

# Tintri VMstore with Oracle Best Practice Guide

For Deploying Oracle Database 11gR2 with VMware  
5.5 on a Tintri VMstore

January 2015

## Contents

Intended Audience .....	4
Executive Summary .....	4
Oracle Database Overview.....	5
Consolidated List of Practices.....	5
Configuration Details.....	7
Preparing to Deploy Oracle with Tintri.....	7
Networking Guidelines.....	8
VMware vSphere (ESX) Host Guidelines .....	9
VMware vSphere Virtual Machine Start Order .....	9
Use VMware vSphere vApps to Control Startup Order .....	9
VMware vSphere High Availability (HA) Settings .....	9
VMware vCenter vSphere App HA.....	10
VMware vCenter DRS Cluster Settings.....	10
VM Guest Guidelines .....	11
VM Guest Operating System Considerations.....	11
VM Guest Machine Version .....	11
VM Guest VMware Tools.....	11
VM Guest Automatic Disk Partition Alignment.....	12
VM Guest Memory.....	12
VM Swap File and Guest OS Swap File .....	13
VMware Snapshots.....	14
Linux Queue Scheduler .....	14
Linux Queue Size .....	14
VM Guest Linux Device Readahead Settings .....	15
VM Guest Testing the Storage Throughput with the “dd” Command.....	15
Deploying Oracle with Tintri .....	16
Types of Oracle Database Deployments Supported by the Tintri VMstore .....	16
Deployment Architecture Options .....	17
VMware PVSCSI Adapter.....	18
Virtual Disk Provisioning Mode.....	18
vDisk Sizing .....	19
Using Oracle OFA with the Tintri VMstore.....	20
Deploying Oracle OFA with a Single VMDK.....	21
Deploying Oracle OFA with one PVSCSI Adapter and Multiple VMDKs .....	23

Deploying Oracle OFA with Multiple PVSCSI Adapters and Multiple VMDKs .....	25
Deploying Oracle ASM with the Tintri VMstore.....	27
Tintri Deduplication and Compression versus Oracle Database Compression .....	30
Oracle DB_BLOCK_SIZE Initialization Parameter .....	30
Oracle Database Protection and Backup Options .....	31
Using Oracle RMAN for Database Backups .....	32
VMware Snapshots.....	32
Tintri SnapVM™ Snapshots .....	32
Tintri SnapVM VM-Consistent Snapshots.....	33
VMware Tools Scripts.....	34
Tintri SnapVM Crash-Consistent Snapshots .....	34
Tintri ReplicateVM for VM-Level Replication .....	35
Performance Analysis.....	36
Tintri VMstore Performance Dashboard .....	36
VMstore Performance counters .....	36
The performance reserves gauge .....	36
The physical space capacity gauge.....	37
The “days till full” counter.....	37
Performance reserves changers.....	37
Space changers.....	37
Virtual Machines Tab .....	38
IOPS Graph.....	38
MBps Graph .....	38
Latency Graph.....	38
Conclusion.....	40
References .....	41
Tintri Links.....	41
Oracle Links.....	41
RedHat Links .....	41
VMware Links.....	41
Appendix A – Support Statements.....	42

## Intended Audience

This Best Practice Guide for deploying Oracle Databases on Tintri VMstore systems will assist individuals who want to architect and deploy production Oracle databases with VMware and Tintri VMstore™ storage systems. Although not required, prior knowledge of Oracle database, RHEL, virtualization, networking, and storage technology will help in understanding the concepts covered in this best practice paper.

## Executive Summary

Per [Gartner](#), at least 70% of workloads in the enterprise datacenter have been virtualized. The successful deployment of virtualized database workloads with VMware is made possible with the right choice of storage. Tintri designed the VMstore storage platform specifically to support the I/O requirements of virtualized workloads. Thus it provides an exceptional storage platform for the deployment of mission-critical databases with the VMware hypervisor.

Tintri smart storage sees, learns and adapts, enabling virtualization administrators and DBAs to focus on running the database instead of managing the storage infrastructure. Tintri application-aware storage eliminates second-guessing and complex troubleshooting by providing VM-level visibility, control, performance insight, and consistent high levels of performance from flash. Today, the Tintri VMstore powers hundreds of thousands of virtual machines running business critical databases, enterprise apps, desktop and mobile apps, and private clouds.

- **Run multiple databases** – Eliminate the performance issues that ordinary all flash arrays suffer when they are used to host multiple databases. Tintri VM-aware performance technology isolates individual VMs into their own QoS queues. This technology prevents high pressure I/O tasks, such as backups and full table scans, which run in one database from robbing performance from other databases.
- **Accelerate application testing and deployment lifecycles** – Tintri space and performance efficient per-VM cloning technology makes it possible to rapidly deploy hundreds of Oracle databases, reducing time to deploy test and development databases, and shortening application development and test cycles.
- **Simplify Database Performance Troubleshooting** – Tintri VMstore performance tools provide detailed per-VM and per-Virtual Disk statistics, including detailed host, network, and storage latency statistics, which make it simple to visualize and identify performance bottlenecks within the application running in the VM. With Tintri the DBA has new insight into the performance of the Oracle database, and individual datafiles, making it simple to monitor and diagnose performance issues with the database.
- **Consistent High Performance** –Based on patented FlashFirst™ design, the Tintri VMstore services over 99 percent of all I/O from flash, providing high levels of IOPS, throughput, and consistent sub-millisecond latencies. VMstore delivers VM-level QoS to guarantee every VM gets the performance it needs without any manual tuning or configuration. The VMstore delivers non-disruptive and automatic VM alignment, which eliminates the performance-robbing overhead that comes from ordinary storage.

Deploying storage into your virtual environment should be a simple and straight forward process. The Tintri VMstore is designed so that IT administrators with a working knowledge of VMware vSphere™ can successfully deploy Tintri's purpose-built storage quickly and easily.

## Oracle Database Overview

Oracle makes powerful, extensible and reliable database application software that is deployed by thousands of companies to store, manage and derive business value from structured data. Oracle can be deployed as a stand-alone instance of the database or with multiple database servers in what is called a Real Application Cluster (RAC). Recent versions of Oracle database software include Oracle 12c, where the “c” stands for cloud computing, and Oracle 11g, where the “g” stands for grid architecture. This paper focuses on stand-alone instances of the Oracle database, but many of the recommendations will apply to Oracle RAC.

## Consolidated List of Practices

This table summarizes the recommended practices. Click the text on any of the recommendations to jump to the section that corresponds to the recommendation.

- DO:** Use vApps to group the Oracle database VM with related application VMs.
- DO:** Set VM Restart Setting to High on the Oracle database VMs.
- DO:** Use DRS rules to ensure that Oracle VMs always have the required resources.
- DO:** Review the VMware Best Practices documents for Oracle databases.
- DO:** Disable unnecessary foreground and background processes within the guest operating system.
- DO:** Schedule backups and virus scanning programs to run in off-peak hours.
- DO:** Use OS kernel parameters recommended by the Oracle Installation guide.
- DO:** Use the latest supported VMware Virtual Machine Version in your guest.
- DO:** Install the latest version of VMware Tools in the guest operating system.
- DO:** Modify VM Options to set the memory reservation equal to, or larger than, the size of the Oracle SGA.
- DO:** Store the Virtual Machine swap file in the same directory as the Virtual Machine.
- DO:** Follow Oracle best practices to set the size of the Guest OS swap = 1.5 times the memory in the VM.
- DO:** Delete VMware snapshots as soon as you are done with them.
- DO:** Set the Linux I/O queue scheduler to NOOP to increase throughput and decrease latency.
- DO:** Increase the Linux Queue size if the database requires large amounts of write I/O.
- DO:** Check whether database performance can benefit by changing the Linux device readahead settings.
- DO:** Use the “dd” command with a block size of 64k.
- DO:** Use Tintri VMstore tools to monitor the performance of Oracle database VMs.
- DO:** Choose either the Oracle ASM or Oracle OFA model for managing the database files.
- DO:** The recommended VM deployment architecture for Oracle with Tintri VMstore is to use multiple VMDK files with multiple PVSCSI adapters.
- DO:** Use the Paravirtual SCSI adapter (PVSCSI) when deploying Oracle databases on Tintri VMstore systems
- DO:** Use thin-provisioned virtual disks.
- DO:** Monitor disk capacity reserves of the Tintri VMstore and migrate VMs if additional space is needed.
- DO:** Use differently sized vDisks to avoid confusion when managing multiple virtual disks.

**DO:** Create the OFA disk layout that supports your database manageability and I/O requirements

**DO:** Turn off Oracle compression technologies.

**DO:** Use Tintri VMstore's built-in deduplication and compression technology with Oracle databases to efficiently reduce the amount of data being stored in flash.

**DO:** Use a DB\_BLOCK\_SIZE set to a multiple of 8K.

**DO:** Use Tintri SnapVM™ snapshots to protect the Oracle database and the datafiles.

**DO:** Install and update to the latest version of Use VMware Tools.

**DO:** Use pre-freeze and post-thaw scripts with VM-Consistent snapshots.

**DO:** Use the VMstore Performance Dashboard to view the performance of the Oracle VM and troubleshoot latency breakdown across the infrastructure at a virtual disk and VM level.

**DO:** When architecting the Oracle database, use multiple Virtual disks to store datafiles. This allows you to use the Tintri performance insight tools to pinpoint the actual cause of latency (host, network, or storage) and quickly indentify and resolve any performance issue with the Oracle database.

## Configuration Details

This paper summarizes the best practices for the use of Tintri VMstores with Oracle and Red Hat Enterprise Linux (RHEL) on virtualized servers running on VMware vSphere. Although it provides specific recommendations for the use with the Tintri VMstore, it is not intended to replace the best practices provided by Oracle, RHEL and VMware for their respective platforms. We recommend that you download and follow the best practices that have been provided by those companies. A list of best practices papers for Oracle, RHEL and VMware are found in the appendix of this document.

These best practices are based on the installation Oracle 11gr2 (11.2.0.3) and RHEL 6.4 with VMware vSphere 5.5 and Tintri OS 3.1.0. A 10GbE network was used to connect the Dell 720R server to the Tintri VMstore. These best practices can also be applied to Oracle 12g and RHEL 7 running on versions of vSphere that are compatible with Tintri VMstores.

This document assumes you are working with a fully configured, highly-available VMware Infrastructure. Configuration of the hosts, networking, storage, and virtual infrastructure components are out of scope of this document. The Tintri Best Practices Guide for NFS (Network File Systems) and vSphere will assist IT administrators who are responsible for the design, deployment and management of virtual infrastructures powered by Tintri VMstore storage appliances. Screenshots in this document are based on VMware vSphere 5.5.

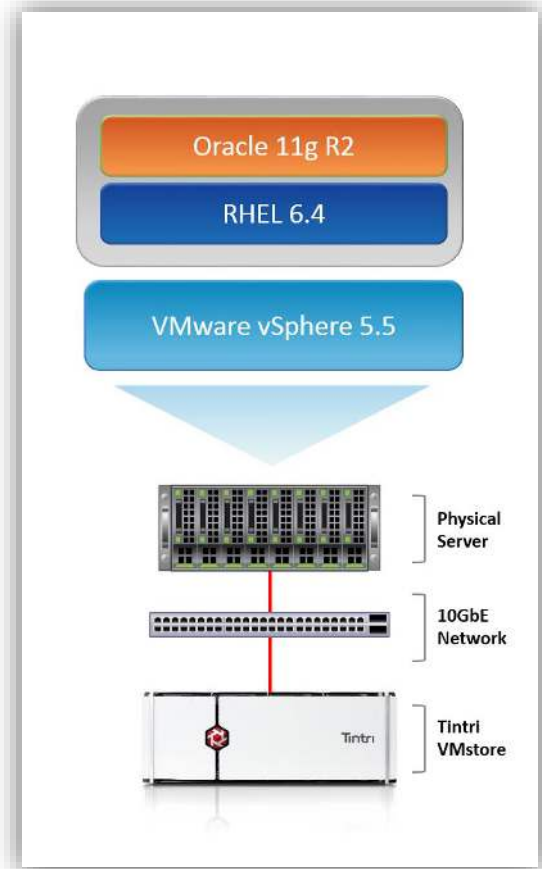


Figure 1 – Configuration Details

## Preparing to Deploy Oracle with Tintri

When deploying Oracle database as a virtual machine, care should be taken to adhere to vendors' best practices, most specifically those made available from Oracle, Red Hat and VMware. Although these Tintri best practices are straightforward, it is important to keep in mind that they are being made without specific knowledge of your unique environment.

