WHITE PAPER

Best Practice Guidelines for SAP Solutions on VMware® Infrastructure

vmware®

Best Practice Guidelines for SAP Solutions on VMware® Infrastructure

Vas Mitra
Solutions Engineer, Alliances
VMware

Note from the Author

The primary audience for this document is SAP customers who are considering or at the beginning phase of deploying SAP solutions on VMware® infrastructure.

Contents

1 Introduction

This paper provides best practice guidelines for deploying SAP software solutions on VMware Infrastructure 3. These recommendations are generic and are not specific to any particular SAP application or to the size and scope of the SAP implementation. The examples and considerations in this document provide guidance only and do not represent strict design requirements as the flexibility of SAP and VMware products allows for multiple variations on valid configurations.

This guide assumes a basic knowledge and understanding of VMware Infrastructure 3 and SAP Basis architecture. It discusses CPU, memory, networking and storage with respect to their impact on SAP installations that run on VMware Infrastructure.

VMware has created best practice documents in the areas of storage, networking and performance (See Appendix 1 for a list of these publications). The guidelines covered in these resources are broadly applicable to SAP environments; in most cases this paper brings these concepts into the SAP context. Refer to these resources for more background.

1.1 Virtualization of SAP Environments

SAP deployments can generate significant server sprawl primarily due to the need to provision separate systems for development (DEV), quality assurance/test (QAS) and production (PRD) environments. Each SAP solution typically has its own system landscape containing at minimum a DEV, QAS and PRD environment. SAP Enterprise environments can also contain multiple layers in the application architecture, including database, application server, and web server layers. In a typical deployment every layer of the environment is hosted on dedicated physical systems that are not fully utilized at all times. VMware virtualization technology contains this server sprawl by running SAP application layers in virtual machines consolidated onto fewer enterprise-class servers, resulting in increased server utilization.

1.2 SAP Platform Overview

SAP ERP is the SAP flagship product. In addition to ERP software, other key SAP products and solutions include business intelligence, customer relationship management, supply chain management, supplier relationship management, human resource management, product life cycle management, enterprise portal software, and knowledge warehouse software.

SAP applications are based on the SAP NetWeaver application and integration platform. SAP enterprise applications can be deployed in a two or three-tier architecture. The three-tier client/server architecture generally consists of a presentation layer, an application layer, and a database layer. These three layers can run separately on different computers or all together on the same computer, depending on the requirements and size of the SAP solution being deployed. In three-tier configurations the database and application services reside on separate OS images whereas in two-tier configurations they co-exist on the same OS image. The three-tier architecture scales to support large numbers of users. The two-tier architecture is usually sufficient for many smaller and midsize companies, as well as for sandbox, development, training and test systems.

The SAP application layer contains the following components:

- Central services, which include messaging (manages client connections and communications) and enqueue (SAP lock management) services.

- Application services that process online and batch workloads, which can be further categorized into the following types:
  - ABAP: processes workloads based on an SAP proprietary programming language.
  - Java: processes Java based workloads.
  - ABAP + Java: processes both ABAP and Java based workloads.

  The specific SAP product determines the type of application service required (ABAP, Java or both).
- SAP enables a distributed architecture that allows services to be centralized or distributed across separate instances or servers, depending on sizing and availability requirements. An instance that includes both central and application services is referred to as the Central Instance (CI). An instance that includes only application services is called a Dialog instance.

Figure 1-1 illustrates the architecture of the SAP NetWeaver application server which includes both the ABAP and Java stack. Note that the two stacks can be installed separately or together.

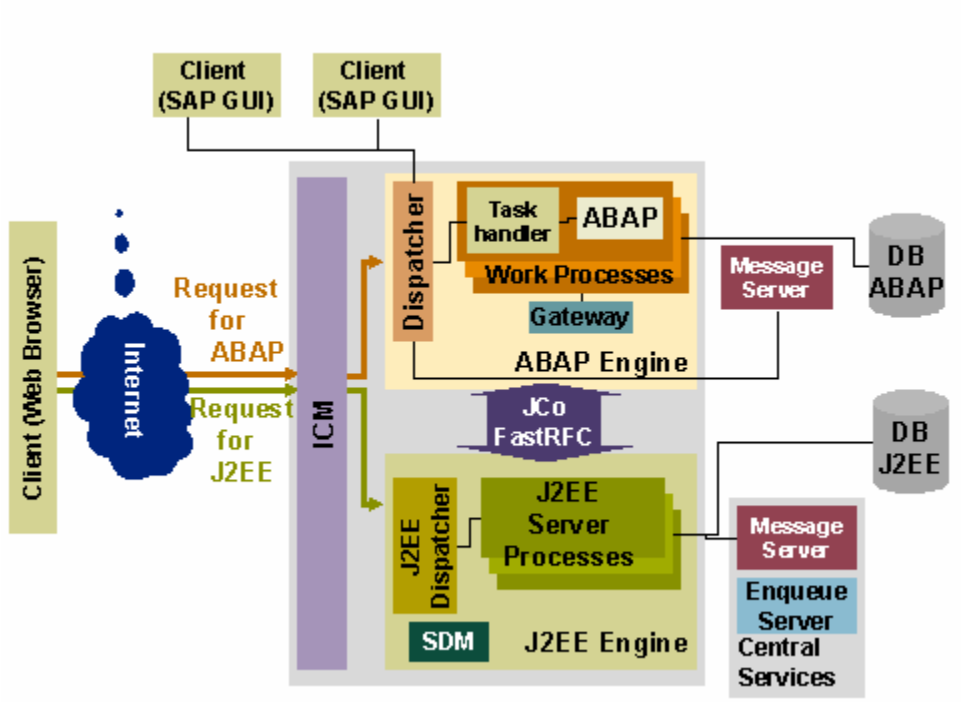

Figure 1-1. Architecture of SAP NetWeaver Application Server

Key components in Figure 1-1 include:

- The Internet Communication Manager (ICM) handles communication between the SAP Application Server and the outside world via HTTP, HTTPS and SMTP protocols.
- Central Services comprises message and enqueue services. There are separate central services for the ABAP and Java stack (hence Figure 1-1 shows two message servers, while the second enqueue server for the ABAP stack is not shown).
- The Dispatcher accepts requests from the ICM or external SAP clients and places them into a queue from which it distributes requests to other SAP processes to execute application logic.

1.3 VMware® Infrastructure 3 Software Suite

The VMware Infrastructure 3 software suite virtualizes servers, storage, and networking, allowing multiple unmodified operating systems and their applications to run independently in virtual machines while sharing physical resources. VMware Infrastructure 3 is available in Foundation, Standard, and Enterprise editions. The Enterprise edition includes the following components shown in Figure 1-1: VMware® ESX Server; VMware® Virtual Machine File System (VMFS), VMware® Virtual SMP (Virtual SMP), VMware® VirtualCenter; the VMware® Infrastructure (VI) Client, VMware® Distributed Resource Scheduler (DRS), VMware® High Availability (HA), and VMware® Consolidated Backup.

