



NETAPP TECHNICAL REPORT

RAPID DEPLOYMENT OF ORACLE® DATABASE 11G USING VMWARE® INFRASTRUCTURE AND NETAPP FLEXCLONE®

Steven Schuettinger, Network Appliance, Inc.
Desmond Chan, VMware Inc.

January 2008 | TR-3646-0108

EVERY DEVELOPER HAS A DATABASE

This paper provides an overview of the first-generation grid infrastructure, describes how advanced server and storage virtualization can enhance existing grids, and includes a case study of a next-generation grid implementation that demonstrates how to solve a pervasive technical problem.



TABLE OF CONTENTS

1	INTRODUCTION	4
1.1	THE FIRST-GENERATION GRID	4
1.2	ENABLING SOFTWARE TECHNOLOGIES	4
1.3	ENABLING HARDWARE TECHNOLOGIES.....	4
1.4	WHAT'S NEXT? VIRTUALIZATION AND THE GRID	4
2	NETAPP FLEXCLONE AND FLEXVOL.....	5
2.1	DYNAMIC, NONDISRUPTIVE STORAGE PROVISIONING.....	5
2.2	SUPERIOR STORAGE UTILIZATION LOWERS COST	6
2.3	INCREASED PERFORMANCE	6
3	VMWARE INFRASTRUCTURE.....	6
3.1	VMWARE ESX SERVER.....	7
3.2	VMWARE VIRTUALCENTER.....	7
3.3	VMWARE INFRASTRUCTURE BENEFITS	7
4	TECHNICAL PROBLEM: CLONING DATABASES FOR TEST AND DEVELOPMENT	8
5	ORACLE DATABASE 11G DEPLOYMENT ON VMWARE INFRASTRUCTURE AND NETAPP STORAGE.....	9
5.1	DEPLOYMENT ENVIRONMENT	9
5.2	SOFTWARE SUMMARY	10
5.3	HARDWARE SUMMARY	10
5.4	DEPLOYMENT APPROACH SUMMARY	10
6	VMWARE INFRASTRUCTURE INSTALLATION.....	11
7	STORAGE CONFIGURATION	12
7.1	STEP 1: CONFIGURE THE VMKERNEL TCP/IP NETWORKING STACK.....	12
7.2	STEP 2: ENABLE THE ISCSI SOFTWARE CLIENT SERVICE AND OPEN THE CORRESPONDING FIREWALL PORT	14
7.3	STEP 3: CONFIGURE THE ISCSI SOFTWARE INITIATOR.....	14
7.4	STEP 4: CREATE A LUN	15
7.5	STEP 5: CONFIGURE AN INITIATOR GROUP ASSOCIATED WITH THE ESX SERVER ISCSI INITIATOR	16
7.6	STEP 6: RESCAN FOR NEW ISCSI LUNS	16
7.7	STEP 7: SET UP THE VMFS DATA STORE	16
8	INSTALL ORACLE DATABASE 11G VIRTUAL MACHINE	17
9	NETAPP FLEXCLONE	18
9.1	CREATE A BASELINE VOLUME FOR FLEXCLONE.....	18
9.2	FLEXCLONE CREATION.....	18
10	REGISTER THE VIRTUAL MACHINES	18

11 CUSTOMIZE THE CLONED VIRTUAL MACHINE.....	19
11.1 SAMPLE CUSTOMIZATION.....	19
11.2 CHANGE SID	19
12 CONCLUSION	19
13 RESOURCES.....	20
13.1 NETAPP	20
13.2 VMWARE	20
13.3 ORACLE	20
14 APPENDIX: EXAMPLE SCRIPTS TO PERFORM FLEXCLONE AND VMWARE VIRTUAL MACHINE REGISTRATION	21

1 INTRODUCTION

Data centers built on a grid infrastructure provision standard commodity servers and networked storage to create a flexible, scalable architecture capable of achieving increased levels of resource utilization as compared with traditional data centers. Since Oracle introduced grid-optimized Oracle Database 10g software, many of its customers have benefited from deploying their applications on a grid infrastructure.

The pace of technological change has accelerated in recent years, and as one would expect, first-generation grid architectures are absorbing the latest technology and continuing to evolve. The grid is now poised to leverage the latest advancements in virtualization in order to further enhance the value of a grid infrastructure.

This paper provides an overview of the first-generation grid infrastructure, describes how advanced server and storage virtualization can enhance existing grids, and includes a case study of a next-generation grid implementation that demonstrates how to solve a pervasive technical problem by creating database clones that are space efficient, rapidly provisioned, and easily managed.

1.1 THE FIRST-GENERATION GRID

An Oracle grid allows IT managers to dynamically share resources such as servers, storage, and networks among different enterprise applications, thereby reducing costs through higher resource utilization. Instead of buying more powerful computers, a grid allows them to add capacity on demand for superior scalability. Implementers of first-generation grids have achieved rapid return on investment while successfully consolidating disparate IT systems, raising management efficiency, and meeting higher service-level agreements (SLAs).

1.2 ENABLING SOFTWARE TECHNOLOGIES

From a software standpoint, the enabling technologies for the first-generation Oracle grid include:

- Oracle Database 10g, the first grid-optimized database platform
- Oracle Real Application Clusters (RAC) technology, which delivers a flexible and resilient database engine that can scale to meet dynamically changing database needs
- Oracle Automatic Storage Management (ASM), which provides a simple storage management interface that is consistent across all server and storage platforms
- Oracle Enterprise Manager Grid Control, which simplifies the management of multiple instances of Oracle applications within a grid environment

1.3 ENABLING HARDWARE TECHNOLOGIES

On the hardware side, the first-generation grid is made possible by the availability of powerful, rack-mounted commodity servers coupled with flexible networked storage from NetApp. By flexibly connecting servers to networked storage IT professionals can quickly add server resources to a particular application and ensure immediate access to the necessary storage volumes.

1.4 WHAT'S NEXT? VIRTUALIZATION AND THE GRID

While first-generation grids have achieved impressive results, IT organizations continue to face a host of challenges. Increased power and cooling costs, physical space constraints, manpower shortfalls, and interconnection complexity contribute significantly to the cost and reduce the feasibility of continued IT expansion.

