

ORACLE®

VM

An Oracle White Paper
August 2011

Oracle VM 3: Server Pool Deployment Planning Considerations for Scalability and Availability

ORACLE®

Note

This whitepaper discusses a number of considerations to be made when planning an Oracle VM 3 server virtualization and management deployment in order to obtain the best results with regards to reliability, performance and scalability, and high-availability. This whitepaper presumes you have at least basic familiarity with the technical capabilities and concepts of Oracle VM 3, and it is recommended that you review the Oracle VM 3 Architecture and Technical Overview whitepaper or the product documentation prior to reading this whitepaper.

Introduction

Server virtualization environments can be complex as they often necessitate deploying new network and storage management strategies, as well as thinking through complex application deployment topologies for enterprise software. As a result, a little up-front work to plan the environment is well worth the investment in time and thinking to assure you can meet your service level commitments without having to deploy excessive hardware or complexity ‘just in case’. Planning ahead also dramatically reduces the risk of unexpected factors affecting your ability to deliver services on time and on budget.

Oracle VM is an enterprise-class server virtualization solution comprised of [Oracle VM Server for x86](#) and [Oracle VM Manager](#). Oracle VM 3 reflects Oracle’s strategic commitment to deliver application-driven virtualization—virtualization that makes the entire enterprise software and hardware “stack” easier to deploy, manage, and support in order to make IT and the business more agile.

Oracle VM has been designed from the ground-up for excellent scalability, manageability, and ease-of-use. The broad feature set of Oracle VM 3 make it faster and easier to rollout operating systems, enterprise applications, and middleware across your entire data center—not just the Oracle software environment —while reducing costs and making your datacenter or cloud environment more highly available and secure.

[Oracle VM Server for x86](#) is the most scalable x86 server virtualization solution in the market today, having been tested to handle mission critical enterprise workloads with support for up to 160 physical CPUs and 2 TB of memory. For virtual machines, Oracle VM 3 can support up to 128 virtual CPUs and 1TB memory per guest VM. Oracle VM supports industry standard x86 operating systems and servers from Oracle and other leading vendors, and it supports a broad range of network and storage devices, making it easy to integrate into your environment.

[Oracle VM Manager](#) provides an easy-use-centralized management environment for configuring and operating your server, network, and storage infrastructure from a browser-based interface (no Java client required) accessible from just about anywhere. Users leverage the virtualization manager for creating management policies as well as, cloning, sharing, configuring, booting and migrating VMs. Oracle VM can also help customers improve server utilization, achieve higher availability and better performance, while reducing costs.

By leveraging the considerations described in this whitepaper, users should be able to plan Oracle VM 3 deployments more accurately to assure they can meet the service level commitments efficiently and to minimize the risks of over (or under) resourcing their environment.

Overview of Oracle VM 3 Concepts and Components

Running on x86 servers, Oracle VM includes Oracle VM Manager and Oracle VM Server for x86.

- Oracle VM Manager** controls the virtualization environment, creating and monitoring Oracle VM servers and the virtual machines. Oracle VM Manager serves as the only administrative interface to the Oracle VM servers, unlike previous-generations of Oracle VM that were jointly administered from the management server, as well as locally from the command-line for each Oracle VM Server. Oracle VM Manager 3 is an Oracle Fusion Middleware application based on the Oracle Weblogic Server application server and Oracle Database. The Oracle VM Manager runs on 64-bit Oracle Linux 5.5 OS or later. For its management repository, Oracle VM Manager uses an Oracle Database, which can be installed either on the same management server or a separate server. Oracle Database 11g Express Edition (XE) is bundled with Oracle VM Manager 3 for customer evaluation purposes, but Oracle Database Standard Edition (SE) or Enterprise Edition (EE) it is required for production support as the management repository. All necessary licenses, including licenses for Weblogic Server and Database (SE or EE), are included at no additional charge.
- Oracle VM Server for x86** installs directly on server hardware with x86 Intel or AMD processors and does not require a host operating system. An Oracle VM Server is comprised of a hypervisor and privileged domain (Dom0) that allows multiple domains or virtual machines (i.e. Linux, Solaris, Windows, etc.) to be run on one physical machine. The Dom0 runs a process called Oracle VM Agent. The Oracle VM Agent receives and processes management requests, and provides event notifications and configuration data to the Oracle VM Manager. Oracle VM Server 3 requires 64-bit x86 hardware, but can support either 64-bit or 32-bit guest virtual machines.

