



An Oracle White Paper
September 2013

Oracle VM 3: Backup and Recovery Best Practices Guide

Contents

Introduction.....	1
Part 1: Product Architecture, Concepts and Tools.....	2
How to use this Guide	2
Understanding Oracle VM 3 Architecture in General	2
Oracle VM Manager.....	3
Oracle VM Servers.....	4
Oracle VM Server Pools	4
Oracle VM Guests	7
Oracle VM Guest Resources	7
Understanding Storage.....	8
Shared Storage.....	8
Local storage	10
Understanding Storage Repository Architecture	10
Understanding Deployment Architecture	13
Example deployment architecture using NAS	14
Example deployment architecture using SAN	15
Design a highly available, highly scalable storage architecture	15
Understanding Oracle VM Guest cloning.....	17
Cold clone.....	19
Hot clone	19
Challenges to Backup and Recovery Solutions for Oracle VM.....	19
Oracle VM Guests cannot directly access tape devices	20
Oracle VM Guests are not always neat packages.....	20
Understanding what needs to be backed up.....	20
Oracle VM Manager.....	21
Pool File System.....	21
Oracle VM Servers.....	21
Oracle VM Guests	21
Oracle VM Guest Resources	22

Backup Strategies	23
Preferred backup methods for Oracle VM	23
A tiered approach to backups	23
A note about restoring snapshots.....	23
Backup frequency	24
Backup retention windows	24
Exporting storage repositories to back up Oracle VM Guests	24
Using cloning to back up Oracle VM Guests	25
Oracle VM tools to help facilitate automated backup and recovery	25
Part 2: Backup & Recovery for Oracle VM Manager	27
Overview	27
Backup Frequency	27
Validate Backup & Restore as Part of Initial Implementation.....	27
A Note about Automated Snapshots	28
A Note about Restoring Snapshots	28
Critical Oracle VM Manager Files & Directories	28
Critical Oracle VM Automated Backup Files & Directories	29
Change Default Location for Automated Backups.....	29
Suggested Backup Architectures.....	30
Building a Multitier Backup Scheme.....	31
Combine the OS and Database Backup Schemes.....	31
Devise a Multitier Backup Scheme for the Operating System	31
Physical Server Backup using Hardware RAID	31
Physical Server Backup using KickStart.....	32
Physical Server Backup using SAN Boot	33
Oracle VM Manager as Guest VM using SAN or NAS	35
Devise a Multitier Backup Scheme for the Manager Database	36
Manager Backup on Physical Server with Internal Storage.....	36
Manager Backup on Physical Server with Shared Storage	37
Manager Backup as a Guest using MySQL Database	38

Oracle VM Manager Backup as a Guest using an Oracle Database	40
Backup Solutions.....	42
How to Create an Ad Hoc Backup using MySQL	42
Recovery Solutions.....	43
How to recover an Ad Hoc Backup using MySQL	43
Conclusion.....	44
Part 3: Backup & Recovery for Oracle VM Guests and Resources	45
Overview	45
Critical Files & Directories for Oracle VM Guests and Resources ...	45
Backup Frequency.....	45
Choosing a backup approach for Oracle VM guests	46
Quiesce Activity during Backups.....	48
Segregating Guest Resources into a Common Repository	49
The Value of Segregating Guest Resources from Guest Images	49
Implementing a Solution for Segregation	50
Multitier Backup for Guest Resources at Storage Level	52
Single Tier Backup for Guest Resources at Repository Level.....	54
Conclusion.....	54
References	55

Introduction

Oracle VM 3 is a scalable and resilient server virtualization solution with built-in clustering technology designed to keep mission critical business systems running under a variety of enterprise workloads. Oracle VM forms the foundation of public and private cloud platforms integrating a full stack of hardware and software solutions engineered to work together. So, a well designed disaster recovery plan that includes a robust backup and recovery solution is a key component for the quick resumption of operations after experiencing disastrous events.

The goal of this guide is to provide knowledge and insights into Oracle VM that will help the reader develop backup and recovery processes that fit the unique requirements of the readers' Oracle VM deployment. The document delves into product architecture, design concepts, best practices and solutions that will help the reader achieve the goal of developing backup and recovery solutions for Oracle VM environments.

An Oracle VM 3 environment is comprised of a collection of software products combined with storage and networking to build a stable platform for virtual machines and corresponding business systems. There are several distinct components of Oracle VM that must be part of a backup plan. This paper is divided into three parts with the first part dedicated to explaining Oracle VM architecture and design considerations for creating and executing successful backup and recovery of Oracle VM. The remaining parts of the guide explain in detail how to backup and recover the significant components that make up an Oracle VM platform.

Part 1: Product Architecture, Concepts and Tools

Oracle VM is a scalable and robust platform intended to form a solid foundation to build both private and public cloud infrastructures for small to large enterprise data centers. However, we realize not everyone will be designing, deploying and managing Oracle VM 3 environments for large enterprise data centers; some deployments of Oracle VM will be very small indeed.

Individuals and small business will still find that the solutions documented herein can easily be adapted to lower cost backup products and even very low tech tools available on any Linux distribution such as dump, tar, cpio, rsync as well as a plethora of other tools and shareware available throughout the Linux community at large. For smaller organizations that don't have access to sophisticated tape or snapshot technologies, simply substitute whatever tool you plan to use in place of any of the tape backup solutions we discuss in this document.

How to use this Guide

The guide is designed to fulfill two purposes:

- *Concepts* – Part 1 is dedicated to architecture and concepts meant for someone that is charged with designing the infrastructure, backup and restoration plan for an initial implementation of Oracle VM. Part 1 also includes details of Oracle VM architecture and internals that are important to understand for those individuals that might not be involved in day-to-day operations of Oracle VM but are responsible for designing and implementing a backup and restoration strategy for Oracle VM.
- *Solutions* – Part 2 and Part 3 of the document are dedicated to explaining the specifics of backing up and restoring individual components of Oracle VM. Readers can skip directly to these parts without exploring Part 1 if they simply want to quickly backup or restore something specific. Of course, Oracle recommends that the reader not skip Part 1 since important concepts for designing a robust and resilient Oracle VM platform are explained in some detail.

The reader will not find any solutions in this document that recommend or explain how to use technologies such as dd, tar, dump, cpio. We feel that employing these applications will only make backup and restores unnecessarily complex, slow and inefficient. All of your data should reside on enterprise class storage where you have access to highly efficient and sophisticated software and applications that make backing up and restoring data simple, exceedingly quick and highly reliable.

Understanding Oracle VM 3 Architecture in General

Oracle VM can be deployed many, many different ways to fit the unique needs and requirements of individual data centers. But, there are always five major components to Oracle VM no matter which deployment architecture is devised to fit your unique requirements.

- **Oracle VM Manager** – Oracle VM Manager is a standalone graphical application used to manage the Oracle VM environment (model) including storage, servers, server pools and virtual machines. (see Item 1 in Figure 1)
- **Oracle VM Servers** – Oracle VM Servers are the physical, bare metal servers that provide shared resources such as CPU, memory, storage and networking needed to run one or more Oracle VM Guests on each server. (see Item 2 in Figure 1)

- Oracle VM server pools – One or more Oracle VM Servers are grouped into server pools to allow any Oracle VM Guest to run on any server in the pool. (see Item 3A, 3B & 3C in Figure 1)
- Oracle VM Guests – Instances of virtual machines that host guest operating systems such as Microsoft Windows, Linux and Solaris. Many different Oracle VM Guests (virtual machines) can run on a single Oracle VM Server. (see Item 4 in Figure 1)
- Storage – Storage can range from internal disks found in each physical server to highly scalable, highly available enterprise class storage arrays for access to multipath shared storage. (see Item 5, 5A & 5B in Figure 1)
- Oracle VM virtual machine resources – These are ancillary tools that can be used to create new Oracle VM guests. Virtual machine resources such as assemblies, ISO images of operating systems and Oracle VM Templates comprise a sixth element that may or may not be incorporated into your particular Oracle VM environment. (see Item 6 in Figure 1)

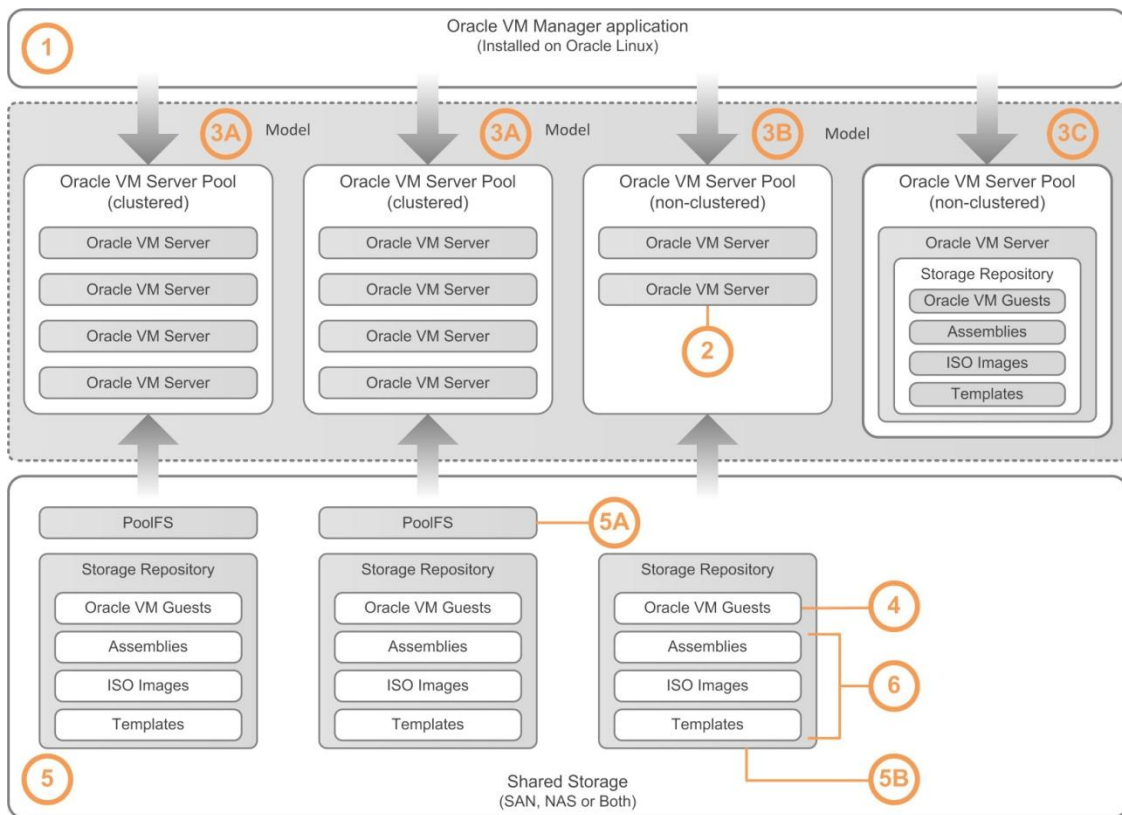


Figure 1: Overview of Oracle VM architecture showing the five major components of shared storage, Oracle VM Manager, servers, pools and guests

Figure 1 above illustrates the relationship between the five major components and a sixth element which may or may not be incorporated into your particular Oracle VM environment. We will use the illustration to explore each of the Oracle VM components in more detail.

Oracle VM Manager

Oracle VM Manager is an application that is installed on Oracle Linux or Red Hat Enterprise Linux (see Item 1 in Figure 1). The Oracle VM Manager is typically installed on a standalone server. The

Oracle VM Manager is used to build the initial infrastructure for the Oracle VM environment (model) and perform ongoing day-to-day management of various objects as well as the infrastructure. It is important to understand that Oracle VM Servers, Guests and high availability features of clustering continue to function even if the Oracle VM Manager is not available due to a complete loss of the Oracle VM database.

Please see [Part 2: Backup & Recovery for Oracle VM Manager](#) for specific information about requirements and suggested backup and recovery processes for the Oracle VM Manager.

Oracle VM Servers

An Oracle VM Server is comprised of the Xen hypervisor and privileged domain (Dom0) that is installed on physical, bare metal machines. The servers constitute the entire runtime platform for Oracle VM Guests providing processing power and other resources such as memory, access to networks, access to storage and cluster capabilities for high availability of Oracle VM Guests.

Oracle VM Servers that belong to non-clustered server pools using **local disk** must be backed up since all critical data about the Oracle VM Guests reside on local disk as opposed to shared storage. Clustered and non-clustered server pools are discussed in more detail below.

Oracle VM Servers that belong to clustered and non-clustered server pools using **shared disk** should be backed up as a matter of practice. However, Oracle VM Servers that are members of server pools with shared storage do not contain any critical custom information can be recovered simply by removing the server from Oracle VM Manager, reinstalling and rediscovering the server; although you may still perform regular backups of the Oracle VM Servers .

Oracle VM Server Pools

One or more Oracle VM Servers are grouped into server pools to allow any Oracle VM Guest to run on any server in the pool. Each Oracle VM Manager can manage multiple server pools which can be either non-clustered or clustered for high availability. Figure 1 above illustrates three slightly different variations on server pools being managed by a single Oracle VM Manager:

All variations boil down to three distinct differences:

- Clustered server pool using shared storage (see Item 3A in Figure 1)
- Non-clustered server pool using shared storage (see Item 3B in Figure 1)
- Non-clustered server pool using local storage (see Item 3C in Figure 1)

The primary difference between a clustered and non-clustered server pool is the fact that non-clustered server pools are created without a pool file system. Non-clustered server pools are not able to take advantage of many high availability features without pool file systems. Also, a non-clustered server pool can scale from 1 to 64 servers using either local or shared storage or a combination of the two for the storage repository.

A clustered server pool can scale from 1 to 32 servers and is the more robust deployment architecture for Oracle VM due to the high availability features built into the product. Clustered server pools must utilize a pool file system which is independent of the Oracle VM Servers but makes them vulnerable to the loss of this storage related object.

Please refer to [Oracle VM 3 user guides](#) for additional information about the differences between clustered and non-clustered server pools.

Clustered server pool with multiple servers

Item 1A in Figure 2 below shows a clustered server pool containing many Oracle VM Servers. The clustered server pool takes advantage of shared storage enabling the full use of all high availability features built into Oracle VM such as auto-restarts, live migration of Oracle VM Guests, Distributed Power Management and Distributed Resource Scheduler. Please refer to [Oracle VM 3 user guides](#) for additional information about the high availability features of Oracle VM.

Notice in Figure 2 item 2 that shared storage is a key component of a clustered server pool and can be either SAN or NAS. Shared storage provides the means of making available the required pool file system (item 2A) and the storage repositories (item 2B) for each server pool.

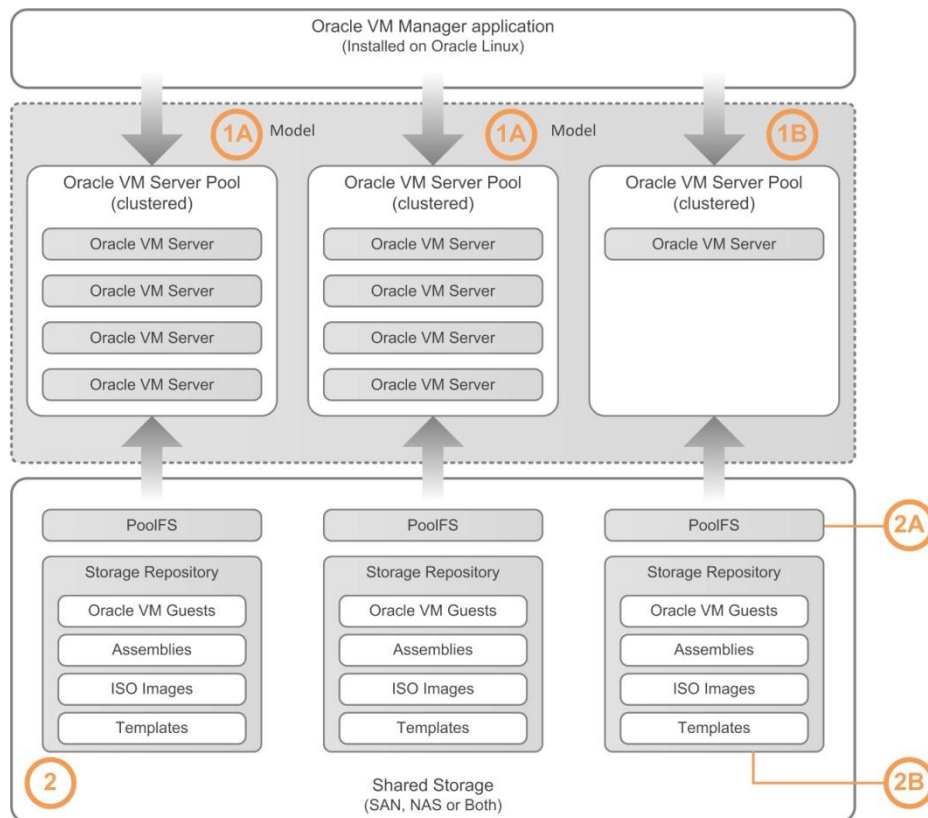


Figure 2: Overview of Oracle VM architecture showing the four major components of Shared storage, Oracle VM Servers, Guests and Manager

Clustered server pool with a single server

Item 1B in Figure 2 illustrates a clustered server pool containing a single Oracle VM Server with the option of enlarging the size of the server pool by adding more Oracle VM Servers at some point in the future. This variation on the clustered server pool is simply shown to demonstrate that it is possible and reasonable to create a single node clustered server pool using shared storage. The value of a clustered single node server pool is simply derived from the fact that additional Oracle VM Servers can

be added at any point without having to change anything; begin with a single server and expand as needed – it is a very scalable solution.

Non-clustered server pool using shared storage

Item 1A in Figure 3 illustrates a non-clustered server pool containing a single Oracle VM Server with the option of enlarging the size of the server pool by adding more Oracle VM Servers at some point in the future. A non-clustered server pool can take advantage of some high availability features built into Oracle VM such as live migration, Distributed Power Management and Distributed Resource Scheduler. Additional servers can be added to the existing pool at any point.

Non-clustered server pools can take advantage of shared storage as well as local storage but it is very important to notice that item 2 in Figure 3 illustrates that only NFS can be used to present the shared storage repository (item 2A) to the Oracle VM Servers. As noted above, non-clustered server pools do not have a pool file system so recovery of the server pool in this case means a successful restoration of individual Oracle VM Servers. Please [Part 3: Backup & Recovery for Oracle VM Guests and Resources](#) for specific information about requirements and suggested backup and recovery processes for physical servers.

