



Technical Report

Oracle Dev/Test on VMware vSphere and NetApp Storage Solutions Overview

Solution guide for rapid deployment of space-efficient Oracle Database copies
using VMware vSphere and NetApp FlexClone

NetApp and VMware, Inc.
July 2010 | TR-3861 | Version 1.0

TABLE OF CONTENTS

1	INTRODUCTION	3
1.1	PURPOSE OF THE DOCUMENT	3
1.2	TARGET AUDIENCE	3
1.3	PREREQUISITES	3
1.4	BACKGROUND	3
2	SOLUTION OVERVIEW	5
2.1	BUSINESS PROBLEM	5
2.2	THE NETAPP AND VMWARE SOLUTION—A BETTER APPROACH	5
2.3	ORACLE DEV/TEST SOLUTION ARCHITECTURE FROM NETAPP—THE BIG PICTURE	6
3	SOLUTION SUMMARY	8
3.1	SOLUTION VERSION 1: ORACLE DEV/TEST ENVIRONMENT IN A SELF-CONTAINED VIRTUAL MACHINE	8
3.2	SOLUTION VERSION 2: PRODUCTION ORACLE DATABASE CLONES MOUNTED ON VIRTUAL MACHINE CLONES	9
3.3	SOLUTION COMPONENTS	10
3.4	NETAPP SOFTWARE PRODUCTS USED IN THE SOLUTION	11
4	SOLUTION ARCHITECTURE—SOLUTION VERSION 1	12
4.1	ARCHITECTURE OVERVIEW	12
4.2	HARDWARE AND SOFTWARE COMPONENTS USED IN THE LAB SETUP	12
4.3	SOLUTION IMPLEMENTATION	13
4.4	SOLUTION VALIDATION—SOLUTION VERSION 1	35
5	SOLUTION ARCHITECTURE—SOLUTION VERSION 2	37
5.1	ARCHITECTURE OVERVIEW	37
5.2	HARDWARE AND SOFTWARE USED IN THE LAB SETUP	37
5.3	SOLUTION IMPLEMENTATION	38
5.4	SOLUTION VALIDATION—SOLUTION VERSION 2	51
6	SUMMARY	52
7	REFERENCES	52
8	ACKNOWLEDGEMENTS	53

1 INTRODUCTION

1.1 PURPOSE OF THE DOCUMENT

This solution overview document presents a rapid and efficient way of deploying Oracle® dev/test environments on virtual server and storage infrastructure consisting of VMware® vSphere™ and NetApp® storage.

1.2 TARGET AUDIENCE

This document is for:

- Oracle DBAs
- VMware administrators
- Storage administrators
- Infrastructure architects and consultants
- Data center managers

1.3 PREREQUISITES

This document assumes familiarity with:

- Server virtualization technologies from VMware and their integration with NetApp storage
- NetApp storage systems and storage efficiency technologies based on Data ONTAP®

The following documents may be helpful for this purpose:

Document Name	Location
"NetApp and VMware vSphere Storage Best Practices"	http://media.netapp.com/documents/tr-3749.pdf
"Helping DBAs Become More Efficient—NetApp Efficiency and Manageability Advantages"	http://media.netapp.com/documents/wp-7079.pdf
"Rapid Database Development and Deployment"	http://media.netapp.com/documents/wp-7014.pdf

1.4 BACKGROUND

1.4.1 NETAPP STORAGE EFFICIENCY TECHNOLOGIES

NetApp offers a full range of manageability and data efficiency technologies to address database management problems, including technologies for thin provisioning, thin cloning, deduplication, replication, and transparent data movement. In particular, technologies that help you manage and use storage more efficiently have become an important focus for NetApp. Table 1 shows the key NetApp storage efficiency technologies that can positively impact database management.

Table 1) NetApp storage efficiency technologies.

NetApp Storage Efficiency Technology	Benefit
Thin provisioning	Storage no longer has to be allocated to a volume up front. Storage is treated as a shared resource, and capacity is consumed only as it is needed.
Thin cloning	Clone complete data sets in seconds without full data copies using NetApp FlexClone®. Storage space is only consumed incrementally as changes occur.
Snapshot™	NetApp Snapshot technology delivers space savings over competing products. Create up to 255 Snapshot copies without excessive space penalties and virtually no performance penalty.

Deduplication	NetApp deduplication identifies and eliminates redundancy at the block level in both primary and secondary storage.
Thin replication	SnapMirror® and SnapVault® software perform only incremental block transfers—Thin transfers—accelerating replication and conserving bandwidth.
Compression	Compression can be used with NetApp's Thin Replication technologies to reduce The size of the replicated data, saving network bandwidth.

Refer to the following documents for a detailed description of the above NetApp storage efficiency technologies and to know how they can be used in the context of an Oracle Database.

Document Name	Location
Storage Efficiency Management Guide in NOW™	http://now.netapp.com/NOW/knowledge/docs/ontap/rel733/pdf/s/ontap/semg.pdf
Storage efficiency in an Oracle environment	http://media.netapp.com/documents/wp-7084.pdf

1.4.2 REDUCE COSTS AND MAXIMIZE IT EFFICIENCY WITH VMWARE VSPHERE

VMware vSphere can transform your IT infrastructure into a private cloud, then bridge on demand to public clouds, delivering IT infrastructure as easily accessible services. It delivers uncompromised control over all IT resources with the highest efficiency and choice in the industry.



Figure 1) VMware vSphere (Graphic provided by VMware).

You can also dramatically reduce capital and operational expenses while minimizing any potential lost revenue associated with downtime, outages, and failures with VMware vSphere. It reduces capital and operational expenses per application by more than 50% through virtualization of server, storage, and networking resources, reducing capital and operational expenses per application by more than 50%.

Organizations can achieve consolidation ratios of more than 15:1 while increasing efficiency through management automation and the dynamic allocation of resources to enterprise applications across private cloud and public cloud infrastructures.

2 SOLUTION OVERVIEW

2.1 BUSINESS PROBLEM

From small and midsize businesses to global corporations, demand for database application development continues to grow at a dramatic rate. Database administrators (DBAs) struggle to meet stricter service-level agreements while managing an ever-expanding, complex database landscape. There is an increased need to support the complex and resource-consuming requirements of database application development, even as these development environments grow increasingly costly and complicated. Full or partial production database copies are required for many activities, including application development and testing (dev/test), maintenance activities (operating system and database upgrades), QA, training and demos, and reporting and data warehouse ETL (extract, transform, load) processes.

The database cloning process starts with a request to the DBA from the application development team for a copy of one or more production databases to support application development and testing. Often, many copies of the production database are required to support multiple code streams as part of a modular development process. In addition, the distribution of development resources and other operational requirements may demand that many database copies be created on separate storage systems, often located in several geographically dispersed locations. DBAs are expected to support these requirements quickly and efficiently, without disrupting the production database environment or affecting overall system performance. They are also expected to create clones with limited storage resources and only during short or nonexistent “nonpeak” periods.

The ability to clone production databases quickly, accurately, and efficiently directly affects application development, test, and deployment. To support database application development and test while creating multiple database copies, DBAs require:

- Significant (100%) additional storage capacity for each development copy
- Long periods of time for manual replication, backup, and migration operations
- No production system downtime or performance degradation during copying
- Regular assistance (budget, time, and expertise) from the storage administrator

The traditional approach to cloning a database is a heavyweight activity. Since every copy requires 100% of the storage capacity of the original database, it places a major burden on storage resources. Getting necessary approvals, working with the storage admin to locate storage capacity, and copying the data also typically take a long time to accomplish. The only alternative to having a full clone, however, is to do testing with only a subset of your production data, and that can often lead to inconsistent results. Because full database copies are an expensive resource, they are always in short supply and often not as current as you would like them to be. DBAs, developers, testers, and line-of-business specialists all queue up waiting on the limited resources.

In addition to the above obstacles for creating clone database copies, another aspect also needs to be considered for implementing a database dev/test environment: provisioning of adequate server and storage resources once the database cloning process has been decided. The servers must be loaded with the appropriate software “stack” (operating system, drivers, database, application servers, and so on), and the database must be restored or created and populated, often using manual cloning scripts that are complex to write and subject to error. This traditional provisioning process is often done manually and takes significant time to complete.

2.2 THE NETAPP AND VMWARE SOLUTION—A BETTER APPROACH

NetApp solutions can dramatically reduce the time, resources, and effort required for rapid and cost-effective database cloning and deployment. DBAs and application developers can leverage powerful, automated data and storage management software from NetApp to easily and quickly create local or remote database clones as needed. NetApp clones require less storage than traditional clones, and they can be created without the need for storage administrator support. The NetApp Unified Storage Architecture, coupled with NetApp V-Series storage controllers, also provides simplified management and support across all storage

topologies and vendor arrays, including Fibre Channel SAN, IP SAN, and NAS (CIFS for Windows[®] environments and NFS for UNIX[®] and Linux[®] environments). In addition, the ability to leverage NetApp's storage efficiency technologies from the VMware vSphere layer itself (through VMware vCenter[™] Server invoking services of the storage layer) has enabled rapid and efficient provisioning of the server and storage resources for the dev/test environment.

Server virtualization technologies provided by VMware vSphere 4 allow you to create a base virtual machine that runs an Oracle Database properly installed and tested. This base virtual machine, which can actually mimic a real production Oracle Database system, is stored as a set of files in a shared storage system, that is, a NetApp FAS. It is self-contained, including a guest operating system, Oracle home, and all the data and log partitions. So cloning the virtual machine using NetApp's storage-efficient thin cloning technologies gives you all you need for a functional dev/test Oracle Database system.

Another approach to this solution can be to create cloned dev/test virtual machines containing just the guest operating system and the Oracle binaries and then to mount the Oracle Database clones created from the production database instance using NetApp thin cloning technologies on the cloned virtual machines.

Oracle DBAs and VMware administrators can use automated and powerful tools from NetApp like SnapManager[®] for Oracle (SMO) and Rapid Cloning Utility (RCU) to easily and quickly create database and virtual machine clones, using negligible additional storage and without the need for extensive storage skills or storage administrator support.

Note:

1. NetApp thin cloning technologies like FlexClone can alleviate the resource contention that results from not having enough database copies. With FlexClone you can create multiple clones using minimal incremental storage space. For example, you can create 30 database clones in minutes for the same space as three database copies. As a result, everyone who needs database copies can gain immediate access, so productivity can increase.
2. Installing the Oracle Database to be cloned inside a virtual machine running on VMware vSphere and then cloning the virtual machine using FlexClone technology can even further simplify the dev/test environment creation process. This eliminates the need for testers or developers to have physical hosts on which the database clones need to be mounted and also can be automated using NetApp Rapid Cloning Utility from the VMware vCenter Server.
3. The NetApp SnapManager for Oracle product can be integrated seamlessly into this solution and can be used to create and mount production Oracle Database clones into the cloned virtual machines created using RCU.

2.3 ORACLE DEV/TEST SOLUTION ARCHITECTURE FROM NETAPP—THE BIG PICTURE

Figure 2 below illustrates a typical deployment scenario of an Oracle dev/test solution using NetApp's storage efficiency technologies. It also shows where the solution described in this document fits into the big picture.

It is important to observe the following points corresponding to this approach to database cloning and development:

- Working clones of very large databases can be created in just a few minutes and with minimal incremental storage by using SnapManager for Oracle and NetApp FlexClone technology. A typical 300GB database can be cloned and brought up in just a couple of minutes.
- Production data can be continually replicated (synchronously or asynchronously) to a secondary storage system with minimal impact on the production environment and used as the basis for the cloning process.
- The production databases residing on other vendors' arrays also can be mirrored to the NetApp dev/test storage system using Oracle Data Guard.
- DBAs can streamline and automate database backup, recovery, and cloning by using SnapManager for Oracle data management capabilities. This significantly reduces the time required to provision database clones for development and operational scenarios, without the need for extensive storage expertise or support.
- When V-Series open storage controllers are used to manage arrays from EMC, IBM, HP, Hitachi Data Systems, and others, FlexClone technology can be used to instantly replicate copies of databases.

NetApp clones increase productivity and save storage space without compromising performance across other vendors' storage.

- The new database servers are dynamically provisioned—they can be either physical (for example, a blade server pool) or virtual (for example, cloned virtual machines on VMware vSphere using NetApp Rapid Cloning Utility). Once created, these dev/test servers can be connected to the database clones on the secondary (dev/test) storage system. One of the solutions described in this document covers this scenario.

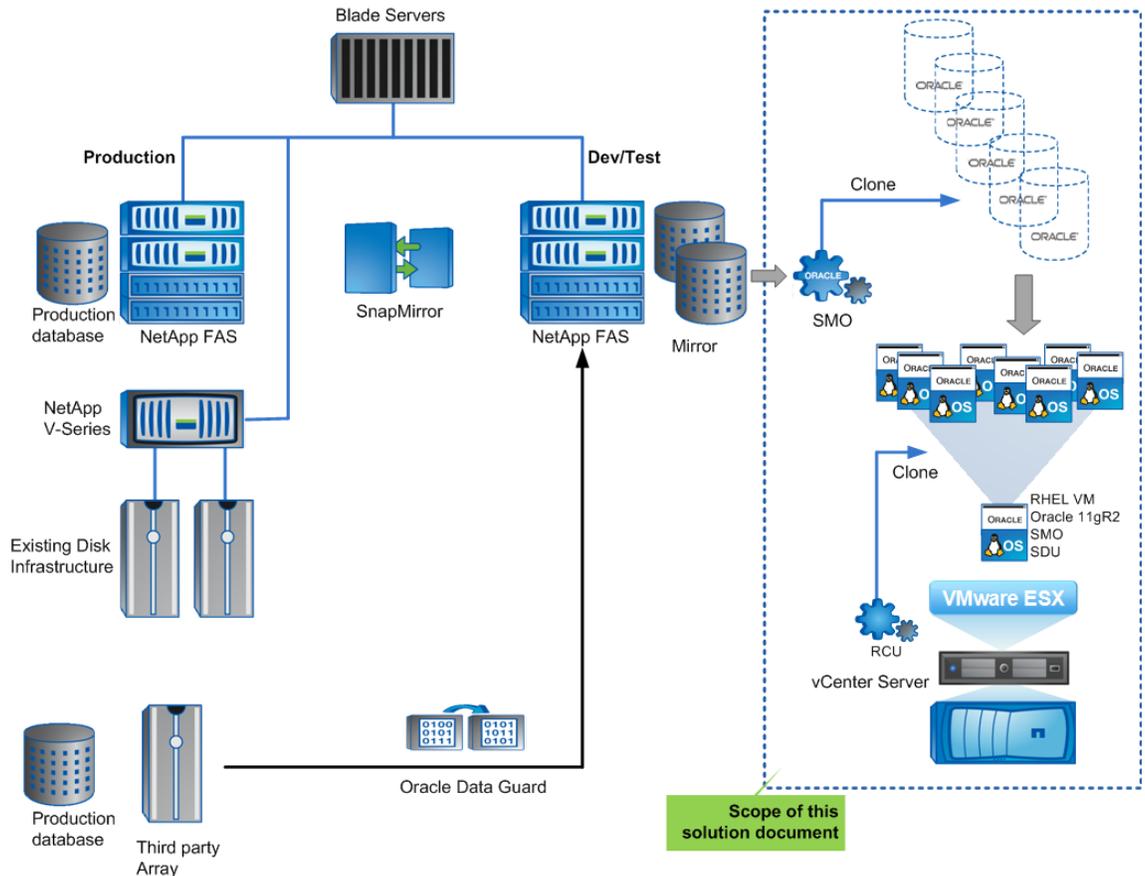


Figure 2) Oracle dev/test solution architecture.

Result: Simplified, rapid, and cost-effective database cloning and development, with faster provisioning times and greater utilization for servers and storage.

The portion of the solution marked within dotted lines in the above diagram is the focus area of this document.

For more details regarding the other parts of this overall dev/test solution landscape, see <http://media.netapp.com/documents/ra-0002.pdf>.

3 SOLUTION SUMMARY

This document presents two versions of an extremely efficient Oracle dev/test solution powered by proven NetApp storage efficiency technologies.

Figure 3 illustrates the scenarios in which different versions of this solution can be used.

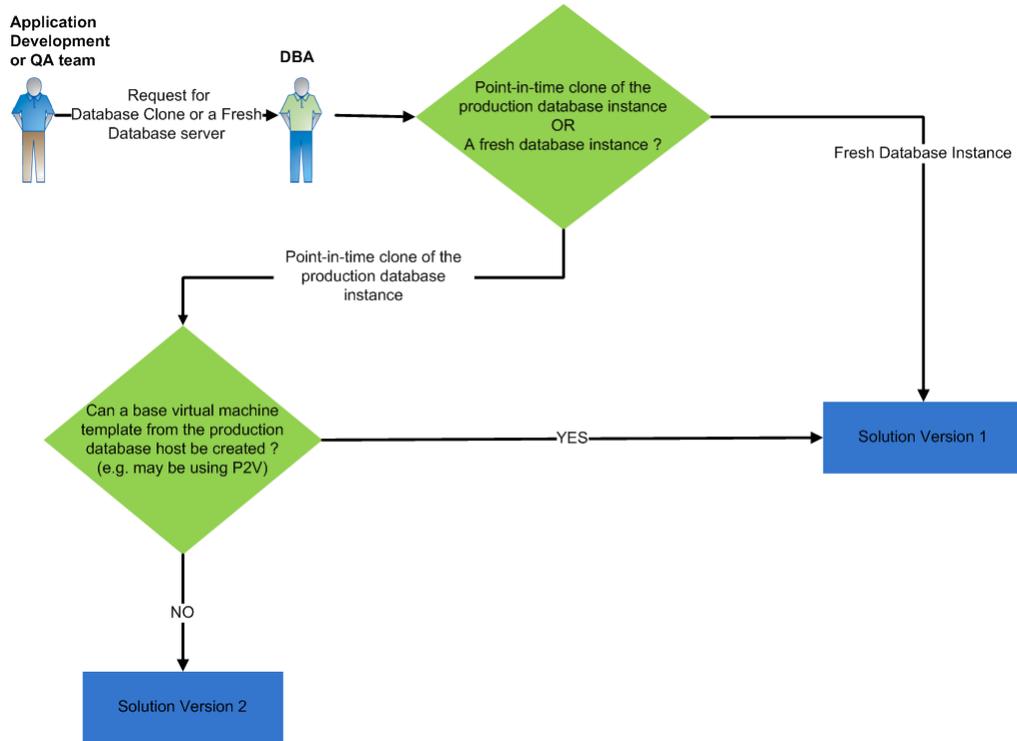


Figure 3) Solution summary.

3.1 SOLUTION VERSION 1: ORACLE DEV/TEST ENVIRONMENT IN A SELF-CONTAINED VIRTUAL MACHINE

Scenario: Oracle Database to be cloned is installed inside a golden virtual machine hosted on a VMware and NetApp infrastructure. Using NetApp Rapid Cloning Utility, which is installed as a plug-in to VMware vCenter Server, clones of that golden virtual machine are created.

Technologies and products used:

- NetApp thin cloning technologies (FlexClone)
- VMware vSphere and vCenter Server
- NetApp RCU plug-in

Figure 4 illustrates how various components fit together in the overall solution architecture for version 1.

Although version 1 of this solution is extremely simple to implement, the Oracle Database clones created may not be as current as the production database (that is, there may be a considerable time lag between the refresh of production data in the database contained inside the golden virtual machine from which cloned dev/test virtual machines are created and the actual production Oracle Database server). This version of the solution is suitable for situations in which fresh Oracle Database servers need to be provisioned for the dev/test teams or when the database inside the dev/test virtual machine need not be made current very frequently.

Version 2 of this solution overcomes this limitation and can create a dev/test environment that can be as current as the production Oracle Database environment.

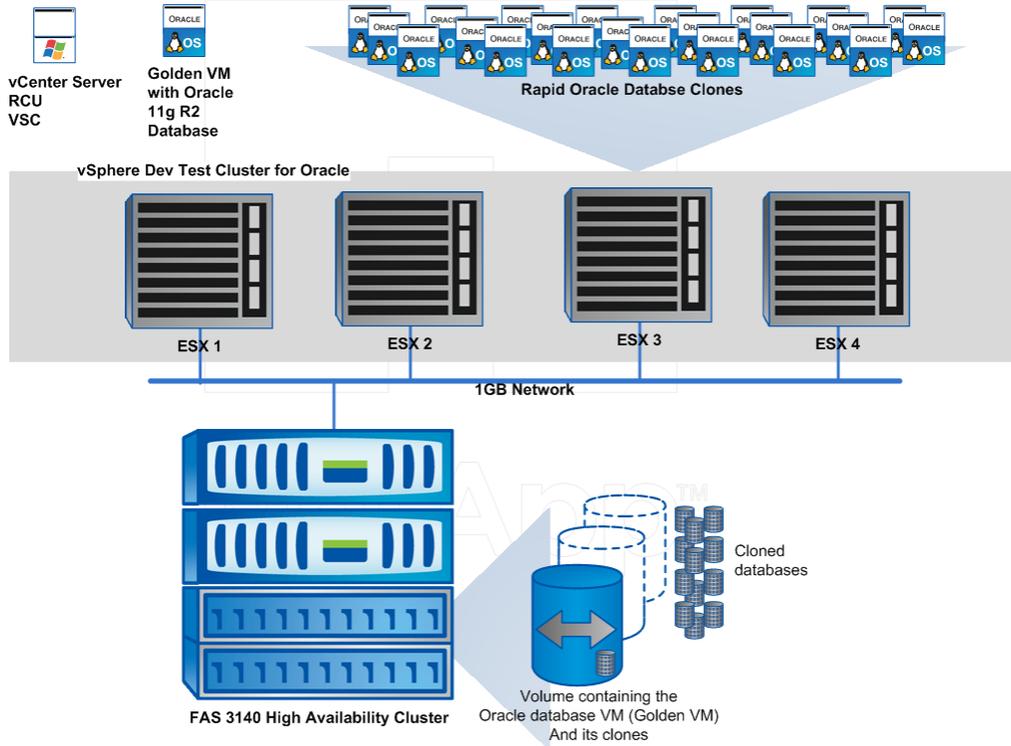


Figure 4) Solution 1 implementation.