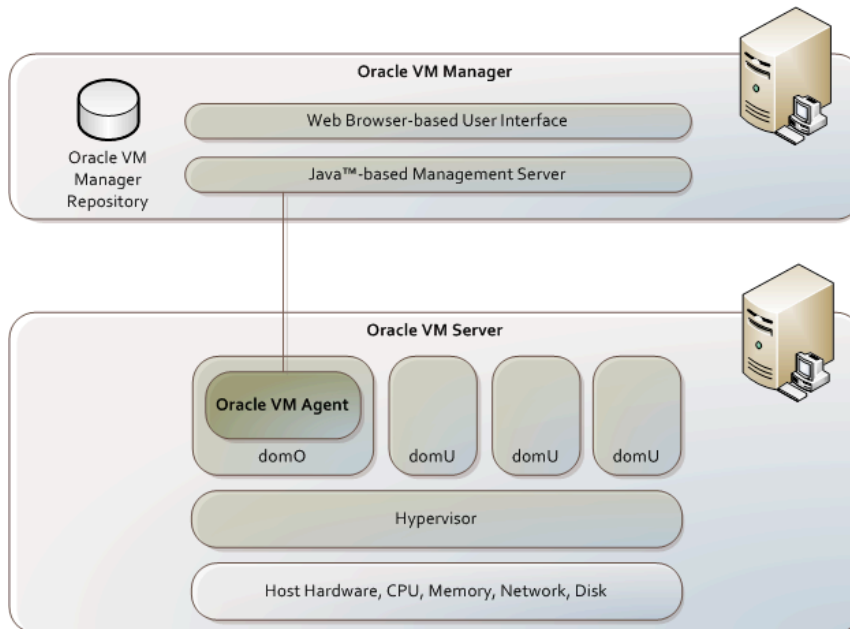


Figure 1. Oracle VM Components

Key Deployment Concepts

From a deployment perspective, multiple Oracle VM Servers are grouped into server pools as shown in Figure 2. Each server pool can have up to 32 physical servers. Every server in a given pool has access to shared storage, which can be NFS, Fibre Channel, or iSCSI (or a combination of these). This allows VMs associated with the pool to start and run on any physical server within the pool. Note that local storage can also be configured, but is often not appropriate for production environment since it prevents or sharply constrains the ability for a VM to run anywhere in the pool in the event of a server failure or live migration event.

When a VM first starts in the server pool, the VM is placed on the server that is the “best” fit based on an algorithm that takes into account the server CPU, memory and network utilization. Typically the server that is assigned to host the VM has the most resources available. Then the VM is associated with the server, but it can be securely migrated to any other server in the server pool by the administrator or can be load balanced across the pool based on a dynamic resource scheduling policy.

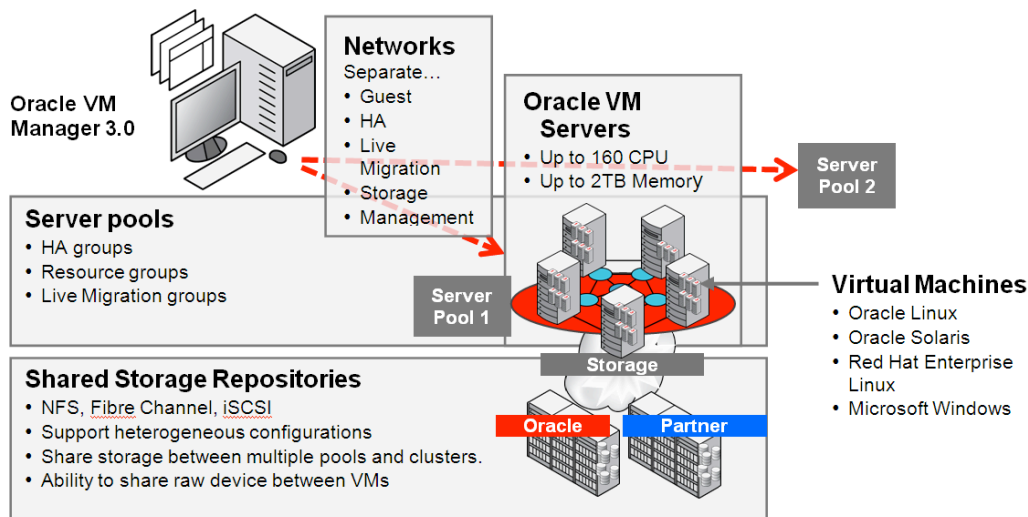


Figure 2. Oracle VM Deployment

As a result of this architecture, VMs can easily start-up, power-off, migrate, and/or restart without being blocked by the failure of any individual server, or by the failure of multiple servers as long as there are adequate resources in the pool to support the requirements for all VMs to run concurrently.

Server Pool Planning

There are a large number of considerations when planning the virtual infrastructure and one size does not fit all. This section provides some considerations and guidelines to help develop a plan that is well-suited to an organization’s unique requirements.