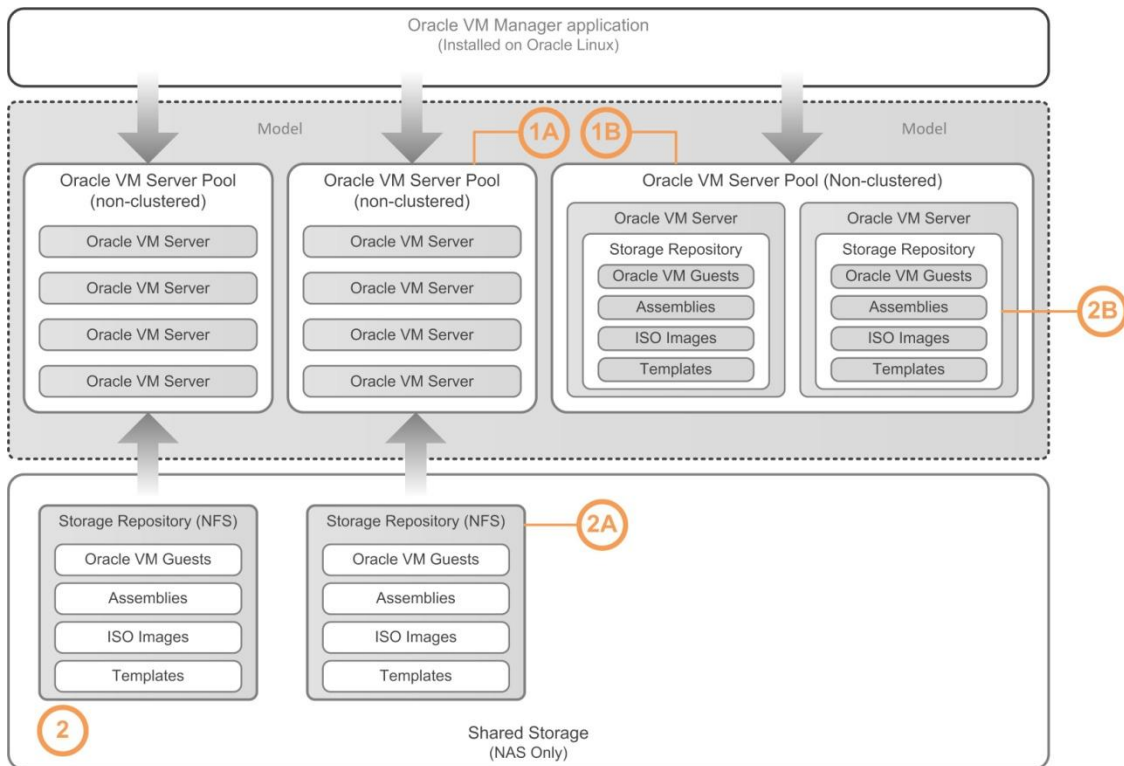


Figure 3: Overview of Oracle VM architecture showing the four major components of local storage, Oracle VM Servers, Guests and Manager

Non-clustered server pool using local storage

Item 1B in Figure 3 shows a non-clustered server pool using local storage available to each individual Oracle VM Server. This variation on server pools utilizes internal local disk instead of shared storage

to create a server pool of one or more Oracle VM Servers. Notice that the storage repository (item 2B) is local to each Oracle VM Server in the server pool shown as item 2B is not part of the shared NAS storage array (item 2).

This variation on the non-clustered server pool is shown to demonstrate that it is possible and reasonable to create a single node server pool using local storage. The value of a non-clustered single node server pool is simply derived from the fact that additional Oracle VM Servers can be added at any point without having change anything; begin with a single server and expand as needed – it is a very scalable solution. Keep in mind that non-clustered server pools cannot be easily converted to clustered server pools without rebuilding each pool.

Oracle VM Guests

Oracle VM Guests are Virtual Machines that host guest operating systems, databases and applications that comprise various business systems within your Oracle VM environment. Note in Figure 2 above that the Oracle VM Guests are shown as part of the storage repositories. This is an important distinction since the Oracle VM Servers are where the virtual machines are running, but the storage repository is where all configuration files and virtual disks reside for each Oracle VM Guest.

Backups and recoveries of Oracle VM guests can be quite complex because any physical disks associated with guest operating systems are artifacts of the Oracle VM Servers and are not part of the storage repositories. In addition, your deployment architecture may include applications and data files that are presented to guest operating systems from NFS exports being served from many different NFS servers instead of residing on virtual disks within storage repositories.

Please see [Part 3: Backup & Recovery for Oracle VM Guests and Resources](#) for specific information about requirements and suggested backup process for virtual machines.

Oracle VM Guest Resources

Virtual machine resources such as assemblies, ISO images of operating systems and Oracle VM Templates comprise a sixth element that may or may not be incorporated into your particular Oracle VM environment. (see Item 6 in Figure 1). These objects can be downloaded from the [Oracle Software Delivery Cloud](#) to a local http server or ftp server at your site and then imported into the into storage repositories using Oracle VM Manager. You can create your own Oracle VM Templates from existing Oracle VM Guests that you initially created using Kickstart or some other provisioning tool. So, you may have Oracle VM Templates even if you like to create virtual machines from scratch.

There are three different types of virtual machine resources for Oracle VM Guests:

- ISO images – ISO images of operating systems can be imported into storage repositories and presented to newly created virtual machines as DVD media. The media can then be used to install guest operating systems of your choice. Refer to the [Oracle VM User Guide](#) for more information about creating, downloading, importing and using ISO images to create virtual machines.
- Oracle VM Templates – Templates are virtual machines that are preconfigured with operating systems and sometimes applications. You can also create your own Oracle VM Template by cloning any virtual machine in your Oracle VM environment; perhaps you have installed Oracle Linux using Kickstart or an ISO image and taken the time to configure the operating system a particular way that you want to use as a template from many other similar types of virtual machines.

Refer to the [Oracle VM User Guide](#) for more information about cloning virtual machines to create your own templates or downloading preinstalled/configured templates.

- Assemblies – Assemblies are used to create Oracle VM Templates. Assemblies are simply zipped files containing all the resources needed to create one or more preinstalled/preconfigured Oracle VM Templates. Assemblies can be downloaded from Oracle or created using Oracle Virtual Assembly Builder. Refer to the [Oracle VM User Guide](#) for more information about downloading and using assemblies to create Oracle VM Templates.

Please see [Part 3: Backup & Recovery for Oracle VM Guests and Resources](#) for specific information about requirements and suggested backup process for virtual machine resources.

Understanding Storage

Shared Storage

Shared storage is basis for the scalability and flexibility inherent in Oracle VM (see item 5 in Figure 1). Shared storage is the foundation on which all else is built and the key to backup and recovery for enterprise class data centers utilizing clustered server pools (see item 1A and 1B in Figure 2). Because of shared storage Oracle VM can live migrate virtual machines from one server to another in the same pool allowing rolling upgrades of Oracle VM Servers, empowering the use distributed resource scheduler and distributed power management all with zero downtime. Shared storage is also perfunctory for being able to move virtual machines from one server pool to another using the move virtual machine feature (this requires downtime). For example, being able to easily promote an Oracle VM Guest from a server pool dedicated to user acceptance testing to a server pool dedicated to running production business systems can be achieved from the Oracle VM Manager with shared storage.

Due to the robustness and feature rich capabilities of today's storage appliances, data centers keep everything of real importance on the storage hardware, abstracting all important data and applications from physical servers. Data centers allocate storage space to LUNs and NFS exports that are in turn presented to Oracle VM Servers and/or directly to Oracle VM Guests containing data and application binaries. The end result is that mission critical applications and data reside on a flexible platform that any server can access rather than on local disks where it is completely useless if the server is temporarily down or unavailable.

To achieve the most reliable, easy to implement backup solution, the following list of storage related objects should all reside on highly reliable, highly scalable storage appliances. Each storage related object listed below is explained in more depth and used in examples throughout the remainder of this document.

Pool file system

The pool file system is a single shared SAN LUN or NFS export that is mounted onto all Oracle VM servers in any given clustered server pool as shown in Figure 2 item 2A. Pool file systems are needed for clustered server pools.

Each clustered server pool will have its own dedicated pool file system that forms the basis of Oracle VM clustering for high availability. The pool file system contributes two very important components

toward the high availability features of Oracle VM and the reason Oracle VM Servers and Guests are not dependent on Oracle VM Manager for clustering.

- *Quorum disk* - The pool file system acts as the quorum disk for OCFS2 (Oracle Cluster Filesystem Release 2). Oracle VM relies on OCFS2 to maintain integrity of clustered server pools ensuring that the same Oracle VM Guest is not running on multiple Oracle VM Servers (split-brain syndrome).
- *Cluster data* - The pool file system also contains a shared Oracle Berkeley DB (BDB) database containing information about the clustered server pool including the Oracle VM Servers and running HA enabled guests that are part of the server pool.

Storage repositories

The primary function of storage repositories is to contain the configuration files and virtual system images for each Oracle VM Guest, plus static guest resources such as assemblies, ISO images and templates. Each server pool should have a dedicated storage repository just for that particular pool that can contain Oracle VM guests and other virtual machine resources such as assemblies, ISO images and templates. Basically, each server pool should have a minimum of one storage repository:

- A dedicated storage repository for a single server pool to contain Oracle VM Guests and guest resources (NAS or SAN)

As you will see in the section below titled [Understanding Storage Repository Architecture](#), virtual disk images for all Oracle VM Guests, Oracle VM Templates and assemblies are all intermingled in the same **VirtualDisks** directory of the storage repository. This can be problematic when attempting to recover guest resources or Oracle VM Guests.

For example, you may have backed up the entire storage repository in a single backup, but later on might want to just restore a single Oracle VM guest. Since all of the virtual disk images for virtual machines and templates are kept in a common directory, you can't just indiscriminately restore all virtual disk images because you will overwrite virtual disks for templates and other virtual machines that you may not want to restore. This means you will need to figure out which of the virtual disk images belong to that particular Oracle VM guest and then restore only those files.

To help reduce the complexity of restores, it might behoove you to consider the following deployment of storage repositories instead:

- A dedicated storage repository to contain only Oracle VM Guests for a single pool (NAS or SAN). For example, if you had three server pools, each server pool would have its own storage repository dedicated to Oracle VM Guests only; three server pools, three repositories.
- A dedicated storage repository to contain only Oracle VM Guest resources such as assemblies, ISOs and templates for multiple server pools (NAS only). For example, if you had three server pools, then all three pools would have access to the same exact repository containing nothing but guest resources.

The above repository deployment options are discussed in more detail in [Part 3: Backup & Recovery for Oracle VM Guests and Resources](#) of this guide.

NFS exports

NFS exports can be used to present file systems directly to Oracle VM Guest operating systems to contain applications and data used by the business systems being hosted on the virtual machines.

SAN physical disks

LUNs on storage arrays can be presented to the Oracle VM Servers and passed through to Oracle VM guest operating systems as physical disks. Just like the NFS exports, the physical disks can contain applications and data used by the business systems being hosted on the virtual machines. Although beyond the scope of this document, physical disks can also be used to contain the system image for Oracle VM Guests outside the confines of the storage repository.

Local storage

There are occasions where individuals and even data centers find that non-clustered server pools are advantageous as illustrated in Figure 3 above (see items 1A & 1B). In this case everything related to Oracle VM Guests will reside on local storage contained within or connected to a single Oracle VM Server.

Pool file system

Non-clustered server pools do not utilize a pool file system at all since access to local disk is limited to a single server.

Storage repositories

Local internal disk can be used for storage repositories on each Oracle VM Server, but has very limited application due to the fact that local physical disks are only available to the in which they reside. Figure 3 item 2B above shows a non-clustered server pool using local disk as the storage repository on each server.

Understanding Storage Repository Architecture

Understanding the structure of the storage repository is important, but in most instances is by no means the only Oracle VM object that will be part of your backup and recovery plan.

If you designed your guest deployment architecture where 100% of all disks associated with each Oracle VM Guest are virtual disks that reside in the storage repository, then your entire backup and recovery plan for Oracle VM Guests will be focused only on capturing the data contained in each storage repository. However, this is not the case in most enterprise class deployments.

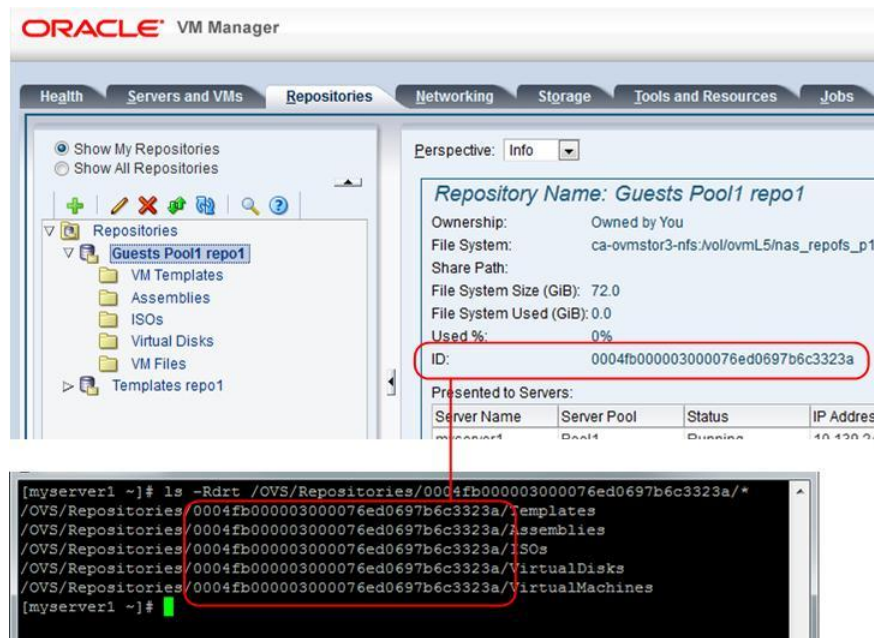


Figure 4: Screen shot showing a storage repository as seen from the perspective of the Oracle VM Manager and the relationship between the Repository ID and the directory name as seen on the Oracle VM Servers

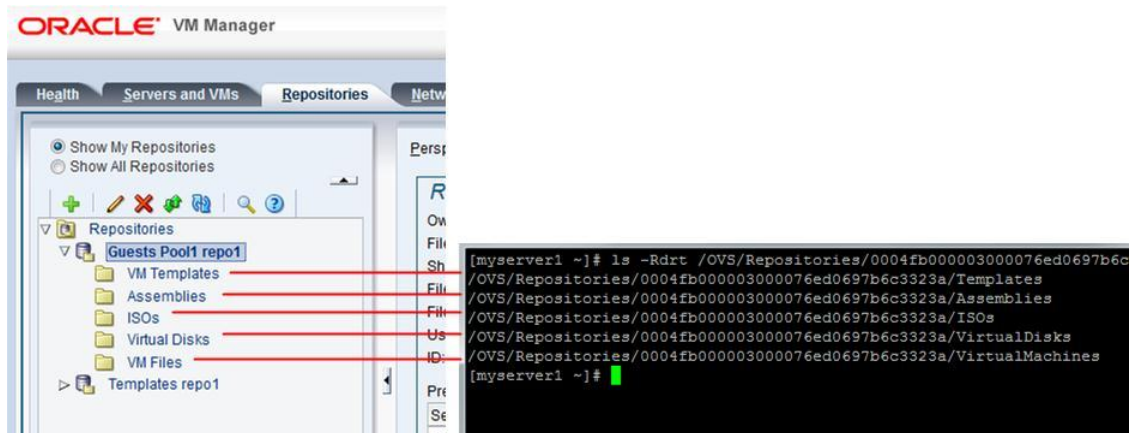


Figure 5: Screen shot showing a storage repository as seen from the perspective of the Oracle VM Manager and Oracle VM Servers

Figure 5 above shows two screen shots. The top screen shot shows a storage repository from the perspective of the Oracle VM Manager and the bottom screen shot shows the same storage repository from the perspective of an Oracle VM Server. The storage repositories are all mounted on each Oracle VM Server under the `/OVS/Repositories<repository ID>` mount point. The repository ID can be seen in the Oracle VM Manager user interface (UI) when the repository is highlighted in the Repositories navigation pane on the left of the UI; the red oval labeled as 1 in top screen shots shows where the repository ID can be found. There is a direct relationship between the repository ID and the mount point for the repository on the Oracle VM server in the red oval labeled as 2 in the bottom screen shot.

Figure 5 above also shows the relationship between the repository folders shown in the navigation pane in the top screen shot and the actual directory names found on the Oracle VM Servers:

- *VM Templates* – this directory contains the vm.cfg files that are associated with each Oracle VM Template that is imported into the model. Only the vm.cfg files reside in this directory, the system images and other virtual disks associated with a template reside with all the other virtual disks in the VirtualDisks directory.
- *Assemblies* – this directory contains the OVA files associated with any assemblies imported into the model. An OVA file is analogous to a zip or tar file: it contains other files such as an XML file needed to construct the vm.cfg files, system images and any other virtual disks needed to create one or more Oracle VM Templates.
- *ISOs* – this directory contains ISO images you might use to install operating systems on when you create Oracle VM Guests without using an Oracle VM Template. The ISO images can be presented to new or existing Oracle VM guests as a CD/DVD during creation or when editing an existing virtual machine.
- *Virtual Disks* – this directory contains all of the virtual disks associated with assemblies, Oracle VM Templates as well as Oracle VM Guests.
- *VM files* – this directory contains the vm.cfg files that are associated with each Oracle VM Guest. Only the vm.cfg files are contained under this directory in each storage repository.