3.2 SOLUTION VERSION 2: PRODUCTION ORACLE DATABASE CLONES MOUNTED ON VIRTUAL MACHINE CLONES

Scenario: Oracle Database clones created using SnapManager for Oracle from the production Oracle Database server are mounted on the virtual machine clones created using NetApp Rapid Cloning Utility.

Technologies and products used:

- NetApp thin cloning technologies (FlexClone)
- VMware vSphere and vCenter Server
- NetApp RCU plug-in
- NetApp SnapManager for Oracle
- SnapDrive® for UNIX (SDU)

Figure 5 illustrates how various components fit together in the overall solution architecture for version 2.

Note:

- The production Oracle Database server is shown as a physical server in the diagram below. However, it also can be a virtual server.
- The SMO repository is shown in the same host as that of the production Oracle Database server. However, the SMO repository and the production Oracle Database server can be located on separate hosts as well.

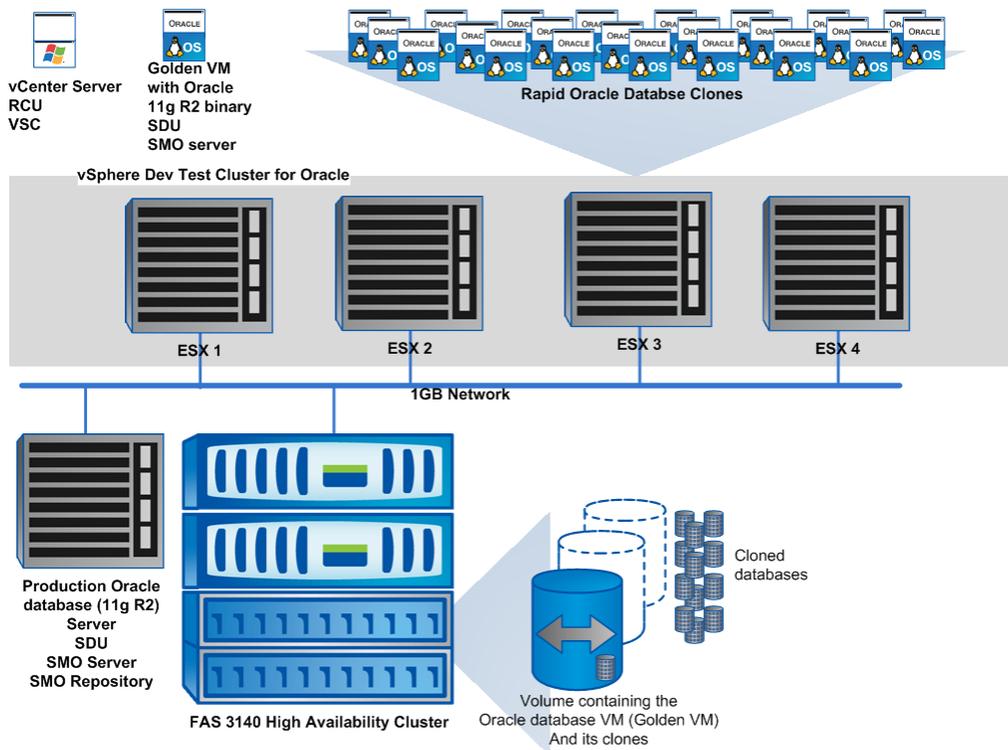


Figure 5) Solution 2 implementation.

3.3 SOLUTION COMPONENTS

Error! Reference source not found.2 summarizes the high-level components of the Oracle dev/test solution on VMware vSphere and NetApp storage described in this document.

Table 2) Solution components.

Component	Details
VMware Version	VMware vSphere 4
Server Platform	Any Intel [®] or AMD server with AMD-v or Inter-VT support
Storage Model	NetApp FAS3140 (active-active), Data ONTAP 7.3.3
Guest Virtual Machines for Dev/Test	Oracle Database 11g [™] R2 Single Instance (64-bit) on RHEL5U4 x86_64 1GB RAM 1 vCPU
NetApp Rapid Cloning Utility 3.0	Plug-in to VMware vCenter Server; needs FlexClone and ASIS license on the storage controller
NetApp Virtual Storage Console (VSC) 1.0	Plug-in to VMware vCenter Server
NetApp SnapManager for Oracle 3.0.3	Utility for creating Oracle Database clones
SnapDrive for UNIX 4.1.1	Utility for providing host-based data management of NetApp Storage from Linux servers

3.4 NETAPP SOFTWARE PRODUCTS USED IN THE SOLUTION

3.4.1 RAPID CLONING UTILITY

The Rapid Cloning Utility provides a means to optimally provision and manage NetApp SAN and NAS datastores along with a means to provision zero cost virtual machines directly from within vCenter. Below is a list of the core abilities of the RCU:

- Provision FC, FCoE, iSCSI, and NFS datastores
- Automated assignment of multipathing policies to datastores with the VMware pluggable storage architecture for LUNs and ALUA-enabled LUNs and distributing NFS path connections based on a path round-robin policy
- Automated storage access control security implemented when datastores are provisioned; access control is in the form of LUN masking and NFS exports
- Dynamically resize FC, FCoE, iSCSI, and NFS datastores on the fly
- Provides a central interface to collect data from storage controller, network switches, and ESX®/ESXi™ hosts in order to aid in the resolution of I/O-related case issues

For more details regarding RCU, refer to:

“Rapid Cloning Utility Installation and Administration Guide” in <http://now.netapp.com>

“NetApp and VMware vSphere—Storage Best Practices Guide” in <http://media.netapp.com/documents/tr-3749.pdf>

3.4.2 VIRTUAL STORAGE CONSOLE

Virtual Storage Console provides optimal availability and performance with ESX/ESXi hosts. The core abilities of the VSC are:

- Identifying and configuring optimal I/O settings for FC, FCoE, iSCSI, and NFS in ESX/ESXi hosts
- Identifying and configuring optimal path selection policies for existing FC, FCoE, and iSCSI datastores
- Monitoring and reporting storage utilization at levels from the datastore to the physical disks
- Provides a central interface to collect data from storage controller, network switches, and ESX/ESXi hosts to aid in the resolution of I/O-related case issues

For more details regarding RCU, refer to:

“Virtual Storage Console Installation and Administration Guide” in <http://now.netapp.com>

“NetApp and VMware vSphere—Storage Best Practices Guide” in <http://media.netapp.com/documents/tr-3749.pdf>

3.4.3 SNAPMANAGER FOR ORACLE

SnapManager for Oracle automates and simplifies the complex, manual, and time-consuming processes associated with the backup, recovery, and cloning of Oracle Databases running on NetApp storage systems.

SnapManager for Oracle software enables efficient clone creation in a matter of minutes, giving DBAs greater flexibility to create many more clones than before, even with very large databases. It enables a more dynamic, fast-paced environment in which clones are easily created and disposed of to support the faster development cycles that customers demand today.

For more details regarding SMO, refer to:

“SnapManager for Oracle Best Practices”: <http://media.netapp.com/documents/tr-3761.pdf>

3.4.4 SNAPDRIVE FOR UNIX

SnapDrive for UNIX provides host-based data management of NetApp storage from Linux servers (either physical or virtual).

For more details regarding SDU, refer to:

“SnapDrive for UNIX Installation and Administration Guide” in [NOW](#)

4 SOLUTION ARCHITECTURE—SOLUTION VERSION 1

4.1 ARCHITECTURE OVERVIEW

Figure 6 illustrates the architecture for this dev/test solution and the process followed to create the clone virtual machines. The numbers in the diagram correspond to the solution implementation steps described in section 4.3.

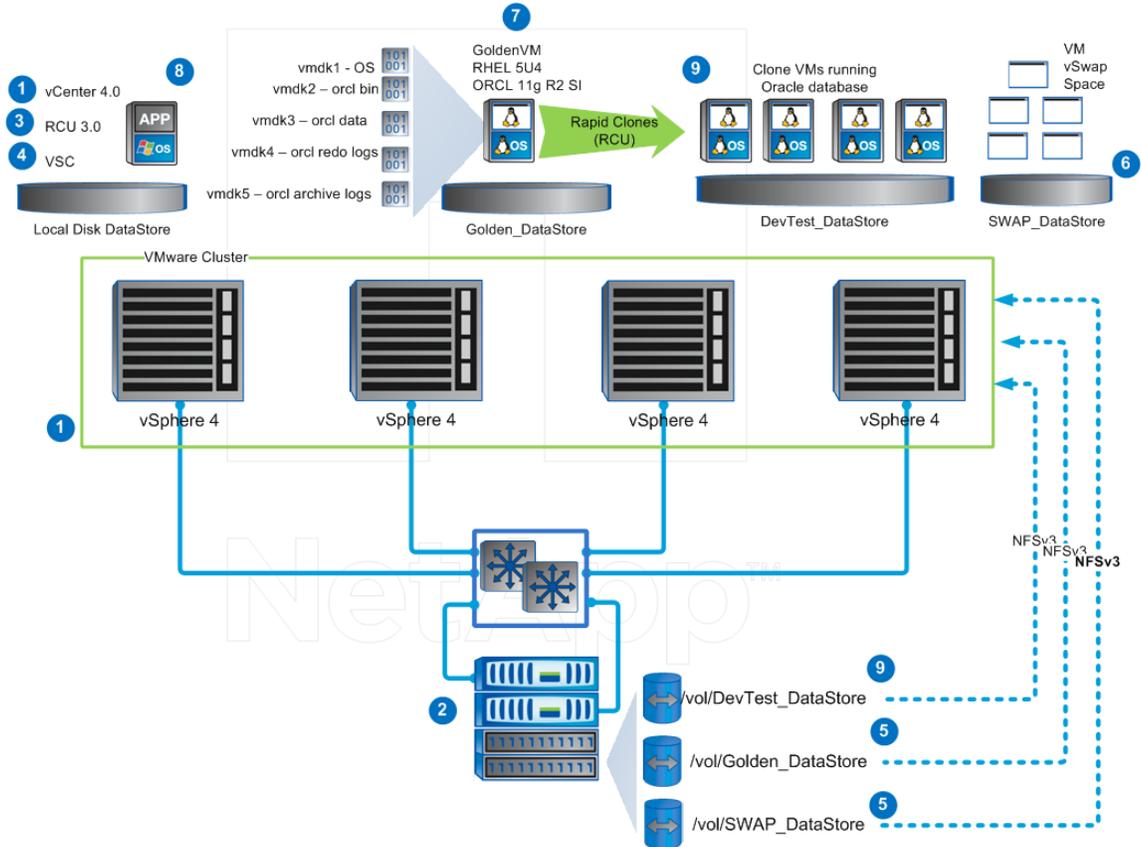


Figure 6) Architecture overview.

4.2 HARDWARE AND SOFTWARE COMPONENTS USED IN THE LAB SETUP

Table 3) Components used.

Infrastructure Component	Vendor	Quantity	Details
Server	IBM	1	IBM x3550 server
		3	IBM x3650 server
			Intel Xeon™ processor (Intel-VT), CPU: 38 GHz total in the cluster Memory: 24GB total in the cluster
Storage	NetApp		FAS3140: Active-active config, RAID-DP®
Network adapter	Broadcom	2/ server	Broadcom NetXtreme II BCM 5708 1000Base-T
HBA	N/A	N/A	N/A
Network switch	Cisco	2	
Software	NetApp		Data ONTAP 7.3.3
	NetApp		Cluster
	NetApp		NFS

	NetApp		FlexClone
	NetApp		ASIS
	NetApp		Rapid Cloning Utility 3.0 (VMware vCenter plug-in)
	NetApp		VSC (VMware vCenter plug-in)
	VMware		VMware ESX v4
	VMware		VMware vCenter Server 4
	Oracle		Oracle Database 11g R2 Single Instance
	Red Hat		RHEL 5 U4

4.3 SOLUTION IMPLEMENTATION

Table 4 illustrates the steps involved in implementing the Oracle dev/test solution on VMware vSphere and NetApp storage.

Table 4) Solution implementation.

#	Steps	Comments
1	VMware ESX and vCenter Server setup and configuration	Installation and configuration of the VMware vSphere 4 ESX servers. Installation and configuration of VMware vCenter Server 4. Configuration of VMware cluster and networking (vDS).
2	NetApp storage setup	Installation and configuration of the NetApp storage systems (NetApp FAS3140 active-active)
3	RCU setup	Installation and configuration of NetApp Rapid Cloning Utility 3.0 in the VMware vCenter Server (vCenter plug-in)
4	VSC setup	Installation and configuration of NetApp VSC 1.0 (VMware Storage Console) in the VMware vCenter Server (vCenter plug-in)
5	NFS datastore creation	Create 'Golden_DataStore' and 'SWAP_DataStore' using RCU from the same aggregate of the NetApp storage. Set correct NFS tunables on the datastores created using VSC.
6	vSwap location setup	Make the changes in the ESX servers so that all virtual machine swap space (vSwap) is stored on 'SWAP_DataStore'
7	Golden virtual machine creation	Create the golden virtual machine with the Oracle Database 11g R2 (single instance) on 'Golden_DataStore.' Align the vmdk files of the virtual machine using mbrscan and mbralign (or fdisk).
8	Custom specification script	Create a "custom specification" on the vCenter Server for the clone virtual machines to be created so that they get a unique host name and the users are prompted to enter the IP address when they boot up.
9	Creation of clone virtual machines	Create 25 clones of the Golden virtual machine using RCU on a new datastore – 'DevTest_DataStore'

4.3.1 STEP 1: VMWARE ESX AND VCENTER SERVER SETUP AND CONFIGURATION

Install and configure the VMware ESX servers that will be the part of Oracle dev/test setup.

Also install the latest version of VMware vCenter Server.

Refer to [VMware documentation](#) for steps for installing and configuring VMware ESX and vCenter Server:

"ESX and vCenter Server Installation Guide, Update 1- ESX 4.0," "vCenter Server 4.0"

http://www.vmware.com/pdf/vsphere4/r40_u1/vsp_40_u1_esx_vc_installation_guide.pdf

For the sample dev/test setup used in this document:

- Four servers were installed with the latest version of VMware vSphere 4.
- vCenter Server version 4.0 was installed as a virtual machine in one of the ESX servers installed on the local storage.
- A VMware Virtual Distributed Switch (vDS) was configured for ESX server networking and shared storage access.

Best Practice: Follow the best practice configurations recommended by [TR- 3749](#) for deploying VMware vSphere servers with NetApp storage systems for production deployments.

- Refer to section 3.1 and 3.2 for the details of hardware and software components used.

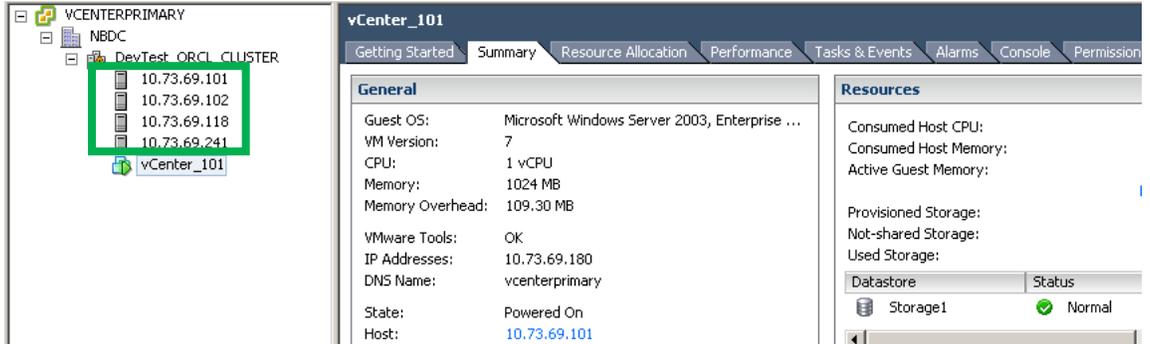


Figure 7) VMware ESX Server and vCenter Server setup and configuration.

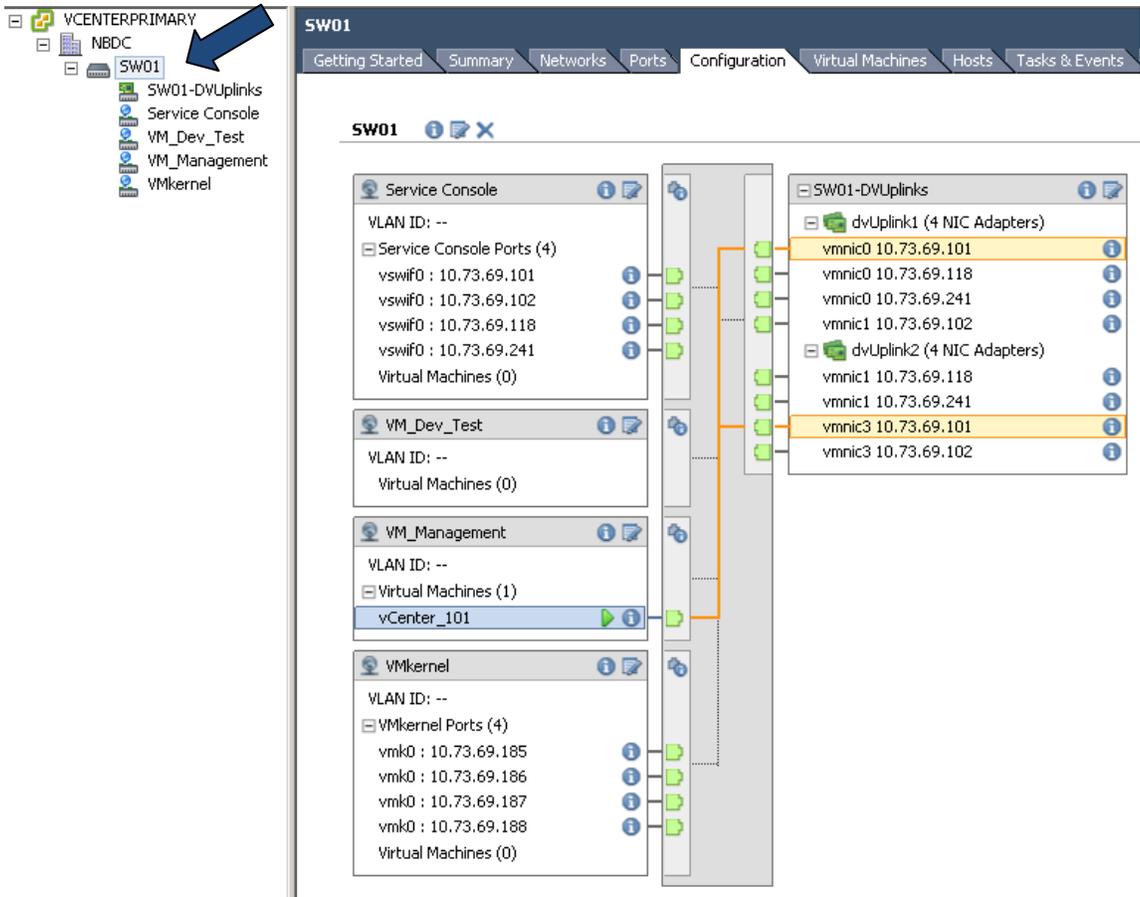


Figure 8) ESX Server and vCenter Server setup and configuration.

4.3.2 STEP 2: NETAPP STORAGE SYSTEM SETUP

Install and configure a NetApp active-active storage pair for the dev/test environment.

For the sample dev/test setup used in this document, a NetApp active-active pair—btcppe-filer35 and btcppe-filer36—is used as shown in the diagram below.

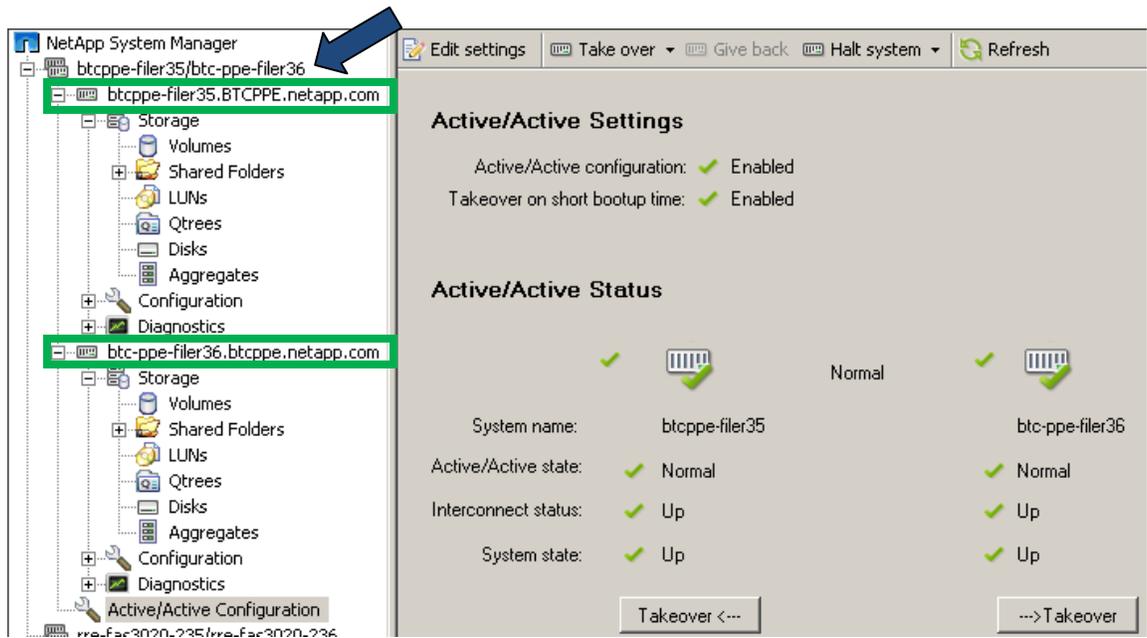


Figure 9) NetApp storage system configuration.