It may be helpful to think of the server pool as if it were one big server with an aggregate amount of CPU, memory, storage, and network bandwidth. As such, planning for deploying VMs into a pool is much like planning for a server consolidation. It involves deciding how much aggregate capacity is

needed to support normal and peak workloads as well as what types of workloads are appropriate to share the pool or server. Workload profiles should consider how predictable or unpredictable the workloads might be.

There are also some significant similarities between server pool planning and physical server planning with regards to node (physical server) size versus overall pool size. For example, in some cases it is better to have relatively fewer but larger servers in a pool. In other cases, a greater number of relatively small servers or blades could be a better fit. Both deployments may provide the same aggregate CPU, memory, storage, and bandwidth but the implications of the deployment in a pool can be different.

High Availability Planning

Oracle VM provides key features to provide maximum up time for VMs running in server pools. These features listed below should be used to collectively maximize up time of guest VMs running in a pool.

- **Guest VM HA.** Auto-restart on server or VM failure.
- **Secure Live Migration.** Move VMs off of servers that are undergoing planned maintenance.
- **Automatic pool load balancing on VM start-up.** At VM power-on, an algorithm dynamically assigns the host server for the VM in order to load balance, but also to avoid a down server blocking VM start.

The following section summarizes best practices and considerations for planning server pools.

Recommendations for Deploying Highly Available Server Pools

The following steps are recommended to achieve high availability:

- Enable the “High Availability (HA)” option at both the server pool level and individual VM level to ensure VMs are automatically restarted after an unplanned failure.
- Create the server pool file system with a size of at least 12GB, either on NFS volume or a physical disk. The server pool file system is used to hold the server pool and cluster data, and is also used for the cluster heartbeating.
- Plan to use the Secure Live Migration feature to migrate VMs in support of planned events like server maintenance to prevent any service outage.
- Plan for enough excess capacity in aggregate across the pool to support running all VMs to an appropriate service level, even when one or more servers in the pool are out of service. Up to 32 nodes are supported within a server pool.
- Plan for multiple server pools when there is a need for more than 32 physical servers.

Pool Capacity, Performance, and Scalability

Capacity planning for a server pool is similar to capacity planning for a physical server. However, the following additional considerations are also important when planning capacity for a server pool:

- **Plan extra capacity to support Guest HA/ auto-restart.** There should be sufficient capacity to support hosting additional VMs on fewer machines in the event that one or more of the servers fails and its VMs end up being restarted on the remaining, healthy servers, if only temporarily.
- **Plan for extra capacity to support Live VM Migration during planned events or to better load-balance using DRS as the pool loading changes.** When performing maintenance on a server (or servers) in the pool, Live Migration allows administrators to migrate VM(s) to another server in the pool without interrupting service. To take advantage of this capability, there should be enough excess capacity in aggregate across the pool so that a server can be taken offline (after migrating its VMs), without inappropriately impacting service levels. Further, it can be desirable to maintain some excess capacity in the pool in order to support VMs moving between servers dynamically as the load changes. For example, changes in VM workload across the day may mean that some servers incur a heavier load while others become less loaded. DRS can help migrate VMs to take advantage of the change in resource demands over time.

Determining How Many Servers or VMs a Pool Should Contain

The number of servers or VMs ideal for a pool depends on a number of factors that vary greatly between datacenters and deployments. There is no single correct answer to this question, but there are several factors that should influence such decisions. Some considerations are described below.

Storage Topologies, Performance, and Implementation

Oracle VM pools require that all servers in a pool have shared access to the same storage so that VMs can be moved around easily. This means that server pools must use shared storage such as NFS, OCFS2 (Oracle's Cluster File System), or SAN (iSCSI or FC) storage. Both the physical make-up of the storage devices and the scalability of the file systems used will dictate how many servers are practical for a given shared-storage pool without adversely affecting I/O performance.

When evaluating how many servers can share a given storage topology, the following questions should be considered:

- How much I/O will each server generate and can the throughput and latency needs be accommodated through the designated NIC or HBA ports?
- How much I/O can the storage device or devices support?
- What are the HA needs (bonding or multipath)?
- Are there any application requirements for directly accessed storage?

Of course, the answers to these storage questions depend on the I/O environment. Is the application I/O intensive? What is the average size of an I/O request?

Determining a realistic number of storage nodes in a storage cluster requires considering similar questions:

- How much I/O will each storage cluster node support?
- How much I/O will each server generate to/from the clustered file system?
- How many servers will be accessing the clustered file system?

Workload Profile

Is the workload flat and stable or variable and peaky? The ideal VM has relatively low utilization and is very flat and stable with minimal peaks. These types of VMs can often be very tightly consolidated, with a large number of VMs per server. Since they are very predictable they require little excess capacity or headroom to be available to accommodate unexpected peaks.