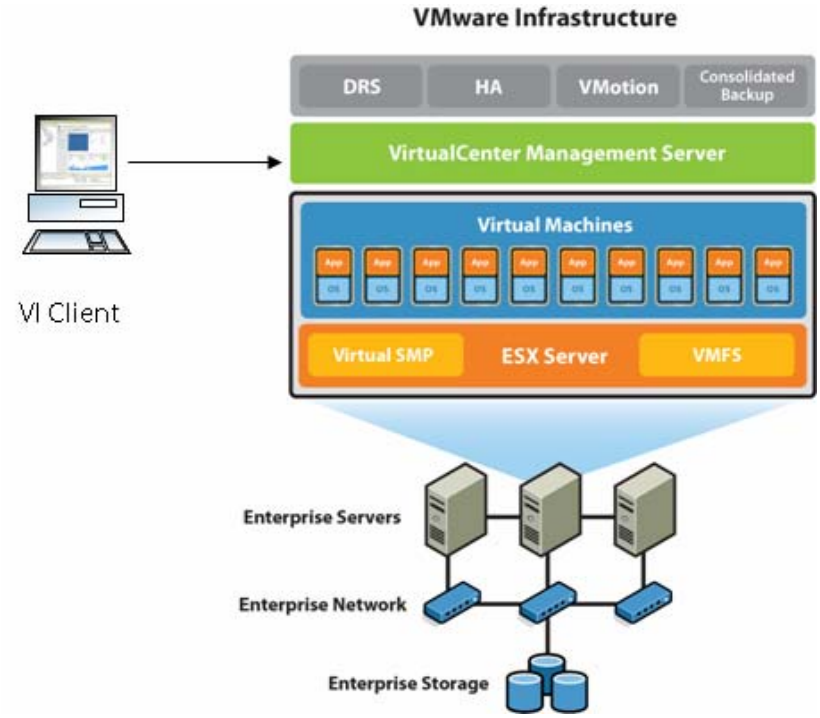

Figure 1-2. VMware Infrastructure 3

VMware ESX Server

ESX Server is the bare-metal hypervisor that is installed directly on the server and abstracts processor, memory, storage, and networking resources into multiple virtual machines, giving IT greater hardware utilization and flexibility. Figure 1-3 depicts the architecture of the ESX Server hypervisor.

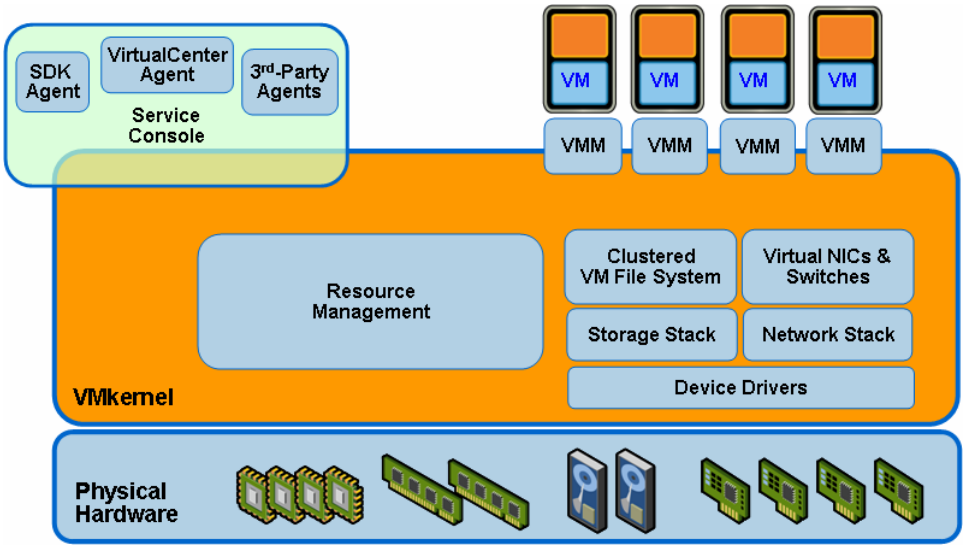

Figure 1-3. VMware Infrastructure 3 - ESX Server Architecture

The VMkernel is the proprietary kernel developed by VMware that runs directly on the ESX Server host. VMkernel controls and manages most of the physical resources on the hardware, including memory, physical processors, storage controller,; networking, keyboard, video, and mouse. The VMkernel includes schedulers for CPU, memory, and disk access, and has full-fledged storage and network stacks.

The virtual machine monitor (VMM) is the component actually responsible for virtualizing the CPUs. When a virtual machine is turned on, control transfers to the virtual machine monitor, which begins executing instructions from the virtual machine. The transfer of control to the virtual machine monitor involves setting the system state so that the virtual machine monitor runs directly on the hardware.

The ESX Server service console provides an execution environment to monitor and administer the entire ESX Server host. The service console operating system is a reduced version of Red Hat Enterprise Linux.

VMware VMFS
A high•performance cluster file system that provides storage virtualization optimized for virtual machines. Each virtual machine is encapsulated in a small set of files that use VMFS as the default storage system on physical SCSI disks and partitions.

Virtual SMP
A product that supports symmetric multiprocessing. Virtual SMP enables you to assign multiple virtual processors to virtual machines on multiprocessor host machines.

VMware VirtualCenter
A software solution for deploying and managing virtual machines across the datacenter. With VirtualCenter, datacenters can instantly provision servers, globally manage resources, and eliminate scheduled downtime for hardware maintenance. VirtualCenter performance monitoring of physical servers and virtual machines, intelligently optimizes resources, ensures high availability

to all applications in virtual machines, and makes IT environments more responsive with distributed services such as VMware® High Availability (HA) and VMware® VMotion.

VI Client:
A user interface that runs locally on a Windows computer and provides access to the virtual machine's display. The VI Client runs on a networked machine that does not need to be the same machine as the VirtualCenter Server. The VI Client requires a monitor for access to the virtual machine's display.

VMware DRS:
A product that dynamically allocates and balances computing capacity across a collection of hardware resources aggregated into logical resource pools. DRS continuously monitors use across resource pools and allocates available resources among the virtual machines based on predefined rules that reflect business needs and changing priorities. When a virtual machine experiences an increased load, VMware DRS automatically allocates additional resources by redistributing virtual machines among the physical servers in the resource pool.

VMware HA:
An optional feature that supports distributed availability services in an environment that includes ESX Server and VirtualCenter. If DRS is configured and one of the hosts that the VirtualCenter Server manages becomes unavailable, all virtual machines on that host are immediately restarted on another host.

VMware Consolidated Backup:
Software that enables virtual machines to be backed up from a centralized Microsoft Windows 2003 proxy server rather than directly from ESX Server.

2 SAP Software Benchmark Certifications

SAP and its hardware partners have developed the SAP Standard Application Benchmarks to test the hardware and database performance of SAP applications and components. You can see certified benchmarks for the SAP core enterprise product, SAP ECC 6.0, running on ESX Server on the Web here: http://www.sap.com/solutions/benchmark/sd2tier.epx.

SAP states that "…benchmarks provide basic sizing recommendations to customers by placing a substantial load upon a system…."[1] These results should be used as a guideline only, as they are based on greater than 90 percent CPU utilization, very near two second response time, and do not factor for batch and business specific requirements. Work with your infrastructure vendors to develop a detailed sizing and architecture plan designed for your individual situation.

Table 2-1 summarizes the benchmark certifications that exist for SAP ECC 6.0 running on ESX Server software completed between May and November 2007. These certifications can be found at http://www.sap.com/solutions/benchmark/sd2tier.epx. Any future benchmark certifications will also be listed on this site.

