Enterprise Vault Best Practice Guide

Implementing Enterprise Vault on VMware

This document provides design and deployment considerations for implementing Enterprise Vault on VMware.

If you have any feedback or questions about this document please email them to **EV-TFE-Feedback@symantec.com** stating the document title.

This document applies to the following version(s) of Enterprise Vault:

10.0

This document applies to the following version(s) of VMware vSphere:

4.1, 5.0

Document Control

Contributors

Who	Contribution
Daniel Strydom	Author

Revision History

Version	Date	Changes
1.0	January 2011	Initial version

Related Documents

Version	Date	Title
EV10	01/08/2011	Enterprise Vault Product Installation and Configuration Guides http://www.symantec.com/docs/DOC4402
EV10	01/08/2011	Enterprise Vault 10 Performance Guide http://www.symantec.com/docs/DOC4553
EV10	01/08/2011	Enterprise Vault 10 Best Practice Guide for Indexing http://www.symantec.com/docs/DOC4250
EV10	01/08/2011	Enterprise Vault 10 Discovery Accelerator Best Practice Guide http://www.symantec.com/docs/DOC2797

Note: All links correct at the time of writing but may be subject to change

Table of Contents

Appendices

Scope of This Document

This document aims to provide guidance on designing and deploying Enterprise Vault 10 on the VMware vSphere 4.1 and 5.0 platform. The recommendations in this document do not apply to earlier versions of Enterprise Vault or VMware.

This document should be used in conjunction with other performance and best practice guides as outlined in the "Related Documents" section of this document.

Intended Audience

This document is aimed at system administrators, solutions architects, and consultants. It is assumed that the reader has a thorough understanding of the architecture and operational aspects of Enterprise Vault 10. It is also assumed that the reader has experience and understanding of VMware vSphere 4.1 and 5.0.

Terminology Used in This Document

Term	Description
Virtual Machine	A virtual machine is an isolated software container that can run its own operating systems and applications as if it were a physical computer
Hypervisor	A hypervisor, also called a virtual machine manager (VMM), is a program that allows multiple operating systems to share a single hardware host
VMware vSphere	Formerly known as VMware Infrastructure, vSphere comprises a suite of tools required to manage and support the virtual infrastructure
vCPU	CPU allocated to a virtual machine
VMware High Availability (HA)	Minimizes downtime with automated restart of Virtual Machines on a different physical host in the event of hardware failure
VMware vMotion	Reduces downtime from planned maintenance by migrating virtual machines between different physical hosts
VMware Fault Tolerance (FT)	Provides continuous downtime protection by keeping a live shadow copy of a virtual machine running on a separate host
Virtual Symmetric Multi-Processing (Virtual SMP)	Enhances virtual machine performance by allowing a virtual machine to use multiple physical processor cores simultaneously
NUMA	Non-Uniform Memory Access (NUMA) is a computer memory design used for Multiprocessing

Choosing the Right Platform for Enterprise Vault

Virtualization technology has helped many customers introduce cost savings both in terms of lowered data center power consumption and cooling requirements. Virtualization typically also simplifies the datacenter landscape through server consolidation, requiring less hardware to provide the same service to end users with the added benefit of application independent high availability.

Application architectures are rapidly evolving towards highly distributed, loosely-coupled applications. The conventional x86 computing model, in which applications are tightly coupled to physical servers, is too static and restrictive to efficiently support most modern applications. With a virtual deployment, the architecture can be as modular as is appropriate, without expanding the hardware footprint. The dynamic nature of virtual machines mean that the design can grow and adapt as required, without the need for an initial "perfect" design. Virtual deployments typically take minutes, can share currently deployed hardware, and can be adjusted "on the fly" when more resources are required.

Certain server applications are less suitable for virtualization, especially those requiring heavy use of physical server resources such as CPU and memory. Traditionally customers have been reluctant to place applications with high service level agreements such as Microsoft Exchange Server and SQL Server on a virtual platform, not only because the application's demand on resources meant that only one or two virtual machines could co-exist on a single server, but also because the server could not offer the same performance it would have on a physical server.

However, more powerful hardware, enhancements in virtualization technology and better support from application vendors now mean that customers are looking to virtualize even the top-end applications.

A number of factors should be considered before deploying Enterprise Vault in a VMware environment:

- Enterprise Vault is heavily dependent on CPU resources. In a typical physical server configuration it is not unusual for the CPU to run at 90% or higher utilization while archiving is being performed. Generally, the more powerful the processor, the better the ingestion and retrieval rates
- The recommended CPU and memory configuration for Enterprise Vault 10 is 8 CPU cores and 16GB RAM
- It is recommended that CPU and Memory resources are dedicated (reserved) to the Enterprise Vault server, and not shared with other virtual machines on the host. This aligns with VMware recommendations for virtualized Microsoft Exchange Server and SQL Server implementations
- Other system components such as network and storage need to be sized accordingly to prevent them from becoming a bottleneck

If the above considerations are acceptable and supported by the customer environment then it is likely that virtualizing the Enterprise Vault environment will be a good fit for the organization.

Sizing Enterprise Vault 10 for VMware

One of the most important considerations when sizing Enterprise Vault is a thorough understanding of the expected workload on each of the Enterprise Vault servers. However, the initial design of Enterprise Vault should be done independent of whether it will run on physical servers or virtual machines, with the main consideration being the customer requirements for archiving and eDiscovery.

It is outside the scope of this document to provide a design and sizing introduction to Enterprise Vault, but in general terms, once the customer requirements are understood, a close look at the archive targets will help determine what server resources will be required to not only archive the backlog but also keep up with the daily change (also known as the “steady-state”).

Enterprise Vault sizing utilities such as Exchange Mailbox Analyzer, Domino Mail File Reporter, File System Analyzer, SharePoint Analyzer and PST Analyzer will help provide a better understanding of what data is held within the archive targets and what the impact of different archiving policies will be. The data collected from these sizing utilities can then be used in the Enterprise Vault 10 Sizing Estimator Excel workbook. From here the tool will provide a number of recommendations, including number of Enterprise Vault servers, expected performance for different CPU configurations, SQL database recommendations and estimated Vault Store and Index storage figures over a 3 year period.

The Sizing Estimator workbook is able to adjust recommendations and expected performance if the option to deploy on a virtualized environment is selected. Choosing this option will reduce the expected performance of each Enterprise Vault server by 30%. This reduction is based on a sub-optimally configured VMware environment - with an optimal VMware configuration the virtual machine performance will be very close to that of a physical server.

The above sections represents an over simplified view of the design process as there will be many environment specific factors that will affect a design.

The most common mistake when designing Enterprise Vault is to **size for capacity**, as opposed to **sizing for performance**. The following sections in this guide will provide detail on how to design the various components for optimal configuration.