Proper configuration contributes to the optimal performance and operation of your Oracle database. Below are some of the most common guidelines to follow when deploying Oracle databases on VMware vSphere.

- Use vSphere ESXi 5.x.
- Optimize your vSphere environment.
- Use NIC teaming for availability and load balancing.
- Use as few virtual vCPUs as possible. More vCPUs are not necessarily better and could be detrimental depending on host configuration and the number of VMs running on a host.
- Use the VMXNET PVSCSI adapters for the Oracle data files.
- Clone the completed RHEL/Oracle VM to create a Golden Image.

## Networking Guidelines

The standard VMware networking best practices apply to running Oracle databases on vSphere. It is recommended that the Tintri VMstore be attached to the ESX Host via a 10GbE Ethernet network. The performance and reliability of your database can be affected by the performance of the network that connects the ESX host and the Tintri VMstore. For high performance, we recommend the following best practices:

- Use the VMXNET3 PVSCSI network adapters in the VM. The VMXNET3 adapter is designed to provide high performance network access for virtual machines.
- Use 10GbE networking between the ESX host and the Tintri VMstore.
- Use multiple physical 10GbE switches with link aggregation for high availability. See [best practices](#) for deploying Tintri VMstore systems in VMware vSphere environments.
- Use NIC teaming for availability and load balancing. A NIC team can share the load of traffic between physical and virtual networks as well as provide passive failover in the event of a hardware failure or network outage.
- Use a subnet to isolate the traffic between Tintri VMstore and the ESX host(s).
- If you elect to use jumbo frames they must be enabled for each vSwitch as well as for the VMkernel network. Jumbo frames must also be enabled throughout the infrastructure including the physical hardware, which includes the network switches and the Tintri storage arrays.

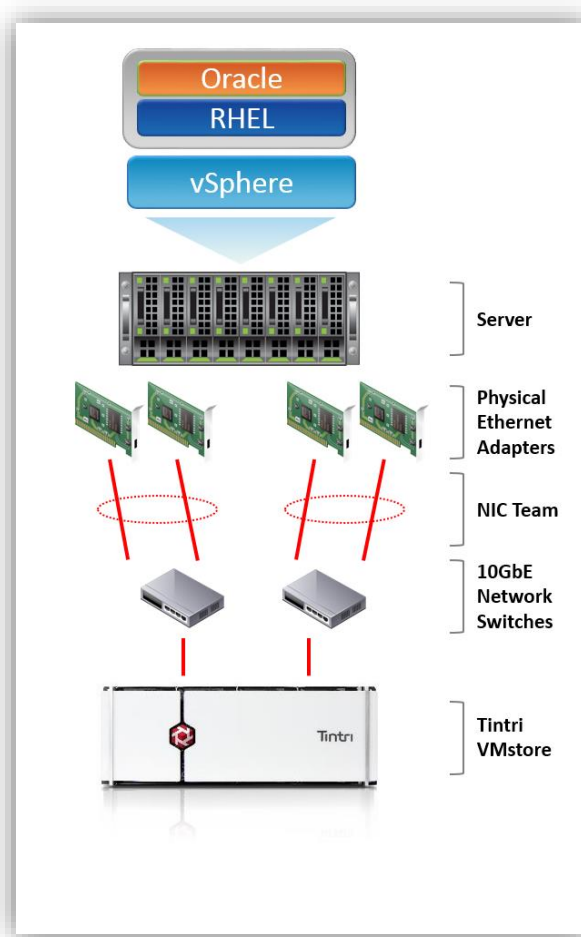


Figure 2 – Example Network Architecture



## VMware vSphere (ESX) Host Guidelines

The VMware Best Practices Guide for running Oracle Databases on VMware is the authoritative source for information on sizing and configuring an ESX host for running Oracle. According to that [document](#), VMware vSphere 5.5 supports a full range of Oracle database System Global Area (SGA) sizes, including new large capacity virtual machines that makes it possible to run Oracle databases with large SGA footprints.

- ESXi 5.5 host can support 320 CPUs, 4TB RAM, 512 virtual machines, and 4096 virtual CPUs.
- Individual virtual machine (VM) can support 64 virtual CPUs, 1TB RAM and 60 Virtual Disks.

## VMware vSphere Virtual Machine Start Order

The startup order and timing for individual VMs can be configured on standalone VMware (ESX) hosts. This capability should be used to ensure that the Oracle database VM gets started before the other VMs that use, or depend upon, the Oracle database. The order used to shut down the VMs is also configurable.

Note: the Virtual Machine Startup and Shutdown (automatic startup) feature is disabled for all virtual machines residing on hosts that are in a vSphere HA cluster. Automatic startup is not supported with vSphere HA.

## Use VMware vSphere vApps to Control Startup Order

When the Oracle database VM is running in a HA cluster use the vSphere vApp feature to group the Oracle database VM with other applications that use the database. VMware uses a vApp as a container that includes one or more virtual machines. The vApp startup order is used with a group of individual VMs to ensure that an application that includes multiple VMs boots up in the correct sequence. An example of this is a 2-Tier web application which requires the Oracle database to be started and running before the applications and the web server are brought online. vApps are an easy way to ensure the correct startup/shutdown order.

---

*DO: Use vApps to group the Oracle database VM with related application VMs.*

---

## VMware vSphere High Availability (HA) Settings

VMware vSphere High Availability (HA) can be used to provide a type of high availability for Oracle database VMs by detecting hardware failures in a cluster of vSphere (ESX) hosts. If a failure is detected, VMware vSphere HA will restart all of the virtual machines on another node. If this method of high availability is used for an Oracle database, then startup scripts are required to restart the Oracle instance.

VMware provides the ability to set a startup order for virtual machines. We recommend setting the “VM Restart Setting” to “High” on Oracle database servers to ensure these VMs are prioritized in the event of a host failure. Oracle database servers are often pre-requisites for other systems in multi-tiered applications, and should be brought back online before other tiers are brought up (ex. Web application front-end that relies upon an Oracle database backend).

Although VMware vSphere High Availability is a cost effective solution to protect against hardware failures, there is no monitoring of Oracle instances inside the virtual machine. This means there will be some downtime for recovery after failure. Time is required to restart the virtual machine, for the OS to boot, for Oracle to start and then complete the database instance recovery.

---

*DO: Set VM Restart Setting to High on the Oracle database VMs.*

---

## VMware vCenter vSphere App HA

When a VM with an Oracle database is up and running, it doesn't necessarily mean that the database instance is available and accepting connections. Fortunately, vSphere App HA can be used to monitor Oracle and other applications running in a VM by defining availability policies and remediation actions.

## VMware vCenter DRS Cluster Settings

VMware vSphere includes a load balancing mechanism called the Distributed Resource Scheduler (DRS). DRS monitors the ESXi hosts and the resource requirements of the individual VMs. If the resources on the ESXi host changes, then DRS can automatically migrate a VM to a host with sufficient resources to support the VM. By using DRS you guarantee that the Oracle database VM will always get the resources it requires for optimal performance.

DRS uses Host Affinity rules to allow you to restrict movement of VMs to a subset of hosts within a cluster, so that only selected hosts in the cluster are "available" to the Oracle database.

---

**DO:** Use DRS rules to ensure that Oracle VMs always have the required resources.

---

## VM Guest Guidelines

VMware vSphere and VM Guest settings can influence the performance and scalability of your Oracle database. There is a wealth of information on VMware and Oracle performance in the VMware Best Practices guides. For information on VMware's best practices, refer to the following Best Practices guides:

### Performance Best Practices for vSphere 5.5

[http://www.vmware.com/pdf/Perf\\_Best\\_Practices\\_vSphere5.5.pdf](http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.5.pdf)

### Best Practices for Performance Tuning of Latency-Sensitive Workloads in vSphere VMs

<http://www.vmware.com/files/pdf/techpaper/VMW-Tuning-Latency-Sensitive-Workloads.pdf>

With regards to general performance best practices, VMware recommends that unnecessary foreground and background processes be disabled in the Virtual Machine, and to schedule long-running backups and virus scanning programs to run in off-peak hours.

---

*DO: Review the VMware Best Practices documents for Oracle databases.*

*DO: Disable unnecessary foreground and background processes within the guest operating system.*

*DO: Schedule backups and virus scanning programs to run in off-peak hours.*

---

## VM Guest Operating System Considerations

From the operating system perspective, a virtual machine is experienced as if it were a physical server. Identical operating system best practices apply whether the machine is running on physical hardware or is virtualized. For this paper, the recommendations in the Best Practices for Red Hat Enterprise Linux 6 were applied.

### Deploying Oracle Database 11g R2 on Red Hat Enterprise Linux 6 – Best Practices

<http://www.redhat.com/en/resources/deploying-oracle-11gr2-on-rhel-6>

---

*DO: Use OS kernel parameters recommended by the Oracle Installation guide.*

---

## VM Guest Machine Version

When creating your VM, use the latest supported VMware virtual machine version to ensure you get the latest VMware features in your VM.

---

*DO: Use the latest supported VMware Virtual Machine Version in your guest.*

---

## VM Guest VMware Tools

Install the latest version of VMware Tools in your VM. This will ensure you have the right drivers installed for virtual components, ensuring optimal performance and stability. VMware Tools also allows the hypervisor to communicate with the guest to properly quiesce the vDisks, a prerequisite for using some backup applications and for taking "VM-Consistent" Tintri snapshots.

---

*DO: Install the latest version of VMware Tools in the guest operating system.*

---

## VM Guest Automatic Disk Partition Alignment

VMware highlights the importance of disk partition alignment and cautions that misaligned disk partitions result in poor performance. When the disk partitions on the storage are correctly aligned on the storage the VM will achieve increased throughput and decreased latency. Fortunately, you never have to worry about disk partition alignment on Tintri VMstore systems as disk partitions are automatically aligned without any user intervention. This feature guarantees that VMs will always achieve the highest throughput and lowest latency from the Tintri VMstore.

The Tintri VMstore UI can also show you whether VMs that were migrated from existing storage to Tintri VMstore systems have suffered from disk partition misalignment problems in the past. In the Virtual Machine view of the Tintri VMstore UI, add the “Aligned IO %” column to the header at the top of the page with a simple right-click of the mouse. Regardless of the number reported, if the VMDK is stored on the Tintri VMstore, no action on your part is required to ensure great performance.

## VM Guest Memory

The VMware Best Practices Guide for running Oracle recommends that memory reservations for a VM Guest running an Oracle database be set to the size of the Oracle SGA. This setting will minimize memory swapping between vSphere (ESX) and the guest OS.