Table 2-1. SAP Software on ESX Server Benchmark Certifications

Date	Hardware Vendor	Server	VM Configuration	No. SD Users	SAPs[2]
11/22/07	Sun	X2200 M2	2 x vCPU Windows Server 2003 / SQL Server	380	2020
10/22/07	Fujitsu Siemens	PRIMERGY Model TX600 S3 | RX600 S3	2 x vCPU Windows Server 2003 / SQL Server	250	1250
7/30/07	Dell	PowerEdge 2900	2 x vCPU, Windows Server 2003 / SQL Server	421	2130
7/12/2007	IBM	System x3755	2 x vCPU, Windows Server 2003 / DB2 9	445	2230
7/12/2007	Fujitsu Siemens	PRIMERGY Model TX300 S3 | RX300 S3	2 x vCPU, Windows Server 2003 / SQL Server	420	2150
5/11/2007	HP	ProLiant BL460c	2 x vCPU, Windows Server 2003 / SQL Server	402	2070

[1] This statement can be found here: http://www.sap.com/solutions/benchmark/benefits/index.epx:
[2] SAPS refers to the SAP Application Performance Standard and is a hardware-independent unit that describes the performance of a system configuration in the SAP environment.

3 Sample High-Level Configurations

Virtualization with VMware technology adds the components of virtual machines and virtual CPU (vCPU) to the solution infrastructure. This section contains example configurations that illustrate how these components contribute to the overall architecture of virtualized SAP solution-based environments. The subsequent chapters cover in more depth the architectural design considerations for CPU, memory, network and storage in a virtual SAP environment.

For the sake of brevity, VMware uses the following acronyms in this document:

- DB: SAP database
- CI: Central Instance - performs central functions (SAP messaging and locking) and online and batch workloads.
- DIA: Dialog instance – performs online and batch workloads.
- ERP: refers to the SAP core enterprise product (e.g. ECC 6.0).
- BW: refers to the SAP data warehousing product.

Note: The virtual machine configurations illustrated in this section include details of virtual CPUs (VCPUs) and memory settings. These details are provided as examples and are not based on any implementation or sizing requirements.

3.1 Two-Tier and Three-Tier Configurations

Figure 3-1 represents a typical small implementation of two-tier and three-tier SAP components in virtual machines. VMware virtualization does not introduce any changes to the installation and technical configuration of SAP software. The application can be scaled and distributed across virtual machines in a similar manner to physical servers with the following criteria which are based on the current configuration limitations of virtual machines on ESX Server 3.0.X:

- Number of vCPUs = 1,2 or 4
- Maximum memory for virtual machine = 16 GB.

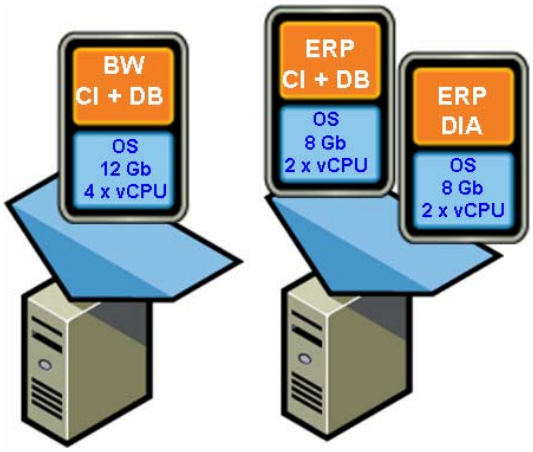

Figure 3-1. Examples of Two-Tier and Three-Tier Implementations

3.2 Three-Tier Mixed Native + Virtual Environments

Figure 3-2 illustrates an architecture in which the SAP application tier is virtualized while database and CI services remain on native hardware. This architecture has the advantages of scaling to large

number of users and enables easier adoption of virtualization in environments where significant investment has already been made in high RAS (Reliability, Availability, Serviceability) installations of the database and CI. Meanwhile the VMware virtualization benefits of flexibility and consolidation are available to the SAP application tier.

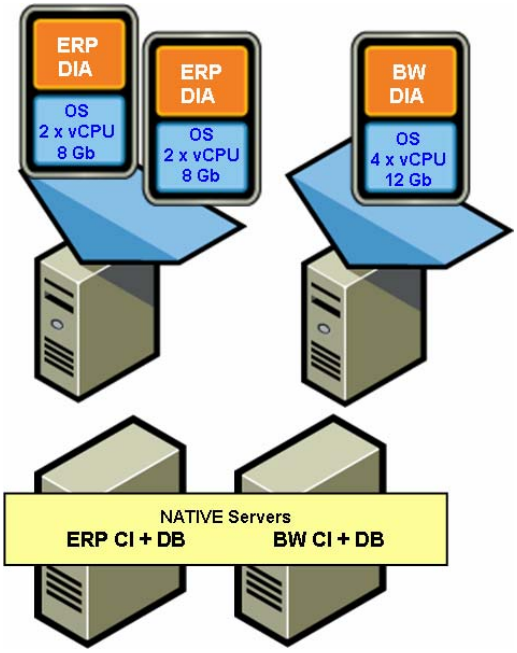

Figure 3-2. Example of Mixed Three-Tier Native and Virtual Implementation

4 Physical and Virtual CPUs

VMware uses the terms virtual CPU (vCPU) and physical CPU to distinguish between the processors within the virtual machine and the underlying physical x86-based processors. Virtual machines with more than one virtual CPU are also called SMP (symmetric multi-processing) virtual machines. The virtual machine monitor (VMM) is responsible for virtualizing the CPUs. When a virtual machine starts running, control transfers to the VMM which begins executing instructions from the virtual machine.

4.1 Virtual SMP

VMware Virtual Symmetric Multi-Processing (Virtual SMP) enhances virtual machine performance by enabling a single virtual machine to use multiple physical processors simultaneously. The biggest advantage of an SMP system is the ability to use multiple processors to execute multiple tasks concurrently, thereby increasing throughput (for example, the number of transactions per second). Only workloads that support parallelization (including multiple processes or multiple threads that can run in parallel) can benefit from SMP. SAP architecture is multi-threaded (NetWeaver JAVA stack) and includes multiple processes (NetWeaver ABAP stack) which makes it a perfect candidate to take advantage of Virtual SMP.

The virtual processors from SMP-enabled virtual machines are co-scheduled. That is, if physical processors are available, the virtual processors are mapped one-to-one onto physical processors and are then run simultaneously. In other words, if one vCPU in the virtual machine is running, a second vCPU is co-scheduled so that they execute nearly synchronously. You should take into consideration the following points when using multiple vCPUs:

- If multiple idle physical CPUs are not available when the virtual machine wants to run, the virtual machine will be in a special wait state. The time a virtual machine spends in this wait state is called "ready time."
- Even idle processors perform a limited amount of work in an operating system. In addition to this minimal amount, ESX Server manages these "idle" processors, resulting in some additional work by the hypervisor. These low-utilization vCPUs compete with other vCPUs for system resources.

As a result, VMware recommends the following practices:

- Only allocate multiple vCPUs to a virtual machine if the anticipated SAP workload can truly take advantage of all the vCPUs.
- If the exact workload is not known, size the virtual machine with a smaller number of vCPUs initially and increase the number later if necessary.
- For performance-critical SAP virtual machines (e.g. production systems), try to ensure the total number of vCPUs assigned to all the virtual machines is equal to or less than the total number of cores on the ESX Server host machine.