Differences in Versions and Editions of VMware

It is important to highlight the core differences between versions and editions of vSphere when deciding whether VMware is a suitable platform for Enterprise Vault. The following table highlights some crucial differences in how much CPU and RAM can be allocated to a single virtual server:

	VMware ESXi 4.1 Standard	VMware ESXi 4.1 Enterprise	VMware ESXi 4.1 Enterprise Plus
vCPU entitlement per virtual server	4 CPU Cores	4 CPU Cores	8 CPU Cores
vRAM entitlement per virtual server	24GB	32GB	48GB
Meets recommended minimum for EV 10 (8 CPU cores, 16GB RAM)	No	No	Yes

Table 1 - VMware vSphere 4.x Edition CPU & Memory Maximums

	VMware ESXi 5.0 Standard	VMware ESXi 5.0 Enterprise	VMware ESXi 5.0 Enterprise Plus
vCPU entitlement per virtual server	8 CPU Cores	8 CPU Cores	32 CPU Cores
vRAM entitlement per virtual server	32GB	64GB	96GB up to 1TB
Meets recommended minimum for EV 10 (8 CPU cores, 16GB RAM)	Yes	Yes	Yes

Table 2 - VMware vSphere 5.0 Edition CPU & Memory Maximums

The Enterprise Vault 10 compatibility guide lists supported versions of VMware. Note that VMware ESXi 4.0 is not supported in for Enterprise Vault 10. At the time of writing the only supported versions of ESXi are 4.1 and 5.0.

ESX Host Configuration

Each ESX host should provide enough physical hardware resources to accommodate the planned workload and provide some headroom in the event of a VMware HA failover or planned VMware vMotion migration of live virtual machines for host hardware maintenance.

If planning to use VMware vMotion, confirm that the CPU supports vMotion (which in turn affects DRS) and VMware Fault Tolerance. Confirm that the selected hardware is supported by searching the VMware product compatibility database.

The following general recommendations can be made regarding the ESX host:

- Confirm that the host is running the latest version of the BIOS available
- Enable the “Turbo Boost” option in the BIOS, if supported by the CPU
- Some NUMA-capable systems provide an option in the BIOS to disable NUMA by enabling node interleaving. Generally the best performance is achieved by disabling node interleaving (in other words, leaving NUMA enabled)
- Ensure that any hardware-assisted virtualization features (VT-x, AMD-V, EPT, RVI, etc.) are enabled in the BIOS
- Disconnect or disable any physical hardware devices that will not be used. These might include devices such as COM ports, LPT ports, USB controllers, floppy drives, network interfaces and storage controllers. Disabling hardware devices can free interrupt resources. Additionally, some devices, such as USB controllers, operate on a polling scheme that consumes extra CPU resources. Some PCI devices reserve blocks of memory, making that memory unavailable to the ESX host
- Set the ESX host power policy to “Maximum performance” or disable power management altogether. It is also recommended to disable C1E and other C-states in BIOS.

Optimizing VMware for Enterprise Vault

The following sections will cover recommendations regarding the VMware and Enterprise Vault components. ESX resource configuration is vitally important to ensure applications such as Enterprise Vault run with optimal performance.

Not all recommendations will be suitable to all environments – these recommendations are aimed at creating the best possible platform specifically tuned for Enterprise Vault.

CPU configuration

VMware Virtual Symmetric Multi-Processing (Virtual SMP) enhances virtual machine performance by allowing a virtual machine to use multiple physical processor cores simultaneously. VMware ESXi 4.1 supports the use of up to 8 virtual CPUs per virtual machine, ESXi 5.0 supports up to 32 virtual cores. Virtual SMP allows virtual machines to use multiple processors to execute multiple tasks concurrently, thereby increasing throughput.

The virtual processors are co-scheduled and run simultaneously, and (providing physical cores are available on the host) each vCPU is mapped one-to-one to physical processors. In practice what this means is that if one vCPU is running, a second vCPU is co-scheduled so that they execute nearly synchronously. Even when the operating system is idle, the vCPUs will still perform a minimal amount of work and needs to be managed by the ESX host, effectively competing with other vCPUs for system resources.

A common misconception is that by assigning more than 2 CPU cores per virtual machine you actually reduce the performance of that server. While this may have been true in earlier versions of VMware, a number of improvements in the VMware 4.1 and 5.0 co-scheduling algorithm means that the virtual machines with multiple vCPUs are now more scalable and perform better, and the effects of idle multi-vCPU machines are reduced. In ESX 4.1 and 5.0 the larger 4-core and 8-core vCPU systems not only scale better, but it also means that it's easier to scale the virtual machines up, compared to scaling out[1].

The following general recommendations can be made regarding vCPU allocation for Enterprise Vault:

- Allocate 8 CPU cores to each Enterprise Vault server. For smaller or non-production environments with less than 500 users 4 CPU cores can be allocated
- Set a CPU Reservation - this guarantees the CPU is dedicated to the virtual machine. Generally this practice is not recommended because the reserved resource is then not available to other virtual machines, but VMware recommends this practice in cases where service level agreements need to be guaranteed. The processor priority and bandwidth should be set to provide the virtual machine with full utilization of the selected CPUs

[1] For more details see ESX CPU consideration sections in VMware Performance Best Practices available at http://www.vmware.com/pdf/Perf_Best_Practices_vSphere4.0.pdf and http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.0.pdf

- Ensure that the total number of vCPUs assigned to the virtual machines is equal or less than the total number of cores on the ESX host
- Do not enable Hyperthreading – in most cases this provides little or no benefit to multi-CPU virtual machines. Internal Symantec testing has shown that Hyperthreading provides no performance benefit to Enterprise Vault.

Memory Configuration

Symantec recommends that each Enterprise Vault 10 server is allocated 16GB of memory. For smaller or non-production environments (less than 500 user) 8GB of memory should be sufficient. The memory requirements do however increase depending on the search requirements; specifically in environments where regular eDiscovery searches are run the memory should be increased to 32GB.

Virtual machines also require memory beyond the amount allocated to account for memory overhead; this memory needs to be available to the physical host. Memory overhead includes space reserved for virtual machine devices, such as SVGA frame buffers and internal data structures. The amount of overhead required depends on the number of vCPUs allocated per virtual machine, the amount of configured memory and whether the operating system is 32 or 64-bit. As an example, a virtual machine with 2 vCPUs and 32GB of memory will consume approximately 500MB of memory overhead. This memory overhead must be available on the ESX host, and any other ESX host potentially hosting the Enterprise Vault virtual server.