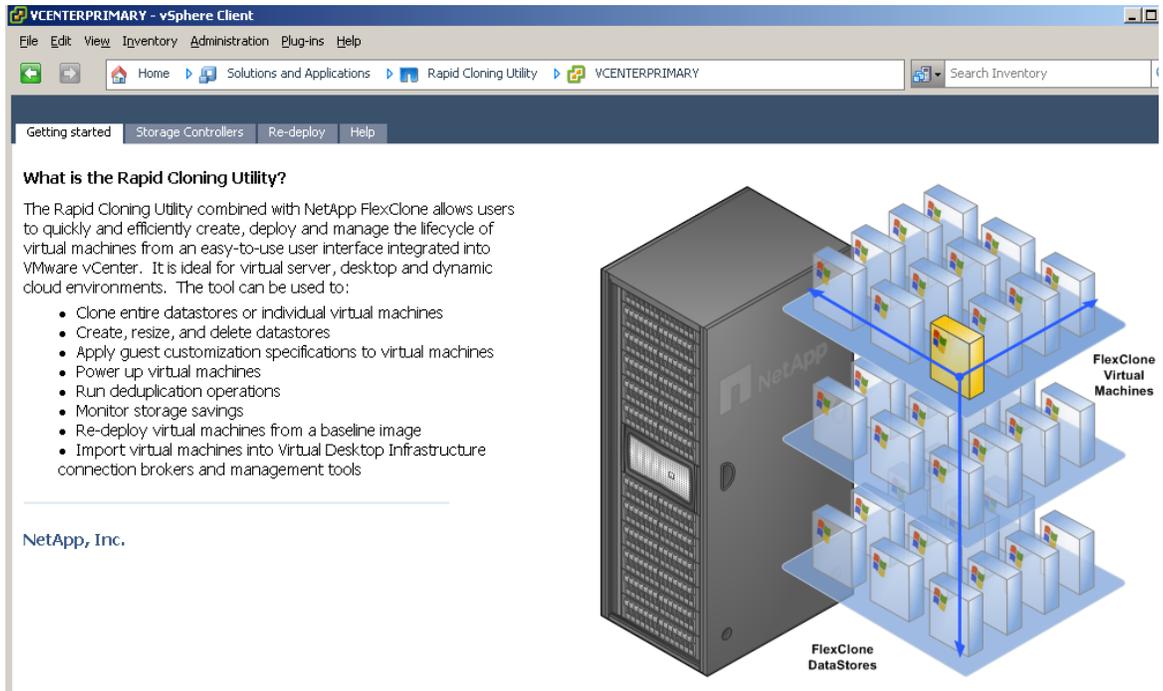
4.3.3 STEP 3: NETAPP RAPID CLONING UTILITY SETUP

Install the RCU 3.0 plug-in in vCenter Server. Configure it using the credentials of the NetApp storage systems setup in step 2. For details refer to the “NetApp Rapid Cloning Utility Installation and Administration Guide” in <http://now.netapp.com>:

http://now.netapp.com/NOW/knowledge/docs/rapid_cloning_utility/reirapidclone30/pdfs/install.pdf

Best Practice: Follow the best practice steps recommended by TR- 3749 for using the NetApp RCU plug-in in VMware vCenter Server.

1. From VMware vCenter Server, click View → Solutions and Applications → Rapid Cloning Utility.



What is the Rapid Cloning Utility?

The Rapid Cloning Utility combined with NetApp FlexClone allows users to quickly and efficiently create, deploy and manage the lifecycle of virtual machines from an easy-to-use user interface integrated into VMware vCenter. It is ideal for virtual server, desktop and dynamic cloud environments. The tool can be used to:

- Clone entire datastores or individual virtual machines
- Create, resize, and delete datastores
- Apply guest customization specifications to virtual machines
- Power up virtual machines
- Run deduplication operations
- Monitor storage savings
- Re-deploy virtual machines from a baseline image
- Import virtual machines into Virtual Desktop Infrastructure connection brokers and management tools

NetApp, Inc.

NetApp

FlexClone DataStores

FlexClone Virtual Machines

Figure 10) Rapid cloning utility setup.

2. Add the storage controllers.

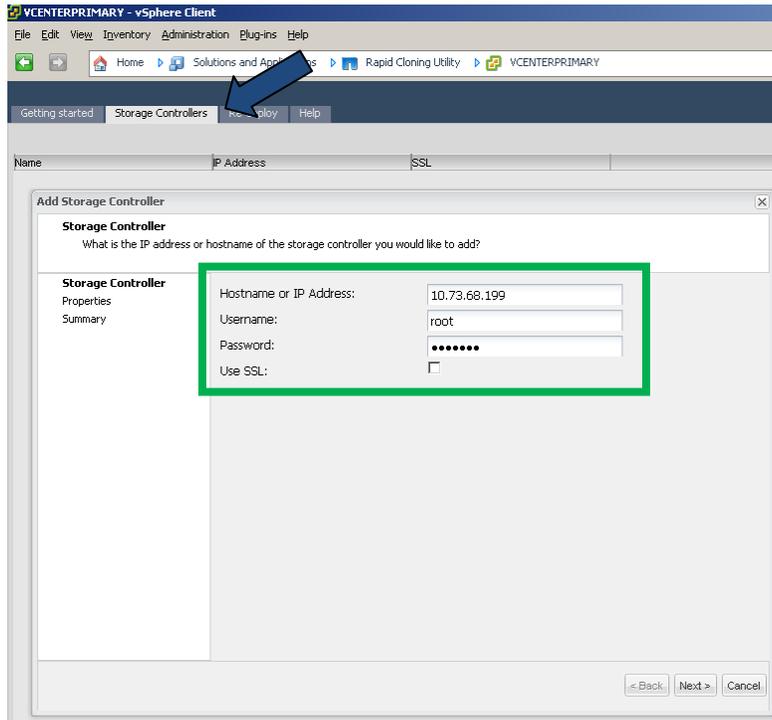


Figure 11) Rapid cloning utility setup.

3. Choose the aggregate and the network interface that will be used for the dev/test setup.

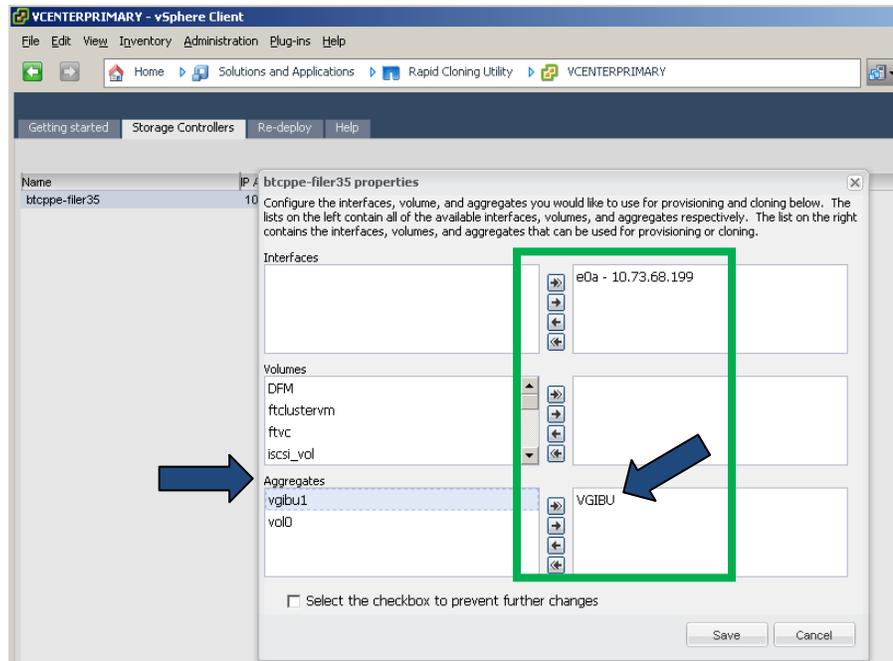


Figure 12) Rapid cloning utility setup.

4.3.4 STEP 4: NETAPP VIRTUAL STORAGE CONSOLE FOR VMWARE VSPHERE SETUP

Install the NetApp VSC 1.0 plug-in in VMware vCenter Server. For details, refer to the [Virtual Storage Console for VMware vSphere Installation and Administration Guide](#).

The VSC plug-in will automatically detect the NetApp storage controllers mapped to the VMware ESX cluster and will provide storage monitoring capabilities.

Best Practice: Follow the best practice configurations recommended by [TR- 3749](#) for using the NetApp Virtual Storage Console plug-in in VMware vCenter Server.

4.3.5 STEP 5: NFS DATASTORE CREATION

1. Create two NFS datastores from the NetApp system using RCU.
 - **SWAP_DataStore:** For storing the vSwap space of all virtual machines to be created
 - **Golden_DataStore:** For storing the golden virtual machine to be created from where all the clone virtual machines will be created

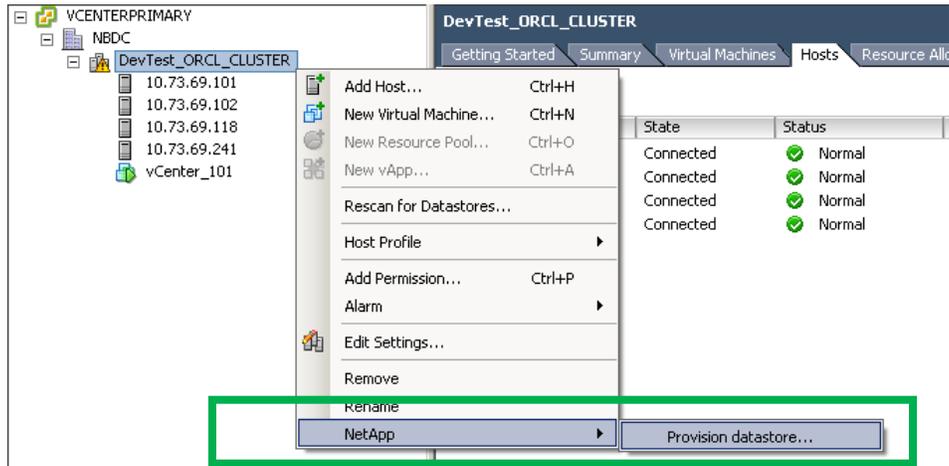


Figure 13) NFS datastore creation.

2. Select the storage controller.



Figure 14) NFS datastore creation.

3. Select the datastore type.

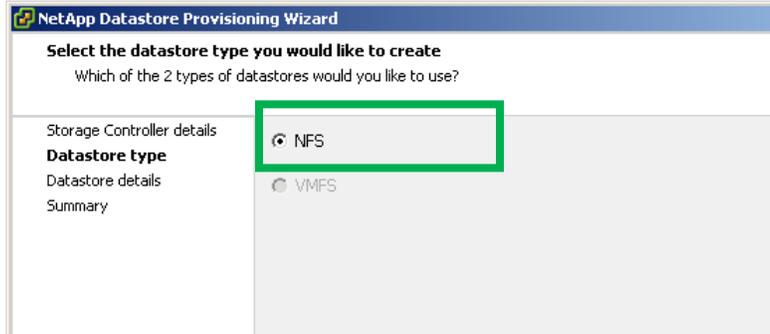


Figure 15) NFS datastore creation.

4. Enter the specifications for the NFS datastores to be created.
- i. Size of the datastore
 - ii. Name of the datastore
 - iii. Aggregate on which the datastore needs to be created
 - iv. Options for enabling the 'Thin Provisioning' and 'Auto Grow' option

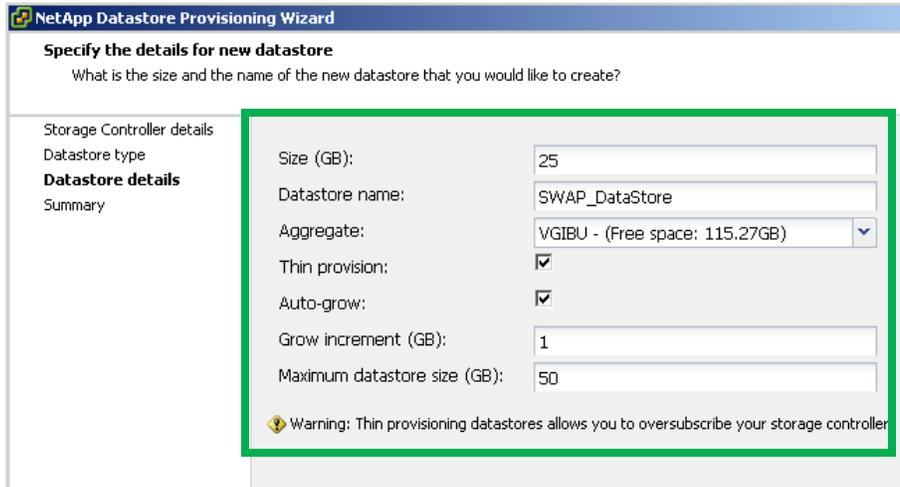


Figure 16) NFS datastore creation.

5. Review the information and submit (complete the wizard).

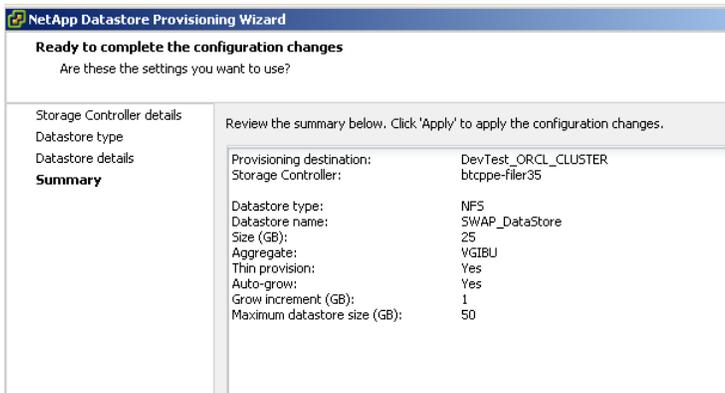


Figure 17) NFS datastore creation.

RCU will automatically create the volume corresponding to the datastore and mount it on all the servers of the cluster.

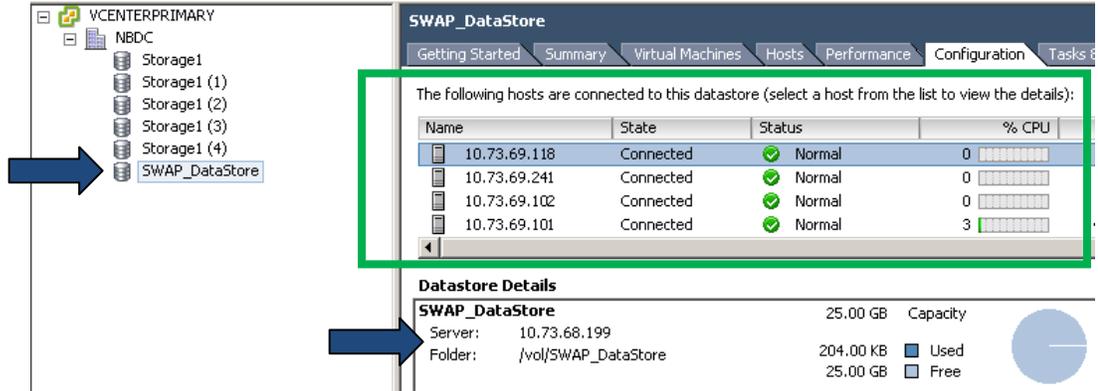


Figure 18) NFS datastore creation.

6. Similarly create another datastore for storing the golden virtual machine.

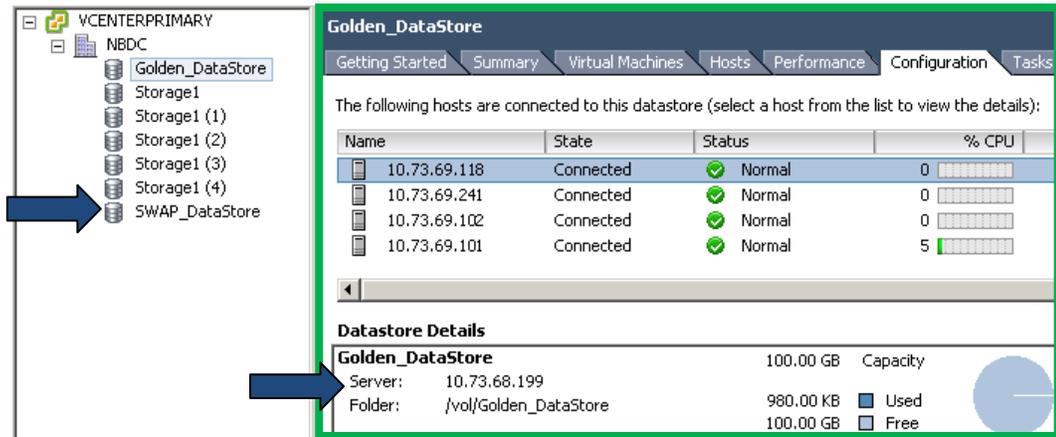


Figure 19) NFS datastore creation.

Now the storage map looks like this:

Custom Map

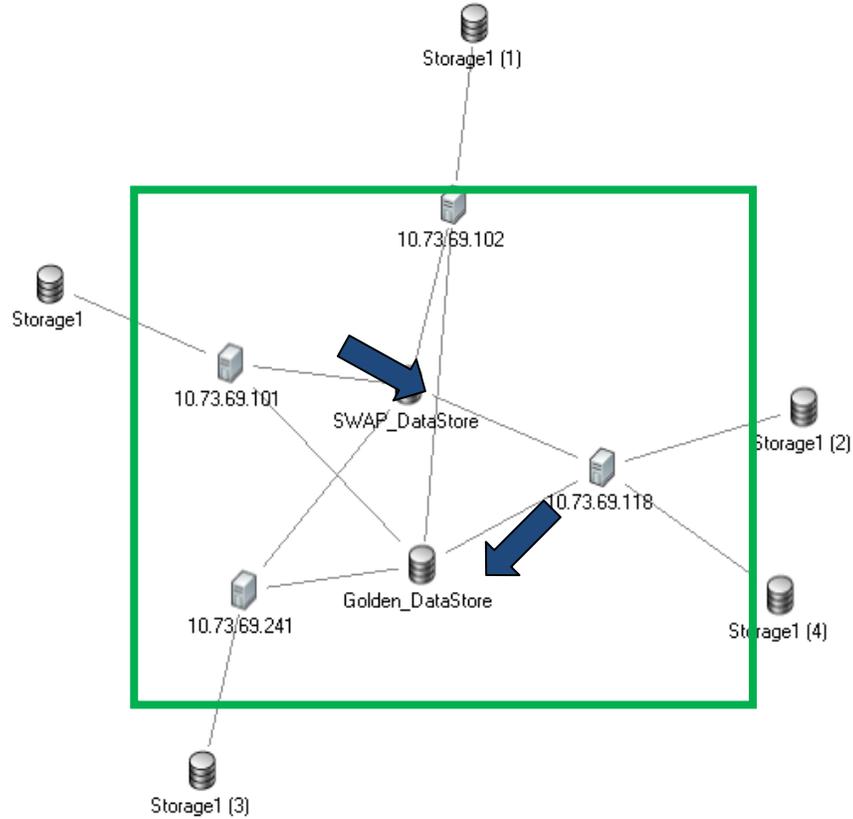


Figure 20) Storage map.

The details regarding the NFS datastore created also can be seen from the VSC console. RCU also makes sure that all the NFS tunables corresponding to NetApp best practices are set. Otherwise these can be set from the VSC console itself.

Controller: btcppe-filer35 (2 Datastores)	Datastore Capacity	NFS pathname	Access Mode	Storage Status
Golden_DataStore	100.00GB	/vol/Golden_DataStore	readWrite	Normal
SWAP_DataStore	25.00GB	/vol/SWAP_DataStore	readWrite	Normal

NFS	Capacity
VMFS pathname: /mnt/volumes/d5738d21-8bef778d NFS pathname: /vol/Golden_DataStore Status: normal File System Security: unix Anonymous Username: N/A	Datastore Usage (0%) Volume Usage (0%) Aggregate Usage (76%)
Host Privileges View Read-Only Hosts... View Read-Write Hosts... View Root Access Hosts...	Volume Name: Golden_DataStore Status: online Type: flex Guarantee: none Aggregate: VGIBU Snapshot Reserve: 0% Autogrow Increment: 1.00GB Autogrow Max. Size: 30.00GB

Figure 21) NFS datastore details.

4.3.6 STEP 6: VSWAP LOCATION SETUP

1. Set up the location of the vSwap file of the virtual machines and set it to SWAP_DataStore.
2. From vCenter Server, right-click the cluster name (DevTest_ORCL_CLUSTER) and select Edit Settings...

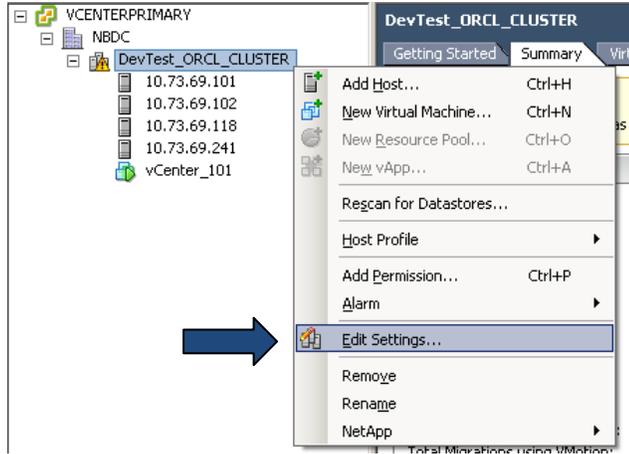


Figure 22) vSwap location setup.