4.2 Service Console

The ESX Server service console provides an execution environment to monitor and administer the entire ESX Server host. The service console operating system (COS) is a reduced version of Red Hat Enterprise Linux. The COS is scheduled on CPU 0 for ESX Server version 3.0.x. Under high loads the COS will compete with virtual machines that are scheduled on CPU 0. VMware recommends you avoid running programs in the COS. The elimination of COS in ESX Server 3i removes this contention.

5 Memory Configuration Guidelines

This section provides guidelines for configuring the memory requirements of multiple SAP virtual machines on a single server in order to maximize performance. The guidelines given here take into account ESX Server memory overhead and the virtual machine memory settings.

ESX Server virtualizes guest physical memory by adding an extra level of address translation. Shadow page tables make it possible to provide this additional translation with little or no overhead. Managing memory in the hypervisor enables the following features:

- Memory sharing across virtual machines that have similar data e.g. same guest operating systems
- Memory over-commitment, which means allocating more memory to virtual machines than is physically available on the ESX Server host
- A memory balloon technique wherein virtual machines that do not need all they have been allocated give memory to virtual machines that are using all of their allocated memory.

For more details on ESX Server memory management concepts please consult the VMware Resource Management Guide (http://www.vmware.com/pdf/vi3_301_201_resource_mgmt.pdf).

5.1 Memory Overhead

ESX Server requires three areas of memory overhead:

- A fixed system-wide overhead for the service console – about 272 Mb for ESX 3.x.
- A fixed system-wide overhead for the VMkernel – depends on number and size of device drivers.
- Additional overhead for each virtual machine. The virtual machine monitor for each virtual machine requires some memory for its code and data.

You can view the fixed system-wide overhead for any ESX Server instance in the VI Client. Select the configuration tab under memory settings. Figure 5-1 shows an example.

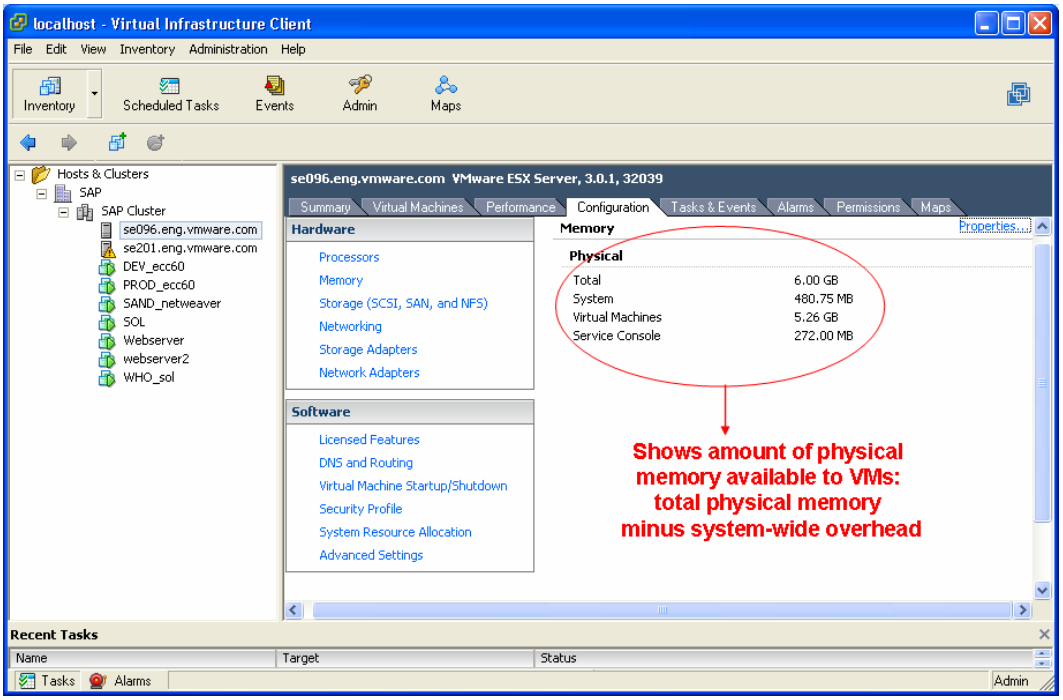

Figure 5-1. VI Client Screenshot Showing System-Wide Memory Overhead

The overhead for each virtual machine, defined in Table 5-1, is dependent on the number of vCPUs, the amount of memory allocated to the virtual machine and whether the virtual machine is 32 or 64 bit.

Table 5-1. Virtual Machine Memory Overhead[3]

Virtual CPUs	Memory (MB)	Overhead for 32-bit virtual machine (MB)	Overhead for 64-bit virtual machine (MB)
1	256	79	174
1	512	79	176
1	1024	84	180
1	2048	91	188
1	4096	107	204
1	8192	139	236
1	16384	203	300
2	256	97	288
2	512	101	292
2	1024	101	300
2	2048	125	316
2	4096	157	349
2	8192	221	413
2	16384	349	541
4	256	129	511
4	512	133	515
4	1024	141	523
4	2048	157	540
4	4096	189	572
4	8192	222	605
4	16384	350	734

5.2 Virtual Machine Memory Settings

The ESX Server memory settings for a virtual machine include the following parameters:

- The configured size, which is the memory size specified at creation of the virtual machine.
- The reservation, which is a guaranteed lower bound on the amount of memory that the host reserves for the virtual machine and cannot be reclaimed by ESX Server for other virtual machines.

[3] (Source: *VMware Resource Management Guide*. Note second column refers to memory allocated to VM)

Figure 5-2 illustrates the memory settings of a virtual machine.

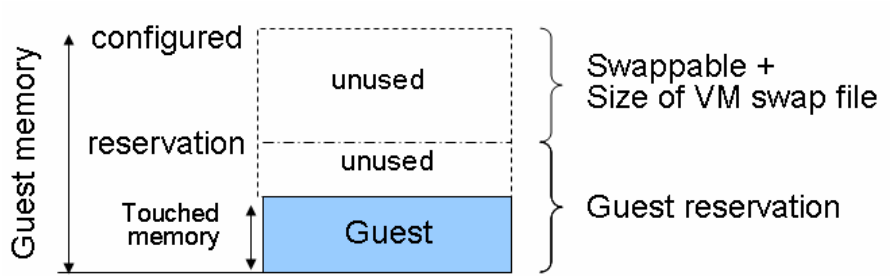

Figure 5-2. Virtual Machine Memory Settings

The terms in Figure 5-2 are defined here:

- Configured memory = memory size of virtual machine assigned at creation.
- Touched memory = memory actually used by the virtual machine. Guest memory is only allocated on demand by ESX Server.
- Swappable = virtual machine memory that can be reclaimed by the balloon driver or worst case by ESX Server swapping. Ballooning will occur before ESX Server swapping. If this memory is in use by the virtual machine (i.e. touched and in use) the balloon driver will cause the guest operating system to swap. Also this value is the size of the per-virtual machine swap file that is created on the VMFS file system (".vswp" file).
- If the balloon driver is unable to reclaim memory quickly enough or it is not installed or is disabled ESX Server will forcibly reclaim memory from the virtual machine using the vmkernel swap file (last resort).