The following diagram describes the use of memory settings available for a virtual machine:

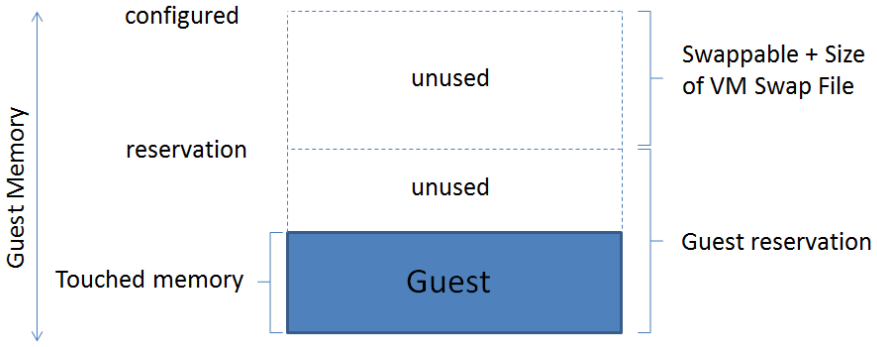

Figure 1 - Memory Settings per virtual machine

The memory settings include the following parameters:

- Configured memory: The amount of memory allocated to the virtual machine when it was created
- Touched memory: The amount of memory actually in use by the virtual machine
- Swappable: Memory that can be reclaimed and reallocated by the vSphere balloon driver. If this memory is being used by the virtual machine the balloon driver will cause the guest operating system to swap. This value is the size of the per virtual machine swap file that is created on the VMFS file system (as a .vswp file).

The following best practice recommendations can be made regarding memory allocation:

- Set the memory reservation to the configured size of the memory allocated to the virtual machine. Note that by setting the reservation to the configured size of the virtual machine the vmkernel swap file will be 0 bytes. This will consume less storage and help increase performance by eliminating ESX host-level swapping. Note that Windows will still maintain its own separate swap/page file on the guest operating system
- Setting memory reservations as described above may limit VMware vMotion as a virtual machine can only be migrated if the target ESX host has sufficient free physical memory
- Use ESX performance counters to measure actual memory usage, see later section in this document on the use of “Active”, “SwapIn” and “SwapOut” counters
- Ensure that the Enterprise Vault server does not reside on an ESX host where the memory have been over committed
- It is generally recommended not to disable the memory balloon driver on the virtual machine (installed with VMware Tools)
- If available, enable vSphere Dynamic Resource Scheduler (DRS) within the ESX cluster. Together with reservations DRS can guarantee the resource availability
- Configure the virtual machine memory size to be greater than the average memory usage of the Enterprise Vault server. This will avoid guest operating system swapping. In a typical Enterprise Vault environment around 6GB of RAM will be used in core activities during the day, increasing to the full available 16GB recommended memory with journaling and searching (end user searches, Virtual Vault activity, etc.). For more details on expected server throughput for different CPU and memory configurations refer to the EV10 Performance Guide[2].
- If the Enterprise Vault virtual server requires more memory than has been allocated, the performance will reduce as the guest operating system will have to perform swap operations. It is recommended that the host is allocated a swap file size 1.5x that of the memory assigned to the virtual machine (as per standard recommendation for physical hosts).

The recommendations relating to memory are purposely conservative to avoid kernel swapping between ESX and the guest operating system. Once the workload is known and predictable, if VMware vCenter reports that steady state active memory usage is below the amount of memory on the ESX host, then the reservation settings may be relaxed to the steady state active memory value. To get a better understanding of this

[2] www.symantec.com/docs/doc4553

scenario, refer to VMworld content "TA2627 – *Understanding "Host" and "Guest" Memory Usage and Related Memory Management Concepts"*[3].

Disk Configuration

Storage performance depends on many factors, including the workload, hardware, RAID level, cache size, stripe size, etc. The storage vendor documentation should always be consulted along with recommendations from VMware.

VMware Virtual Machine File System (VMFS) is a high performance cluster file system developed to scale beyond the use of a single system. VMFS provide multiple virtual machines shared access to a consolidated pool of clustered storage, and provide the foundation for enabling services such as vMotion, Distributed Resource Scheduler (DRS) and VMware High Availability. Each virtual machine is encapsulated in a small set of files; in most cases these files are stored using the VMFS file system on physical SCSI disks.

VMware supports Fibre-Channel, iSCSI, and NAS shared-storage protocols. VMware storage virtualization can be categorized into three layers of storage technology.

[3] http://www.vmworld.com/docs/DOC-3817 (VMworld access is required, apply for a free account)

