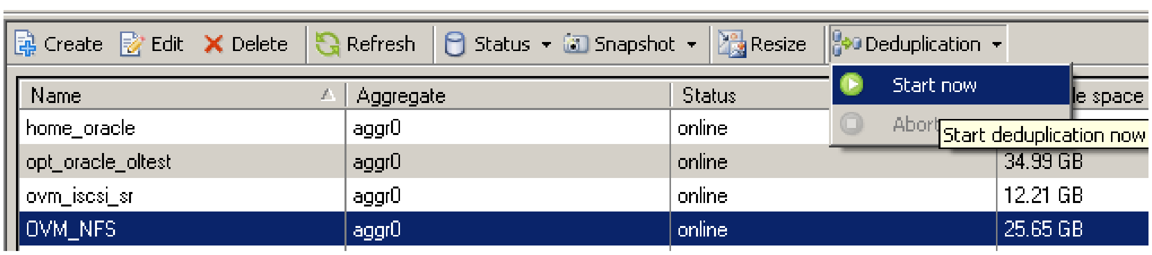

6. After the deduplication process is over, check the space consumed by the volume and the storage savings realized.

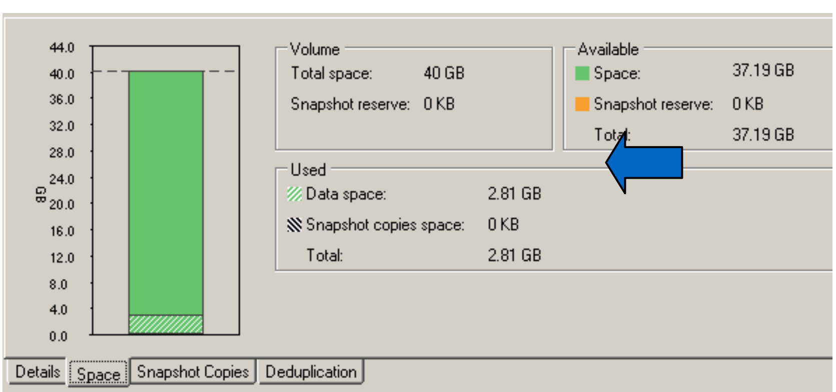

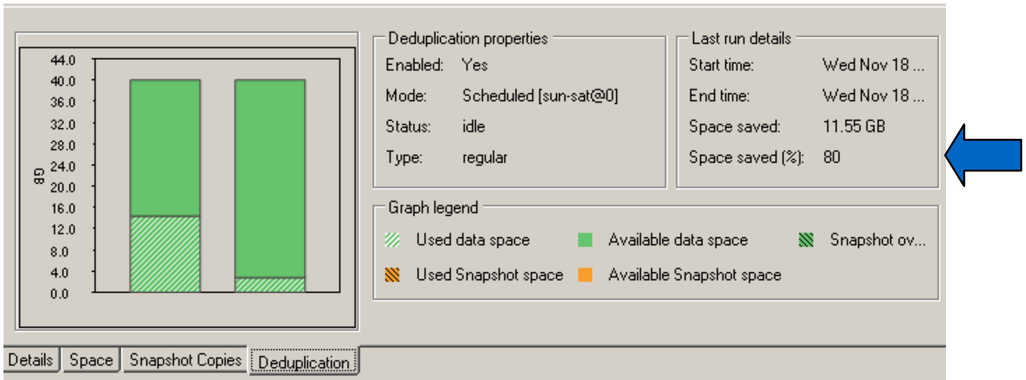

4.3 NETAPP VIRTUAL CLONING (FLEXCLONE & SIS CLONE) IN ORACLE VM ENVIRONMENT

The following procedure describes NetApp thin cloning (file and volume level), which corresponds to the design illustrated in section 3.5.