---

**DO:** Modify VM Options to set the memory reservation equal to, or larger than, the size of the Oracle SGA.

---

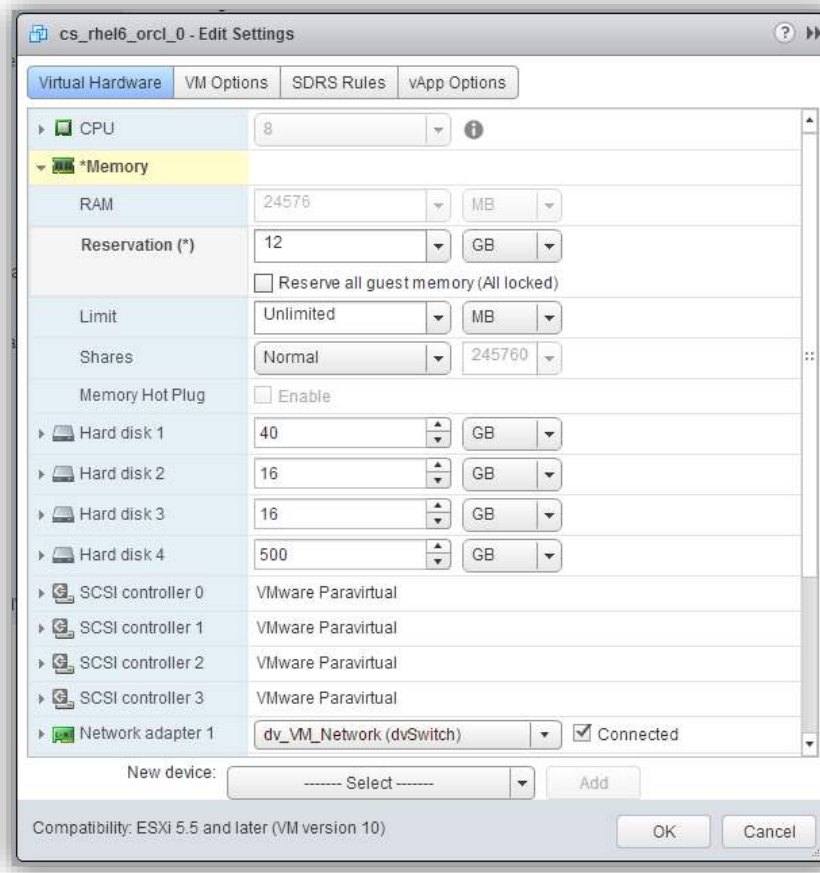


Figure 3 – Setting Memory Reservations in the VM Options Edit Settings Dialog Box

## VM Swap File and Guest OS Swap File

There are two swap files to consider – the VM swap file and the Linux Guest swap. The ESX host stores the Virtual Machine swap file in the same directory as the virtual machine. We recommend this configuration.

The guest operating system in the virtual machine also uses a swap file. VMware recommends sizing and placing the Linux swap file using the guidelines provided by RHEL and Oracle. According to Oracle best practices, the Linux swap space should be 1.5x the size of the memory assigned to the virtual machine.

Linux operating systems refer to swap space as swap files. The following commands manage swap files.

`mkswap` — Sets up a Linux swap area.

`swapon` — Enables devices and files for paging and swapping.

---

**DO:** Store the Virtual Machine swap file in the same directory as the Virtual Machine.

**DO:** Follow Oracle best practices to set the size of the Guest OS swap = 1.5 times the memory in the VM.

---

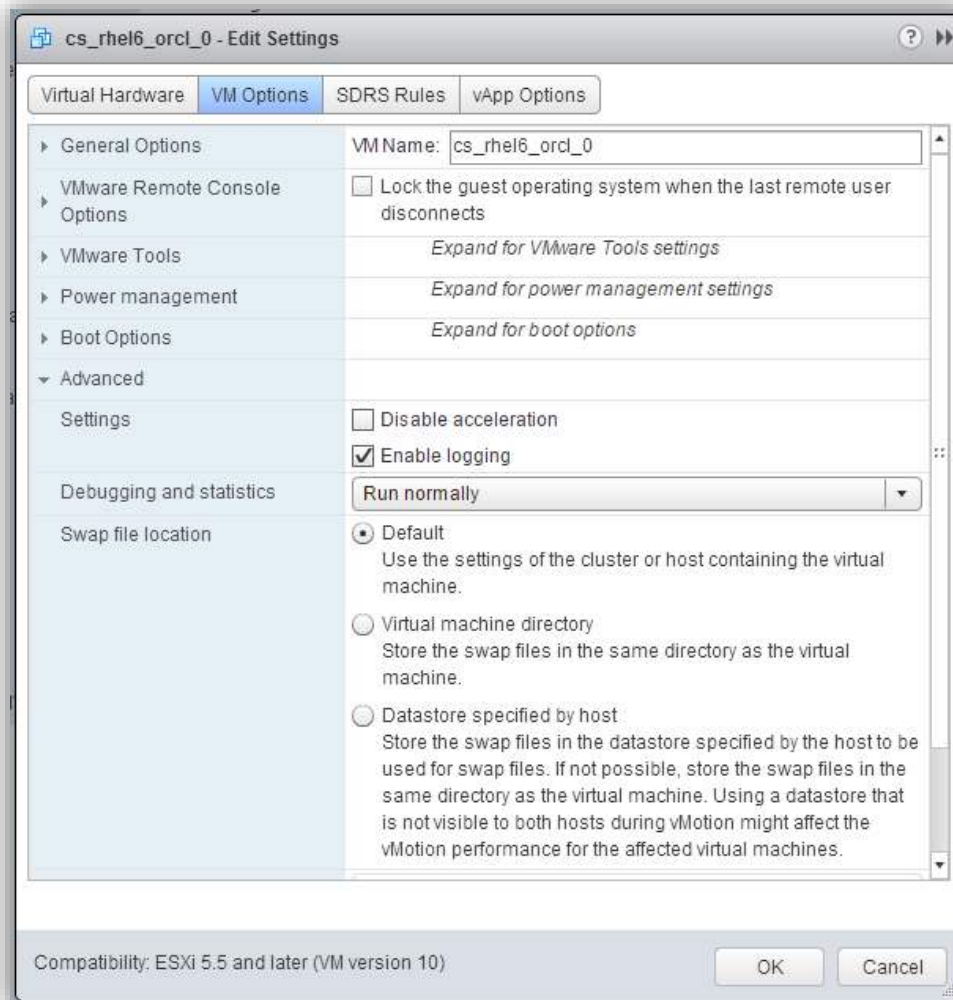


Figure 4 – Location of the VM Swap File

## VMware Snapshots

According to [VMware](#), retaining a VMware snapshot for more than 24-72 hours can create performance problems. We recommend using VMware snapshots sparingly and remember to delete all snapshots as soon as you are finished using them. We strongly recommend that you follow the VMware's best practice guidelines when using VMware snapshots.

On the other hand, Tintri VMstore snapshots can be used as often as desired. They are designed to work transparently and without affecting the performance of individual vDisks or Virtual Machines.

---

*DO: Delete VMware snapshots as soon as you are done with them.*

---

## Linux Queue Scheduler

The Linux I/O queue scheduler supports multiple types of scheduling algorithms. The NOOP I/O scheduler provides increased throughput and reduced I/O latency for flash storage. We recommend changing the I/O queue scheduler type to NOOP for all vDisks (VMDKs) stored on the Tintri VMstore.

The type of I/O scheduler can be viewed with the following commands.

```
$ grep '' /sys/block/sd*/queue/scheduler
/sys/block/sda/queue/scheduler: noop anticipatory deadline [cfq]
```

To change the type of I/O queue scheduler to NOOP, enter the following commands.

```
$ echo noop > /sys/block/sda/queue/scheduler
$ grep '' /sys/block/sd*/queue/scheduler
/sys/block/sda/queue/scheduler: [noop] anticipatory deadline cfq
```

---

*DO: Set the Linux I/O queue scheduler to NOOP to increase throughput and decrease latency.*

---

## Linux Queue Size

For databases that issue large amounts of write I/O, we recommend increasing the Linux queue size. A large queue size provides the opportunity for Linux to merge multiple I/O requests and improve write ordering. We recommend confirming that `nr_request` and `max_sectors_kb` are set to 128 and 512 respectively.

```
$ grep '' /sys/block/sd*/device/model
$ grep '' /sys/block/sd*/queue/nr_requests
$ grep '' /sys/block/sd*/queue/max_sectors_kb
```

In our testing with various Linux releases, `nr_request` and `max_sectors_kb` are usually already set large enough, to 128 and 512 respectively. If it is necessary to set these values, execute the following lines:

```
echo 128 > /sys/block/$dev/queue/nr_requests
echo 512 > /sys/block/$dev/queue/max_sectors_kb
```

---

*DO: Increase the Linux Queue size if the database requires large amounts of write I/O.*

---

## VM Guest Linux Device Readahead Settings

Check to see if the performance of your database would benefit from increasing or decreasing the Linux block device readahead settings. An OLTP database would benefit from a small readahead setting, while a DSS database would benefit from a larger readahead setting. For example, Tintri testing with sequential read workloads shows good results with a readahead setting of 10000 sectors (5MB).

The following commands generates a list of file systems and devices and increases the device readahead to 10000 sectors (5MB).

```
mount
blockdev --report
blockdev --getra /dev/sd?
blockdev --setra 10000 /dev/sd?
blockdev --getra /dev/sd?
```

On systems using LVM, you need to increase the read-ahead on the corresponding LVM device. The following commands sets the readahead values for all LVM devices (vDisks) in the virtual machine.

```
mount
blockdev --report
blockdev --getra /dev/mapper/XXX
blockdev --setra 10000 /dev/mapper/XXX
blockdev --getra /dev/mapper/XXX
```

---

**DO:** Check whether database performance can benefit by changing the Linux device readahead settings.

---

## VM Guest Testing the Storage Throughput with the “dd” Command

Having made the recommended VM Guest Performance changes, care must be taken when using the “dd” command to measure throughput. The default block size is 512 bytes and will be noticeably slower than with larger block sizes. We recommend using the use of the following “dd” options when testing throughput.

```
dd if=/dev/zero of=testfile bs=64k
dd if=testfile of=/dev/null bs=64k
```

---

**DO:** Use the “dd” command with a block size of 64k.

---



## Deploying Oracle with Tintri

Deploying a virtualized Oracle database with Tintri is very similar to deploying Oracle on a physical server. An Oracle DBA who designs, implements, and maintains Oracle databases running on physical hardware will use the same set of tools and knowledge when deploying Oracle in a virtual machine with a Tintri VMstore.

The performance resources of a single Tintri VMstore can effectively support one database or hundreds of databases. Tintri's FlashFirst technology delivers high levels of throughput and consistent sub-millisecond latencies to every VM and VMDK. The QoS technology effectively isolates the performance of one VM from all other VMs on the system to ensure that a "noisy neighbor" has no effect on your VM.

The VMstore monitors every I/O request at the vDisk level to track IOPS, MB/sec and latency at the host, network and storage levels. Performance graphs allow you to view the performance and headroom of the appliance, and can be used to drill down into performance details for individual VMs and VMDKs. This performance visibility allows you to confidently run multiple databases on the Tintri VMstore.

When installing Oracle in a VM with a Tintri VMstore, install Oracle exactly as you would for a physical environment and choose either the Oracle ASM or the Oracle OFA model for managing the database files.

---

*DO: Use Tintri VMstore tools to monitor the performance of Oracle database VMs.*

*DO: Choose either the Oracle ASM or Oracle OFA model for managing the database files.*

---

## Types of Oracle Database Deployments Supported by the Tintri VMstore

Tintri can be used to support many different types of Oracle databases, each with different I/O patterns. The Tintri VMstore supports all of these deployment options.

- Multiple OLTP databases with high IOPS requirements can be deployed on a single VMstore.
- Databases with heavy read/write operations that run around the clock will benefit from Tintri's ability to monitor performance and adapt to changing workload profiles.
- Development and test environments can leverage Tintri's built-in cloning technology to quickly clone and run hundreds of test and development instances.
- Using Oracle replication tools, Oracle production databases can be replicated from physical server environments to a virtualized environment. Oracle running in the virtual machine can be used as the failover system; used to run database reports; and rapidly cloned for QA, test and development tasks.
- Databases that support web applications can be virtualized along with the associated suite of application and web servers. As the workload on the application layer increases, various portions of the environment can be cloned and used to extend the performance of the system.
- Decision Support databases and Data Warehouses with storage requirements that exceed a single Tintri VMstore can be deployed across multiple VMstores.



## Deployment Architecture Options

The Tintri VMstore is designed to support a mixed workload of hundreds of VMs, each with a unique IO configuration and performance profile. Whether a simple, or a sophisticated, vDisk architecture is used with a Virtual Machine, Tintri’s all flash level performance delivers consistent flash-based performance to each VM and vDisk.