3. Choose Swapfile Location and select the Store the swapfile in the datastore specified by the host button.

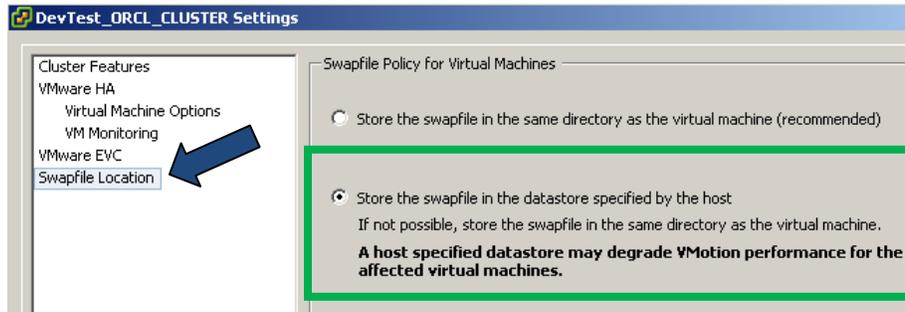


Figure 23) vSwap location setup.

- Now for each ESX server in the cluster, click the Configuration tab and select SWAP_DataStore corresponding to the Virtual Machine Swapfile Location option.

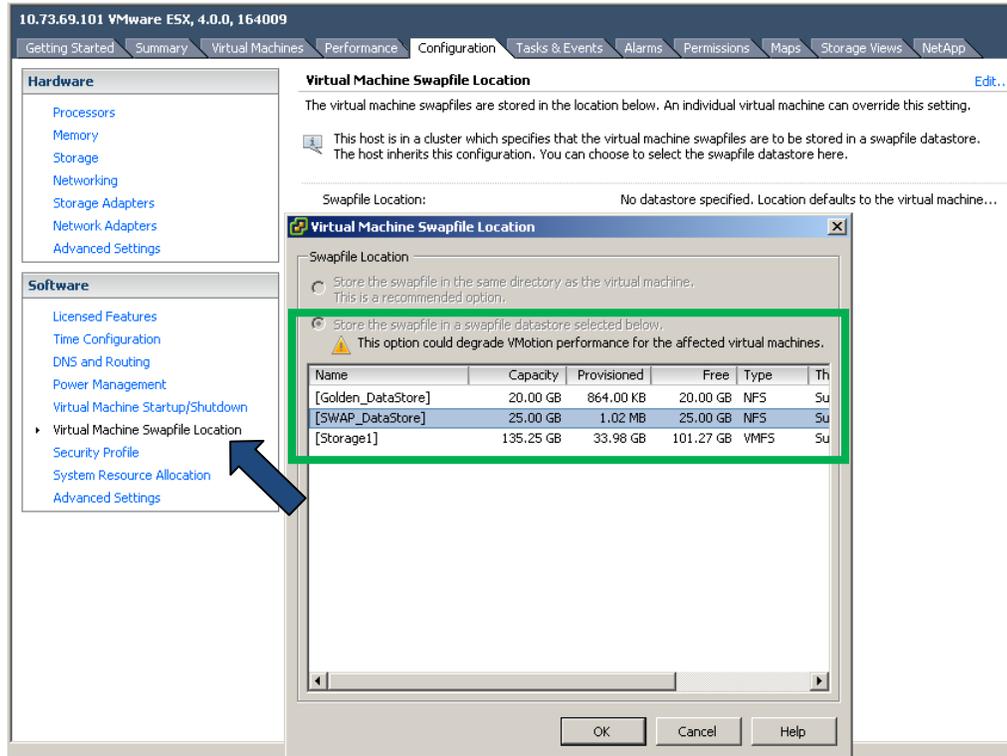


Figure 24) vSwap location setup.

4.3.7 STEP 7: GOLDEN VIRTUAL MACHINE CREATION

Various approaches can exist for creating the golden virtual machine from which the dev/test virtual machines are cloned. Two common ones are:

1. Creating a golden virtual machine from scratch by creating a new virtual machine on Golden_DataStore and then installing the required operating system and the Oracle binary and database (fresh database server instance).
2. Doing a P2V conversion from a production Oracle Database server.

Both these processes are illustrated in the following sections.

1. Creating the golden virtual machine from scratch

- a. Create Golden_VM on Golden_DataStore. Assign the required processor, memory, network resources ('VM_Dev_Test' port group – figure # in section 3.3.1), and a hard disk (vmdk-1) of size 20GB from Golden_DataStore.
- b. Boot the virtual machine with an ISO image of the operating system (RHEL 5U4 64-bit) assigned to its CD drive and install the operating system on the first hard disk (vmdk-1).
- c. Create and assign four additional hard disks (that is, vmdk files) to Golden_VM from Golden_DataStore.

#	vmdk	Size	Purpose	Corresponding Path Inside the Virtual Machine
1	vmdk-2	15GB	Oracle binary	/dev/sdb1
2	vmdk-3	20GB	Oracle data	/dev/sdc1
3	vmdk-4	20GB	Redo logs	/dev/sdd1
4	vmdk-5	20GB	Archive	/dev/sde1

Align all the VMDKs corresponding to the virtual machine using `fdisk` utility. Refer to [TR-3747](#) for details.

- d. Once the virtual machine boots up, discover the other hard disks corresponding to the vmdk-2 to vmdk-5.

```
cat /proc/partitions
```

- e. Format the hard disks using `fdisk` and create ext3 file systems on them using `mkfs`.

```
fdisk /dev/hdd  
mkfs -t ext3 /dev/hdd1
```

Mount the hard disks and install the Oracle Database 11g R2 single instance.

2. Creating the golden virtual machine from a production Oracle Database server using P2V conversion

Use any VMware P2V conversion [method](#) for creating the virtual machine on Golden_DataStore.

Once Golden_VM has been created, align the vmdk files of the virtual machines using either `fdisk` or `mbrscan` and `mbralign` utilities provided by NetApp. They can be downloaded from the VSC console. For details about the virtual machine alignment process refer to [TR-3747](#).

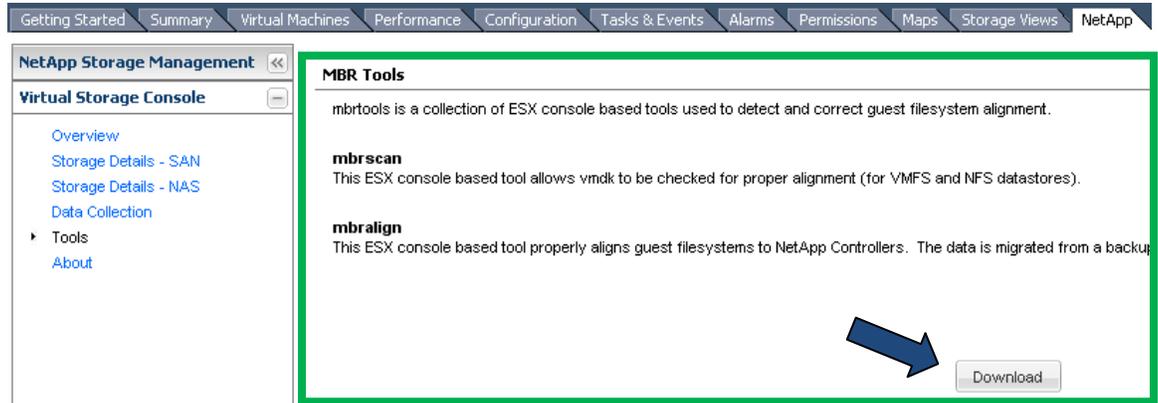


Figure 25) MBR download.

The following figure shows Golden_VM created with the RHEL5U4 operating system and Oracle Database 11g R2 installed.

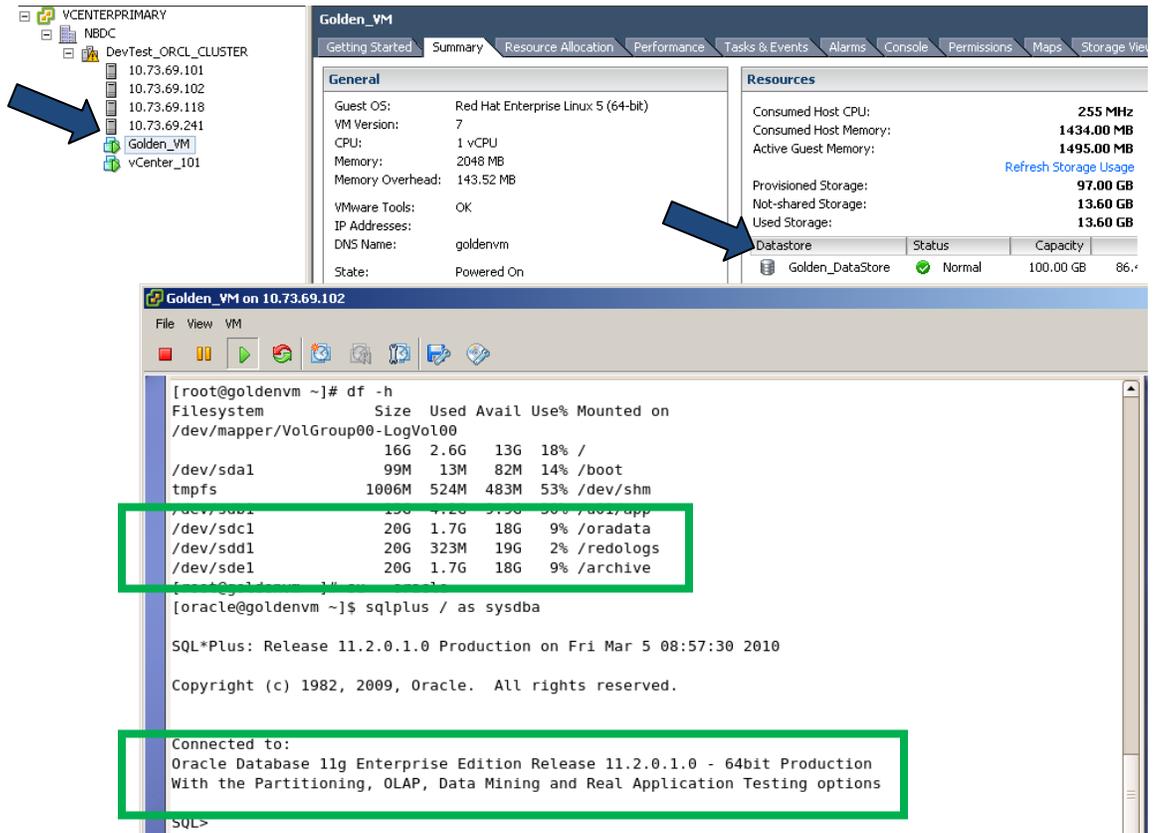


Figure 26) Golden VM details.

4.3.8 STEP 8: CUSTOM SPECIFICATION SCRIPT

To make sure that the clone dev/test virtual machines come up with a unique host name when they boot up, a custom specification script needs to be created from vCenter Server.

1. From VMware vCenter Server, choose: Home → Customization Specifications Manager → New.

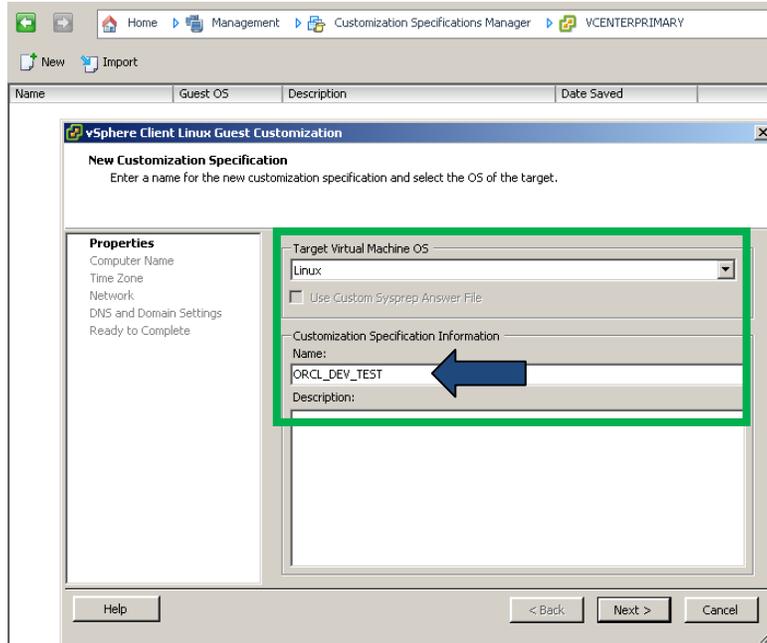


Figure 27) vSphere client guest customization.

2. Select the option to set the same virtual machine host name as that of the virtual machine name.

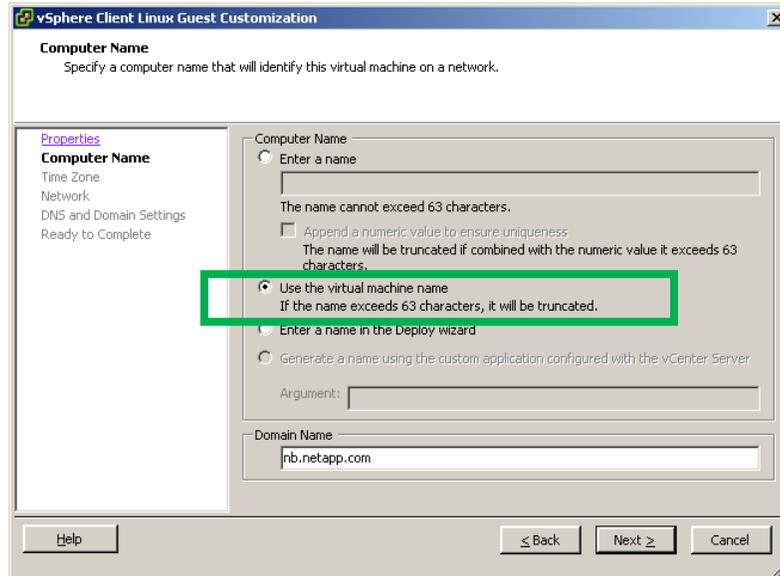


Figure 28) vSphere client guest customization.

3. Select the time zone.

4. Select a customer setting for network.

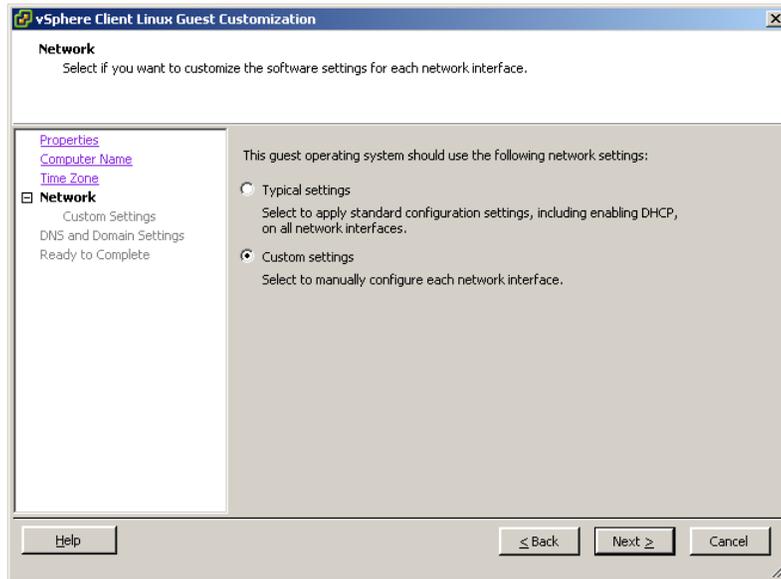


Figure 29) vSphere client guest customization.

5. Select the option to use DHCP for the IP address.

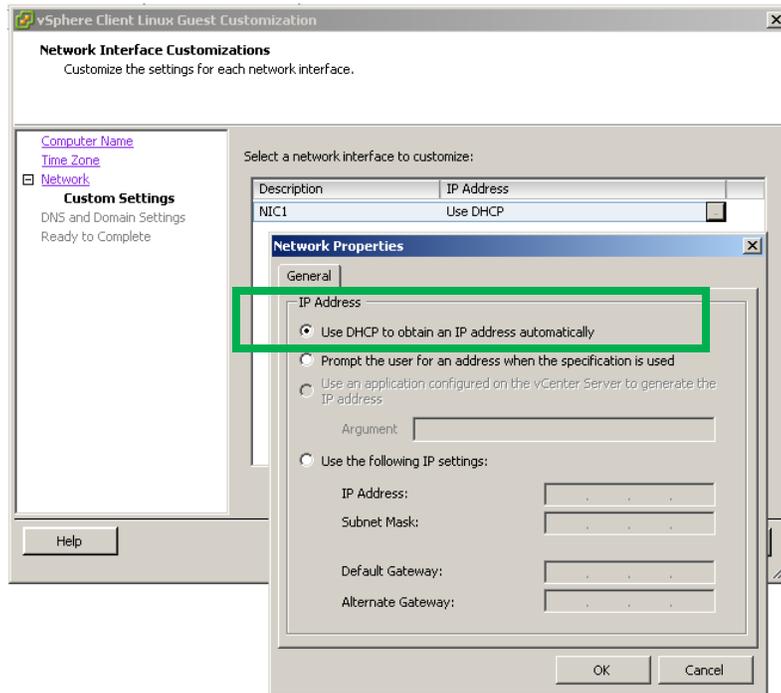


Figure 30) Network interface customization.

6. Enter the DNS server and domain information.

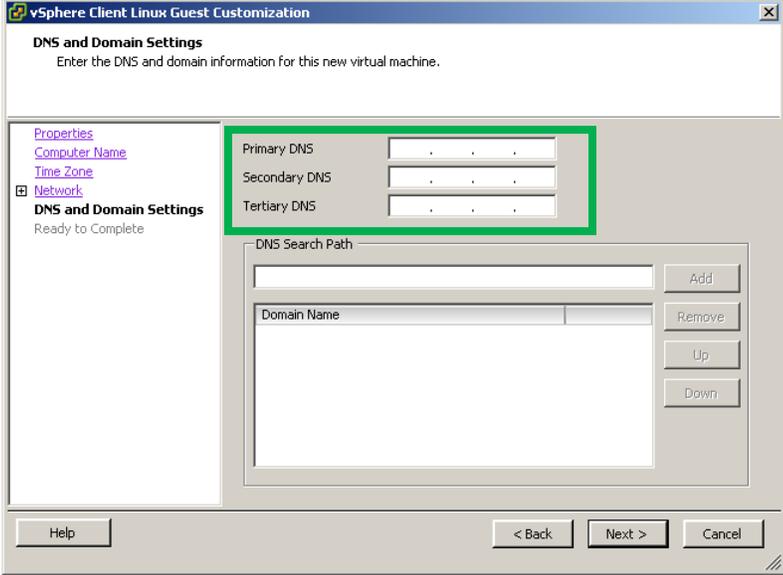


Figure 31) vSphere client guest customization.

7. Finish the wizard and create the customization setting.

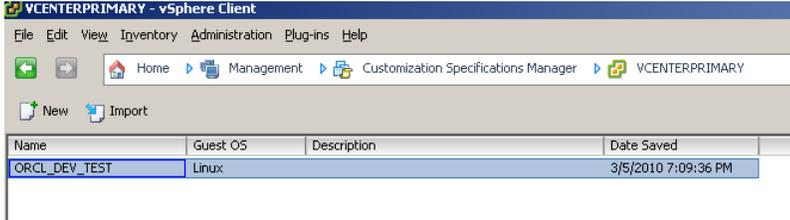


Figure 32) vSphere client guest customization.

4.3.9 STEP 9: CREATION OF CLONE VIRTUAL MACHINES

1. Right-click Golden_VM and click NetApp → Create rapid clones...

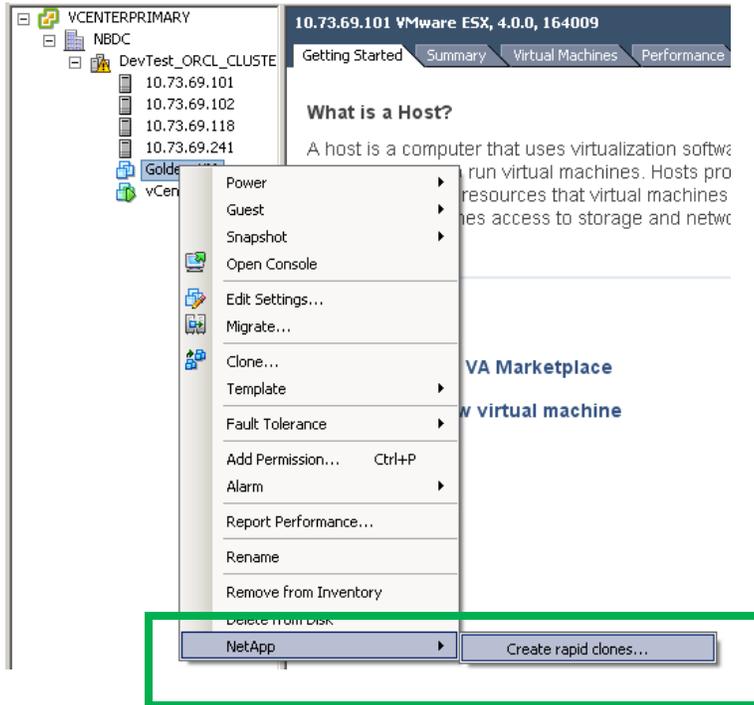


Figure 33) Create rapid clones.