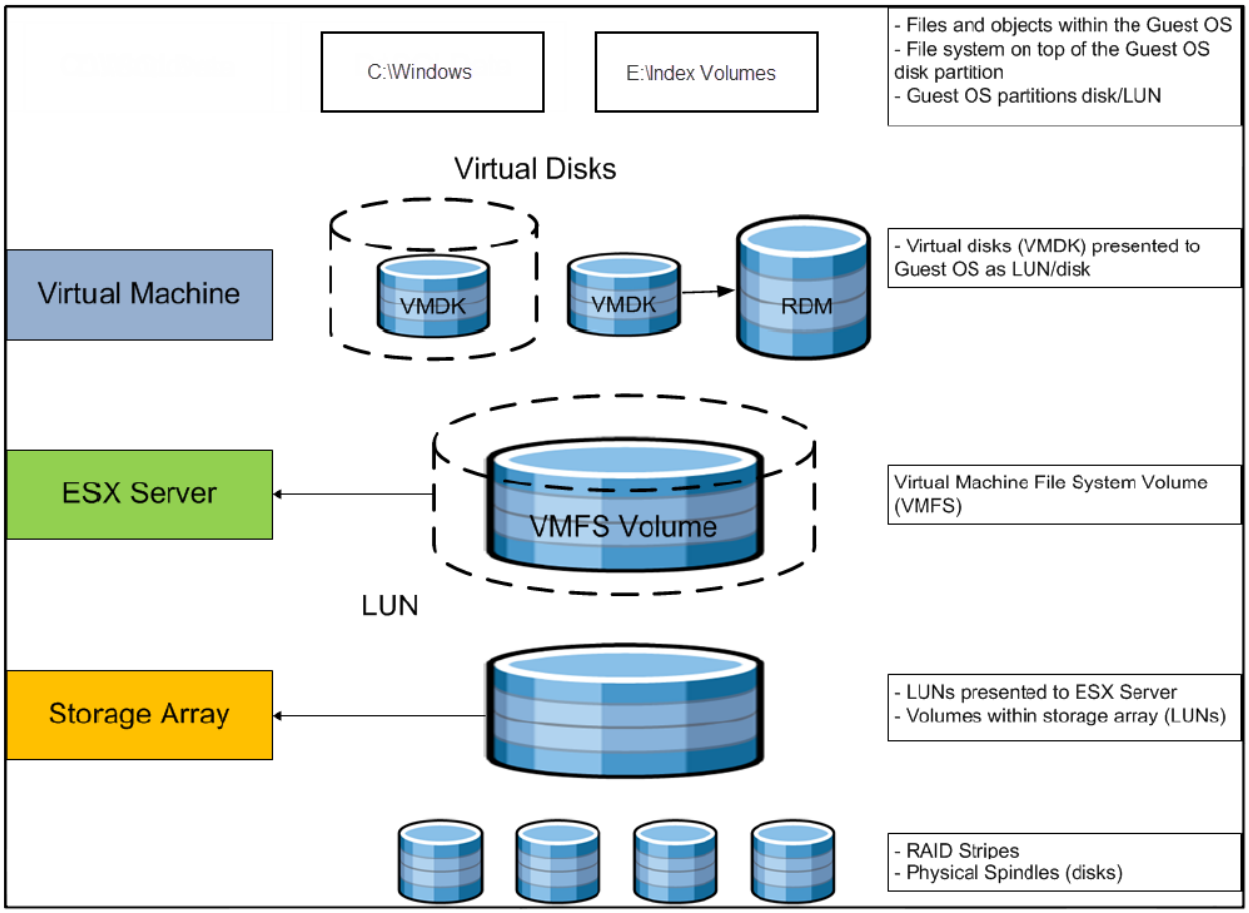

Figure 2 - VMware Storage Virtualization Stack

The storage array (consisting of physical disks) is presented as logical disks in the form of storage array volumes or LUNs to the ESX server. Storage array LUNs are then formatted as VMFS volumes in which virtual disks can be created. Virtual machines consist of virtual disks that are presented to the guest operating system as disks that can be partitioned and used in file systems.

The best way to configure a LUN for a given VMFS volume is to size for throughput first and capacity second. That is, you should aggregate the total I/O throughput for all applications or virtual machines that might run on a given shared pool of storage; then make sure you have provisioned enough back-end disk spindles (disk array cache) and appropriate storage service to meet the requirements.

Aside from VMFS, ESXi also supports raw device mapping (RDM) as referenced in the above diagram. RDM effectively allows management and access of raw SCSI disks or LUNs as VMFS files. RDM is a special file on a VMFS volume that acts as a proxy for a raw device. The file contains metadata used to manage the disk, and is able to communicate direct access to the physical device. The mapping file, not the raw volume, is referenced in the virtual machine configuration.

RDMs can be used in virtual or physical compatibility mode. Virtual mode specifies full virtualization of the mapped drive, allowing the guest operating system to use the RDM like any other virtual disk file in a VMFS volume. Physical mode specifies minimal SCSI virtualization of the mapped drive, giving the most flexibility to SAN management software.

It is often a difficult decision when it comes to choosing between the VMFS and RDM for data volumes. They both provide similar performance characteristics - random access (50% read, 50% write) performance is very similar, but for sequential read operations RDM provides more I/O per second. Test results vary depending on the data size used, for a more detailed analysis refer to the "Performance Characterization of VMFS and RDM Using a SAN"[4] whitepaper by VMware.

The following table described some of the main differences between VMFS and RDM.

VMFS	RDM
Many virtual machines can be hosted on one volume (can be dedicated to one virtual machine)	Maps a single LUN to one virtual machine
Better flexibility, easier management	Requires more LUNs, consideration of 256 LUN limit per ESX host
Large third-party VMware ecosystem with V2P products to aid in certain support scenarios	RDM volumes can help facilitate migrating physical servers to virtual machines using the LUN swing method
Cannot be used for in-guest clustering, does not support quorum disks as required by MSCS	Required for in-guest clustering. Cluster data and quorum disks should be configured using RDM
Supports VMware VMotion, HA and Distributed Resource Scheduler (DRS) and VMware Site Recovery Manager (SRM)	RDM guarantees no other virtual machine is able to use the LUNs
	Required to leverage array-level backup and replication tools (VSS) integrated
	Supports VMware VMotion, HA and Distributed Resource Scheduler (DRS) and VMware Site Recovery Manager (SRM)

Table 3 - VMFS and RDM Trade-offs

[4] http://www.vmware.com/files/pdf/performance_char_vmfs_rdm.pdf

Depending on existing storage practices it may be advantageous to mix VMFS and RDM under the following circumstances:

- Where 3rd party storage management software is already in use RDM can be used to leverage that infrastructure, for example if using storage-based backups to disk
- RDM is required when using clustering technologies like Microsoft Clustering Services (MSCS)
- One advantage of using RDM volumes is that it gives the administrator the ability to point both virtual and physical machines to the same storage.

Both VMFS and a mixed VMFS/RDM storage configuration are suitable for Enterprise Vault. In a mixed configuration the guest operating system is installed on VMFS, while the Index volumes can reside on RDM. The Vault Store partitions can also be stored on RDM, or if preferred on a NAS or other iSCSI storage partition.

For more information and a more detailed comparison between VMFS and RDM see the VMware VMFS Best Practices Whitepaper[5].

Number of virtual machines per LUN

You can assign multiple virtual machines to a VMFS LUN, however for best performance it is recommended that a single LUN is dedicated to a virtual machine. The following diagram discusses the differences.

