Technical Report

Oracle VM and NetApp Storage Best Practices Guide

Karthikeyan Nagalingam, Reference Architect, NetApp
Michele Resta, Director Alliances Oracle Linux and Virtualization, Oracle
Doan Nguyen, Senior Principal Product Marketing Director, Oracle
Ronen Kofman, Director Product Management Oracle VM , Oracle
Greg King, Principal Best Practices Consultant, Oracle

March 2014 | TR-3712

Executive Summary

NetApp's innovative technologies enable organizations to extract benefits from their virtual infrastructures by seamlessly integrating advanced virtualized storage alongside 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) backups 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 and NetApp deployment case study.

TABLE OF CONTENTS

1 Introduction 4

1.1 Purpose and Scope of the Document 4

1.2 Intended Audience and Assumptions 4

2 Overview of the Oracle VM and NetApp Storage Solution 4

2.1 Advantages of Virtualization 4

2.2 Oracle VM Application-Driven Virtualization for the Enterprise 4

2.3 NetApp Storage Solutions for Application-Driven Virtualization 5

2.4 Oracle VM Validated Configuration with NetApp Storage 6

2.5 Shared Storage Options for Oracle VM 7

2.6 Plug-Ins 7

2.7 Choose Your Plug-In 8

2.8 Oracle VM Server Repositories 9

3 Oracle VM and NetApp Storage Best Practices 10

3.1 NetApp Shared Storage Best Practices for High Availability and Resiliency of Oracle VM Infrastructure 10

3.2 NetApp Storage Networking Best Practices 13

3.3 NetApp Storage Provisioning Best Practices for Oracle VM Server 18

3.4 Oracle VM Storage Best Practices—Create, Import, Live Migrate 21

3.5 Increasing Storage Utilization Using Deduplication, Thin Cloning, Thin Provisioning 23

3.6 Disk-Based Snapshot Backups for Oracle VMs Using NetApp Snapshot and SnapRestore 28

3.7 Disaster Recovery of an OVM Virtual Infrastructure Using NetApp SnapMirror 30

3.8 Designing a Highly Available OVM Virtual Infrastructure Using OVM HA and MetroCluster 30

4 Sample Implementations and Best Practices 31

4.1 Provisioning Storage for Clustered Data ONTAP and 7-Mode 31

4.2 Configuring NFS Storage on Oracle VM Server 34

4.3 Configuring FC Shared Storage and VM Operations on Oracle VM Server 39

4.4 Configuring iSCSI Shared Storage on Oracle VM Server 52

4.5 NetApp Deduplication in an Oracle VM Environment 57

4.6 NetApp Virtual Cloning (FlexClone and SIS Clone) in an Oracle VM Environment 62

5 Case Study—Oracle VM in a Data Center Environment with NetApp Storage 69

6 Appendix 70

6.1 NFS—Storage Repository Creation in Clustered Data ONTAP Using the CLI from OVM Manager 70

6.2 SAN—Storage Repository Creation in Clustered Data ONTAP Using the CLI from OVM Manager 71

7 Conclusion 73

8 Revision History 73

LIST OF FIGURES

Figure 1) Oracle VM architecture. 5
Figure 2) Physical and logical storage configuration of NetApp FAS controllers. 7
Figure 3) RAID-DP. 11
Figure 4) High-availability cluster. 11
Figure 5) NetApp SyncMirror. 11
Figure 6) Multimode VIF. 14
Figure 7) Single-mode VIF. 14
Figure 8) Second-level single-mode or multimode VIFs. 15
Figure 9) Storage provisioning. 20
Figure 10) Traditional array. 23
Figure 11) NetApp FAS array. 24
Figure 12) Thin-cloning deployment with Oracle VM Server. 27
Figure 13) Thin-cloning deployment with Oracle VM Server. 28
Figure 14) Contents of the system.img file. 29
Figure 15) SnapMirror deployment with Oracle VM. 30
Figure 16) Stretch MetroCluster. 31
Figure 17) Fabric MetroCluster. 31
Figure 18) Configuring NFS storage on Oracle VM Server in clustered Data ONTAP. 34
Figure 19) Configuring NFS storage on Oracle VM Server in 7-Mode. 37
Figure 20) FC storage configuration using NetApp plug-in for OVM. 39
Figure 21) Configuring iSCSI shared storage on Oracle VM server. 53
Figure 22) Thin cloning. 63

1 Introduction

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

1.1 Purpose and Scope of the Document

This document 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 who are familiar with concepts of Oracle VM Server v2.2 and higher and the NetApp Data ONTAP® 7G operating system and clustered Data ONTAP.

2 Overview of the Oracle VM and NetApp Storage Solution

2.1 Advantages of Virtualization

Oracle's application-driven server virtualization strategy is to integrate virtualization into every layer of the Oracle stack offering, from operating systems with Oracle Linux® and Oracle Solaris to engineered systems to Oracle Database and software, including middleware and applications and management. Virtualization has to address the entire compute stack in order to provide IT with the control, flexibility, and agility needed for modern data centers or the cloud.

Oracle VM is at the center of Oracle Server virtualization strategy. Designed for rapid application deployment and simplified lifecycle management, Oracle VM supports Windows®, Solaris, and Linux workloads and is optimized for Oracle Database and Oracle applications.

Oracle VM's integrated approach allows simplified integration with storage and other infrastructure solutions to deliver:

- Improved server consolidation
- Rapid enterprise application deployment
- Reduction of deployment and management complexities
- Simplification of data management tasks
- Access to capacity on demand

2.2 Oracle VM Application-Driven Virtualization for the Enterprise

Oracle VM is Xen hypervisor based and fully supports both Oracle and non-Oracle applications. Oracle VM is designed with a highly scalable architecture to support enterprise applications. Oracle VM consists of two components: the Oracle VM Server, the server software; and the Oracle VM Manager, the management console.

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, so that resources are fully utilized, and to move live virtual machines from one physical server to another without reconfiguring 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 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.

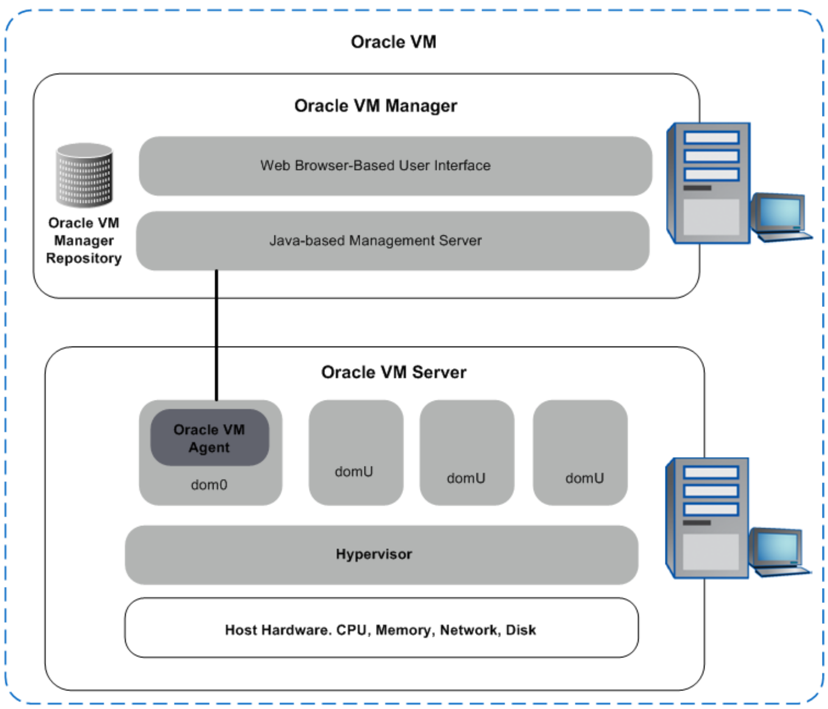

2.3 NetApp Storage Solutions for Application-Driven Virtualization

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) and NAS, including pNFS

- 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
- Compression, thin provisioning, and cloning, which provide additional benefits for storage efficiency
- Spending less time and fewer resources on running business-critical workloads
- Ability to achieve peace of mind with nondisruptive operations
- Ability to deliver on-demand responses to changing application needs
- Scalability for performance and capacity
- Cost and performance using Flash Pool™ and Flash Cache™ intelligent caching, solid-state drives

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® technology (a licensable feature) is another storage virtualization approach provided by NetApp that 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, refer to TR-3462, Storage Virtualization and DR Using vFiler.

Storage management and monitoring are critical to the successful operation of a virtual infrastructure. The NetApp Operations Manager suite of products can be used for day-to-day activities on storage systems, such as discovering storage systems, monitoring device health, determining the capacity utilization and performance characteristics of a storage system, configuring alerts and thresholds for event management, 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 in which an application (such as Oracle Database) administrator, virtual server (such as an Oracle VM) administrator, and storage (NetApp) administrator may have to work in tight synergy, role-based control from 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

The Oracle Linux and Oracle VM Hardware Certification List (HCL) provides users with information about x86 servers, x86 64-bit servers, and storage arrays certified for Oracle Linux with Unbreakable Enterprise Kernel or Oracle VM.

NetApp collaborated with Oracle to create a certification list for Oracle VMs with NetApp storage controllers. For more details, refer to the Oracle VM Hardware Certification List with NetApp Storage.

For more information on the Oracle VM Hardware Certification List, refer to the HCL website.

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.

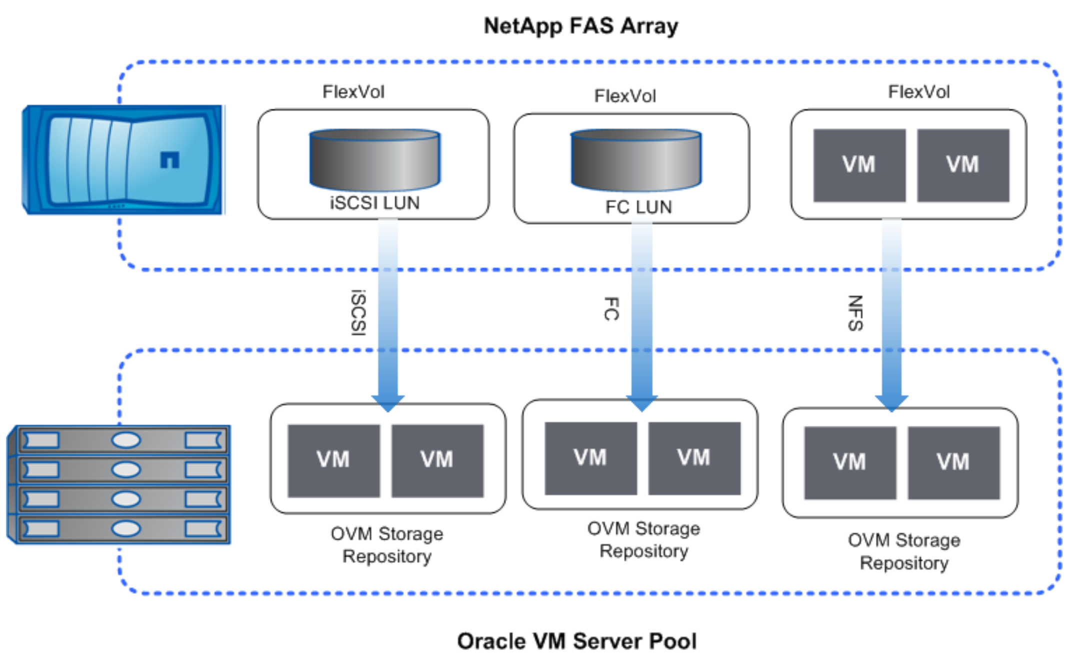

2.6 Plug-Ins

NetApp Plug-Ins

NFS Deployment with Clustered Data ONTAP showmount Plug-In

The `showmount -e` tool enables Oracle VM and other Linux operating systems that require s`howmount -e` to function.

Oracle's generic NFS plug-in requires the `showmount -e` command during storage registration. In a NetApp storage controller operating in 7-Mode, `showmount -e` works without any additional plug-ins. However, in a NetApp clustered Data ONTAP storage controller, the `showmount -e` command is not enabled at the storage level. You need to install the showmount plug-in in the OVM Server.

This tool provides 7-Mode functionality in clustered Data ONTAP for the `showmount -e` command. It is executed by a generic Oracle NFS plug-in as a workaround solution until the feature is added to Data ONTAP 8.3.

The `showmount -e` tool is a set of scripts that must be copied to the client machines on which the `showmount -e` command will be executed.

Download the showmount plug-in from http://support.netapp.com/NOW/download/tools/showmount_plugin_cdot/.

SAN Deployment—NetApp Plug-In for Oracle VM

The NetApp Plug-in for Oracle VM enables provisioning and cloning of Oracle VMs quickly and space efficiently from the Oracle VM Manager. The NetApp Plug-in for Oracle VM significantly reduces the cost and complexity of managing VMs that rely on NetApp storage and Oracle technologies.

The aggregate and volume provisioning are the prerequisites for the plug-in, which is explained in section 4.0.

Download and install the NetApp Plug-in for Oracle VM from here: http://mysupport.netapp.com/NOW/download/tools/ntap_plugin_ovm/. You can check the detailed deployment video at https://communities.netapp.com/videos/3533.

Oracle Plug-Ins

Oracle Generic NFS Plug-In and Oracle Generic SCSI Plug-In

Generic plug-ins offer a limited set of standard storage operations on virtually all storage hardware, such as discovering and operating on existing storage resources. We categorize these operations as "passive" in the sense that they do not interact with the storage management but simply detect the available storage architecture and allow it to be used in the Oracle VM environment.

Refer to http://docs.oracle.com/cd/E35328_01/E35332/html/vmusg-storage-plugin.html.

2.7 Choose Your Plug-In

There is no third-party-vendor storage connect plug-in for NFS, and all the vendors have to use an Oracle generic network file system plug-in for NFS storage. Table 1 compares the Oracle generic NFS plug-in with the NetApp SAN plug-in.

Table 1) Oracle generic NFS plug-in versus NetApp SAN plug-in.

Feature	NetApp Plug-in for OVM	Oracle Generic SCSI Plug-in
Allow custom clone names	Yes	No
Allow custom Snapshot copy names	Yes	No
Allow custom clone names	Yes	No
Allow custom Snapshot copy names	Yes	No
Storage name must be set to access the storage server	No	No
Support access controls	Yes	No
Support clone	Fully supported	No

Feature	NetApp Plug-in for OVM	Oracle Generic SCSI Plug-in
Support clone from Snapshot copy	Fully supported	No
Support clone splitting	Fully supported	No
Support clone splitting while clones are open and active	Yes	No
Support resize	Fully supported	No
Support Snapshot technology	Fully supported	No
Synchronous clone	No	No
Synchronous clone from Snapshot copy	Yes	No
Synchronous clone splitting	No	No
Synchronous resize	No	No
Synchronous Snapshot copy	Yes	No

2.8 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 machine 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 v3.2.X**

  File Server

  – Find the NetApp NFS volume (volume group) from 7-Mode or clustered Data ONTAP (file system) that you want to use to create the storage repository. Then create the storage repository, grant access to OVM Server, and refresh.

```
[root@stlrx200s5-26 ~]# ssh -l admin <localhost | ovm manager > -p 10000
admin@localhost's password:

OVM> list filesystem
OVM> refresh FileSystem name=nfs:<exported volume or junchtion path>
OVM> create Repository name=<RepoName> fileSystem=nfs:< exported volume or junchtion path>
OVM> add Server name=<OVM Server> to Repository name=<RepoName>
OVM> refresh Repository name=<RepoName>
```

  – /nfsmnt holds the server pool and its ovspoolfs image.

  SAN Server

  – Find the NetApp SAN volume (volume group) from 7-Mode or clustered Data ONTAP (file system) that you want to use to create the storage repository. Then create the physical disk, map the physical disk with the access group, create the repository, grant access to OVM Server, and refresh.

```
[root@stlrx200s5-26 ~]# ssh -l admin <localhost | ovm manager > -p 10000
admin@localhost's password:
OVM> list volumegroup
OVM> create physicaldisk size=<in gb> name=<name of physical disk> shareable=<yes|no> on
volumegroup name=<volume groupname>
OVM> list physicaldisk
OVM> list serverpool
OVM> list accessgroup
OVM> add physicaldisk name=<physicaldisk name> to accessgroup name=<accessgroup name>
OVM> create Repository name=<repository name> serverPool=<serverpool name> physicalDisk=<physical
diskname>
OVM> list server
OVM> add server name=<OVM server name> to repository name=<repository name>
OVM> refresh repository name=<repository name>
```

- The /OVS directory is the location of the default storage repository created when Oracle VM Server is installed with the default partition layout, which is common for both SAN and NAS.
- /poolfsmnt/<UUID>/db is the location of the default server pool for SAN.

- 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 such as NFS, iSCSI, and FC, refer to section 4.1.

Section 0 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 Resiliency 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 are 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 (such as Oracle VM) are greater than those in a traditional physical server infrastructure. Therefore, 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 you should consider when deploying Oracle VM Server virtualization on NetApp storage solutions.

RAID Data Protection

RAID-DP® technology 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, refer to TR-3505: NetApp Deduplication for FAS and V-Series Deployment and Implementation Guide.

Figure 3) RAID-DP.

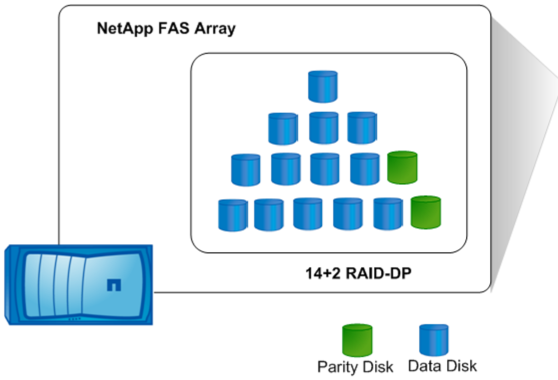

NetApp RAID-DP
– *Protects against double disk failures*
– *Provides high storage utilization (87.5% by default)*
– *No performance impact*
– *Least probability of data loss over a period of time*

	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, refer to:

- 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.

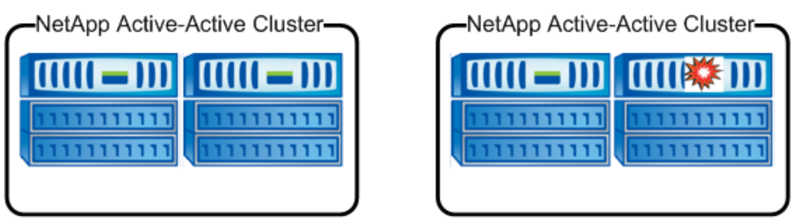

NetApp active-active provides high-availability from controller failures and up to two-disk failure (RAID-DP)

The NetApp HA cluster model can be enhanced by synchronously mirroring data at the RAID level using NetApp SyncMirror® technology in 7-Mode. 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.

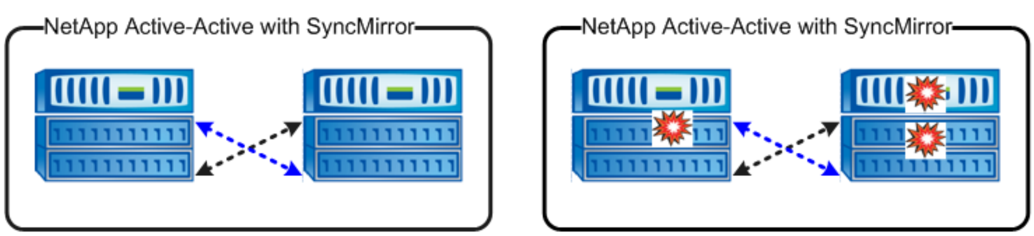

SyncMirror with NetApp active-active cluster provides high availability from complete RAID groups and disk shelves failure in addition to controllers.

Note: The NetApp HA cluster model with clustered Data ONTAP can be enhanced by asynchronously mirroring data through SnapMirror.

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 intershelf 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, refer to TR-3437: Storage Subsystem Resiliency Guide.

Remote LAN Management Card

The Remote LAN Management (RLM) card provides 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, refer to http://support.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 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® technology 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, refer to:

- 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 for a truly redundant architecture.

Best Practice

For designing an FC or IP storage network infrastructure, refer to the “FC/iSCSI Configuration Guide” on the NetApp Support site.

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 network connection fault tolerance 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) Multimode VIF.

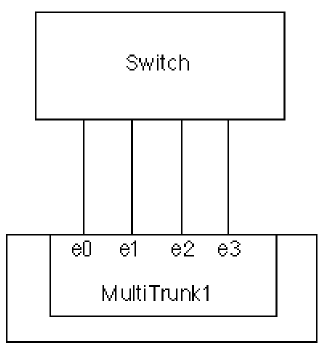

In a single-mode VIF, only one of the physical connections is active at a time. If the storage controller detects a fault in the active connection, a standby connection is activated. No configuration is necessary on the switch to use a single-mode VIF, and the physical interfaces that constitute the 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.

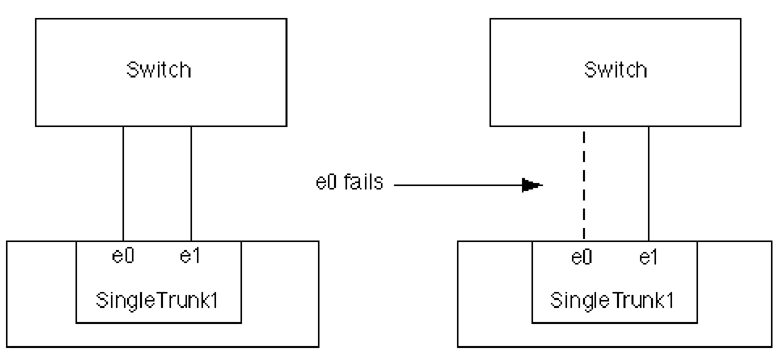

It is also possible to create second-level single-mode or multimode VIFs. By using second-level VIFs, you can take advantage of both the link aggregation features of a multimode VIF and the failover capability of a single-mode VIF.

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, refer to the "Data ONTAP Network Management Guide" available at the NetApp Support site.

Figure 8) Second-level single-mode or multimode VIFs.

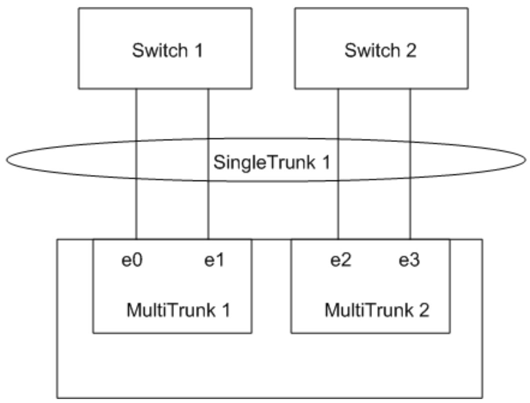

iSCSI Traffic Security

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 Practices

- 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 in 7-Mode

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 with accessibility across all storage ports. NetApp also strongly recommends using Data ONTAP version 7.3 or higher for 7G and 8.0 or higher for 7-Mode.

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

1. Log in to the NetApp console using either SSH, telnet, or the 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, refer to 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. Three of the common values of mode used are:

- **From OVM 3, Mode 0 - balance-rr.** Round-robin policy not supported and mode1 (Linux active/standby), mode4 (802.3ad), and mode6 (Linux balance-alb) are supported**.**
- **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. It’s the default mode.
- **The 802.3ad policy (mode 4) is the preferred mode for Oracle VM,** but this requires specialized connectivity (all interfaces in the same bond/aggregate must be connected to the same switch) and configuration on the switches (LACP/EtherChannel).

Steps for Creating Network Bonding in OVM Server

To create a network binding in the OVM Server, complete the following steps.

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`

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.

Network Flexibility/Portability

NetApp VIFs support multiple IP addresses per SVM and you can maintain individual IP addresses for storage-related resources on a per-pool basis. For example, if a single Oracle VM Manager controls three server pools, the NetApp VIF would have a canonical IP/host name and three additional IPs/host names such as mynfspool1-nfs, mynfspool2-nfs, and mynfspool3-nfs. The NFS pool file systems and repositories for each pool would use a unique IP/host name for those storage objects related to each individual pool (PoolFS and repositories). This would allow the storage resources for a single pool to be relocated to another storage controller without affecting the other two server pools.

Best Practice
The storage repositories created through NFS can be shared between multiple server pools, which provide Oracle VM Guest resources such as ISO images, templates, and assemblies that can be shared across multiple server pools.

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 datasets 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 Practices
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.