2. Choose the storage controller to use.

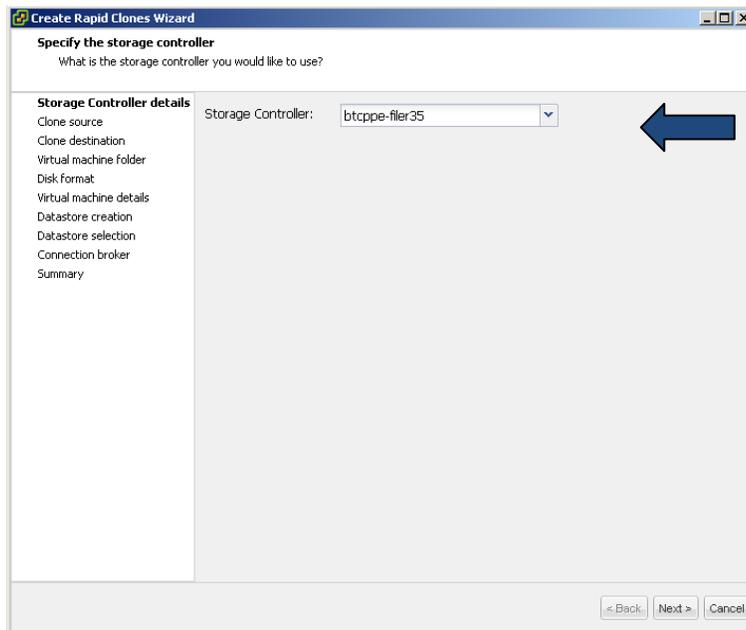


Figure 34) Create rapid clones wizard.

3. Choose the ESX server cluster as the destination in which the clone virtual machines will be created. The cloned virtual machines will be evenly distributed across the ESX servers constituting the cluster.

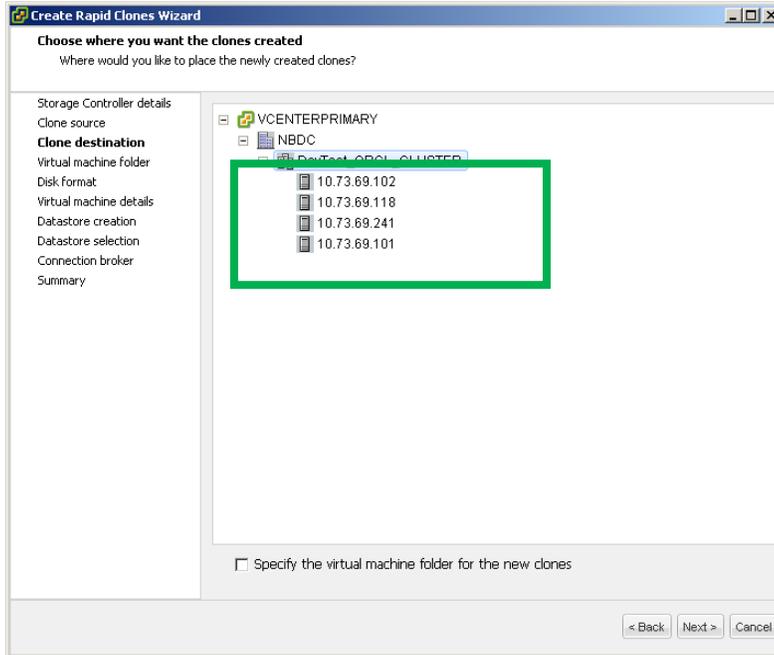


Figure 35) Create rapid clones wizard.

4. Choose the disk format to be used for the clone virtual machines. In this example, the option Same format as source is selected.

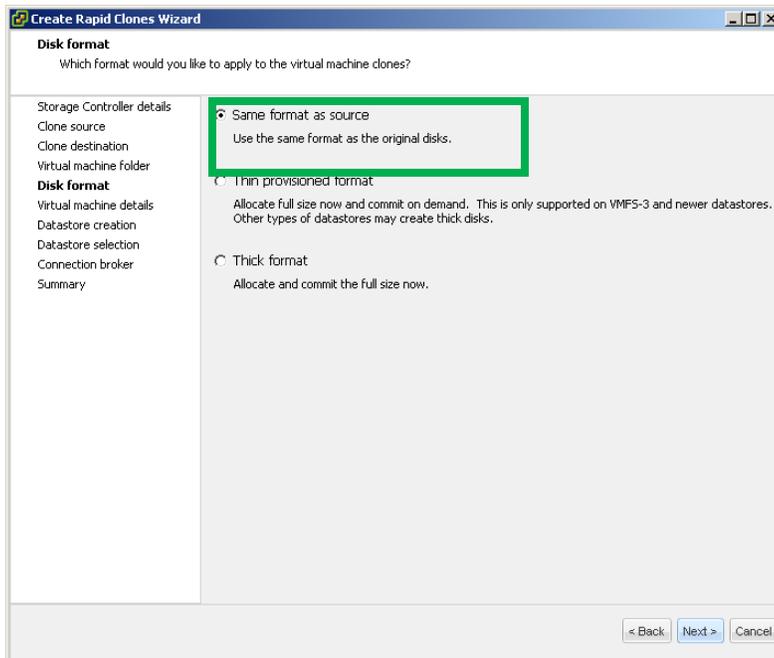


Figure 36) Create rapid clones wizard.

5. Choose the details of the virtual machines to be cloned:
 - Number of clone virtual machines
 - Name of clone virtual machines
 - Custom specification to be applied
 - Whether the clone virtual machines need to be created in a new datastore

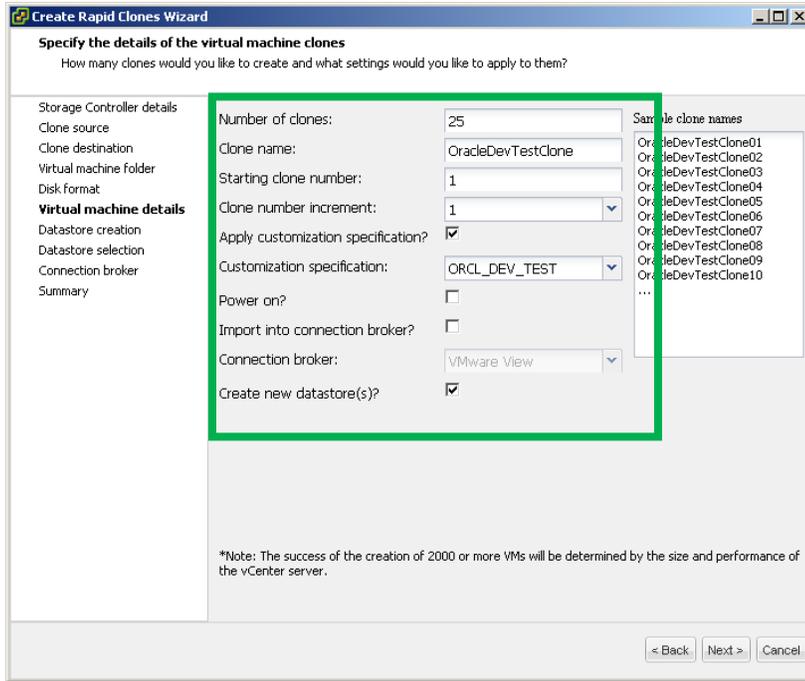


Figure 37) Create rapid clones wizard.

6. Create the new NFS datastore to store the clone virtual machines.

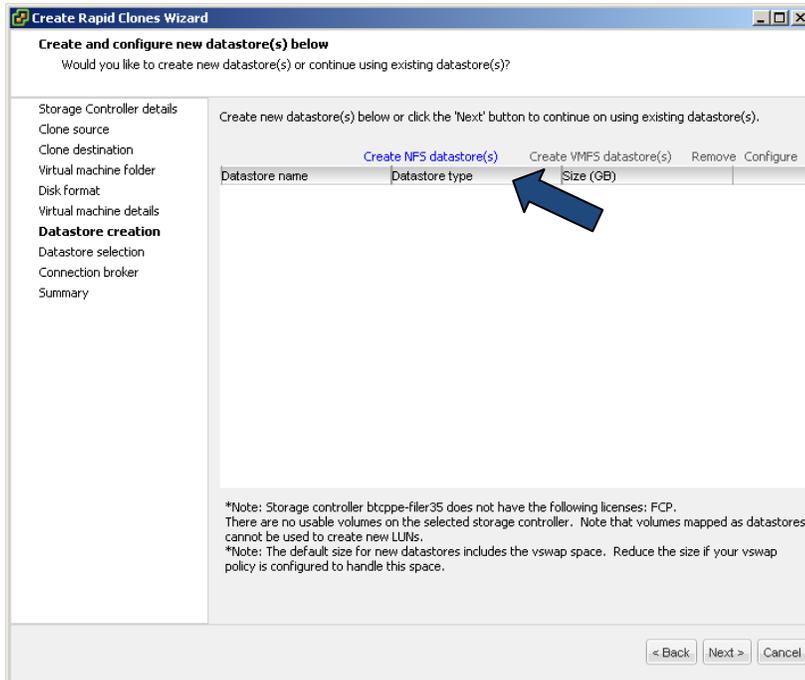


Figure 38) Create rapid clones wizard.

- Enter the name of the new datastore to be created (DevTest_DataStore). The size of the new datastore is automatically calculated by the RCU depending on the Golden_VM size and the number of clone virtual machines to be created.

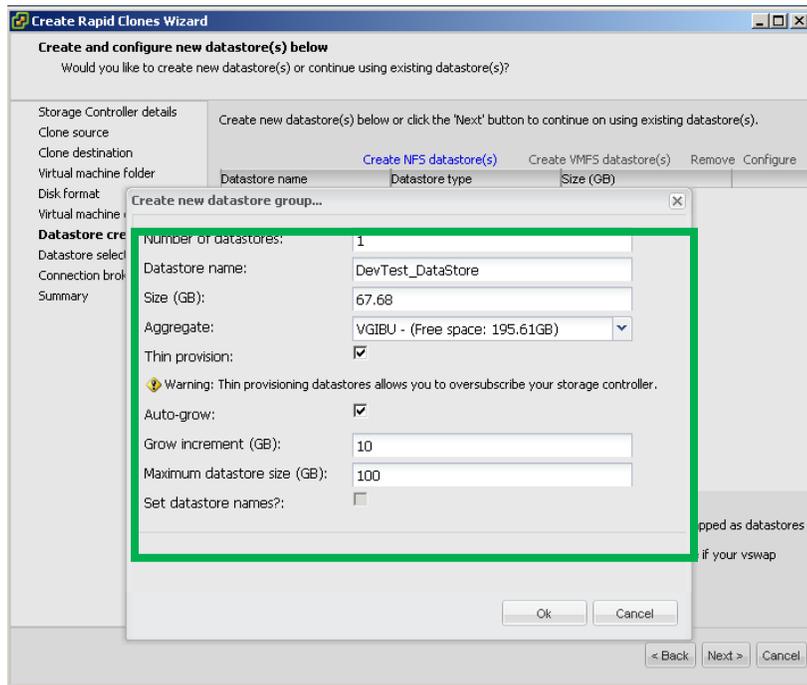


Figure 39) Create rapid clones wizard.

- Choose the datastore on which the clone virtual machines will be stored. For this example, choose DevTest_DataStore, the datastore you just created.

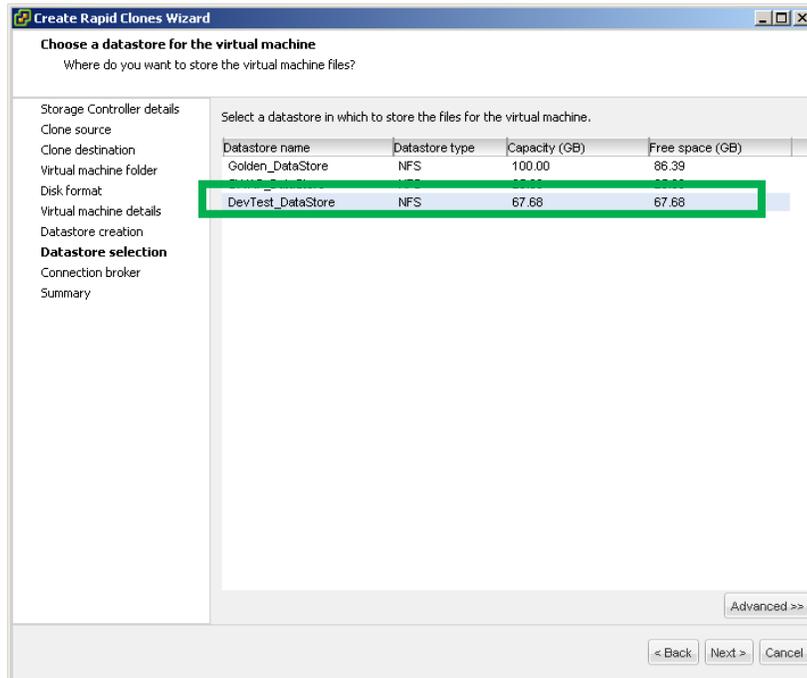


Figure 40) Create rapid clones wizard.

9. Complete the wizard to start the RCU clone creation process.

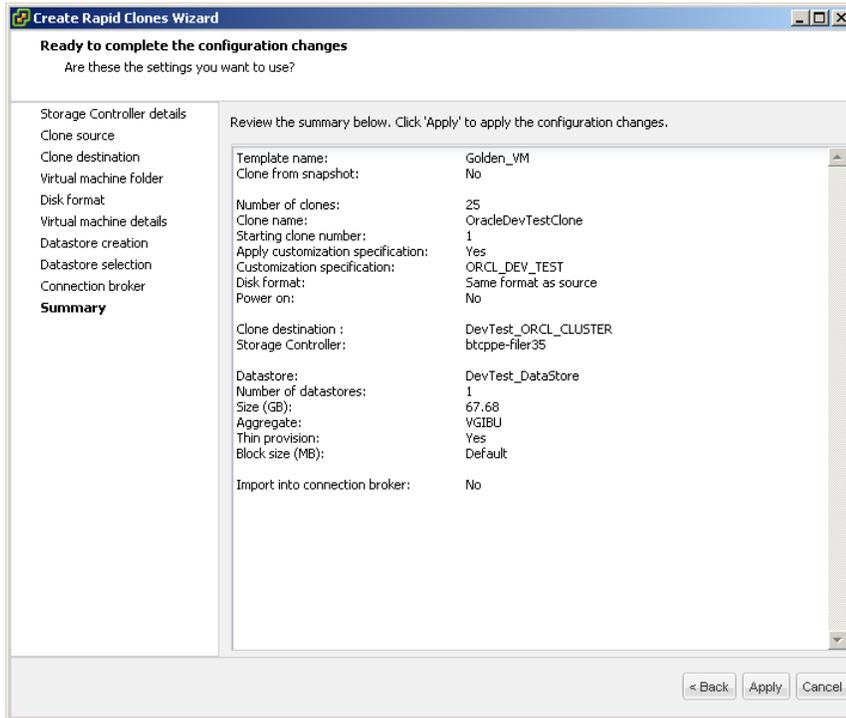


Figure 41) Create rapid clones wizard.

The progress of the cloning operation can be viewed from vCenter Recent Tasks.

Recent Tasks		
Name	Target	Status
Reconfigure virtual machine	OracleDevTestClone17	Completed
Reconfigure virtual machine	OracleDevTestClone16	Completed
Reconfigure virtual machine	OracleDevTestClone15	Completed
Customize virtual machine guest OS	OracleDevTestClone04	Completed
Clone virtual machine	OracleDevTestClone02	Completed
Clone virtual machine	OracleDevTestClone02	Completed
Clone virtual machine	OracleDevTestClone02	Completed
Customize virtual machine guest OS	OracleDevTestClone07	Completed
Reconfigure virtual machine	OracleDevTestClone19	Completed
Reconfigure virtual machine	OracleDevTestClone10	Completed
Reconfigure virtual machine	OracleDevTestClone14	Completed
Customize virtual machine guest OS	OracleDevTestClone06	Completed
Clone virtual machine	OracleDevTestClone02	Completed
Clone virtual machine	OracleDevTestClone02	Completed
Customize virtual machine guest OS	OracleDevTestClone09	Completed

Figure 42) Recent tasks.

RCU will automatically create the clone virtual machines and allocate them across the ESX servers in the cluster.

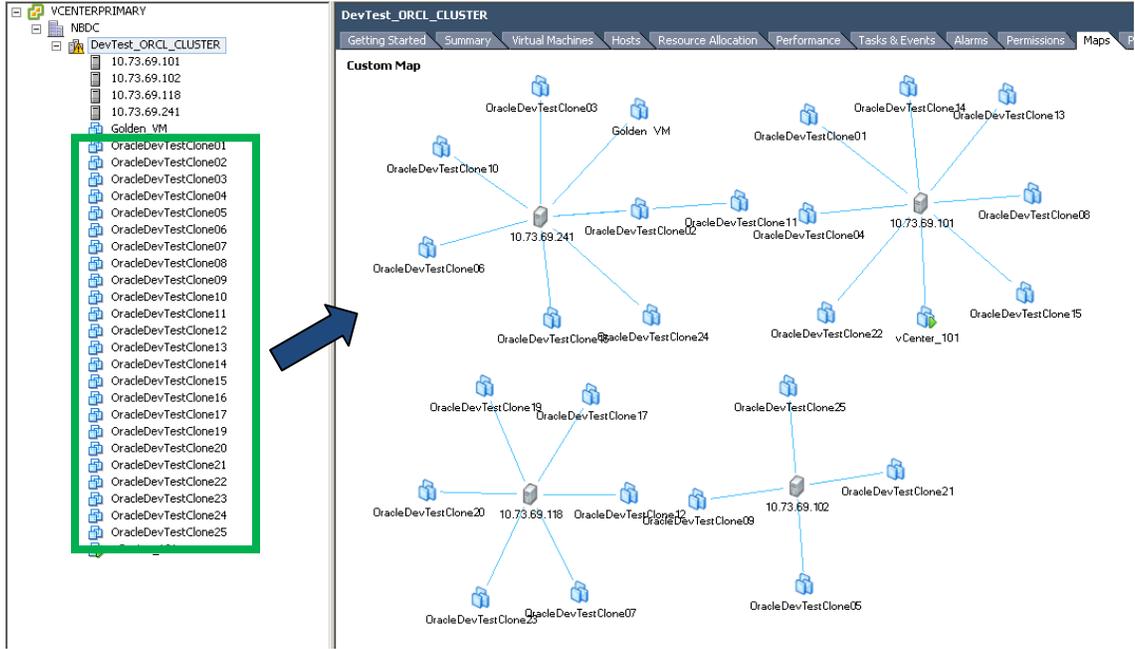


Figure 43) Custom map.

4.4 SOLUTION VALIDATION—SOLUTION VERSION 1

Table 5) Solution validation.

#	Steps	Comments
1	Verify application	When the clone virtual machines come up, verify that they all start running the Oracle Database.
2	Verify storage efficiency	Verify from the VSC console that no extra space in the NetApp storage has been consumed because of the cloned virtual machines and Databases

4.4.1 VERIFY APPLICATION

1. Boot up the clone virtual machines and verify that the Oracle Database inside the virtual machines comes up successfully.
2. The clone virtual machine's host name will be same as that of the virtual machine name as designed in the customization script.

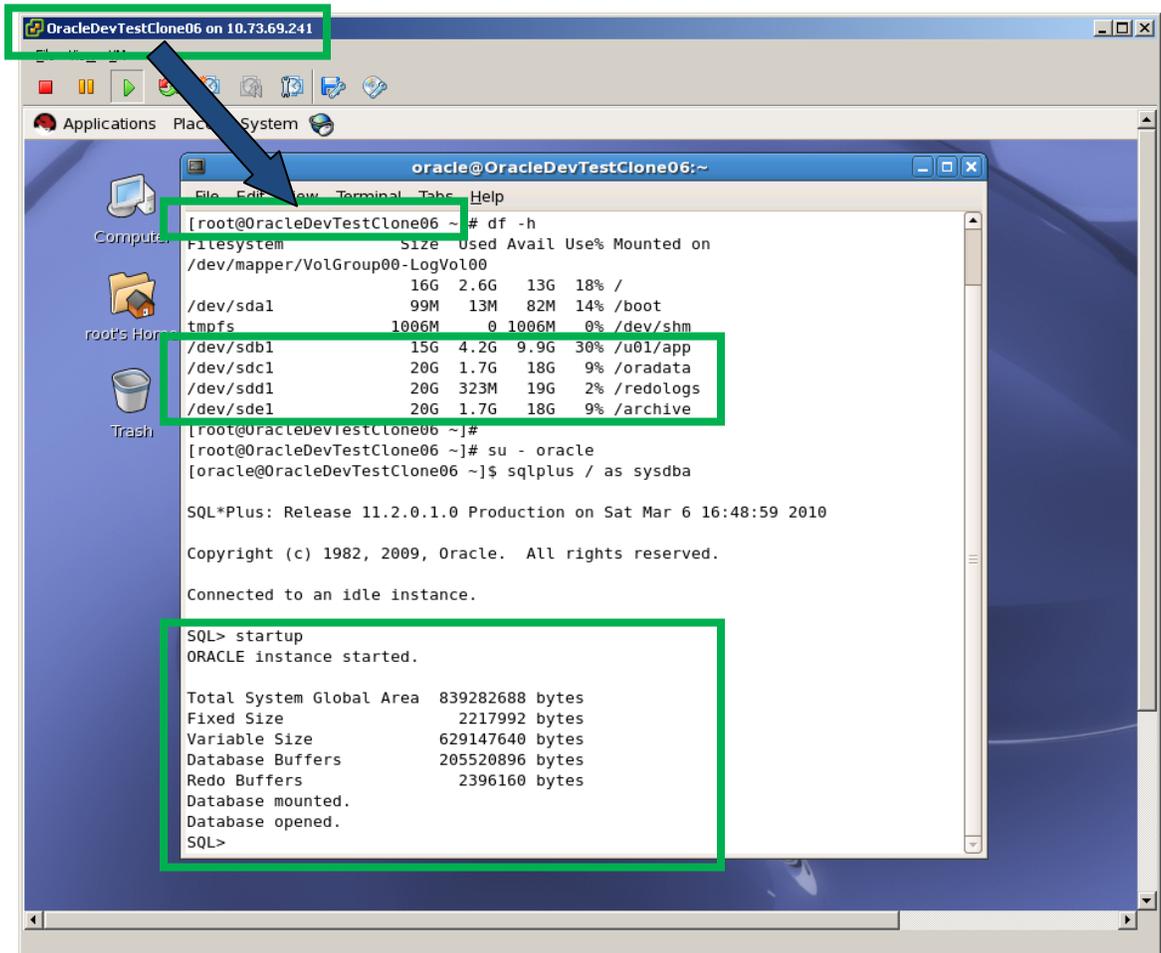


Figure 44) Verifying application.