[5] www.vmware.com/pdf/vmfs-best-practices-wp.pdf

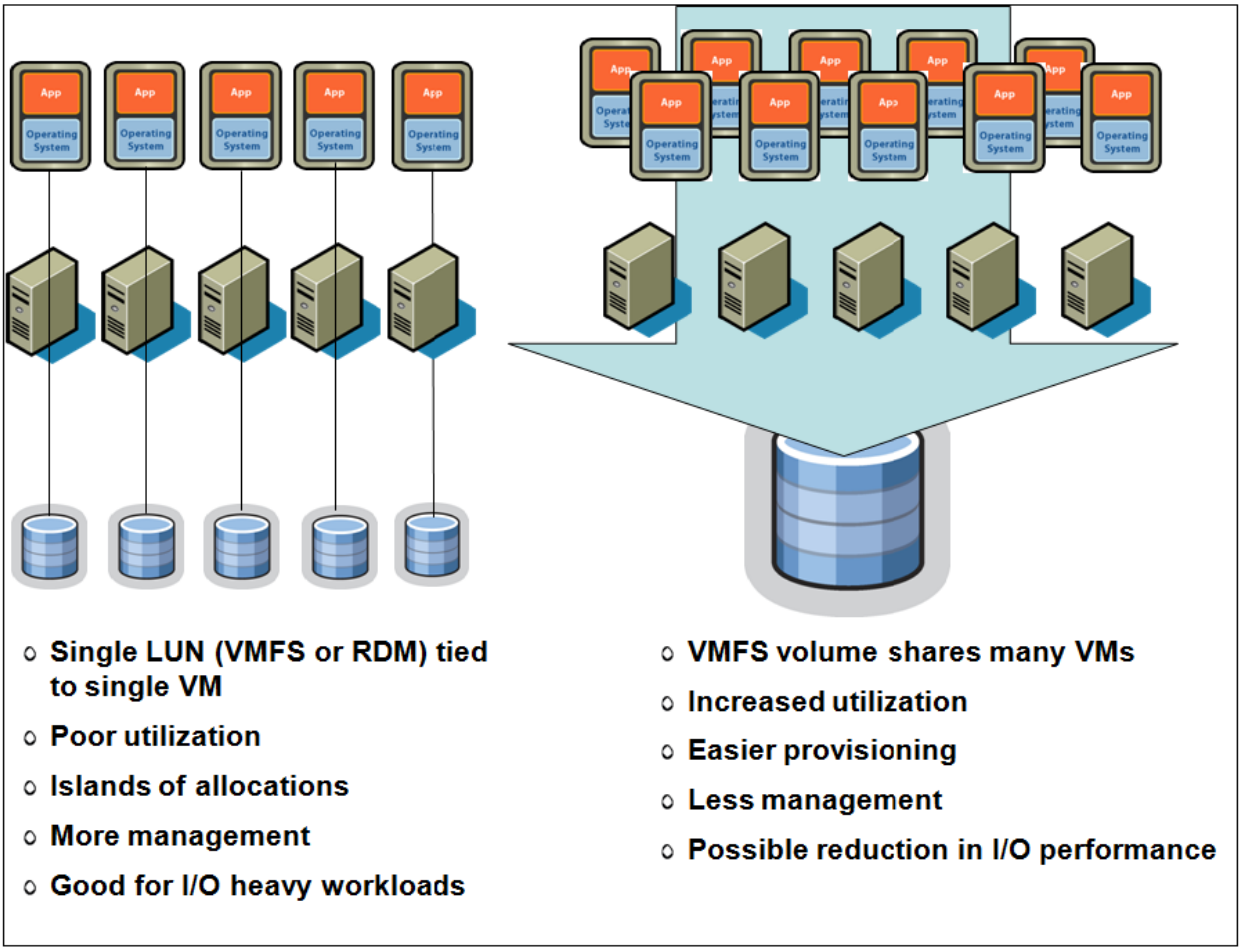

Figure 3 - One versus many virtual machines in a LUN

It's considered best practice to separate heavy I/O workloads from the shared pool of storage in order to optimize the performance of applications requiring high transactional throughput. This approach is best characterized as "consolidation with some level of isolation".

When looking at the number of VMs per LUN, there is no exact rule to determine the limits of performance and scalability, as workloads can vary significantly. These limits also depend on the number of ESX Servers sharing concurrent access to a given VMFS volume. The key is to remember the upper limit of 256 LUNs per ESX Server and consider that this number can limit the consolidation ratio if you take the concept of "1 LUN per VM" too far.

Enterprise Vault per Server Storage Locations

The following table details storage areas within Enterprise Vault that can affect the application performance. For optimal performance, each of the storage areas should be treated as separate areas and their individual needs catered for.

Per server storage area	Type of access	Expected Size	Recommendation
Index Locations	Sequential Access. Large flat file area consisting of many small files	Depending on Indexing Level – Brief 3% of total information archived, Full Indexing will require 13%	Fast SAN based storage – FC or iSCSI, can be mapped as VMFS or RDM. Storage should be capable of supporting 400+ IOPS, 2000+ IOPS for eDiscovery environments. NAS storage not considered appropriate in most scenarios (refer to Indexing Best Practice paper for more information)
Vault Store Partitions	Random Access. Large flat file area consisting of many small files	Very large but split up in smaller “vault partitions”.	Slower, lower tier storage such as NAS is appropriate. Low IOPS requirement
MSMQ Data Folder	Random Access	40GB	Fast SAN based or local RAID1 disk. Ideally split onto a different spindle set from MSMQ Log area
MSMQ Log Folder	Sequential Access	10GB	Fast SAN based or local RAID1
EMC Centera Collections	Random Access	100GB	Fast SAN based protected disk, RAID1 or better performing stripe
Vault Server Cache Area	Random Access	50GB	Fast SAN based disk, RAID1 or better performing stripe. Storage should be capable of supporting 400 IOPS for typical environment. Virtual Vault environments can expect up to 1,000 IOPS during busy periods
Windows Temp Directory	Random Access	2GB	Fast SAN based disk, used for temporary EV operations RAID1 or better performing stripe. Used during archiving to process large files

Table 4 - Enterprise Vault per Server Storage Locations

Choosing Suitable Storage for Index Locations

When choosing a suitable storage device the whole storage solution should be considered. The supported IOPS is just one aspect of performance and other areas should also be looked at such as connectivity – for example Fiber and iSCSI are preferred over CIFS.

As per previous versions of Enterprise Vault, a high speed storage device is recommended for Index locations. Lower tier NAS devices are generally not recommended for index locations, and should definitely not be used to host indexes where any type of eDiscovery search is used.

NAS devices connected over CIFS shares is not suitable to host index volumes for certain environments due to the slower connectivity speed.

The best devices are local storage, direct attached storage, or SAN LUNs.

In the case of local or direct attached storage:

- Use multiple controllers supporting multiple channels to distribute the load between index file locations and provide sufficient throughput
- Provide battery-backed read and write cache to aid performance.

Before using partitions on a SAN, consider the I/O load together with any other applications that are already using the SAN to ensure that the performance can be maintained. Ideally, the implementation should be discussed with the SAN hardware vendor to ensure that optimum performance is achieved. Typically LUNs should be created across as many suitable disks as possible, using entire disks rather than partial disks to prevent multiple I/O-intensive applications from using the same disks.

General ESXi Storage Recommendations

- Monitor storage on a regular basis. I/O latency statistics can be monitored using esxtop (or resxtop), which reports device latency, time spent in the kernel, and latency seen by the guest operating system
- Check that the average latency for storage devices is not too high. Latency can be seen in esxtop (or resxtop) by looking at the GAVG/cmd metric. A reasonable upper value for average latency depends on the storage subsystem. This value should be confirmed with the storage vendor
- For iSCSI and NFS, if the network switch deployed for the data path supports VLAN, it might be beneficial to create a VLAN just for the ESXi host's vmknic and the iSCSI/NFS server. This minimizes network interference from other packet sources
- Local storage performance might be improved with write-back cache. Ensure that write-back cache installed and contains a functional battery module
- To optimize storage array performance, spread I/O loads over the available paths to the storage (that is, across multiple host bus adapters (HBAs) and storage processors)
- The default virtual storage adapter in ESXi 5.0 is either BusLogic Parallel, LSI Logic Parallel, or LSI Logic SAS, depending on the guest operating system and the virtual hardware version. However, ESXi also includes a paravirtualized SCSI storage adapter, PVSCSI (also called VMware Paravirtual). The PVSCSI adapter offers a significant reduction in CPU utilization as well as potentially increased throughput compared to the default virtual storage adapters, and would therefore be the best choice for environments with very I/O-intensive guest applications. In order to use PVSCSI, your virtual machine must be using virtual hardware version 7 or later
- If you choose to use the BusLogic Parallel virtual SCSI adapter, and are using a Windows guest operating system, you should use the custom BusLogic driver included in the VMware Tools package
- The depth of the queue of outstanding commands in the guest operating system SCSI driver can significantly impact disk performance. A queue depth that is too small, for example, limits the disk bandwidth that can be pushed through the virtual machine. See the driver-specific documentation for more information on how to adjust these settings – this is a very common performance bottleneck
- Ensure that an adequate number of spindles is available to support the I/O requirements with a acceptable latency
- Make sure the disk partitions within the guest are aligned
- Use up-to-date HBA drivers recommended by the storage vendor
- Ensure that the storage array firmware is up to the latest recommended level