- For 7-Mode:

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

- For clustered Data ONTAP:

```
volume modify -vserver <vservername> -volume <volumename>
-percent-snapshot-space 0 -fractional-reserve 0 -read-realloc on -snapshot-policy none
```

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—refer to section 4.1.

Figure 9) Storage provisioning.

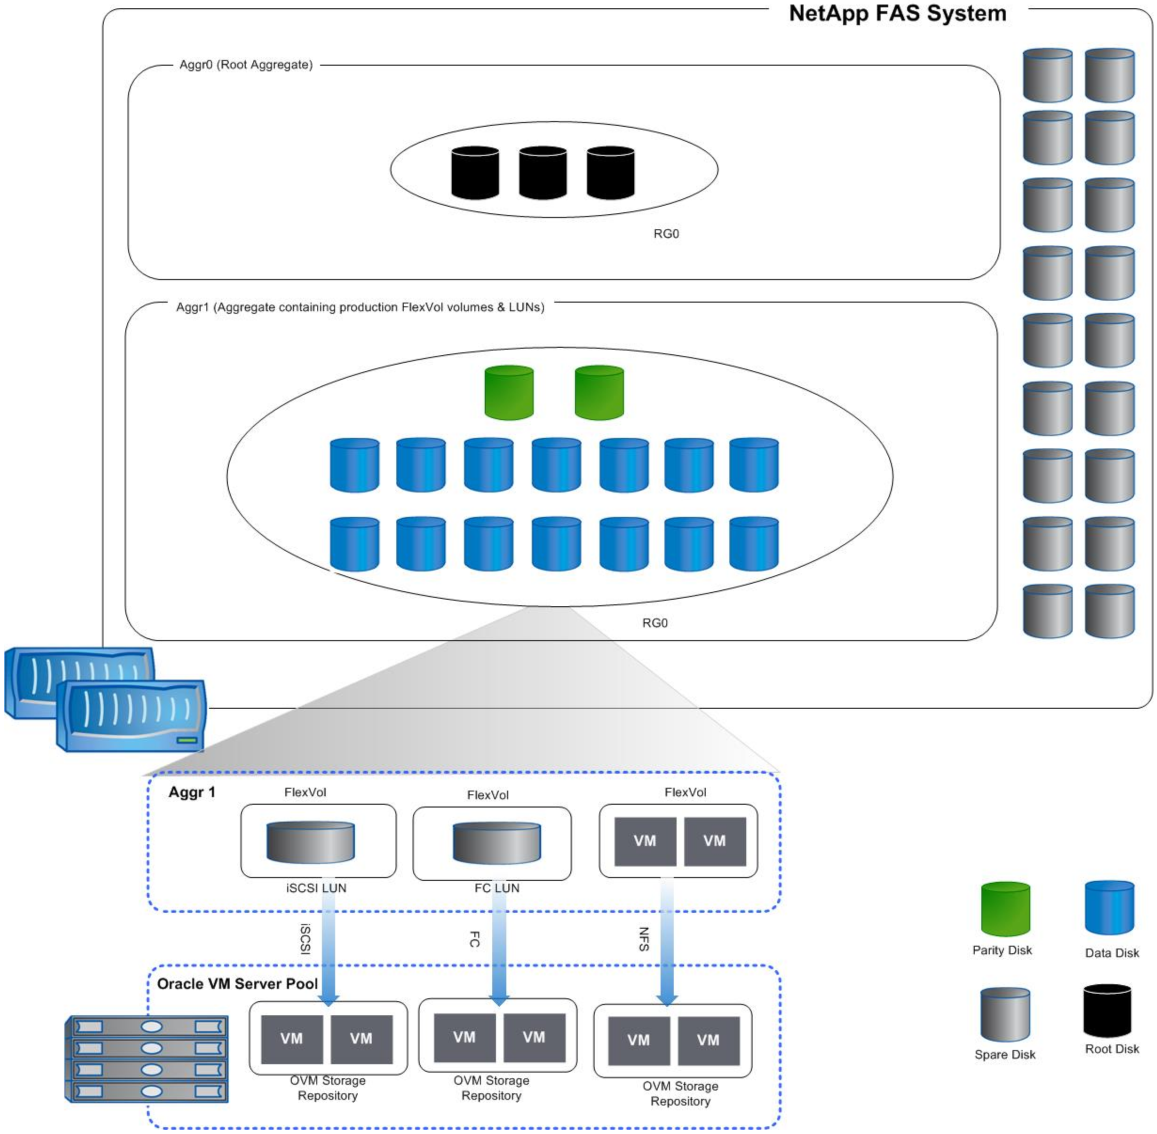

Recommendations for Effectively Using Volume Separation

Best Practices

- Never combine the pool file systems (NFS, iSCSI, or FCP) from multiple pools into a single volume. Overallocation of space in one server pool can cause other unrelated server pools to fail. If you fill a volume 100%, the pool file system residing in that volume will stop functioning, which will reboot all servers in the pool and then fail to start any Oracle VM Guests. So, if you have a single volume containing the pool file system for multiple server pools, then all of the server pools will be affected the same.
- Create a separate volume for each server pool to contain a single pool file system and any qtrees/LUNs used for repositories.
- Do not attempt to create a storage repository for each individual Oracle VM Guest. This is not a scalable solution and dramatically limits the number of virtual machines managed in a single server pool.

3.4 Oracle VM Storage Best Practices—Create, Import, Live Migrate

VM Creation (NFS)—3.2.x

1. Create and start a VM from the template:

```
OVM> list vm
Command: list vm
Status: Success
Time: 2013-12-10 11:28:41,928 EST
Data:
  id:0004fb00000600004153be5456636d62  name:tvm1
  id:0004fb00000600009307cb973a7c8c09  name:VM1
  id:0004fb0000140000cefd404f83bb5c2a  name:OVM_EL5U5_X86_64_PVM_10GB.tgz
OVM> list serverpool
Command: list serverpool
Status: Success
Time: 2013-12-10 11:28:45,791 EST
Data:
  id:0004fb0000020000345b946c0aa210d3  name:NetAppCDOT_NFS_ServerPool
OVM> clone vm name=OVM_EL5U5_X86_64_PVM_10GB.tgz destType=Vm destName=TemplateVM
serverPool=NetAppCDOT_NFS_ServerPool
Command: clone vm name=OVM_EL5U5_X86_64_PVM_10GB.tgz destType=Vm destName=TemplateVM
serverPool=NetAppCDOT_NFS_ServerPool
Status: Success
Time: 2013-12-10 11:31:47,361 EST
Data:
  id:0004fb0000060000cc16548b15f1d3c3  name:TemplateVM
OVM>
OVM> list vm
Command: list vm
Status: Success
Time: 2013-12-10 11:33:25,189 EST
Data:
  id:0004fb0000060000cc16548b15f1d3c3  name:TemplateVM
  id:0004fb00000600004153be5456636d62  name:tvm1
  id:0004fb00000600009307cb973a7c8c09  name:VM1
  id:0004fb0000140000cefd404f83bb5c2a  name:OVM_EL5U5_X86_64_PVM_10GB.tgz
OVM>
```

2. Add the networking to the virtual machine:

```
OVM> add Vnic name=00:21:f6:00:00:07 to Vm name=TemplateVM
Command: add Vnic name=00:21:f6:00:00:12 to Vm name=TemplateVM
Status: Success
Time: 2013-12-10 13:51:11,919 EST
OVM>
OVM> edit Vnic name=00:21:f6:00:00:07 network=10.61.173.0
Command: edit Vnic name=00:21:f6:00:00:12 network=10.61.173.0
Status: Success
```

```
Time: 2013-12-10 13:58:55,618 EST
OVM>
```

3. Start the virtual machine:

```
OVM> start Vm name=TemplateVM
Command: start Vm name=TemplateVM
Status: Success
Time: 2013-12-10 14:01:56,993 EST
OVM>
```

VM Creation (SAN)—3.2.x

To create a VM and to use the thin-clone procedure, refer to section 4.3, “Configuring FC Shared Storage and VM Operations on Oracle VM Server.”

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—OVM 3.2.x

Select the VM to be migrated, then right-click and select Migrate or click the Migrate icon. Select the target-compatible OVM Server and click Confirm.

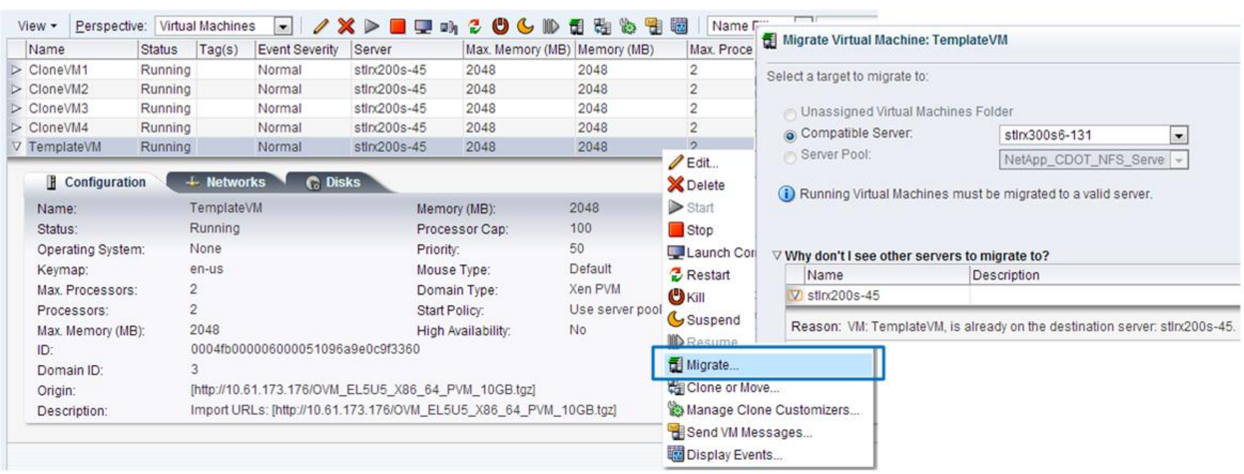

From the Oracle VM Server 2.x and 3.x CLI

Enter one of the following commands:

```
xm migrate <vmname/Domin ID> <destination-OVM-Server> --live
```

Or:

```
OVM> migrate Vm name=TemplateVM destServer=stlrx300s6-131
Command: migrate Vm name=TemplateVM destServer=stlrx300s6-131
Status: Success
Time: 2013-12-10 16:53:18,003 EST
OVM>
```

3.5 Increasing Storage Utilization Using 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.

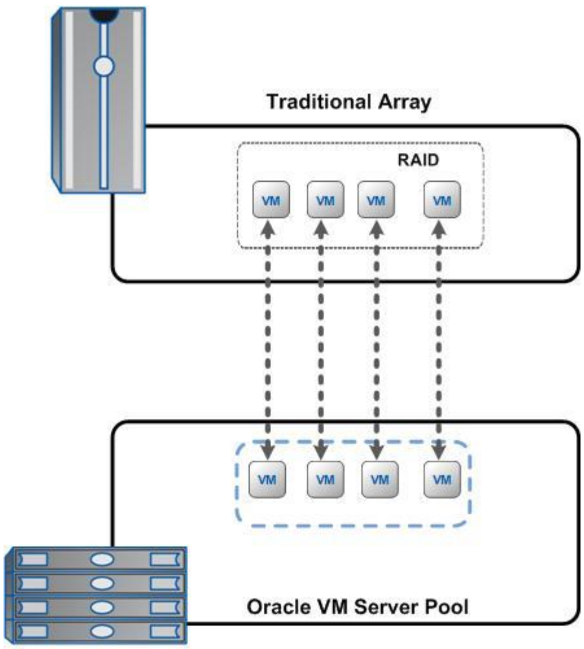

Deduplication technologies from NetApp assist Oracle VM deployments in eliminating duplicate data in their environment, enabling greater storage utilization on the production environment.

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.

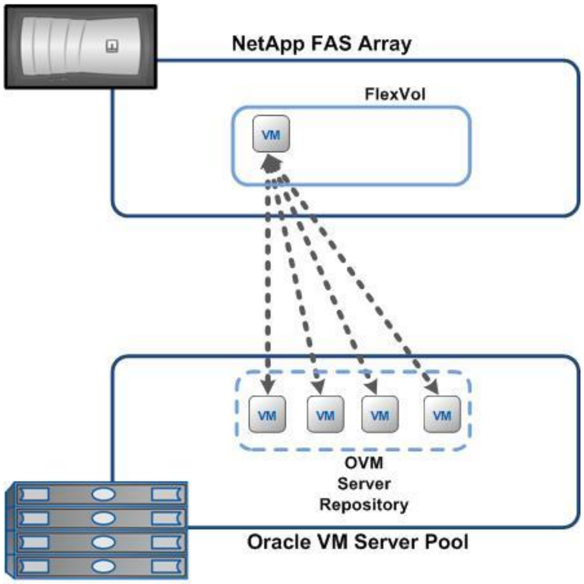

Deduplication is enabled on a volume; the amount of data deduplication realized is based on the commonality of the data stored in a deduplication-enabled volume.

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. This is 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, refer to 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 4.3.

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 the 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; that is, NFS, iSCSI, and FC.

Best Practices

- If using NFS storage, NetApp flexible volumes are thin-provisioned by default. No extra configuration steps are necessary.
- When using iSCSI or FC storage, make sure that the Space Reserved checkbox in the LUN wizard is not selected.
- 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, complete the following steps.

1. Log in to the NetApp console.
2. Set the volume autosize policy.
   – For 7-Mode, enter:

```
vol autosize <vol-name> [-m <size>[k|m|g|t]] [-i <size>[k|m|g|t]] on
```

   – For clustered Data ONTAP, enter:

```
vol autosize –vserver <vservername> -volume <volname> -maximum-size <size> [KB|MB|GB|TB|PB] -
increment-size <size> [KB|MB|GB|TB|PB]} -mode grow
```

```
vol autosize <volname>
example:

TESO::> volume autosize -vserver vs2_dnfs_rac -volume ovm_nfs_repository -maximum-size 500GB -
increment--mode grow
    -increment-size    -increment-percent
TESO::> volume autosize -vserver vs2_dnfs_rac -volume ovm_nfs_repository -maximum-size 500GB -
increment-size 10GB -mode grow
vol autosize: Flexible volume "vs2_dnfs_rac:ovm_nfs_repository" autosize settings UPDATED.

Volume modify successful on volume: ovm_nfs_repository

TESO::> volume autosize -vserver vs2_dnfs_rac -volume ovm_nfs_repository
Volume autosize is currently ON for volume "vs2_dnfs_rac:ovm_nfs_repository".
The volume is set to grow to a maximum of 500g in increments of 10g when the volume used space is
above 90%.
Volume autosize for volume 'vs2_dnfs_rac:ovm_nfs_repository' is currently in mode grow.

TESO::>
```

Snapshot Autodelete 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.

1. Log in to the NetApp console.
2. Set the Snapshot copy autodelete policy and modify the volume autodelete policy.
   - For 7-Mode, enter:

```
snap autodelete <vol-name> commitment try trigger volume target_free_space 5 delete_order
oldest_first
vol options <vol-name> try_first volume_grow
```

   - For clustered Data ONTAP, enter:

```
TESO::> snapshot autodelete modify -vserver <vservername> -volume <volumename> -enabled true -
commitment try -delete-order oldest_first -target-free-space 5% -trigger volume

TESO::> volume modify -vserver <vservername> -volume <volumename> -space-mgmt-try-first
volume_grow
```

Example:

```
TESO::> volume snapshot autodelete modify -vserver vs2_dnfs_rac -volume ovm_nfs_repository -
enabled true -commitment try -delete-order oldest_first -target-free-space 5% -trigger volume

Volume modify successful on volume: ovm_nfs_repository

TESO::>

TESO::> volume modify -vserver vs2_dnfs_rac -volume ovm_nfs_repository -space-mgmt-try-first
volume_grow

Volume modify successful on volume: ovm_nfs_repository

TESO::>
```

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.

1. Log in to the NetApp console.
2. Set the volume Snapshot fractional reserve.

– For 7-Mode, enter:

```
vol options <vol-name> fractional_reserve 0
```

– For clustered Data ONTAP, enter:

```
TESO::> volume modify -vserver <vservername> -volume <volname> -fractional-reserve 0%

Example:

TESO::> volume modify -vserver vs2_dnfs_rac -volume ovm_nfs_repository -fractional-reserve 0%

Volume modify successful on volume: ovm_nfs_repository

TESO::>
```

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.

Different flavors of virtual cloning technology are available from NetApp: volume-level cloning (or FlexClone® volumes), LUN-level cloning, and file-level cloning.

Best Practices

- Depending on the requirement and necessity, decide the design of the cloning methodology (volume, LUN, or file level) and other NetApp technologies (such as deduplication) to be applied. There are many possibilities for achieving the same end result.
- File-level cloning can be used only if the cloned VMs need to reside on NFS storage repositories.
- 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 12 and 13 depict some common ways of using NetApp thin-cloning technologies with Oracle VM Server for an NFS storage repository. For the step-by-step process, refer to section 4.2.

Figure 12) Thin-cloning deployment with Oracle VM Server.

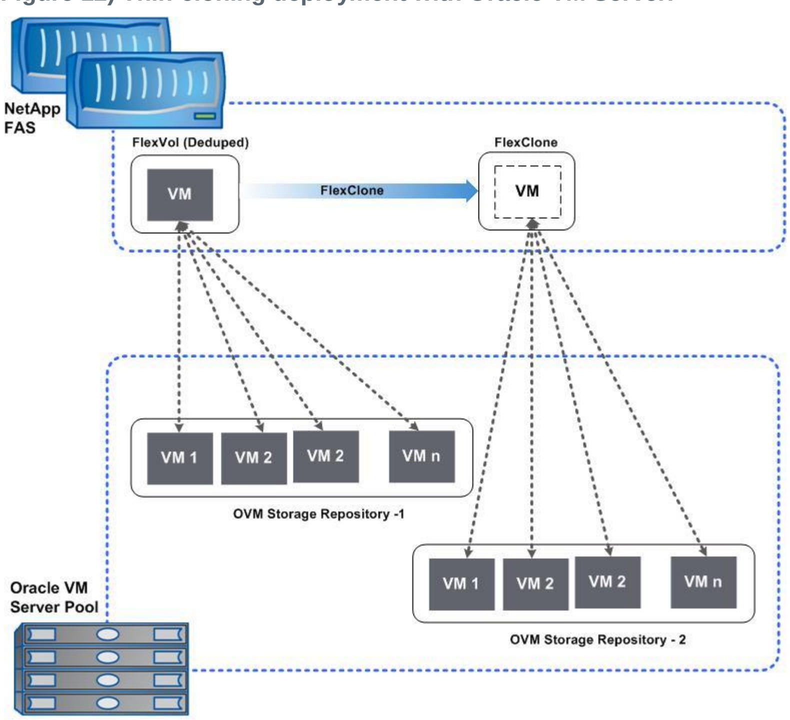

Figure 13) Thin-cloning deployment with Oracle VM Server.

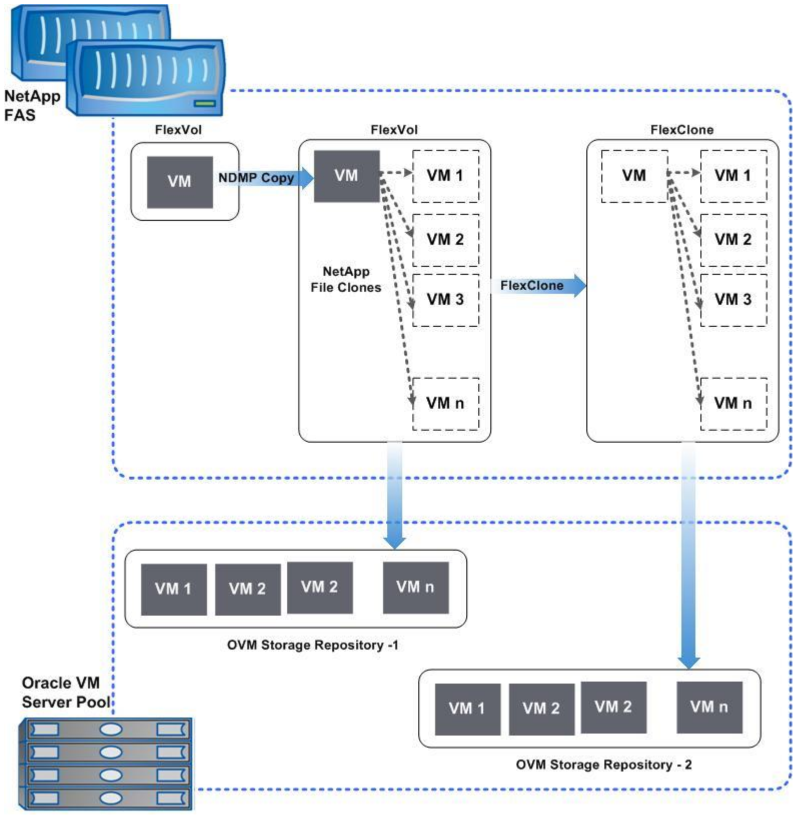

3.6 Disk-Based Snapshot Backups for Oracle VMs Using NetApp Snapshot and SnapRestore

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® technology along with SnapDrive® technology can be used to create application-consistent backups of the applications running inside the virtual machines. NetApp provides SnapManager products for several enterprise applications and databases, including Oracle, Microsoft® Exchange, MS SQL, SAP®, and so on.

To create a crash-consistent backup of an Oracle VM Server repository, complete the following steps.

1. Consider that an NFS volume from a NetApp FAS system has been mounted as a repository in the OVM Server (3.2.x).

```
OVM> show filesystem name=nfs:/ovm_nfs_repository
Command: show filesystem name=nfs:/ovm_nfs_repository
Status: Success
Time: 2014-01-06 15:43:01,114 EST
Data:
  Name = nfs:/ovm_nfs_repository
  Id = 308864d5-6d32-4e89-a24b-1df2a8748509
  Total Size (GiB) = 76.0
  Free Size (GiB) = 67.48
```

```
  Used Size (GiB) = 8.519999999999996
  Used % = 11.0
  Refreshed = Yes
  Path = 10.63.164.18:/ovm_nfs_repository
  Repository 1 = 0004fb0000030000c6cfa8ad15f46620  [NetApp_CDOT_NFS_Repository]
  FileServer = 0004fb00000900003f183c8b8b9e99f2  [vs2_dnfs_rac]
OVM>
```

2. Create a Snapshot copy of the volume in the NetApp FAS system using NetApp OnCommand® System Manager, the FilerView® tool, or the CLI.

   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 Image Files or Virtual Disks Using SnapRestore

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>
```

An example of the above command that uses a Snapshot copy named "OVMNFSSNAPSHOT" is as follows:

```
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 (such as `System.img`). Any lost file inside the virtual disk can then be copied for use 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.

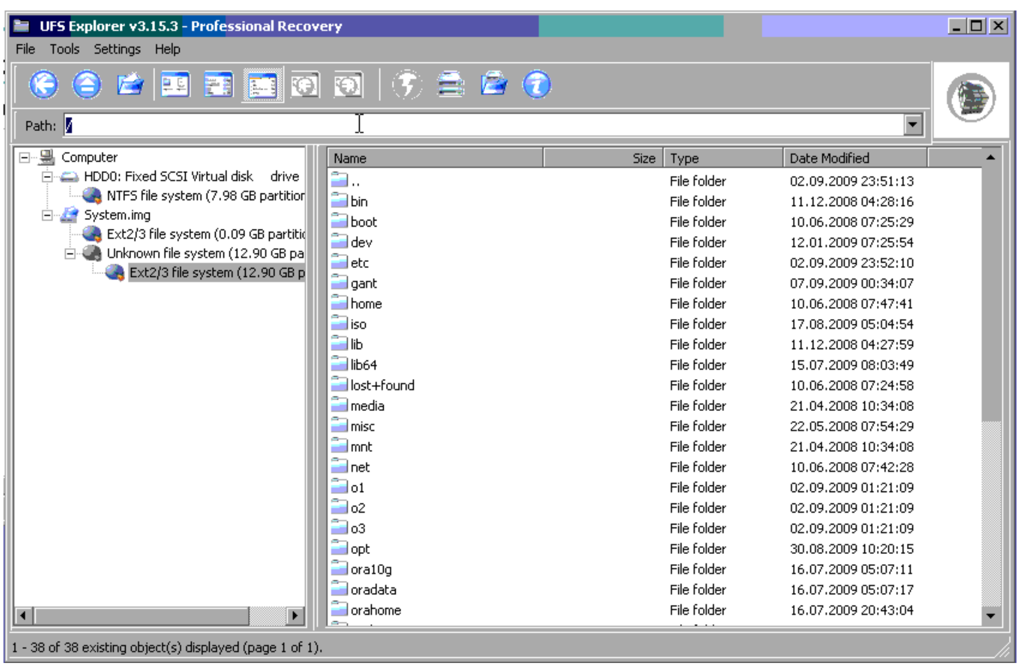

3.7 Disaster Recovery of an OVM Virtual Infrastructure Using NetApp SnapMirror

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 Data Protection Online Backup and Recovery Guide.