1. The source VM resides inside an already existing OVM repository.

```
[root@ovm22nb /]# /opt/ovs-agent-2.3/utils/repos.py -l
[   ] 8480ed5e-0f7c-451e-bbf4-c9b205f9f98a => 10.73.68.199:/vol/ovm_nfs_sr_1
```

```
[root@ovm22nb running_pool]# pwd
/var/ovs/mount/8480ED5E0F7C451EBBF4C9B205F9F98A/running_pool
[root@ovm22nb running_pool]#
[root@ovm22nb running_pool]# ll
total 711800
drwxr-xr-x 2 root root       4096 Nov  3  2009 OVM_EL5U3_X86_64_PVM_4GB
```

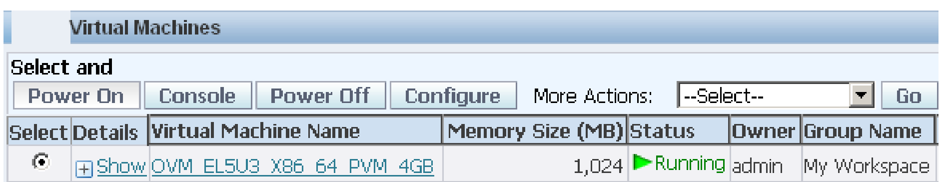

2. Create another OVM repository from a NFS volume on the NetApp FAS system on which the clones of the source VM need to be created.

```
[root@ovm22nb running_pool]# /opt/ovs-agent-2.3/utils/repos.py -l
[   ] 8480ed5e-0f7c-451e-bbf4-c9b205f9f98a => 10.73.68.199:/vol/ovm_nfs_sr_1
[   ] 718bf39f-af48-4922-97fa-f50861634995 => 10.73.68.199:/vol/ovm_nfs_sr_dev_test
```

3. In the example below, five clones of the source VM need to be created.

   First create the directories corresponding to these clone VMs inside the `running_pool` directory of the storage repository. Also create a directory corresponding to the golden VM from which these VMs will be cloned.

```
[root@ovm22nb 718BF39FAF48492297FAF50861634995]# cd running_pool/
[root@ovm22nb running_pool]# ll
total 0
[root@ovm22nb running_pool]# mkdir OVM_EL5U3_X86_64_PVM_4GB_golden
[root@ovm22nb running_pool]# mkdir OVM_EL5U3_X86_64_PVM_4GB_clone1
[root@ovm22nb running_pool]# mkdir OVM_EL5U3_X86_64_PVM_4GB_clone2
[root@ovm22nb running_pool]# mkdir OVM_EL5U3_X86_64_PVM_4GB_clone3
[root@ovm22nb running_pool]# mkdir OVM_EL5U3_X86_64_PVM_4GB_clone4
[root@ovm22nb running_pool]# mkdir OVM_EL5U3_X86_64_PVM_4GB_clone5
[root@ovm22nb running_pool]#
```

4. From the NetApp FAS system console, use `ndmpcopy` to copy the images of the source VM to the directory created for the golden VM.

```
FAS> ndmpcopy /vol/<volume corresponding to OVM repository holding source VM>/running_pool/<VM directory>/   /vol/<volume corresponding to OVM repository holding golden VM directory>/running_pool/<Golden VM directory>
```

```
ndmpcopy /vol/ovm_nfs_sr_1/running_pool/OVM_EL5U3_X86_64_PVM_4GB/System.img
         /vol/ovm_nfs_sr_dev_test/running_pool/OVM_EL5_X86_PVM_4GB_golden/
```