4.4.2 VERIFY STORAGE EFFICIENCY

1. Check the storage usage of the datastore created for storing the clone virtual machines (DevTest_DataStore). Although the datastore now contains 25 clone virtual machines, it consumes space for only 1 virtual machine.

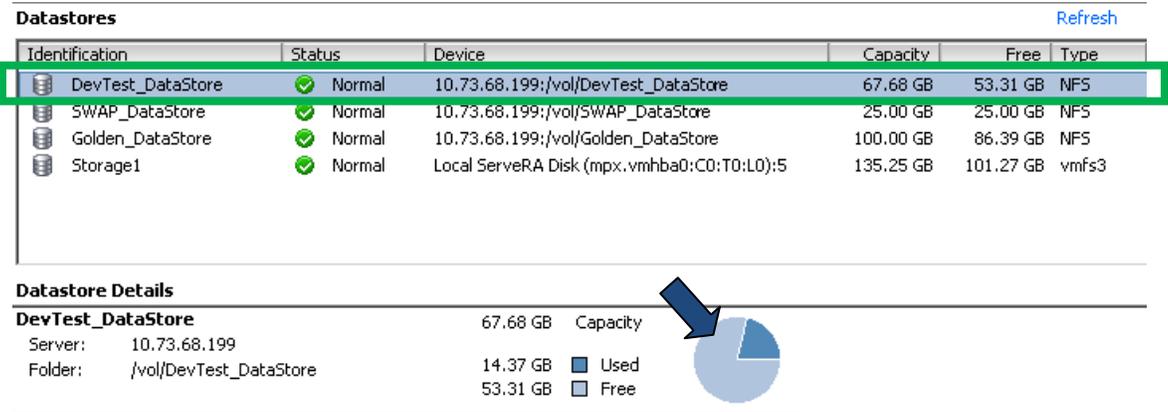


Figure 45) Verifying storage efficiency.

2. Also from the vCenter Server NetApp VSC console, the deduplication savings (96%) and the datastore usage can be tracked.

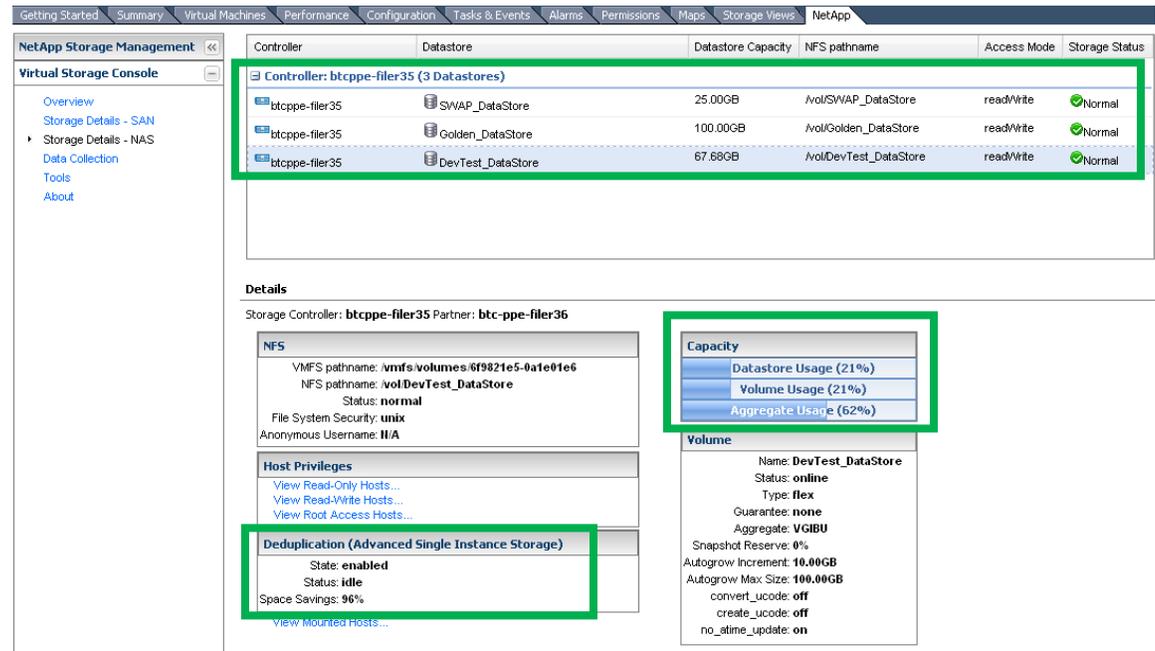


Figure 46) Verifying storage efficiency.

5 SOLUTION ARCHITECTURE—SOLUTION VERSION 2

5.1 ARCHITECTURE OVERVIEW

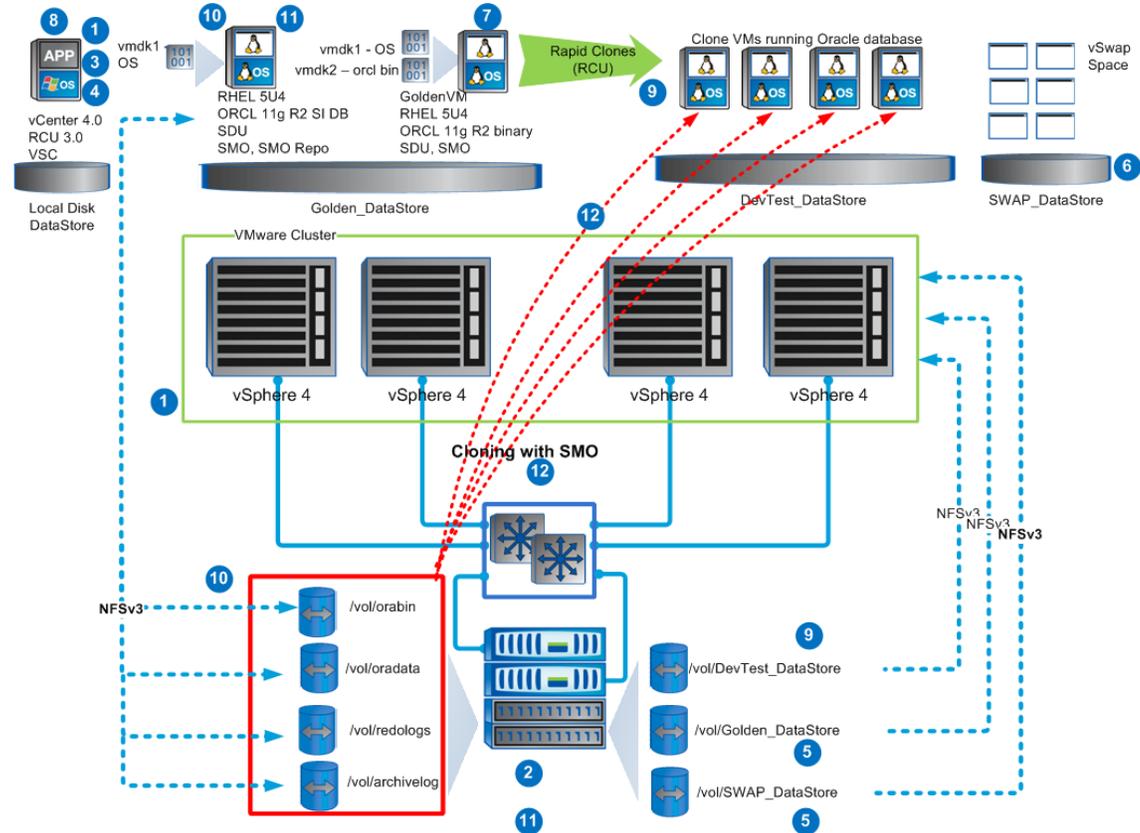


Figure 47) Storage architecture.

The diagram above illustrates the architecture this dev/test solution and the process followed to create the clone virtual machines. The numbers in the diagram corresponds to the solution implementation steps described in section 5.3.

5.2 HARDWARE AND SOFTWARE USED IN THE LAB SETUP

Table 6) Hardware and software used.

Infrastructure Component	Vendor	Quantity	Details
Server	IBM	1	IBM x3550 server
		3	IBM x3650 server
			Intel Xeon processor (Intel-VT), CPU: 38 GHz total in the cluster Memory: 24GB total in the cluster
Storage	NetApp		FAS3140: Active-Active config, RAID-DP
Network adapter	Broadcom	2/ server	Broadcom NetXtreme II BCM 5708 1000Base-T
HBA	N/A	N/A	N/A
Network switch	Cisco	2	
Software	NetApp		Data ONTAP 7.3.3 or higher
	NetApp		Cluster
	NetApp		NFS

	NetApp		FlexClone
	NetApp		ASIS
	NetApp		Rapid Cloning Utility 3.0 (VMware vCenter plug-in)
	NetApp		VSC (VMware vCenter plug-in)
	NetApp		SDU 4.1.1
	NetApp		SMO 3.0.3
	NetApp		Operations Manager 3.7.1
	VMware		VMware ESX v4
	VMware		VMware vCenter Server 4
	Oracle		Oracle Database 11g R2 Single Instance
	Red Hat		RHEL 5 U4

5.3 SOLUTION IMPLEMENTATION

Table 7 illustrates the steps involved in implementing this solution.

Table 7) Solution implementation.

#	Steps	Comments
1	VMware ESX and vCenter Server setup and configuration	Installation and configuration of the VMware vSphere 4 ESX servers. Installation and configuration of VMware vCenter Server 4. Configuration of VMware cluster and networking (vDS).
2	NetApp storage setup	Installation and configuration of the NetApp storage systems (NetApp FAS3140 active-active)
3	RCU setup	Installation and configuration of NetApp Rapid Cloning Utility 3.0 in VMware vCenter Server (vCenter plug-in)
4	VSC setup	Installation and configuration of NetApp VSC 1.0 (VMware Storage Console) in VMware vCenter Server (vCenter plug-in)
5	NFS datastore creation	Create 'Golden_DataStore' and 'SWAP_DataStore' using RCU from the same aggregate of the NetApp storage. Ensure correct NFS tunables on the Datastores created using VSC.
6	vSwap location setup	Make the changes in the ESX servers so that all virtual machine swap space (vSwap) is stored on 'SWAP_DataStore.'
7	Golden virtual machine creation	Create the golden virtual machine with the Oracle Database 11g R2 binary (single instance) on 'Golden_DataStore.' This virtual machine needs to have the SDU and SMO server installed. Align the vmdk files of the virtual machine using mbrscan and mbralign (or fdisk).
8	Custom specification script	Create a "custom specification" on vCenter Server for the clone virtual machines to be created so that they get a unique host name and the users are prompted to enter the IP address when they boot up.
9	Creation of clone virtual machines	Create 25 clones of the golden virtual machine using RCU on a new datastore – 'DevTest_DataStore.'
10	Set up the SMO repository environment for the production database to be cloned	The production database server (which can be either a physical or a virtual machine) needs to be set up for an SMO environment.
11	Provide role based access to the storage system hosting the production database from the virtual machine clones	Using the role-based access control of Operations Manager, set up the access to the storage system hosting the production database to be cloned from the cloned virtual machines.
12	Clone the production database and mount the database clones on the cloned virtual machines using SMO	Using SMO, clone the production database and mount the clones on the Cloned virtual machines

Note:

- Steps 1 to 9 in Table 7 are very similar to that of the first solution described in section 4 while steps 10, 11, and 12 are specific to this solution.
- In step 7 (golden virtual machine creation), only the Oracle Database binary is required and the SDU and SMO server needs to be installed in the golden virtual machine.

5.3.1 STEP 1: VMWARE ESX AND VCENTER SERVER SETUP AND CONFIGURATION

Refer to [section 4.3.1](#).

5.3.2 STEP 2: NETAPP STORAGE SYSTEM SETUP

Refer to [section 4.3.2](#).

5.3.3 STEP 3: NETAPP RAPID CLONING UTILITY SETUP

Refer to [section 4.3.3](#).

5.3.4 STEP 4: NETAPP VIRTUAL STORAGE CONSOLE FOR VMWARE VSPHERE SETUP

Refer to [section 4.3.4](#).

5.3.5 STEP 5: NFS DATASTORE CREATION

Refer to [section 4.3.5](#).

5.3.6 STEP 6: VSWAP LOCATION SETUP

Refer to [section 4.3.6](#).

5.3.7 STEP 7: GOLDEN VIRTUAL MACHINE CREATION

The golden virtual machine used in this solution doesn't contain the database to be cloned (unlike the solution described in section 4); rather, it is used for creating the cloned virtual machines on which SMO will mount the production database clones.

For this solution, the golden virtual machine needs to have the following software installed:

- Same version of the operating system and the Oracle Database binary as that of the production database server
- Same level of the SDU and SMO server as that of the production database server

The steps followed for creating the golden virtual machine used in this solution:

1. Create the 'Golden_VM_smo' virtual machine on 'Golden_DataStore.' Assign the required processor, memory, network resources ('VM_Dev_Test' port group – Figure 7 in [section 4.3.1](#)), and a hard disk (vmdk-1) of size 20GB from 'Golden_DataStore.'
2. Boot the virtual machine with an ISO image of the operating system (RHEL 5U4 64-bit) assigned to its CD drive and install the operating system on the first hard disk (vmdk-1). Note that this is the same version of the operating system as that of the production database server.
3. Install the same level of Oracle binary as that of the production Oracle Database server. For this example, the Oracle Database binary along with the sample database that comes with the installable will be installed on this golden virtual machine. Create four additional hard disks (VMDKs) in the same data store (Golden_DataStore) for the golden virtual machine for holding the database binary and the related files.

#	vmrk	Size	Purpose	Corresponding Path Inside the Virtual Machine
1	vmrk-2	15GB	Oracle binary	/dev/sdb1
2	vmrk-3	20GB	Sample Oracle Database	/dev/sdc1
3	vmrk-4	20GB	Redo logs	/dev/sdd1
4	vmrk-5	20GB	Archive	/dev/sde1

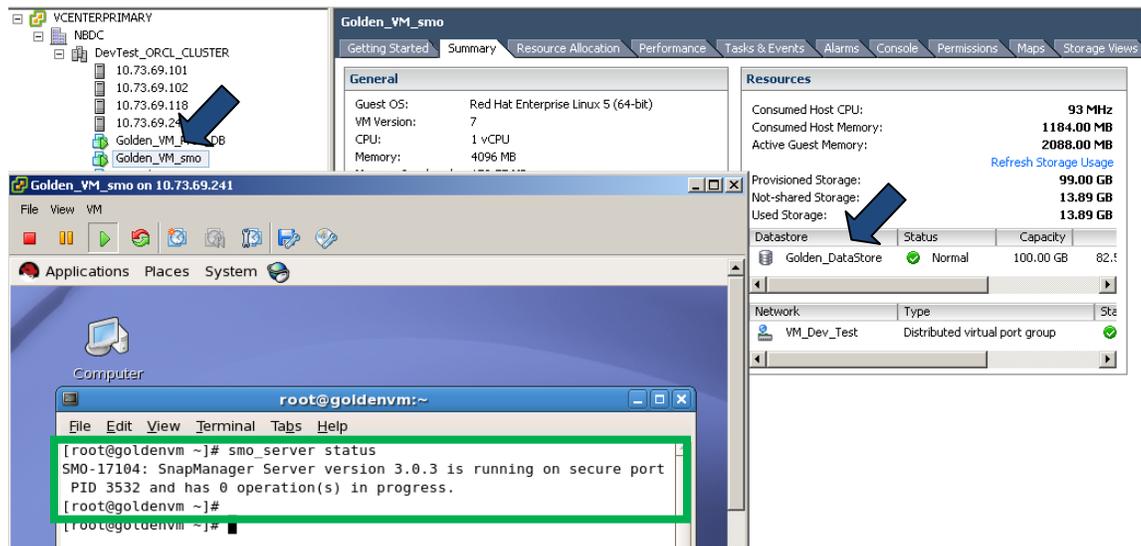
Note: Before installing the OS or the database binary, align all the VMDKs corresponding to the virtual machine using the 'fdisk' utility. Refer to [TR-3747](#) for details.

- Once the virtual machine boots up, discover the other hard disks corresponding to the vmrk-2 to vmrk-5.


```
cat /proc/partitions
```
- Format the hard disks using 'fdisk' and create 'ext3' file systems on them using 'mkfs.'


```
fdisk /dev/hdd
```

```
mkfs -t ext3 /dev/hdd1
```
- Mount the hard disks and install the Oracle Database 11g R2 single instance.
- Install the same level of SnapDrive for UNIX and SMO server as that of the production Oracle Database server on the golden virtual machine.



5.3.8 STEP 8: CUSTOM SPECIFICATION SCRIPT

Refer to [section 4.3.8](#).

5.3.9 STEP 9: CREATION OF CLONE VIRTUAL MACHINES

Refer to [section 4.3.9](#).

Following the same process as shown in [section 4.3.9](#), the required number of clone virtual machines are created from the golden virtual machine installed in step 7 using the Rapid Cloning Utility on a new NFS datastore (OrcIDevTestDataStore).

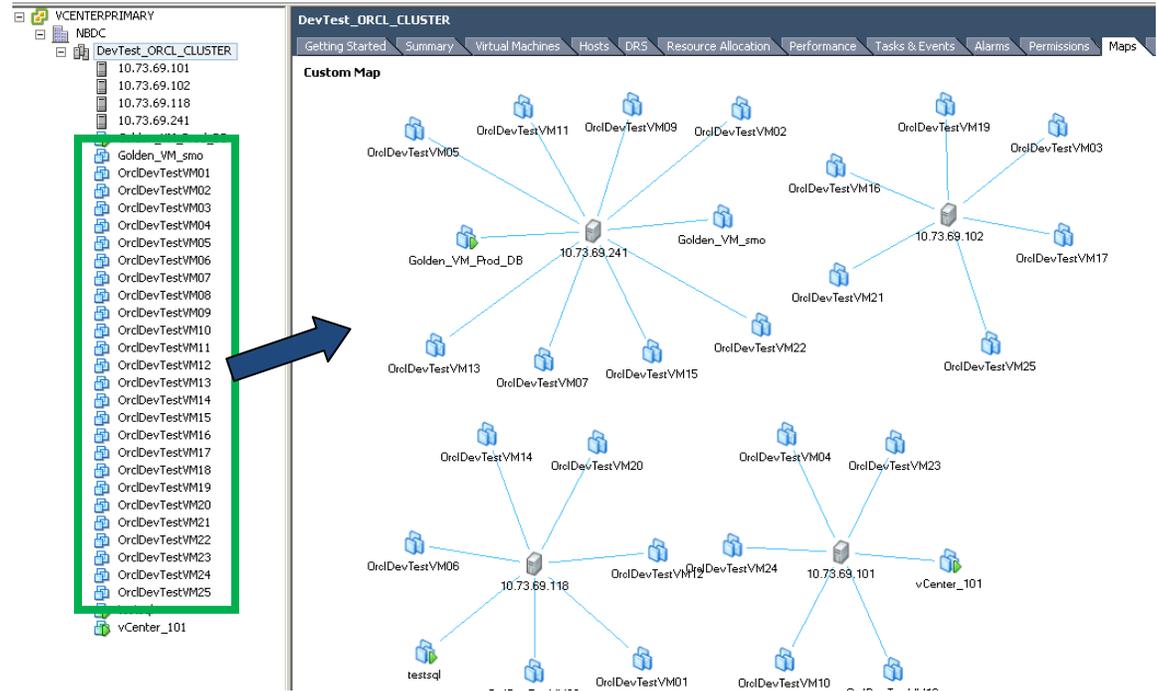


Figure 48) Custom map.

Boot one clone virtual machine to check if it comes up in the network as expected and if SMO services are running. As an example, this virtual machine will be used to mount the production database clones by SMO in step 12, below.