5.3 Sizing Memory of Multiple SAP Virtual Machines

Based on the memory overhead defined above, the available memory for SAP applications (and the guest operating systems) in virtual machines is:

- Available physical memory for SAP applications in virtual machines = total ESX Server host physical memory minus system-wide memory overhead, virtual machine memory overhead, and a user-defined "memory buffer".

Note: the "memory buffer" is not a VMware parameter but a user-defined value designed to provide headroom and flexibility to manage more virtual machines than the initial estimate; for example, virtual machines migrated using VMotion onto the system from another ESX Server host machine. *This will depend on customer specific design requirements.*

ESX Server page sharing makes more physical memory available but this additional memory is not counted here in order to provide a more conservative estimate.

Generally, as SAP applications are memory intensive and for situations where performance is a key factor (e.g. in production environments) VMware recommends the following practices:

- Do not over-commit memory on the ESX Server host until VirtualCenter reports that steady state memory usage is below the amount of memory on the server.
- Set the memory reservation to the configured size of the virtual machine, which results in a per-virtual machine vmkernel swap file of zero bytes. The guest operating system within the virtual machine will still have its own separate swap/page file.
- Do not disable the balloon driver (which is installed with VMware® Tools – see section 8).

- Size the memory of each virtual machine so that the total configured memory for all the virtual machines is equal or less than the memory available for SAP applications calculated above.

To minimize guest operating system swapping, the configured memory size of the virtual machine should be greater than the average memory usage of the SAP application running in the guest. If the SAP application in the virtual machine needs more memory than it has been allocated, the guest operating system paging/swapping mechanisms will be invoked as they are in normal native operations. Memory and swap/page file configuration of the SAP application in the virtual machine follows the same guidelines as for native environments and generally these should be set to minimize any guest operating system swapping. Existing SAP documentation and recommendations should be followed, as provided in the SAP notes identified here:

- 88416 – Zero Administration Memory Management as of 4.0A/Windows
- 1009493 – abap/heap_area* parameter Defaults Changed (64-Bit Windows)
- 723909 – Java virtual machinesettings for J2EE 6.40/7.0
- 941735 - SAP memory management for 64-bit Linux systems

NUMA Systems

NUMA (non-uniform memory access) memory architecture links several small, cost effective nodes by way of a high performance connection. Each node contains both processors and memory. An advanced memory controller allows a node to use memory on all other nodes, creating a single system image. When a processor accesses memory that does not lie within its own node (remote memory), the data must be transferred over the NUMA connection, which is slower than accessing local memory. Thus, memory access times are non uniform depending on the location of the memory. As a result, increasing the memory of virtual machines on a NUMA server may eventually force some memory to be allocated from a remote node, which will decrease performance.

ESX Server uses its own NUMA scheduler to dynamically balance processor load and memory locality or processor load balance. For more details on using NUMA systems with ESX Server please refer to Appendix A of the *VMware Resource Management Guide*.

Follow these practices when you run SAP virtual machines on NUMA based systems:

- Whenever possible, size the virtual machine so both vCPU and memory fit on one node.
- Although it is possible via the VI Client to set CPU and NUMA node affinities to force a virtual machine to run on a specific NUMA node, this is generally not a recommended practice. The preferred practice is to let the ESX Server NUMA scheduler automatically manage the virtual machines.
- Use ESX Server 3.5, the most recent version, which includes various NUMA scheduling improvements.

6 Network Design and Configuration Guidelines

This section covers design guidelines for the virtual networking environment and provides configuration examples at the ESX Server host level for SAP installations. **NOTE: The examples do not represent strict design requirements and do not cover all possible SAP network design scenarios.**

The virtual networking layer consists of the virtual network devices through which virtual machines and the service console interface with the rest of the network and users. In addition, ESX Server hosts use the virtual networking layer to communicate with iSCSI SANs and NAS storage.

The virtual networking layer includes virtual network adapters and the virtual switches. Virtual switches are the key networking components in VMware Infrastructure 3. They are "built to order" at run time, are implemented in much the same way as a modern Ethernet switch, and support functions equivalent to VLANs based on the IEEE 802.1Q protocol.

6.1 SAP Networking

The three-tier client/server architecture of the SAP NetWeaver platform includes two separate traffic and network requirements for SAP application and database servers:

1. between users and the application servers
2. between the application and database servers

The figure below identifies the two networks. (Other tiers and networks exist for web servers and DMZ setups, but for simplicity these are not shown).

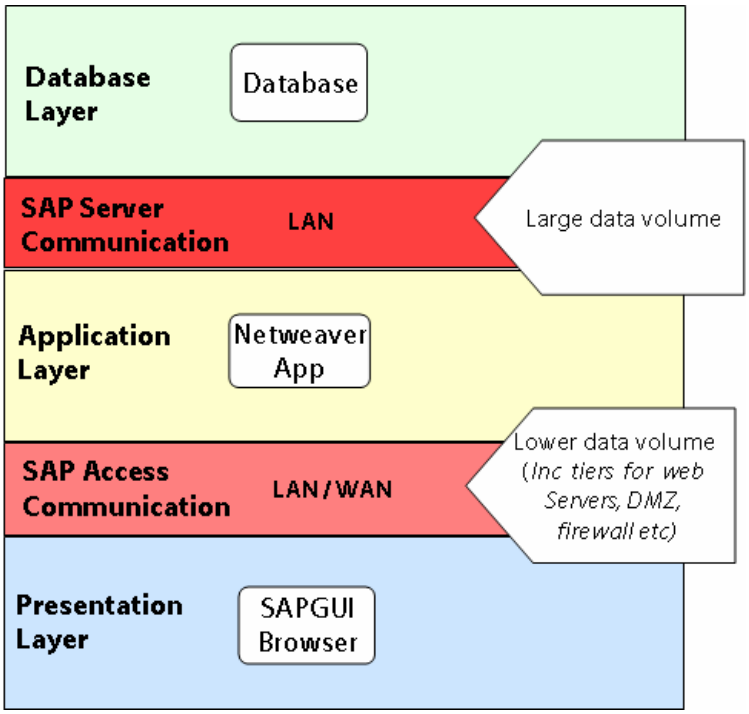

Figure 6-1. SAP Three-Tier Networking Architecture Requirements

You can design the network in two ways:

- One network for all SAP servers and user access communication (viable for smaller SAP implementations)
- Two separate networks (typical for larger installations):
  - Private: for SAP server traffic (called private as there is no access from end users).
  - Public: for user access communication.

While these requirements are typically factored into the network design of non-virtualized SAP environments, they equally apply to ESX Server hosts running SAP virtual machines. Figure 6-2 depicts the two-network architecture for ESX Server hosts running SAP applications.

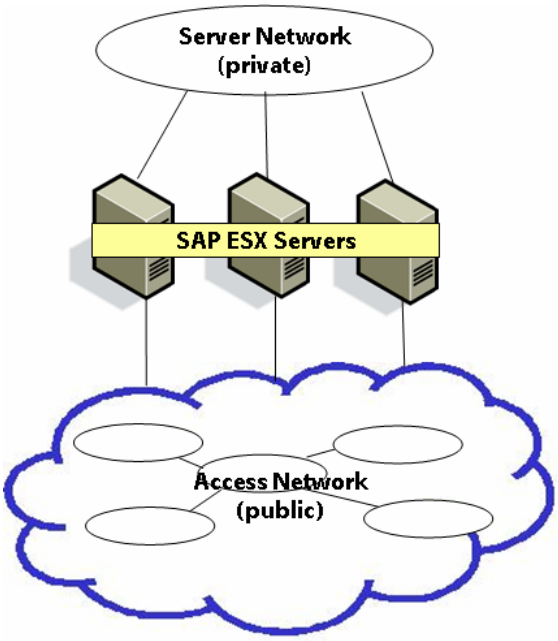

Figure 6-2. Private and Public Networks for SAP ESX Servers

For more details on networking architecture for SAP environments please consult the SAP paper *Network Integration for SAP Servers* which is attached to *SAP note 21151 – Multiple Network adapters in SAP Servers*.