As the processors in commodity servers continue to gain in power, companies upgrading to the latest technology frequently see CPU utilization drop. Studies have shown that newer generation servers are at 15% utilization and lower, sometimes as low as 3% to 5%. This underutilization represents a major opportunity for consolidation.

The next-generation grid will utilize the latest energy-efficient blade server technology, advanced server virtualization, and the latest storage virtualization to:

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

The pairing of server virtualization with storage virtualization in a grid environment can raise both server and storage utilization to maximize data center efficiency. Companies with existing grids can begin to take advantage of Oracle software, NetApp storage virtualization, and VMware Infrastructure capabilities to enhance the value of existing infrastructure.

2 NETAPP FLEXCLONE AND FLEXVOL

NetApp FlexClone technology enables instant cloning of data volumes and data sets without requiring additional storage space at the time of creation. Each cloned volume is a transparent, virtual copy that you can use for essential enterprise operations, such as testing and bug fixing, platform and upgrade checks, multiple simulations against large data sets, remote office testing and staging, and market-specific product variations.

FlexClone is a powerful and unique feature introduced in Data ONTAP® 7G that adds a new level of agility and efficiency to storage operations by allowing an individual to create an instant clone of a flexible volume (FlexVol™ volume). A FlexClone volume is a writable point-in-time image of a FlexVol volume or the FlexVol volume itself. With FlexClone, it takes only a few seconds to create a clone of a FlexVol volume without interrupting access to the parent volume the clone is based on. The clone volume uses space very efficiently, allowing both the original FlexVol volume and the cloned volume to share common data, storing only the data that changes between the original volume and the clone. This provides a huge potential saving in storage space, resources, and cost. In addition, a FlexClone volume has all the features and capabilities of a regular FlexVol volume, including the ability to be grown or shrunk and the ability to be the source of another FlexClone volume.

NetApp FlexVol technology enables true storage virtualization to lower overhead and capital expenses, reduce disruption and risk, and provide the flexibility to adapt quickly and easily to the dynamic needs of an organization. FlexVol technology pools storage resources automatically and enables the creation of multiple flexible volumes on a large pool of disks. This flexibility means applications and storage administrators can focus on managing data, not hardware, and make changes quickly and seamlessly. As a result, storage can be added when and where it is needed without disruption and at the lowest incremental cost.

2.1 DYNAMIC, NONDISRUPTIVE STORAGE PROVISIONING

FlexVol is a breakthrough technology in which volumes are logical data containers that can be managed and moved independently from the underlying physical storage. Since FlexVol volumes have access to all available physical storage in the NetApp system, FlexVol technology provides substantial improvements in storage provisioning, utilization, and volume sizing. Data volumes can be sized and resized¹ quickly and dynamically as application requirements change, and sizing may be automated based on space management policies.

¹ VMware recommends that you consult with VMware documentation and follow the specific procedure for storage resizing. At the time of this writing, VMware Infrastructure does not support dynamic resizing of VMFS volumes.

2.2 SUPERIOR STORAGE UTILIZATION LOWERS COST

Since FlexVol technology allows all data volumes to share a single pool of physical disk storage, storage utilization is dramatically improved. Compared to alternative technologies, FlexVol can double storage utilization, which means that you can significantly reduce the cost of disk drive acquisition. Additionally, the availability of thin provisioning provides the added benefit of "just-in-time" storage, which further reduces acquisition costs.

2.3 INCREASED PERFORMANCE

With FlexVol, I/O performance is also improved, since volumes can be distributed across all available disk drives. This is particularly important for "spindle-bound" applications such as those with heavy database utilization. In addition, since I/O bottlenecks are dramatically reduced, less expensive SATA disk drives may be utilized for a wider range of applications.

3 VMWARE INFRASTRUCTURE

VMware Infrastructure simplifies IT environments so that you can leverage your storage, network, and computing resources to control costs and respond quickly to changing business needs. The VMware Infrastructure approach to IT management creates virtual services out of the physical infrastructure, enabling administrators to allocate these virtual resources quickly to the business units that need them most.

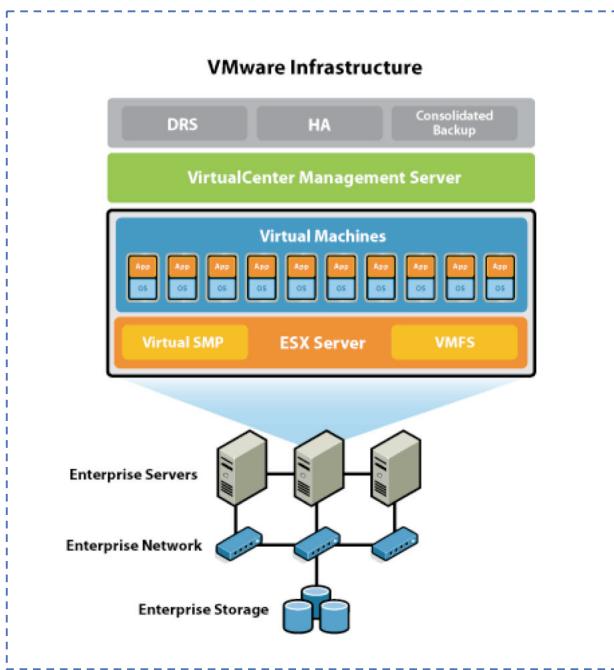


Figure 1) VMware Infrastructure.

VMware Infrastructure 3 is the next generation of industry-leading infrastructure virtualization software that virtualizes servers, storage, and networking, allowing multiple unmodified operating systems and their applications to run independently in virtual machines while sharing physical resources. VMware Infrastructure 3 software components include VMware ESX Server and VMware VirtualCenter.

3.1 VMWARE ESX SERVER

VMware ESX Server abstracts processor, memory, storage, and networking resources into multiple virtual machines, giving IT greater hardware utilization and flexibility.

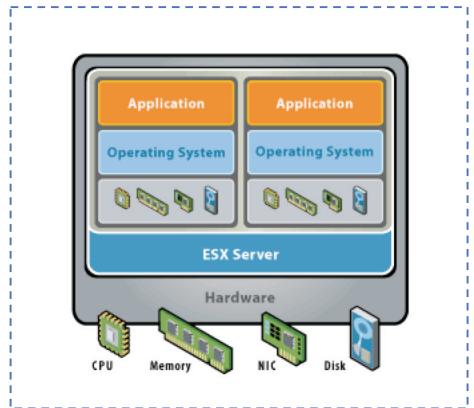


Figure 2) VMware Infrastructure 3: ESX Server.