For production databases, Tintri recommends that individual vDisks be used to host various database components and individual tablespaces. The size of the vDisk has no effect its performance, so vDisks can be as small or as large (up to 16TiB, the largest size of a virtual disk currently supported on Tintri VMstore) as is required. When separate vDisks are used, latency and performance graphs can be displayed for each vDisk. The Tintri performance visualization graphs make it simple to find the component of the database that is involved in the performance issue.

For databases that are being architected to provide the fastest response times and the highest throughputs will benefit from the use of multiple I/O paths and multiple vDisks. For these databases VMware recommends the use of multiple PVSCSI adapters along with multiple vDisks.

Architecture	PVSCSI	VMDK	Pros and Cons
Single VMDK	One	One	<p><b>Cons:</b> Tintri’s performance insight is captured and viewed on a VMDK basis. Using a single VMDK means limited insight into the performance of the individual Oracle datafiles.</p> <p><b>Cons:</b> According to VMware, the use of a single PVSCSI adapter can create a storage I/O bottleneck in the Virtual Machine.</p>
Multiple VMDK	One	Multiple	<p><b>Pros:</b> Multiple VMDK’s provide performance insights into the individual datafiles and components used by the database.</p> <p><b>Cons:</b> According to VMware, the use of a single PVSCSI adapter can create a storage I/O bottleneck in the Virtual Machine.</p>
Multiple VMDK Multiple PVSCSI  <b>BEST PRACTICE</b>	Multiple	Multiple	<p><b>Pros:</b> Multiple VMDKs provide performance insight into individual segments of the database.</p> <p><b>Pros:</b> The use of multiple PVSCSI adapters eliminates storage I/O bottlenecks in high performance databases.</p> <p><b>Pros:</b> This is the most effective deployment architecture for Oracle on the Tintri VMstore.</p>

Figure 5 – Pros and Cons of Three Virtual Machine Architectures

**DO:** The recommended VM deployment architecture for Oracle with Tintri VMstore is to use multiple VMDK files with multiple PVSCSI adapters.

## VMware PVSCSI Adapter

VMware created the Paravirtual SCSI (PVSCSI) adapter to support I/O intensive guest applications, such as Oracle databases. The PVSCSI driver leverages ESXi kernel-level storage stack optimizations to dramatically improve storage I/O performance by creating a separate SCSI queue and supporting the execution of parallel storage I/O operations inside the guest operating system. Up to 4 PVSCSI adapters can be used in a single VM, and 15 disks can be attached to each PVSCSI adapter for a total of 60 virtual disks per virtual machine.

VMware recommends using a PVSCSI adapter for the root of the file system (Linux root) and additional PVSCSI adapters for the Oracle database files. VMware also suggests using a separate PVSCSI controller to support the Redo and Archive Log I/O.

Refer to VMware KB article [1010398](#), **Configuring disks to use VMware Paravirtual SCSI (PVSCSI) adapters**.

---

*DO: Use the Paravirtual SCSI adapter (PVSCSI) when deploying Oracle databases on Tintri VMstore systems*

---

## Virtual Disk Provisioning Mode

When creating a New Hard Disk in vSphere you have three choices of Disk Provisioning:

1. Thick provision lazy zeroed
2. Thick provision eager zeroed
3. Thin provision

The Tintri VMstore has been designed to work most efficiently when the **thin provisioned disk** option is selected. Thin provisioned disks consume space on the VMstore only when data is written to by the application. Tintri deduplication and compression technologies work with thin provisioned disks and reduce the actual amount of data being stored by the application.

The Virtual Machine will report that the full amount of provisioned space is available, but with thin provisioning, deduplication, and compression technologies, the actual amount of storage consumed by the VM will be considerably less.

VMware has tested Oracle database workloads with thin-provisioned disks and found that the thin provisioned disks have the same level of performance as thick provisioned disks, even under I/O intensive workloads. For more information on performance with Thin provisioned disks, refer to the VMware [white paper](#), **Performance Study of VMware vStorage Thin Provisioning**.

By using thin provisioning, and keeping in mind that space is only consumed as Oracle writes data to the database files, you will observe the following.

- Redo logs – created at run time and used immediately, generate their own full provisioning.
  - Archive Log Files – written out completely and thus generate their own full provisioning.
  - Datafiles – storage is consumed as data is written to the tablespaces.
- 

*DO: Use thin-provisioned virtual disks.*

*DO: Monitor disk capacity reserves of the Tintri VMstore and migrate VMs if additional space is needed.*

---

## vDisk Sizing

Managing a VM with multiple vDisks can get confusing during VMware vCenter operations such as adding, growing, and removing vDisks. While some operations are relatively harmless, such as extending the wrong vDisk to increase capacity, other operations can be catastrophic, such as deleting the wrong vDisk. One way to avoid confusion is to size vDisks with noticeably different sizes.

For example, if you were going to make two vDisks of 100GB, make one of the disks 105GB. This small difference will make it easier to differentiate the two disks with RHEL OS commands, and throughout the vSphere Web Client, and the Tintri VMstore Dashboard.

Note that this recommendation applies to OFA style database deployments. When deploying ASM Oracle recommends using vDisks of exactly the same size and Tintri supports that recommendation.

---

**DO:** Use differently sized vDisks to avoid confusion when managing multiple virtual disks.

---

## Using Oracle OFA with the Tintri VMstore

Oracle Optimal Flexible Architecture (OFA) provides a consistent naming convention that makes it simple to organize and identify database data files. DBAs can choose to place the individual data files together into a single vDisk, or to place individual files onto separate vDisks to provide the highest level of control and performance insight into the database.

Whether you choose a simple or an advanced layout for your database, the Tintri VMstore is designed to deliver consistent performance to each VM and vDisk and to isolate that performance from other VMs on the system. Tintri VMstore monitors each virtual disk to provide a graphical display of the performance and latency at the hypervisor, network and storage level. This performance monitoring enables you to quickly identify performance bottlenecks, should they occur.

---

*DO: Create the OFA disk layout that supports your database manageability and I/O requirements*

---

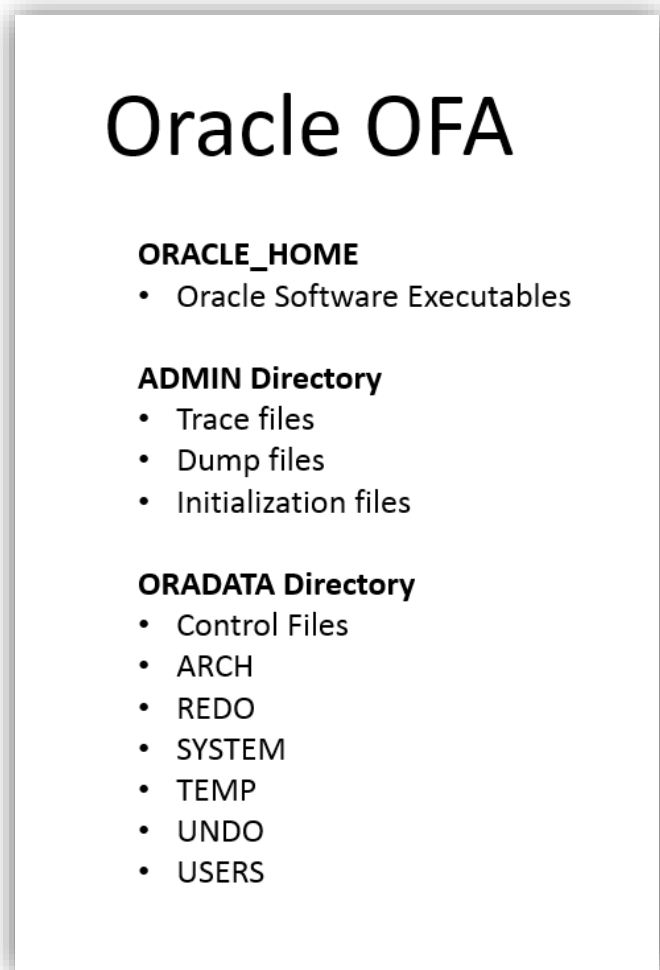


Figure 6 – OFA Data Placement Conventions

## Deploying Oracle OFA with a Single VMDK

This example database layout uses one PVSCSI adapter and one vDisk. The database name is “orcl”.




Figure 7 – Logical Diagram for the Single VMDK and Single PVSCSI Adapter Architecture

<b>vDISK</b>	<b>PVSCSI ID</b>	<b>MOUNT POINT and PATH</b>	<b>DESCRIPTION</b>
Hard disk 1	SCSI(0:0)	/	File system root
Hard disk 1	SCSI(0:0)	/dev/sda2	Linux swap
Hard disk 1	SCSI(0:0)	/u01/app/oracle/	ORACLE_BASE
Hard disk 1	SCSI(0:0)	/u01/app/oracle/product/11.2.0/dbhome_1/	ORACLE_HOME
Hard disk 1	SCSI(0:0)	/u01/app/oracle/oradata/orcl/control01 /u02/app/oracle/oradata/orcl/control02	Control files
Hard disk 1	SCSI(0:0)	/u01/app/oracle/admin/oracle/arch/	Archived Logs
Hard disk 1	SCSI(0:0)	/u01/app/oracle/oradata/orcl/logs/	Redo Logs
Hard disk 1	SCSI(0:0)	/u02/app/oracle/oradata/orcl/logs/	
Hard disk 1	SCSI(0:0)	/u01/app/oracle/oradata/orcl/	SYSTEM
Hard disk 1	SCSI(0:0)	/u01/app/oracle/oradata/orcl/	SYSAUX
Hard disk 1	SCSI(0:0)	/u01/app/oracle/oradata/orcl/	TEMP
Hard disk 1	SCSI(0:0)	/u01/app/oracle/oradata/orcl/	USERS
Hard disk 1	SCSI(0:0)	/u01/app/oracle/oradata/orcl/	UNDO
Hard disk 1	SCSI(0:0)	/u02/app/oracle/oradata/orcl/	Data and Indexes

**Figure 8 – Deployment Details for the Single VMDK and Single PVSCSI Adapter Deployment Architecture**

## Deploying Oracle OFA with one PVSCSI Adapter and Multiple VMDKs

This example database architecture uses multiple vDisks. The Tintri VMstore provides detailed performance graphs for each VMDK (vDisk). The database name is “orcl”.




Figure 9 – Logical Diagram for the Multiple VMDK and Single PVSCSI Adapter Architecture

<b>vDISK</b>	<b>PVSCSI ID</b>	<b>MOUNT POINT and PATH</b>	<b>DESCRIPTION</b>
Hard disk 1	SCSI(0:0)	/	File system root
Hard disk 1	SCSI(0:0)	/dev/sda2	Linux swap
Hard disk 2	SCSI(0:1)	/u01/app/oracle/	ORACLE_BASE
Hard disk 2	SCSI(0:1)	/u01/app/oracle/product/11.2.0/dbhome_1/	ORACLE_HOME
Hard disk 2	SCSI(0:1)	/u01/app/oracle/oradata/orcl/control01	Control files
Hard disk 3	SCSI(0:2)	/u02/app/oracle/oradata/orcl/control02	
Hard disk 4	SCSI(0:3)	/u03/app/oracle/oradata/orcl/control03	
Hard disk 3	SCSI(0:2)	/u02/app/oracle/admin/oracle/arch/	Archived Logs
Hard disk 2	SCSI(0:1)	/u01/app/oracle/oradata/orcl/logs/	Redo Logs
Hard disk 3	SCSI(0:2)	/u02/app/oracle/oradata/orcl/logs/	
Hard disk 4	SCSI(0:3)	/u03/app/oracle/oradata/orcl/logs/	
Hard disk 4	SCSI(0:3)	/u03/app/oracle/oradata/orcl/	SYSTEM
Hard disk 5	SCSI(0:4)	/u04/app/oracle/oradata/orcl/	TEMP + UNDO
Hard disk 6	SCSI(0:5)	/u05/app/oracle/oradata/orcl/	USERS
Hard disk 7	SCSI(0:6)	/u06/app/oracle/oradata/orcl/	Data
Hard disk 8	SCSI(0:7)	/u07/app/oracle/oradata/orcl/	Indexes

**Figure 10 –Deployment Details for the Multiple VMDK and Single PVSCSI Adapter Architecture**



## Deploying Oracle OFA with Multiple PVSCSI Adapters and Multiple VMDKs

This example database architecture uses multiple PVSCSI adapters and multiple vDisks. The PVSCSI adapters provide multiple I/O paths for the database. The Tintri VMstore provides detailed performance graphs for each VMDK (vDisk). The database name is “orcl”.