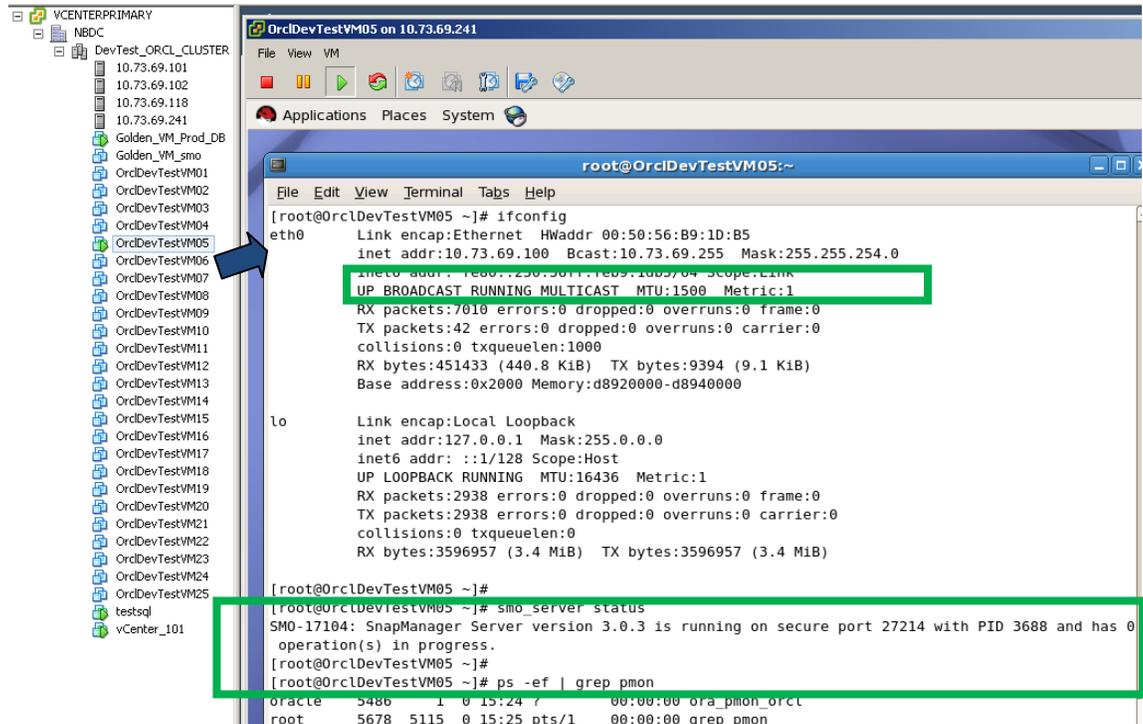


Figure 49) Verifying clone virtual machine.

5.3.10 STEP 10: SETUP OF SMO SERVER AND REPOSITORY FOR THE PRODUCTION DATABASE

1. The production Oracle Database that needs to be cloned can be running either in a physical database server host or in a virtual machine.
2. In this example, the production Oracle Database server uses NFS mount points from a NetApp storage system for storing the Oracle Database binary, data, and related files.

```
10.73.68.199:/vol/orahome on /u01/app type nfs (rw,bg,hard,rsi
wsize=32768,nfsvers=3,nointr,timeo=600,tcp,addr=10.73.68.199)
10.73.68.199:/vol/oradata on /oradata type nfs (rw,bg,hard,rsi
wsize=32768,nfsvers=3,nointr,timeo=600,tcp,addr=10.73.68.199)
10.73.68.199:/vol/redo on /redologs type nfs (rw,bg,hard,rsi
wsize=32768,nfsvers=3,nointr,timeo=600,tcp,addr=10.73.68.199)
10.73.68.199:/vol/archive on /archive type nfs (rw,bg,hard,rsi
wsize=32768,nfsvers=3,nointr,timeo=600,tcp,addr=10.73.68.199)
```

- In this example setup, the SMO repository has been configured in the production Oracle Database server itself. As shown in the screenshot below, the Oracle Database instance called 'orcldemo' is being used for the SMO repository; the other instance, called 'orclsmo,' is the production database instance to be cloned.

```
[root@productionDB ~]#
[root@productionDB ~]# tail --lines 2 /etc/oratab
orcldemo:/u01/app/oracle/product/11.1.0/db_1:Y
orclsmo:/u01/app/oracle/product/11.1.0/db_1:Y
[root@productionDB ~]#
[root@productionDB ~]# ps -ef | grep pmon
oracle      4846      1  0 05:30 ?        00:00:00 ora_pmon_orcldemo
oracle      5061      1  0 05:31 ?        00:00:00 ora_pmon_orclsmo
root        8676    6895  0 06:15 pts/2    00:00:00 grep pmon
[root@productionDB ~]#
[root@productionDB ~]#
```

Best Practice: For best practices and different options related to configuring the SMO server and the SMO repository for a production Oracle Database server, refer to the SMO Best Practices Guide: <http://media.netapp.com/documents/tr-3761.pdf>.

- Create an SMO profile (called 'BACKUP_ORCLSMO' in the example below) for backing up and cloning the 'orclsmo' database instance.

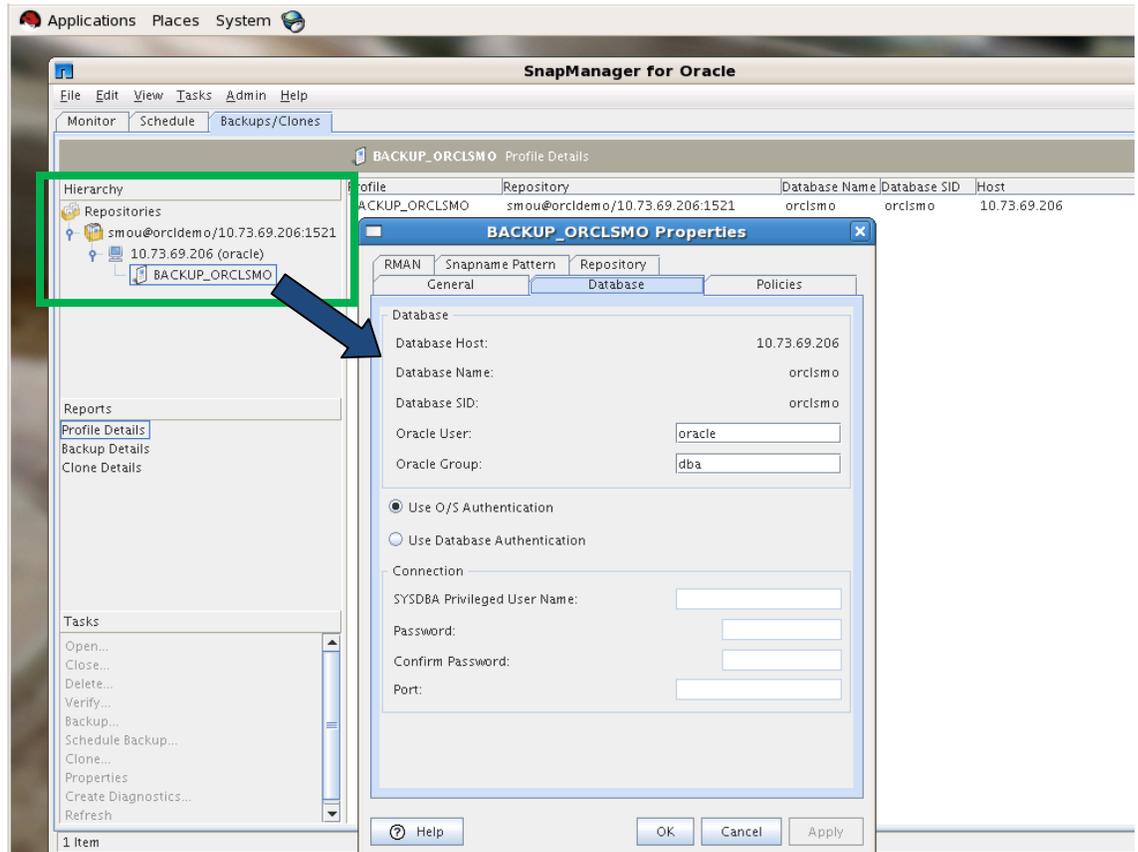


Figure 50) SMO server and repository setup.

5.3.11 STEP 11: ROLE-BASED ACCESS CONTROL TO THE STORAGE SYSTEM HOSTING PRODUCTION DATABASE FROM SNAPMANAGER FOR ORACLE AND CLONE VIRTUAL MACHINES

SnapManager 3.0 for Oracle now allows database administrators to control access based on their roles by leveraging the RBAC capabilities of Operations Manager. This integration of SnapManager with Operations Manager is actually via SnapDrive 4.1 for UNIX. Older versions of SnapDrive provided limited access control and allowed only the root user to perform SnapDrive operations. SnapDrive 4.1 for UNIX now provides controlled access to even nonroot local or NIS users by using the RBAC infrastructure of Operations Manager 3.7.1. Operations Manager provides granular access to storage objects such as LUNs, qtrees, volumes, and aggregates. Because of this integration, SnapManager for Oracle administrators can now control which SnapManager operations each database administrator can perform. Note that SnapManager for Oracle requires Operations Manager only if role-based access control is desired.

Refer to Appendix B and Appendix C of the [SMO Best Practices Guide](#) for a detailed step-by-step process of creating role-based access control configurations with SMO and Operations Manager.

Once the users and roles have been created by Operations Manager in conjunction with the storage system, the following screenshot shows:

5. How to register the credentials of the “sd-admin” user created on the DFM server host with the SnapDrive of the target database server (i.e., the cloned virtual machine)
6. How the DFM-created user called “sd-tardb_host1” on the storage system (10.73.68.199) is registered with the SnapDrive of the cloned virtual machine on which the database clones will be mounted

```
[root@OrclDevTestVM05 ~]# snapdrive config set -dfm sd-admin 10.73.65.173
Password for sd-admin:
Retype password:
[root@OrclDevTestVM05 ~]#
[root@OrclDevTestVM05 ~]# snapdrive config set sd-tardb_host1 10.73.68.199
Password for sd-tardb_host1:
Retype password:
[root@OrclDevTestVM05 ~]#
[root@OrclDevTestVM05 ~]# snapdrive config list
username          appliance name    appliance type
-----
sd-tardb_host1    10.73.68.199     StorageSystem
sd-admin          10.73.65.173     DFM
[root@OrclDevTestVM05 ~]#
```

5.3.12 STEP 12: CLONE THE PRODUCTION DATABASE AND MOUNT IN ON CLONE VIRTUAL MACHINES USING SMO

1. Right-click the SMO profile and click Backup.

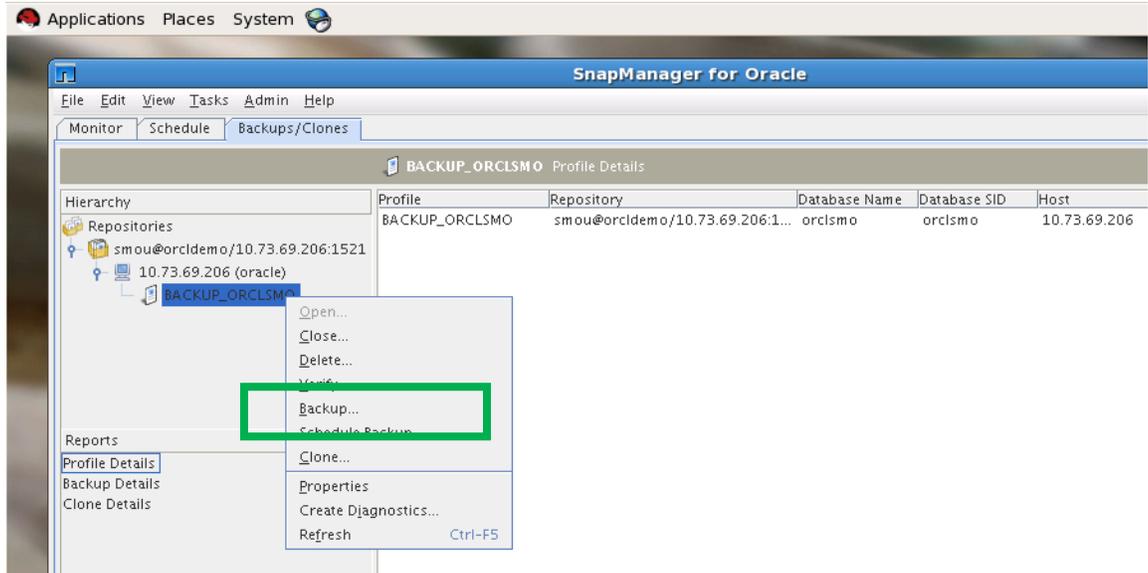


Figure 51) Cloning database and mounting it to virtual machines using SMO.

2. The SMO Backup wizard opens. Click Next.
3. Enter a label name for the backup and click Next.

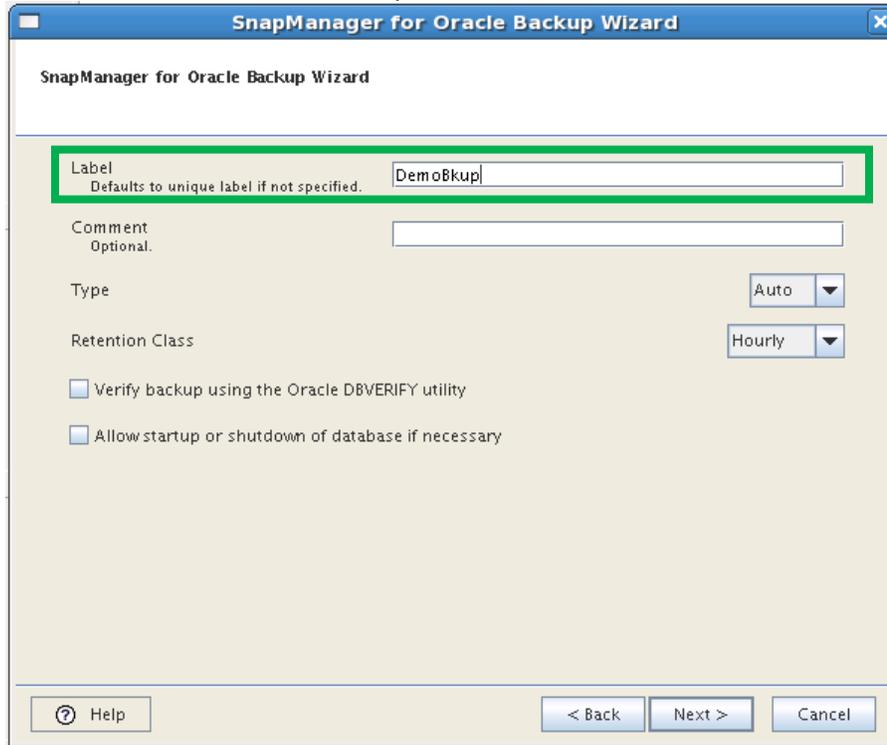


Figure 52) Cloning database and mounting it to virtual machines using SMO.

4. Select Full Backup.

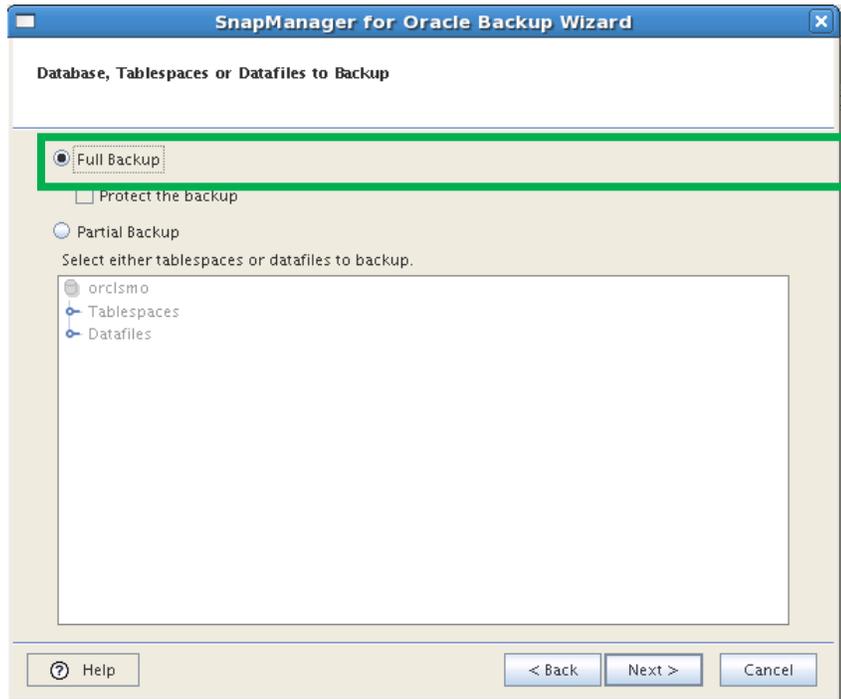


Figure 53) Cloning database and mounting it to virtual machines using SMO.

5. Commit and complete the backup.

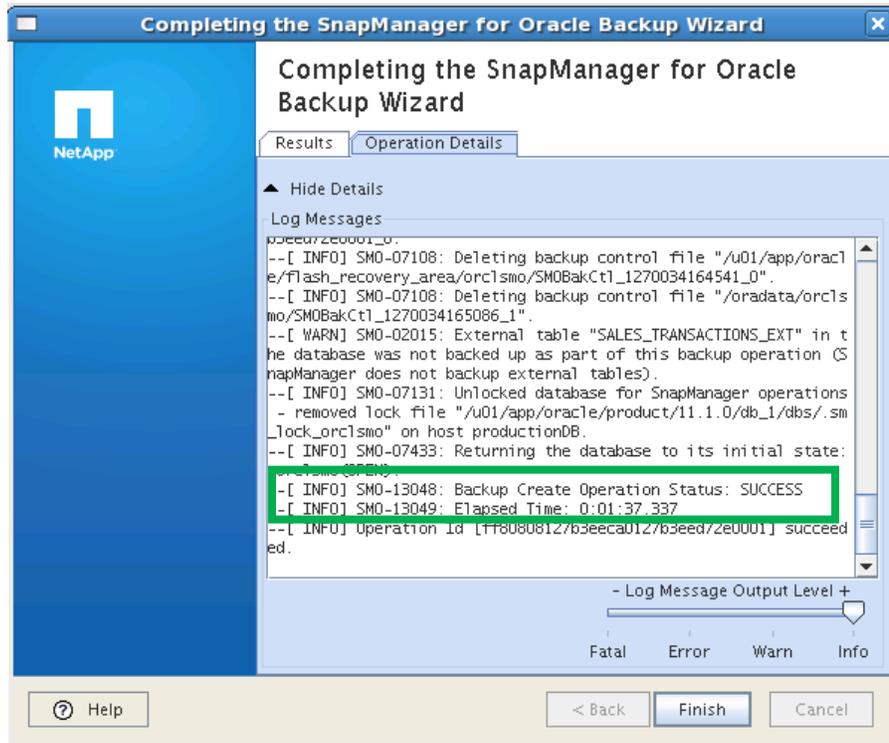


Figure 54) Cloning database and mounting it to virtual machines using SMO.

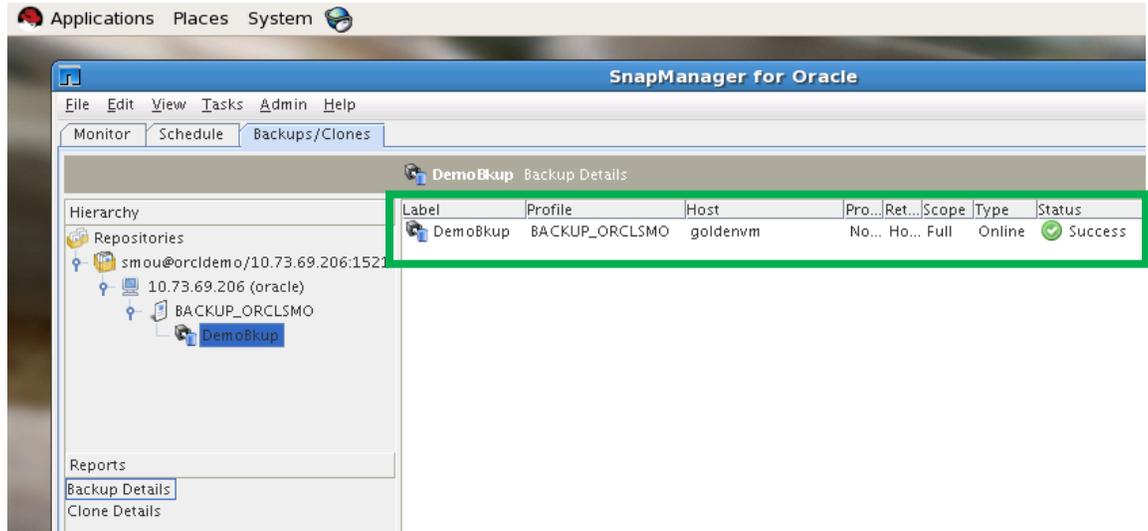


Figure 55) Cloning database and mounting it to virtual machines using SMO.

6. To create a clone of the database from the backup just taken, right-click the backup and click Clone.

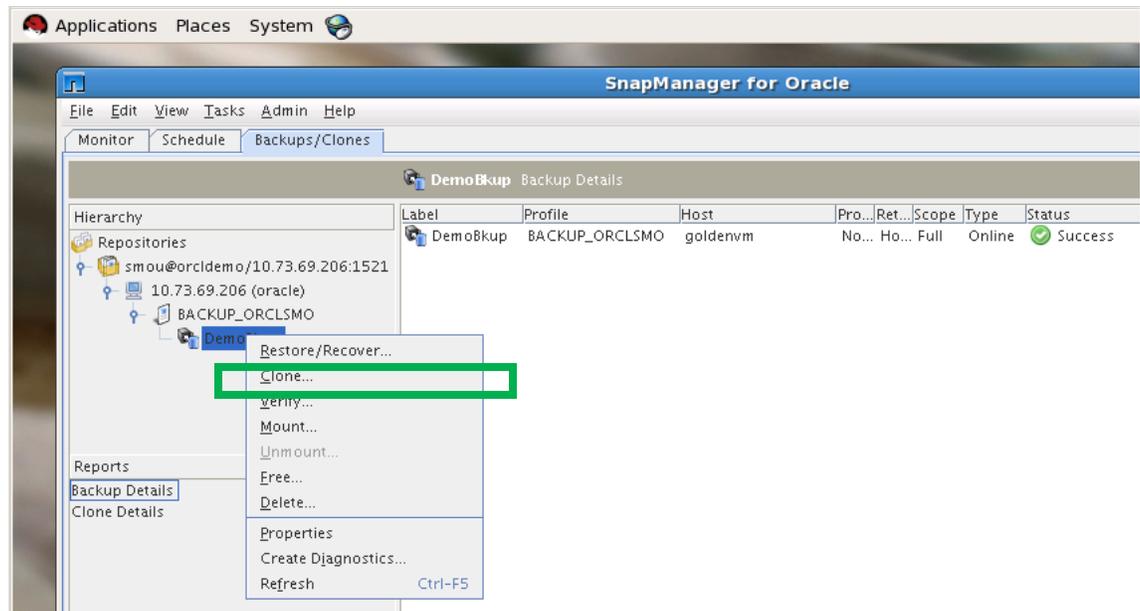


Figure 56) Cloning database and mounting it to virtual machines using SMO.