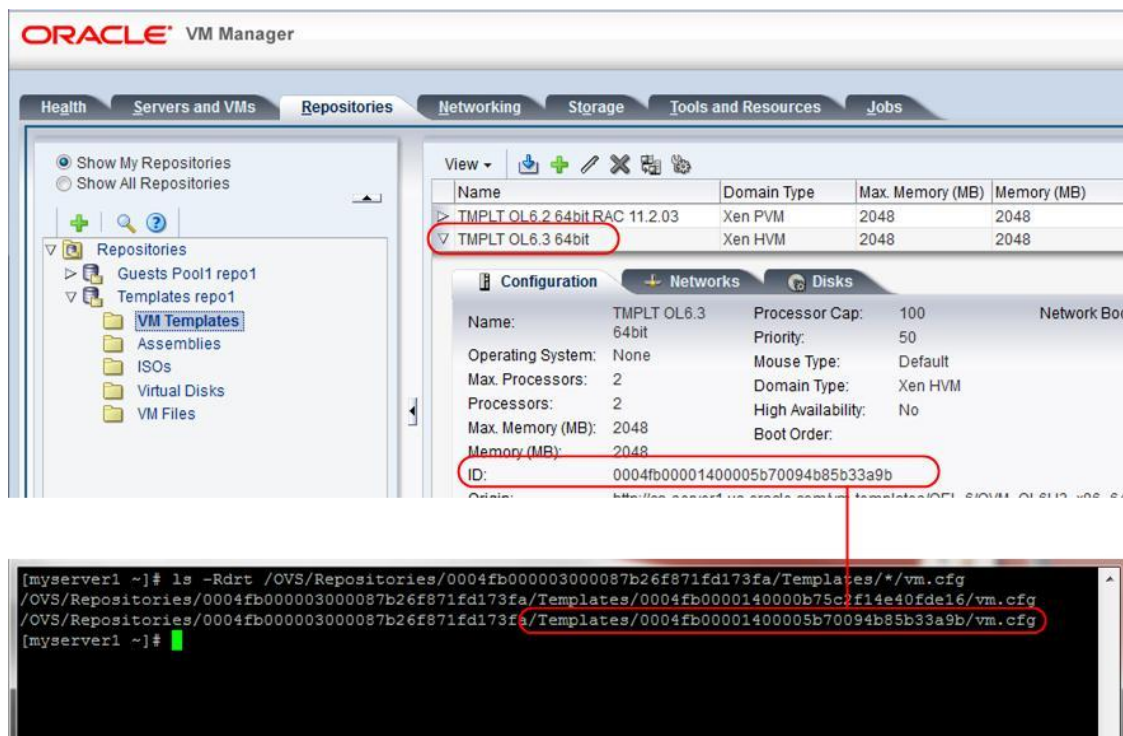


Figure 6: Screen shot showing a specific Oracle VM Template and the location of the corresponding configuration file on the Oracle VM servers

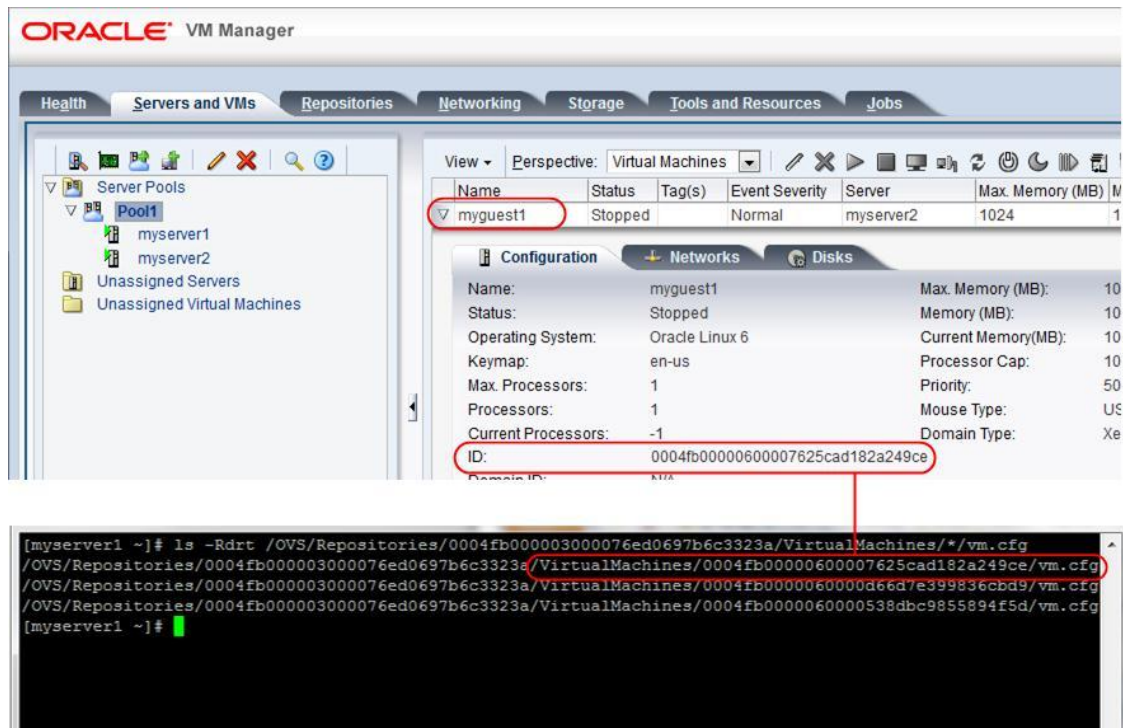


Figure 7: Screen shot showing a specific Oracle VM Guest and the location of the corresponding configuration file on the Oracle VM servers

Understanding Deployment Architecture

Deployment architecture drives the backup and recovery solution. Deployment architecture refers to the concept of how you use servers, storage and networking to build your Oracle VM environment to manage all of the major product components such as the Oracle VM Manager, Oracle VM Servers, and Oracle VM Guests.

For example, do the system images for your Oracle VM Guests reside on virtual disks that are contained in storage repositories or do the system images reside on physical disks that are not part of any storage repository? Do you install applications binaries on each Oracle VM guest operating system or do all application and data files reside on NFS mounted directly to each guest. There are a plethora of ways to design your deployment architecture and the decisions made during this phase of design and implementation impact the way you backup and restore data.

Depending on your deployment architecture, Oracle VM guest objects such as configuration files, system images, data and application disks can be found in a variety of locations within your Oracle VM environment. If you choose to install your guest operating systems on physical disks instead of virtual disk files contained within storage repositories, then your backup and recovery solution will need to adjust for backing up entire disks as well as individual configuration files always contained in the storage repositories.

Example deployment architecture using NAS

There are a multitude of choices that can be made when deciding how to deploy servers, NAS storage (file level protocols such as NFS) and network to provide a virtual server environment for running Oracle VM Guests. The illustration in Figure 8 below shows an example that will be used throughout this document to illustrate backup and recovery concepts. Keep in mind that the illustration shows only one of many possible variations but an exhaustive explanation of all the variations is beyond the scope of this document. The example is built around NAS storage since NFS provides a highly flexible and scalable solution that is easy to backup and restore components and objects associated with Oracle VM.

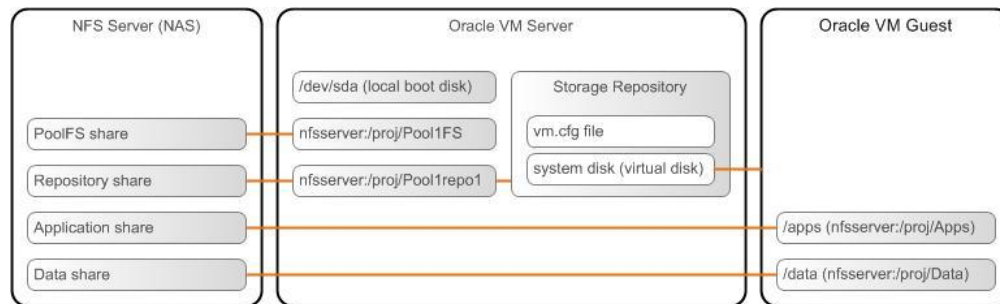


Figure 8: Example Oracle VM deployment architecture showing all storage related objects being presented from NAS servers using a file-level protocol such as NFS

Beginning with the left side of the illustration in Figure 8 above and working toward the right, we start with a shared storage solution where all objects associated with a server pool reside. As noted earlier, centralized enterprise class storage is the key to a robust and easy solution because it provides the ability to take advantage of snapshot technology for quick, reliable backups and restores. Notice that the applications and data NFS exports are passed directly to the Oracle VM Guests.

Oracle of course recommends enabling and using DNFS with Oracle databases for higher performance while maintaining a flexible storage infrastructure based on NFS. Notice that all Oracle VM components reside on the NAS appliance where the execution of few snapshot commands can in theory capture 100% of the data associated with an entire Oracle VM environment.

Note in the middle box above representing an Oracle VM Server that the operating system and hypervisor are installed on a local boot disk; everything else related to the server pool such as pool file system and repositories all reside on the NAS server. Anything of importance and value has been abstracted from the Oracle VM Server which helps make the bare metal server nothing but an easily replaced part.

Also notice in the box representing the Oracle VM Server that the storage repository containing the system image for the Oracle VM Guest is simply an NFS mount from the NAS server. Also like the bare metal Oracle VM Server, the applications and data being used by the guest operating system are abstracted from the operating system through the NFS mounts being passed directly to the virtual machine. So taking a snapshot on the storage array captures the entire guest operating system along with the configuration file, applications and data. Please refer to [Part 3: Backup & Recovery for Oracle VM Guests and Resources](#) for more detailed information about Oracle VM Guests.

Example deployment architecture using SAN

Figure 9 below shows a deployment architecture using SAN (block level protocols such as FCP or iSCSI). It is almost identical to the architecture shown in Figure 8 above. The only difference being that all storage related objects are presented to both the Oracle VM Server and Guests as physical disks. In both cases it is important to note that everything associated with the Oracle VM environment is located on the storage arrays where snapshots can quickly and easily capture all the data during a backup. So, it's the same basic architecture using the same basic approach to backups and recovery using snapshots on the storage array... simple.

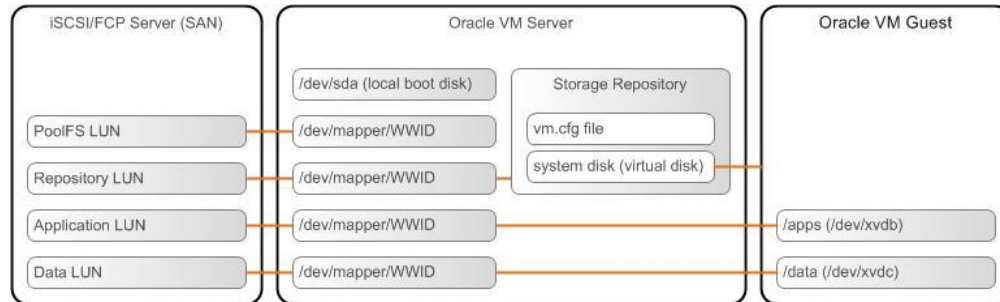


Figure 9: Example Oracle VM deployment architecture showing all storage related objects being presented from SAN servers using block-level protocols such as FCP or iSCSI

Design a highly available, highly scalable storage architecture

The solutions discussed in this document for backing up and restoring various Oracle VM objects all assume you are using enterprise class storage arrays with snapshot capabilities. With the exception of Oracle VM Servers, the significance of this approach to storage is that 100% of an Oracle VM environment resides on a platform that is protected from single points of failures and provides the ability to backup in seconds and restore in minutes. This means the backup and restore methodology is consistent and simple without having to resort to using various standard Linux tools such as tar, dd, dump, etc.

Taking a few moments to explore the storage architecture used in the examples above and throughout this document should help tie it all together. As noted earlier, storage is really the key to designing a robust, flexible and scalable Oracle VM environment. Well designed storage architecture will not only make the Oracle VM environment more resilient, it will also make it easier to backup and restore.

The previous diagrams depicting deployment architectures show a single source of storage for purposes of expediency. In reality, most storage infrastructures for an Oracle VM platform will be comprised of more than one storage array. It is highly likely that storage related to Oracle VM infrastructure will be completely separate from storage related the applications and data for business systems being hosted on the Oracle VM Guests.

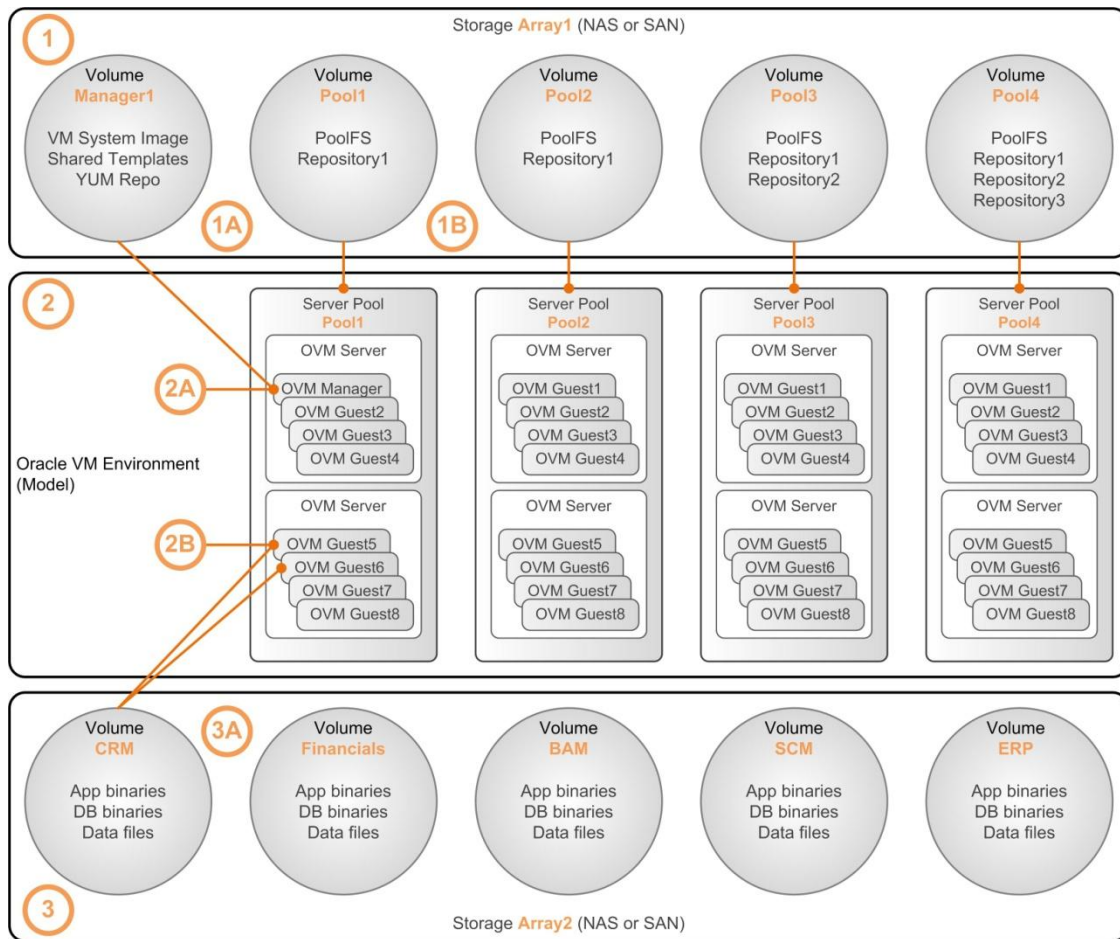


Figure 10: Segregation of data for Oracle VM infrastructure and business systems fosters simpler backup & recovery

Figure 10 above depicts a simplistic example of scalable and resilient storage architecture. This is by no means the only way to deploy storage and is only presented here as a means of explaining the concept of designing a resilient architecture. The deployment architecture for storage in your data center should be designed around the limits of the storage platform and unique requirements particular to IT/data governance in your data center.

The following list describes some of the most important lessons that can be derived from the sample architecture shown in Figure 10 above. The principles of the proffered architecture are applicable using either NAS (NFS) or SAN (FCP, iSCSI); you can even use a combination of all three storage protocols. There are multiple benefits to ensuring you have designed a robust storage architecture that uses at least some of the best practices described in the following list:

- *Multiple storage arrays* – note that the diagram depicts two different storage arrays. In this case, we are showing all storage related to Oracle VM infrastructure including pool file systems and storage repositories are presented from **storage array1**, while storage related to applications and data for individual business systems resides on **storage array2**. There is no requirement what-so-ever to have more than one storage array, but it is very common for large data centers to present storage from many different sources. (See items 1 & 3 in Figure 10 above)

- *Data abstraction* – Most data centers will abstract the application binaries, data, index files, redo logs, etc from the guest operating systems for each virtual machine; the binaries for the applications running on the guest operating system and the data being read and written to data stores all reside on either NFS mounts or physical disks being presented from storage arrays. (See items 2B & 3A in Figure 10 above)
- *Storage segregated by server pool* – this is an important best practice that should be the basis for any Oracle VM storage architecture. Notice in the container for **storage array1** that each server **pool1**, **pool2** and **pool3** each have a dedicated container for all storage related to a single server pool making each server pool independent of one another. For example, running out of space for a pool file system will render the OCFS2 region to become un-writable which will in turn cause all the servers in the same server pool to reboot. Only the servers in a single server pool will be impacted since storage is segregated by server pool. Another very important aspect of this model is restoring snapshots is a little more granular allowing restores of entire snapshots to impact only a single server pool. (See item 1B in Figure 10 above)
- *Oracle VM Manager data on storage array* – at the very least ensure that /u01 for the Oracle VM Manager resides on a storage array. If you are using the MySQL database that comes with Oracle VM, then this allows you to easily backup and restore the Oracle VM database using the automated daily backups that reside in /u01/app/oracle/mysql/dbbackup. If using a remote or local Oracle database instead of the MySQL database, then that data should reside on a storage array somewhere and also be included in your backup plan. (See item 1A in Figure 10 above)
- *Oracle VM Manager as a VM guest* – most of our examples show the entire server with the Oracle VM Manager application residing on a guest operating system. This is not a requirement at all, but a significant number of Oracle customers install Oracle VM Manager on a guest operating system. This means the entire Oracle VM Manager and database are contained with a single virtual disk image, which in turn means a single point-in-time snapshot will capture 100% of the Oracle VM Manager; recovering from a complete loss of the Oracle VM Manager is as simple as restoring a single file. (See item 2A in Figure 10 above)
- *Snapshot technology* – Without a doubt, access to snapshot technology is one of the most important aspects of using a robust enterprise class storage solution as the platform for all data related to an Oracle VM environment. Snapshot capability is an invaluable tool in relation to backups and recoveries since this provides the ability to quickly backup and restore various storage related objects with minimal effort.