5. Now create clones of the golden VM inside the cloned VM directories already created using the `clone` command on the NetApp FAS system console.

```
clone start /vol/ovm_nfs_sr_dev_test/running_pool/OVM_EL5U3_X86_64_PVM_4GB_golden/System.img
            /vol/ovm_nfs_sr_dev_test/running_pool/OVM_EL5U3_X86_64_PVM_4GB_clone1/System.img
```

```
clone start /vol/ovm_nfs_sr_dev_test/running_pool/OVM_EL5U3_X86_64_PVM_4GB_golden/System.img
            /vol/ovm_nfs_sr_dev_test/running_pool/OVM_EL5U3_X86_64_PVM_4GB_clone2/System.img
```

```
clone start /vol/ovm_nfs_sr_dev_test/running_pool/OVM_EL5U3_X86_64_PVM_4GB_golden/System.img
            /vol/ovm_nfs_sr_dev_test/running_pool/OVM_EL5U3_X86_64_PVM_4GB_clone3/System.img
```

```
clone start /vol/ovm_nfs_sr_dev_test/running_pool/OVM_EL5U3_X86_64_PVM_4GB_golden/System.img
            /vol/ovm_nfs_sr_dev_test/running_pool/OVM_EL5U3_X86_64_PVM_4GB_clone4/System.img
```

```
clone start /vol/ovm_nfs_sr_dev_test/running_pool/OVM_EL5U3_X86_64_PVM_4GB_golden/System.img
            /vol/ovm_nfs_sr_dev_test/running_pool/OVM_EL5U3_X86_64_PVM_4GB_clone5/System.img
```

6. Create a virtual machine config file (`vm.cfg`) inside each of the cloned VM directories. A typical `vm.cfg` file may look like this:

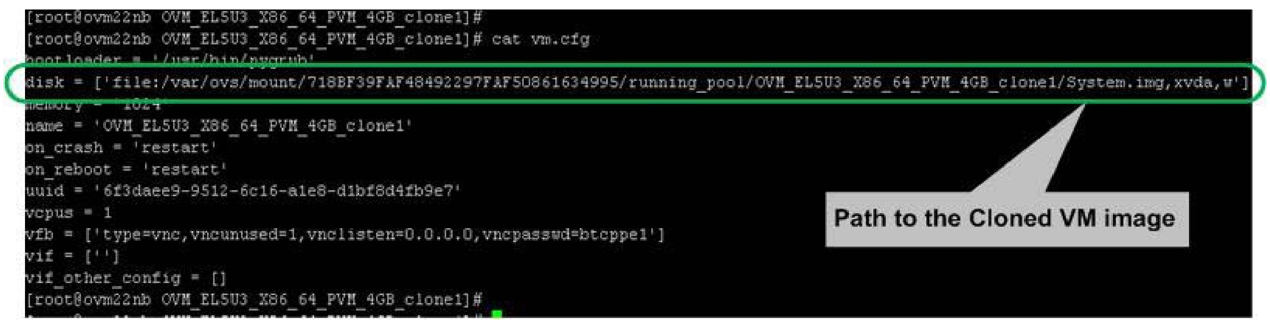

7. Now the cloned VMs can be discovered from Oracle VM Manager and can be powered on.

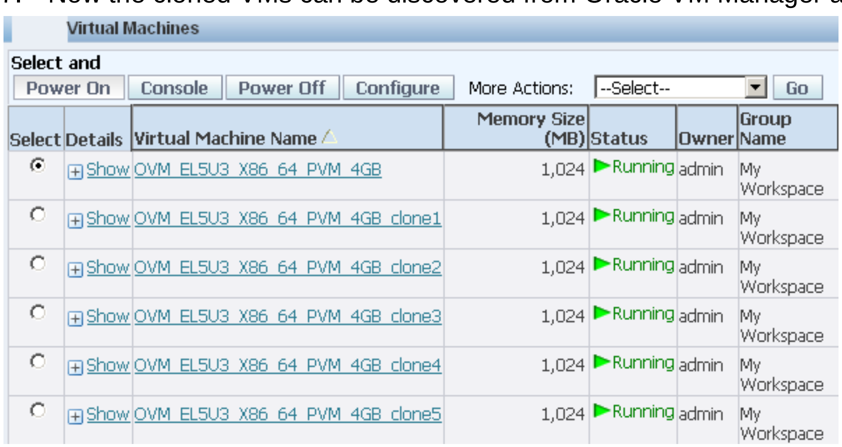

Select	Details	Virtual Machine Name	Memory Size (MB)	Status	Owner	Group Name
◉	Show	OVM_EL5U3_X86_64_PVM_4GB	1,024	Running	admin	My Workspace
○	Show	OVM_EL5U3_X86_64_PVM_4GB_clone1	1,024	Running	admin	My Workspace
○	Show	OVM_EL5U3_X86_64_PVM_4GB_clone2	1,024	Running	admin	My Workspace
○	Show	OVM_EL5U3_X86_64_PVM_4GB_clone3	1,024	Running	admin	My Workspace
○	Show	OVM_EL5U3_X86_64_PVM_4GB_clone4	1,024	Running	admin	My Workspace
○	Show	OVM_EL5U3_X86_64_PVM_4GB_clone5	1,024	Running	admin	My Workspace