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

Figure 15) SnapMirror deployment with Oracle VM.

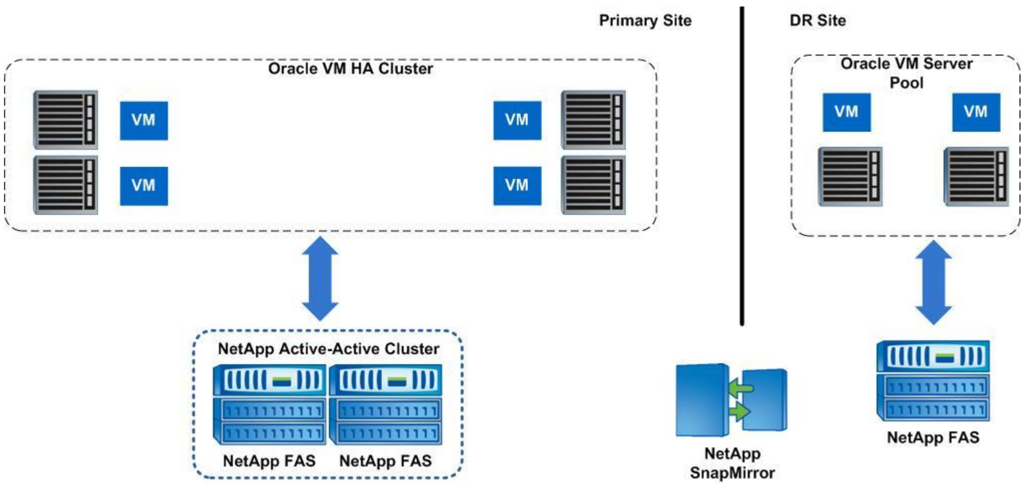

Best Practices
• NetApp SnapMirror Async best practices: TR-3446: SnapMirror Async Overview and Best Practices Guide • NetApp SnapMirror Sync and Semi-Sync best practices: TR-3326: SnapMirror Sync and SnapMirror Semi-Sync Overview and Design Considerations

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™ technology, refer to the NetApp Active-Active Configuration Guide.

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

Figure 16) Stretch MetroCluster.

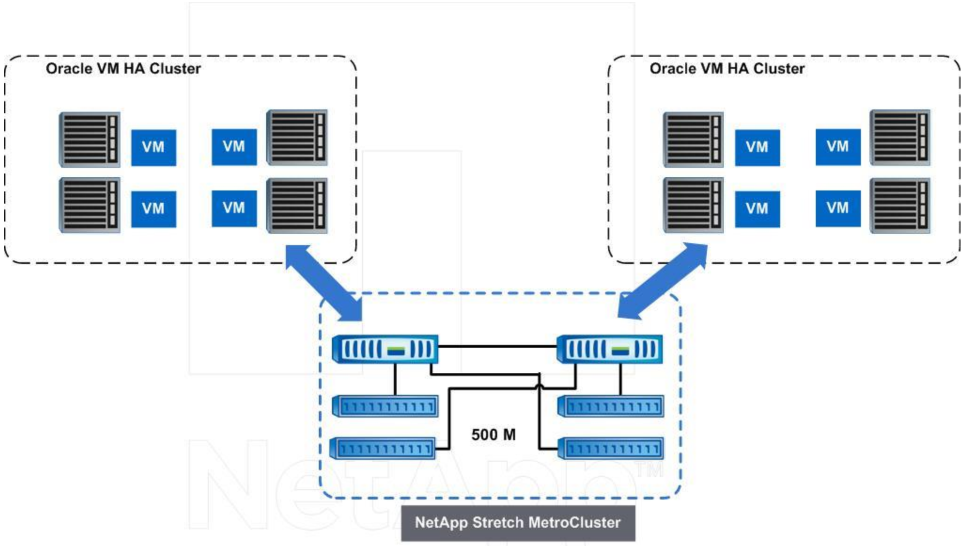

Figure 17) Fabric MetroCluster.

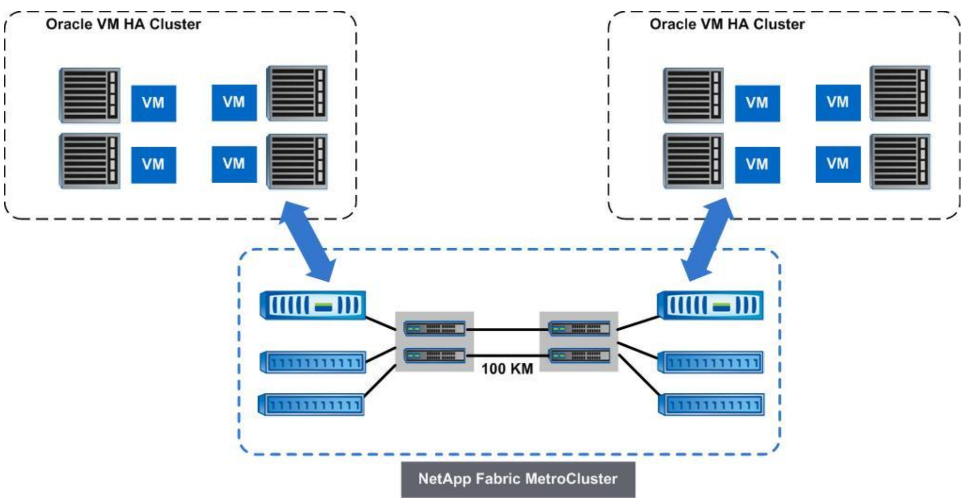

Best Practices
NetApp MetroCluster best practices: http://media.netapp.com/documents/tr-3548.pdf.

4 Sample Implementations and Best Practices

This section describes sample implementation and also provides best practice recommendations such as provisioning, configuration, deduplication, and thin cloning in OVM with a NetApp storage controller in 7-Mode and clustered Data ONTAP.

4.1 Provisioning Storage for Clustered Data ONTAP and 7-Mode

Creating an Aggregate in Clustered Data ONTAP

To configure an aggregate in clustered Data ONTAP, complete the following steps.

1. From NetApp OnCommand System Manager, click Cluster > Storage > Aggregates > Create.
2. Enter the aggregate name and select Dual Parity as the RAID type.

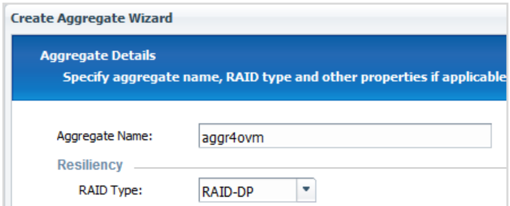

3. Select the Node (Storage Controller, Disk Selection) and Create Aggregate.

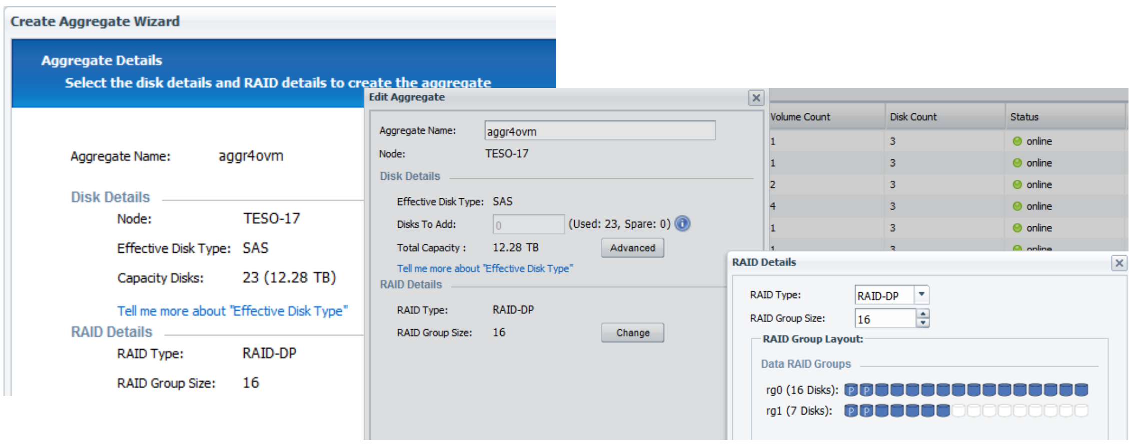

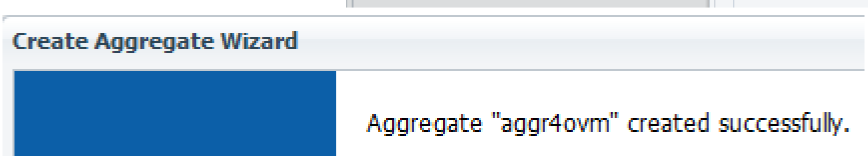

Creating an Aggregate in 7-Mode

Note: The following screenshots are based on the previous version of NetApp System Manager. Use the latest NetApp OnCommand System Manager as you would for storage provisioning with clustered Data ONTAP in the previous section.

1. From NetApp System Manager, start the Create Aggregate Wizard.
2. Set dual parity as the RAID type and make disk selection automatic.

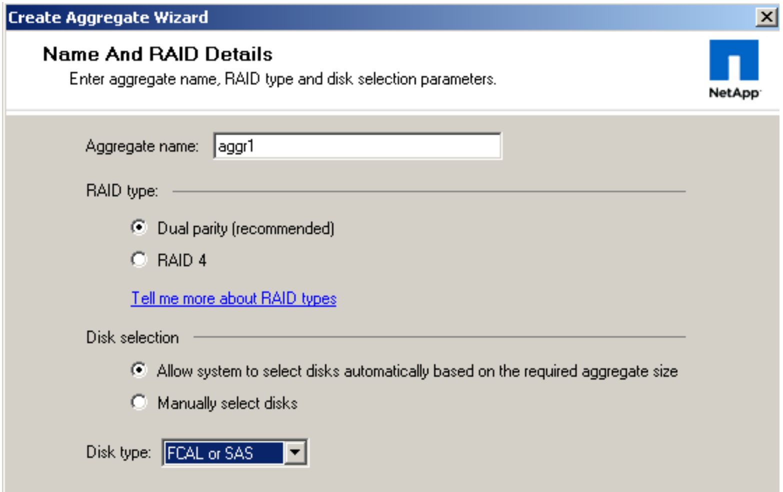

3. Select the size of the aggregate depending on the number of disks.

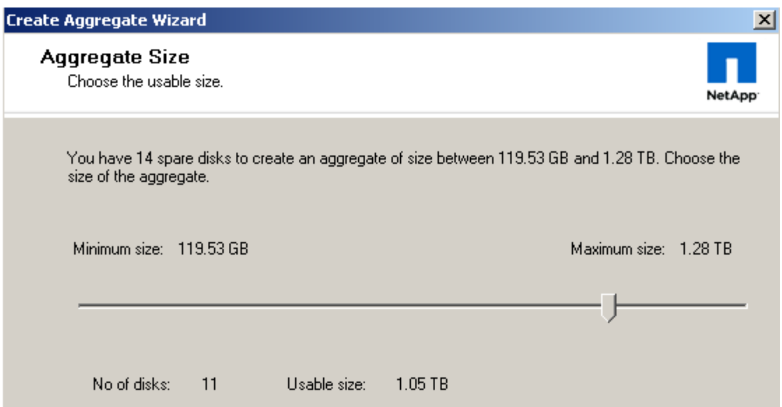

4. Commit the selection to complete the aggregate creation process.

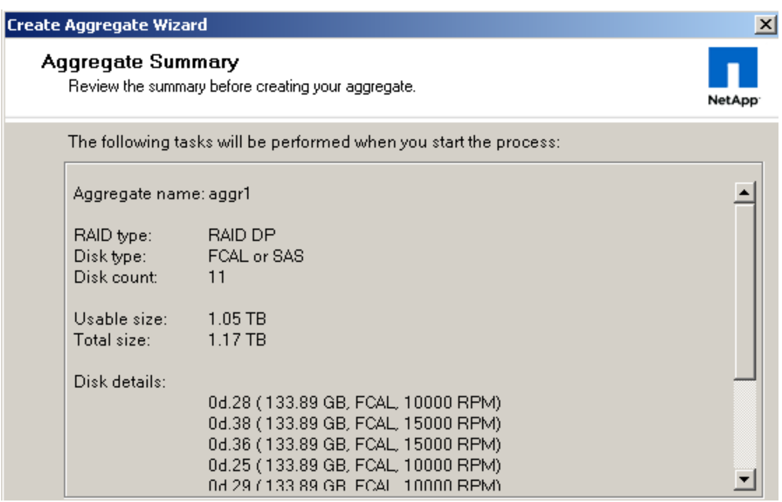

4.2 Configuring NFS Storage on Oracle VM Server

Configuring NFS Storage in Clustered Data ONTAP

Figure 18 shows an overview of the steps required to configure NFS storage on Oracle VM Server in a clustered Data ONTAP environment.

Figure 18) Configuring NFS storage on Oracle VM Server in clustered Data ONTAP.

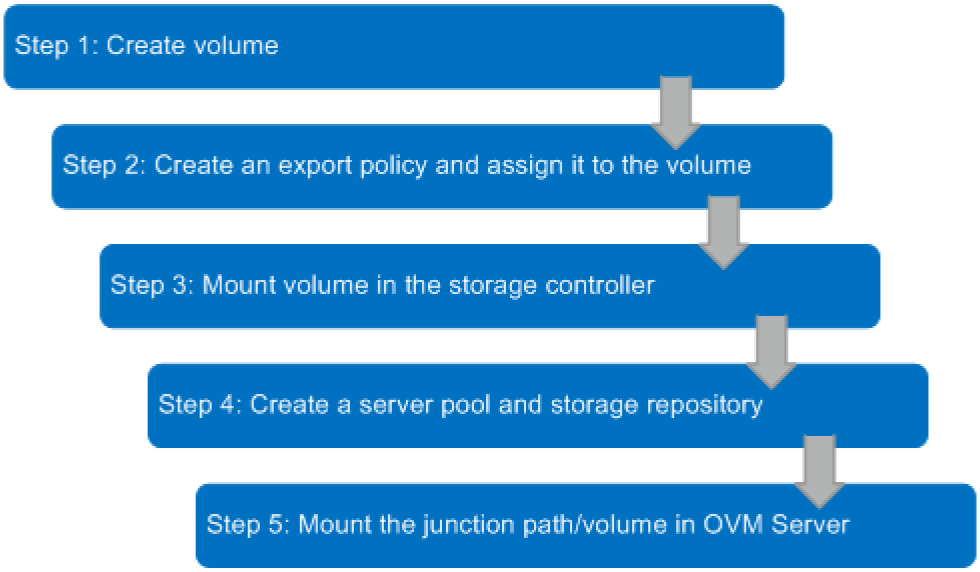

Step 1: Create Volume

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

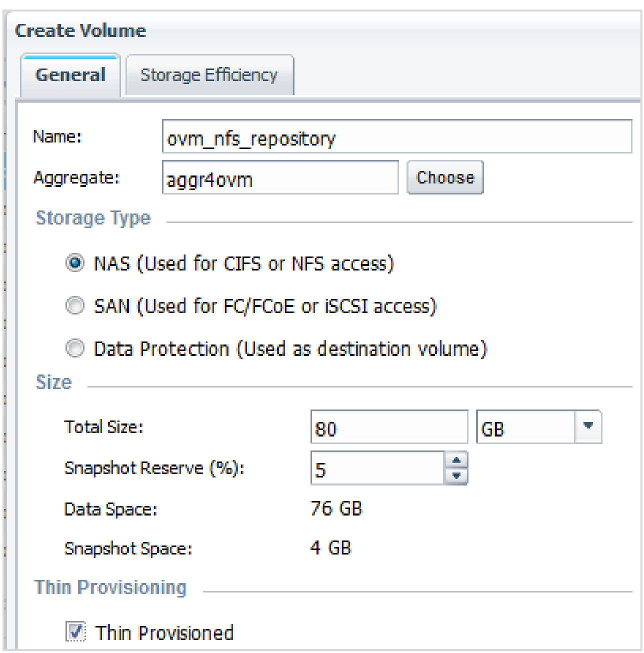

Step 2: Create an Export Policy

1. Create a policy with the required permission, protocol, and network access.

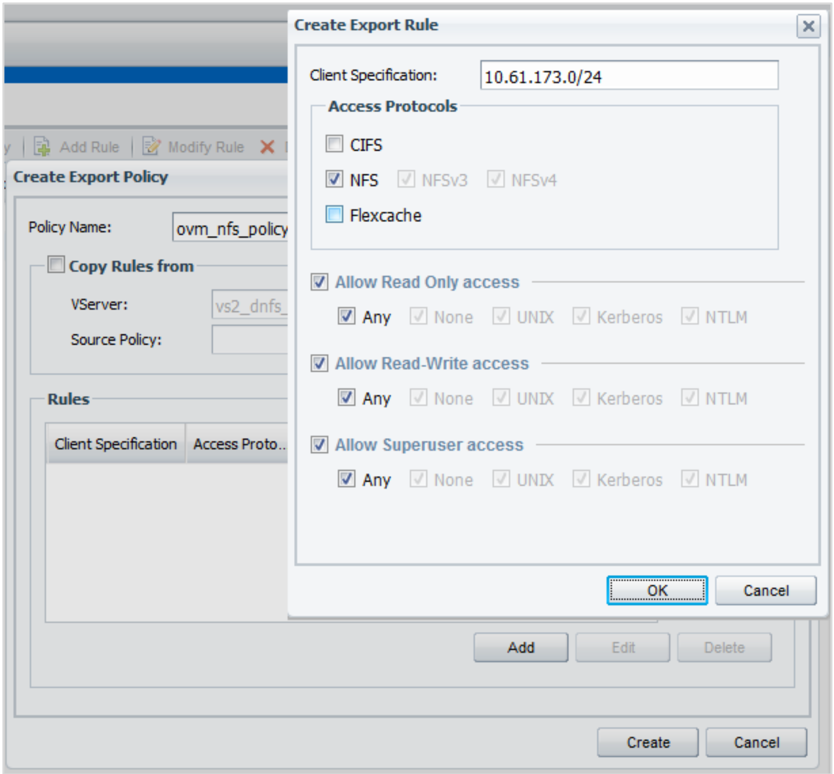

2. Assign the export policy to the volume.

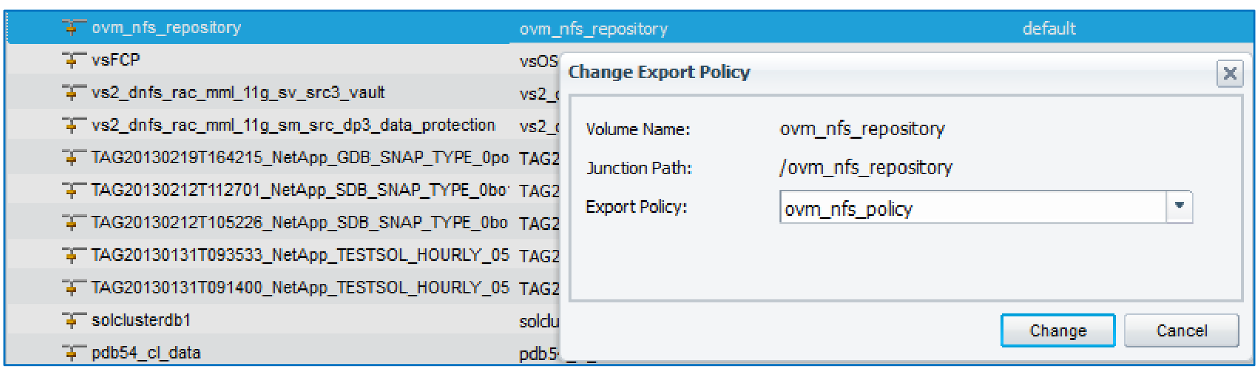

Step 3: Mount the Volume in the Storage Controller

Mount the volume in storage using the CLI or NetApp OnCommand System Manager by entering the following command.

```
TESO::> mount -vserver vs2_dnfs_rac -volume ovm_nfs_repository -junction-path /ovm_nfs_repository
  (volume mount)
```

Note: Follow the same procedure for the server pool as you would for storage repository volumes such as volume create. Use the exported policy and mount the volume in storage.

Step 4: Create a Server Pool and Storage Repository

1. Create a server pool using OVM Manager and the storage repository from the CLI.

In OVM Manager, navigate to Home > Servers and VMs > Server Pools. Right-click the Create Server Pool option.

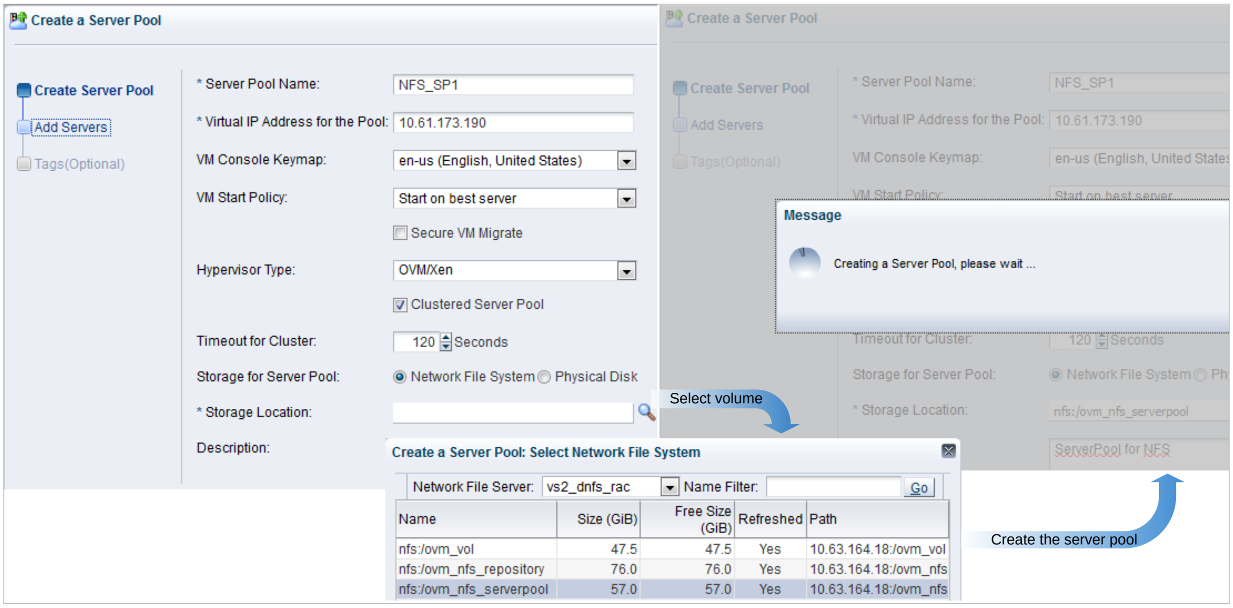

2. Create an NFS storage repository from the CLI or by using OVM Manager.

```
OVM> create Repository name=NetAppCDOT_NFS_Repository filesystem=nfs:/ovm_nfs_repository
Command: create Repository name=NetAppCDOT_NFS_Repository filesystem=nfs:/ovm_nfs_repository
Status: Success
Time: 2013-12-09 16:48:13,576 EST
Data:
  id:0004fb0000030
```

Step 5: Mount the Junction Path/Volume in OVM Server

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

```
10.63.164.18:/ovm_nfs_repository  /OVS/Repositories/0004fb0000030000c6cfa8ad15f46620 nfs
rw,vers=3,rsize=65536,wsize=65536,hard,proto=tcp,timeo=600 0 0
10.63.164.18:/ovm_nfs_serverpool  /nfsmnt/5b474bd3-da9a-4157-8b72-7a7c9ad28fc3 nfs
rw,vers=3,rsize=65536,wsize=65536,hard,proto=tcp,timeo=600 0 0
```

Configuring NFS Storage on Oracle VM Server in 7-Mode

Figure 19 shows an overview of the steps required to configure NFS storage on Oracle VM Server in 7-Mode.

Figure 19) Configuring NFS storage on Oracle VM Server in 7-Mode.