Understanding Oracle VM Guest cloning

Oracle VM allows cloning of virtual machines and templates to create new copies which can then be backed up or customized. There are two basic choices for cloning virtual machines:

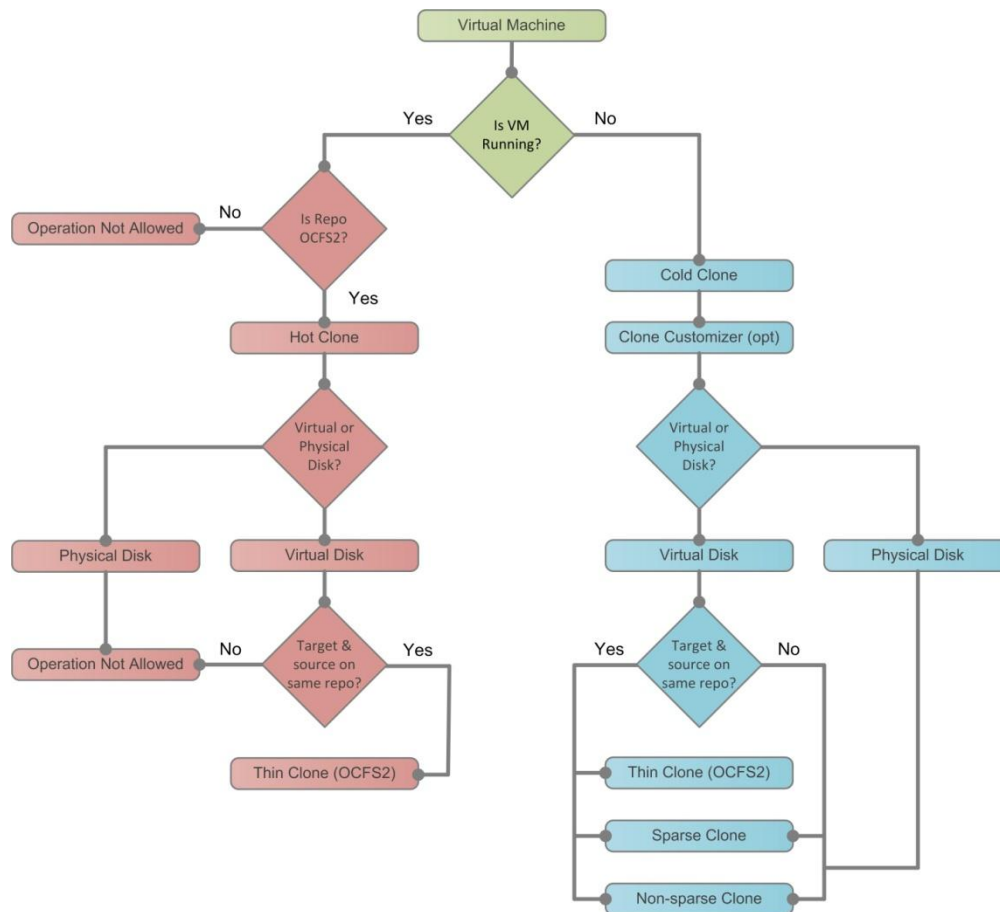
- *Cold clone* – in this case the virtual machine is completely stopped. This choice offers widest number of supported conditions.
- *Hot clone* – in this case the virtual machine is running. This choice is only supported for Oracle VM Guests with no running databases and must reside on virtual disks within an OCFS2 formatted storage repository.

Depending on the type of storage used to contain the virtual machines, there are three different methods that can be employed during the cloning operation.

- *Non-sparse clones* – Disks are duplicated by allocating the same space as the original physical or virtual disk

- Sparse clones – Disks are duplicated by allocating only the space used by the original physical or virtual disks
- Thin clones – Thin clones utilize the OCFS2 ReLink feature to quickly create a snapshot of virtual disk images while the source virtual machine is running

So, cloning will capture the configuration file associated with an Oracle VM Guest for sure, the system image and any other virtual disks associated with the guest but not any physical disks or NFS exports mounted directly to the guest operating system. Depending on your deployment architecture, the system disk may be a physical disk instead of a virtual disk and the only thing you end up cloning is the configuration file – which of course will not allow you to back out of an operating system or application update that was not successful.



In Oracle VM 3, multiple Virtual Machines or templates can be created from a source Virtual Machine through cloning for backup purposes. Users can choose to clone a Virtual Machine before doing maintenance work such as Operating System patching. Clones can be taken online (hot clone) or offline (cold clone). Changes can be reverted immediately by booting up the clone Virtual Machines. If a template is created instead, Virtual Machines can be deployed from those templates.

Virtual Machine Clones can also be booted up concurrently with the source Virtual Machines and therefore individual files can also be recovered from them. However, if this is to be done, precautionary steps would have to be taken to avoid an IP conflict between the source Virtual

Machines and the clones. Application or database consistencies have to be considered before taking hot clones. For database Virtual Machines, it is still recommended to use tools like RMAN for backup.

Please refer to **Cloning a Virtual Machine or Template** in chapter 7 of the latest Oracle VM User's Guide for additional information about the cloning process. The Oracle VM User Guide can be found on the [Oracle Technology Network documentation site](#).

Cold clone

Cold clones can be created when the source virtual machine is in stopped state where data consistency is assured. There are three different ways that cold clones; namely, thin clones, sparse clones and non-sparse clones. Thin clones will be discussed in details under hot cloning in the following section. Oracle VM 3 creates sparse and non-sparse clones through the DD (data description) copy function. Sparse and non-sparse clones can be created across repositories. A move Virtual Machine function also uses the DD function to file copy virtual disk images across different repositories.

Hot clone

Thin clones can be created while the source Virtual Machine is running. Thin clones utilize the Oracle Clustered File System 2 (OCFS2) REFLINK technology to create hard links between the source and target virtual disks. Thin clones are basically inodes pointing to the same data blocks as the source Virtual Machines. When a thin clone is booted up and starts writing data, new delta data blocks are created through the copy-on-write technology of OCFS2.

Thin clones do not guarantee data consistency and disks might need to be filesystem checked before they are usable. Database Virtual Machines and Virtual Machines with heavy transactions are not suitable candidates for thin cloning.

Although thin clones take up negligible disk space initially, it can grow up to the original size of the source Virtual Machine. Therefore, enough disk space on the storage repositories have to be catered for.

However, Oracle VM cloning relies on the following assumptions about your guest deployment architecture:

- *OCFS2 format* – Oracle VM cloning uses the ref link feature of OCFS2 to create a copy of the configuration file and the virtual disks associated with a particular Oracle VM Guest. Therefore, your storage repository must be presented using either FCP or iSCSI and not NFS.
- *Virtual disks only* – the other challenge is the cloning feature only works with those objects related to an Oracle VM Guest that reside in storage repositories.

Challenges to Backup and Recovery Solutions for Oracle VM

The problem of backup and recovery cannot simply be approached from the vantage point of backing up singular objects such as a server, pool file system, storage repository or the management database for Oracle VM. The deployment architecture you developed for Oracle VM in your data center determines the complexity or ease of the backup/recovery solution for your environment. In other words, there is no single simple backup and recovery solution that can be applied consistently from

data center to data center; a backup/recovery plan is predicated entirely on the deployment architecture and the tools you plan to use to capture the data such as tape or snapshot technology.

Oracle VM Guests cannot directly access tape devices

At the time of this writing, Oracle VM Server does not support PCI pass-through SCSI devices for Oracle VM Guests. This means that controllers for robotic arms in tape libraries or SCSI tape devices cannot be presented directly to the Oracle VM guests.

Oracle VM Guests are not always neat packages

Oracle VM Guests are virtual machines containing a guest operating system such as Linux, Solaris or Windows. Each Oracle VM Guest can consist of a number of files, but require a minimum of the vm.cfg file and a system image. However, only the vm.cfg file is guaranteed to be located in a storage repository.

This is where the challenge comes into the picture: Oracle VM Guests are not always neat little packages bundled into a single container for easy backup and recovery. System images, application binaries and data files can be presented from many different sources of storage which can make finding, backing up and restoring individual virtual machines quite complex.

Please see [Understanding Storage Repository Architecture](#) for specific information about files associated with Oracle VM Guests.

Understanding what needs to be backed up

Different components of the entire Oracle VM environment require different approaches to backup and restores. Figure 11 below shows a map of five major components of Oracle VM that will require a backup plan, including Oracle VM Servers which are optional.

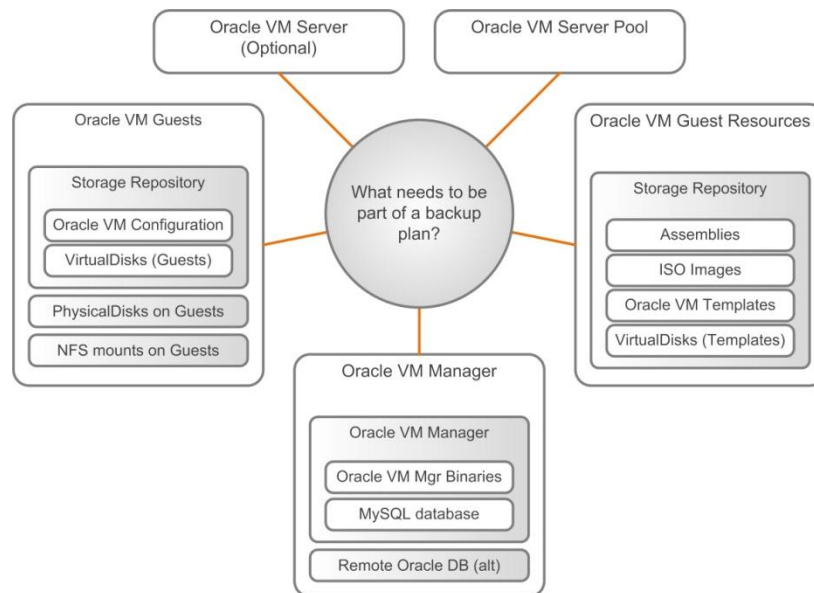


Figure 11: A map showing the Oracle VM components and objects that need to be included in a backup plan

Oracle VM Manager

Although the Oracle VM Manager is used for day-to-day monitoring and management of the environment, it is not critical to the successful continuance of your Oracle VM environment (model) in the event of disaster or loss of the server where Oracle VM Manager is installed. The Oracle VM Servers, Guests and high availability features of clustered server pools will continue to function for short or extended periods of times during the temporary loss of the Oracle VM Manager.

However, it is still very important that the Oracle VM Manager database is backed on a daily or perhaps even more frequent basis if a lot of changes are being made. The Oracle VM management database is the only data store where all of the information about every aspect of the model is maintained. In particular, simple names and descriptions of physical and virtual disks, elements of the network infrastructure, Oracle VM Servers and Guests are not maintained anywhere except in the Oracle VM Manager database. In addition, information about the storage arrays, access groups and other relationships between objects are not maintained anywhere else.

Please see [Part 2: Backup & Recovery for Oracle VM Manager](#) for specific information about requirements and suggested backup process for the Oracle VM Manager and Oracle VM database.

Pool File System

The pool file system is a very important component of Oracle VM. However, the pool files system doesn't contain any user application or virtual machine data. Losing a pool file system severely impacts both the Oracle VM Servers and Guests in the server pool that it supports, but does not impact any other servers in any other pools being managed by Oracle VM Manager. If the pool file system is destroyed, a clean way is to create a new pool and add the existing resources back to it. Please consult Oracle VM Documentation of how to create a server pool.

Oracle VM Servers

Oracle does not recommend backing up the Oracle VM Servers. It is the goal of the product to make the physical servers an easily replaced commodity component of the overall platform. Instead of spending the time backing up and recovering a server, simply delete the server from the Oracle VM Manager, reinstall and discover the Oracle VM Server, then return it to the server pool.

Oracle VM Guests

Backing up Oracle VM Guests can be very complicated because all of the objects associated with VM guests can be found either in a single location or perhaps even a variety of locations depending on your deployment architecture. The location of all storage related objects associated with each Oracle VM Guest are contained in the vm.cfg file of each virtual machine. In a nutshell, the objects that must be included in a backup plan are simply the vm.cfg file itself and all the virtual and or physical disks that are presented to each guest operating system (Oracle VM guest).

There are three basic types of data that must be part of the backup plan, but the three types of data can potentially reside within a variety of storage related objects associated with each Oracle VM Guest – so the location and method of backing up and restoring the data can vary widely. The following list shows the three basic types of data that must be part of the backup plan for virtual machines:

- *Oracle VM Guest configuration file* – The Oracle VM Guest configuration file contains all the meta-data needed by Xen to start an Oracle VM Guest. In particular, the `vm.cfg` file contains the location of all storage related objects that you will need to include as part of any backup and recovery plan. The configuration file happens to be the only file associated with Oracle VM Guests that always resides in a storage repository regardless of the deployment architecture.
- *Oracle VM Guest system image* – The system image for the Oracle VM Guest (boot/OS file) will be either one of the following disk types. The location of the system image depends on how you design your guest deployment architecture but the location can always be found in the configuration file above.
- *Applications & data storage* – Applications and application data can reside within the Oracle VM Guest system image as part of the root file system or additional storage objects such as virtual disks, physical disks as well as NFS mounts. It is Oracle recommended best practice to abstract all applications and data from the Oracle VM Guest system image.

Oracle VM Guest Resources

Oracle VM related objects such as assemblies, ISO images and Oracle VM Templates are known collectively as Oracle VM Guest Resources. Each type of image performs a supporting role in the way you deploy Oracle VM Guests. You may or may not use any Oracle VM Guest Resources depending on how you deploy virtual machines.

For example, if you present virtual machine boot image as a physical disk, then use KickStart to install and configure the guest OS for each virtual machine, then you are not likely to have any guest resources residing your storage repositories. However, if you create a virtual machine image as a virtual disk residing on one of your storage repositories, then you might have ISO images.

The following list provides a short description of the role each type of static image plays in the Oracle VM environment:

- *Assemblies* – Assemblies are used to create Oracle VM Templates in the Oracle VM Manager and are simply OVA files imported into the model. An OVA file is analogous to a zip or tar file: it contains other files such as an XML file needed to construct the `vm.cfg` file(s), system image(s) along with any other virtual disks needed to create one or more Oracle VM Templates.
- *ISOs* – ISO images can be used to install a guest operating system on an Oracle VM Guest after you create a virtual machine without using an Oracle VM Template. The ISO images can be presented to new or existing Oracle VM guests as a CD/DVD during creation or when editing an existing virtual machine.
- *VM Templates* – VM templates include two types of files that need to be backed up: the `vm.cfg` files, the system images and other virtual disks associated with each.

The Oracle VM User's Guide contains more information about the static images listed above in Chapters [7.4 Virtual Machine Installation Media](#) and [7.5 Virtual Machine Resources](#) at the time of this writing. Please refer to the User's Guide for more information on how these files are used in day-to-day operations and requirements and suggested backup process for virtual machine resources in the storage repositories.

Please see [Part 3: Backup & Recovery for Oracle VM Guests and Resources](#) of this guide for specific about requirements

Backup Strategies

This document discusses how to determine which Oracle VM components and objects need to be backed up in order to quickly and reliably recover from a system failure or disaster. The document does not attempt to explain backup and recovery concepts or procedures since organizations will already have well defined backup products, strategies and procedures defined. Oracle makes no suggestions about retention windows, recovery time objectives, multiplexing, multi-streaming or full, incremental and differential backups. It is also assumed by this document that the reader already knows how to perform hot backups and how to quiesce applications, databases, Oracle VM servers or guests.

Preferred backup methods for Oracle VM

The choice of backup methods and technology used in your environment is completely up to you. For the purposes of this guide, products basically fall into three categories of backup technologies:

- *Linux tools* – This includes such system applications such as cpio, dd, dump, tar, gzip, zip etc.
- *Tape backup* – This includes such products as Oracle Secure Backup, Symantec NetBackup, Tivoli Storage Manager, etc.
- *Snapshot* – This includes snapshot capabilities inherent with various storage solutions such as Oracle ZFS Storage Appliance, Oracle Pillar Storage Systems, Dell, EMC², Hitachi Data Systems, IBM, NetApp, and so on

A tiered approach to backups

Most data centers will approach backups using a tiered approach to provide maximum protection while at the same time minimizing the time spent backing up or recovering an Oracle VM 3 environment. Such a backup architecture might include snapshots as the first tier of backups, moving snapshot data to near online storage such as a virtual tape library and finally using streaming media such as tape to capture the snapshots for longer retention periods.

For example, most data centers will quiesce applications and databases related to a business system for a minute or two and then take a point-in-time snapshot of all related data on the storage array. The backup is accomplished very quickly allowing applications to be backed up with minimal downtime and in most cases zero downtime when combined with hot backups of Oracle databases. Since data contained in snapshots are static, tape backups of the data for the purposes of archiving can be taken with little concern for time and with no impact to running applications.

Ideally, your Oracle VM deployment architecture will include robust, enterprise class shared storage with snapshot technology allowing you to put all databases into hot backup mode, prepare middleware and applications running on Oracle VM guest operating systems according to methods proscribed by the software vendors, then take a snapshot using the features of your NAS or SAN storage appliance.

A note about restoring snapshots

We use snapshot technology as the first tier for all backup and recovery solutions presented in this guide, but snapshot technology is not required to implement a tiered approach. You can use tar, cpio,

dump or any other available tool in place of any solution where this guide indicates taking or restoring a snapshot.

How a snapshot is actually restored is completely dependent on your choice of storage vendor hardware and software along with established best practices and standards for your data center.

For those readers not entirely familiar with snapshot technology, in most cases snapshots are restored on the storage array not the Oracle VM Servers. However, a lot of data centers use the practice of temporarily mounting both the export containing the snapshot and the corrupted pool file system to temporary mount points such as /snapshot_of_data and /original_data (whatever you want to name the temporary mount points), then rsync or copy the files from the snapshot location to the temporary mount point.

Once all the files are restored from the snapshot location, the file system containing the snapshot data is un-mounted from the temporary mount point and the administrator can then move on to the next step in any of the recovery recipes described in this document. The storage array will handle all the magic of restoring the files back to the point-in-time when the snapshot was created even though standard Linux commands like cp or rsync are being used – not all storage vendor solutions will have this capability.

Backup frequency

The frequency in which periodic backups are performed is completely dependent upon existing backup policies for your data center and functional business units. Generally, the frequency of backups is predicated on the tolerance for risk and recovery time objective (RTO) for both the business systems and the Oracle VM environment. Here are some general guidelines for backup frequency to give you a feel for how you might choose to design your backup/recovery plan. The backup requirements and recommendations of your individual business systems should always win over any general guidelines suggested in this guide pertaining to the frequency of backups.

Backup retention windows

The length of time to retain snapshots and tape backups is completely dependent upon existing backup policies for your data center and functional business units. Retention periods for snapshots are usually measured in terms of a few days to a week, necessarily short due to the way the software tracks changes between the live data and static data. However, restores from snapshots can be accomplished in minutes if any of the data need to be restored within the relatively short retention period. Tape backups are retained for much longer periods, so if data need to be restored after a snapshot has been destroyed, then they can be restored from tape. Tape restores of course are measured in hours.

Exporting storage repositories to back up Oracle VM Guests

Storage repositories presented using SAN block level protocols such as iSCSI or FCP can be exported as a Network File System (NFS) to a backup server. Perhaps you have a Windows, Linux or Unix server that has backup software installed that can write to remote or local tape, or perhaps you have something as simple as a USB disk that you want use to keep a periodic copy of the entire contents of the repository. The storage repository can be mounted to such a server and the contents can be copied

to tape or external near on-line storage using almost any tool such backup software, a copy command, rsync, tar, zip, etc.

Like cloning, simply backing up storage repositories as a means of backup and recovery relies on all objects associated with Oracle VM Guests being contained within the storage repositories, and like cloning this is not always the case. Exporting the storage repositories to be mounted on a tape media server is superfluous if the Oracle VM Guest also includes the use of physical disks and NFS exports mounted directly to the guest operating system.

However, if 100% of Oracle VM Guests including the system image, applications and data all reside on virtual disks contained in storage repositories, then this might be an excellent and simple means of backing up and recovering the business systems being hosted on Oracle VM Guests. Exporting and mounting the storage repositories to a media server would be most applicable to backup plans that primarily use tape backup software or other more primitive system tools like tar, cpio or dump to copy the files to tape or other disks.