8. The NFS volume on the NetApp FAS system holding all the cloned VMs will consume only the space for the golden VM; no space will be consumed for the cloned VMs till the users make changes inside the cloned VMs.
9. Create a volume-level FlexClone volume from the NFS volume holding the cloned VMs either from the FAS system console or from NetApp System Manager.

```
vol clone create ovm_nfs_sr_dev_test_flex -s none -b ovm_nfs_sr_dev_test
```

10. Mount the just-created FlexClone volume on the Oracle VM Server on a temporary directory and delete the already existing `.ovsrepo` file inside the FlexClone volume.

```
[root@ovm22nb /]#
[root@ovm22nb /]# mount 10.73.68.199:/vol/ovm_nfs_sr_dev_test_flex /mnt
[root@ovm22nb /]#
[root@ovm22nb /]# cd /mnt
[root@ovm22nb mnt]# ls -la
total 36
drwxrwxrwx  8 root root 4096 Nov  3  2009 .
drwxr-xr-x 31 root root 4096 Nov  1 18:29 ..
drwxrwxrwx  2 root root 4096 Nov  3  2009 iso_pool
-rw-rw-rw-  1 root root   84 Nov  3  2009 .ovsrepo
drwxrwxrwx  2 root root 4096 Nov  3  2009 publish_pool
drwxrwxrwx  8 root root 4096 Nov  3  2009 running_pool
drwxrwxrwx  2 root root 4096 Nov  3  2009 seed_pool
drwxrwxrwx  2 root root 4096 Nov  3  2009 sharedDisk
drwxrwxrwx  3 root root 4096 Nov  4  2009 .snapshot
[root@ovm22nb mnt]#
[root@ovm22nb mnt]# cat .ovsrepo
OVS_REPO_UUID=718BF39FAF48492297FAF50861634995
OVS_REPO_SHARED=1
OVS_REPO_VERSION=1
[root@ovm22nb mnt]#
[root@ovm22nb mnt]# rm .ovsrepo
rm: remove regular file `.ovsrepo'? y
```

11. Unmount the FlexClone volume and then create a new Oracle VM repository with it. The new repository will have all the cloned VMs and the golden VM created earlier.

```
[root@ovm22nb /]#
[root@ovm22nb /]# /opt/ovs-agent-2.3/utils/repos.py -l
[   ] 8480ed5e-0f7c-451e-bbf4-c9b205f9f98a => 10.73.68.199:/vol/ovm_nfs_sr_1
[   ] 54a687cc-cc8f-4e9a-8da4-314567032b65 => 10.73.68.199:/vol/ovm_nfs_sr_dev_test_flex
[   ] 718bf39f-af48-4922-97fa-f50861634995 => 10.73.68.199:/vol/ovm_nfs_sr_dev_test
```

```
[root@ovm22nb running_pool]# pwd
/var/ovs/mount/B7A708F7FA574D099829DB16EB7A4111/running_pool
[root@ovm22nb running_pool]#
[root@ovm22nb running_pool]# ll
total 24
drwxr-xr-x 2 root root 4096 Nov  3  2009 OVM_EL5U3_X86_64_PVM_4GB_clone1
drwxr-xr-x 2 root root 4096 Nov  3  2009 OVM_EL5U3_X86_64_PVM_4GB_clone2
drwxr-xr-x 2 root root 4096 Nov  3  2009 OVM_EL5U3_X86_64_PVM_4GB_clone3
drwxr-xr-x 2 root root 4096 Nov  3  2009 OVM_EL5U3_X86_64_PVM_4GB_clone4
drwxr-xr-x 2 root root 4096 Nov  3  2009 OVM_EL5U3_X86_64_PVM_4GB_clone5
drwxr-xr-x 2 root root 4096 Nov  3 12:08 OVM_EL5U3_X86_64_PVM_4GB_golden
```