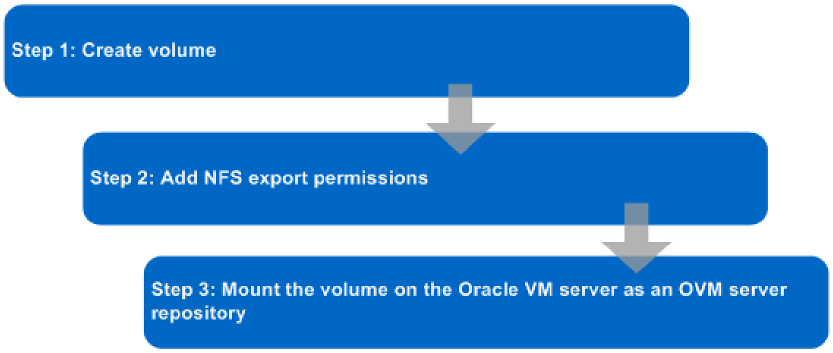

Note: Although screenshots used in this section are based on the previous version of NetApp System Manager, the procedure remains unchanged. Refer to the section "Configuring NFS Storage in Clustered Data ONTAP" to use the latest NetApp OnCommand System Manager.

Step 1: Create Volume

1. From NetApp OnCommand 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.

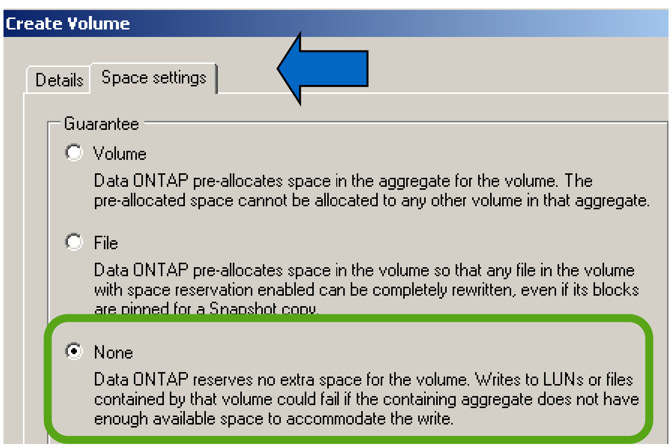

3. After the volume is created, its properties can be further modified by using the Volume autogrow settings.

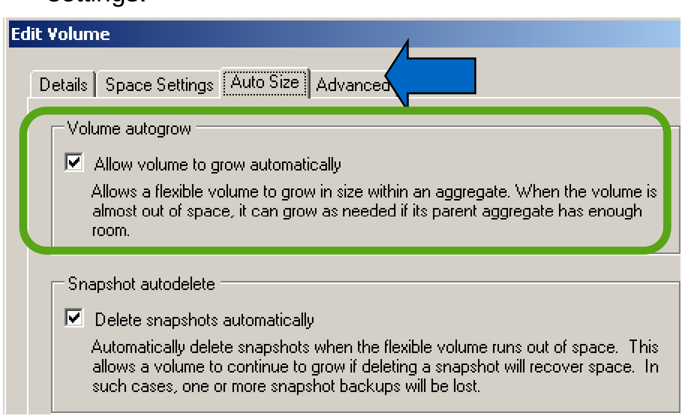

Step 2: Add NFS Export Permission

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

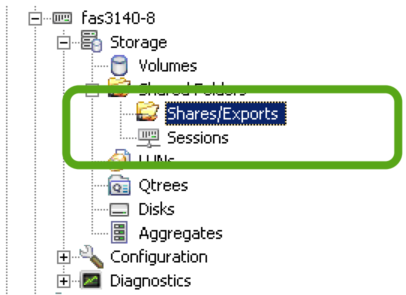

2. Select the volume on the right-hand pane and click Edit.

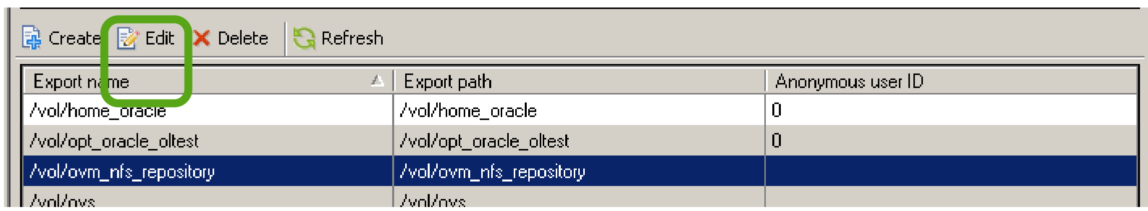

3. Add root access to the Oracle VM Server IP address and set the security setting.

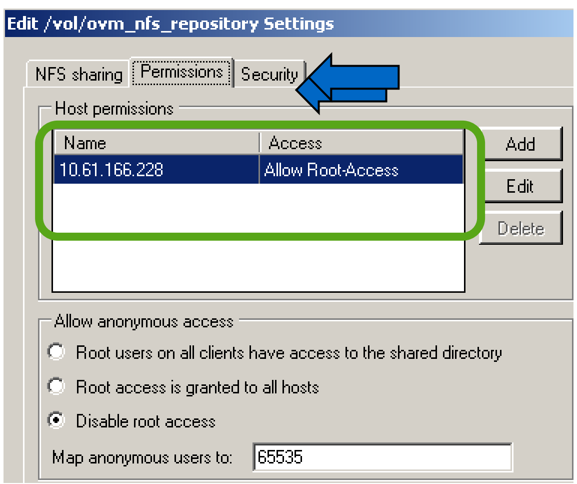

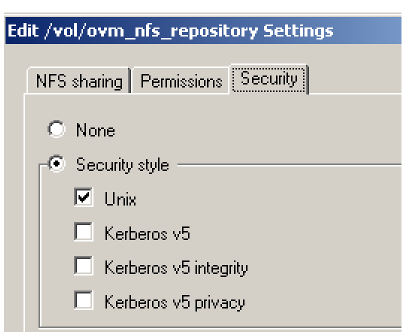

Step 3: Mount the Volume on the Oracle VM Server

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.

```
[root@AMDLoaner-1 /]#
[root@AMDLoaner-1 /]# /usr/lib/ovs/ovs-makerepo 10.61.166.224:/vol/ovm_nfs_repository 1 nfsrepo
Initializing NEW repository 10.61.166.224:/vol/ovm_nfs_repository
SUCCESS: Mounted /OVS/FB256A26FB284E13A83C0CEC61EB7EE2
Updating local repository list.
ovs-makerepo complete
```

```
[root@ovm22nb /]# /opt/ovs-agent-2.3/utils/repos.py  -n 10.61.166.224:/vol/ovm_nfs_repo
[ NEW ] 474da8ab-3fde-4a8d-91ff-f2c0b91feac2 => 10.61.166.224:/vol/ovm_nfs_repo
```

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
```

4.3 Configuring FC Shared Storage and VM Operations on Oracle VM Server

NetApp Plug-in for OVM

Figure 20 shows an overview of the procedure to configure FC shared storage and VM operations on Oracle VM Server using the NetApp plug-in for OVM. This procedure is applicable to both clustered Data ONTAP and 7-Mode for the NetApp storage controller.

Figure 20) FC storage configuration using NetApp plug-in for OVM.

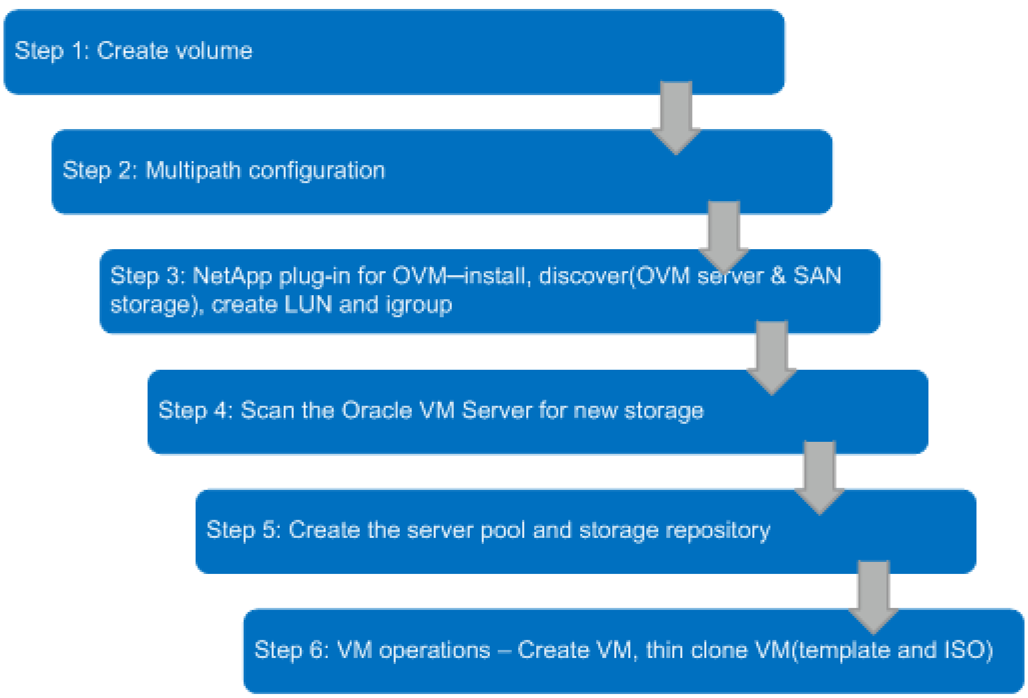

Step 1: Create Volume

1. Navigate to Cluster > Vserver > Storage > Volumes > Create. To create a thin-provisioned volume, select the Thin Provisioned checkbox.

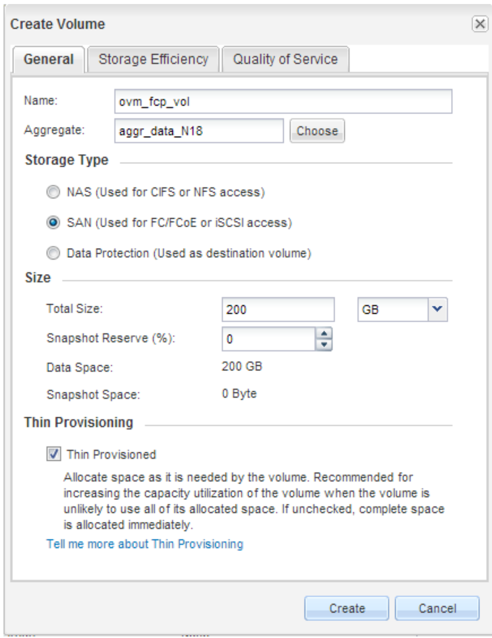

2. After the volume is created, its properties can be further modified. In the Advanced tab, the volume autogrow and fractional reserve settings can be changed. Use the Storage Efficiency tab to enable deduplication.

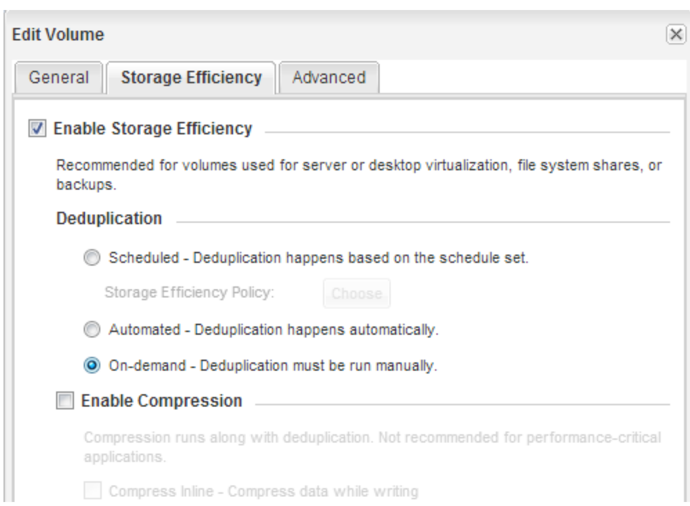

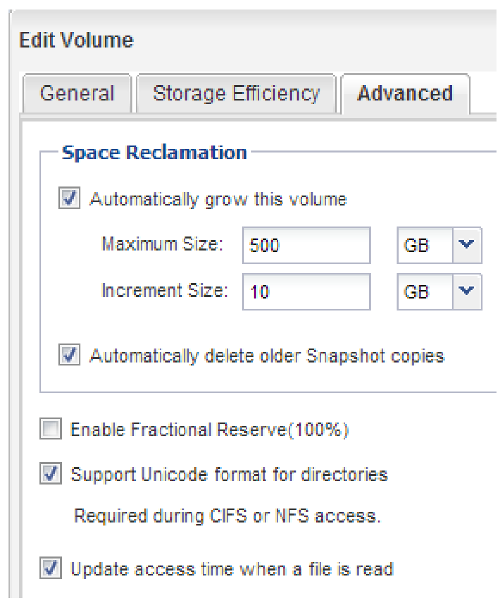

Step: 2: Configure Multipathing in the Oracle VM Server

1. Use the native multipathing—DM-MP (Device Mapper Multipath)—support provided by the Oracle VM Server to configure multipathing.
2. First check if the DM-MP is installed by entering:

```
rpm -q device-mapper
```

3. Check if the DM-MP services are running:

```
[root@stlrx300s6-131 ~]# lsmod | grep dm_
dm_nfs                  7013  2
dm_round_robin          2466  45
dm_multipath           19157  46 dm_round_robin
dm_snapshot            33013  0
dm_zero                 1281  0
dm_mirror              15220  0
dm_region_hash         11129  1 dm_mirror
dm_log                 10076  2 dm_mirror,dm_region_hash
dm_mod                 81987  152 dm_nfs,dm_multipath,dm_snapshot,dm_zero,dm_mirror,dm_log
[root@stlrx300s6-131 ~]#
```

4. Download and install the NetApp [Linux Host Utilities Kit](#).
5. Create or modify the `/etc/multipath.conf` file as described in the NetApp Linux Host Utilities Kit documentation. Check the documentation that has samples for OVM.

   A sample /etc/multipath.conf file for 3.2.1 with Asymmetric Logical Unit Access (ALUA) is shown below:

```
defaults {
user_friendly_names no
max_fds max
flush_on_last_del no
queue_without_daemon no
}
# All data under blacklist must be specific to your system.
blacklist {
devnode "^hd[a-z]"
wwid "<wwid_of_the_local_disk>"
devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
devnode "^cciss.*"
}
devices {
device {
vendor "NETAPP"
product "LUN.*"
path_grouping_policy group_by_prio
features "3 queue_if_no_path pg_init_retries 50"
prio "alua"
path_checker tur
no_path_retry "queue"
failback immediate
hardware_handler "1 alua"
rr_weight uniform
rr_min_io 128
getuid_callout "/lib/udev/scsi_id -gus /block/%n"
}
}
```

The <DevID> refers to the 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 the SCSI ID of each multipath device. This step can be executed after getting the SCSI ID of the devices from step 3. This will enable all multipath devices to 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>
}
}
```

6. Start the multipath service:

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

7. Configure the multipath service:

```
multipath
```

8. Add the multipath service to the boot sequence:

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

9. Verify the multipath configuration:

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

10. As shown in Step 7, 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 3: NetApp Plug-in for OVM—Install, Discover (OVM Server and SAN Storage), Create LUN and Igroup

1. Download the NetApp plug-in from http://community.netapp.com/t5/FAS-Data-ONTAP-and-Related-Plug-ins-Articles-and-Resources/NetApp-Plug-in-For-Oracle-VM/ta-p/87038#Download and follow the NetApp Plug-In 2.0.1 for Oracle VM Installation and Administration Guide for SAN (iSCSI and FCP) deployment.

Note: A video demonstration of how to deploy OVM with the NetApp plug-in with OVM for SAN is available from https://communities.netapp.com/videos/3533.

2. Install the plug-in using “install.sh” from the downloaded tar file.
3. Update /etc/hosts.

```
[root@stlrx200s-45 netapp]# cat /etc/hosts
# Do not remove the following line, or various programs
# that require network functionality will fail.
127.0.0.1               localhost.localdomain localhost
::1                     localhost6.localdomain6 localhost6
10.61.172.69            vsISCSI2
[root@stlrx200s-45 netapp]#
```

4. Discover the OVM Servers from OVM Manager.

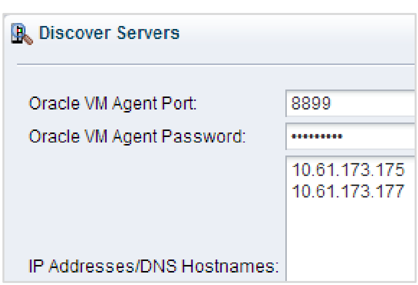

5. Discover the SAN storage controller. Navigate to Home > Storage tab >select SAN Servers > Perspective: (SAN Servers) > and click the button.

Select the NetApp OSC plug-in from the Storage Plug-in drop-down list, select the storage type, provide the storage virtual machine credential for API communication, and move the available OVM Servers to the Selected Server(s) section.

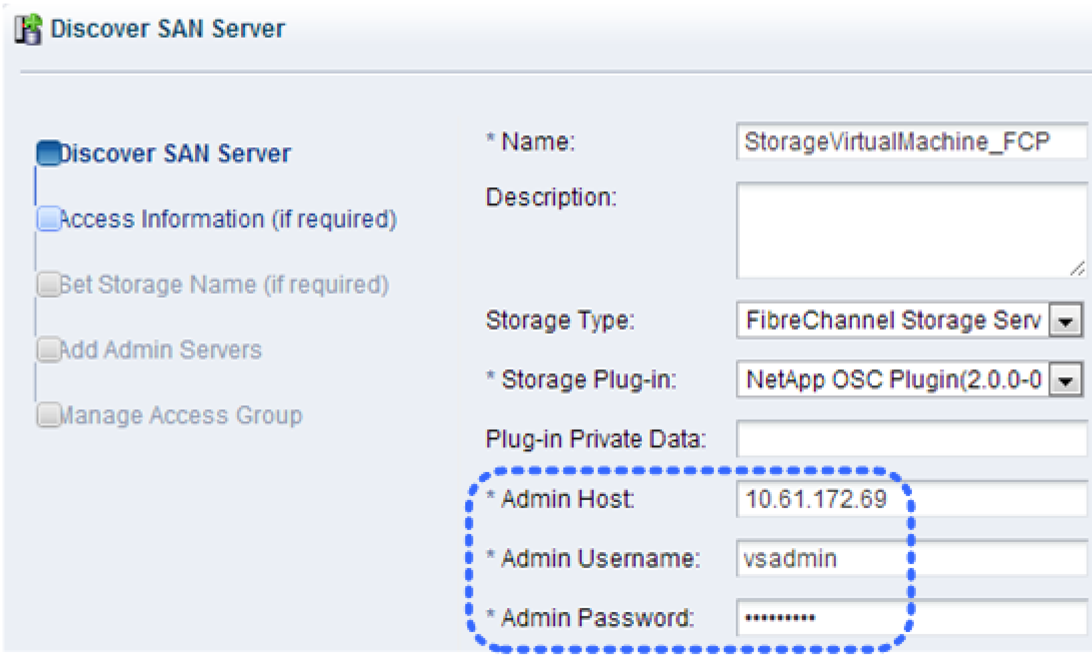

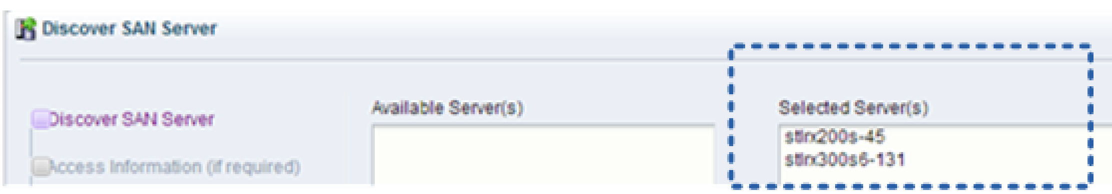

6. Create the physical disk (LUN) for the server pool, repository, and VM disk from OVM Manager. You can also resize and delete the LUN (physical disk). Navigate to Home > Storage > Select the Discovered SAN Server > Perspective :(Physical Disks).

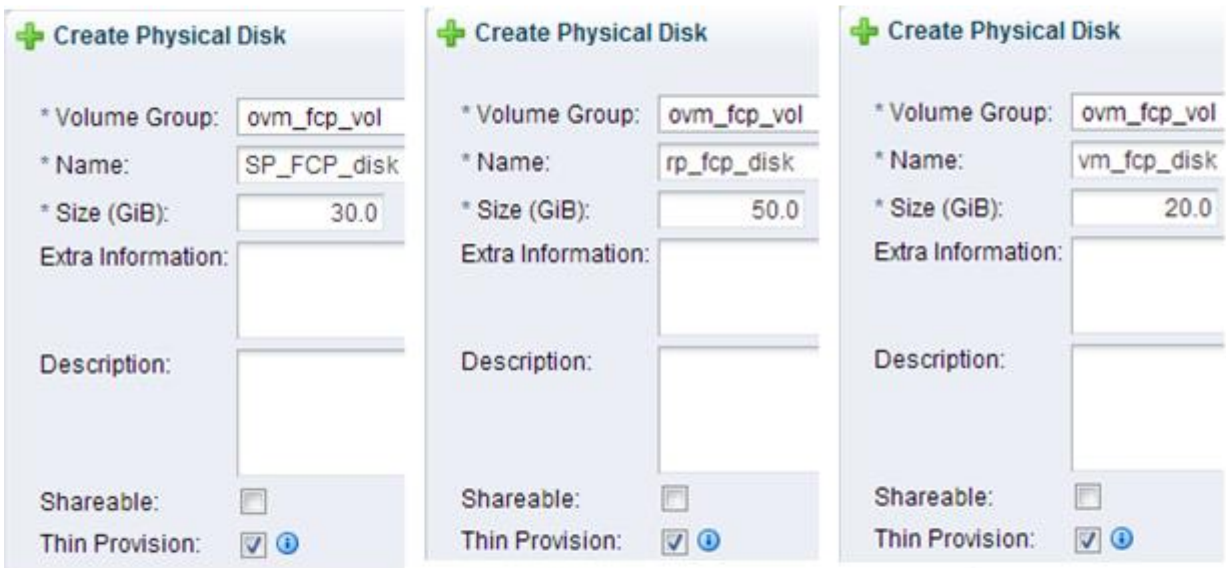

Disk for Server Pool **Disk for Repository** **Disk for Virtual Machine**

7. Create the access group either in the previous screen or create it separately after using Discover the SAN Server.

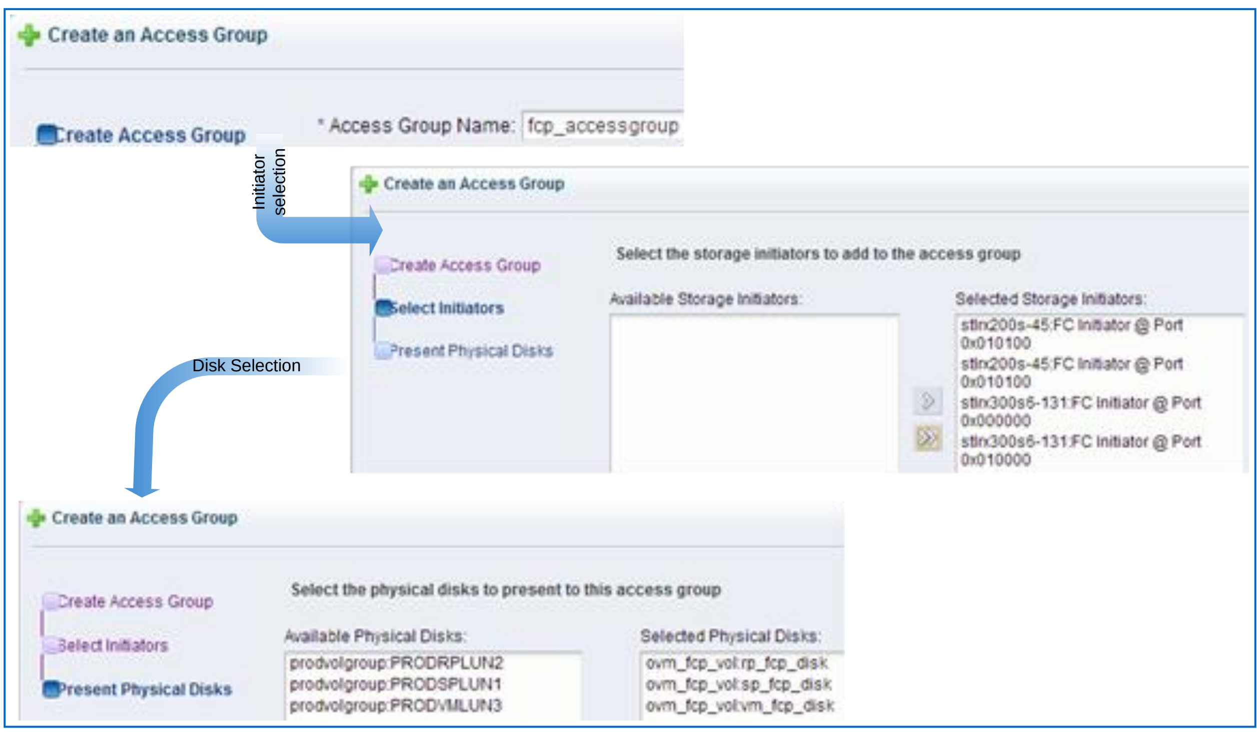

Step 4: Scan the Oracle VM Server for New Storage

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

```
[root@stlrx300s6-131 ~]# rescan-scsi-bus.sh
```