Figure 11 – Logical Diagram for the Multiple VMDK and Multiple PVSCSI Adapter Architecture

vDISK	PVSCSI ID	MOUNT POINT and PATH	DESCRIPTION
Hard disk 1	SCSI(0:0)	/	File system root
Hard disk 1	SCSI(0:0)	/dev/sda2	Linux swap
Hard disk 2	SCSI(0:1)	/u01/app/oracle/	ORACLE_BASE
Hard disk 2	SCSI(0:1)	/u01/app/oracle/product/11.2.0/dbhome_1/	ORACLE_HOME
Hard disk 2	SCSI(0:1)	/u01/app/oracle/oradata/orcl/control01	Control files
Hard disk 4	SCSI(1:0)	/u03/app/oracle/oradata/orcl/control02	
Hard disk 7	SCSI(2:0)	/u06/app/oracle/oradata/orcl/control03	
Hard disk 9	SCSI(3:0)	/u08/app/oracle/oradata/orcl/control04	
Hard disk 3	SCSI(0:2)	/u02/app/oracle/admin/oracle/arch/	Archived Logs
Hard disk 2	SCSI(0:1)	/u01/app/oracle/oradata/orcl/logs/	Redo Logs
Hard disk 4	SCSI(1:0)	/u03/app/oracle/oradata/orcl/logs/	
Hard disk 7	SCSI(2:0)	/u06/app/oracle/oradata/orcl/logs/	
Hard disk 4	SCSI(1:0)	/u03/app/oracle/oradata/orcl/	SYSTEM + SYSAUX
Hard disk 5	SCSI(1:1)	/u04/app/oracle/oradata/orcl/	TEMP + UNDO
Hard disk 6	SCSI(1:2)	/u05/app/oracle/oradata/orcl/	USERS
Hard disk 7	SCSI(2:0)	/u06/app/oracle/oradata/orcl/	Data
Hard disk 8	SCSI(2:1)	/u07/app/oracle/oradata/orcl/	Data
Hard disk 9	SCSI(3:0)	/u08/app/oracle/oradata/orcl/	Data
Hard disk 10	SCSI(3:1)	/u09/app/oracle/oradata/orcl/	Indexes

Figure 12 – Deployment Details for the Multiple VMDK and a Multiple PVSCSI Adapters Architecture

## Deploying Oracle ASM with the Tintri VMstore

Oracle Automatic Storage Management (ASM) is a volume manager and file system that ships with Oracle. ASM stores data in a collection of disks that it calls a disk group. Disks can be added or removed from the disk groups and ASM takes care of redistributing the database files.

To use ASM for database storage, you must create one or more ASM disk groups. Oracle maintains the following recommendations for ASM disks.

- ASM disk groups are populated with disks of equal size and I/O characteristics.
- A minimum of two ASM disk groups are recommended—one for log files, which are sequential in nature, and another for datafiles, which are random in nature.
- The Tintri VMstore provides redundancy and RAID protection. Therefore, when creating the disk group the redundancy option should be set to EXTERNAL REDUNDANCY.

Using ASM when deploying Oracle on Tintri is simple. Simply create multiple virtual disks in the virtual machine and then create the ASM groups. The following recommendations apply:

- Create multiple virtual disks on the VM and ensure the VMDKs are located on the Tintri VMstore.
- Separate types of disk I/O by isolating common I/O patterns to separate PVSCSI adapters.
- Create a single partition on each disk.
- Use a redundancy level of EXTERNAL REDUNDANCY with the CREATE DISKGROUP command.
- Using the external redundancy option removes the disk group redundancy overhead from ASM, thus increasing the CPU cycles available to the Oracle database.
- The Tintri VMstore performance monitoring enables you to quickly identify performance bottlenecks with individual ASM disks.

The CREATE DISKGROUP SQL statement is used to create disk groups. When creating a disk group, you assign a unique name to the disk group, specify the redundancy level for the disk group, and specify the disks that are to be formatted. Disks must have a partition table and Oracle recommends creating exactly one partition on each disk. Since the Tintri VMstore provides RAID protection and redundancy you will be able to leverage the EXTERNAL REDUNDANCY option with the create diskgroup command.

The CREATE DISKGROUP SQL statement also requires us to specify an ASM allocation unit. The ASM Allocation Unit is a fundamental unit of storage and management in any ASM disk group. The default ASM AU size is 1MB which is appropriate for OLTP databases, but a larger AU size can improve the performance of the sequential read I/O that data warehouses require.

More information on Performance and Scalability Considerations for ASM Disk Groups is found [here](#) in the **Oracle ASM Administrator's guide**.

# Virtual Machine




Figure 13 – Logical Diagram for an Oracle ASM Deployment Architecture that uses Multiple VMDKs and Multiple PVSCSI Adapters

VM DISK	PVSCSI ID	MOUNT POINT and PATH	DESCRIPTION
Hard disk 1	SCSI(0:0)	/	LINUX root
Hard disk 1	SCSI(0:0)	/dev/sda2	swap
Hard disk 2	SCSI(0:1)	/u01/app/oracle/	ORACLE_BASE
Hard disk 2	SCSI(0:1)	/u01/app/oracle/product/11.2.0/dbhome_1/	ORACLE_HOME
Hard disk 3	SCSI(1:0)	/dev/sdc	“data” disk group
Hard disk 4	SCSI(1:1)	/dev/sdd	“data” disk group
Hard disk 5	SCSI(2:0)	/dev/sde	“data” disk group
Hard disk 6	SCSI(2:1)	/dev/sdf	”data” disk group
Hard disk 7	SCSI(3:0)	/dev/sdg	“logs” disk group
Hard disk 8	SCSI(3:1)	/dev/sdh	“logs” disk group

**Figure 14 – Deployment Details for an Oracle ASM Deployment Architecture that uses Multiple VMDKs and Multiple PVSCSI Adapters**

## Tintri Deduplication and Compression versus Oracle Database Compression

The Tintri VMstore has built-in deduplication and data compression technologies that reduce the amount of data being stored in flash. By deduplicating and compressing data, Tintri makes the use of solid-state flash storage for Oracle databases truly cost effective. Furthermore, database read operations can realize additional benefit from data compression as a single read operation now returns more data blocks, which results in improved performance for database sequential scan operations.

Oracle 11g and 12c offers compression technology as a licensed feature. Oracle has created multiple types of database compression tools, some of which won't automatically re-compress data after database records have been updated. Maintaining a reasonable compression ratio will require manual intervention.

Compare that with Tintri deduplication and compression that is automatically applied to all types of data, including application executables, database data files, and unstructured file data. Allowing the Tintri VMstore deduplicate and compress the database files is more efficient as the VMstore has been purposely designed to accomplish these tasks. Using Oracle "add-on" compression tools consumes CPU cycles on the server which are better used for other database operations.

Note: When Oracle compression is enabled, the Tintri VMstore will not be able to deduplicate or compress the data being stored by the database.

---

*DO: Turn off Oracle compression technologies.*

*DO: Use Tintri VMstore's built-in deduplication and compression technology with Oracle databases to efficiently reduce the amount of data being stored in flash.*

---

## Oracle DB\_BLOCK\_SIZE Initialization Parameter

The DB\_BLOCK\_SIZE is the size (in bytes) of the data block used by the Oracle database. Typical values are 8K for OLTP databases, and 32K for data warehouse databases. Set the value of this parameter to a multiple of the physical block size used by the Tintri VMstore, which is 8K.

---

*DO: Use a DB\_BLOCK\_SIZE set to a multiple of 8K.*

---



## Oracle Database Protection and Backup Options

The Tintri VMstore provides two VM protection tools: SnapVM and ReplicateVM. Tintri SnapVM creates instantaneous snapshot backups, and Tintri ReplicateVM replicates snapshots to another VMstore, providing the necessary second offsite copy of the snapshot backups.

Detailed best practices for Tintri VMstore backups are covered in the [white paper, Backup and Recovery Best Practices with Tintri VMstore](#).

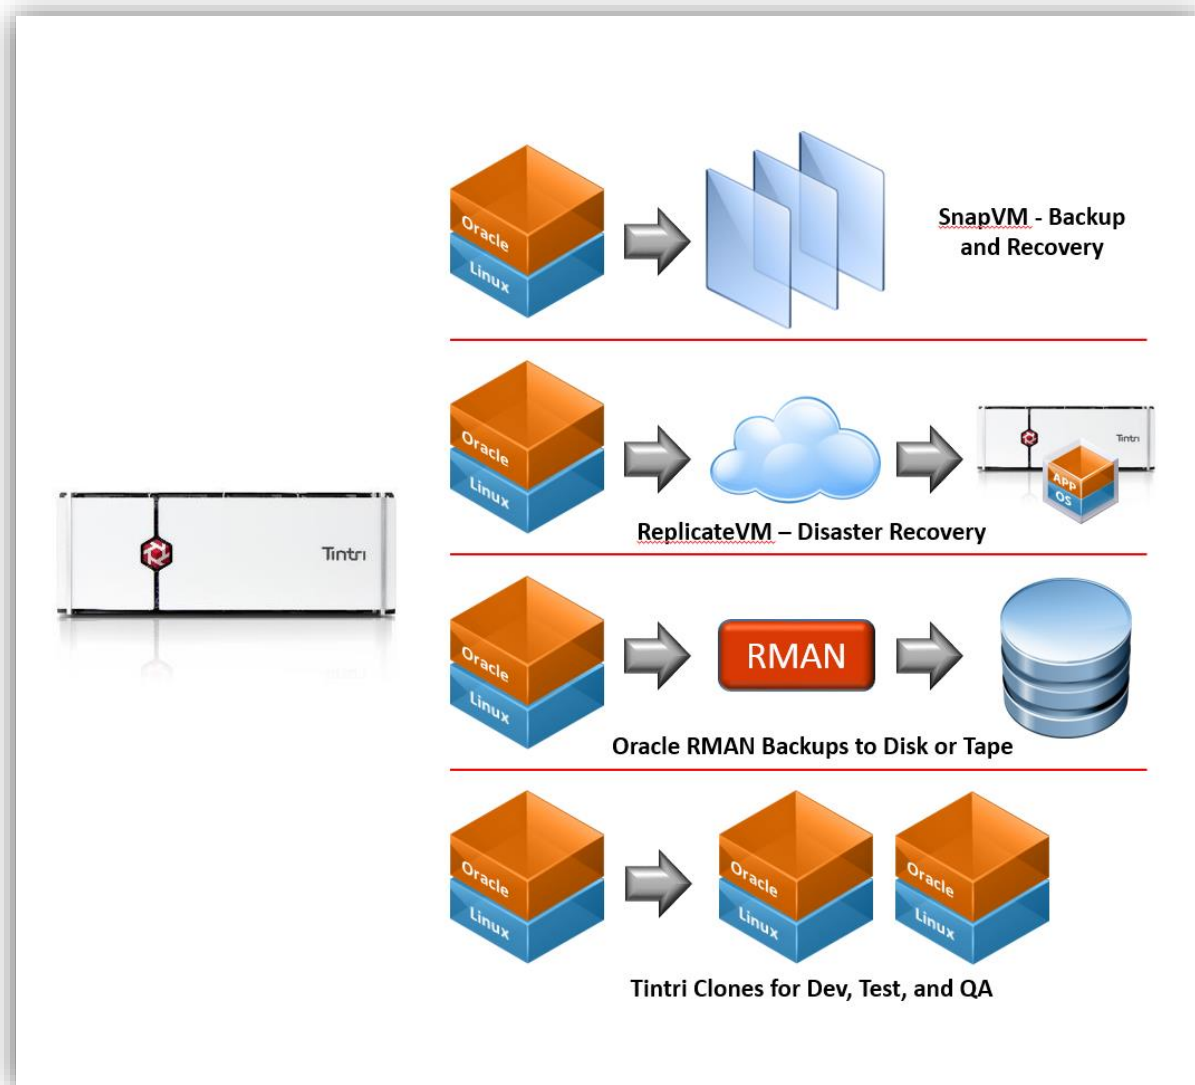


Figure 15 – Tintri Data Protection Options for Oracle Databases

## Using Oracle RMAN for Database Backups

The Oracle Recovery Manager (RMAN) tool is used to backup and restore databases, tablespaces and datafiles to offline disk or tape storage. It can be used to backup databases in their entirety, or to create incremental backups. DBAs have come to rely on its restore features, including the ability to validate backups and to repair block corruptions in the database.