3.2 VMWARE VIRTUALCENTER

VMware VirtualCenter enables rapid provisioning of virtual machines and performance monitoring of physical servers and virtual machines. VirtualCenter intelligently optimizes resources, ensures enhanced high availability to all applications in virtual machines, and makes IT environments more responsive with virtualization-based distributed services such as VMware High Availability (HA) and VMware VMotion.

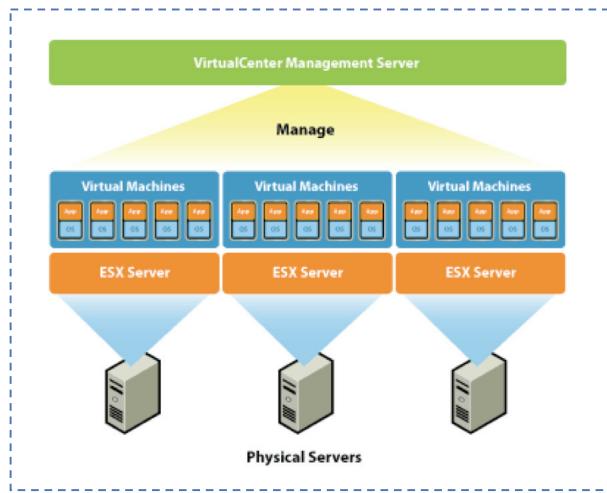


Figure 3) VMware VirtualCenter Management Server.

3.3 VMWARE INFRASTRUCTURE BENEFITS

Deploying an Oracle Database on VMware Infrastructure drives tangible benefits through **server containment**. Traditional Oracle Database deployments generate significant server sprawl due to the need to provision separate systems for development, test, and production environments. Each developer requires a dedicated environment, and each test cycle requires dedicated servers assigned for the duration of the tests. VMware virtualization technology contains server sprawl and increases server utilization by running multiple Oracle Database virtual machines, consolidated onto fewer, highly scalable, reliable enterprise-class systems.

VMware virtualization solutions significantly reduce the time required to provision new instances (development, test, or production) of Oracle Database environments. Users of VMware Infrastructure can take advantage of virtual machine libraries and templates to provision new preconfigured Oracle Database environments in minutes. **Rapid provisioning** and **change management** in Oracle enterprise deployments increase IT flexibility, allowing for timely response to rapidly changing business needs, while data center optimization enables efficient resource pooling, maximizing hardware utilization. Implementing **business continuity** solutions for Oracle software on VMware Infrastructure delivers enhanced high availability without major investments in one-to-one mapping of production and DR hardware.

Used together, Oracle and VMware technologies reduce cost and deployment time while providing increased operational efficiencies and accelerated time to market, driving lower TCO and increased ROI for customers.

4 TECHNICAL PROBLEM: CLONING DATABASES FOR TEST AND DEVELOPMENT

Database cloning is the process of creating an exact copy of an Oracle Database. This process is performed frequently by database administrators to provide near-production instances for various business needs such as application development, QA testing, and report generation. Traditional methods of cloning a database pose several challenges, including extended system downtime and degraded system performance during the cloning process. Additionally, a large amount of storage space is required to store each clone.

In this era of high system availability, organizations cannot afford extended downtime and degraded performance for their production systems. Therefore, the ability to create a usable database clone quickly and with virtually no impact on the production system is extremely important. Moreover, depending on the number of database clones needed, storage consumption can be huge, which can prohibit organizations from deploying the desired number of database clones.

VMware Infrastructure 3 allows you to create a base virtual machine that runs an Oracle Database 11g properly installed and tested. This base virtual machine is stored as a set of files in a storage system, e.g. NetApp FAS. It is self-contained, including a guest operating system, Oracle home, and all the data and log partitions. Cloning the virtual machine, along with minimal configuration such as network parameters configuration, gives you all you need for a functional database system.

To clone the virtual machine, you can leverage the FlexClone feature of a NetApp storage system to create an instantaneous point-in-time copy of the storage system volumes that hold the virtual machine files. By using FlexClone technology, the cloned database virtual machines share data with the original one, saving a tremendous amount of storage space.

The following sections describe a hypothetical Oracle Database 11g deployment on VMware Infrastructure, with a NetApp FAS. In this environment, you can easily leverage NetApp FlexClone technology to create more volume clones that contain the deployed Oracle Database 11g virtual machine, for testing and development purposes.

5 ORACLE DATABASE 11G DEPLOYMENT ON VMWARE INFRASTRUCTURE AND NETAPP STORAGE

You can streamline deployment of Oracle Database 11g instances by leveraging the features in VMware Infrastructure and NetApp storage. This section outlines the procedures of a deployment approach practiced in the NetApp lab.

Note: The approach illustrated in this section serves as a suggestion and is by no means the recommended approach for Oracle Database 11g deployment. Refer to Oracle Database 11g documentation for more information.

5.1 DEPLOYMENT ENVIRONMENT

Figure 4 shows the deployment environment in the NetApp lab located in Sunnyvale, California. Two VMware ESX Server hosts are deployed on industry-standard servers. One server hosts a virtual machine that runs a VirtualCenter Server. The other ESX Server host, connecting to a NetApp FAS3070, contains the base Oracle Database virtual machine. A VMware Infrastructure Client is installed on an IBM Lenovo X60 laptop to provide access to the virtual machine's display.