Make sure to give thought to the growth strategy up front. As the Vault Store and Index volume grows, how will the growth of data files / LUNs / RAID groups be managed? It is much better to design for this up front than to rebalance data files or LUN(s) later in a production deployment

Network Configuration

The following general best practice recommendations can be made regarding VMware network configuration for Enterprise Vault:

- Separate virtual machine and infrastructure traffic – Keep virtual machine and VMkernel or service console traffic separate. This can be accomplished physically using separate virtual switches that uplink to separate physical NICs, or virtually using VLAN segmentation
- Use NIC Teaming – Use two physical NICs per vSwitch, and if possible, uplink the physical NICs to separate physical switches. Teaming provides redundancy against NIC failure and, if connected to separate physical switches, against switch failures. NIC teaming does not necessarily provide higher throughput
- Enable PortFast on ESX host uplinks – Failover events can cause spanning tree protocol recalculations that can set switch ports into a forwarding or blocked state to prevent a network loop. This process can cause temporary network disconnects. To prevent this situation, set the switch ports connected to ESX hosts to PortFast, which immediately sets the port back to the forwarding state and prevents link state changes on ESX hosts from affecting the STP topology. Loops are not possible in virtual switches
- Converge Network and Storage I/O with 10Gbps Ethernet – When possible consolidating storage and network traffic can provide simplified cabling and management over having to maintain separate switching infrastructures.

Performance Impact of Upgrading to Enterprise Vault 10 from Earlier Versions

Enterprise Vault 10 introduces a new 64-bit indexing engine, offering many advantages over the 32-bit engine found in earlier versions. The new indexing engine is more scalable and provides better performance but, as is common with any 64-bit application, it requires more powerful hardware – the recommended CPU and memory requirements have increased to 8 CPU cores and 16GB of RAM as mentioned earlier in this document.

If you are upgrading from Enterprise Vault 9 on an existing VMware environment it is important that these requirements are met, and supported by the version of VMware ESXi[6]. Customers using this recommended hardware should expect to see equal or better archiving performance when compared to a server running EV9 on the recommended hardware specified for that version.

[6] 8 CPU cores are supported by VMware ESXi 4.1 and newer

Virtualizing SQL Server

Follow the recommendations of your VMware and Microsoft when you size and configure the environment for SQL Server. The following general recommendations can be made when virtualizing SQL for Enterprise Vault:

- In a typical virtualized infrastructure, local disks would be used for the hypervisor and SAN-based storage for the guest operating system images and data file locations. The operating system and data storage partitions should be independent dedicated locations.
- Virtual hard disks should be created as fixed size and not dynamic
- The required memory capacity should be dedicated and prioritized to the virtual machine to prevent dynamic allocation or sharing
- Avoid the use of hyper-threading by the hyper-visor
- The processor cores should be exclusively dedicated to the virtual machine, and the processor priority and bandwidth set to provide the virtual machine with full utilization of the selected CPUs

Other Deployment Recommendations

General VMware Recommendations

Generally it is recommended that any unnecessary programs or graphic effects (such as screen savers and Window animations) and disable in virtual machines.

The following recommendations can be made regarding VMware servers:

- Install the latest version of VMware Tools in the guest operating system. Make sure to update VMware Tools after each ESXi upgrade. Installing VMware Tools in Windows guests updates the BusLogic SCSI driver included with the guest operating system to the VMware-supplied driver. The VMware driver has optimizations that guest-supplied Windows drivers do not
- Schedule backups and virus scanning programs in virtual machines to run at off-peak hours. Avoid scheduling them to run simultaneously in multiple virtual machines on the same ESXi host. For workloads such as backups and virus scanning, where the load is predictable, this is easily achieved by scheduling the jobs appropriately
- For the most accurate timekeeping, consider configuring your guest operating system to use NTP, Windows Time Service or the VMware Tools time-synchronization option. On any particular machine it is recommended that either the VMware Tools time-synchronization option is used or another timekeeping utility, but not both
- Unused or unnecessary virtual hardware devices can impact performance and should be disabled. For example, Windows guest operating systems poll optical drives (that is, CD or DVD drives) quite frequently. When virtual machines are configured to use a physical drive, and multiple guest operating systems simultaneously try to access that drive, performance could suffer. This can be reduced by configuring the virtual machines to use ISO images instead of physical drives, and can be avoided entirely by disabling optical drives in virtual machines when the devices are not needed
- ESXi 5.0 introduces virtual hardware version 8. By creating virtual machines using this hardware version, or upgrading existing virtual machines to this version, a number of additional capabilities become available. Some of these, such as support for virtual machines with up to 1TB of RAM and up to 32 vCPUs, support for virtual NUMA, and support for 3D graphics, can improve performance for some workloads

The following general recommendations apply to Enterprise Vault index locations:

- If indexes are stored on NetApp devices, and possibly other NAS systems, opportunistic locking must be turned off for volumes that contain indexes
- Disable Windows file indexing on the drives that contain Enterprise Vault indexes.

Disable Anti-Virus for Index Locations

Ensure that there is no Anti-Virus scanning the index locations, real-time or scheduled. Failing to exclude the indexes from scanning can cause index corruption and will result in reduced server performance. A new

requirement in EV10 is to exclude the Indexing meta-data folder, default location "<install path>\EVIndexing\data\indexmetadata".

Pagefile Size

The Pagefile should conform to Microsoft's recommendation of 1.5 times the amount of memory. It is recommended that the page file is placed on a different partition and different physical hard disk drive from the system partition so that Windows can handle multiple I/O requests more efficiently.

Additionally, if you set the "Initial size" of the Pagefile to the same as the "Maximum size" it will prevent unnecessary fragmentation file. An unfragmented paging file leads to faster virtual memory access.

Index Volume Fragmentation

The index files quickly become fragmented on disk, even if there is a large volume of free storage capacity. This file fragmentation can cause severe performance problems which need to be managed on any index storage device. Either an automated background file defragmentation product or scheduled device defragmentation must be employed.

Ongoing Maintenance

Performance Monitoring

When measuring performance from within virtual machines any time-based performance counter is likely to be inaccurate, especially when the processor is overcommitted. Generally it is safe to assume the results are no more than 10% in error if CPU utilization stays below 80%.

To monitor the CPU usage on the host, use vSphere esxtop or resxtop. To interpret the esxtop data:

- If the load average on the first line of the esxtop CPU panel is equal to or greater than 1, this indicates that the system is overloaded
- The usage percentage for the physical CPUs on the PCPU line can be another indication of a possibly overloaded condition. In general, 80% usage is a reasonable ceiling and 90% should be a warning that the CPUs are approaching an overloaded condition.