6.2 Virtual Networking Concepts

Virtual Switches

Virtual switches work like Ethernet switches and support VLAN segmentation at the port level. VLANs in ESX Server allow logical groupings of switch ports to communicate as if all ports were on the same physical LAN segment. VLANs require tagging of Ethernet frames with the 802.1Q tag (based on IEEE protocol standards) and ESX Server enables port-based VLAN tagging based on the switch ports. The VMware Virtual Networking Concepts paper discusses three different configuration modes to tag (see Appendix A for URL):

- Virtual Switch Tagging (VST mode): virtual switch port group adds and removes tags.
- Virtual Machine Guest Tagging (VGT mode): an 802.1Q VLAN trunking driver is installed in the virtual machine.

- External Switch Tagging (EST mode): external switches perform VLAN tagging so Ethernet frames moving in and out of the ESX Server host are not tagged with VLAN IDs.

The most common configuration, shown here, is VST mode. VST mode requires provisioning one port group on a virtual switch for each VLAN and attaching the virtual machine's virtual adapter to the port group of the virtual switch. The virtual switch port group tags all outbound frames and removes tags for all inbound frames. It also ensures that frames on one VLAN are isolated from other VLANs. VST mode requires that the physical switch provide a trunk (trunking is the technology that allows information from multiple VLANs to be carried over a single link).

Port groups in ESX Server are templates for creating virtual ports with a particular set of specifications. In ESX Server, there are three types of port group / virtual switch connections:

- Service console port group: ESX Server management interface
- VMkernel port group: VMotion, iSCSI and/or NFS/NAS networks
- Virtual machine port group: virtual machine networks

More than one connection type can exist on a single virtual switch, or each connection type can exist on its own virtual switch.

NIC Teaming

A single virtual switch can be connected to multiple physical Ethernet adapters using the VMware Infrastructure feature called NIC teaming. This provides redundancy and/or aggregation.

6.3 SAP Virtual Network Configurations

Network traffic is classified into the following types:

- SAP virtual machine traffic: data sourced and received from virtual machines. This traffic can be organized in two ways and is dependent upon the size of the SAP implementation (as explained above):
    - Two networks connected to each SAP virtual machine: traffic from users to application virtual machines (public network); database traffic between application and database virtual machines (private network).
    - One network: communication between all SAP virtual machines and front-end users.
- Administrative traffic:
    - VMotion traffic: data sent when moving a virtual machine from one ESX Server host to another. This traffic should be dedicated and isolated.
    - Management traffic: should be isolated from virtual machine traffic (note: management traffic includes heartbeats if VMware HA is enabled)
    - iSCSI traffic: if using the software iSCSI initiator

Given the traffic classifications and the virtual networking concepts discussed above, use these configuration guidelines when configuring virtual networks for SAP solution-based environments:

- Set up a minimum of four NIC ports per server (not including iSCSI)
- Use additional NICs if iSCSI is required .
- Create two virtual switches.
- Assign a pair of NIC ports to each virtual switch and configure NIC teaming for load balancing and failover.
- Use virtual switch tagging mode for enforcement of VLAN tagging at the virtual switch port.
- If separate SAP application / dialog and database virtual machines reside on the same ESX Server host they should be connected to the same virtual switch.

Figure 6-3 illustrates the example networking configuration of an ESX Server host machine based on these recommendations (iSCSI is not included in this case). The figure assumes a private and public network for SAP virtual machines as discussed above.

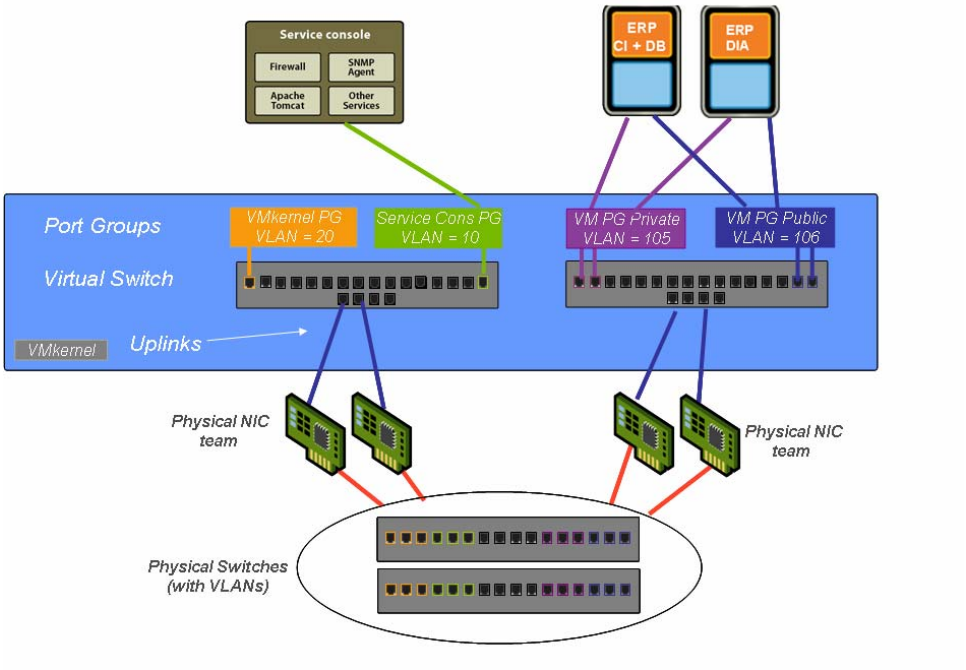

Figure 6-3. ESX Server Host System Virtual Networking Design Using VLAN Tagging and Four NIC Ports (includes redundancy)

Smaller Environments: two NIC ports per ESX Server host

For smaller environments where only two NIC ports are available on each ESX Server host machine (e.g. in a blade environment) and when only one network is required for the SAP virtual machines (the Private and Public networks are combined), follow these guidelines:

- Team both NICs.
- Create one virtual switch.
- Assign port groups for the system console and VMotion (assign first NIC as active).
- Assign a port group for SAP virtual machine traffic (assign second NIC as active).
- Use virtual switch tagging mode for enforcement of VLAN tagging at the virtual switch port.

Figure 6-4 illustrates this configuration.