Oracle recommends Recovery Manager (RMAN) and Fast Recovery Area (FRA) as the supported solution for creating and managing Oracle database backups. Tintri fully supports the RMAN tool for creating and managing database backups.

## VMware Snapshots

VMware snapshot technology preserves the state of a virtual machine at a point in time. The snapshot includes the virtual machine's state and the files that make up the virtual machine.

However, a virtual machine snapshot is not a complete copy of the original VMDK disk files; rather, VMware maintains changes to the VMDK in the snapshot. On high-transaction virtual machines, such as database servers, the existence of a snapshot will affect the performance of the virtual machine. For this reason, VMware recommends that no single snapshot exist for more than 24-72 hours and should then be deleted.

More information on VMware snapshots can be found in the following VMware KB articles:

**Understanding virtual machine snapshots in VMware ESXi and ESX** ([KB Article 1015180](#))

**Best practices for virtual machine snapshots in the VMware environment** ([KB Article 1025279](#))

## Tintri SnapVM™ Snapshots

Tintri's SnapVM snapshots provide a simple and powerful on-line backup tool for Oracle VM and databases. The SnapVM snapshot technology is unique in being able to operate on individual Virtual Machines to create point-in-time, read-only copies of individual VMs and the associated VMDKs. Tintri SnapVM snapshots make it easy to backup, clone, restore, and mirror entire Oracle databases.

Tintri's SnapVM snapshot technology is:

- **Granular** - SnapVM snapshots apply to individual VMs, not the entire storage appliance.
- **High performance** - SnapVM snapshots do not affect the performance of the VM.
- **Space efficient** - Storage is consumed only when the VM writes changes its VMDK file(s).
- **Flexible** - Restoring a snapshot doesn't delete newer or older snapshots. All recovery points continue to be available.
- **Automated** – Configured hourly, daily, weekly, monthly, on-demand or called via a script.

Tintri supports two types of SnapVM snapshots: VM-Consistent and Crash Consistent. VM-Consistent snapshots leverage VMware Tools to quiesce the Oracle database and the virtual machine before creating the snapshot. Crash-Consistent snapshots are created instantly and ignore the state of the virtual machine.

When applied to an Oracle database, SnapVM snapshots quickly create point-in-time copies of the entire Oracle VM and associated VMDK files, including all database executables and database datafiles. Consider using the SnapVM snapshots in the following ways:

- Use VM-Consistent snapshots to create a “golden copy” of a production database prior to upgrading Oracle database executables.
- Use VM-Consistent snapshots to create a source snapshot for making clones of the database.



- Use Crash-Consistent snapshots to create on-line point-in-time copies of the database which provide multiple on-line restore options. The SnapVM snapshots incur no performance penalty so they can remain on the storage for long periods of time with no effect on the production database.

*DO: Use Tintri SnapVM™ snapshots to protect the Oracle database and the datafiles.*

## Tintri SnapVM VM-Consistent Snapshots

A VM-Consistent snapshot is taken when the virtual machine has been quiesced. The ability to quiesce a virtual machine is made possible when VMware Tools is installed in the virtual machine. A quiesce operation will flush in-guest memory buffers and quiesce the storage I/O in the virtual machine, making it possible to use a storage-based snapshot to capture the complete state of the virtual machine and all of its data files.

VM-Consistent snapshots can be scheduled or invoked from the Tintri VMstore UI. When the Tintri SnapVM command is initiated, the following actions are executed:

1. A VMware snapshot is taken with the “Quiesce Guest File System” option selected.
2. A Tintri SnapVM snapshot is created.
3. The VMware snapshot is deleted so there is no performance impact on the virtual machine.

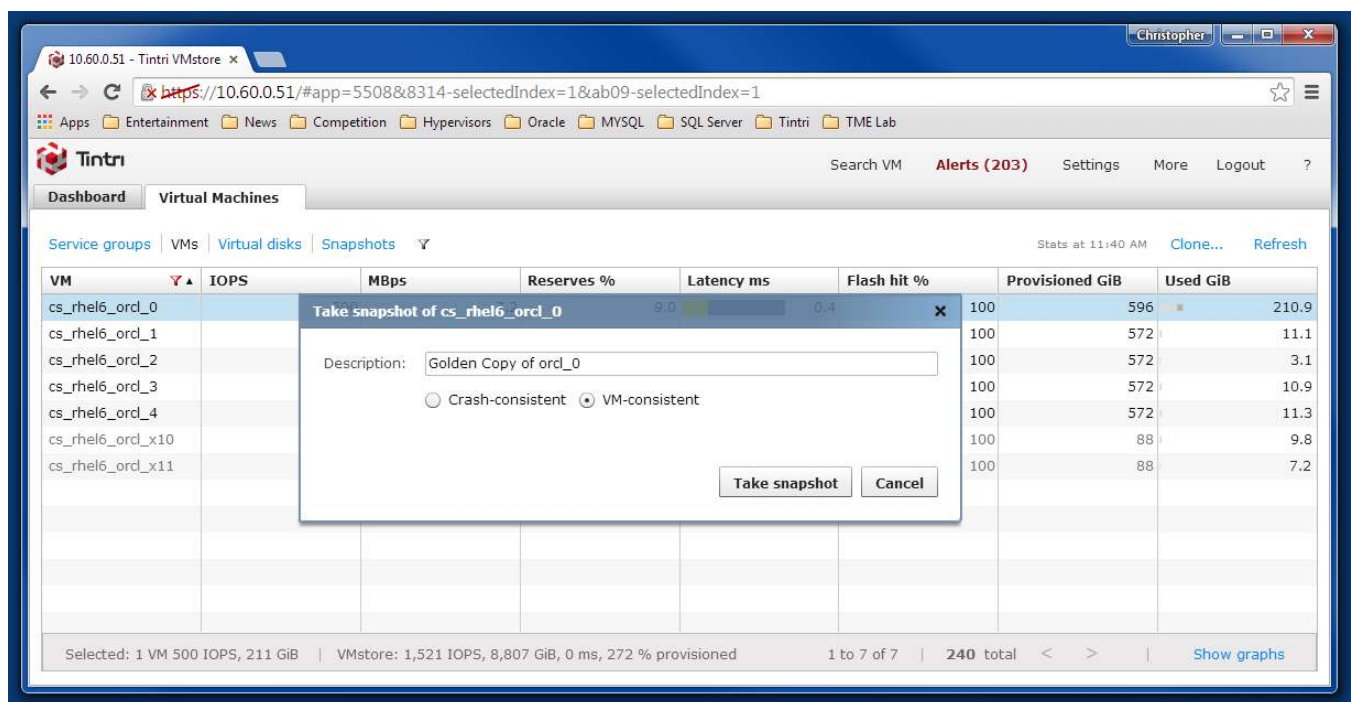


Figure 16 – VMware Snapshot Dialog, Showing the Quiesce Guest File System Option

## VMware Tools Scripts

VMware tools will automatically run scripts before and after freezing I/O in the guest. This capability can be used to execute scripts that place the Oracle tablespaces into hot backup mode. According to the VMware backup guide, freeze and thaw scripts have pre-defined names, and are located in the following directory.

```
/usr/sbin/pre-freeze-script  
/usr/sbin/post-thaw-script
```

To test this capability with sample pre-freeze-script and post-thaw-script on Linux, do the following:

### 1. Create sample pre-freeze and post-freeze scripts

```
$ echo "echo \"pre script\" >> /tmp/freeze.log" > /usr/sbin/pre-freeze-script  
$ echo "echo \"post script\" >> /tmp/freeze.log" > /usr/sbin/post-thaw-script
```

### 2. Change the owner to root and add executable permissions to the scripts

```
$ chown root.root /usr/sbin/pre-freeze-script /usr/sbin/post-thaw-script  
$ chmod 0700 /usr/sbin/pre-freeze-script /usr/sbin/post-thaw-script
```

### 3. Using the vCenter Web Client, right-click the VM and select the "Take Snapshot" action.

4. In the "Take VM Snapshot for guest-name" dialog box, make sure that the "Snapshot the virtual machine's memory" option is NOT checked, and then check the "Quiesce guest file system" option.

5. Press OK and then review the contents of the log file to confirm that the scripts executed.

---

*DO: Install and update to the latest version of Use VMware Tools.*

*DO: Use pre-freeze and post-thaw scripts with VM-Consistent snapshots.*

---

## Tintri SnapVM Crash-Consistent Snapshots

When Tintri SnapVM Crash-Consistent snapshots are used, no attempt is made to flush in-memory I/O buffers for guest applications or guest operating systems. When this type of snapshot is used with an active Oracle database, the resulting snapshot will contain an image of the database that is equivalent to an Oracle database that has crashed.

Oracle databases are designed to be recoverable from a crashed state, and experience shows that SnapVM Crash-Consistent will be completely recoverable and usable. This value of this type of snapshot comes from the fact that it executes very quickly and does not disrupt the I/O to the production database. It executes quickly because it does not freeze the I/O on the virtual machine, nor does it execute the pre-freeze or post-freeze scripts.

Because they execute so quickly, DBAs can take advantage of its non-disruptive nature to create multiple snapshot copies of the database each day using a round-robin schedule. These snapshots provide an online safety net and provide multiple on-line recovery points for the production database. Plus, these crash-consistent snapshots of the database can be cloned to create copies of the database for use as reporting, testing, or development instances.

## Tintri ReplicateVM for VM-Level Replication

Tintri's ReplicateVM performs WAN-efficient replication of a VM from one VMstore to another. If a VM is configured for replication, any snapshots created by that VM will be automatically replicated to the destination VMstore. Only changed blocks between snapshots are replicated after deduplication and compression reducing the WAN bandwidth utilization by up to 95%. A replicated VM is called a synthetic VM on the destination VMstore and can be materialized into a virtual machine from any of the snapshots available.

When ReplicateVM is used to replicate a production Oracle database, the result is a powerful disaster recovery solution that is simple to setup. Tintri's replication technology allows you to view and use any of the VM-Consistent and Crash-Consistent snapshots at the remote site. Clones of the replicated database can be used as reporting, testing and development instances at the remote site.