12. Rename the VM directories with unique VM names and change the `vm.cfg` files inside the VM directories to reflect the new names and the correct disk path.
    a. Change the name of the VM directory name inside `/var/ovs/mount/<repository>/running_pool`
    b. Modify the `vm.cfg` file inside `running_pool` as follows:
        i. Change the VM name.

ii. Modify the disk path.

iii. Delete the VIF information, if any.

```
[root@ovm22nb running_pool]# pwd
/var/ovs/mount/B7A708F7FA574D099829DB16EB7A4111/running_pool
[root@ovm22nb running_pool]#
[root@ovm22nb running_pool]# ll
total 24
drwxr-xr-x 2 root root 4096 Nov  4  2009 OVM_EL5U3_X86_64_PVM_4GB_clone10
drwxr-xr-x 2 root root 4096 Nov  4  2009 OVM_EL5U3_X86_64_PVM_4GB_clone6
drwxr-xr-x 2 root root 4096 Nov  4  2009 OVM_EL5U3_X86_64_PVM_4GB_clone7
drwxr-xr-x 2 root root 4096 Nov  4  2009 OVM_EL5U3_X86_64_PVM_4GB_clone8
drwxr-xr-x 2 root root 4096 Nov  4  2009 OVM_EL5U3_X86_64_PVM_4GB_clone9
drwxr-xr-x 2 root root 4096 Nov  3 12:08 OVM_EL5U3_X86_64_PVM_4GB_golden
```

The VMs can now be discovered from Oracle VM Manager and can be powered on.

Virtual Machines

Select and

Power On | Console | Power Off | Configure | More Actions: --Select-- | Go

Select	Details	Virtual Machine Name	Memory Size (MB)	Status	Owner	Group Name
⦿	Show	OVM_EL5U3_X86_64_PVM_4GB	1,024	Running	admin	My Workspace
○	Show	OVM_EL5U3_X86_64_PVM_4GB_clone1	1,024	Running	admin	My Workspace
○	Show	OVM_EL5U3_X86_64_PVM_4GB_clone10	1,024	Powered Off	admin	My Workspace
○	Show	OVM_EL5U3_X86_64_PVM_4GB_clone2	1,024	Running	admin	My Workspace
○	Show	OVM_EL5U3_X86_64_PVM_4GB_clone3	1,024	Running	admin	My Workspace
○	Show	OVM_EL5U3_X86_64_PVM_4GB_clone4	1,024	Running	admin	My Workspace
○	Show	OVM_EL5U3_X86_64_PVM_4GB_clone5	1,024	Running	admin	My Workspace
○	Show	OVM_EL5U3_X86_64_PVM_4GB_clone6	1,024	Powered Off	admin	My Workspace
○	Show	OVM_EL5U3_X86_64_PVM_4GB_clone7	1,024	Powered Off	admin	My Workspace
○	Show	OVM_EL5U3_X86_64_PVM_4GB_clone8	1,024	Running	admin	My Workspace
○	Show	OVM_EL5U3_X86_64_PVM_4GB_clone9	1,024	Powered Off	admin	My Workspace

Note that there are now 10 cloned VMs running in the Oracle VM Server pool using storage for just 1 VM—the golden VM created in Step 4.

The process described above is for illustration purposes only and, depending on the requirement of the user, clones of the VM can be created in many other ways, combining the NetApp technologies volume-, LUN-, and file-level cloning and deduplication. See section 3.5 for details.