Please refer to Oracle VM User's Guide [Enabling Storage Repository Back Ups](#) for instructions on how to export storage repositories as NFS.

Using cloning to back up Oracle VM Guests

The Oracle VM Manager has the ability to clone Oracle VM Guests, which is roughly analogous to a snapshot. An obvious use case for the cloning capability would be in a scenario where a point-in-time copy of an Oracle VM Guest would allow you to apply updates to the guest operating system and have an easy way to back out of the update in case things go awry.

Cloning may work under some limited circumstances as a backup and recovery solution for Oracle VM Guests, but it is probably going to be most effective as a means of making periodic ad-hoc backups of a virtual machine before making some sort of change that might have unforeseen catastrophic results.

Oracle VM tools to help facilitate automated backup and recovery

Oracle VM has a variety of tools to help accomplish automated backup and recovery solutions that can be integrated with tape and snapshot products. Some of these tools can be integrated with other Oracle and third party backup solutions to create seamless, automated backup and recovery solutions. Creating automated solutions using the following tools is beyond the scope of this document, but it is important to know that highly sophisticated custom backup and recovery methods can be developed with simple tools that can be used by systems administrators as well as product developers.

- *Oracle VM Manager Command Line Interface (CLI)* - Oracle VM Manager 3 CLI provides a command line interface for the Oracle VM Manager 3. You can use the CLI to perform the same functions as the Oracle VM Manager Web Interface, such as managing all your server pools, servers and guests. The CLI commands can be scripted and run in conjunction with the Web Interface, thus bringing more flexibility to help you deploy and manage an Oracle VM environment.
- *Oracle VM Guest Additions* - Oracle VM Guest Additions is a set of packages that can be installed on the *guest* operating system of a virtual machine running in the Oracle VM environment. These packages provide the tools to allow bi-directional communication directly between the *Oracle VM Manager* ; and the operating system running within the virtual machine. This is a powerful tool that

provides administrators fine-grained control over the configuration and behavior of components running within the virtual machine directly from Oracle VM Manager.

Please review Oracle VM [documentation](#) for further details about the above tools.

Part 2: Backup & Recovery for Oracle VM Manager

Overview

The Oracle VM Manager database is relatively easy to backup and restore. The product is resilient enough to allow recovery from disasters even if no valid backups are available simply by rediscovering the Oracle VM Servers. Although recovering the database from server discovery returns the database to a working model, there are few things that are not restored. For example, the custom simple names for objects such as physical and virtual disks along with a few other object types are not restored using just the server discovery. It is essential that regular and reliable backups are performed.

A very important note to keep in mind is that the loss of the Oracle VM Manager database has no impact on running Oracle VM Servers or Oracle VM Guests. Hopefully this knowledge should help alleviate panic and anxiety in some small way when you are confronted with the loss of the Oracle VM database.

This guide is written specifically for Oracle VM 3.2. Beginning with Oracle VM 3.2, Oracle VM Manager began shipping with MySQL Enterprise Edition as the default database which is supported for deployment in large scale production environments. Oracle VM Manager also began shipping with MySQL Enterprise Backup product which is used to create full, hot backups of the running Oracle VM environment.

The automated daily backups are supported in product environments and can be relied on as the sole solution for full backups of the database provided you have included a means of storing the daily backups on another medium such as tape or storage array for longer term retention and protection.

Backup Frequency

Daily full backups are recommended for Oracle VM environments that are running in production as stable environments where strict change control policies are in place. More frequent full backups are recommended for Oracle VM environments that are stable, but are undergoing frequent changes such as adding and removing servers and virtual machines, migrating virtual machines from legacy Oracle VM 2 environments or where change control is less restrictive.

You should always perform full backups before any major change such as upgrading the Oracle VM Manager, adding or removing servers to an existing server pool or making changes to the storage and network infrastructure.

Validate Backup & Restore as Part of Initial Implementation

It is critical that the backup and restoration process be validated during the initial implementation of Oracle VM and for each additional server pool as they are added to the Oracle VM Manager. The entire backup and restoration process should be executed once the server pool is up, running, validated and one or two Oracle VM Guests exist for validation purposes. Do not begin adding Oracle VM Guests to server pools until you have validated that you can recover from a complete loss of the database.

It is important that the individuals responsible for the day-to-day management of Oracle VM become familiar and at ease with the process of recovering from the loss of the Oracle VM database.

A Note about Automated Snapshots

Most snapshot products from storage vendors provide the ability to automate snapshots on a periodic basis such hourly, daily, weekly, etc. However, the Oracle VM Manager database must be completely quiesced or put into hot backup mode before taking a snapshot. Make sure you coordinate the automated snapshots while the Oracle VM Manager database is quiescent.

A Note about Restoring Snapshots

How a snapshot is actually restored is completely dependent on your choice of storage vendor hardware and software along with established best practices and standards for your data center.

For those readers not entirely familiar with snapshot technology, in most cases snapshots are restored on the storage array not the Oracle VM Servers. However, a lot of data centers use the practice of temporarily mounting both the export containing the snapshot and the corrupted pool file system to temporary mount points such as `/snapshot_of_u01` and `/u01`, then `rsync` or copy the files from the snapshot location back to the `/u01` mount point.

Once all the files are restored from the snapshot location, the pool file system is un-mounted from the temporary mount point and the administrator can then move on to the next step in any of the recovery recipes described in this document. The storage array will handle all the magic of restoring the files back to the point-in-time when the snapshot was created even though standard Linux commands like `cp` or `rsync` are being used – not all storage vendor solutions will have this capability.

Critical Oracle VM Manager Files & Directories

Oracle VM Manager is built on top of Oracle WebLogic server. Almost everything about the product is contained in a single directory with a couple of exceptions noted below.

- Oracle VM Manager application – Everything about the Oracle VM Manager including the WebLogic binaries, JAR files, WAR files, scripts, etc. is contained within the `/u01` directory structure. The only two pieces of Oracle VM Manager that are contained outside of `/u01` are the Oracle VM Manager configuration file and MySQL Enterprise Backup product (see above).
- Oracle VM Manager data files – If you are using the MySQL database, then all data related to Oracle VM model is contained along with everything else under `/u01`. The specific location of the data files is `/u01/app/oracle/mysql/data`. If you are using Oracle Standard or Enterprise editions then your database administrator will be able to tell you where the data file reside and how to back them up.
- Oracle VM Manager configuration file for database – The file `/u01/app/oracle/ovm-manager-3/.config` is created during the install process by `runInstaller.sh` and contains key information about accessing the database and core product being managed by the WebLogic server. This file is removed when you uninstall Oracle VM Manager.
- Oracle VM Manager configuration file for backups – The file `/etc/sysconfig/ovmm` file is created by the `runInstaller.sh` script during the initial install and contains persistent information about the Oracle VM Manager UUID as well as information needed by MySQL Enterprise Backup to perform the daily automated backup of the Oracle VM database when using MySQL. This file

remains in place if you uninstall Oracle VM Manager. It is used by the `runInstaller.sh` script if you reinstall Oracle VM Manager to ensure the product is reinstalled using the UUID from a previous install.

Critical Oracle VM Automated Backup Files & Directories

If you installed Oracle VM Manager using the **simple** option, then the default database is MySQL. Oracle VM Manager performs full automated hot backups of the Oracle VM database on a daily basis using the MySQL Enterprise Backup application which is also included as part of the Oracle VM Manager install. The automated backups are rotated daily and have a retention period of 15 days. Please refer to [Oracle VM 3 Installation and Upgrade Guide](#) for the latest information about automated backups for Oracle VM.

- MySQL backup data – The default location for each backup on the Oracle VM management server is `/u01/app/oracle/mysql/dbbackup`; each daily backup is completely self contained under `.../dbbackup/AutoFullBackup-<date>-<time>`. The daily backups are the key to a successful recovery and essentially the only data that really needs to be backed up and saved to a secure location. Everything else about the Oracle VM Manager can be reinstalled easily from scratch, but the data contained in the individual backups is critical for a full and painless recovery. Each backup is completely self contained and includes all data and transaction logs needed to accomplish a full restore.
- MySQL backup product – The automated daily backups are accomplished using MySQL Enterprise Backup which resides in `/opt/mysql/meb-3.8`. You may or may not want to back up this product depending on your approach to backups since this is always installed whenever the Oracle VM Manager is installed/reinstalled and does not contain any user customizable information.

Change Default Location for Automated Backups

Everything under `/u01` is normally deleted whenever you uninstall Oracle VM Manager using the `runInstaller.sh` script. Changing the location for the backup directory will protect the daily backups from being inadvertently deleted if Oracle VM Manager is uninstalled and then reinstalled again.

How to change the default location

To change the default location for the automated backups, simply edit the `/etc/sysconfig/ovmm` file and change **DBBACKUP** to use a new path of your choosing. For the purpose of this document we show the changed location in all of the examples as **/ovmm-backups**. Please devise a directory name/location that best fits your unique requirements; ensure that the new path you choose is not contained within `/u01`.

Using external storage for default location

Oracle recommends that MySQL Enterprise Backup write all backups to an NFS export mounted to the default backup directory specified in the `/etc/sysconfig/ovmm` file. Figure 17 below illustrates the concept in more detail. If you decide to use such a solution, then the file system should be sized to accommodate a rotation of fifteen daily backups, plus some space for additional growth and periodic ad hoc backups.

To size the file system, simply use the following formula:

```
dbsize * growth * ([15 daily backups] + [ad hoc backups]) = file system size
```

For example, if your database is 600 Megabytes and you anticipate the environment might grow twice in size over the next two years. Let's also assume you need some space to run ad hoc backups just before you upgrade the manager or make other changes to the environment – perhaps enough space to accommodate five ad hoc backups.

You would then use that information to calculate the file system size of a 600 Megabyte database as follows: $600M * 2 * (15 + 5) = 24,000M$. So, you would need a 24 Gigabyte file system.

A quick and easy way of determining the size of the database is to use the following command:

```
[root@mymanager ~]# du -hs /u01/app/oracle/mysql/data
600M /u01/app/oracle/mysql/data
```

Suggested Backup Architectures

The deployment architecture you implement for Oracle VM Manager drives the backup and recovery solution. So, the way you install Oracle VM Manager is an important consideration when designing a backup and recovery plan. The following examples of popular deployment architectures for Oracle VM Manager are explained in much more detail throughout the remainder of Part 2.

- Oracle VM Manager on a physical server
- Oracle VM Manager on a physical server using NFS
- Oracle VM Manager as a guest appliance using MySQL
- Oracle VM Manager as a guest using Oracle database

Keep in mind as you are reading that each backup and recovery solution associated with a deployment architecture is simply a suggestion. The various backup solutions noted above are not exhaustive and only examples of how your company might best deploy your Oracle VM Manager; perhaps they will give you ideas for a different solution that better fits your requirements. The important point to keep in mind is all of the following backup solutions can be adapted in many ways to better fit your particular needs - let your skills and imagination guide you.

All of the backup solutions for Oracle VM Manager discussed in Part 2 are predicated on just capturing the schema and data associated with the Oracle VM database. Everything else about the host operating system and Oracle VM Manager can be reinstalled quite easily. However, your situation may differ so please feel free to adapt any of the solutions to something that better fits your needs.

For example, you may have enabled the HTTP server on the Oracle VM management server to act as a YUM repository, or perhaps a repository for staging Oracle VM Templates, ISO images or Assemblies for importing into your Oracle VM environment. Perhaps you are using the Oracle VM management server as a KickStart server or any other number of roles or uses. In these cases, you would also want to include all of those directories and files as part of a backup solution.

Building a Multitier Backup Scheme

All backup solutions for Oracle VM Manager use a multitier approach to backups and recovery. This provides a more granular approach, allowing the operating system and Oracle VM Manager to be recovered independently of each other. A multitier approach to backing up and recovering Oracle VM Manager simply means that several different types of data protection and backups are incorporated into a single backup plan to handle different types of failures or catastrophic events.

For example, the operating system is protected and backed up independently of Oracle VM Manager. If you lose the physical server where Oracle VM Manager is running, then you will restore the operating system first. If you lose or corrupt just the Oracle VM database, then you will restore only the Oracle VM database. To accomplish this, backups are divided into two major categories. You will need to devise a backup scheme for the operating system and then combine that with a backup scheme for the Oracle VM database. The remainder of Part 2 is devoted to the following two topics to help accomplish the task of creating a full spectrum, multitier backup plan for Oracle VM Manager:

- Devising a multitier backup scheme for the operating system
- Devising a multitier backup scheme for Oracle VM database

Combine the OS and Database Backup Schemes

You will need to read and understand solutions for both the operating system and the Oracle VM database in order to build a complete backup solution. Once you have determined which two solutions best fit your needs, combine or adapt a backup scheme for the operating system with a backup scheme for the Oracle VM database into single, cohesive backup and recovery plan.

Devise a Multitier Backup Scheme for the Operating System

All of the backup solutions for Oracle VM Manager use a multitier approach to backups and recovery. The operating system is protected independent of Oracle VM Manager. If you lose the physical server where Oracle VM Manager is running, then you will restore the operating system first. If you lose or corrupt just the Oracle VM database, then you will restore only the Oracle VM database.

It is very important to incorporate one of the following multitier backup schemes for protecting the operating system where the Oracle VM Manager is installed and running.

Physical Server Backup using Hardware RAID

A lot of data centers will adopt the practice of protecting the Oracle VM management server with a multilevel backup scheme using hardware RAID. This solution incorporates a complete protection and business continuance in event of a catastrophic disk failure or operating system corruption. This solution relies on hardware RAID, an offline copy of the OS and finally a tape backup in case the first two levels of OS protection fail. Figure 12 below shows the basic premise for the solution with three levels of protection for the operating system.

Ensure you have or order a physical server equipped with the following features:

- The server should have an embedded RAID capable SCSI controller for the internal disks

- The server should come with four internal disks of equal size ranging anywhere between 36 to 172 Gigabytes in size. Any more than that will be overkill since additional needed space should really come from NAS or SAN.

Use the four physical disks to create two different logical volumes using RAID 1 in order to devise a hardware RAID solution similar to the one shown below.

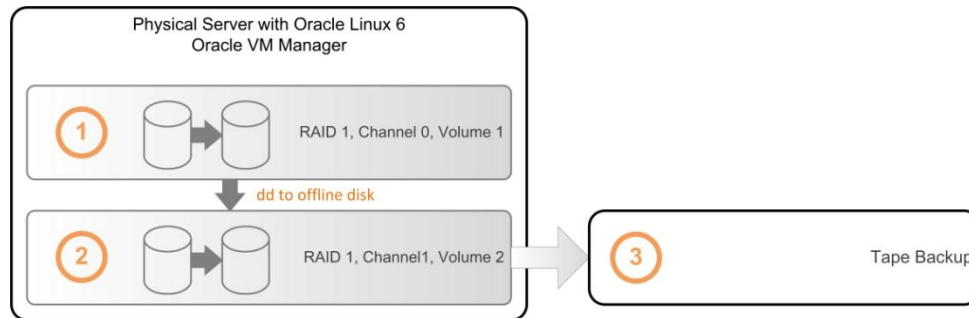


Figure 12: Create a multitier protection scheme using hardware RAID

Referring to Figure 12 above the first level of backup for the operating system uses hardware RAID to protect from a failed disk. In this case, the first two of four disks are incorporated into a RAID 1 mirrored logical volume (this refers to hardware level logical volume, not LVM). If the primary disk fails, the hardware RAID continues using the second, mirrored disk without interruption. In most cases, modern RAID controllers allow hot swapping the failed disk for a new disk.

The second level of backup utilizes the Linux `dd` command to create an offline copy of the primary mirrored disk. In this case, the remaining two of four disks are incorporated into a RAID 1 mirrored logical volume (this refers to hardware level logical volume, not LVM). This is an offline copy of the primary boot disk to protect from unintentional human error. The second pair of mirrored disks has the same advantage as the first pair since the boot disk is protected from a disk failure, but it has the added advantage of not being subject to system administrator mistakes being propagated immediately.

Normally, a daily cron job is created to automatically execute a synchronization script that accomplishes the image copy of the primary mirrored boot disk to the offline mirrored disk, creating any new partitions that may have been created since the last backup. Keep in mind that a compromised OS can still be propagated to the second set of disks if the error is not caught until after the cron job has run the synchronization script. This is the reason for having the third level of the backup solution.

The third level of backup utilizes a tape backup of the second set of offline disks before the daily cron job is called to execute the disk synchronization. It is advisable to make the daily tape backup a full backup rather than an incremental. In this case the retention windows for OS level backups should be sufficient to recover from week or two running on a compromised OS.

Physical Server Backup using KickStart

A few data centers will adopt the practice of protecting the Oracle VM management server using a multilevel backup scheme incorporating both hardware RAID and KickStart. This solution incorporates a complete automated reinstall of the operating system in case of catastrophic disk failure or operating system corruption. This solution relies on a complete reinstall of the operating system

using KickStart with custom post install scripts to rebuild the server in a matter of minutes. Figure 13 below shows the basic premise of the solution with two levels of protection for the operating system.

Ensure you have or order a physical server equipped with the following capabilities:

- The server should have an embedded RAID capable SCSI controller for the internal disks
- The server should come with two internal disks of equal size ranging anywhere between 36 to 172 Gigabytes in size. Any more than that will be overkill since additional needed space should really come from NAS or SAN.

Use the two physical disks to create a single logical volumes using RAID 1 in order to devise a hardware RAID solution similar to the one shown below.




Figure 13: Create a multilevel protection scheme using KickStart

Referring to Figure 13 above the first level of backup for the operating system uses hardware RAID to protect from a failed disk. In this case, the only two internal disks are incorporated into a RAID 1 mirrored logical volume (this refers to hardware level logical volume, not LVM). If the primary disk fails, the hardware RAID continues using the second, mirrored disk without interruption. In most cases, modern RAID controllers allow hot swapping the failed disk for a new disk.

Hardware RAID does not protect the operating system from unintentional human error that might render the OS useless; whatever error is made on the primary disk is immediately propagated to the mirrored disk. In this case, the only recourse to get the server up and running again is to boot the server using PXE to initiate the KickStart process via TFTP.

The second level of backup is not really protection, but rather a reinstall of the entire OS using KickStart. This method has the advantage of relatively quick rebuild of the Oracle VM server without needing to perform a backup of any type. The disadvantage of this approach is the fact that you will need to understand how to implement and manage a KickStart server, plus take the time to create and maintain custom install scripts if you want to automation the reinstallation of Oracle VM Manager after KickStart has completed the OS install.

Physical Server Backup using SAN Boot

SAN booting is a much less common way of protecting the Oracle VM management server using a multilevel level backup scheme incorporating enterprise class storage arrays for diskless booting of servers. This solution relies on the features inherent in storage arrays with Fibre Channel capabilities supporting SAN boot.