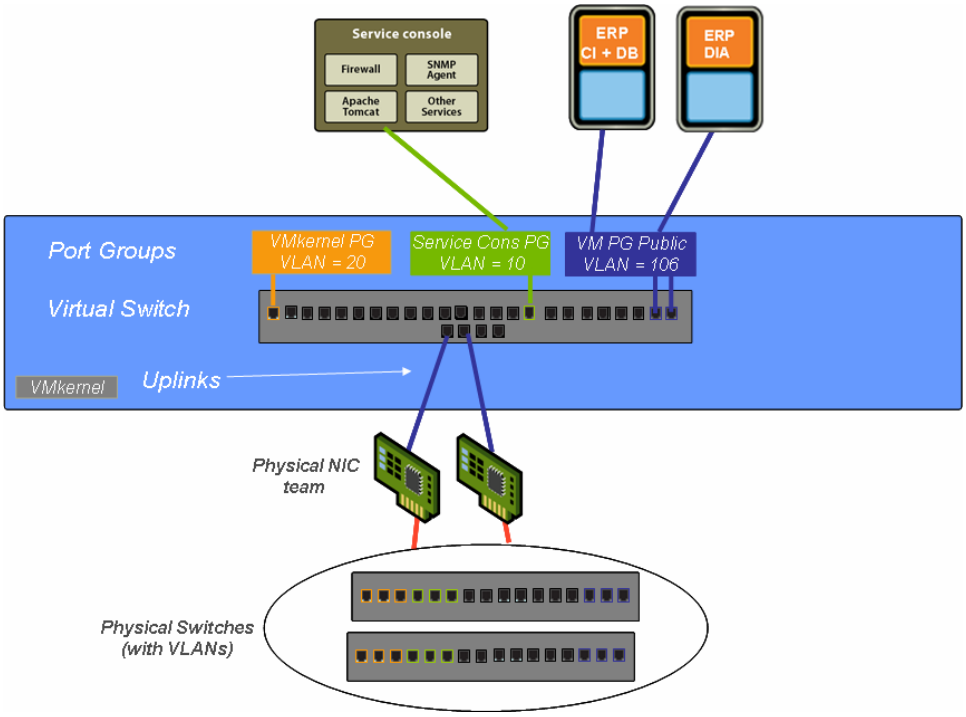

Figure 6-4. ESX Server Host System Virtual Networking Design Example using VLAN Tagging and Two NIC Ports (includes redundancy)

7 Storage Virtualization

VMware® Virtual Machine File System (VMFS) is a cluster file system that provides storage virtualization optimized for virtual machines. Each virtual machine is encapsulated in a small set of files and VMFS is the default storage system for these files on physical SCSI disks and partitions. VMware supports Fibre-Channel, iSCSI and NAS shared storage protocols.

It is preferred practice to deploy virtual machines files on shared storage to take advantage of VMotion, VMware HA and DRS. This practice aligns well with mission-critical SAP solution-based deployments, which are typically installed on third party shared storage management solutions.

VMware storage virtualization can be categorized into three layers of storage technology, as illustrated in Figure 7-1. The storage array is the bottom layer, consisting of physical disks that are presented as logical disks (storage array volumes or LUNs) to the layer above, the virtual environment occupied by ESX Server. Storage array LUNs are 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.

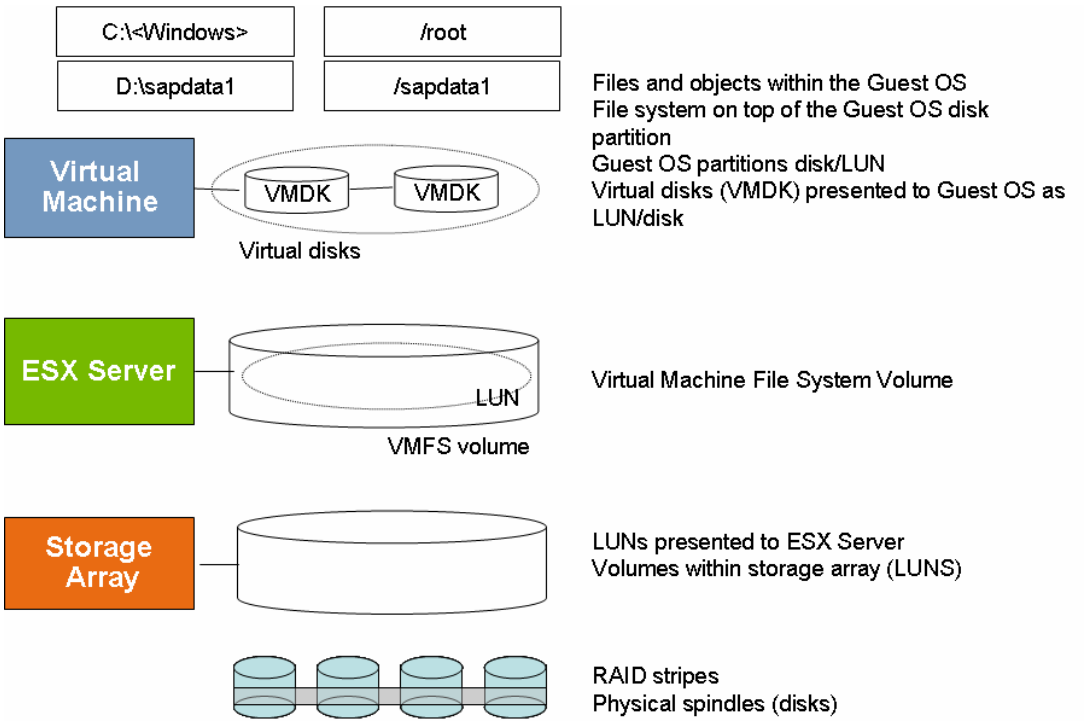

Figure 7-1. VMware Storage Virtualization Stack

7.1 Storage Multipathing

VMware recommends you set up a minimum of four paths from an ESX Server host to a storage array, which means the host requires at least two HBA ports. See Figure 7-2.

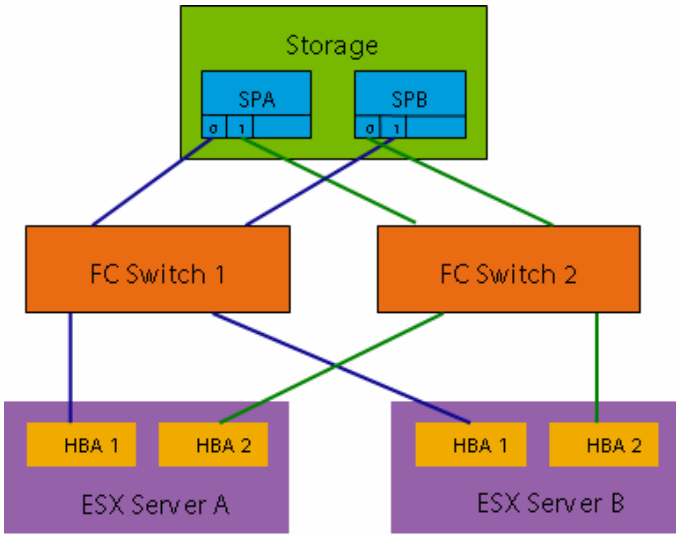

Figure 7-2. Storage Multipathing Requirements for ESX Server

The terms used in Figure 7-2 above are defined here:

- HBA: Host Bus Adapter, a device that connects one or more peripheral units to a computer and manages data storage and I/O processing.
- FC: Fibre Channel:, a gigabit-speed network technology used to build storage area networks and to transmit data.
- SP: storage processor, a SAN component that processes HBA requests routed through an FC switch and handles the RAID/volume functionality of the disk array.