5 CASE STUDY—ORACLE VM IN A DATA CENTER ENVIRONMENT WITH NETAPP STORAGE

Oracle on Demand, Oracle's software- as-a-service business, introduced server virtualization within Oracle on Demand Managed Services Grid using Oracle VM. In addition to the guaranteed consolidation benefits, the use of Oracle VM in the Oracle on Demand data center also provides additional benefits like capacity on demand, rapid provisioning, and high availability of virtual machines. NetApp storage solutions play a major role in achieving these benefits.

Use of NetApp storage with Oracle VM in the Oracle on Demand environment delivers significant benefits, including:

- **Capacity on demand**. The hardware resources allocated to a particular application can be scaled without disruption.
- **Rapid provisioning**. Downtime resulting from hardware failures is nearly eliminated. Virtual machines can be restarted rapidly in the event of hardware failure.
- Decreased impact of necessary hardware maintenance: When maintenance is necessary, virtual machines can simply be migrated to another physical server, eliminating downtime.

NetApp storage when used with Oracle VM also helps making the transition to the virtualized environment from the physical environment seamless and it does not require any major process changes. That is, the strategies and processes used for backup and restore, disaster recovery, software upgrade, and so on that are used in the data center do not go through a complete makeover, because the underlying infrastructure has been transformed to a virtual environment.

Some of the innovative technologies from NetApp that are exploited by Oracle in deploying Oracle VM in the Oracle on Demand data center are:

- **FlexClone**

  NetApp FlexClone technology is used to rapidly provision virtual machines and storage for test and development environments in Oracle on Demand. NetApp FlexClone technology creates true clones—instantly replicated data volumes and data sets, without requiring additional storage space. For information on using FlexClone with Oracle VM, see section NetApp Virtual Cloning.
- **Over provisioning of storage through thin provisioning**

  For nonproduction environments, Oracle on Demand uses NetApp thin provisioning technology to overprovision storage to the virtual machines. For details on thin provisioning technology in the context of Oracle VM, see section NetApp Thin Provisioning.
- **Deduplication**

  NetApp deduplication is being used successfully in the production environment. For more details about NetApp deduplication in an Oracle VM environment, see section Deduplication.
- **Online backup/restore**

  The Oracle on Demand online backup and restore strategy of Oracle VM virtual machine images is based on NetApp Snapshot technology. For more details about NetApp Snapshot technology in an Oracle VM environment, see section 3.6.

For more details about the Oracle VM implementation in Oracle on Demand, see the following white paper from Oracle: http://www.oracle.com/ondemand/collateral/virtualization-oracle-vm-wp.pdf.

6 CONCLUSION

With the introduction of Oracle VM, Oracle is now the only software vendor to combine the benefits of server clustering and server virtualization technologies, delivering integrated clustering, virtualization, storage, and management for grid computing. This report provides detailed guidance on how best to implement Oracle VM Server virtualization solutions on NetApp storage. NetApp has been at the forefront of solving complex business problems with its innovative technology and end-to-end-solutions approach.

This report is not intended to be a definitive implementation or solutions guide. Additional expertise may be required to solve specific deployments. Contact your local NetApp sales representative to speak with one of our Oracle VM solutions experts.

Please forward any errors, omissions, differences, new discoveries, or comments about this paper to the authors.

REVISION HISTORY

Date	Name	Description
Dec 2009	Preetom Goswami Antonio Jose Rodrigues Neto Padmanabhan Sadagopan	Created and Updated for Oracle VM 2.2
April 2011	Padmanabhan Sadagopan	Updated Revision History and added note.
Nov 2011	Padmanabhan Sadagopan	Updated the `o2cb_idle_timeout_ms` value to 160000 and updated the note.

Note: This document will be updated shortly for Oracle VM 3.0.

ACKNOWLEDGEMENT

Steven Schuettinger—Technical Alliance Manager, NetApp

Lynne Thieme—Senior Manager, NetApp

Uday Shet—Senior Manager, NetApp

Mani Kannan—Reference Architect, NetApp

Ganesh Kamath—Rapid Response Engineer, NetApp

NetApp provides no representations or warranties regarding the accuracy, reliability or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®

www.netapp.com