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 user-friendly manner. Similar information is also displayed by the `multipath -ll` command.

```
sanlun lun show
```

A sample output of the `sanlun lun show` command is as follows:

```
[root@stlrx300s6-131 ~]# sanlun lun show
controller(7mode)/                                     device           host                  lun
vserver(Cmode)          lun-pathname                   filename         adapter    protocol   size    mode
------------------------------------------------------------------------------------------------------------
vsISCSI2                /vol/ovm_fcp_vol/rp_fcp_disk /dev/sdg         host5      FCP        50g     C
vsISCSI2                /vol/ovm_fcp_vol/sp_fcp_disk /dev/sdf         host5      FCP        30g     C
vsISCSI2                /vol/ovm_fcp_vol/vm_fcp_disk /dev/sdc         host5      FCP        20g     C
vsISCSI2                /vol/ovm_fcp_vol/vm_fcp_disk /dev/sdb         host5      FCP        20g     C
vsISCSI2                /vol/ovm_fcp_vol/sp_fcp_disk /dev/sdd         host5      FCP        30g     C
vsISCSI2                /vol/ovm_fcp_vol/rp_fcp_disk /dev/sde         host5      FCP        50g     C
[root@stlrx300s6-131 ~]#
```

A sample output of the `multipath -ll` command is as follows:

```
[root@stlrx300s6-131 ~]# sanlun lun show -v -p

                    ONTAP Path: vsISCSI2:/vol/ovm_fcp_vol/vm_fcp_disk
                           LUN: 1
                      LUN Size: 20g
                          Mode: C
                   Host Device: 3600a09804176376f735d425161747765
              Multipath Policy: round-robin 0
                DM-MP Features: 3 queue_if_no_path pg_init_retries 50
              Hardware Handler: 1 alua
            Multipath Provider: Native
----------- --------- ---------- ------- ------------ ---------------------------------------------- -------
            host      vserver            host:
dm-mp       path      path       /dev/   chan:        vserver                                        major:
state       state     type       node    id:lun       LIF                                            minor
----------- --------- ---------- ------- ------------ ---------------------------------------------- -------
active      up        primary    sdb     5:0:2:1      vsSCSI2_fcp3                                   8:16
active      up        secondary  sdc     5:0:5:1      vsISCSI2_fcp1                                  8:32

                    ONTAP Path: vsISCSI2:/vol/ovm_fcp_vol/rp_fcp_disk
                           LUN: 5
                      LUN Size: 50g
                          Mode: C
                   Host Device: 3600a09804176376f735d425161747764
              Multipath Policy: round-robin 0
                DM-MP Features: 3 queue_if_no_path pg_init_retries 50
              Hardware Handler: 1 alua
            Multipath Provider: Native
----------- --------- ---------- ------- ------------ ---------------------------------------------- -------
            host      vserver            host:
dm-mp       path      path       /dev/   chan:        vserver                                        major:
state       state     type       node    id:lun       LIF                                            minor
----------- --------- ---------- ------- ------------ ---------------------------------------------- -------
active      up        primary    sde     5:0:2:5      vsSCSI2_fcp3                                   8:64
active      up        secondary  sdg     5:0:5:5      vsISCSI2_fcp1                                  8:96

                    ONTAP Path: vsISCSI2:/vol/ovm_fcp_vol/sp_fcp_disk
                           LUN: 4
                      LUN Size: 30g
                          Mode: C
                   Host Device: 3600a09804176376f735d425161747763
              Multipath Policy: round-robin 0
                DM-MP Features: 3 queue_if_no_path pg_init_retries 50
              Hardware Handler: 1 alua
            Multipath Provider: Native
----------- --------- ---------- ------- ------------ ---------------------------------------------- -------
            host      vserver            host:
dm-mp       path      path       /dev/   chan:        vserver                                        major:
state       state     type       node    id:lun       LIF                                            minor
----------- --------- ---------- ------- ------------ ---------------------------------------------- -------
active      up        primary    sdd     5:0:2:4      vsSCSI2_fcp3                                   8:48
active      up        secondary  sdf     5:0:5:4      vsISCSI2_fcp1                                  8:80
[root@stlrx300s6-131 ~]#
```

```
[root@stlrx300s6-131 ~]# multipath -ll
3600a09804176376f735d425161747765 dm-3 NETAPP,LUN C-Mode
size=20G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=enabled
| `- 5:0:2:1 sdb 8:16 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  `- 5:0:5:1 sdc 8:32 active ready running
3600a09804176376f735d425161747764 dm-5 NETAPP,LUN C-Mode
size=50G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=enabled
| `- 5:0:2:5 sde 8:64 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  `- 5:0:5:5 sdg 8:96 active ready running
3600a09804176376f735d425161747763 dm-4 NETAPP,LUN C-Mode
size=30G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=enabled
| `- 5:0:2:4 sdd 8:48 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  `- 5:0:5:4 sdf 8:80 active ready running
[root@stlrx300s6-131 ~]# ll /dev/mapper/
total 0
brw-rw---- 1 root disk 252,   4 Jan  9 10:47 3600a09804176376f735d425161747763
brw-rw---- 1 root disk 252,   5 Jan  9 10:48 3600a09804176376f735d425161747764
brw-rw---- 1 root disk 252,   3 Jan  9 10:47 3600a09804176376f735d425161747765
crw------- 1 root root  10, 236 Jan  8 12:45 control
[root@stlrx300s6-131 ~]#
```

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

Step 5: Create the Server Pool and Storage Repository

The multipath device can be used either as an Oracle VM shared server pool and storage repository (OCFS2) or as a standalone storage repository (ext3). While being using as a shared storage repository, OCFS2 first needs to be configured for the Oracle VM Server nodes in the server pool. Let's create the server pool using the NetApp plug-in, which will create the server pool and repository in an OCFS file system and be shared among the OVM Servers.

1. Create the server pool using the NetApp plug-in for OVM-SAN. Navigate to Home > Server and VMs > click the  button.
2. Set the following options:
   - Enable secure VM migrate.
   - Choose physical disk and select newly created disk for server pool.
   - Set 190 sec for Timeout for cluster (NetApp's recommendation based on cluster failover and takeover).
   - Move the OVM Servers to the Selected Server(s) section.

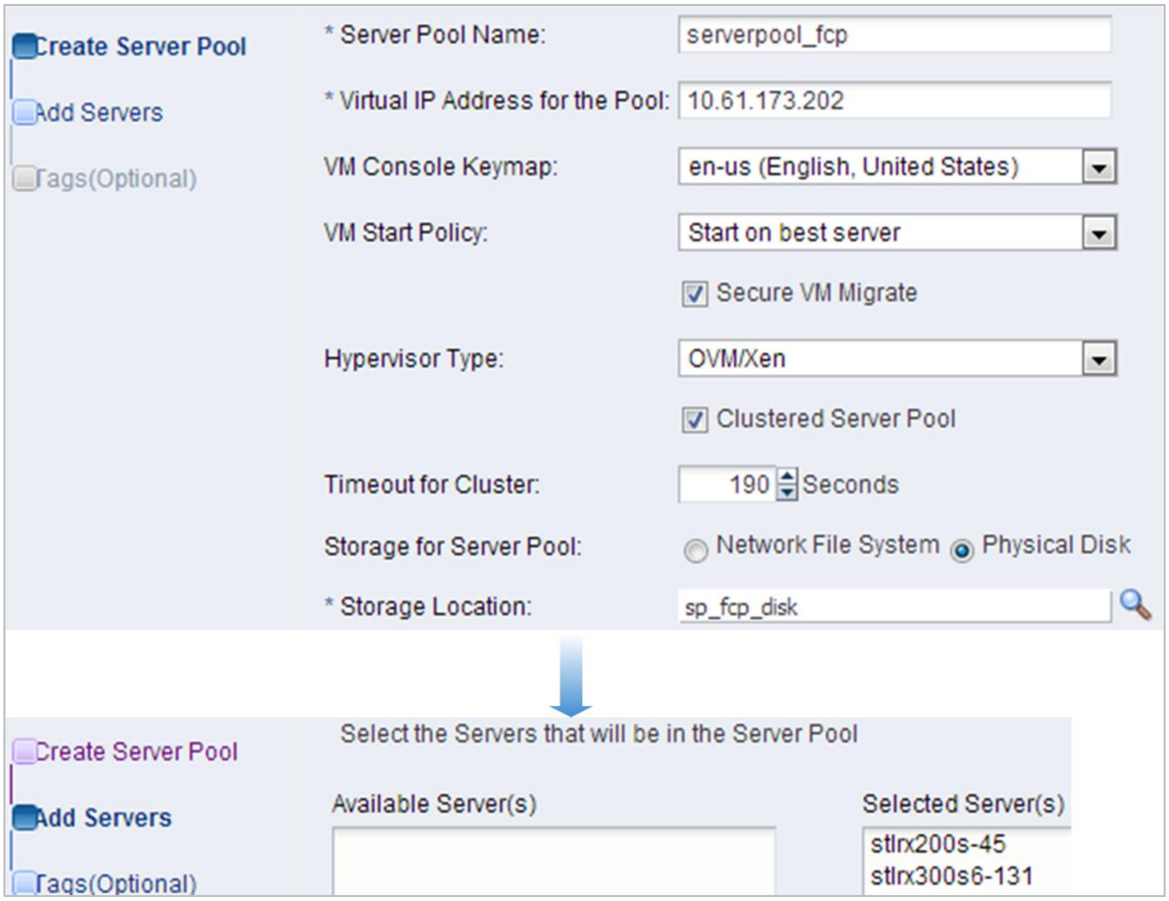

1. NetApp recommends the following values (applicable for both 7-Mode and clustered Data ONTAP storage controllers):

```
O2CB_HEARTBEAT_THRESHOLD=96, O2CB_IDLE_TIMEOUT_MS=190000, O2CB_KEEPALIVE_DELAY_MS=4000,
O2CB_RECONNECT_DELAY_MS=4000
```

2. Update `/etc/sysconfig/o2cb` with the following and reboot the OVM Server or restart the ocfs2 service:

```
[root@stlrx300s6-131 ~]# cat /etc/sysconfig/o2cb
O2CB_HEARTBEAT_THRESHOLD=96
O2CB_RECONNECT_DELAY_MS=4000
O2CB_KEEPALIVE_DELAY_MS=4000
O2CB_BOOTCLUSTER=34a289117ee06652
O2CB_IDLE_TIMEOUT_MS=190000
O2CB_ENABLED=true
O2CB_STACK=o2cb
[root@stlrx300s6-131 ~]#
```

Note: Setting `O2CB_HEARTBEAT_THRESHOLD` to 65 is sufficient. However, since in an active-active cluster cf giveback sometimes reboots the OVM Server, NetApp recommends setting it from 60 to 120 seconds. Our lab setup used 96 and `O2CB_IDLE_TIMEOUT_MS` was set to `190000`.

3. The plug-in creates the OCFS file system and mounts it on the OVM Server for server pools.

4. A sample `/etc/ocfs2/cluster.conf` file that needs to be present in each node of the server pool for the OCFS2 configuration may appear as follows:

```
[root@stlrx300s6-131 ~]# cat /etc/ocfs2/cluster.conf
heartbeat:
        region = 0004FB0000050000BDE2BB49F8C8A2CA
        cluster = 34a289117ee06652

node:
        ip_port = 7777
        ip_address = 10.61.173.177
        number = 0
        name = stlrx200s-45
        cluster = 34a289117ee06652

node:
        ip_port = 7777
        ip_address = 10.61.173.175
        number = 1
        name = stlrx300s6-131
        cluster = 34a289117ee06652

cluster:
        node_count = 2
        heartbeat_mode = global
        name = 34a289117ee06652

[root@stlrx300s6-131 ~]#
```

5. Create a storage repository using the NetApp plug-in for OVM–SAN.

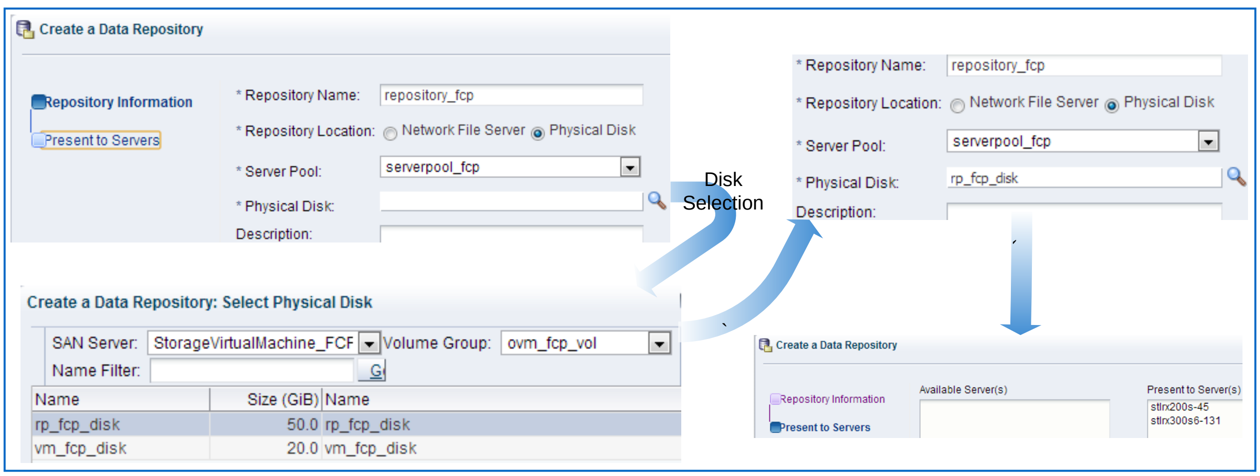

6. Check the server pool and repository in the OVM Server:

```
[root@stlrx300s6-131 ~]# df -h
Filesystem            Size  Used Avail Use% Mounted on
/dev/sda4              48G  998M   44G   3% /
tmpfs                 330M     0  330M   0% /dev/shm
10.60.132.21:/vol/rtpdist/swdist
                      5.0T  4.1T  937G  82% /swdist
none                  330M   40K  330M   1% /var/lib/xenstored
/dev/mapper/3600a09804176376f735d425161747763
                       30G  264M   30G   1% /poolfsmnt/0004fb0000050000bde2bb49f8c8a2ca
/dev/mapper/3600a09804176376f735d425161747764
                       50G  4.3G   46G   9% /OVS/Repositories/0004fb0000030000b266f3a97befe0af
[root@stlrx300s6-131 ~]#
```

7. To manually configure, create, and start the OCFS2 cluster service:

```
service o2cb status
service o2cb load
```

```
service o2cb online
service o2cb start
service o2cb configure
```

8. Now the multipath device can be formatted and mounted as an Oracle VM Server repository:

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

For OVM v3.2.x:

a. Server pool

```
ssh -l admin 10.61.173.178 -p 10000
OVM> create ServerPool ….
```

b. Repository

```
ssh -l admin 10.61.173.178 -p 10000
OVM> create repository….
```

Step 6: VM Operations—Create VM, Thin-Clone VM

Create a VM Using the Template

To create the VM using the template, complete the following steps.

1. Download the ova files from https://edelivery.oracle.com/linux and copy them to the HTTP/FTP server.
2. Import the assembly by navigating to Home > Repositories > Assembles > Import & Home > Repositories > Assembles.
3. Right-click the .ova file and click Create VM Template.

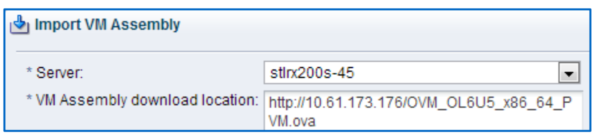

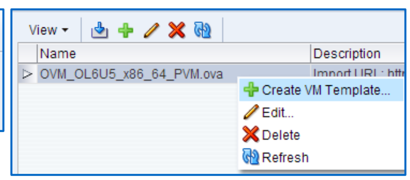

4. Create a VM or clone a VM from a template.

   Select the OVM Navigate to Home > Servers and VMs> Server Pools > select OVM server.

   Add networking. Go to Home > Networking>Virtual NICs> Auto Fill> and click the create  button.

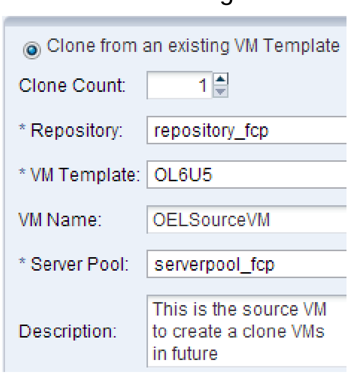

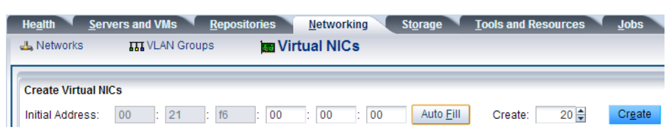

5. Create a clone customizer to create a thin-cloned VM using the NetApp plug-in. Navigate to Home > Servers and VMs and click the  button.

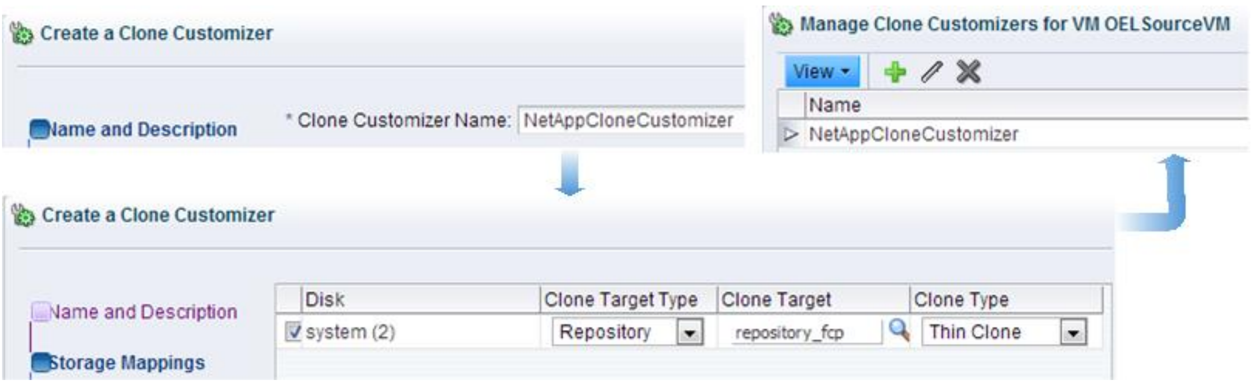

6. Check the repository storage usage before cloning the VM:

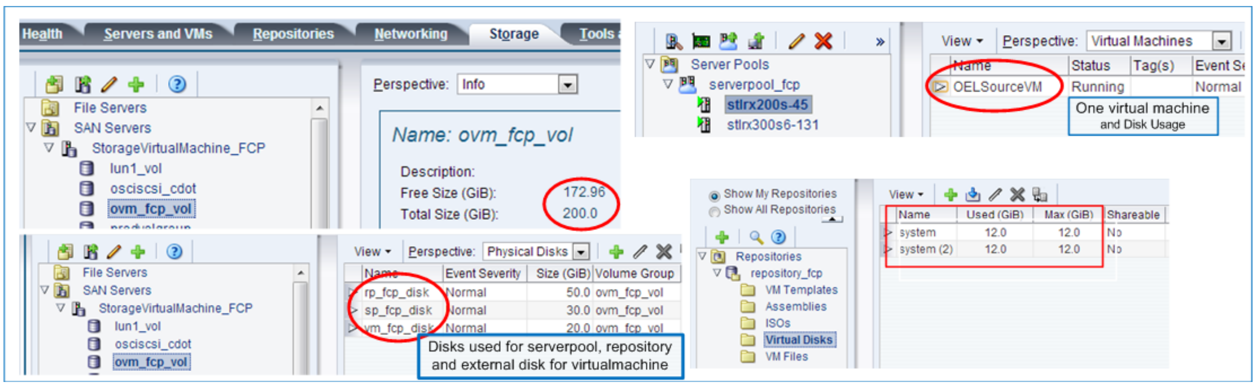

7. Create 10 clone VMs from the source VM:

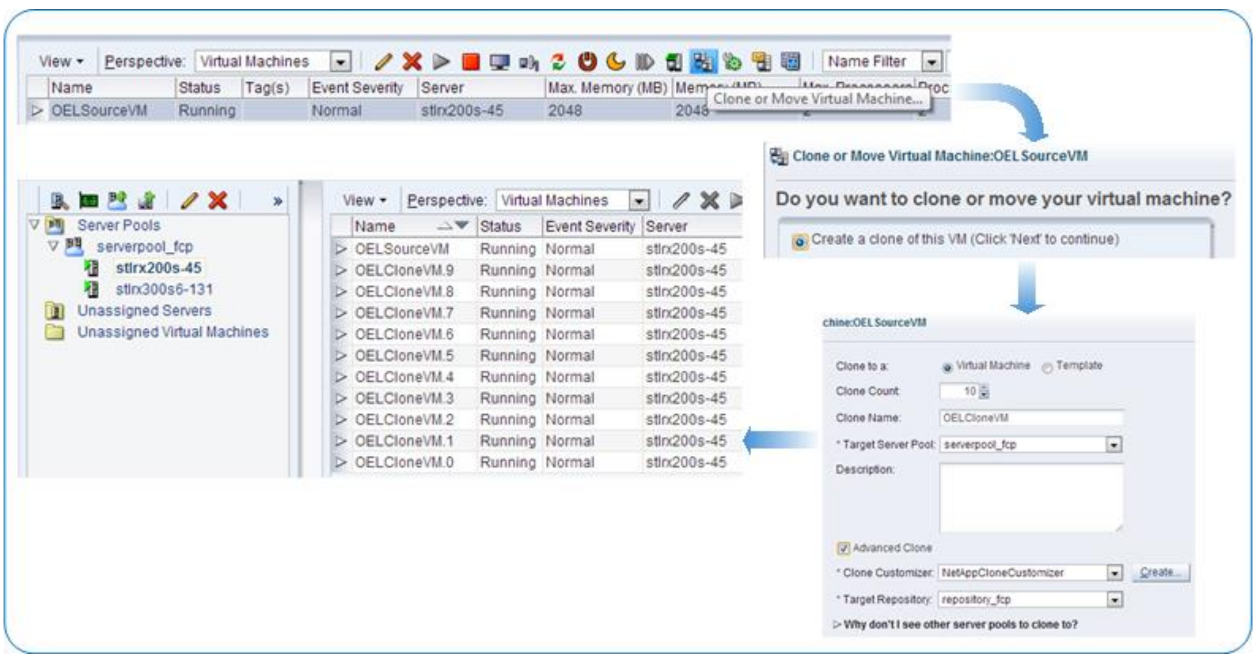

8. Check the storage space after the cloned VMs' creation:

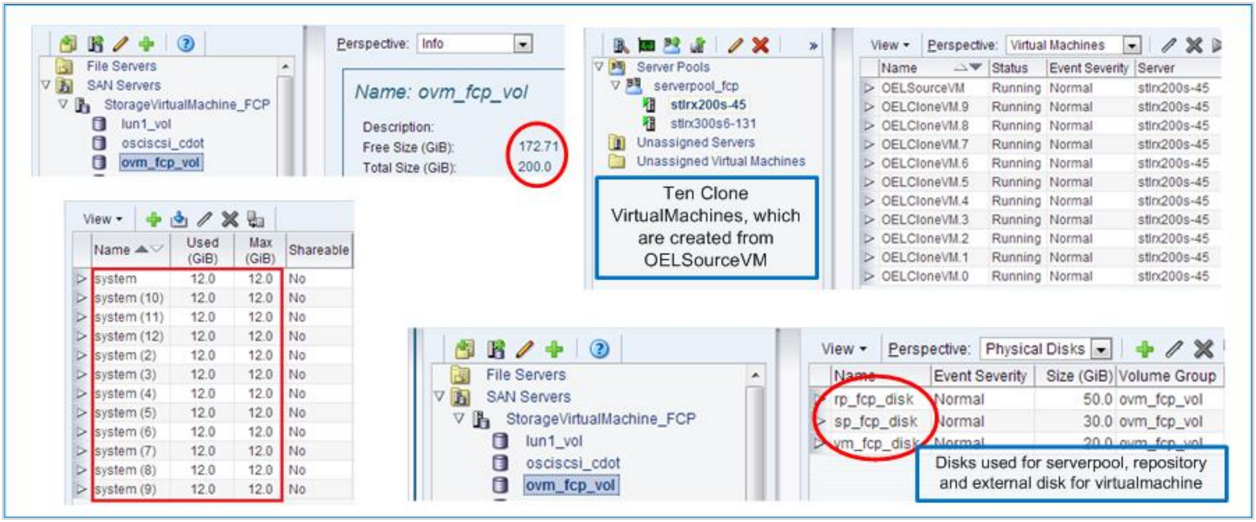

The free space in the ovm_fcp_vol volume is 172.96GB before the thin-clone VM creation; the free space is 172.71GB after thin-clone VM creation. Each virtual machine uses 12GB, but the 10 virtual machines do not occupy additional space except for the source virtual machine (OEKSourceVM) and for template image space usage.