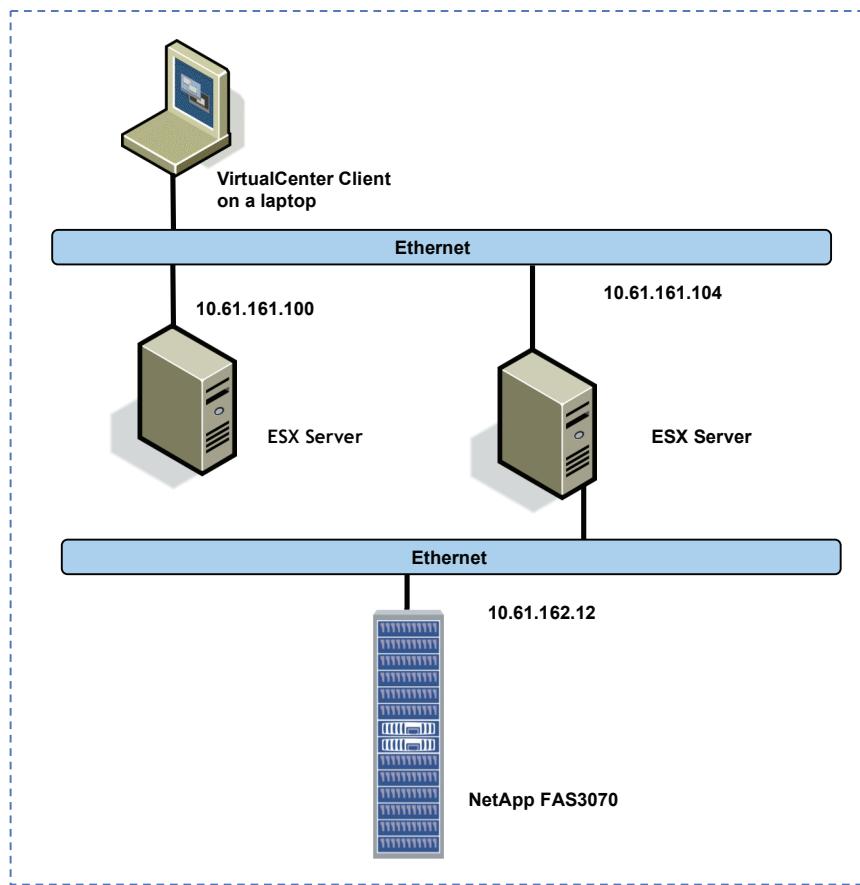


Figure 4) Sample VMware Infrastructure environment, using NetApp array, for Oracle Database 11g deployment.

5.2 SOFTWARE SUMMARY

The table below lists the software used in the deployment outlined in this guide.

Category	Software
Infrastructure	VMware Infrastructure 3 (VI3)
Applications/Software	Oracle Database 11g
Guest Operating System	Oracle Enterprise Linux®
NetApp Array	Data ONTAP Release 7.2.2

5.3 HARDWARE SUMMARY

The table below outlines the hardware used to support the virtualized Oracle Database 11g installation.

Category	Hardware
ESX Server Host	Fujitsu Siemens D1889
VirtualCenter Management Server	N/A (hosted on a virtual machine)
VMware Infrastructure Client	IBM X60 Lenovo laptop computer
Storage	NetApp FAS3070

5.4 DEPLOYMENT APPROACH SUMMARY

Note: The snapshots referred to in the following sections are **NetApp volume Snapshot™ copies**. They are not to be interpreted as VMware virtual machine snapshots.

This section explains how to streamline rapid deployment of an Oracle Database 11g environment on VMware Infrastructure 3, leveraging the NetApp FlexClone feature. It tells you how to install a baseline virtual machine containing an Oracle Database and stored in a NetApp FlexVol volume. Next you will create a Snapshot copy of the NetApp FlexVol volume, which will serve as the baseline from which clones of volumes can be produced with NetApp FlexClone technology. This section also explains how to register the virtual machines in these cloned volumes with VMware Infrastructure by executing ESX service console commands. At the end of this exercise, you will have a baseline volume image that you can use with FlexClone technology to create Oracle Database 11g virtual machines, within minutes and without using extra storage.

Deployment steps include:

1. Prepare a VMware Infrastructure environment.
 - a Install ESX Server on the host machines.
 - b Install VirtualCenter Management Server in a virtual machine.
 - c Install VMware Infrastructure Client.
2. Enable software-based iSCSI initiator on ESX Server.
3. Create and Configure NetApp FlexVol volume.
 - a Create a FlexVol volume.
 - b Create a LUN in the volume.
 - c Associate the LUN with an initiator group.

4. Create a VMware Virtual Machine File System (VMFS) data store² on the LUN configured in step 3.
5. Create an Oracle Database 11g virtual machine.
 - a Create a virtual machine on one ESX Server host.
 - b Add the necessary virtual disk(s) to the virtual machine, in the data store created in step 4.
 - c Install guest operating system (Oracle Enterprise Linux 5) on the virtual machine.
6. Install and configure Oracle Database 11g environment.
 - a Install and configure Oracle Database 11g, using Oracle ASM.
7. Create a baseline image of the NetApp volume.
 - a Create a Snapshot copy of the volume using NetApp Snapshot.
8. Clone the virtual machine created in step 5 using NetApp FlexClone.
 - a Launch the cloned virtual machine.
 - b Configure the cloned virtual machine.

6 VMWARE INFRASTRUCTURE INSTALLATION

The VMware Infrastructure consists of a number of components:

- ESX Server
- VirtualCenter Management Server
- VMware Infrastructure Client
- VMware License Server

The installation procedures for each component are shown in detail in the guide located here:

www.vmware.com/pdf/vi3_installation_guide.pdf.

² VMware recommends that you evaluate the requirements before making the decision to use VMFS or raw device mapping (RDM) in your data stores. This exercise uses VMFS.

7 STORAGE CONFIGURATION

This section shows you how to configure the NetApp storage to expose a LUN via an iSCSI interface to the ESX Server host.

Steps	ESX Server	NetApp FAS	VMware Infrastructure Client	Expected Results/Effects
Configure VMKernel TCP/IP networking stack			Step 1	This step sets up VMKernel TCP/IP networking stack for iSCSI to function.
Enable the iSCSI software client service and open the corresponding firewall port	Step 2			The iSCSI software client service should be enabled on the ESX Server side to support proper communication with the NetApp FAS.
Configure the iSCSI software initiator			Step 3	The iSCSI software initiator should be properly configured, with an iSCSI name. The iSCSI name will be used to configure an initiator group on NetApp FAS in step 5. Any LUN associated with the initiator group will be visible to ESX Server.
Create a LUN		Step 4		An aggregate needs to be created on the NetApp FAS. On the aggregate, a flexible volume is created. Within the flexible volume, a LUN is created, as an iSCSI target that will be made visible to ESX Server in the next steps.
Configure an initiator group associated with the ESX Server iSCSI initiator and map the LUN to the initiator group		Step 5		An initiator group needs to be configured, using the iSCSI name noted in step 3. The LUN needs to be mapped to the initiator group in order to be visible to ESX Server.
Rescan for new iSCSI LUNs			Step 6	With the LUN properly mapped to the initiator group, ESX Server should be able to detect it through a rescan process.
Set up the VMFS data store			Step 7	In this step, a VMFS data store is set up on the newly provisioned LUN to host virtual disks of virtual machines.