- SMO Clone Wizard opens. Click Next.
- Enter the name of the new SID and enter the IP address/hostname of one of the cloned virtual machines created in step 9 on which the database clones will be mounted.

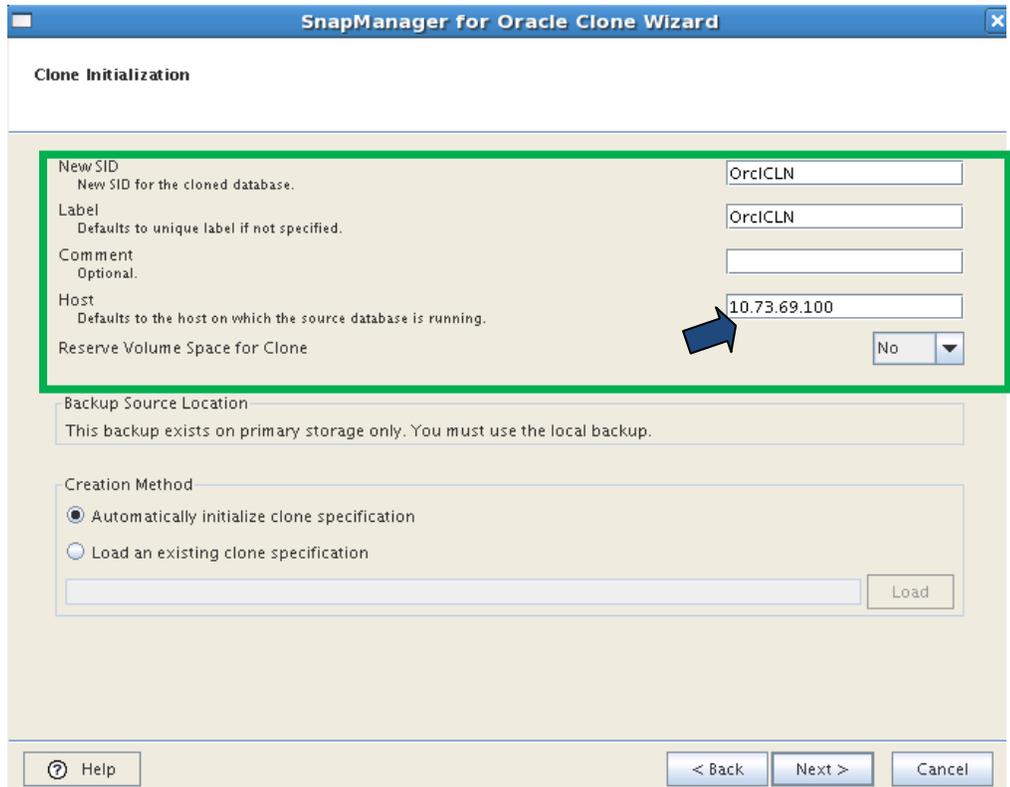


Figure 57) Cloning database and mounting it to virtual machines using SMO.

7. Enter the host credentials for the clone virtual machine. Use the `oracle` user for authentication.

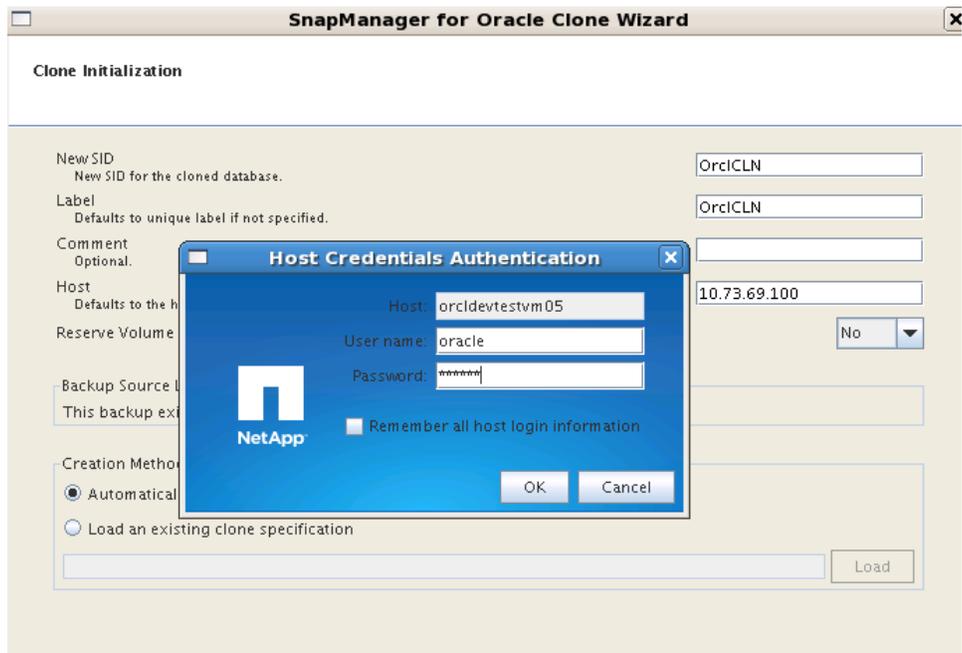


Figure 58) Cloning database and mounting it to virtual machines using SMO.

8. Verify the clone specifications.

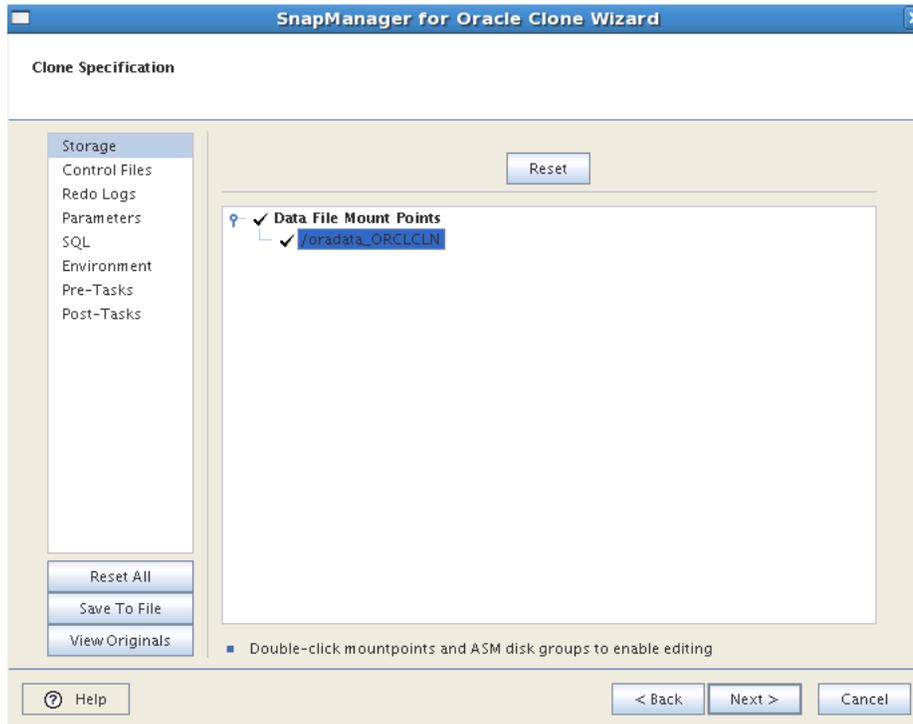


Figure 59) Cloning database and mounting it to virtual machines using SMO.

9. Perform the clone operation by clicking 'Clone.'

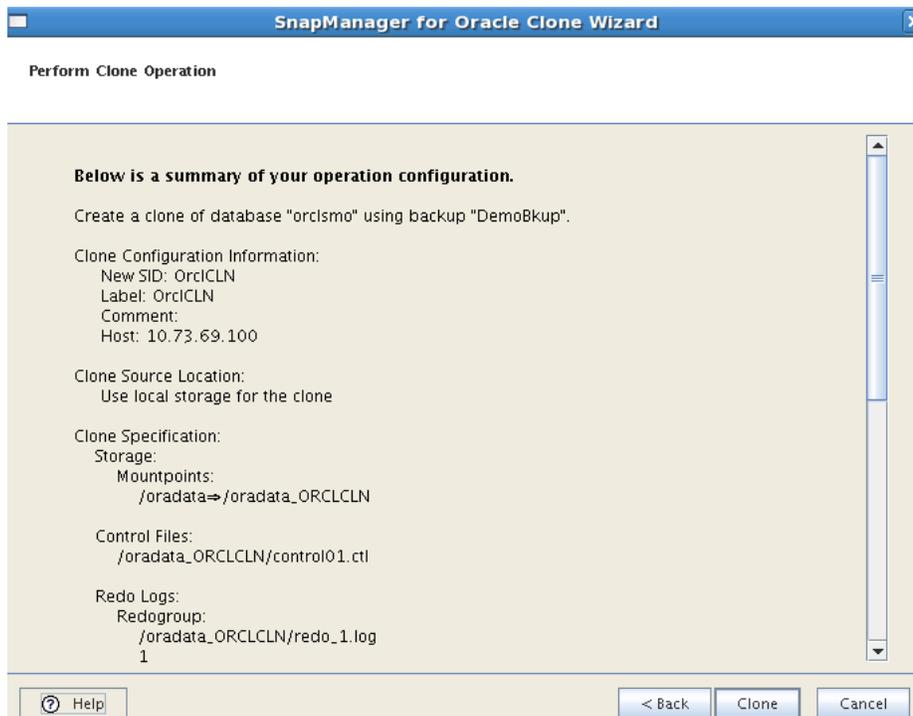


Figure 60) Cloning database and mounting it to virtual machines using SMO.

10. Verify that the clone operation is successful.

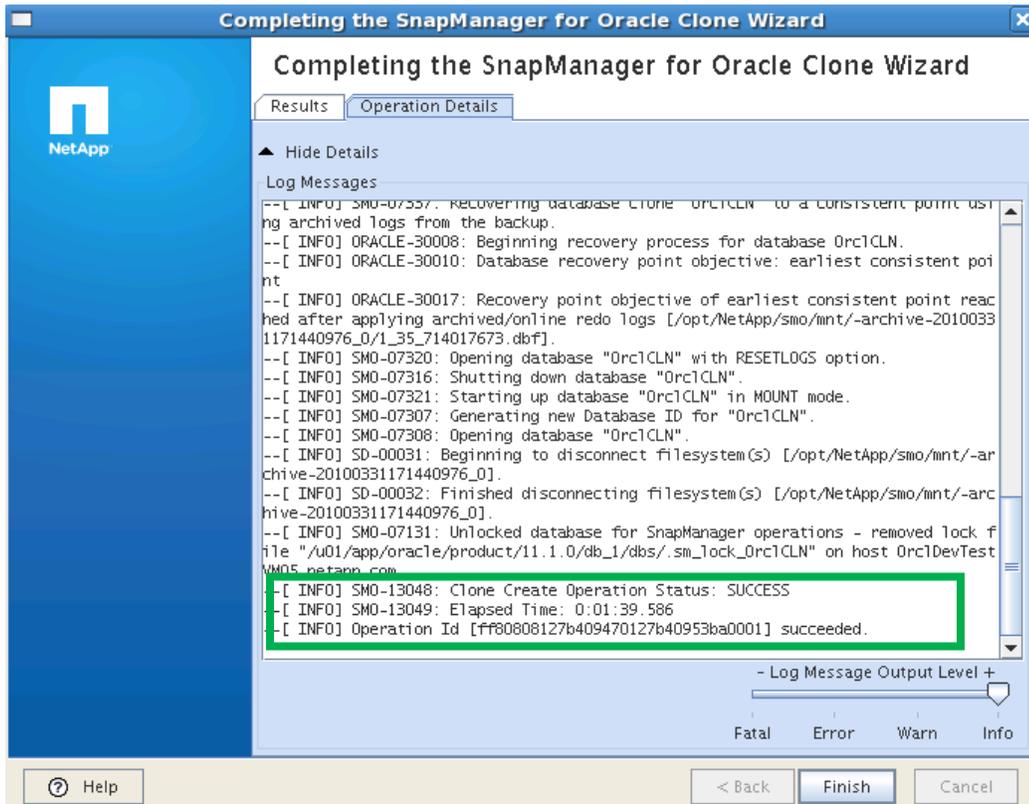


Figure 61) Cloning database and mounting it to virtual machines using SMO.

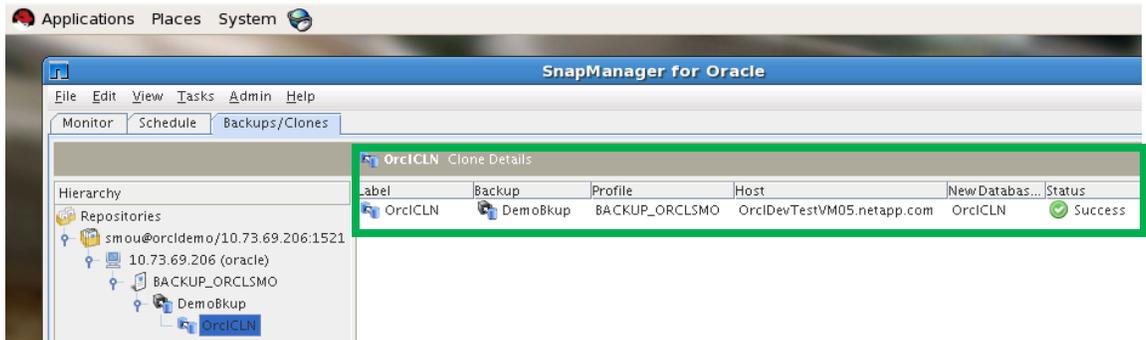


Figure 62) Cloning database and mounting it to virtual machines using SMO.

5.4 SOLUTION VALIDATION—SOLUTION VERSION 2

Table 8) Solution validation.

#	Steps	Comments
1	Verify application	Verify that the database cloned by SMO was properly mounted on the cloned virtual machine
2	Verify storage efficiency	Verify from the VSC console that no extra space in the NetApp storage has been consumed because of the cloned virtual machines and Databases

5.4.1 VERIFY APPLICATION

1. Log in to the clone virtual machine used with the SMO clone operation and check if SMO could successfully mount the cloned database on it.

```

root@OrclDevTestVM05:~
[oracle@OrclDevTestVM05 ~]# ps -ef | grep pmon
oracle      5486      1    0 15:24 ?        00:00:00 ora_pmon_orcl
oracle     17223      1    0 17:15 ?        00:00:00 ora_pmon_OrclCLN
root       18529 14195    0 17:36 pts/2    00:00:00 grep pmon
[oracle@OrclDevTestVM05 ~]#
[oracle@OrclDevTestVM05 ~]# ps -ef | grep OrclCLN
oracle     17223      1    0 17:15 ?        00:00:00 ora_pmon_OrclCLN
oracle     17225      1    0 17:15 ?        00:00:00 ora_vktm_OrclCLN
oracle     17229      1    0 17:15 ?        00:00:00 ora_gen0_OrclCLN
oracle     17231      1    0 17:15 ?        00:00:00 ora_diag_OrclCLN
oracle     17233      1    0 17:15 ?        00:00:00 ora_dbwr_OrclCLN
oracle     17235      1    0 17:15 ?        00:00:00 ora_psp0_OrclCLN
oracle     17237      1    0 17:15 ?        00:00:00 ora_dia0_OrclCLN
oracle     17239      1    0 17:15 ?        00:00:00 ora_mman_OrclCLN
oracle     17241      1    0 17:15 ?        00:00:00 ora_dbw0_OrclCLN
oracle     17243      1    0 17:15 ?        00:00:00 ora_lgwr_OrclCLN
oracle     17245      1    0 17:15 ?        00:00:00 ora_ckpt_OrclCLN
oracle     17247      1    0 17:15 ?        00:00:00 ora_smon_OrclCLN
oracle     17249      1    0 17:15 ?        00:00:00 ora_reco_OrclCLN
oracle     17251      1    0 17:15 ?        00:00:01 ora_mmon_OrclCLN
oracle     17253      1    0 17:15 ?        00:00:00 ora_mmon1_OrclCLN

```

2. Verify that you can connect to the cloned database and then proceed to execute your dev/test exercise.

```

oracle@OrclDevTestVM05:~
[oracle@OrclDevTestVM05 ~]# su - oracle
[oracle@OrclDevTestVM05 ~]$
[oracle@OrclDevTestVM05 ~]$ export ORACLE_SID=OrclCLN
[oracle@OrclDevTestVM05 ~]$
[oracle@OrclDevTestVM05 ~]$ sqlplus / as sysdba

SQL*Plus: Release 11.2.0.1.0 Production on Wed Mar 31 17:39:49 2010

Copyright (c) 1982, 2009, Oracle. All rights reserved.

Connected to:
Oracle Database 11g Enterprise Edition Release 11.2.0.1.0 - 64bit Production
With the Partitioning, OLAP, Data Mining and Real Application Testing options

SQL>

```

5.4.2 VERIFY STORAGE EFFICIENCY

3. Now check the storage usage of the datastore created for storing the clone virtual machines (OrclDevTestDataStore). Although the datastore now contains 25 clone virtual machines, it consumes space for only 1 virtual machine.
4. The deduplication savings (96%) and the datastore usage can also be tracked from the vCenter Server NetApp VSC console.
5. Check the storage consumed by the Oracle Database clones (FlexClone volumes) on the storage system.

6 SUMMARY

Implementation of an Oracle dev/test environment is often thought of as a series of trade-offs between increasing development/testing speed, improving developer/tester service levels, and lowering costs. That's not the case if you implement the Oracle dev/test environment on a virtual infrastructure leveraging the powerful capabilities of storage efficiency technologies from NetApp and their seamless integration with server virtualization technologies from VMware. Furthermore, companies with non NetApp storage can use V-Series controllers to utilize the efficient data management and data protection capabilities from NetApp. Many organizations use these solutions to streamline database application development, test, QA, and migration while saving time, dramatically improving efficiency and service levels, and reducing storage and administrative costs.

As a summary, efficient Oracle Database cloning with NetApp and VMware:

- Improves DBA productivity
 - Space-efficient, on-demand database cloning
 - Near-instant database recovery
- Enables rapid provisioning of server and storage resources for the dev/test environment
- Speeds up the dev/test cycle
- Enables rapid application development models
- Supports geographically distributed teams

NetApp also offers database and VMware administrators automated and powerful tools like SnapManager for Oracle and Rapid Cloning Utility to easily create nearly instantaneous database clones on demand, using negligible additional storage and without the need for extensive storage expertise or support.

This document is not intended to be a definitive implementation or solutions guide for Oracle dev/test on VMware vSphere and NetApp. Many factors related to specific customer environments are not addressed in this document. Contact NetApp support to speak with one of our virtualization solutions experts for any deployment requirement.

Please forward any errors, omissions, differences, new discoveries, and comments about this paper to preetom@netapp.com.

7 REFERENCES

[NetApp and VMware vSphere—Storage Best Practices](#)

[Rapid Database Development and Deployment](#)

[Creating an Architecture for Cloning Oracle Databases on Demand](#)

[Storage Efficiency in an Oracle Environment](#)

[Storage Efficiency Management Guide](#) in [NOW](#)

[Helping DBAs Become More Efficient—NetApp Efficiency and Manageability Advantages](#)

[Oracle Database Dev/Test Reference Architecture Using Data Guard and SnapManager for Oracle Deployment Guide](#)

8 ACKNOWLEDGEMENTS

Various teams from NetApp and VMware helped to a great extent to develop this paper. Their invaluable guidance and participation in all phases made sure of a technical report that relied on their real-world experience and expertise. We acknowledge the contributions of the following people:

Kannan Mani	Technical Solutions Architect, Oracle Solutions, VMware, Inc.
Bob Goldsand	Sr. Technical Alliance Manager, VMware, Inc.
Wen Yu	Sr. Technical Alliance Manager, VMware, Inc.
Bill L. Heffelfinger	Cloud, Database, and Business Apps Global Solutions Lead, NetApp
Uday Shet	Senior Manager, TME, NetApp
Lynne Thieme	Senior Manager, Oracle Alliance, NetApp
Steven Schuettinger	Technical Alliance Manager, NetApp
Preetom Goswami	Technical Marketing Engineer, NetApp
Abhinav Joshi	Reference Architect, NetApp
M. Vaughn Stewart	Virtualization Evangelist, NetApp
Bikash Roy Choudhury	Solution Architect, NAS BU, NetApp
Srikanth S. Raghavan	Reference Architecture Specialist, NetApp

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

© Copyright 2010 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. NetApp, the NetApp logo, Go further, faster, Data ONTAP, FlexClone, NOW, RAID-DP, SnapDrive, SnapManager, SnapMirror, and Snapshot are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Oracle is a registered trademark and Oracle 11g is a trademark of Oracle Corporation. VMware is a registered trademark and VMotion, vCenter, and vSphere are trademarks of VMware, Inc. Linux is a registered trademark of Linus Torvalds. Intel is a registered trademark and Xeon is a trademark of Intel Corporation. Windows is a registered trademark of Microsoft Corporation. UNIX is a registered trademark of The Open Group. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-3861