Create a VM Using ISO

To create the VM using ISO, complete the following steps.

1. Download the ISO from https://edelivery.oracle.com/linux and copy to HTTP/FTP server.
2. Import the ISO by navigating to Home > repositories > ISO > import.
3. Create a VM from an ISO. Navigate to Home > Servers and VMs > Server Pools > select OVM server and click the button.
4. Select Create Virtual Machine.
5. Enter the VM name and choose an operating system, domain type, memory, and CPU.
6. Add networking. Navigate to Home > Networking > Virtual NICs > Auto Fill > create.

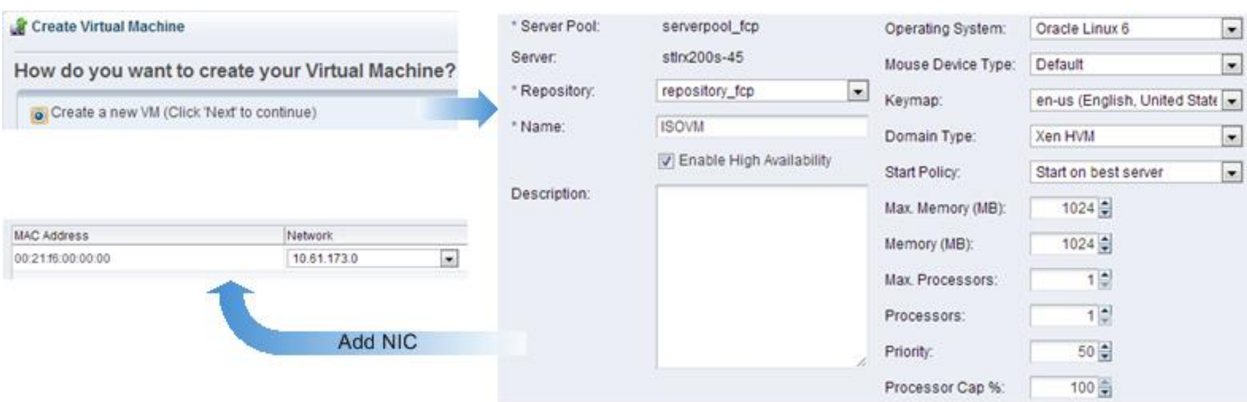

7. Arrange disks for the source VM.
8. For the boot order sequence, make sure to set the boot order to boot from the disk first.

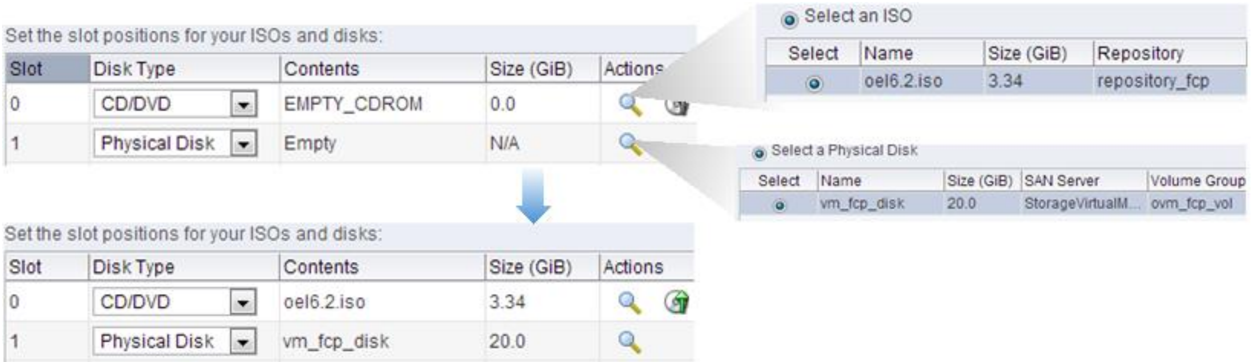

9. Create a clone customizer to create a thin-cloned VM using the NetApp plug-in. Navigate to Home > Servers and VMs and click the button.

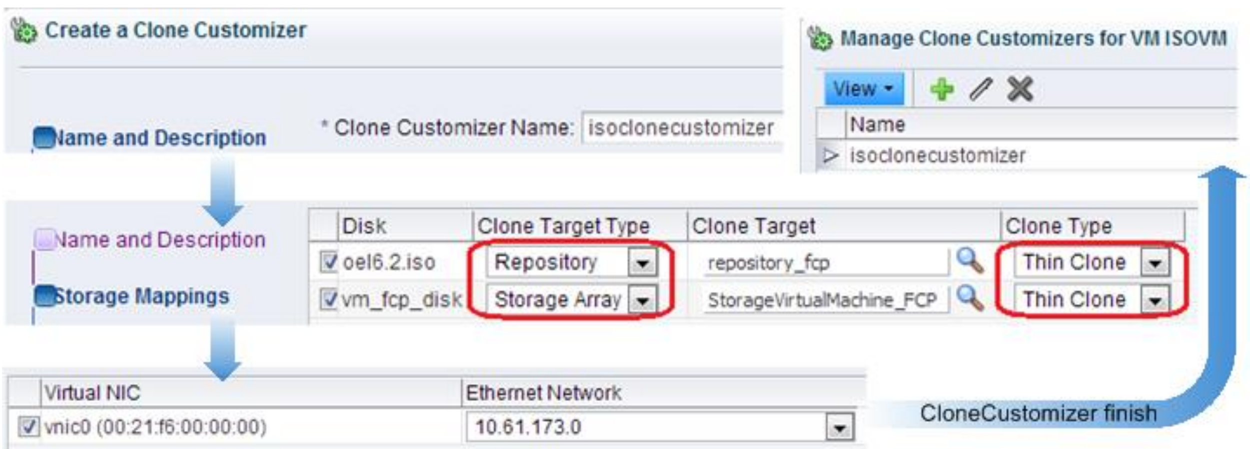

10. Execute the following as you did in the previous section.
    a. Check the repository storage usage before cloning the VM.
    b. Create 10 cloned VMs from the source VM.
    c. Check the storage space after cloning the VMs.

Oracle Generic NFS Plug-In

To use the Oracle generic NFS plug-in, refer to
http://docs.oracle.com/cd/E35328_01/E35332/html/index.html.

4.4 Configuring iSCSI Shared Storage on Oracle VM Server

Figure 21 shows an overview of the procedure to configure iSCSI shared storage on the Oracle VM Server.

Note: Follow the same procedure as described in section 4.3, “Configuring FC Shared Storage and VM Operations on Oracle VM Server.”

Figure 21) Configuring iSCSI shared storage on Oracle VM Server.

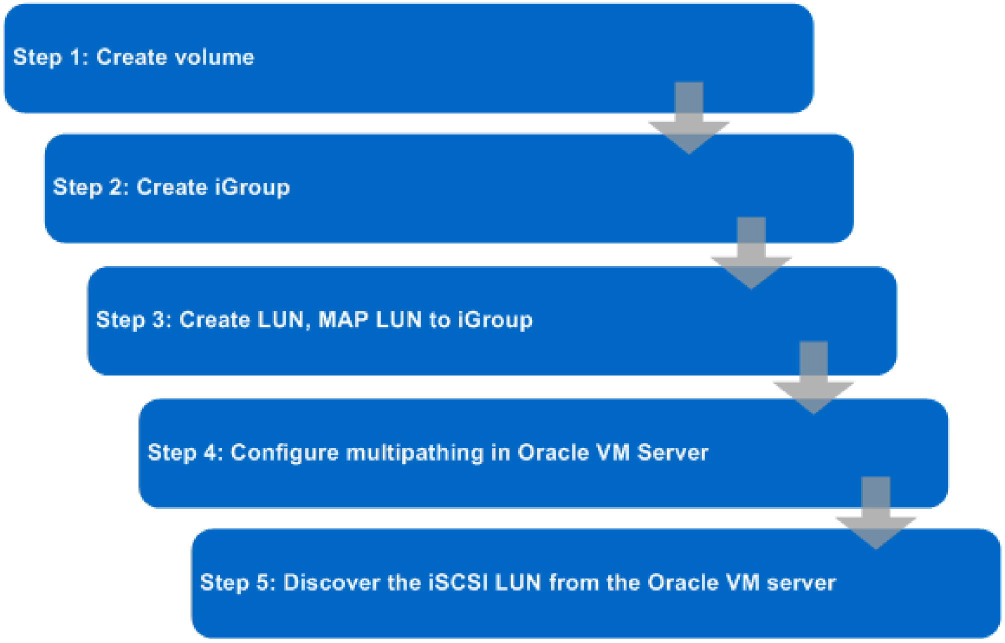

Step 1: Create Volume

This is the same as in step 1 in section 4.3, “Configuring FC Shared Storage and VM Operations 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.

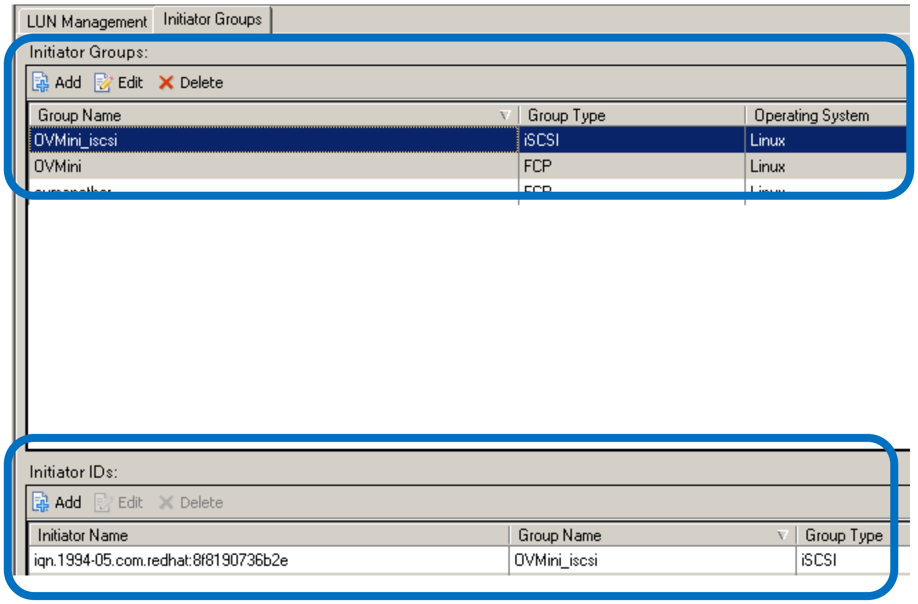

Step 3: Create LUN, Map Igroup

1. From the LUN Wizard in 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.

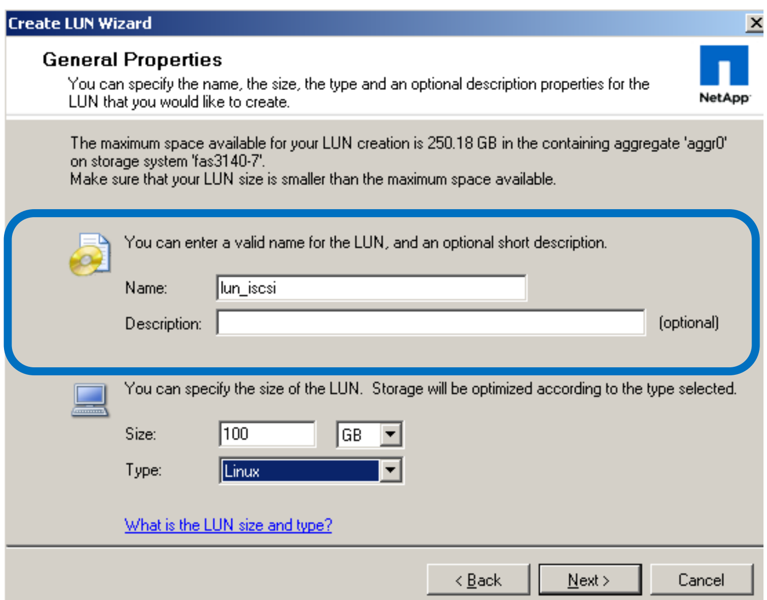

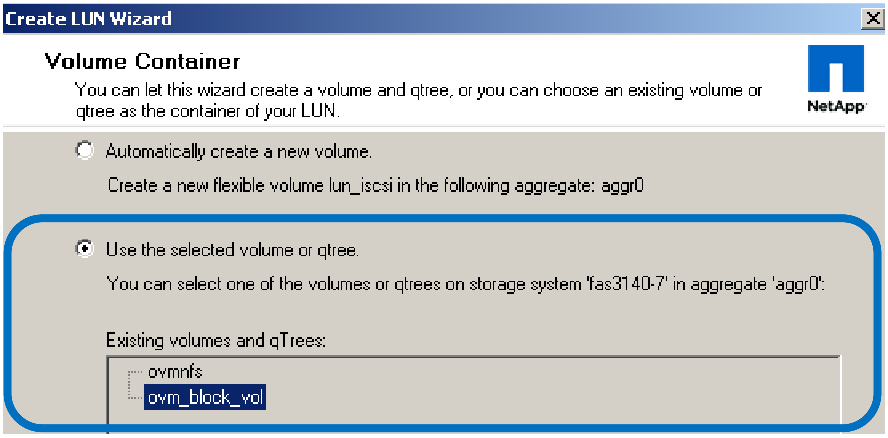

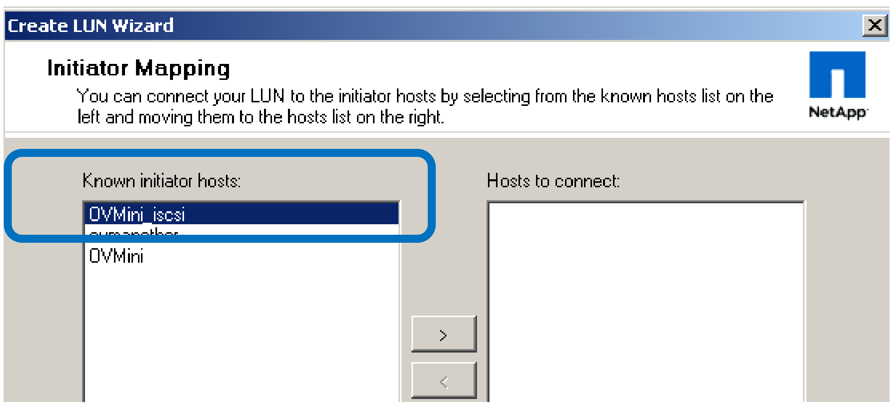

2. Complete the steps in the LUN Wizard to finish the LUN creation process.

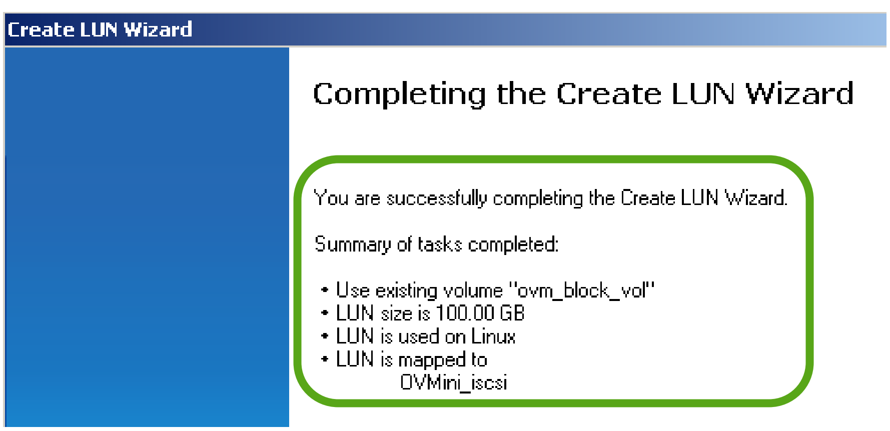

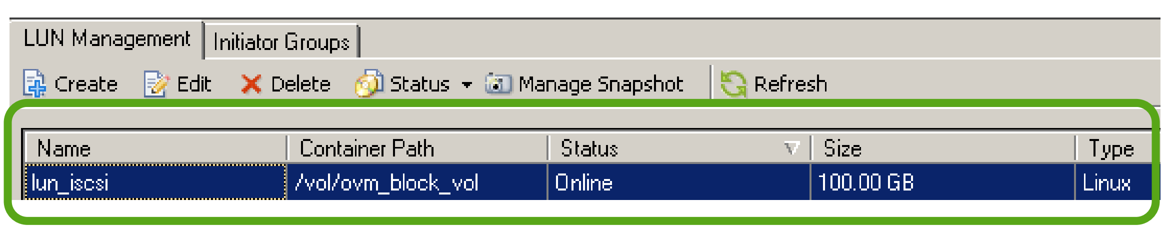

3. If you use thin provisioning, clear the Space Reserved checkbox in LUN Properties.

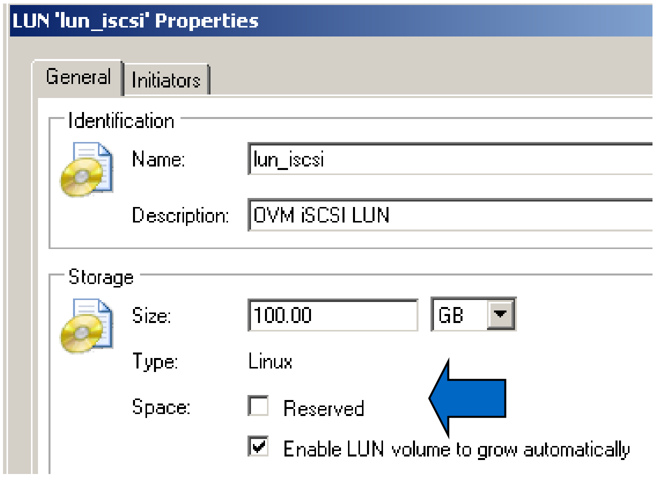

Step 4: Configure Multipathing on the Oracle VM Server

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
  fas3140-7:  /vol/ovm_block_vol/lun_iscsi  /dev/sdj          host12   iSCSI        100g (107374182400)    GOOD
```

3. After 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 as that mentioned for the FC LUN in the previous section.

4.5 NetApp Deduplication in an Oracle VM Environment

Deduplication in Clustered Data ONTAP

1. Ten virtual machines and one template reside in the /OVS/Repositories folder.
2. Check the space consumed from the OVM Server.

```
[root@stlrx300s6-131 VirtualMachines]# df -h
Filesystem            Size  Used Avail Use% Mounted on
/dev/sda4              48G  985M   44G   3% /
tmpfs                 330M     0  330M   0% /dev/shm
10.60.132.21:/vol/rtpdist/swdist
                      5.0T  3.8T  1.2T  77% /swdist
none                  330M   88K  330M   1% /var/lib/xenstored
10.63.164.18:/ovm_nfs_serverpool
                       57G  140M   57G   1% /nfsmnt/5b474bd3-da9a-4157-8b72-7a7c9ad28fc3
/dev/mapper/ovspoolfs
                       10G  369M  9.7G   4% /poolfsmnt/0004fb000005000037c9216df83d74cc
10.63.164.18:/ovm_nfs_repository
                       76G   72G  4.4G  95% /OVS/Repositories/0004fb0000030000217d0559cfc41604
[root@stlrx300s6-131 VirtualMachines]#
```

The space consumed can also be observed from NetApp System Manager.

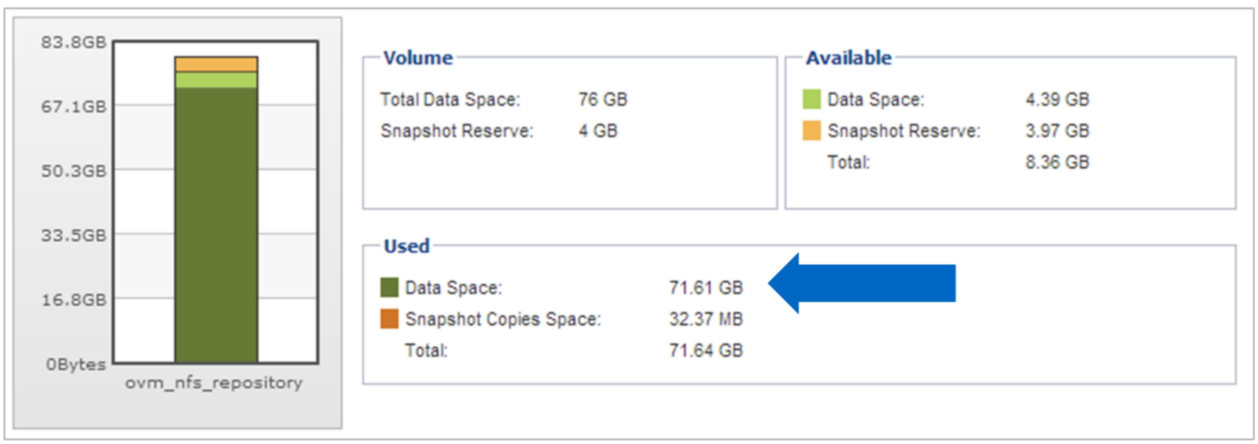

3. Enable deduplication by using one of the following methods.
   – From the CLI, run the following command:

```
cluster1::> volume efficiency on -vserver <vservername> -volume <volname>
```

   – In OnCommand System Manager, enable dedupe in the Storage Efficiency tab.

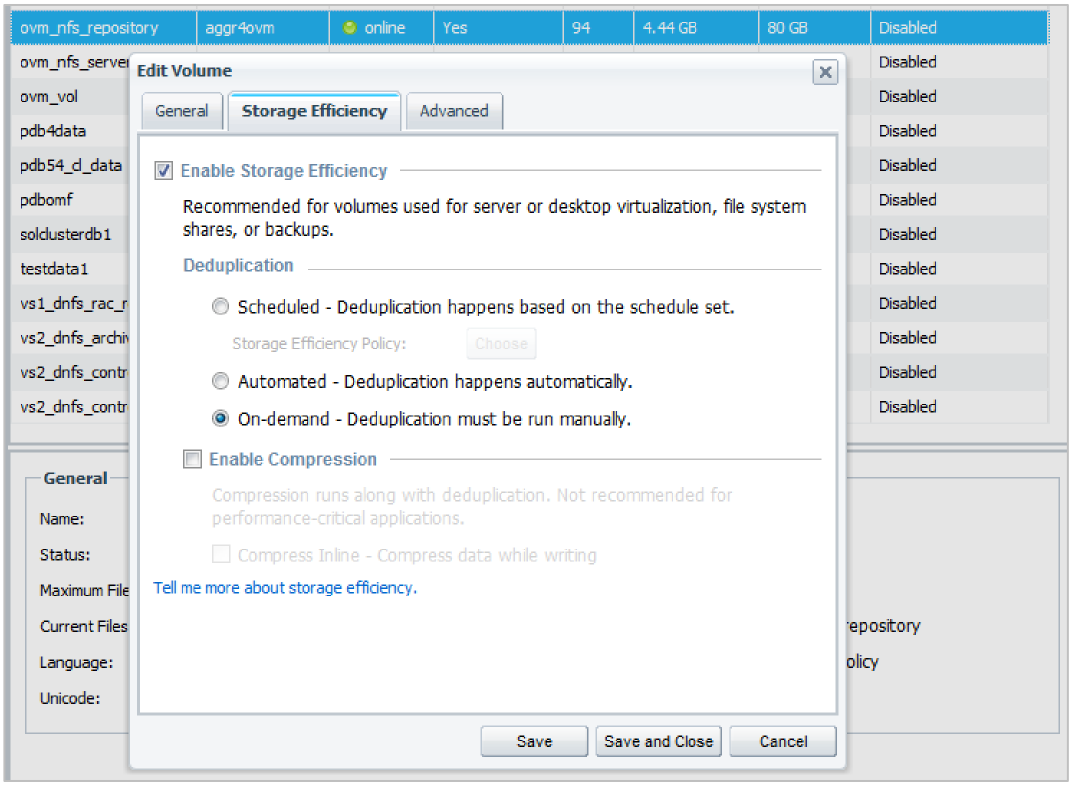

4. After deduplication is enabled on the volume, the deduplication process can be started manually or automatically or it can be scheduled in one of the following ways.
   – From the CLI, run the following command:

```
cluster1::> volume efficiency start -vserver <vservername> -volume <volname>
```

   – From System Manager:

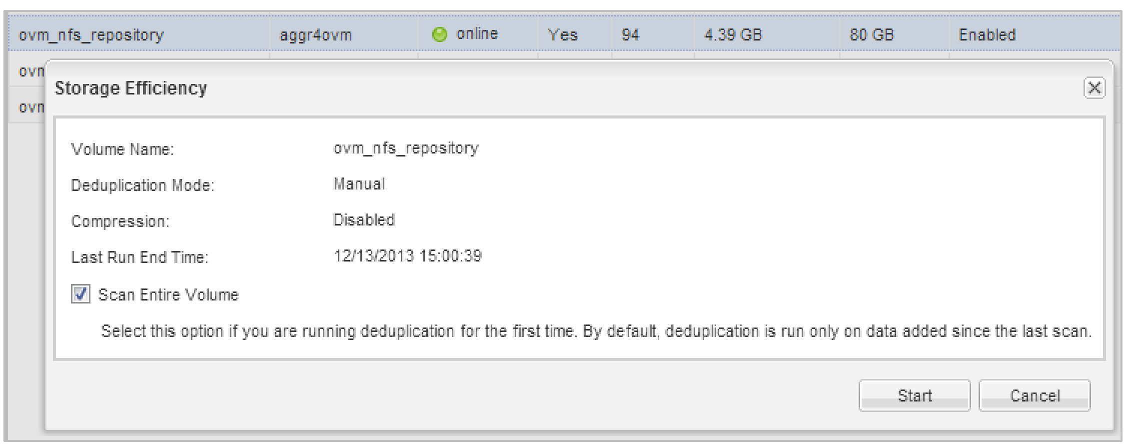

5. Check the progress:

```
TESO::> volume efficiency show -vserver vs2_dnfs_rac -volume ovm_nfs_repository

                        Vserver Name: vs2_dnfs_rac
                         Volume Name: ovm_nfs_repository
                         Volume Path: /vol/ovm_nfs_repository
                               State: Enabled
                              Status: Active
                            Progress: 8536140 KB (12%) Done
                                Type: Regular
                            Schedule: -
              Efficiency Policy Name: -
               Blocks Skipped Sharing: 0
                Last Operation State: Success
      Last Success Operation Begin: Fri Dec 13 15:00:39 EST 2013
        Last Success Operation End: Fri Dec 13 15:00:39 EST 2013
                Last Operation Begin: Fri Dec 13 15:00:39 EST 2013
                  Last Operation End: Fri Dec 13 15:00:39 EST 2013
                 Last Operation Size: 0B
                Last Operation Error: -
                     Changelog Usage: 0%
                   Logical Data Size: 71.60GB
                  Logical Data Limit: 640TB
                Logical Data Percent: 0%
                          Queued Job: -
      Stale Fingerprint Percentage: 0
                         Compression: false
                  Inline Compression: false
    Incompressible Data Detection: false
                  Constituent Volume: false
Compression Quick Check File Size: 524288000

TESO::>
```

6. After completing the deduplication process, verify the space consumed.

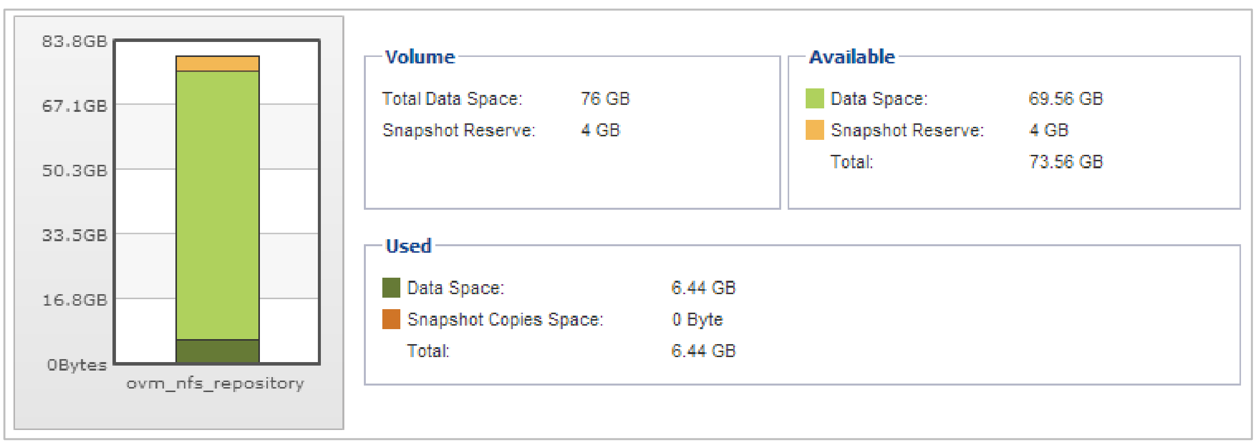

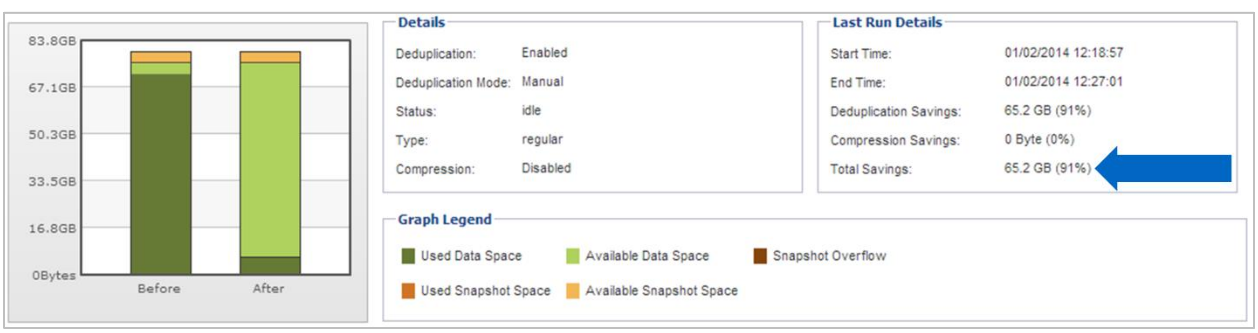

Deduplication in 7-Mode

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 was 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 NetApp System Manager or FilerView or from the NetApp FAS system command line.

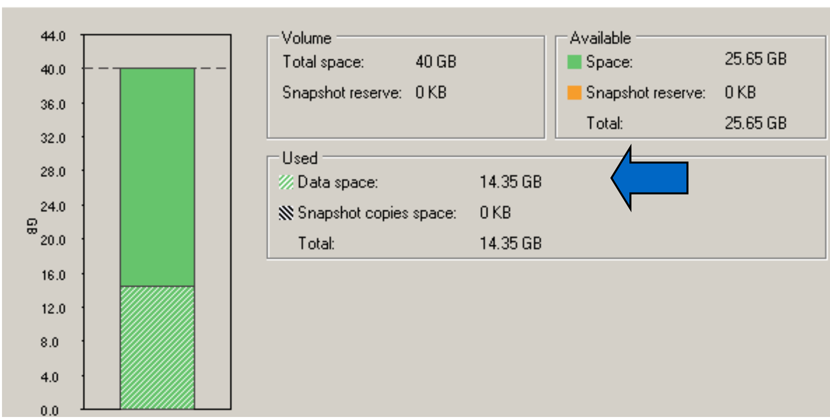

4. Enable deduplication on the volume from the NetApp FAS CLI using the `sis on` command or from NetApp System Manager.

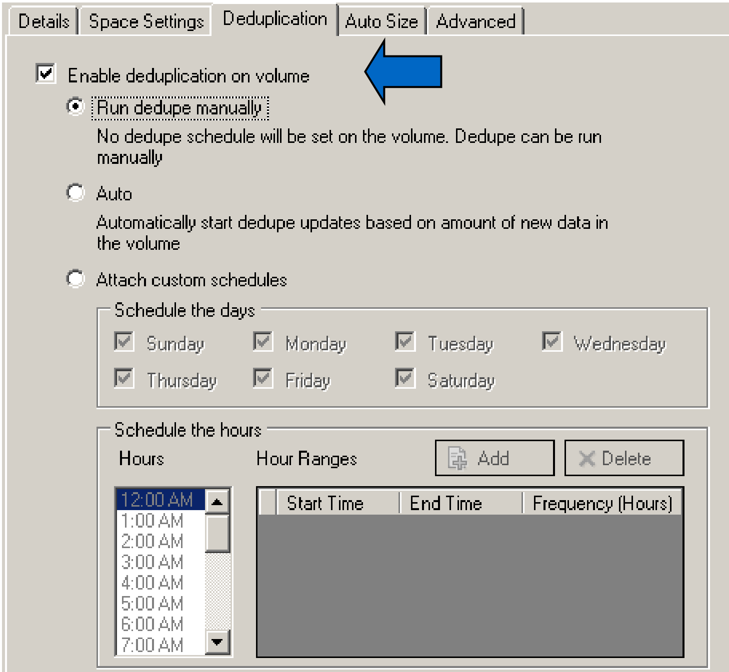

5. Once deduplication is enabled on the volume, the deduplication process can be started manually or automatically or it can be scheduled either from the NetApp FAS CLI using the `sis start` command or from NetApp System Manager.

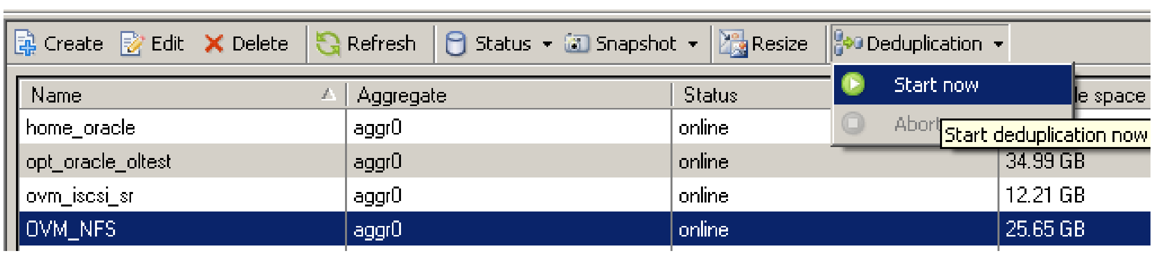

6. After the deduplication process has been completed, check the space consumed by the volume and the storage savings realized.

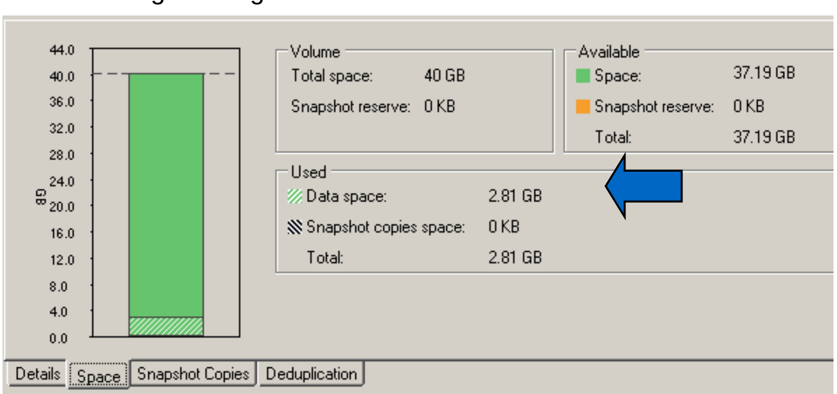

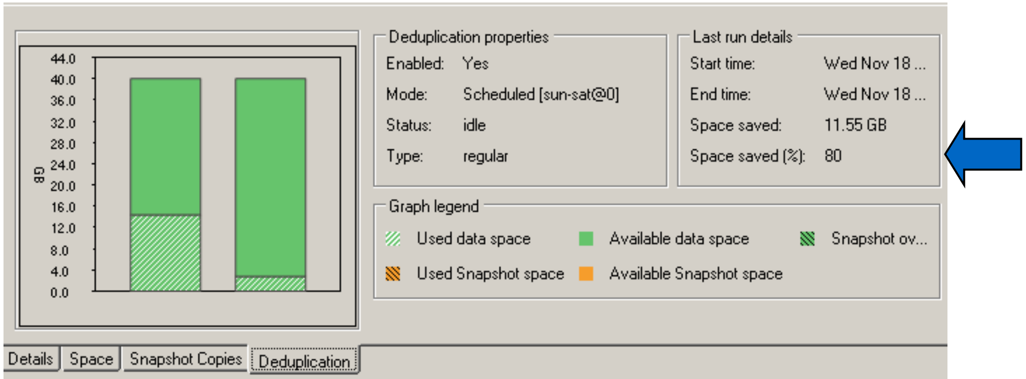

4.6 NetApp Virtual Cloning (FlexClone and SIS Clone) in an Oracle VM Environment

Clustered Data ONTAP and OVM 3.x (NFS)

The following procedure describes NetApp thin cloning (file), which corresponds to the process illustrated below for NFS. The (sis clone) file cloning helps to clone the VMs.

Figure 22) Thin cloning.

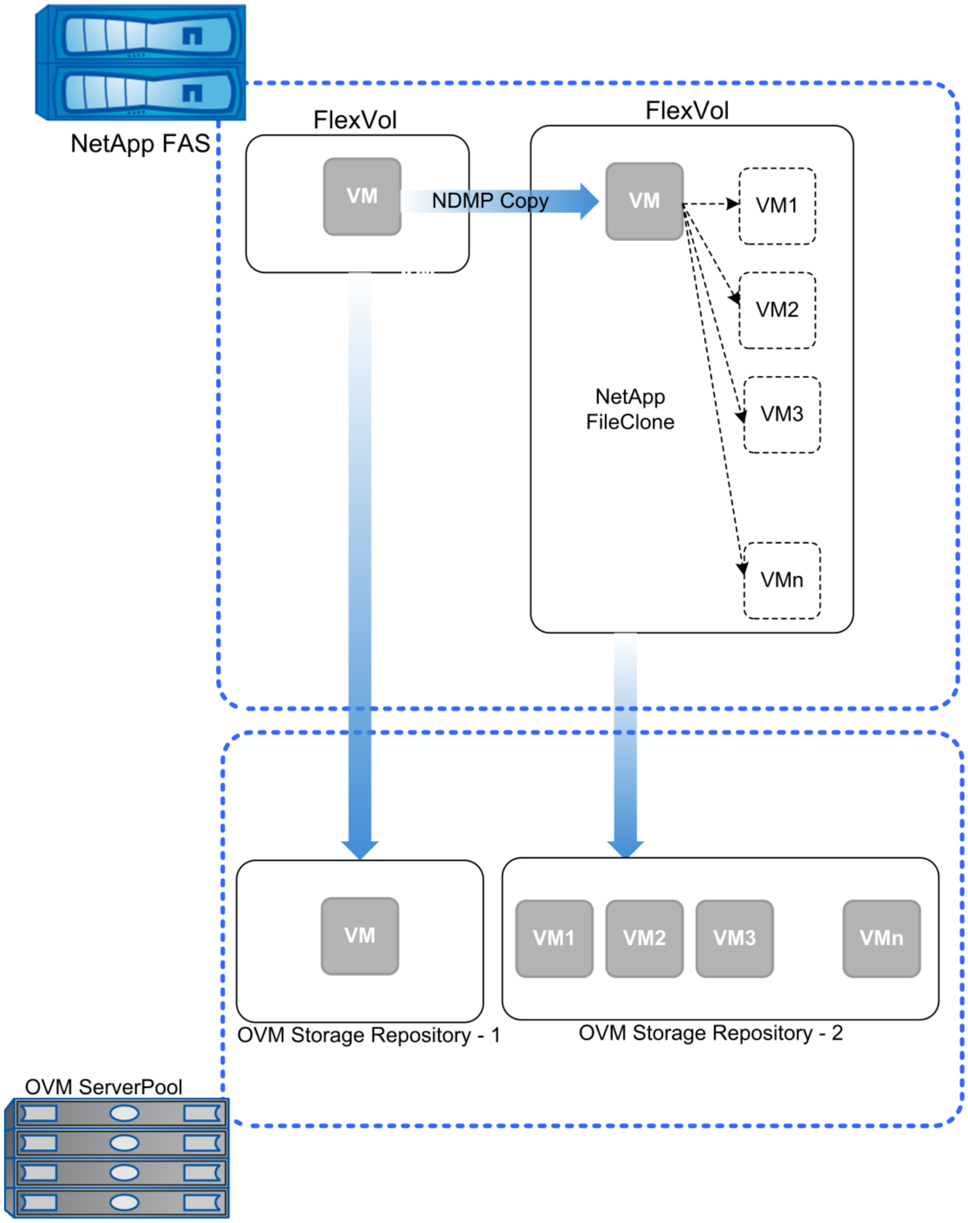

1. The source VM resides inside an already existing OVM repository.

Note: Connect to the "OVM>" prompt using `ssh -l admin localhost -p 10000` in the OVM Manager Server.

```
OVM> list repository
Command: list repository
Status: Success
Time: 2014-01-03 16:39:40,519 EST
Data:
  id:0004fb0000030000c6cfa8ad15f46620  name:NetApp_CDOT_NFS_Repository
OVM>
OVM> list vm
Command: list vm
Status: Success
```

```
Time: 2014-01-03 16:41:22,541 EST
Data:
  id:0004fb000006000051096a9e0c9f3360  name:TemplateVM
  id:0004fb000014000093285ef818aa259e  name:OVM_EL5U5_X86_64_PVM_10GB.tgz
OVM> show vm name=TemplateVM
Command: show vm name=TemplateVM
Status: Success
Time: 2014-01-03 16:41:47,030 EST
Data:
  Name = TemplateVM
  Id = 0004fb000006000051096a9e0c9f3360
  Status = Running
  Memory (MB) = 2048
  Max. Memory (MB) = 2048
  Max. Processors = 2
  Processors = 2
  Priority = 50
  Processor Cap = 100
  High Availability = No
  Operating System = None
  Mouse Type = Default
  Domain Type = Xen PVM
  Keymap = en-us
  description = Import URLs: [http://10.61.173.176/OVM_EL5U5_X86_64_PVM_10GB.tgz]
  Server = 00:00:00:00:00:00:00:00:00:00:26:2d:04:d2:8e  [stlrx200s-45]
  Repository = 0004fb0000030000c6cfa8ad15f46620  [NetApp_CDOT_NFS_Repository]
  Vnic 1 = 0004fb000007000020ebf1da76d3787d  [00:21:f6:00:00:08]
  VmDiskMapping 1 = 0004fb00001300001276b635e4def750
OVM>
```

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

```
OVM> list repository
Command: list repository
Status: Success
Time: 2014-01-03 16:46:23,712 EST
Data:
  id:0004fb0000030000c6cfa8ad15f46620  name:NetApp_CDOT_NFS_Repository
  id:0004fb0000030000a8cd28901056dbb2  name:NetApp_CDOT_Dev_Test_Repository
OVM>

[root@stlrx300s6-131 ~]# df -h
Filesystem            Size  Used Avail Use% Mounted on
/dev/sda4              48G  987M   44G   3% /
tmpfs                 330M     0  330M   0% /dev/shm
10.60.132.21:/vol/rtpdist/swdist
                      5.0T  4.0T  961G  81% /swdist
10.63.164.18:/ovm_nfs_serverpool
                       57G  141M   57G   1% /nfsmnt/5b474bd3-da9a-4157-8b72-7a7c9ad28fc3
10.63.164.18:/ovm_nfs_repository
                       76G  8.6G   68G  12% /OVS/Repositories/0004fb0000030000c6cfa8ad15f46620
none                  330M   80K  330M   1% /var/lib/xenstored
/dev/mapper/ovspoolfs
                       10G  263M  9.8G   3% /poolfsmnt/0004fb000005000008686e2a728bbc1c
10.63.164.18:/ovm_nfs_repository_dev_test
                       95G  6.2G   89G   7% /OVS/Repositories/0004fb0000030000a8cd28901056dbb2
[root@stlrx300s6-131 ~]#
```

Note: NetApp_CDOT_Dev_Test_Repository is created in ovm_nfs_repository_dev_test volume and NetApp_CDOT_NFS_Repository is created in ovm_nfs_repository.

3. In the example below, four clone VMs are created from one source VM.
    a. First, from the NetApp FAS system console, use `ndmpcopy` to copy the images of the source VM to the directory created for the golden VM (in the example, the name of the VM is “TemplateVM”).

```
ndmpcopy -d -sa <sourcefiler user>:<password> -da <destfiler user>:<password> -st text -dt text
<source vserver ip>:/<source vserver name>/<volumename>/”.img file location” <dest vserver
ip>:/<dest vserver name>/<volumename>/<VirtualDisk folder in repository>
```

b. To find the .img location, view the vm.cfg file of the “TemplateVM” VM.

```
[root@stlrx300s6-131 ~]# cat
/OVS/Repositories/0004fb0000030000c6cfa8ad15f46620/VirtualMachines/0004fb000006000051096a9e0c9f33
60/vm.cfg
vif = ['mac=00:21:f6:00:00:08,bridge=0a3dad00']
OVM_simple_name = 'TemplateVM'
disk =
['file:/OVS/Repositories/0004fb0000030000c6cfa8ad15f46620/VirtualDisks/0004fb000012000037d313718b
df0da2.img,xvda,w']
bootargs = ''
uuid = '0004fb00-0006-0000-5109-6a9e0c9f3360'
on_reboot = 'restart'
cpu_weight = 27500
memory = 2048
cpu_cap = 0
maxvcpus = 2
OVM_high_availability = False
OVM_description = 'Import URLs: [http://10.61.173.176/OVM_EL5U5_X86_64_PVM_10GB.tgz]'
on_poweroff = 'destroy'
on_crash = 'restart'
bootloader = '/usr/bin/pygrub'
name = '0004fb000006000051096a9e0c9f3360'
guest_os_type = 'default'
vfb = ['type=vnc,vncunused=1,vnclisten=127.0.0.1,keymap=en-us']
vcpus = 2
OVM_os_type = 'None'
OVM_cpu_compat_group = ''
OVM_domain_type = 'xen_pvm'
[root@stlrx300s6-131 ~]#
```

c. Check the disk parameter value, in this example: “:/OVS/Repositories/0004fb0000030000c6cfa8ad15f46620/VirtualDisks/0004fb000012000037d313718bdf0da2.img”

```
TESO-17> ndmpcopy -d -sa root:netapp123 -da root:netapp123 -st text -dt text
10.61.172.164:/vs2_dnfs_rac/ovm_nfs_repository/VirtualDisks/0004fb000012000037d313718bdf0da2.img
10.61.172.164:/vs2_dnfs_rac/ovm_nfs_repository_dev_test/VirtualDisks/systemndmpcopy.img
Ndmpcopy: Starting copy [ 14 ] ...
Ndmpcopy: 10.61.172.164: Notify: Connection established
Ndmpcopy: 10.61.172.164: Notify: Connection established
Ndmpcopy: 10.61.172.164: Connect: Authentication successful
Ndmpcopy: 10.61.172.164: Connect: Authentication successful
Ndmpcopy: 10.61.172.164: Log: DUMP: creating
"/vs2_dnfs_rac/ovm_nfs_repository/../snapshot_for_backup.5" snapshot.
Ndmpcopy: 10.61.172.164: Log: DUMP: Using Partial Volume Dump of selected subtrees
Ndmpcopy: 10.61.172.164: Log: DUMP: Using snapshot_for_backup.5 snapshot
Ndmpcopy: 10.61.172.164: Log: DUMP: Date of this level 0 dump: Fri Jan  3 11:38:24 2014.
Ndmpcopy: 10.61.172.164: Log: DUMP: Date of last level 0 dump: the epoch.
Ndmpcopy: 10.61.172.164: Log: DUMP: Dumping
/vs2_dnfs_rac/ovm_nfs_repository/VirtualDisks/0004fb000012000037d313718bdf0da2.img to NDMP
connection
Ndmpcopy: 10.61.172.164: Log: DUMP: mapping (Pass I)[regular files]
Ndmpcopy: 10.61.172.164: Log: DUMP: Reading file names from NDMP.
Ndmpcopy: 10.61.172.164: Log: DUMP: mapping (Pass II)[directories]
Ndmpcopy: 10.61.172.164: Log: DUMP: estimated 2548713 KB.
Ndmpcopy: 10.61.172.164: Log: DUMP: dumping (Pass III) [directories]
Ndmpcopy: 10.61.172.164: Log: DUMP: dumping (Pass IV) [regular files]
Ndmpcopy: 10.61.172.164: Log: RESTORE: Fri Jan  3 11:38:26 2014: Begin level 0 restore
Ndmpcopy: 10.61.172.164: Log: RESTORE: Fri Jan  3 11:38:26 2014: Reading directories from the
backup
Ndmpcopy: 10.61.172.164: Log: RESTORE: Fri Jan  3 11:38:26 2014: Creating files and directories.
Ndmpcopy: 10.61.172.164: Log: RESTORE: Fri Jan  3 11:38:26 2014: Writing data to files.
Ndmpcopy: 10.61.172.164: Log: ACL_START is '2591247360'
Ndmpcopy: 10.61.172.164: Log: DUMP: dumping (Pass V) [ACLs]
Ndmpcopy: 10.61.172.164: Log: RESTORE: Fri Jan  3 11:38:40 2014: Restoring NT ACLs.
Ndmpcopy: 10.61.172.164: Log: DUMP: 2530522 KB
Ndmpcopy: 10.61.172.164: Log: DUMP: DUMP IS DONE
Ndmpcopy: 10.61.172.164: Log: DUMP: Deleting
"/vs2_dnfs_rac/ovm_nfs_repository/../snapshot_for_backup.5" snapshot.
Ndmpcopy: 10.61.172.164: Log: RESTORE: RESTORE IS DONE
```

```
Ndmpcopy: 10.61.172.164: Notify: restore successful
Ndmpcopy: 10.61.172.164: Log: DUMP_DATE is '5683734400'
Ndmpcopy: 10.61.172.164: Notify: dump successful
Ndmpcopy: Transfer successful [ 0 hours, 0 minutes, 27 seconds ]
Ndmpcopy: Done
TESO-17>
```

d. Move the `systemndmpcopy.img` to the proper location.