7.1 STEP 1: CONFIGURE THE VMKERNEL TCP/IP NETWORKING STACK

Refer to the following guidelines when configuring VMkernel networking.

The IP address that you assign to the service console during installation must be different from the IP address that you assign to VMkernel's TCP/IP stack from the Configuration > Networking tab of the Virtual Infrastructure Client.

Unlike other VMkernel services, iSCSI has a service console component, so networks that are used to reach iSCSI targets must be accessible to both service console and VMkernel TCP/IP stacks.

To set up the VMkernel:

1. Log on to the VMware Infrastructure Client and select the server from the inventory panel. The hardware configuration page for this server opens.
2. Click the **Configuration** tab, and click **Networking**.
3. Click the **Add Networking** link.

The **Add Network** wizard opens.

4. Select **VMkernel** and click **Next**.

Selecting **VMotion** and **IP Storage** lets you connect the VMkernel, which runs services for VMotion and IP storage (NFS or iSCSI), to the physical network.

The **Network Access** page opens.

5. Select the virtual switch you would like to use, or select the **Create a virtual switch** radio button to create a new virtual switch.
6. Select the checkboxes for the network adapters your virtual switch will use.

Your choices appear in the **Preview** pane.

Select adapters for each virtual switch so that virtual machine or other services that connect through the adapter can reach the correct Ethernet segment. If no adapters appear under **Create a new virtual switch**, all the network adapters in the system are being used by existing switches. You can either create a new virtual switch without a network adapter or select a network adapter used by an existing virtual switch.

7. Click **Next**.

The **Connection Settings** page opens.

8. Under **Port Group Properties**, select or enter a network label and a VLAN ID.
 - **Network Label:** A name that identifies the port group that you are creating. This is the label that you specify when configuring a virtual adapter to be attached to this port group, when configuring VMkernel services such as VMotion and IP storage.
 - **VLAN ID:** Identifies the VLAN that the port group's network traffic will use.

9. [Optional] Select the **Use this port group for VMotion** checkbox to enable this port group to advertise itself to another ESX Server instance as the network connection where VMotion traffic should be sent.

You can enable this property for only one VMotion and IP storage port group for each ESX Server host. If this property is not enabled for any port group, migration with VMotion to this host is not possible.

10. Under **IP Settings**, click **Edit** to set the **VMkernel Default Gateway** for VMkernel services such as VMotion, NAS, and iSCSI.

The **DNS and Routing Configuration** dialog box appears. Under the **DNS Configuration** tab, the name of the host is entered into the name field by default. The DNS server addresses that were specified during installation are also preselected, as is the domain.

Under the **Routing** tab, the service console and the VMkernel each need their own gateway information. A gateway is needed if connectivity to machines not on the same IP subnet such as the service console or VMkernel is needed.

Static IP settings is the default.

11. Click **OK** to save your changes and close the **DNS Configuration and Routing** dialog box.
12. Click **Next**.
13. Use the **Back** button to make any changes.
14. Review your changes on the **Ready to Complete** page and click **Finish**.

7.2 STEP 2: ENABLE THE iSCSI SOFTWARE CLIENT SERVICE AND OPEN THE CORRESPONDING FIREWALL PORT

To open the corresponding firewall port, execute:

```
esxcfg-firewall -e swISCSIClient
```

To ensure that the iSCSI software client service is running, execute the query:

```
esxcfg-swiscsi -q
```

If the iSCSI software client service is not running, execute:

```
esxcfg-swiscsi -e
```

7.3 STEP 3: CONFIGURE THE iSCSI SOFTWARE INITIATOR

To view the iSCSI software initiator properties:

1. Log in to the VMware Infrastructure Client and select the desired ESX Server host from the inventory panel.
2. Click the **Configuration** tab, and click **Storage Adapters** under hardware.
The list of available storage adapters appears.
3. Under **iSCSI Software Adapter**, choose the available software initiator.
The details for the initiator display, including the model, IP address, iSCSI name, discovery methods, iSCSI alias, and any discovered targets.
4. Click **Properties**.
The **iSCSI Initiator Properties** dialog box opens. The **General** tab displays additional characteristics of the software initiator.
5. Click **Configure**.
The **General Properties** dialog box opens, displaying the initiator's status, default name, and alias.
6. Ensure that the **Enabled** checkbox is checked.
7. To change the default iSCSI name for your initiator, enter the new name. (**Note:** Make a note of the iSCSI name for initiator group configuration on NetApp FAS.)
8. Enter the iSCSI alias or change the existing one.
The alias is a friendly name that you use to identify the iSCSI hardware initiator.
9. Click **OK** to save your changes.

To set up target discovery addresses for the software initiator:

1. Open the **iSCSI Initiator Properties** dialog box by performing the steps listed above.
 2. Click the **Dynamic Discovery** tab.
 3. To add a new iSCSI target your ESX Server host can use for a SendTargets session, click **Add**.
The **Add Send Targets Server** dialog box opens.
 4. Enter the Send Targets Server IP address (i.e., the NetApp FAS IP address) and click **OK**.
 5. To change or delete a Send Targets server, select **Edit** or **Remove**.
- Note:** When configuring your iSCSI software initiator, make sure that CHAP parameters are enabled for the initiator. If they are not enabled, you need to set them up.

To set up CHAP parameters for the software initiator:

1. Open the iSCSI Initiator Properties dialog box by performing the steps listed above.
2. Click the **CHAP Authentication** tab.