Enterprise Vault administrators should pay particular attention to the following CPU performance metrics:

Esxtop Metric	Description	Implication
%RDY	The percentage of time a vCPU in a run queue is waiting for the CPU scheduler to let it run on a physical CPU.	A high %RDY time (use 20% as a starting point) may indicate the virtual machine is under resource contention. Monitor Enterprise Vault performance, if the archiving rates are as expected a higher threshold may be tolerated.
%MLMTD	Percentage of time a vCPU was ready to run but was deliberately not scheduled due to CPU limits.	A high %MLMTD time may indicate a CPU limit is holding the virtual machine in a ready to run state. If the application is running slow consider increasing or removing the CPU limit.
%CSTP	Percentage of time a vCPU spent in read, co-descheduled state (only applicable for virtual machines with more than one vCPU).	A high %CSTP time usually indicates that vCPUs are not being used in a balanced way – evaluate the necessity for multiple vCPUs in smaller environments.

Windows Performance Counters

The following counters may be useful to monitor within the virtual machine:

Objects and Counters	Description
Processor	
% Processor Time	Shows processor usage over a period of time. If this counter is consistently too high it is likely that the system performance will be impacted. You can measure the utilization on each processor to achive balanced performance between cores.
Disk	
Avg. Disk Queue Length	Shows the average number of read and write requests that were queued for the selected disk during the sample interval. A bigger disk queue length may not be a problem as long as disk reads/writes are not suffering and the system is working in a steady state without expanding queuing.
Avg. Disk Read Queue Length	The average number of read requests that are queued.
Avg. Disk Write Queue Length	The average number of write requests that are queued.
Disk Reads/sec	The number of reads to disk per second.
Disk Wites/sec	The number of writes to disk per second.
Memory	
Available Mbytes	Shows the amount of physical memory available for allocation. Insufficient memory is likely to cause excessive use of the page file and an increase in the number of page faults per second.
Cache Faults/sec	Shows the rate at which faults occur when a page is sought in the file system cache and is not found. This may be a soft fault, when the page is found in memory or a hard fault when the page is on disk. The use of cache for read and write operations can have a significant impact on server performance. Monitor for increased cache failures, indicated by a reduction in the Async Fast Reads/sec or Read Aheads/sec.
Pages/sec	Shows the rate at which pages are read from or written to disk to resolve hard page faults. If this rises, it indicates system-wire performance problems.

Objects and Counters	Description
Paging File	
% Used and % Used Peak	The server paging file holds “virtual” memory addresses on disk. Page faults occur when a process has to stop and wait while required “virtual” resources are retrieved from disk into memory. These are more frequent if the physical memory is inadequate.
Network	
Total Bytes/sec	The rate at which data is sent and received on the network interface. If the rate is over 50% capacity you should investigate for issues. To troubleshoot monitor Bytes Received/sec and Bytes Sent/sec.
Avg. Disk Read Queue Length	The average number of read requests that are queued.
Avg. Disk Write Queue Length	The average number of write requests that are queued.
Disk Reads/sec	The number of reads to disk per second.
Disk Wites/sec	The number of writes to disk per second.

Appendix A – VMware Enhancements to Deployment and Operations

VMware vMotion, VMware DRS and VMware HA

VMware VMotion technology enables the migration of virtual machines from one physical server to another without service interruption this migration allows you to move Exchange virtual machines from a heavily-loaded server to one that is lightly loaded, or to offload them to allow for hardware maintenance without any downtime.

VMware Distributed Resource Scheduler (DRS) takes the VMware VMotion capability a step further by adding an intelligent scheduler. DRS allows you to set resource assignment policies that reflect business needs. VMware DRS does the calculations and automatically handles the details of physical resource assignments. It dynamically monitors the workload of the running virtual machines and the resource utilization of the physical servers within a cluster.

VMware VMotion and VMware DRS perform best under the following conditions:

- The source and target ESX hosts must be connected to the same gigabit network and the same shared storage
- A dedicated gigabit network for VMware VMotion is recommended
- The destination host must have enough resources
- The virtual machine must not use physical devices such as CD ROM or floppy
- The source and destination hosts must have compatible CPU models, or migration with VMware VMotion will fail. For a listing of servers with compatible CPUs, consult VMware VMotion compatibility guides from specific hardware vendors
- To minimize network traffic it is best to keep virtual machines that communicate with each other together (e.g., Enterprise Vault Mailbox and SQL Server) on the same host machine
- Virtual machines with smaller memory sizes are better candidates for migration than larger ones.

With VMware High Availability (HA), Enterprise Vault virtual machines on a failed ESX host can be restarted on another ESX host. This feature provides a cost-effective failover alternative to expensive third-party clustering and replication solutions. If you use VMware HA, be aware that:

- VMware HA handles ESX host hardware failure and does not monitor the status of the Enterprise Vault services—these must be monitored separately
- Proper DNS hostname resolution is required for each ESX host in a VMware HA cluster
- VMware HA heartbeat is sent via the vSphere VMkernel network, so redundancy in this network is recommended.

Templates

VMware template cloning can increase productivity of system administration and testing in Enterprise Vault environments. A VMware template is a golden image of a virtual machine that can be used as a master copy to create and provision new virtual machines. It includes the guest operating system and Enterprise Vault application data. You can use virtual machine templates to provision a new preconfigured Enterprise Vault system. In native environments, this process can consume significant time, requiring you to procure hardware and install the operating system. Cloning ensures a controlled virtual machine configuration so deployment is less error prone and less time-consuming.

VMware vCenter Lab Manager

Patching and upgrading Enterprise Vault can be a time-consuming process. Most environments change control procedures require that any hotfix or patch go through some form of testing before production deployment. With a multi-tiered system like Enterprise Vault this requires a separate lab environment that might not exactly mimic your production environment. This can result in flawed test scenarios which can lead to failed upgrades to production and extended downtime. VMware vCenter Lab Manager can streamline the testing of configuration changes, patches, or upgrades to your Enterprise Vault infrastructure. When you need to make changes to the production Enterprise Vault systems, Lab Manager allows you to take a clone of the current environment and apply the changes to an identically configured, running test bed to validate the installation.

The Enterprise Vault environment clone is an exact replica of the production system, including all network settings, host names, and IP addresses. VMware vCenter Lab Manager deploys the clone in a fenced network to prevent network collisions. The fenced networking feature allows simultaneous deployment of multiple instances of the exact same Enterprise Vault configuration, which allows multiple teams to work in parallel without interrupting or conflicting with one another. After the patches or upgrades have been applied and validated, the same upgrade procedure can be reproduced in the production environment.