Although a fantastically flexible and powerful architecture, SAN boot can be quite quirky and problematic at times. The initial configuration and implementation of a SAN boot solution is daunting and requires someone with very advanced storage and storage networking skills and experience. Figure

14 below shows the basic premise of the solution with three levels of protection for the operating system.

Implementing SAN boot is beyond the scope of this document, but ensure you have or order a physical server equipped with the following capabilities:

- The server will have no internal disks in the chassis
- The server should have at least one dual port Fibre Channel host bus adapter. The Fibre Channel adapter comes with bios capable of fooling a server into thinking that it is being presented a local disk housed in the server, but in reality the “disk” is a sparse file residing on a storage array and containing a bootable operating system and presented to the physical server as a block level physical disk – just like a boot image for a virtual machine, but in this case you are booting a physical server instead of a virtual machine. Both Emulex and Qlogic Fibre Channel adapters are both supported.

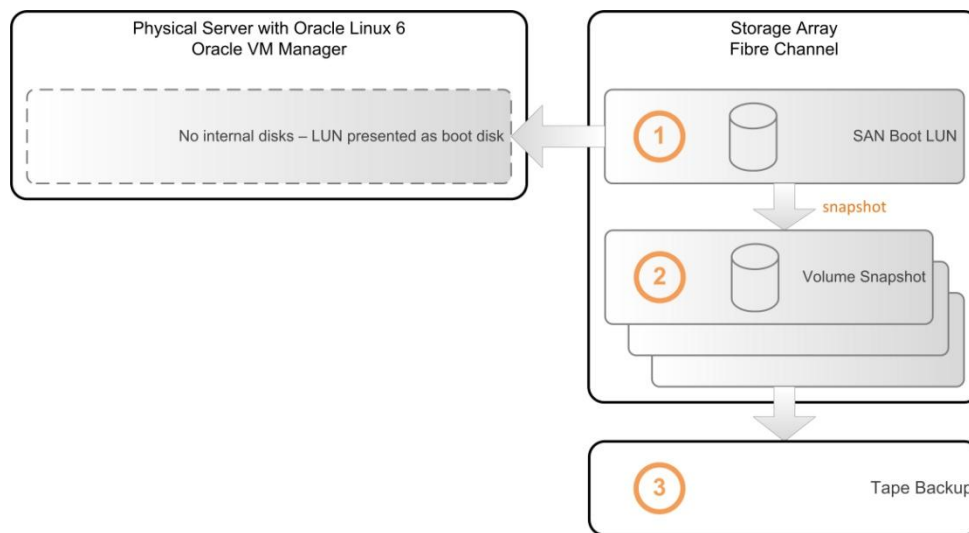


Figure 14: Create a multilevel protection scheme using SAN boot

Referring to Figure 14 above the first level of protection is the SAN boot LUN itself. Enterprise storage arrays normally employed for SAN include an impressive collection of hardware RAID and automated features to protect from failed disks. Some storage arrays can lose up to two disks per volume, incorporating spare disk technology that automatically enables hot spares to rebuild parity data to allow the array to continue serving data without worrying about replacing failed disks right away.

In this case, the a SAN boot LUN is created within a volume on a storage array and presented to the physical server as if it were a local disk. As noted above, the tolerance for failed disks on an enterprise class storage array is impressive.

The second level of protection comes from the capabilities of snapshot software that is inherent in most storage arrays. Snapshots will allow you to recovery from unintended errors or compromises to a running operating system. Snapshots are usually completed in seconds and can be used to restore an operating system to a point-in-time in minutes. Snapshots can quickly fill space on the storage array as the differences between the original volume and the snapshot continue to grow; this leads to the third level of protection.

The third level of protection with this solution incorporates backup to tape normally using NDMP for serverless, LAN free backups. The snapshots can be taken to tape periodically which captures 100% of the data associated with a volume including the SAN boot LUNs – it automatically combines the data from the original volume as well as the changed data contained in the snapshot to form a full backup. The snapshot can be deleted once the data from a snapshot has been captured on tape. Tape allows for retention windows of snapshots to be measured in weeks, months and years.

Oracle VM Manager as Guest VM using SAN or NAS

This solution should not be confused with backup solutions related to backing up Oracle VM Guests covered in Part 6 of this guide. Installing the Oracle VM Manager on a guest operating system is probably one of the most robust, flexible and easiest deployment architectures to backup and recovery. In this case; the Oracle VM Manager itself is installed on a virtual machine running Oracle Linux. Figure 15 below shows the basic premise of the solution with three levels of protection for both the operating system and Oracle VM Manager.

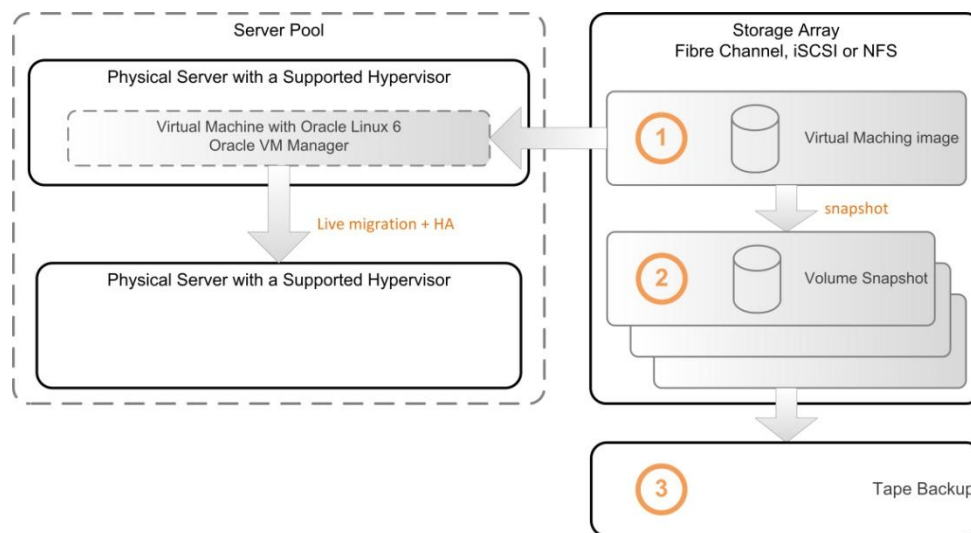


Figure 15: Create a multilevel protection scheme using SAN or NAS for a virtual machine

The operating system has the same multitier backup scheme as the SAN boot solution shown in Figure 14 above but with additional benefits not available to any of the other server backup schemes. The virtual machine hosting the Oracle VM Manager can take advantage of any high availability features the virtual server product has to offer: including:

- Automated restart of the virtual machine if the physical server fails
- Live migration of the virtual machine
- Portability from one virtual server environment to another

One additional advantage to this solution is the fact that both the operating system and the Oracle VM Manager are protected with a single backup operation. So, there is no need to develop another backup scheme for the Oracle VM database – it is all accomplish in one solution. Please refer to the backup solution for running **Oracle VM Manager as a guest appliance** for additional specific information about accomplishing a backup and recovery.

There are some limitations to supported virtual server products, but running an Oracle VM Manager as a guest using either Oracle VM 2 or Oracle VM 3 is fully supported. Please refer to Oracle Linux support policies to fully understand what virtual server products are supported and what level of support can be expected. Keep in mind that installing Oracle VM Manager on a virtual machine running in an Oracle VirtualBox environment is supported for non-production learning or proof of concept projects only.

Devise a Multitier Backup Scheme for the Manager Database

Once you have decided on a backup scheme for the physical, bare metal server that will host Oracle VM Manager, then you will need to devise a backup scheme for the Oracle VM Manager database. The following four backup solutions are simply starting points for building a multitier backup solution for the database. Please feel free to adapt any of the following solutions to better fit your unique requirements.

Manager Backup on Physical Server with Internal Storage

The simplest method for deploying Oracle VM Manager is to install Oracle Linux on a standalone physical server then install the Oracle VM Manager application on local disk within the server chassis. However, this is probably the least robust solution for backup and recovery since it relies on traditional file level backups and restores. Figure 16 below shows the basic premise of the solution with two levels of protection for the Oracle VM database.




Figure 16: Using file level backup to capture key Oracle VM files and directories

The first level of protection for the Oracle VM database is accomplished by the automated daily backups performed by MySQL Enterprise Backup. Figure 16 above shows that the automated backups are being saved in a custom location arbitrarily named /ovmm-backups. Remember that MySQL Enterprise Backup is creating daily, full hot backups of the database.

If you have installed Oracle VM Manager using the Oracle Standard or Enterprise edition database, then you will need to have your DBA help devise a custom automated daily backup scheme, saving the dumps or exports to a location other than /u01 on the Oracle VM management server. You will need to quiesce Oracle VM Manager before taking a backup of the database if you or your DBA are not able to devise a hot backup scheme.

The second level of protection for the Oracle VM database is to simply take the daily backups to tape. You could use Oracle Secure Backup or any other vendor software product capable of writing to tape. You can also use any other tool at your disposal including cpio, dump, tar etc. to copy the daily backups to an NFS server.

Manager Backup on Physical Server with Shared Storage

This is a variation of installing Oracle VM Manager on a physical server. This solution relies on critical files and directories related to the automated backups being contained on a storage array and presented to the Oracle VM management server. Figure 17 below shows the basic premise of the solution with four levels of protection for the Oracle VM database.

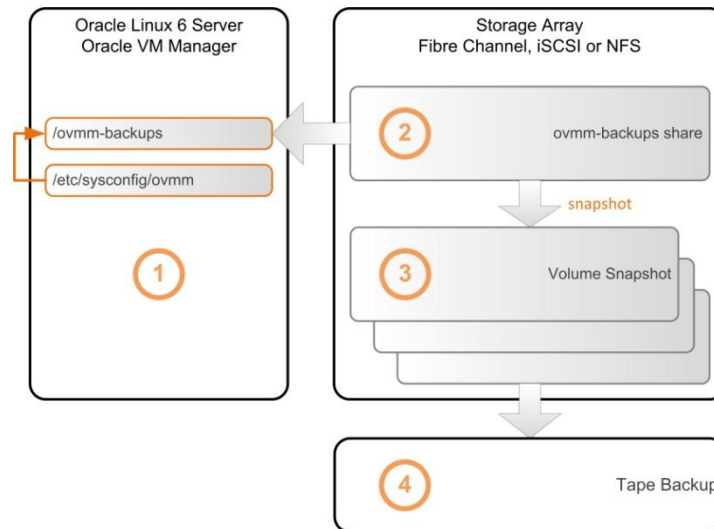


Figure 17: Using image level backup to capture automated backups of the Oracle VM database

Once again we are not attempting to protect or capture the Oracle VM Manager binaries since they will be captured during the operating system backup. The goal of this solution is simply to capture the essential data and file objects needed to recover the Oracle VM database.

The real advantage of this solution is that both the operating system and everything associated with Oracle VM Manager are protected at the same time with a single snapshot. For example, you would simply stop the virtual machine where Oracle VM is running, take the snapshot and then restart the virtual machine. This entire backup process would take a few minutes to accomplish. Recovering either the operating system or Oracle VM or both can be accomplished in a few minutes with very little trouble other than stopping the virtual machine, restoring the snapshot and starting the virtual machine again.

Another advantage to this solution is the fact that the virtual machine running Oracle VM can take advantage of all high availability features the virtual server product has to offer including live migration and automatic restarts of the virtual server if the physical server it is running on fails for any reason.

Install the Oracle VM Manager using the **simple** install option to ensure it is using MySQL for the database engine. MySQL is supported with small to very large production deployments of Oracle VM and makes the Manager component almost a hands-off appliance – administration becomes very simple and painless when using the MySQL database. Recovery of the MySQL database is very quick and simple and can all be done with a single command by the systems administrator rather than involving application and database support resources. Using MySQL for the database engine also allows the systems administrator to perform very quick and simple ad hoc backups before performing upgrades or making changes to the environment.

Of course, the same thing can be accomplished when using Oracle SE or EE as the database engine if it is installed locally on the same virtual machine as the Oracle VM Manager. The decision is really up to you and IT governance policies in your data center.

The first level of protection comes from the MySQL automated backups on the Oracle VM management server. In this case, the backups will be written to an NFS export residing on an external storage array. Figure 17 above shows that you will need to create a mount point named /ovmm-backups and then mount an NFS export to that directory. The reason for changing the location of default backups is explained in subsection **Change Default Location for Automated Backups** above.

Although the /etc/sysconfig/ovmm file should already be captured along with the operating system level backup you have devised, you might want also capture this file along with the MySQL backups. You would have to create a cron job to copy the file to the /ovmm-backup directory to ensure the latest version is always captured.

The second level of protection is provided by the storage array. Enterprise storage arrays normally employed for SAN or NAS include an impressive collection of hardware RAID and automated features to protect from failed disks. Some storage arrays can lose up to two disks per volume, incorporating spare disk technology that automatically enables hot spares to rebuild parity data to allow the array to continue serving data without worrying about replacing failed disks right away.

The third level of protection comes from the capabilities of snapshot software that is inherent in most storage arrays. Snapshots will allow you to recovery from unintended errors or compromises to a running operating system. Snapshots are usually completed in seconds and can be used to restore an operating system to a point-in-time in minutes. Snapshots can quickly fill space on the storage array as the differences between the original volume and the snapshot continue you to grow; this leads to the fourth level of protection.

The fourth level of protection with this solution incorporates backup to tape normally using NDMP for serverless, LAN free backups. The snapshots can be taken to tape periodically which captures 100% of the data associated with a volume including the SAN boot LUNs – it automatically combines the data from the original volume as well as the changed data contained in the snapshot to form a full backup. The snapshot can be deleted once the data from a snapshot has been captured on tape. Tape allows for retention windows of snapshots to be measured in weeks, months and years.

Manager Backup as a Guest using MySQL Database

This solution builds on the operating system level backup solution explained above in the section titled **Oracle VM Manager as Guest VM using SAN or NAS**. The Oracle VM Manager is installed on a virtual machine disk image being presented from an enterprise class storage array using NFS, iSCSI or FCP providing four levels of protection for the operating system, Oracle VM Manager and the Oracle VM database. This is a highly flexible and robust solution offering quick and easy backup and recovery of every aspect of the Oracle VM Manager. Figure 18 below shows the basic premise of the solution with four levels of protection for the operating system as well as the Oracle VM Manager binaries and database.

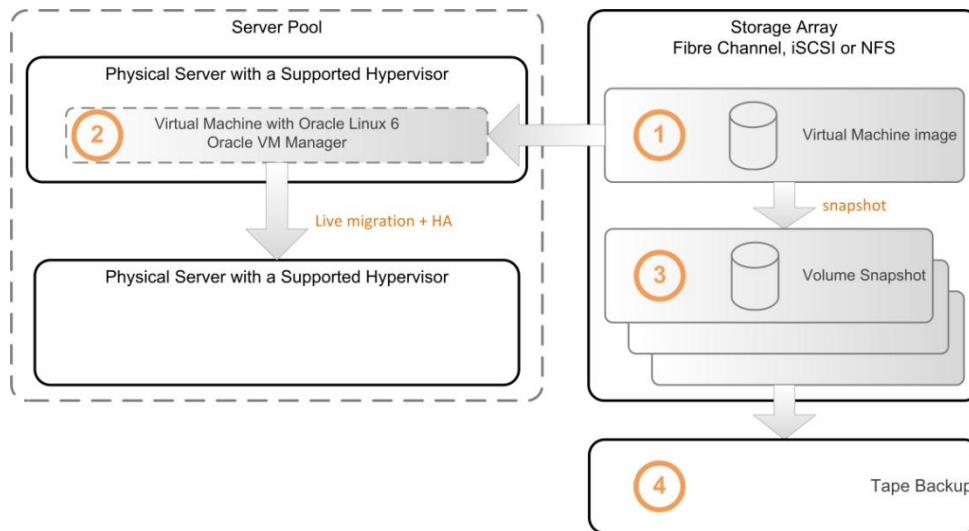


Figure 18: Using image level backup to capture entire virtual machine hosting Oracle VM Manager

The real advantage of this solution is that both the operating system and everything associated with Oracle VM Manager are protected at the same time with a single snapshot. For example, you would simply stop the virtual machine where Oracle VM is running, take the snapshot and then restart the virtual machine. This entire backup process would take a few minutes to accomplish. Recovering either the operating system or Oracle VM or both can be accomplished in a few minutes with very little trouble other than stopping the virtual machine, restoring the snapshot and starting the virtual machine again.

Another advantage to this solution is the fact that the virtual machine running Oracle VM can take advantage of all high availability features the virtual server product has to offer including live migration and automatic restarts of the virtual server if the physical server it is running on fails for any reason.

Install the Oracle VM Manager using the **simple** install option to ensure it is using MySQL for the database engine. MySQL is supported with small to very large production deployments of Oracle VM and makes the Manager component almost a hands-off appliance – administration becomes very simple and painless when using the MySQL database. Recovery of the MySQL database is very quick and simple and can all be done with a single command by the systems administrator rather than involving application and database support resources. Using MySQL for the database engine also allows the systems administrator to perform very quick and simple ad hoc backups before performing upgrades or making changes to the environment.

Of course, the same thing can be accomplished when using Oracle SE or EE as the database engine if it is installed locally on the same virtual machine as the Oracle VM Manager. The decision is really up to you and IT governance policies in your data center. Please refer to [Oracle VM 3 Installation and Upgrade Guide](#) for more detail about installation options for Oracle VM Manager.

The first level of protection is provided by the storage array since the virtual machine hosting Oracle VM Manager is being presented via NFS (iSCSI and FCP are also supported). Enterprise storage arrays normally employed for SAN or NAS include an impressive collection of hardware RAID and automated features to protect from failed disks. Some storage arrays can lose up to two disks per

volume, incorporating spare disk technology that automatically enables hot spares to rebuild parity data to allow the array to continue serving data without worrying about replacing failed disks right away.

The second level of protection comes from the MySQL automated backups on the Oracle VM management server. In this case, the backups will be written to local directory named /ovmm-backups. There is no need to mount an NFS export to this directory in this case since everything about the virtual machine image is already contained on a storage array. But, you will still want to change the default backup directory to ensure the database backups are not inadvertently removed if someone reinstalls Oracle VM Manager. Please read about changing the default location and why it should be changed in section **Change Default Location for Automated Backups** above.

The third level of protection comes from the capabilities of snapshot software that is inherent in most storage arrays. Snapshots will allow you to recovery from unintended errors or compromises to a running operating system. Snapshots are usually completed in seconds and can be used to restore an operating system to a point-in-time in minutes. Snapshots can quickly fill space on the storage array as the differences between the original volume and the snapshot continue you to grow; this leads to the fourth level of protection.

The fourth level of protection with this solution incorporates backup to tape normally using NDMP for serverless, LAN free backups. The snapshots can be taken to tape periodically which captures 100% of the data associated with a volume including the SAN boot LUNs – it automatically combines the data from the original volume as well as the changed data contained in the snapshot to form a full backup. The snapshot can be deleted once the data from a snapshot has been captured on tape. Tape allows for retention windows of snapshots to be measured in weeks, months and years.