```
[root@stlrx300s6-131 ~]# mv
/OVS/Repositories/0004fb0000030000a8cd28901056dbb2/VirtualDisks/systemndmpcopy.img/0004fb00001200
0037d313718bdf0da2.img
/OVS/Repositories/0004fb0000030000a8cd28901056dbb2/VirtualDisks/systemndmpcopy.img_temp
[root@stlrx300s6-131 ~]# rm -fr
/OVS/Repositories/0004fb0000030000a8cd28901056dbb2/VirtualDisks/systemndmpcopy.img
[root@stlrx300s6-131 ~]# mv
/OVS/Repositories/0004fb0000030000a8cd28901056dbb2/VirtualDisks/systemndmpcopy.img_temp
/OVS/Repositories/0004fb0000030000a8cd28901056dbb2/VirtualDisks/systemndmpcopy.img
[root@stlrx300s6-131 ~]# find /OVS/Repositories/0004fb0000030000a8cd28901056dbb2
/OVS/Repositories/0004fb0000030000a8cd28901056dbb2
/OVS/Repositories/0004fb0000030000a8cd28901056dbb2/.generic_fs_stamp
/OVS/Repositories/0004fb0000030000a8cd28901056dbb2/Assemblies
/OVS/Repositories/0004fb0000030000a8cd28901056dbb2/Templates
/OVS/Repositories/0004fb0000030000a8cd28901056dbb2/ISOs
/OVS/Repositories/0004fb0000030000a8cd28901056dbb2/VirtualDisks
/OVS/Repositories/0004fb0000030000a8cd28901056dbb2/VirtualDisks/systemndmpcopy.img
/OVS/Repositories/0004fb0000030000a8cd28901056dbb2/VirtualMachines
/OVS/Repositories/0004fb0000030000a8cd28901056dbb2/.ovsrepo
/OVS/Repositories/0004fb0000030000a8cd28901056dbb2/.snapshot
[root@stlrx300s6-131 ~]#
```

e. Create the clone file in storage from `systemndmpcopy.img`.

```
TESO::*> volume file clone create -volume ovm_nfs_repository_dev_test -source-path
/VirtualDisks/systemndmpcopy.img -destination-path /VirtualDisks/clone2.img -vserver vs2_dnfs_rac
```

f. Create the uuid from the `uuidgen` command.

```
[root@stlrx300s6-131 ~]# uuidgen
f314471e-2a1c-4e47-a3aa-af776f166265
```

g. Copy the `TemplateVM vm.cfg` file to the `/tmp/` folder that will be modified for the cloned VM (in this example, CloneVM4).

```
[root@stlrx300s6-131 ~]# cp
/OVS/Repositories/0004fb0000030000c6cfa8ad15f46620/VirtualMachines/0004fb000006000051096a9e0c9f33
60/vm.cfg /tmp/vm.cfg_clone4
```

h. Modify the following parameters in `/tmp/vm.cnf_clone4 –`

disk, uuid, name, `OVM_simple_name`—as below.

```
[root@stlrx300s6-131 ~]# cat /tmp/vm.cfg_clone4
OVM_simple_name = 'CloneVM4'
disk = ['file:/OVS/Repositories/0004fb0000030000a8cd28901056dbb2/VirtualDisks/clone4.img,xvda,w']
bootargs = ''
uuid = 'f314471e-2a1c-4e47-a3aa-af776f166265'
on_reboot = 'restart'
cpu_weight = 27500
memory = 2048
cpu_cap = 0
maxvcpus = 2
OVM_high_availability = False
OVM_description = 'Import URLs: [http://10.61.173.176/OVM_EL5U5_X86_64_PVM_10GB.tgz]'
on_poweroff = 'destroy'
on_crash = 'restart'
bootloader = '/usr/bin/pygrub'
name = 'f314471e2a1c4e47a3aaaf776f166265'
guest_os_type = 'default'
vfb = ['type=vnc,vncunused=1,vnclisten=127.0.0.1,keymap=en-us']
```

```
vcpus = 2
OVM_os_type = 'None'
OVM_cpu_compat_group = ''
OVM_domain_type = 'xen_pvm'
[root@stlrx300s6-131 ~]#
```

i. Create the folder for the cloned VM in the VirtualDisk folder based on the UUID result.

```
[root@stlrx300s6-131 ~]# mkdir
/OVS/Repositories/0004fb0000030000a8cd28901056dbb2/VirtualMachines/f314471e2a1c4e47a3aaaf776f1662
65
[root@stlrx300s6-131 ~]#
```

j. Copy the `/tmp/vm.cnf_clone4` to: "/OVS/Repositories/0004fb0000030000a8cd28901056dbb2/VirtualMachines/f314471e2a1c4e47a3aaaf776f166265/vm.cfg

```
[root@stlrx300s6-131 ~]# cp /tmp/vm.cfg_clone4
/OVS/Repositories/0004fb0000030000a8cd28901056dbb2/VirtualMachines/f314471e2a1c4e47a3aaaf776f1662
65/vm.cfg
```

k. Refresh the new repository NetApp_CDOT_Dev_Test_Repository.

```
OVM> list repository
Command: list repository
Status: Success
Time: 2014-01-03 17:32:51,434 EST
Data:
  id:0004fb0000030000c6cfa8ad15f46620  name:NetApp_CDOT_NFS_Repository
  id:0004fb0000030000a8cd28901056dbb2  name:NetApp_CDOT_Dev_Test_Repository
OVM> refresh repository name=NetApp_CDOT_Dev_Test_Repository
Command: refresh repository name=NetApp_CDOT_Dev_Test_Repository
Status: Success
Time: 2014-01-03 17:33:04,706 EST
OVM>
```

l. Add the CloneVM4 to one of the OVM Servers.

```
OVM> add vm name=CloneVM4 to server name=stlrx300s6-131
Command: add vm name=CloneVM4 to server name=stlrx300s6-131
Status: Success
Time: 2014-01-03 17:35:00,295 EST
OVM>
```

m. To start the CloneVM4, assign the vnic and network.

```
OVM> addAvailableVnic Vm name=CloneVM2
Command: addAvailableVnic Vm name=CloneVM2
Status: Success
Time: 2014-01-03 13:23:07,559 EST
Data:
  id:0004fb00000700003efeae4f8ac099e9  name:00:21:f6:00:00:05
OVM> list Network
Command: list Network
Status: Success
Time: 2014-01-03 17:34:43,704 EST
Data:
  id:0a3dad00  name:10.61.173.0
OVM>
OVM> edit vnic name=00:21:f6:00:00:05 Network=10.61.173.0
Command: edit vnic name=00:21:f6:00:00:05 Network=10.61.173.0
Status: Success
Time: 2014-01-03 13:23:29,099 EST
OVM>
```

n. Start the cloned VM.

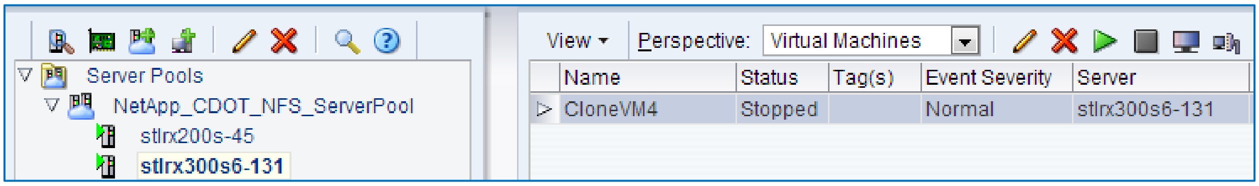

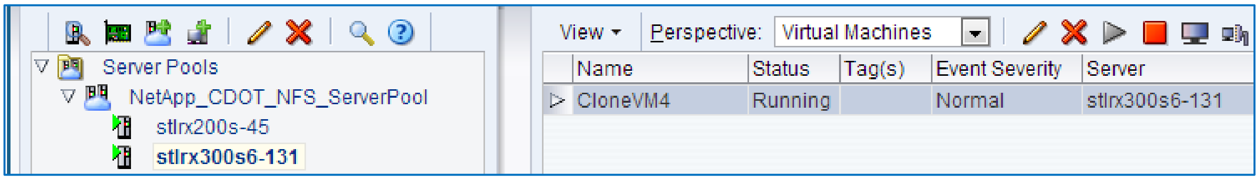

o. We created four cloned VMs through the procedure above and migrated them to another OVM Server.

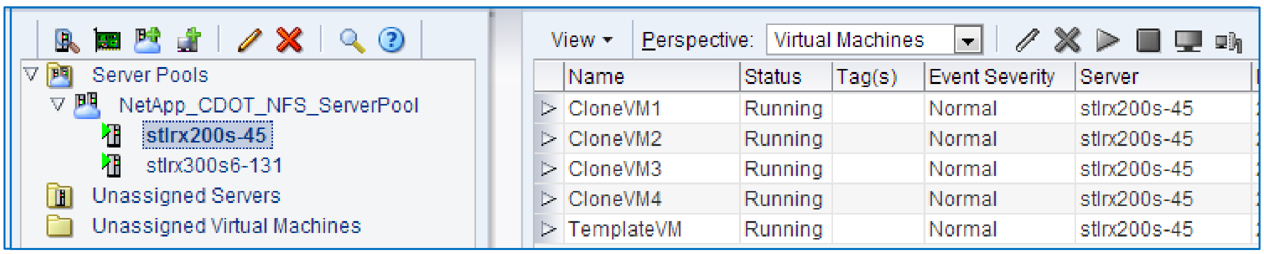

Note: The procedure above is applicable for NetApp storage controller (7-Mode) with OVM 3.x.

Clustered Data ONTAP and OVM 3.x (SAN–iSCSI and FCP)

Refer to section 4.3, “Configuring FC Shared Storage and VM Operations on Oracle VM Server.”

You can also check the “Create Access Groups, LUNs During VM Operations” section in http://community.netapp.com/fukiw75442/attachments/fukiw75442/fas-data-ontap-and-related-plugins-articles-and-resources/131/1/NetApp+Plug-in+2+0+1+for+Oracle+VM.pdf (download software and documentation, view and download the installation and administration guide) for detailed steps to create the access group and physical disk and to use VM Create and VM Clone.

5 Case Study—Oracle VM in a Data Center Environment with NetApp Storage

Oracle Managed Cloud Services, Oracle's software-as-a-service business, introduced server virtualization within Oracle Managed Cloud Services using Oracle VM. In addition to the consolidation benefits, the use of Oracle VM in the Oracle on Demand data center also provides additional benefits such as capacity on demand, rapid provisioning, and high availability of virtual machines.

Use of NetApp storage with Oracle VM in the Oracle Managed Cloud Services 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 to make seamless the transition to the virtualized environment from the physical environment, 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 described here.

- **FlexClone**

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

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

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

  The Oracle Managed Cloud Services 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, refer to section 3.6.

For more details about Oracle VM implementation in Oracle Managed Cloud Services, refer to the white paper from Oracle at http://www.oracle.com/ondemand/collateral/virtualization-oracle-vm-wp.pdf.

6 Appendix

6.1 NFS—Storage Repository Creation in Clustered Data ONTAP Using the CLI from OVM Manager

1. Check the available file server:

```
OVM> list FileServer
Command: list FileServer
Status: Success
Time: 2013-11-25 14:24:10,791 EST
Data:
  id:0004fb0000090000c452c6c91c7f2061  name:vs2_dnfs_rac
```

2. List the file system:

```
OVM> list FileSystem
Command: list FileSystem
Status: Success
Time: 2013-11-25 14:20:43,925 EST
Data:
  id:06ce34d9-af20-4ad0-af0d-4cc8f1759e78  name:nfs:/ovm_generic_repo
  id:b60b13c6-c2ff-4ad8-90c6-9cdc7c8d3652  name:nfs:/ovm_generic_sp
  id:0885467b-2769-4748-ba76-556d2e82c789  name:nfs:/ovm_vol
  id:5d9247cd-6a53-4a57-bf61-22fb0e48d71b  name:nfs:/pdb54_cl_data
  id:d287749e-3d62-41e1-a21e-85072689d679  name:nfs:/solclusterdb1
  id:98ee516a-c6f4-4257-8619-01c89ca08bba  name:nfs:/vsFCP
  id:8fd889d2-adad-40e6-b4ba-d0cc9c946956  name:nfs:/mml_11g_linux_poc_db
  id:8988660a-531d-49e5-92e3-5c30710387a1  name:nfs:/ovm_generic_repo2
```

3. Refresh the exported junction path before creating the repository:

```
OVM> refresh filesystem name=nfs:/ovm_generic_repo2
Command: refresh filesystem name=nfs:/ovm_generic_repo2
Status: Success
Time: 2013-11-25 14:22:20,749 EST
```

4. Create a repository:

```
OVM> create Repository name=NetAppCDOTNFSRepo filesystem=nfs:/ovm_generic_repo2
Command: create Repository name=NetAppCDOTNFSRepo filesystem=nfs:/ovm_generic_repo2
Status: Success
Time: 2013-11-25 14:23:26,265 EST
Data:
  id:0004fb0000030000e4b4cd7cd503145f  name:NetAppCDOTNFSRepo
```

5. List the OVM Servers:

```
OVM> list Server
Command: list Server
Status: Success
Time: 2013-11-25 14:24:23,768 EST
Data:
  id:00:00:00:00:00:00:00:00:00:00:26:2d:04:d2:8e  name:stlrx200s-45
  id:20:7e:32:67:89:a9:11:df:a8:47:00:19:99:83:3b:9a  name:stlrx300s6-131
  id:42:2a:52:5e:49:66:ab:c7:38:42:0a:ed:05:23:9a:9b  name:ovmvmserver1
```

6. Present the OVM Server to the repository:

```
OVM> add Server name=stlrx300s6-131 to Repository name=NetAppCDOTNFSRepo
Command: add Server name=stlrx300s6-131 to Repository name=NetAppCDOTNFSRepo
Status: Success
Time: 2013-11-25 14:25:15,153 EST
OVM> refresh Repository name=NetAppCDOTNFSRepo
Command: refresh Repository name=NetAppCDOTNFSRepo
Status: Success
Time: 2013-11-25 14:25:41,157 EST
OVM> exit
OVM> Connection to 10.61.173.178 closed.
```

7. Verify that the newly created repository is mounted in the OVM Server:

```
[root@stlrx300s6-131 ~]# df -h
Filesystem            Size  Used Avail Use% Mounted on
/dev/sda4              48G  853M   45G   2% /
tmpfs                 330M     0  330M   0% /dev/shm
none                  330M   80K  330M   1% /var/lib/xenstored
10.63.164.18:/ovm_generic_sp
                       48G  137M   48G   1% /nfsmnt/b60b13c6-c2ff-4ad8-90c6-9cdc7c8d3652
/dev/mapper/ovspoolfs
                       10G  369M  9.7G   4% /poolfsmnt/0004fb0000050000bf5f8c2855437096
10.63.164.18:/ovm_generic_repo
                       38G  7.0G   32G  19% /OVS/Repositories/0004fb0000030000a008989dbce7f526
10.63.164.18:/ovm_generic_repo2
                       38G  128K   38G   1% /OVS/Repositories/0004fb0000030000e4b4cd7cd503145f
[root@stlrx300s6-131 ~]#
```

You can use the OVM Manager to verify the repository.

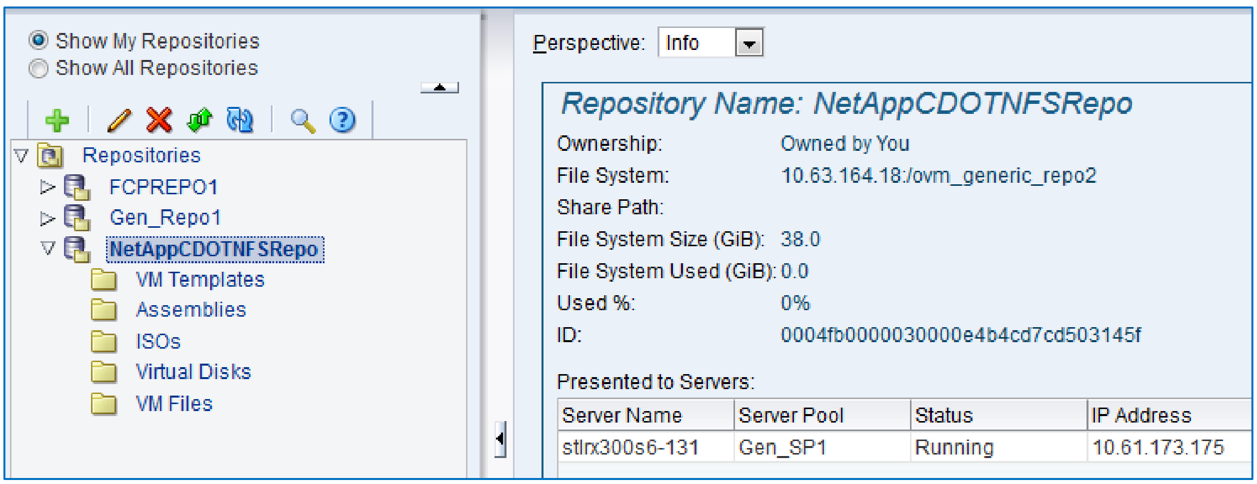

6.2 SAN—Storage Repository Creation in Clustered Data ONTAP Using the CLI from OVM Manager

Install the NetApp Plug-in for Oracle VM Plug-in, which can be downloaded from http://community.netapp.com/fukiw75442/attachments/fukiw75442/fas-data-ontap-and-related-plugins-articles-and-resources/131/1/NetApp+Plug-in+2+0+1+for+Oracle+VM.pdf.

1. List the available NetApp volumes:

```
OVM> list volumegroup
Command: list volumegroup
Status: Success
Time: 2013-11-25 14:51:47,346 EST
Data:
  id:osciscsi_cdot @ 0004fb0000090000b84070a9d38a6f29  name:osciscsi_cdot
  id:prodvolgroup @ 0004fb0000090000b84070a9d38a6f29  name:prodvolgroup
  id:Generic_iSCSI_Volume_Group @ Unmanaged iSCSI Storage Array  name:iSCSI Volume Group
OVM>
```

2. Create the physical disk in the volume:

```
OVM> create physicaldisk size=20 name=cDOTFCPdisk shareable=yes on volumegroup name=prodvolgroup
Command: create physicaldisk size=20 name=cDOTFCPdisk shareable=yes on volumegroup
name=prodvolgroup
Status: Success
Time: 2013-11-25 15:00:38,428 EST
Data:
  id:0004fb00001800001bcf58db4cce23d9  name:cDOTFCPdisk
OVM>
```

3. Check the newly created physical disk:

```
OVM> list physicaldisk
Command: list physicaldisk
Status: Success
Time: 2013-11-25 15:02:18,885 EST
Data:
  id:0004fb00001800007718cdd713ed8858  name:PRODSPLUN1
  id:0004fb000018000031d02332ebe427e3  name:PRODRPLUN2
  id:0004fb00001800006e33ca9a5692270e  name:PRODVMLUN3
  id:0004fb0000180000ea6a88bc103b5c50  name:PRODVMLUN3_20131021172909
  id:0004fb00001800001bcf58db4cce23d9  name:cDOTFCPdisk
```

4. Check the available server pool:

```
OVM> list serverpool
Command: list serverpool
Status: Success
Time: 2013-11-25 15:04:04,960 EST
Data:
  id:0004fb000002000041cb579bb09c649a  name:FCPSP1
  id:0004fb000002000041d559d1015915fd  name:Gen_SP1
OVM>
```

5. Check the available access group:

```
OVM> list accessgroup
Command: list accessgroup
Status: Success
Time: 2013-11-25 15:16:50,586 EST
Data:
  id:FCPAG @ 0004fb0000090000b84070a9d38a6f29  name:FCPAG
  id:ISCSIAG @ 0004fb0000090000d1bf8d2758fd6452  name:ISCSIAG
OVM>
```

6. Add the newly created physical disk to the access group:

```
OVM> add physicaldisk name=cDOTFCPdisk to accessgroup name=FCPAG
Command: add physicaldisk name=cDOTFCPdisk to accessgroup name=FCPAG
Status: Success
Time: 2013-11-25 15:20:12,176 EST
OVM>
```

7. Create a repository in the newly created physical disk:

```
OVM> create Repository name=cDOTFCPRepo2 serverPool=FCPSP1 physicalDisk=cDOTFCPdisk
Command: create Repository name=cDOTFCPRepo2 serverPool=FCPSP1 physicalDisk=cDOTFCPdisk
Status: Success
Time: 2013-11-25 15:28:21,199 EST
Data:
  id:0004fb0000030000cd43031532b19e75  name:cDOTFCPRepo2
OVM>
```

8. Present the OVM Server in the repository:

```
OVM> list server
Command: list server
Status: Success
Time: 2013-11-25 15:29:20,977 EST
Data:
  id:00:00:00:00:00:00:00:00:00:00:00:26:2d:04:d2:8e  name:stlrx200s-45
  id:20:7e:32:67:89:a9:11:df:a8:47:00:19:99:83:3b:9a  name:stlrx300s6-131
  id:42:2a:52:5e:49:66:ab:c7:38:42:0a:ed:05:23:9a:9b  name:ovmvmserver1
OVM>
OVM> add server name=stlrx200s-45 to repository name=cDOTFCPRepo2
Command: add server name=stlrx200s-45 to repository name=cDOTFCPRepo2
Status: Success
Time: 2013-11-25 15:43:16,041 EST
OVM>
```

9. Refresh the repository:

```
OVM> refresh repository name=cDOTFCPRepo2
Command: refresh repository name=cDOTFCPRepo2
Status: Success
Time: 2013-11-25 15:49:47,522 EST
OVM>
```

7 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.

8 Revision History

Date	Name	Description
December 2009	Preetom Goswami Antonio Jose Rodrigues Neto Padmanabhan Sadagopan	Created and updated for Oracle VM 2.2
April 2011	Padmanabhan Sadagopan	Updated the note
November 2011	Padmanabhan Sadagopan	Updated the `o2cb_idle_timeout_ms` value to 160000 and updated the note
March 2014	Karthikeyan Nagalingam	Updated the document for OVM 3.2.x and clustered Data ONTAP