- The tab displays the default CHAP parameters.
3. To make any changes to the existing CHAP parameters, click **Configure**.
The **CHAP Authentication** dialog box opens.
 4. To keep CHAP enabled, make sure **Use the following CHAP credentials** is selected.
 5. To use a new CHAP name, deselect the **Use initiator name** checkbox and enter the name of your choice.
 6. If needed, specify the CHAP secret.
All new targets will use the CHAP secret to authenticate the initiator. Any established sessions are not affected.
7. Click **OK** to save changes.

7.4 STEP 4: CREATE A LUN

On the NetApp storage side, a LUN needs to be provisioned for a VMware virtual machine to store its corresponding data. This section familiarizes you with some fundamental NetApp storage concepts before showing you the steps to create a LUN.

NETAPP STORAGE CONCEPTS

Within a NetApp storage appliance, storage is organized in a hierarchy of units: aggregates, (flexible) volumes, and LUNs.

Aggregates

An aggregate is simply a collection of RAID groups across which flexible volumes are striped. Aggregates add a layer of abstraction between the physical disks and the volumes. As a result, the volumes are not tightly coupled to the physical disks. You can allocate a large number of disks to a large aggregate, which is built on RAID groups. There is no file system within an aggregate. It is a construct of allocatable space.

Flexible Volumes

NetApp flexible volumes are created on top of aggregates and thus are abstracted from the underlying disks. They give users the flexibility to resize the volume without regard for the physical layout.

Flexible volumes contain LUNs or virtual disk files that are accessed by VMware ESX Server host.

LUNs

LUN stands for logical unit number. It is an address that uniquely identifies a SCSI target that can be accessed by servers.

In a NetApp FAS system, LUN is implemented as a file in the file system of a volume. In other words, you need to create a volume before you can create a LUN within it.

A VMware ESX Server host accesses a LUN in two ways. The first and most common method is as storage to hold virtual disk files for multiple virtual machines. This type of usage is referred to as a Virtual Machine File System (VMFS) LUN. The second method is raw device mapping (RDM). With RDM, the LUN is connected to the ESX Server host and is passed directly to a virtual machine to use with its native file system, such as NTFS or EXT3.

Note: Our scenario uses VMFS.

To create a LUN in a NetApp FAS, perform the following:

1. Log in to the NetApp FAS command line interface and create an aggregate if required:

```
aggr create <new aggregate name> <number of disks>
```

2. Create a flexible volume:

```
vol create <aggregate name> <volume size>
```

3. Create a qtree to store each LUN:

```
qtree create <path to the parent directory of the LUN>
```

4. Create a LUN:

```
lun create -s <size> -t linux <path to LUN>3
```

7.5 STEP 5: CONFIGURE AN INITIATOR GROUP ASSOCIATED WITH THE ESX SERVER ISCSI INITIATOR

To create an initiator group, execute:

```
igroup create -i -t linux <igroup name> <initiator name>
```

To map the newly created lun to this initiator, execute:

```
lun map <LUN path> <igroup name>
```

7.6 STEP 6: RESCAN FOR NEW ISCSI LUNS

To perform a rescan:

1. In the VMware Infrastructure Client, select an ESX Server host and click the **Configuration** tab.
2. Choose **Storage Adapters** in the Hardware panel and click **Rescan** above the **Storage Adapters** panel.
3. To discover new disks or LUNs, select **Scan for New Storage Devices**.

If new LUNs are discovered, they appear in the disk/LUN list.

7.7 STEP 7: SET UP THE VMFS DATA STORE

In order to add a shared storage system, such as a new LUN, after ESX Server has been installed, follow the procedure below to rescan and add the shared storage system to the infrastructure using the VMware Infrastructure Client.

1. Select the desired ESX Server host in Inventory View.
2. Click the Configuration tab.
3. Click **Add Storage...** on the top-right corner.
4. In the **Add Storage** wizard, enter the following:

³ Since, in NetApp FAS, a LUN is a file within the volume file system, you can construct a LUN path the way you construct the path to the file representing the LUN.

Disk/LUN	Select disk/LUN.	
Device Location	A list of available LUNs should appear in the window to the right. Select the device to add.	
Current Disk Layout	(Indication that the disk is blank should be shown.) Click Next	
Properties	Datastore Name	ORA11GDBSTORE
Formatting	Keep the defaults. Make sure that the checkbox for Maximize Capacity has been checked.	
Ready to Complete	Click Finish.	

8 INSTALL ORACLE DATABASE 11G VIRTUAL MACHINE

In this exercise, you will create a virtual machine with the following configurations⁴:

vCPU	2
Memory	1024MB
Storage	4 hard disks: 15GB – OEL, 15GB – Oracle home, 20GB – ASM data disk group, 20GB – ASM log disk group

The virtual machine and the virtual disks should be stored in the data store **ORA11GDBSTORE**. The volume that contains this data store will later be cloned using NetApp FlexClone. With this practice, you ensure that all the needed data for the virtual machine is contained in ORA11GDBSTORE. As a result, by cloning the volume containing ORA11GDBSTORE, you have a clone of a self-contained database virtual machine.

For more details on how to install an Oracle Database 11g virtual machine, refer to the [VMware and Oracle Databases Software Solutions Deployment Guide](#):
www.vmware.com/files/pdf/customers/07Q2_OracleDatabasesSoftwareSolutionsDeployment_wp_tiff_061207v2.pdf.

⁴ This configuration is a sample. You can configure your Oracle Database virtual machine to suit your requirements.

9 NETAPP FLEXCLONE

This section shows how to leverage NetApp FlexClone to create duplicate Oracle Database virtual machines without extra storage consumption.

9.1 CREATE A BASELINE VOLUME FOR FLEXCLONE

NetApp FlexVol volumes allow you to create a Snapshot copy of the volume. The Snapshot copy serves as the baseline volume from which clones of volumes can be created. In our scenario, a virtual machine with a full installation of Oracle Database 11g is deployed in a NetApp volume. In other words, the NetApp volume stores an Oracle Database 11g virtual machine that is self-contained. When the volume is cloned, a full Oracle Database 11g virtual machine is cloned as well.