Oracle VM Manager Backup as a Guest using an Oracle Database

This solution is a slight variation on the preceding solution. The Oracle VM Manager is installed on a virtual machine disk image and the database server resides on a separate virtual machine both being presented from an enterprise class storage array using NFS, iSCSI or FCP. This solution provides four levels of protection for the operating systems, the Oracle VM Manager and the Oracle VM database. Although this is a little more complex and not bundled into a single “appliance”, it is still a highly flexible and robust solution offering quick and easy backup and recovery of every aspect of the Oracle VM Manager. Figure 19 below shows the basic premise of the solution with four levels of protection for the operating system as well as the Oracle VM Manager binaries and database.

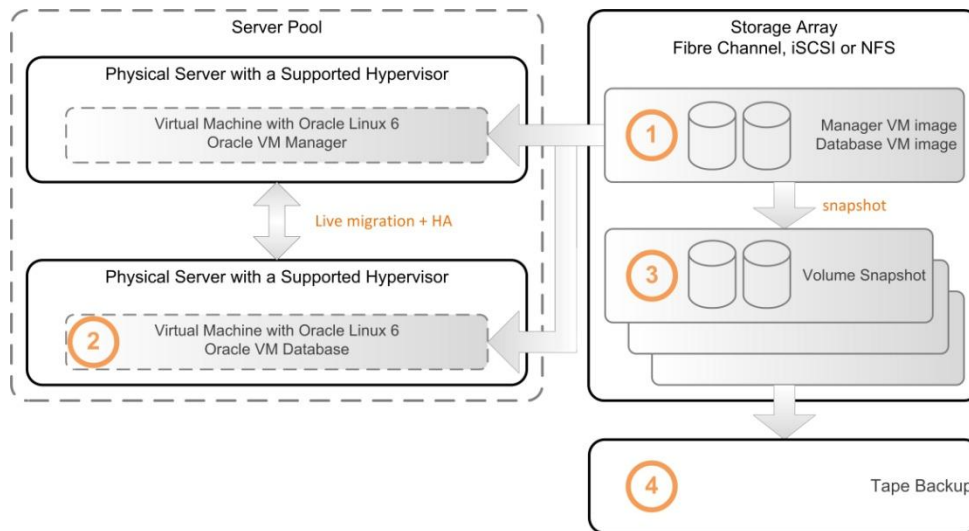


Figure 19: Using image level backup to capture entire virtual machines hosting Oracle VM Manager & database

This solution simply adds another virtual machine where an Oracle SE or EE database is installed remote to the virtual machine running Oracle VM. The above illustration shows the Oracle database server as another virtual machine in the same server pool, but in fact the database server could be any remote Oracle database including a RAC implementation running on any other servers in the data center.

The real advantage of this solution is that both the operating system and everything associated with Oracle VM Manager are protected at the same time with a single snapshot. For example, you would simply stop the two virtual machines where Oracle VM Manager and database are running, take the snapshot and then restart the virtual machines. This entire backup process would take a few minutes to accomplish. Recovering either the operating system or Oracle VM or both can be accomplished in a few minutes with very little trouble other than stopping the virtual machines, restoring the snapshot and starting the virtual machines again.

Install the Oracle VM Manager using the **advanced** install option – this will allow you to provide all of the connection information including host, listener port and SID for the database engine. The Oracle database is supported with small to very large production deployments of Oracle. Depending on your data center policies implementing and maintaining the Oracle database engine may necessitate involving application and database support resources. Please refer to [Oracle VM 3 Installation and Upgrade Guide](#) for more detail about installation options for Oracle VM Manager.

The first level of protection is provided by the storage array since the virtual machine hosting Oracle VM Manager is being presented via NFS (iSCSI and FCP are also supported). Enterprise storage arrays normally employed for SAN or NAS include an impressive collection of hardware RAID and automated features to protect from failed disks. Some storage arrays can lose up to two disks per volume, incorporating spare disk technology that automatically enables hot spares to rebuild parity data to allow the array to continue serving data without worrying about replacing failed disks right away.

The second level of protection would come from automating daily backups of the Oracle database. Oracle VM does not come with any tools to configure or set up automated dumps/exports of the Oracle VM database, so this is something you and your DBA would have to devise and script.

If you devise your own automated hot database backup then it should probably do something similar to the MySQL Enterprise Backup by writing to local directory name/location of your choice. There is no need to mount an NFS export to this directory in this case since everything about the virtual machine image is already contained on a storage array.

The third level of protection comes from the capabilities of snapshot software that is inherent in most storage arrays. Snapshots will allow you to recovery from unintended errors or compromises to a running operating system. Snapshots are usually completed in seconds and can be used to restore an operating system to a point-in-time in minutes. Snapshots can quickly fill space on the storage array as the differences between the original volume and the snapshot continue you to grow; this leads to the fourth level of protection.

The fourth level of protection with this solution incorporates backup to tape normally using NDMP for serverless, LAN free backups. The snapshots can be taken to tape periodically which captures 100% of the data associated with a volume including the SAN boot LUNs – it automatically combines the data from the original volume as well as the changed data contained in the snapshot to form a full backup. The snapshot can be deleted once the data from a snapshot has been captured on tape. Tape allows for retention windows of snapshots to be measured in weeks, months and years.

Backup Solutions

No matter which backup architecture you follow, the first level backup should always be a full backup of the MySQL or Oracle database.

How to Create an Ad Hoc Backup using MySQL

Oracle VM using MySQL as the database engine makes it exceedingly simple to initiate an ad hoc backup of the database. An ad hoc backup simply means you are taking backup that is outside the normal regularly scheduled backups being performed on a daily basis by MySQL Enterprise Backup.

You might create a backup of the database just before you begin an upgrade of the Oracle VM Manager. Perhaps you going to make some changes to the Oracle VM network infrastructure or try deleting a server pool and need a way to back out in case things go awry.

The Oracle VM product comes with a script for generating backups on the fly when you chose the **simple** install option when installing Oracle VM Manager. This is the same script that is called automatically on a daily basis by Oracle VM, so this can be executed while the Oracle VM Manager and database server are running to create a full, hot backup.

You should always read the [Oracle VM 3 Installation and Upgrade Guide](#) for the latest information about performing backups and recoveries for the Oracle VM Manager. As of this writing Appendix B in the Installation and Upgrade Guide contains all the information about performing recoveries.

Create a Full Backup

Assuming you have changed the default location of the backup directory defined in the `/etc/sysconfig/ovmm` file to `/ovmm-backups`, you would execute the following command on the Oracle VM management server:

The command below uses a custom backup name of **adhoc-backup** as an example - you should provide a custom name of your choosing; the backup script will automatically append the date and time to the file name `adhoc-backup-<date_time>`. If you don't provide a custom name for the backup, the script will automatically use: `AutoFullBackup-<date_time>`. The script will also prompt you for a user name and password if you don't provide one on the command line which will be a little more secure since the username/password will not show up in the process table while the script is executing.

```
[root@mymanager ~]# cd /u01/app/oracle/ovm-manager-3/bin
[root@mymanager ~]# ./createBackup.sh -n adhoc-backup

Please enter the Oracle VM manager user name: Admin

Please enter the Oracle VM manager user password:

Backing up the Oracle VM Manager MySQL Database...

INFO: Successfully backed up database as adhoc-backup-20130819_103004
[root@mymanager ~]#
```

The full, hot backup is complete at this point.

Recovery Solutions

No matter which backup architecture you follow, the first level backup should always be a full backup of the MySQL or Oracle database. This means you will be restoring a file level copy of the backup data created by the backup process.

How to recover an Ad Hoc Backup using MySQL

Restoring the Oracle VM database using an automatic or ad hoc backup is also exceedingly simple and quick.

Restore the Full Backup

Assuming you have changed the default location of the backup directory defined in the `/etc/sysconfig/ovmm` file to `/ovmm-backups`, you would execute the following command on the Oracle VM management server:

```
[root@ mymanager ~]# service ovmm stop; service ovmm_mysql stop
[root@ mymanager ~]# su - oracle
```

```

$ cd /u01/app/oracle/ovm-manager-3/ovm_shell/tools
$ ./RestoreDatabase.sh adhoc-backup-20130819_103004
INFO: Expanding the backup image...
INFO: Applying logs to the backup snapshot...
INFO: Restoring the backup...
INFO: Success - Done!
INFO: Log of operations performed is available at:
      /ovmm-backups/adhoc-backup-20130819_103004/Restore.log

IMPORTANT:

As 'root', please start the OVM Manager database and application using:
service ovmm_mysql start; service ovmm start

$ exit
[root@ mymanager ~]# service ovmm_mysql start; service ovmm start

```

Refresh All Objects after Restore

You must refresh all objects after the database restore has completed. Simply log into the Oracle VM Manager after it has restarted, select the **Server Pools** folder in the navigation pane and then choose the **Refresh All** icon from the toolbar in the management pane.




Figure 20: Screen shot showing how to refresh all objects in Oracle VM Manager

You should now have a fully functioning Oracle VM Manager with all objects restored, including any custom simple names you originally added to various storage objects.

Conclusion

Part 2 of this guide has discussed separate multitier backup schemes for both the operating system and the Oracle VM database. You now need to combine or adapt a backup scheme for the operating system with a backup scheme for the Oracle VM database into single, cohesive backup and recovery plan.

Part 3: Backup & Recovery for Oracle VM Guests and Resources

Overview

Oracle VM related objects such as assemblies, ISO images and Oracle VM Templates are known collectively as Oracle VM Guest Resources. Each type of image performs a supporting role in the way you deploy Oracle VM Guests. You may or may not use any Oracle VM Guest Resources depending on how you deploy virtual machines.

For example, if you create Oracle VM virtual machines without using an Oracle VM Template and then use KickStart to install the guest operating system, then you probably won't have any templates or ISO images. If you create Oracle VM virtual machines without using an Oracle VM Template, but boot an ISO image in the repository to install the guest operating system, then you might have a collection of ISO images, but no templates. So the type and extent of guest resources is completely dependent upon the way you create and deploy Oracle VM Guests.

Critical Files & Directories for Oracle VM Guests and Resources

Of course, Oracle recommends you perform image level backups of the repository objects on the storage array, but for file level backups simply capture all files and subdirectories under the following directories on any server to capture assemblies, ISO images, templates, or virtual machines:

- /OVS/Repositories/<UUID>/Assemblies. This subdirectory contains a subdirectory for each downloaded assembly. Each assembly subdirectory includes the original unprocessed assembly file plus another directory containing the exploded contents of the unprocessed assembly file.
- /OVS/Repositories/<UUID>/ISOs. The ISO images are completely self-contained in this subdirectory.
- /OVS/Repositories/<UUID>/Templates. This subdirectory contains a subdirectory for each downloaded template. Each template subdirectory contains only the vm.cfg file for the template; the virtual disk images are in a common directory with all other virtual disks.
- /OVS/Repositories/<UUID>/VirtualDisks. This directory contains all virtual disks associated with templates and assemblies. It also contains all virtual disks associated with virtual machines, so if you choose not to follow the best practice of segregating guest resources from virtual machine images as suggested in the next section, then you will also be capturing the virtual disk images for Oracle VM Guests which will be useless if the virtual machines are running when you take this backup.
- /OVS/Repositories/<UUID>/VirtualMachines/<VirtualMachineUUID>/vm.cfg. This directory contains the virtual machine configuration file.

Backup Frequency

If you are performing regular backups of entire storage repositories, then the guest resources are captured along with everything else - you might not need a specialized backup for guest resources. However, if you adopt the best practice of segregating guest resources as suggested below in the section titled [Segregating Guest Resources into a Common Repository](#), then you will need to develop a specialized backup of guest resources.

Oracle VM Guest resources are relatively static once assemblies, ISO images and template are imported and customized for use in the environment. However, templates do change periodically as people create other templates based on running virtual machines or existing templates are cloned and customized for other purposes; assemblies get updated and additional ISO images are downloaded.

The following list of suggested backup frequencies may helpful when adopting your own backup policies for guest resources:

- Daily backups – Daily backups are recommended during the initial implementation period while there is a lot of activity making changes to the Oracle VM environment. Generally, you will have a lot of people with varying skills and experience working with virtual machines, templates and ISO images. The potential for people inadvertently overwriting or destroying custom assemblies, ISO images and templates is very likely; you will want to protect against losing a lot of custom work that may have been put into customer
- Weekly backups – You might step down the frequency of backups once active creation of virtual machines and changes to the Oracle VM environment begin to die down.
- Monthly backups – You can probably adopt a policy of monthly backups once the Oracle VM environment has gone into production and strict change control policies are in effect.
- Ad hoc backups – You might want to adopt a policy of creating an ad hoc backup after any new assemblies and ISO are added or templates are added or created from running virtual machines.

Choosing a backup approach for Oracle VM guests

Except for the Oracle VM Guests, backup and recovery solutions for Oracle VM are relatively straight forward, quite easy to accomplish and are discussed in more detail in subsequent parts of this document. Let’s take a closer look at a couple approaches to backing up virtual machines.

Storage level backup

In this case, all Oracle VM related objects reside on a SAN array where all backups can occur using a single, consistent strategy such as remote or local serverless backups using protocols such as NDMP.

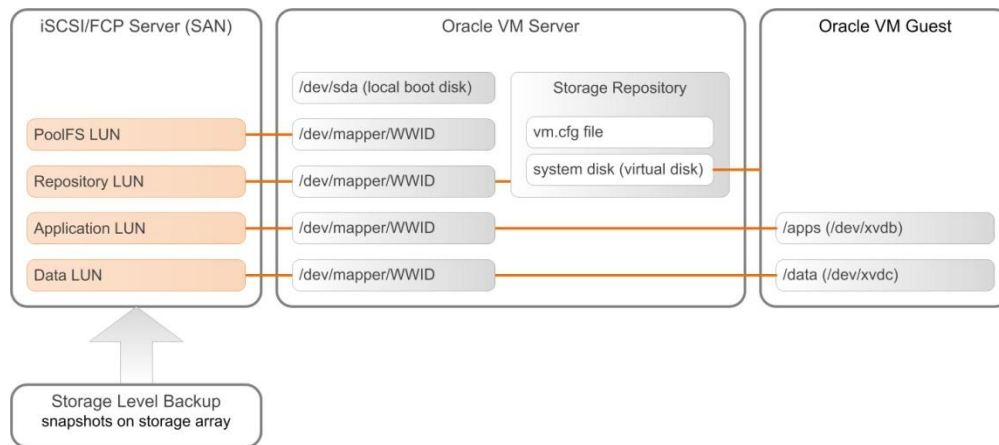


Figure 24: A SAN based deployment architecture where all Oracle VM related objects are backed up directly from the storage array

Serverless backup solutions can be very challenging to restore individual Oracle VM Guests. This is due to the fact that entire storage objects are captured in whole during a backup, but only a small subset normally needs to be restored. For example, a storage repository might contain hundreds of Oracle VM Guests that are all captured during a backup, but an individual business unit might only want to restore one or two Oracle VM Guests out of the hundreds of virtual machines that were captured during the backup.

In the event you are recovering from a disaster, then restoring all objects captured during a backup is fine. But in normal day-to-day operations you will most likely be restoring an individual Oracle VM Guest. An even more likely scenario is that an individual business unit may want to restore only a few specific files that reside on a guest operating system and not an entire virtual machine.

So, it is important to design a backup and recovery plan that includes many levels of backups, perhaps using different backup techniques or software to capture data in ways that allow for either wholesale or fine grained restoration depending on the needs of your user community. Perhaps snapshots of entire volumes on a storage array combined with tape backup of individual Oracle VM Guests at the level of the guest operating system like you would perform on normal bare metal servers.

Repository level backup

An Oracle VM Server can be configured to enable third party applications to perform a back up of the contents of a storage repository. To enable this, an Oracle VM Server is configured to provide an NFS share that a third party back up tool can use to access the contents of the repository. The Oracle VM Server must be in a clustered server pool and have the OCFS2-based storage repository presented to it. For further details, please refer to Oracle VM User Guide: [Enabling Storage Repository Back Ups](#).

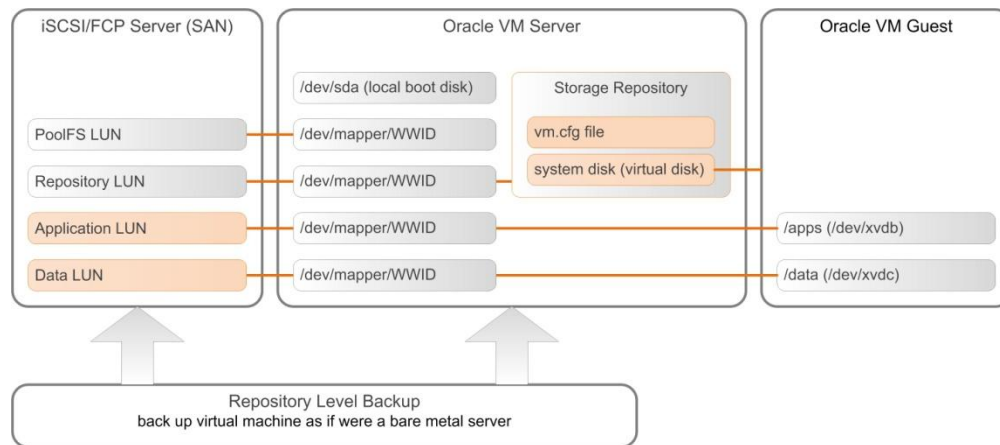


Figure 25: A SAN based deployment architecture where all Oracle VM related objects are backed up from the perspective of storage repository

When you have created a repository export, use the **Repository Path** (displayed in the management pane table) and the Oracle VM Server hostname or IP address to connect to the NFS mount point from the third party backup software.

For NFS-based repositories, the backup could be run from any server that can mount the NFS share. Since Oracle VM servers are primarily for running virtual machines, backup schedule should be managed accordingly without impacting the regular workloads.

When you perform guest VM backup at the repository level, please make sure that the virtual machines are not in any “powered on” state.

Oracle VM Guest OS level backup

To backup at the guest VM OS level, you can install standard backup agent (Oracle Secure Backup, NetBackup, etc.) into the virtual machine, and treat the guest VM as a regular host to perform backup and restore. In this case, the backup and restore is no different than any regular OS backup and restore.

