



NETAPP AND ORACLE WHITE PAPER

ENABLING VIRTUALIZED GRIDS WITH ORACLE AND NETAPP

Generosa Litton, Network Appliance, Inc.

Monica Kumar, Frank Martin, Don Nalezty, Oracle

March 2008 | WP-7037-0208

ORACLE®

EXECUTIVE SUMMARY

Oracle® Grid architectures are absorbing new technologies and are continuing to evolve. The recent introduction of Oracle VM for advanced server virtualization, in combination with storage virtualization features of NetApp such as FlexVol®, FlexClone®, and FAS deduplication, make it possible to significantly increase server and storage utilization on Oracle Grids for improved efficiency and reduced operating costs. This white paper provides an overview of these technologies and highlights a case study of their adoption.

TABLE OF CONTENTS

1	INTRODUCTION.....	4
2	THE GRID AND SERVER CLUSTERING.....	4
	ENABLING SOFTWARE TECHNOLOGIES.....	4
	ENABLING HARDWARE TECHNOLOGIES.....	5
3	WHAT'S NEXT? VIRTUALIZATION AND THE GRID.....	5
	ORACLE VM.....	5
	NETAPP STORAGE VIRTUALIZATION.....	6
4	CASE STUDY: ORACLE UNIVERSITY.....	6
	MOVING TO VIRTUALIZATION.....	7
	A FAST, AUTOMATED PROVISIONING ENVIRONMENT LEVERAGING NETAPP FLEXCLONE.....	7
	BENEFITS OF VIRTUALIZATION.....	7
	ORACLE ON DEMAND.....	8
5	CONCLUSION.....	9

1 INTRODUCTION

Grid-based data center infrastructures use commodity servers and network storage to create a flexible, scalable architecture that is capable of achieving much higher levels of resource utilization than traditional data centers. Since Oracle introduced grid-optimized Oracle Database 10g™ software, many customers have benefited from Oracle applications running on Grids.

The pace of technological change has only quickened since that time, and as you would expect, grid architectures are absorbing the latest technology and continuing to evolve. The Oracle Grid is now poised to leverage the latest advancements in virtualization to further enhance the value of the Grid infrastructure.

This white paper provides an overview of the Grid infrastructure using server clustering, describes how advanced server and storage virtualization can enhance existing grids, and includes a case study about the use and benefits of these technologies.

2 THE GRID AND SERVER CLUSTERING

An Oracle Grid (Figure 1) enables you 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, with a Grid you can add capacity on demand for superior scalability. Implementers of server clustering Grids have achieved rapid return on investment while successfully consolidating disparate IT systems, raising management efficiency, and meeting higher service-level agreements (SLAs).




Figure 1) Elements of an Oracle Grid.

ENABLING SOFTWARE TECHNOLOGIES

The enabling software technologies for Oracle Grid include:

- Oracle Database 10g, the first grid-optimized database platform
- Oracle Real Application Clusters (RAC), which deliver a flexible and resilient database engine that can scale out 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

ENABLING HARDWARE TECHNOLOGIES

From a hardware standpoint, an Oracle Grid is the direct result of the availability of powerful, rack-mounted commodity servers coupled with network storage such as that from NetApp. By flexibly connecting servers to network storage, you can add server resources to a particular application and ensure immediate access to the necessary storage resources.

The results speak for themselves. Many of today's most innovative companies (see <http://www.oracle.com/customers/technologies/grid.html>) rely on Grid computing to drive their IT infrastructures and deliver a competitive advantage.

3 WHAT'S NEXT? VIRTUALIZATION AND THE GRID

Today's IT teams continue to face a host of problems. Increased power and cooling costs, physical space constraints, manpower shortfalls, and interconnection complexity all contribute significantly to the cost and feasibility of continued IT expansion.

As commodity servers support higher core counts and each processor core continues to grow in power, companies deploying the latest server technology may find that CPU utilization is well below optimal levels. This represents a major opportunity for improvement. At the same time, supporting legacy applications on the latest hardware platforms may be difficult or impossible.

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

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 virtualization software and NetApp storage virtualization capabilities to enhance the value of their existing infrastructure.

ORACLE VM

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

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

Oracle VM is fully server-centric, backed by the Oracle worldwide technical support organization, and designed for data center use. With Oracle VM you can easily load balance to ensure that resources are fully utilized and move live virtual machines from one physical server to another without any downtime or reconfiguration. A number of versions of Linux® and Windows® are supported as guest operating systems on either 32-bit or 64-bit server platforms (Table 1).