To create a baseline volume for FlexClone, you will need to create a Snapshot copy of the volume. Follow the procedure below:

1. Power off the virtual machine.
2. telnet to the NetApp array.
3. Execute the following command:

```
snap create -V <volume name> <snapshot name>
```

9.2 FLEXCLONE CREATION

From the Snapshot volume, you can create as many clones as you desire with the NetApp FlexClone technology, without consuming extra storage. Follow the procedure below to create a clone:

1. Create a volume clone:

```
vol clone create <volume clone name> -b <volume name> <snapshot name>
```

2. Map the LUN in the cloned volume to an initiator group that is associated with the ESX Server host:

```
lun map <LUN path>5 <initiator group name>
```

3. Make the LUN available online:

```
lun online <LUN path>
```

Repeat the procedure above as many times as the number of clones you would like to create. After the procedure the cloned LUNs are properly mapped to the initiator group associated with the ESX Server host and brought online. As a result, the ESX Server host will be able to detect the cloned LUNs after a rescan.

10 REGISTER THE VIRTUAL MACHINES

The steps above allow for efficient cloning of the LUN where an Oracle Database virtual machine is located. The entire Oracle Database virtual machine is self-contained within the LUNs, including the configuration (.vmx) file and the virtual disk (.vmdk) files.

The next step is to make ESX Server detect the cloned LUNs and register each of the virtual machines in them. The registration of the virtual machines allows you to use VirtualCenter to manage and operate on the virtual machines.

Preparing VMware ESX Server hosts for cloned databases:

When connecting to cloned data stores, by default the ESX Server hosts do not serve the cloned data. To enable the use of clones, you must enable VMware ESX3 Logical Volume Manager (LVM) to accept cloned data stores.

⁵ Since, in NetApp FAS, a LUN is a file within the volume file system, you can construct a LUN path the way you construct the path to the file representing the LUN.

Follow these steps to enable this functionality:

1. Log on to the ESX Server service console.
2. As root, execute:

```
esxcfg-advcfg -s 1 /LVM/EnableResignature
```

After the volumes are cloned, you have to follow a procedure to make the VMware Infrastructure software recognize the virtual machines contained in the cloned volumes.

The procedure is below:

1. Rescan the storage using VMware Infrastructure Client.
2. In the ESX Server service console, issue the following command to register each of the vmx files:

```
vmware-cmd -s register <path to the .vmx file>
```

11 CUSTOMIZE THE CLONED VIRTUAL MACHINE

In order to use a cloned virtual machine, minimal customization such as network parameter configuration needs to be performed on it.

11.1 SAMPLE CUSTOMIZATION

You can choose to leave the Oracle Database SID unchanged and follow the procedure below to ensure that the cloned database is accessed separately from other databases.

1. Configure the network parameters of the new virtual machine (e.g., IP address, subnet masks, and hostname).
2. Add the IP address to /etc/hosts.
3. Execute the following command to ensure that ASM works properly:

```
$ORACLE_HOME/bin/localconfig reset
```
4. Configure the listener.
Edit listener.ora to reflect the new IP address.

11.2 CHANGE SID

Refer to [MetaLink Note: 15390.1, How to Determine and Change DB_NAME or ORACLE_SID](#).

12 CONCLUSION

VMware Infrastructure provides features for IT administrators to create a self-contained virtual machine for an Oracle Database system. A VMware ESX Server host machine can host a number of virtual machines simultaneously, depending on the requirements of the users. You can easily deploy many Oracle Database virtual machines on VMware Infrastructure simultaneously by leveraging the VMware features in addition to NetApp FlexClone. NetApp FlexClone instantaneously clones a volume containing the virtual machine files without consuming extra storage. With VMware Infrastructure and NetApp FlexClone, you can reduce the provisioning time for database clones tremendously while reaping the benefit of greatly reduced storage consumption.

13 RESOURCES

You can find more information about NetApp, VMware, and Oracle products at the Web sites listed here:

NetApp official Web site: www.netapp.com

VMware official Web site: www.vmware.com

Oracle official Web site: www.oracle.com

13.1 NETAPP

Network Appliance and VMware Virtual Infrastructure 3 Storage Best Practices:

www.netapp.com/library/tr/3428.pdf

DB2: Cloning a Database Using NetApp FlexClone Technology:

www.netapp.com/library/tr/3460.pdf

Network Appliance and VMware Virtual Desktop Infrastructure:

www.netapp.com/library/tr/3612.pdf

13.2 VMWARE

VMware Infrastructure 3 product Web site:

www.vmware.com/products/data_center.html

VMware Infrastructure 3 download Web site:

www.vmware.com/download/vi/eval.html

VMware support Web site:

www.vmware.com/vmtn/

VMware best practices papers:

www.vmware.com/pdf/vsmp_best_practices.pdf

www.vmware.com/pdf/vi_performance_tuning.pdf

VMware products:

www.vmware.com/products/

VMware Infrastructure 3 system compatibility guide:

www.vmware.com/pdf/vi3_systems_guide.pdf

VMware converter:

www.vmware.com/download/converter/

13.3 ORACLE

Oracle Enterprise Linux download Web site:

<http://edelivery.oracle.com/linux>

Oracle support Web site:

www.oracle.com/technology/about/index.html

14 APPENDIX: EXAMPLE SCRIPTS TO PERFORM FLEXCLONE AND VMWARE VIRTUAL MACHINE REGISTRATION

vol_clone.sh – connects to a NetApp storage system and creates COUNT number of volume clones using the FlexClone feature.

```
#!/bin/bash
#
# File Name: vol_clone.sh
#
# This sample code connects to a NetApp storage system and creates COUNT
# number of
# volume clones using the FlexClone feature
#
# This sample code is provided AS IS, with no support or warranties of any
# kind, including but not limited to warranties of merchantability or
# fitness of any kind, expressed or implied.
#
# This script requires supplemental perl scripts, snap_vol.pl and
# clone_vol.pl, which are listed separately
#
# 2007 Desmond Chan, VMware Inc.
#
VERBOSE=0
SUFFIX=0
VERBOSE=0
COUNT=0
STORAGESYSTEM=""
VOL=""
LUN=""
SNAPNAME=""
#
# Name: snap_volume
# Func: Sends a command to the NetApp storage system to create a Snapshot
# copy in the provided name
# of the volume provided.
#
snap_volume()
{
    ${PERL} snap_vol.pl ${STORAGESYSTEM} "snap create -V $VOL $SNAPNAME"
}
#
# Name: clone_volume
# Func: Clones the volume as many time as the value of COUNT with the suffix
# incremented by 1 each time.
#
clone_volume()
{
    while [ ${COUNT} -gt 0 ]; do
        NEWVOLNAME="${VOL}fclone${SUFFIX}"
        NEWLUNPATH="/vol/${NEWWOLNAME}/${LUN}"
        echo $NEWWOLNAME $NEWLUNPATH
        let SUFFIX+=1
        let COUNT-=1
        ${PERL} clone_vol.pl ${STORAGESYSTEM} "vol clone create
$NEWWOLNAME -s none -b $VOL $SNAPNAME" "lun map $NEWLUNPATH esxriker" "lun
online $NEWLUNPATH" ;
        echo $NEWWOLNAME "has been cloned on NetApp filer" ${STORAGESYSTEM}
    done
}
```

```

}

#
#MAIN SCRIPT
#
while getopts "f:o:c:v:l:s:" flag
do
    case "$flag" in
        f)      STORAGESYSTEM=$OPTARG;;
        o)      VOL=$OPTARG;;
        c)      COUNT=$OPTARG;;
        l)      LUN=$OPTARG;;
        s)      SUFFIX=$OPTARG;;
        v)      VERBOSE=1;;
    esac
done
SNAPNAME="${VOL}_dbsnp01"
snap_volume
clone_volume
echo "Script completed"

```

snap_vol.pl – Perl script to send a snap command to the NetApp storage system, via telnet. The snap command and the storage system address are provided in the arguments.

```

#!/usr/bin/perl
#
# File Name: snap_vol.pl
#
# This sample code sends a snap command to a NetApp storage system, via
# telnet. The snap command and
# the storage system address are provided in the arguments
#
# This sample code is provided AS IS, with no support or warranties of any
# kind, including but not limited to warranties of merchantability or
# fitness of any kind, expressed or implied.
#
# 2007 Desmond Chan, VMware Inc.
#
use Expect;
$host=$ARGV[0];
$snap_cmd=$ARGV[1];
print "host=$host\n";
print "snap_cmd=$snap_cmd\n";
$session=Expect->spawn("telnet $host");
$session->expect(30, "logged in") || die "Error waiting for filer prompt,
".$session->exp_error()."\n";
$session->send($snap_cmd."\n");
$session->expect(30, ">") || die "Error waiting for filer prompt,
".$session->exp_error()."\n";
$session->close();

```

clone_vol.pl – Perl script to send commands to the NetApp storage system, via telnet, to clone a volume, map the cloned LUN, and bring the cloned LUN online.

```

#!/usr/bin/perl
#
# File Name: clone_vol.pl
#
# This sample code sends commands to the NetApp storage system, via telnet,
# to clone a volume,
# map the cloned LUN, and bring the cloned LUN online.
#
#
# This sample code is provided AS IS, with no support or warranties of any

```

```

# kind, including but not limited to warranties of merchantability or
# fitness of any kind, expressed or implied.
#
# 2007 Desmond Chan, VMware Inc.
#
use Expect;
$host=$ARGV[0];
$vol_clone_prompt=$ARGV[1];
$lun_map_prompt=$ARGV[2];
$lun_online_prompt=$ARGV[3];
print "host=$host\n";
print "vol_clone_prompt=$vol_clone_prompt\n";
print "lun_map_prompt=$lun_map_prompt\n";
print "lun_online_prompt=$lun_online_prompt\n";
$session=Expect->spawn("telnet $host");
$session->expect(30, "logged in") || die "Error waiting for filer prompt,
".$session->exp_error()."\n";
# clone the LUN
$session->send($vol_clone_prompt."\n");
$session->expect(30, "has completed.") || die "Error waiting for filer
prompt, ".$session->exp_error()."\n";
$session->expect(30, "or move operation") || die "Error waiting for filer
prompt, ".$session->exp_error()."\n";
# map the cloned LUN to the initiator group
$session->send($lun_map_prompt."\n");
$session->expect(30, "auto-assigned") || die "Error waiting for filer
prompt, ".$session->exp_error()."\n";
$session->expect(30, "> ") || die "Error waiting for filer prompt,
".$session->exp_error()."\n";
# bring the cloned LUN online
$session->send($lun_online_prompt."\n");
$session->expect(30, "> ") || die "Error waiting for filer prompt,
".$session->exp_error()."\n";
$session->close();

```

register_vms.sh – connects to an ESX Server and registers the virtual machines whose configuration file names match the pattern provided.

```

#!/bin/bash
#
# File Name: register_vms.sh
#
# This sample code connects to an ESX Server and registers the virtual
machines whose configuration
# file names match the pattern provided.
#
# This sample code is provided AS IS, with no support or warranties of any
# kind, including but not limited to warranties of merchantability or
# fitness of any kind, expressed or implied.
#
# 2007 Desmond Chan, VMware Inc.
#
RSH=/usr/bin/ssh
VERBOSE=0
RSH_USER=root
RSH_CMD="${RSH} -l ${RSH_USER}"
SUFFIX=0
VERBOSE=0
VOL=""
LUN=""
ESXHOST=""
#
# Name: esx_register_vms

```

```

# Func: Registers the virtual machines, whose configuration file names match
# the pattern provided
# in VMX_NAME to the ESX Server
#
esx_register_vms()
{
    VMX_NAME=$1
    VMXES=`$RSH_CMD ${ESXHOST} "find /vmfs/volumes -mindepth 3 -maxdepth 3
-name '${VMX_NAME}' -print"`
    if [ "${VMXES}" != "" ]; then
        for VMX in ${VMXES}; do
            ${RSH_CMD} ${ESXHOST} "/usr/bin/vmware-cmd -s register
${VMX}"
            if [ $? != 0 ]; then
                echo "esx_register_vms(): ERROR: register VM ${VMX}
failed."
            fi
        done
    fi
}
#
#MAIN SCRIPT
#
while getopts "ve:x:" flag
do
    case "$flag" in
        e)      ESXHOST=$OPTARG;;
        x)      VMX_NAME=$OPTARG;;
        v)      VERBOSE=1;;
    esac
done
esx_register_vms $VMX_NAME
echo "Script completed"

```