Information on ReplicateVM is found at: <http://www.tintri.com/products/replicatevm>

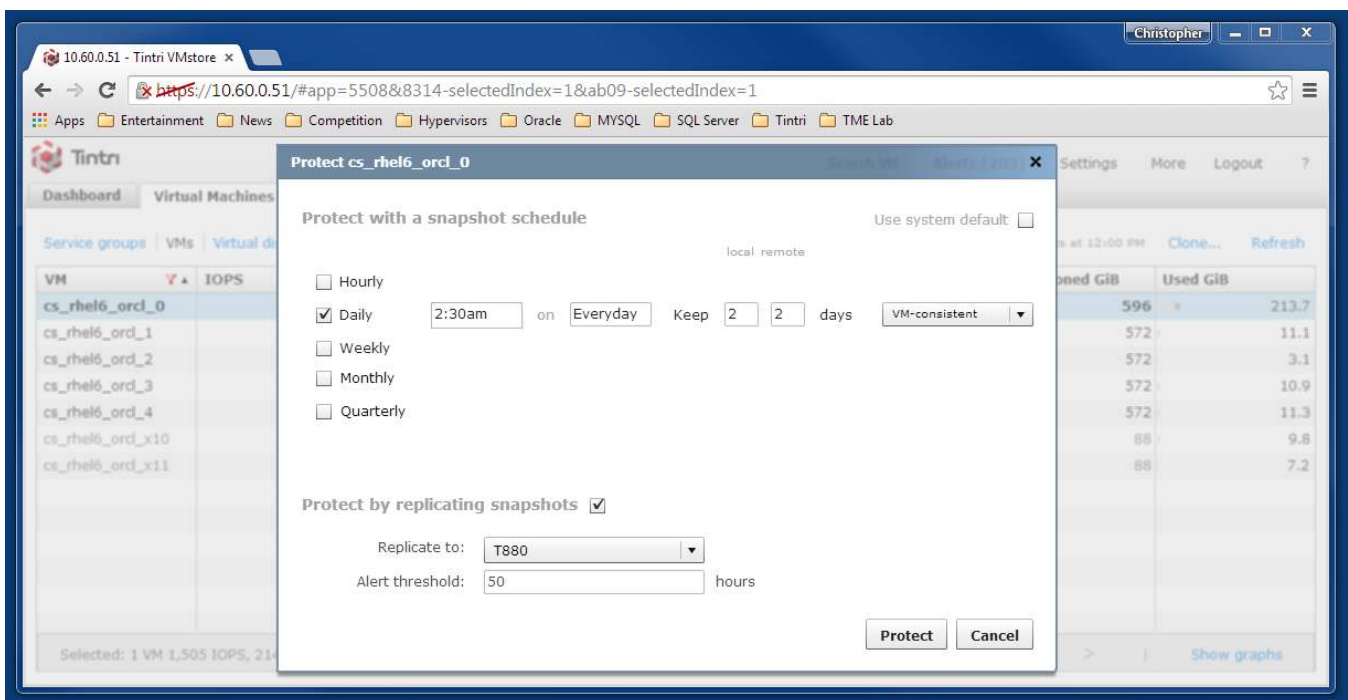


Figure 17 – The Tintri VM Protection Dialog Box

## Performance Analysis

The Tintri VMstore is the first storage product designed to support virtual machines (VMs) without forcing the administrator to deal with low-level storage details. A Tintri VMstore T880 has the ability to support up to 3,500 VMs and 10,000 virtual disks (VMDKs). With this level of VM performance density it is critically important to be able to view both the overall performance of the system and to be able to drill down into the performance details of individual Virtual Machines.

Tintri's goal is to simplify the deployment and management of virtual machines and to that end the VMstore provides several ways to view the performance of the system.

1. The Performance Dashboard provides a graphical view of the overall performance of the system, including IOPS, Throughput, Latency, Performance Reserves, and Physical Space consumed.
2. The Virtual Machines (VMs) tab provides performance details for individual VMs, including IOPS, MBps, Reserves, Latency, and much more.
3. The Virtual Machines (Virtual Disks) tab provides performance details on the VMDKs in each VM.
4. The Virtual Machines (Snapshots) tab provides space and creation details for individual snapshots, including the Source VM, Created Date, Change MB, Cloned Count, and Hypervisor Type.

The system gathers performance statistics every 10 seconds, physical space information every 10 minutes, and keeps the data for seven days. When you first load a graph, statistics from the latest collection are displayed, thereafter the statistics are captured every 10 seconds.

The Virtual Machines tab provides the most granular view of resources and entities on the VMstore. By reviewing the VM and VMDK performance data provided by the Virtual Machines tab, you can easily access latency data that will help you pinpoint the source of performance issues experience by Oracle running in a VM.

---

*DO: Use the VMstore Performance Dashboard to view the performance of the Oracle VM and troubleshoot latency breakdown across the infrastructure at a virtual disk and VM level.*

---

### Tintri VMstore Performance Dashboard

The Dashboard is designed to help you draw quick conclusions about your VMstore's health, identify problems and help you make informed resource management decisions. The system-wide performance graphs include a Performance Reserves gauge, a Physical Space gauge, and individual counters for IOPS, Throughput, Latency, and Flash Hit Rate.

#### VMstore Performance counters

The VMstore performance counters, displayed at the top of the VMstore Dashboard, provide an average rate of operations for the system. Clicking on an individual counter will open a graph of real-time and historic trends for that performance counter. Real-time counters are based on 10-second averages, and historical data is averaged over 10-minute intervals.

#### The performance reserves gauge

The performance reserves gauge shows the headroom available on the VMstore, expressed as a percentage of all performance resources available in the VMstore. By design, the VMstore automatically allocates reserves to a virtual disk based on its observed I/O characteristics. Such self-tuning ensures that all VMs get the performance resources needed to maintain optimum performance.

## The physical space capacity gauge

The space capacity gauge breaks down the types of data being stored in the VMstore. Live data is an accounting of the data consumed by individual VMs. The Snapshots portion of the gauge indicates the space consumed by Tintri snapshots. The Other section of the gauge shows space consumed by descriptors, logs, configuration files, and VMs that are not included in the vCenter inventory. Provisioned calculates the space you have promised to the VMs, and is expressed as a percentage of the system total.

## The “days till full” counter

The VMstore estimates the number of days till full, based on space usage trends.

## Performance reserves changers

Performance reserves changers lists the VMs that are experiencing the largest change in performance for the last week.

## Space changers

Space changers are VMs that are experiencing the largest change in space consumed for the last week.

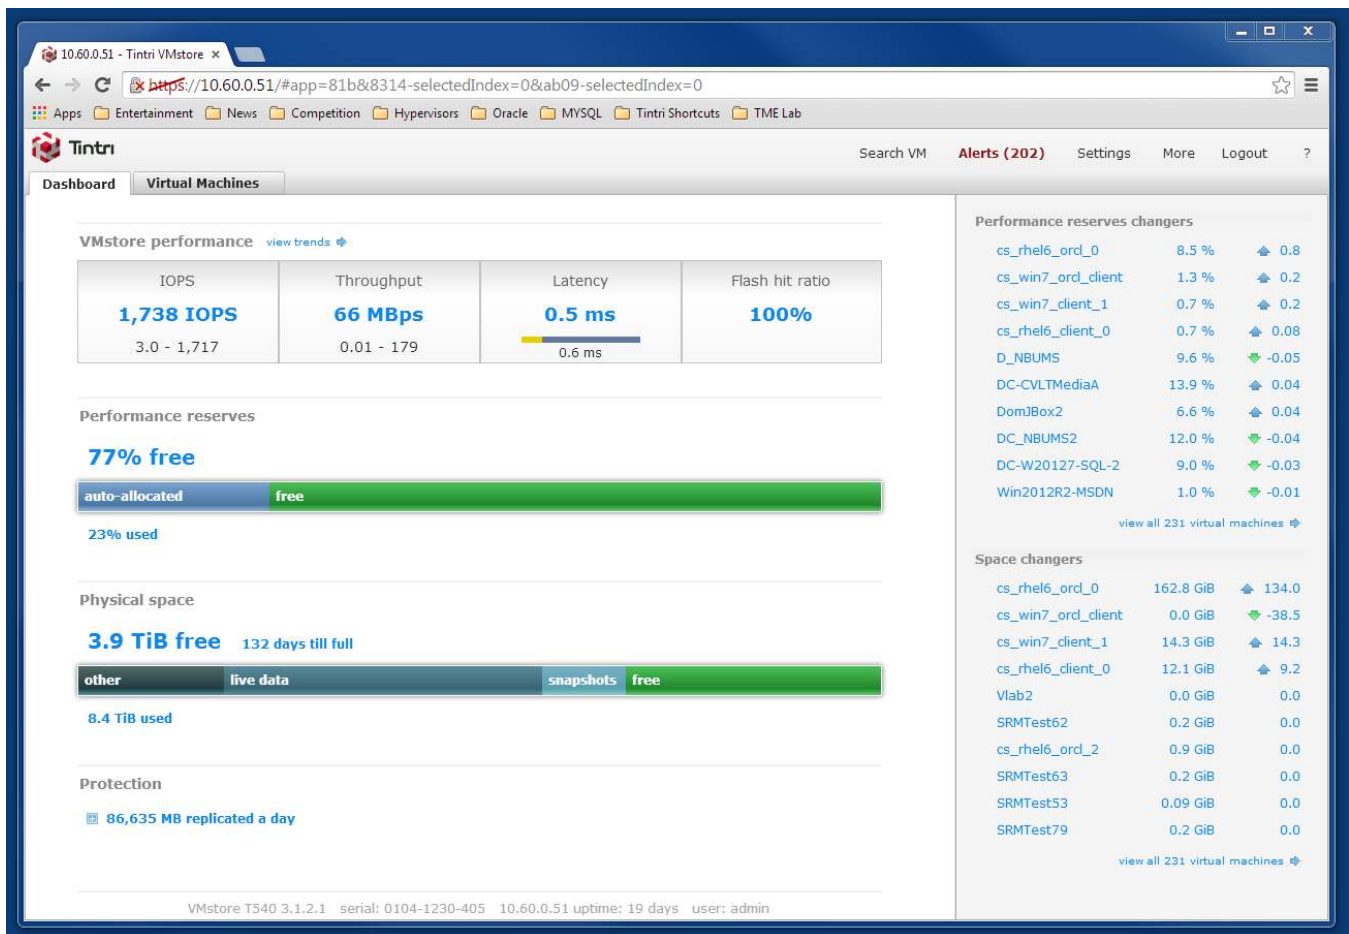


Figure 18 – Tintri VMstore Performance Dashboard

## Virtual Machines Tab

The Virtual Machines tab provides the most granular view of resources and entities on the VMstore. By reviewing VM and Virtual Disks performance data provided by the Tintri VMstore you can easily pinpoint the source of VM performance issues. For example, the Virtual Machines Tab displays performance of individual VMs and Virtual Disks and includes counters for IOPS, MBps, Performance Reserves, Latency, and Flash Hit Rates. These counters are useful to understanding the performance profile of individual VMs on the Tintri VMstore.

When you employ a VM deployment architecture with multiple virtual disks, each being used to store a separate tablespace, you will be able to correlate the performance counters with the performance of a specific tablespace in the database.

### IOPS Graph

The IOPS Graph displays the read and write IOPS served by the specific VM or VMDK.

### MBps Graph

The MBps Graph displays the read and write IOPS served by the specific VM or VMDK.

### Latency Graph

The Latency counter is of particular interest to anyone running an Oracle database as it offers an end-to-end breakdown of the source of latency for individual VMs and vDisks (VMDKs).

**Host** - If you see a lot of green in the latency breakdown, the ESX host may be overloaded. You may need to reduce the load on the host or tune the allocated host CPU or memory for the VM. You can examine the Host graph of %Ready and %Swap-wait to help you with the tuning. VMstore collects these statistics. If %Ready is high, the VM benefits from additional CPU allocation. If %Swap-wait is high, the VM needs more memory.

**Network** - If you see a lot of yellow in the latency breakdown, the network may have congestion issues or be misconfigured. Examine the network switch, DNS, or VLAN configuration. Look for rogue scripts or applications that may be using up bandwidth. Network latency is the amount of time it takes to “recv” a complete a complete NFS RPC request to the storage.

**Storage** - If you see a lot of blue in the latency breakdown, this shows that data is being stored and read from flash on the VMstore. This is the operational latency of the VMstore.

**Disk** - If you see a lot of orange in the latency breakdown, then data is not being read from the working set in flash, but is being read from the second performance tier, which stores cold data.