For current information, see <http://www.oracle.com/technology/tech/virtualization/index.html>.

Table 1) Oracle VM hardware requirements and supported operating environments.

Oracle VM	
Hardware Requirements	Supported Operating Environments
<ul style="list-style-type: none"> • A 32- or 64-bit server, with at least an i686-class processor. This includes all Intel® Pentium® Pro or later and all AMD Athlon/Duron processors or later. • Pentium 4 or Athlon CPU is recommended. • For unmodified guest operating systems, a CPU with hardware virtualization support is required. This includes some Intel Pentium D, Core, Core2, and Xeon™ models, and some AMD Athlon and Opteron models. 	<ul style="list-style-type: none"> • Oracle Enterprise Linux 4 and 5, RHEL3, RHEL4, and RHEL5 • Windows 2000, Windows Server 2003, and Windows XP on HV-capable hardware only

Unlike other virtualization solutions, Oracle VM uses native Linux device drivers so you never have to wait for the latest hardware to be supported by your virtualization solution. Two types of virtual machines are supported:

- **Hardware virtualized.** No modifications to the guest operating system are required. Available only on Intel VT and AMD SVM CPUs.
- **Paravirtualized.** The guest operating system is recompiled for the virtual environment for optimal performance.

Oracle VM consists of two components:

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

NETAPP STORAGE VIRTUALIZATION

NetApp delivers a complete set of complementary storage virtualization capabilities for all its storage platforms through the NetApp operating environment, Data ONTAP 7G®, specifically designed with the needs of Grid environments in mind. NetApp flexible volumes (FlexVol) ensure that storage volumes are always spread across all disk spindles for optimal performance. Storage volumes can grow or shrink easily without disruption. NetApp Snapshot™ copies make it easy to create a point-in-time, space-efficient view of an active volume in a matter of seconds and form the basis of NetApp's backup, rapid recovery, and disaster recovery capabilities, which provide significant benefits in Grid environments.

NetApp FlexClone technology makes it possible to clone an existing storage volume without creating a complete duplicate of the volume. Additional storage space is consumed only as changes are made to the clone. FlexClone has proven particularly useful in busy Oracle environments and offers unique advantages for virtual server environments and large Grid infrastructures.

4 CASE STUDY: ORACLE UNIVERSITY

Oracle University is responsible for delivering education and training classes on Oracle products. As the result of rapid business growth, Oracle University was facing some big challenges. Most Oracle classes require one or more servers configured with the product under study for demonstration purposes and hands-on activities. Recent acquisitions have resulted in a dramatic increase in the variety of classes that Oracle

University provides, creating a need for more and different server configurations. At the same time, the total number of classes per week of all types is growing dramatically, creating a need for more servers.

The result was significant problems with power and space. The Oracle University data center simply couldn't add any more physical servers to meet growing classroom demands. An additional problem was that it could take as much as 48 hours to fully provision a new classroom environment. That's a long time when you're hosting multiple events per week in locations all over the world. Provisioning limits were actually hampering the ability of Oracle University to grow.

MOVING TO VIRTUALIZATION

The solution for Oracle University was a fully virtualized Grid environment utilizing Oracle VM for server virtualization and NetApp for storage virtualization. Oracle University relies on a "stateless" model in which no particular physical device is essential. Virtual machines (running under Oracle VM connect through virtual private network (VPN) connections to back-end storage.

The first step in migration was to move to reliable network storage. Having the whole server environment on network storage made it possible to quickly move that environment from one physical server to another. The addition of server virtualization with Oracle VM supercharges that capability by making it possible to carve up the resources of a single physical server for use by multiple VMs at one time, leading to a significant reduction in overall server count.

The combination of Oracle VM and network storage makes it possible to move a running VM from one physical machine to another while the VM is in use. Oracle chose NetApp storage for this deployment because of key technical capabilities. Also, the level of service that NetApp was able to provide was head and shoulders above other vendors, and was key to selling the concept of virtualized storage in the Oracle University organization.

A FAST, AUTOMATED PROVISIONING ENVIRONMENT LEVERAGING NETAPP FLEXCLONE

Oracle University maintains about 2,000 course modules. In any given week there are about a dozen of each type of event running, although some heavy hitters such as Oracle E-business Suite may have 150 to 160 events per week. That's a lot of events to provision; it quite possibly represents the worst-case provisioning scenario. The demonstration database used for a course such as E-business Suite has grown to about 250GB. Creating a physical copy for 150 simultaneous events would consume more than 37TB of storage.