The next best scenario is when consolidating multiple VMs that may have peaks, but the peaks are very predictable in both timing and magnitude. For instance, some VMs may contain applications that always peak at the end of the week or month. If those VMs could be consolidated with other VMs that peak at exactly the opposite time—workloads that peak at the beginning of the week or month—they can potentially be packed tightly to maximize the number of VMs per server and per pool.

The worst-case scenario is when the VMs are highly variable in load and in timing. In this situation, it is likely that a comparatively large amount of extra headroom will be needed on the servers, thus fewer VMs will be able to be accommodated per server and per pool.

Service Level Support Strategy

Sometimes service level objectives may dictate that there is enough planned excess capacity to support normal service levels even if all the VMs peak at the same time. This is the most conservative option, but also the most expensive since it requires extra hardware that may not be utilized much of the time. Another alternative is to plan for the average load and accept any performance hit based on resource contention if there is too much of a peak. This certainly reduces hardware expense, but may not provide acceptable service levels if the workloads are either too unpredictable or mission-critical. As a result, many datacenters plan their capacity to support some percentage of the aggregate peak load. For example, they may plan for 40-60 percent of the peak above average. This is often a good compromise between meeting service levels and having reasonable hardware utilization. However, it clearly depends on how critical service levels are for the given workloads.

Dynamic Server Pools

Understanding that the server pool is dynamic, keep in mind that VMs are not associated with any one physical server until they are powered-on and placed on a server. Typically they are placed on the server that currently has the most available memory. Based on HA events or Live Migrations (manually-initiated or based on DRS policies), VMs can be moved around within the pool, thus a given VM can end-up sharing a server with any other VM in the pool unless policies are implemented to restrict which servers can host which VMs. This means that the capacity plan really needs to be at the

pool level, not the individual server level. So it may be best to consider keeping highly volatile VMs in their own pool or restricting them to a subset of the pool where a relatively large amount of excess capacity can be maintained for handling unpredictable peaks.

Conversely, highly predictable VMs should be restricted to a separate pool where the resources can be very tightly planned for high utilization without the need for much excess capacity.

Recommendations for Server Pool Capacity, Performance, and Scalability Planning

The following guidelines are recommended for designing and sizing the server pool:

- Plan for excess resource capacity at the pool level to support advanced features such as Live Migration, DRS, and guest VM HA.
- When determining the number of nodes in a server pool, consider storage topologies and their characteristics as well as network requirements, workload characteristics, and HA needs.
- Plan excess capacity according to business requirements for meeting peak loads versus only a proportion of the peak load.
- Memory capacity is the most critical resource. I/O capacity and then CPU capacity are the second and third priorities respectively. CPU and I/O should be balanced given that I/O activity is often CPU intensive.
- The amount of memory required for all running VMs must never exceed the amount physically available on the server(s) in the pool. However, the total amount of memory required by all VMs assigned to the pool (running + powered off) may exceed the total physical memory since powered off VMs do not consume memory.
- CPU over commitment is supported (e.g. more virtual CPUs configured than physically present). The over commitment ratio depends upon workload requirements (see Workload Profiles section). However, the recommendation is to keep the ratio of virtual CPU to physical CPU at 2:1 or less.
- Determining the best physical server node size for a pool is a combination of factors depending on workload characteristics as well as the operational and budgetary issues of the datacenter.
- Use identically configured server nodes throughout the pool to support consistent performance and feature set regardless of individual server failure(s). In order to support the best agility and flexibility, it is also recommended that every server node in the pool be identical in capacity and configuration. This approach ensures that no matter where live migration or (re-) start occurs, the same performance and features are uniformly available.

Conclusion

As these guidelines and considerations show, a little up-front work to plan your virtualization environment is well worth the investment in time to assure you can meet your service level commitments. Planning ahead improves your ability to deliver services on time and on budget.

Oracle VM 3 demonstrates Oracle's commitment to deliver application driven server virtualization solutions. The new capabilities introduced in Oracle VM provide increased levels of scalability, manageability, and ease-of-use to help customers to make most demanding enterprise applications easier to deploy, management, and support.

The best way to experience all the benefits of Oracle VM 3 is to download the software from Oracle's E-Delivery site (<http://edelivery.oracle.com/oraclevm>) and try it out in your environment.

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



Oracle VM 3: Server Pool Deployment Planning
Considerations for Scalability and Availability
August 2011

Author: Honglin Su
Contributing Authors: Adam Hawley, Susan
Roberts, Chris Kawalek

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com



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

Copyright © 2011, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

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

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0811

Hardware and Software, Engineered to Work Together