---

**DO:** When architecting the Oracle database, use multiple Virtual disks to store datafiles. This allows you to use the Tintri performance insight tools to pinpoint the actual cause of latency (host, network, or storage) and quickly indentify and resolve any performance issue with the Oracle database.

---



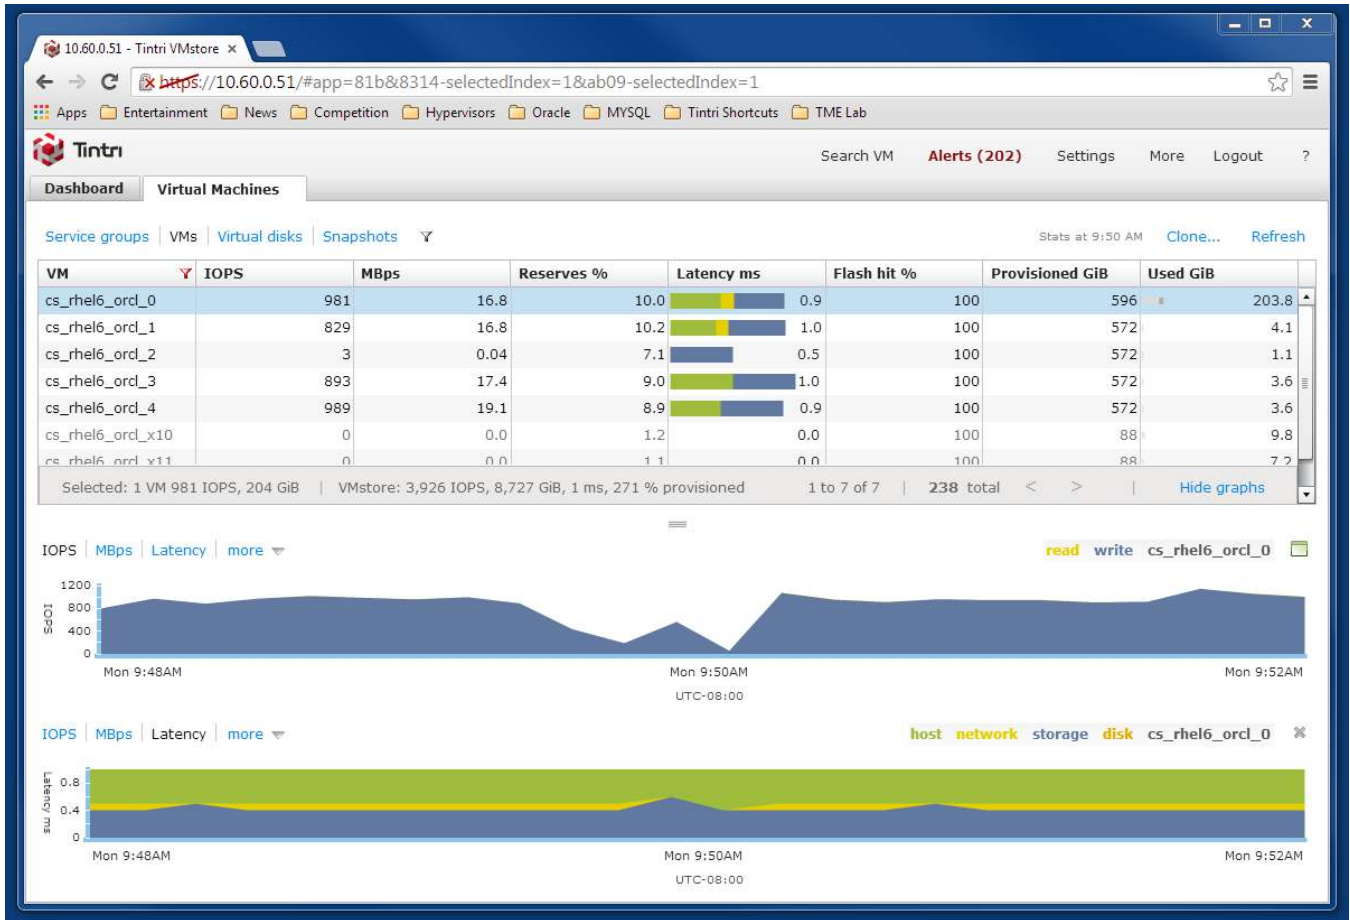


Figure 19 – Virtual Machines Tab with per-VM Performance Statistics Showing Latency (ms)

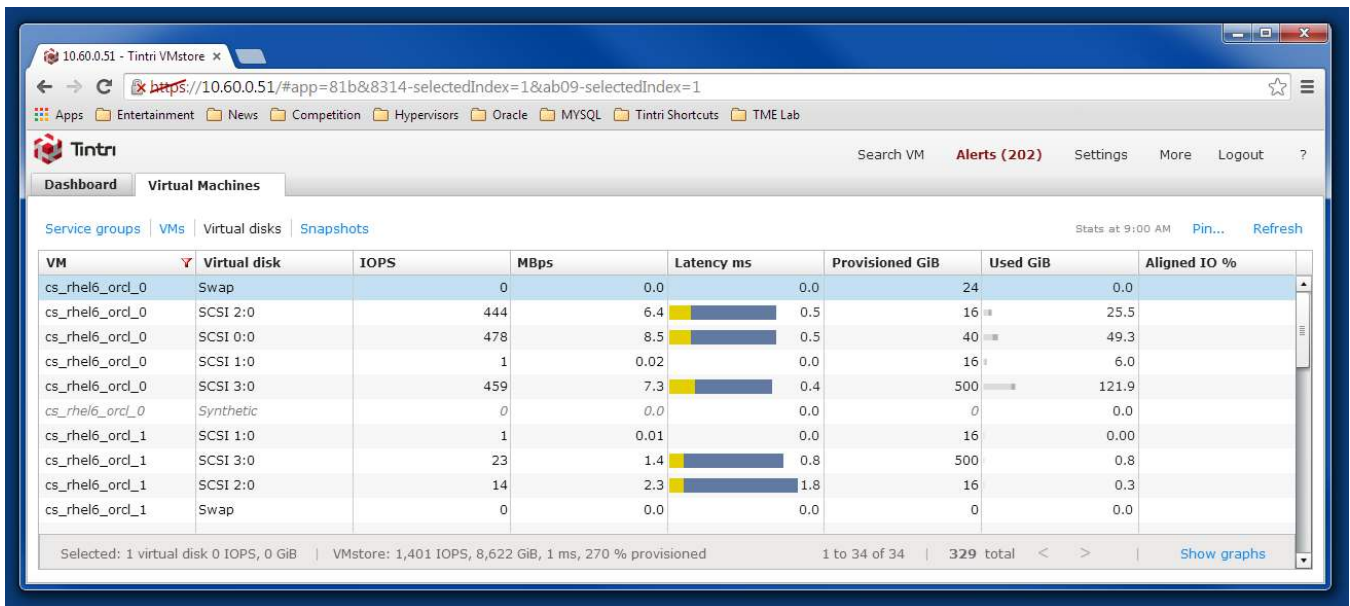


Figure 20 – Virtual Machines Tab with per-VMDK Performance Graphs showing Latency (ms)

## Conclusion

These Tintri best practices provide the information required for a successful deployment of a virtualized Oracle database with the Tintri VMstore. Tintri fully supports the use of virtualized Oracle databases so administrators can leverage Tintri's native per virtual machine data management capabilities for deploying, maintaining and troubleshooting.

The Tintri VMstore's feature rich UI provides a virtualized data center administrator the tools to effectively and efficiently monitor, manage and troubleshoot virtual machines. The ability to visualize performance bottlenecks across all layers of the infrastructure, the ability to track performance and capacity utilization on a per virtual machine and per virtual disk basis gives virtualization and database administrators the tools necessary to simplify and scale Oracle database deployments.

The use of Tintri VMstore's SnapVM, CloneVM, and ReplicateVM functionalities ensures that virtualized Oracle databases can be protected locally and remotely between multiple VMstores and across data centers. Given the hypervisor agnostic nature of Tintri VMstore systems, the same VM-level data management functionality can be used across multiple hypervisors deployed on a single Tintri VMstore. We hope the information in this document and the best practice guidelines helps you get the most from your Oracle deployments on Tintri VMstore systems.



## References

### Tintri Links

- [Tintri NFS and vSphere Best Practices](#)
- [Ensuring Robust Data Integrity with Tintri](#)
- [Backup and Recovery Best Practices with Tintri VMstore](#)

### Oracle Links

- [Overview of Oracle Automatic Storage Management](#)
- [Oracle ASM storage concepts](#)
- [Oracle ASM FAQ](#)

### RedHat Links

- [Deploying Oracle Database 11g R2 on Red Hat Enterprise Linux 6 – Best Practices](#)

### VMware Links

- [Oracle Databases on VMware Best Practices Guide](#)
- [Oracle Databases on VMware Frequently Asked Questions \(FAQ\)](#)
- [Oracle Databases on VMware vSphere](#)
- [Oracle Databases on VMware High Availability](#)
- [Best Practices for Performance Tuning of Latency-Sensitive Workloads in vSphere VMs](#)
- [VMware Networking Guide](#)
- [Performance Best Practices for VMware vSphere 5.5](#)
- [VMware vSphere 5.5 Configuration Maximums](#)
- [VMware vSphere App HA 1.1 Users Guide](#)
- [Performance Study of VMware vStorage Thin Provisioning](#)
- [KB Article 1001805 “Choosing a network adapter for your virtual machine”](#)
- [KB Article 1010398: Configuring disks to use VMware Paravirtual SCSI \(PVSCSI\) adapters](#)
- [KB Article 844: Configuring Linux Host for Optimal Performance and Virtual Machine Stability](#)
- [KB Article 1015180: Understanding virtual machine snapshots in VMware ESXi and ESX](#)
- [KB Article: Best practices for virtual machine snapshots in the VMware environment \(1025279\)](#)

## Appendix A – Support Statements

Oracle Support will assist customers running Oracle Products on VMware if they can demonstrate that the issue is not related to VMware. In turn, VMware will fully support customers running Oracle with their hypervisor technology, and they augment Oracle's policy with a total ownership policy for support issues. VMware Support will drive any issue to resolution regardless of which vendor (VMware, Oracle or others) is responsible for the resolution.

Tintri stands by VMware's statement of support, and we will assist customers with resolving issues with Tintri and VMware that are related to Oracle.

### [VMware Support for Oracle](#)

#### [Oracle Support for VMware Virtualized Environments](#)

#### [Oracle Support Policy for VMware by VMware is based on Total Ownership](#)

VMware Support will accept accountability for any Oracle-related issue reported by a customer. By being accountable, VMware Support will drive the issue to resolution regardless of which vendor (VMware, Oracle, or others) is responsible for the resolution. In most cases, reported issues can be resolved via configuration changes, bug fixes, or feature enhancements by one of the involved vendors.

In the rare situation that another vendor is unable or unwilling to provide a satisfactory technical resolution, VMware Support will immediately notify the customer, assist in escalation, and explore other potential technical workarounds with the customer.

VMware will also assist its customers with technical issues for other Oracle software products, besides the Oracle Database, and provide similar escalation assistance if needed.

Besides technical assistance, VMware Support will advocate on the customer's behalf to:

- Provide any relevant evidence that virtualization does not play a part in the Oracle product technical problem
- Engage Oracle Support in resolving the customer's technical issue, escalating management attention as appropriate

This document was created using the official VMware icon and diagram library. Copyright © 2012 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>.

VMware does not endorse or make any representations about third party information included in this document, nor does the inclusion of any VMware icon or diagram in this document imply such an endorsement.

Tintri, the Tintri logo, Tintri VMstore, and FlashFirst are trademarks or registered trademarks of Tintri, Inc. All other trademarks or service marks are the property of their respective holders and are hereby acknowledged.

Copyright © 2015 Tintri, Incorporated. All Rights Reserved.

[www.tintri.com](http://www.tintri.com)



303 Ravendale Drive  
Mountain View CA 94043  
+1 650.810.8200  
[info@tintri.com](mailto:info@tintri.com)