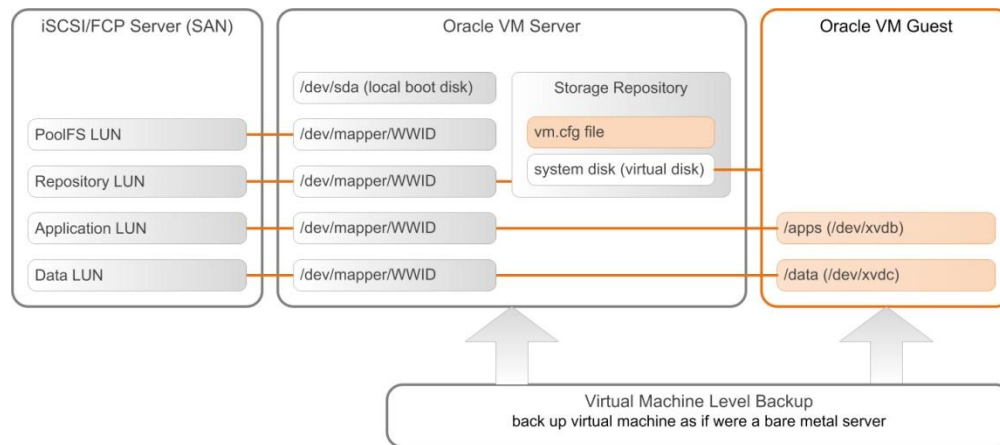


Figure 26: A SAN based deployment architecture where all Oracle VM related objects are backed up from the perspective of the Oracle VM Guest operating systems

Quiesce Activity during Backups

The backup solutions outlined do not require you to quiesce any applications or servers during the backup. However, there is always the chance that someone might be modifying a template, creating a new template or importing some other type of guest resource during the backup.

You shouldn't need to worry about activity during backups if your data center is only staffed during normal business hours. However, if you have a data center with a 24/7 operations model then you will need to devise some sort of policy to ensure modifications to the storage repository containing guest resources are not occurring during a backup. This will be challenging.

The most obvious way to quiesce activity during a backup is to stop the Oracle VM Manager and then restart it after the backup has completed. However, stopping the Oracle VM Manager will also have some impact on systems administrators in a data center with a 24/7 operations model that may be working on Oracle VM. So, if you choose to stop the Oracle VM Manager during a backup, then you will need to make sure this policy is understood by everyone to ensure no one spends time trying to figure out why the Manager stopped working while they were in the middle of importing an ISO image or template.

Segregating Guest Resources into a Common Repository

An Oracle recommended best practice is to segregate guest resources from storage repositories containing Oracle VM Guests. Each server pool is presented one or more storage repositories that are dedicated each server pool and would only contain Oracle VM Guests for a given server pool. A separate storage repository is then presented to all Oracle VM Servers in all server pools and would only contain guest resources that all server pools could access.



Figure 27: Screen shot showing both server pool repositories and a common template repository

Figure 27 above shows the basic premise for the solution. The storage repository containing only guest resources is named **Templates repo1** and is presented to all server pools. A storage repository spanning multiple server pools must be presented using NFS. This is due to the fact that each server pool belongs to a unique OCFS2 cluster and a storage object containing an OCFS2 file system cannot belong to more than one cluster.

The storage repositories named **Guests mypool1 repo1** and **Guests mypool2 repo1** each belong to a single server pool and only contain Oracle VM Guests images running on Oracle VM Servers within their respective pools. These repositories can use either NAS or SAN as the storage protocol since each repository only belongs to a single pool.

The Value of Segregating Guest Resources from Guest Images

The practice of segregating the type of objects provides several advantages in terms of saving space as well as backup and recovery:

- *Significantly reduce duplication* – Creating a centralized repository containing common guest resources that all server pools can access reduces storage space needed to store duplicate resources. Without this solution you would be forced to import a lot of the same assemblies, ISO images and templates needed to support Oracle VM Guest creation in each server pool. For example, assuming you had four different server pools, each requiring the same template to create RAC clusters; this alone would require you to consume 12 Gigabytes for a total of 48 Gigabytes across all four server pools. Using a centralized repository would reduce this number by 36 Gigabytes.
- *Faster backup and restores* – Capturing files and directories for backups of guest resources would be faster since you are not also capturing virtual disks for Oracle VM. In most cases, capturing virtual disks belonging to running virtual machines will be useless since databases and applications may not be quiesced at the time of backup.
- *Less space storing backups* – The guest resources will change a lot less than running virtual machines. Capturing the virtual disks of running virtual machines along with the more static guest resources

will cause snapshots to grow as the snapshot software tracks changes between the snapshot and the active original images. Less space on tape will be consumed since you are not also storing oftentimes useless virtual disk files.

- *Better organization of space* – Less time is spent determining which virtual disks belong to templates and which belong to running virtual machines.
- *Easier to quiesce activity* – Separate storage repositories for guest resources will make it easier to develop and manage policies to limit system administrator activity to prevent people from making changes to the files contained in the repository during a backup.

Implementing a Solution for Segregation

This is a very easy solution to implement and really depends on presentation to Oracle VM Servers or server pools. The following series of screen shots show a little more information about how the solution works. The screen shot in Figure 28 below shows that the repository named **Guests mypool1 repo1** is presented to the only two servers that are members of **mypool1**.

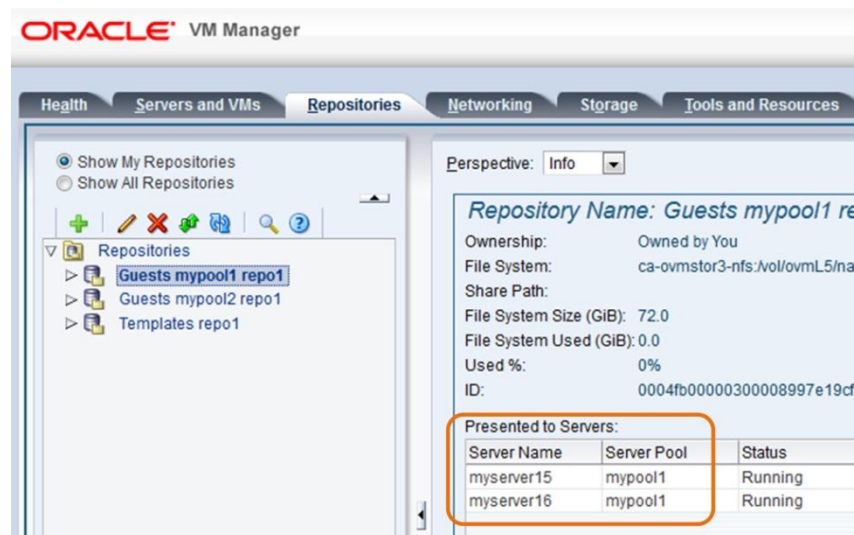


Figure 28: Screen shot showing the repository for mypool1 is presented only to Oracle VM Servers in mypool1

The screen shot in Figure 29 below shows that the repository named **Guests mypool2 repo1** is presented to the only two servers that are members of **mypool2**.

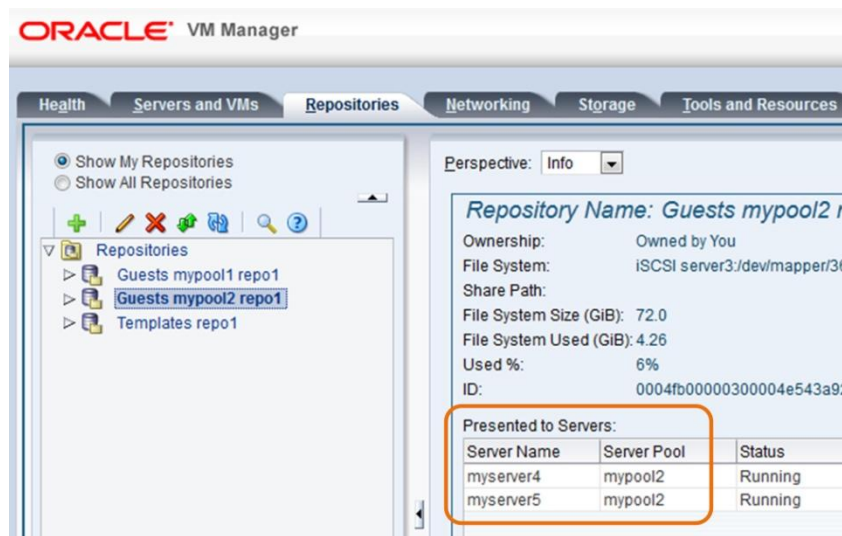


Figure 29: Screen shot showing the repository for mypool2 is presented only to Oracle VM Servers in mypool2

The screen shot in Figure 30 below shows that the repository named **Templates repo1** is presented to all Oracle VM Servers belonging to **mypool1** as well as **mypool2**. This of course is the repository containing only guest resources such as assemblies, ISOs and templates which is accessible by all Oracle VM Servers in the Oracle VM model.

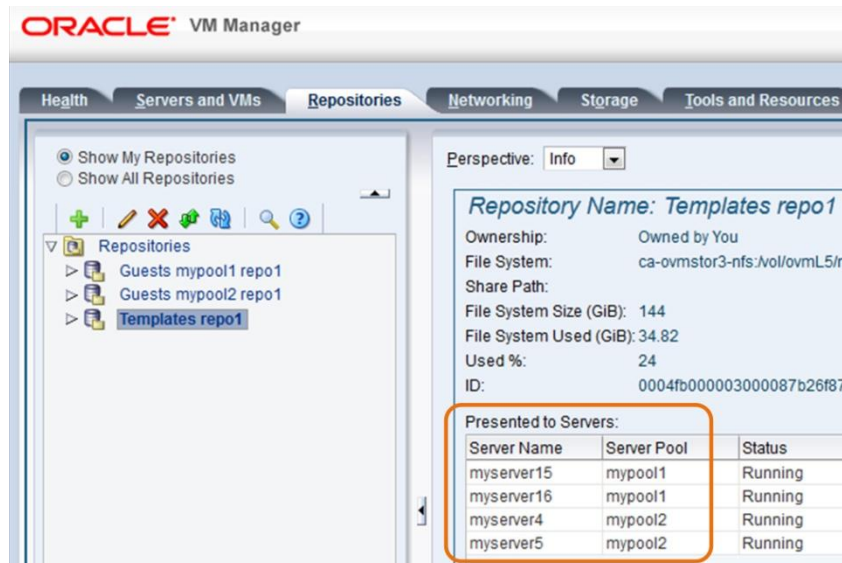


Figure 30: Screen shot showing the template repository is presented to all Oracle VM Servers in all pools

A Clear Repository Naming Convention is Essential

The only weakness with this solution is the fact all of the repositories can inadvertently be used for cross purposes – that is there is nothing that prevents an inattentive system administrator from

importing guest resources into one of the repositories for Oracle VM Guests images only. You can see in Figure 3131 below that all of the repositories have the same directory structure which means there is no difference between the various repositories and there is nothing in the Oracle VM Manager to limit how a repository should be used or the type of images that can be added to any repository.

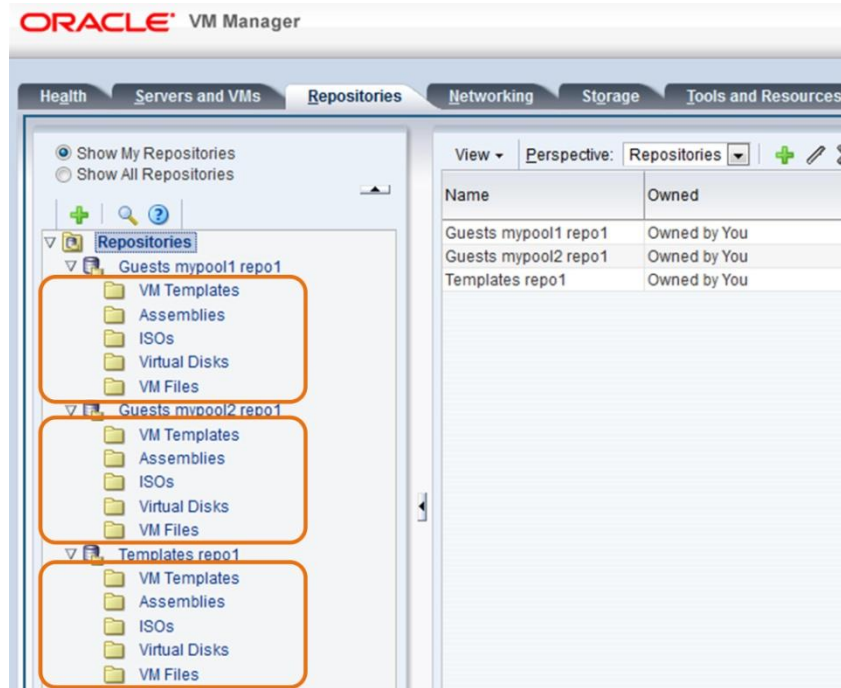


Figure 31: Screen shot showing storage repositories have same directory structure

The solution for segregating guest resources from running virtual machines is very simple to implement but does require some discipline to maintain correctly. You will need to develop a naming convention to help prevent people from inadvertently mixing Oracle VM Guest resources with repositories meant for Oracle VM Guests and vice versa. For the purpose of this document we chose to use the conventions of prefacing repository names with the intended roll: **Templates** for repositories meant to contain only assemblies, ISOs and template and **Guests** for repositories meant to contain only Oracle VM Guests for each server pool. You will need to devise something that works for your unique requirements that makes the roll of a repository very obvious when you are performing various tasks in the Oracle VM Manager

Tailor a Solution that Fits Your Needs

Oracle believes this is a simple yet very useful best practice for making your Oracle VM platform easier to manage and more scalable while saving a significant amount of space on shared storage as well as tape. However, segregation of guest resources is simply a suggestion for your consideration. Use your experience and skills to create a repository deployment scheme that fits your unique requirements.

Multitier Backup for Guest Resources at Storage Level

Backups for Oracle VM Guest resources are very simple and straight forward. A backup scheme involving a couple different levels of backups to different media will provide the most robust solution

for recovering from different levels of catastrophic events. This particular solution will protect guest resources from the corruption or loss of individual files as well as the corruption or loss of an entire NFS export.

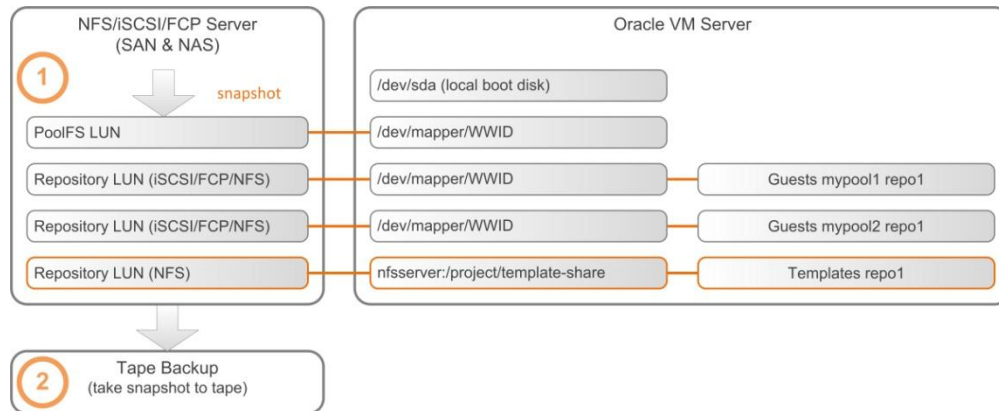


Figure 32: Create a multitier protection scheme using snapshots plus tape backup

The first level of protection for Oracle VM Guest resources will be an image based snapshot on the storage array where the storage object containing the NFS image resides. This will allow you to quickly recover from the corruption or loss of individual files simply by restoring the entire snapshot or individual files and directories contained in a snapshot.

This solution assumes that any repositories meant to span multiple server pools will reside in a volume/project dedicated just this purpose. Looking back at the storage deployment architecture discussed in Part 1 [Understanding Deployment Architecture](#), the repository dedicated to presenting Oracle VM Guest resources across multiple server pools would reside in its own volume/project separate from storage objects destined for individual server pools.

The second level of protection for Oracle VM Guest resources will be periodic backups of the snapshots to tape. This will allow you to recover from the corruption or loss of an entire NFS export or storage object on shared storage.

Single Tier Backup for Guest Resources at Repository Level

Although single tier backups using only tape backup is not as robust as the multitier solution, some data centers may not have access to enterprise class storage arrays with built-in hardware RAID, hot spares and sophisticated snapshot technology. This solution will still provide protection against corruption or loss of individual files as well as the corruption or loss of an entire NFS export.

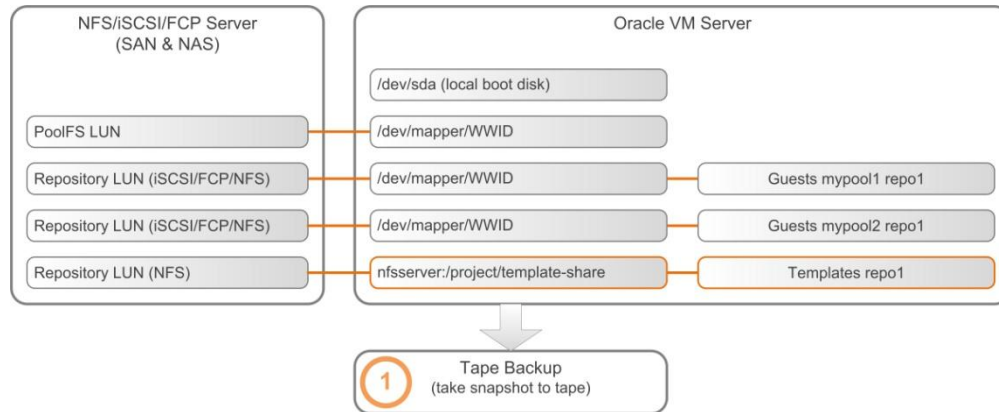


Figure 32: Create a single tier protection scheme using tape or other Linux system tool

The only level of backup in this case is regular scheduled backups to tape. It simply requires that the guest resources repository be mounted to a backup server, then begin the backup.

Conclusion

Part 3 of this guide has discussed how to backup and restore the Oracle VM Guests and Resources. You now have learned the options to protect the valuable assets of the Oracle VM infrastructure.

References

For more information about Oracle's virtualization, visit www.oracle.com/virtualization.

- Oracle VM at Oracle Technology Network: <http://www.oracle.com/technetwork/server-storage/vm/overview/index.html>
- Oracle VM Documentation: <http://www.oracle.com/technetwork/documentation/vm-096300.html>
- My Oracle Support Doc ID 1519114.1: Oracle VM Resources - Backup & Restore Considerations
- My Oracle Support Doc ID 1477421.1: How To Backup And Restore A VM Guest (domU) Domain On Oracle VM 3.x



Oracle VM 3: Backup and Recovery Best Practices Guide

August 2013, Version 1.0.1

Author: Gregory King

Contributors:

Honglin Su, Chee Weng Hey

Oracle Corporation

World Headquarters

500 Oracle Parkway

Redwood Shores, CA 94065

U.S.A.

Worldwide Inquiries:

Phone: +1.650.506.7000

Fax: +1.650.506.7200

oracle.com



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

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

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

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

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0113

Hardware and Software, Engineered to Work Together