Oracle University uses NetApp FlexClone to significantly reduce its overall disk footprint. One or two master server images containing all necessary software and demonstration data sets are created for each course. When an event is scheduled, a clone of the master image is created by using FlexClone and the clone is assigned to a VM for use during the event. The unique advantage of FlexClone is that a clone consumes additional disk space only as changes are made. That might add up to 6 to 10GB of storage being used for a class—a huge improvement over the 250GB required for a physical copy.

Scripts have been written that create all the FlexClones necessary for a particular set of courses at one time. After the clones are created, additional scripts set up the logistical information. The whole process is extremely fast and efficient.

BENEFITS OF VIRTUALIZATION

The move to a Grid infrastructure using Oracle VM and NetApp storage yielded dramatic improvements for Oracle University. Today, the entire education infrastructure is supported by 234 Dell 1950 servers. Roughly 4,000 VMs are active at any given time. As shown in Table 2, floor space consumption (as measured by data center tiles) has been reduced by 50% and data center power consumption has been reduced by 40%. Although the number of environments needed to support events has grown by over 200% since 2005, the entire infrastructure is supported by only three full-time-equivalent administrators.

The provisioning process—which used to take 48 hours—has been reduced to 30 minutes because of Oracle VM and FlexClone. Oracle University currently hosts about 450 events per week on the Grid and has

experienced a significant increase in annual revenues. Without FlexClone, Oracle estimates that only 150 events per week would be possible.

In terms of storage, total storage capacity has grown by 140% since 2005 to 108TB. However, by oversubscribing storage using FlexClone, today's storage infrastructure is the equivalent of 245TB of physical storage—saving Oracle University from having to purchase and maintain an additional 137TB of storage. NetApp AutoSupport (ASUP) shows Oracle University's storage utilization to be 377%, revealing a utilization change from 2005 to 2007 to be at 403%.

Table 2) Changes in key metrics between 2005 and 2007 at Oracle University after the adoption of Oracle VM and NetApp storage including FlexClone.

	2005	2007	Change
TB Procured	44	108	140%
TB Logical	32	245	665%
Data-Center Tiles	8	4	-50%
Data-center power	28 kW	16.8 kW	-40%
# of FTE Sys. Admins	3	3	0%
TB Per Sys Admin	32	245	665%
% Utilization	75%	377%	403%
Refresh Time	48 hours	30 minutes	-99%
Events	150	450	200%

Although the provisioning needs of Oracle University may seem extreme, many other NetApp customers leverage a similar approach to accomplish test and development work, data mining, and a host of other database-related work without making full physical copies of a database.

ORACLE ON DEMAND

Oracle On Demand, Oracle's application outsourcing business, has introduced server virtualization within the Oracle On Demand Managed Services Grid. This move is expected to deliver significant benefits:

- Decreased impact of necessary hardware maintenance. When maintenance is necessary, VMs can simply be migrated to another physical server, eliminating downtime.
- Capacity on Demand. The hardware resources allocated to a particular application can be scaled up without disruption.
- Downtime resulting from hardware failures can be almost eliminated. Virtual machines can be restarted extremely rapidly in the event of hardware failure.

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

5 CONCLUSION

Adding advanced virtualization capabilities to new or existing Grid infrastructures brings all the benefits of virtualization to the Grid. Server virtualization that uses Oracle VM provides a flexible, extensible, and easy-to-manage server infrastructure that helps you squeeze the most from your server hardware to decrease power and space consumption in your data center while improving resiliency and speeding fault recovery.

NetApp storage virtualization provides a complementary set of data management tools that can simplify the administration of enterprise-class Oracle Grids and deliver a high level of storage utilization by eliminating the allocation of storage for duplicate copies of the same data.

© 2008 NetApp. All rights reserved. Specifications are subject to change without notice. NetApp, the NetApp logo, Data ONTAP, FlexClone, and FlexVol are registered trademarks and Snapshot is a trademark of NetApp, Inc. in the U.S. and other countries. Intel and Pentium are registered trademarks and Xeon is a trademark of Intel Corporation. Linux is a registered trademark of Linus Torvalds. Windows is a registered trademark of Microsoft corporation. Oracle is a registered trademark and Oracle 10g is a trademark of Oracle Corporation. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.
WP-7037-0208