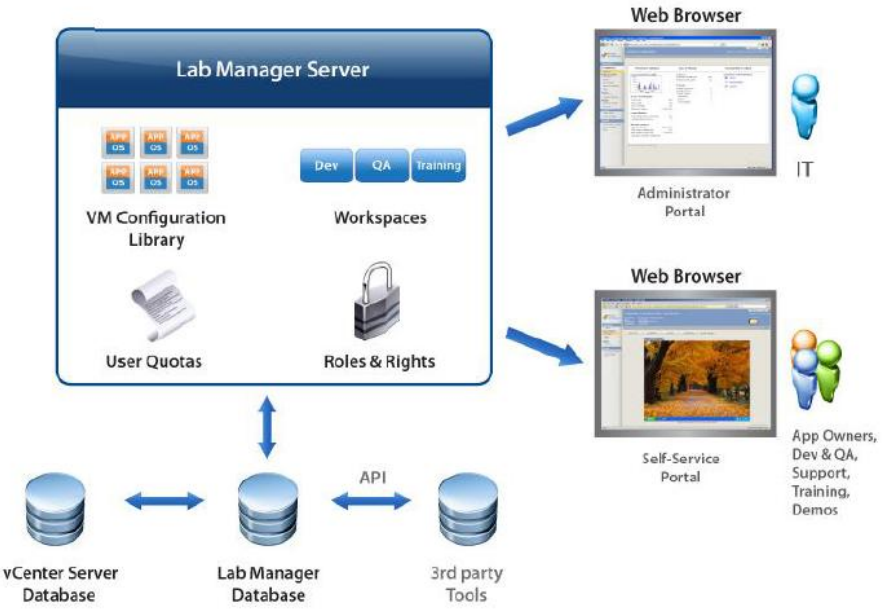

VMware vCenter CapacityIQ

Many organizations deploy tools to monitor Enterprise Vault and the underlying OS and hardware for faults and capacity warnings. Standard monitoring tools from Microsoft and other third-party vendors do a great job of providing application and OS level details. Unfortunately, when virtualized, these tools do very little to monitor the capacity of the underlying environment. Often Enterprise Vault servers are built with a fixed amount of memory and CPU, based on the recommendations from Symantec. While deficiencies in this design can quickly be observed by the users, over-provisioning is usually overlooked. Application owners usually do not complain that their application has more resources than it really requires, however, this can lead to inefficient use of the resources and lower the possible levels of consolidation.

VMware vCenter CapacityIQ brings vSphere-aware capacity monitoring and reporting to vCenter. By monitoring the usage and performance characteristics of virtual machines, CapacityIQ can provide forecasting of resource consumption over time. In many cases, development and test environments are built to mimic production at the request of developers. CapacityIQ can provide usage details to help you determine if these virtual machines are over-provisioned and help reclaim unused capacity.

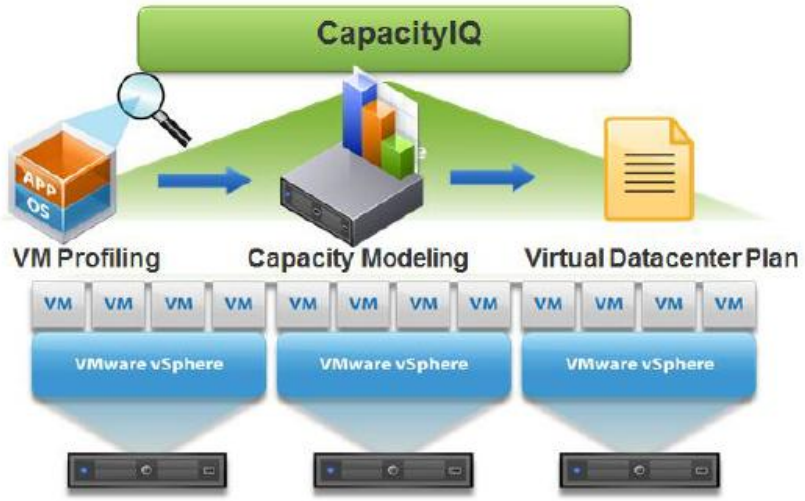

VMware vCenter Site Recovery Manager

VMware vCenter Site Recovery Manager (SRM) takes advantage of virtual machine encapsulation to make testing and initiating DR fail-over a simple, integrated vCenter process. vCenter SRM runs alongside VMware vCenter Server to provide planning, testing and automated recovery in the case of a disaster. By using replication technology from leading storage vendors vCenter SRM eliminates the manual steps required during a fail-over scenario to ensure consistent and predictable results each time. At a high-level the steps that can be performed during a fail-over test or actual run are as follows:

- Shutdown production virtual machines (fail-over)
- Promote recovery storage to primary (fail-over)
- Take and mount snapshot of recovery storage in read/write mode (test only)
- Rescan ESX hosts to make storage visible
- Register recovery virtual machines
- Power-on virtual machines at recovery site
- Reconfigure IP settings and update DNS if required
- Verify VMware tools starts successfully on recovered virtual machines
- Power-off recovered virtual machines (test only)
- Un-register virtual machines (test only)
- Remove storage snapshot from recovery side (test only)

SRM provides integration of the storage replication solution, VMware vSphere and customer-developed scripts to ensure a simple, repeatable and reportable process for disaster recovery of the entire virtual environment, regardless of the application.

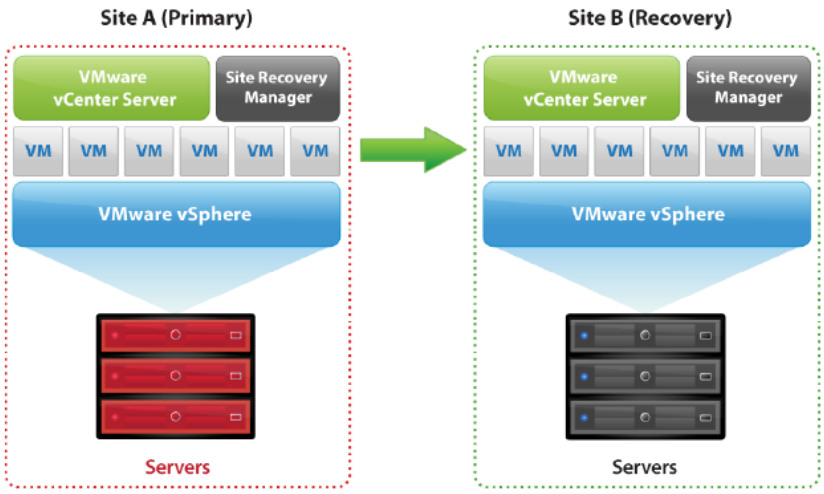

About Symantec:

Symantec is a global leader in providing storage, security and systems management solutions to help consumers and organizations secure and manage their information-driven world.

Our software and services protect against more risks at more points, more completely and efficiently, enabling confidence wherever information is used or stored.

For specific country offices and contact numbers, please visit our Web site: **www.symantec.com**

Symantec Corporation
World Headquarters
350 Ellis Street
Mountain View, CA 94043 USA
+1 (650) 527 8000
+1 (800) 721 3934