7.2 Raw Device Mapping

VMFS also supports Raw Device Mapping (RDM). RDM provides a mechanism for a virtual machine to have direct access to a volume on the physical storage subsystem and can only be used with Fibre Channel or iSCSI. RDM can be thought of as providing a symbolic link from a VMFS volume to a raw volume. The mapping makes volumes appear as files in a VMFS volume. The mapping file, not the raw volume, is referenced in the virtual machine configuration.

There are no "hard-and-fast" recommendations to use VMFS or RDM for SAP deployments. The following table summarizes some of the options and trade-offs. For a more complete discussion please consult the VMware SAN System Design and Deployment Guide (URL in Appendix A).

Table 7-1: VMFS and Raw Disk Mapping Trade-offs

VMFS	RDM
• Volume can host many virtual machines (or can be dedicated to one virtual machine). • Increases storage utilization, provides better flexibility, easier administration and management.	• Maps a single LUN to one virtual machine so only one virtual machine is possible per LUN. • More LUNs are required, so it is easier to hit the LUN limit of 256 that can be presented to ESX Server.
• Does not support Quorum disks required for third party clustering software.	• Required for third party clustering software (e.g. MSCS). Cluster data and quorum disks should be configured with RDM.

An environment that mixes VMFS and RDM is practical for SAP environments:

- Where existing systems already make use of third party storage management software, RDM can be used to leverage practices based on these products such as storage-based backups to disk.
- RDM is required when using third party clustering software.
- RDM is required when migrating from a physical server to a virtual one., This allows data on the physical environment to be accessed just as they are, without the need for data copy, from the new virtual machine.
- Deploying multiple non-production SAP systems on VMFS facilitates easier management and administration of template cloning, snapshots and storage consolidation.
- A mixed storage configuration is viable for an SAP virtual machine: the guest operating system is installed with VMFS and the SAP database files with RDM. VMware template cloning can be used for the guest operating system and database files can be managed by third party storage management software.

7.3 Number of Virtual Machines per LUN

The number of virtual machines allocated to a VMFS LUN influences the final architecture. Figure 7-3 illustrates the concepts and highlights the differences between a one-to-one and many-to-one virtual machine to LUN assignment.

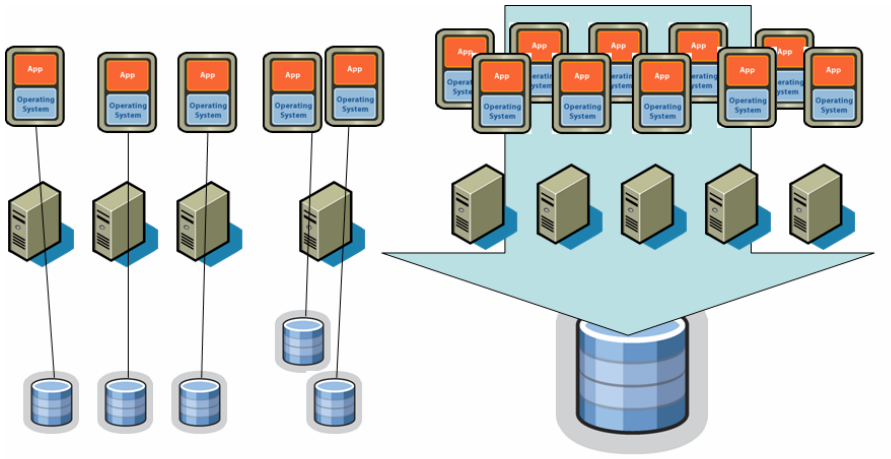

Figure 7-3. One versus Many Virtual Machines in a LUN

You would normally use a mixed approach for SAP deployments:

- Spread database files for performance-critical SAP systems into separate VMFS (or RDM) volumes for better I/O performance (see next section for more information).
- For systems that have less stringent performance service level requirements (e.g. non-production systems), place virtual machines on the same VMFS volume. Be aware, however, that heavily-used virtual machines accessing the same LUN may cause reduction in I/O performance.

Storage Layout for Performance-Critical SAP Systems

SAP and infrastructure vendors provide rules and recommendations for laying out file systems for the SAP database on shared storage in native environments. These same guidelines apply to SAP solutions that run on ESX Server.

The following example maps database layout guidelines from the publicly available whitepaper *SAP with Microsoft SQL Server 2005: Best Practices for High Availability, Maximum Performance, and Scalability SQL Server Technical Article* to ESX Server storage requirements (http://download.microsoft.com/download/d/9/4/d948f981-926e-40fa-a026-5bfcf076d9b9/SAP_SQL2005_Best%20Practices.doc). According to the Microsoft white paper, storage requirements for a mid size SAP solution-based environment are as follows:

- For data files (assuming RAID 1 or similar): four partitions for each disk with 8-14 disks minimum
- For transaction log (assuming RAID 1 or similar): one partition for each disk with 8-12 disks minimum
- For tempdb and log (assuming RAID 1 or similar): one partition for each disk with 4-8 disks

The following table maps these requirements into VMFS volumes.

Table 7-2: Storage Requirements for a Midsize SAP Solution-Based Environment

SAP Database Files	VMFS Volume	Storage Array RAID Requirement
Datafile	VMFS 1	LUN 1 (RAID 1 8-14 disk stripe set)
Datafile	VMFS 2	LUN 2 (RAID 1 8-14 disk stripe set)
Datafile	VMFS 3	LUN 3 (RAID 1 8-14 disk stripe set)
Datafile	VMFS 4	LUN 4 (RAID 1 8-14 disk stripe set)
transaction Log	VMFS 5	LUN 5 (RAID 1 8-12 disk stripe set)
tempdb	VMFS 6	LUN 6 (RAID 1 4-8 disk stripe set)

8 Deployment and Operations

You can leverage VMware Infrastructure to provide significant benefits in a virtualized SAP solution-based data center, including:

- Increased operational flexibility and efficiency: rapid software applications and services deployment in shorter times.
- Efficient change management: increased productivity when testing the latest SAP software patches and upgrades.
- Minimized risk and enhanced IT service levels: zero-downtime maintenance capabilities, rapid recovery times for high availability, and streamlined disaster recovery scenarios across the data center.

This chapter describes these features in more detail.

8.1 Templates and Snapshots

VMware template cloning and snapshot features can increase productivity of system administration and testing in SAP environments.

Templates

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 SAP application data. You can use virtual machine templates to provision a new pre-configured SAP 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 time-consuming.

Snapshots

VMware snapshots preserve the state of a virtual machine. This feature captures the entire state of a virtual machine at the time the snapshot is taken - including the memory state, the settings state (virtual machine settings) and the disk state. Reverting to a snapshot returns these items to the state they were in at the time the snapshot was taken. Using this feature allows tests to be run with just one virtual machine.

The snapshot feature can increase productivity of testing SAP code changes and /or SAP kernel or guest operating system patches in test and development systems. Changes can be tested and, in the event of unsuccessful tests, the SAP virtual machine can be rolled back to a previous snapshot/state prior to the change.

Snapshots can also increase productivity of SAP upgrade testing on test and development systems. An SAP upgrade project normally requires multiple test runs to determine upgrade duration and fine tune the upgrade process. Using snapshots makes it easy to roll back to previous phases of the upgrade in order to retest scenarios. Figure 8-1 illustrates the process. (Prepare, uptime, downtime and post refer to different stages of the SAP upgrade process).

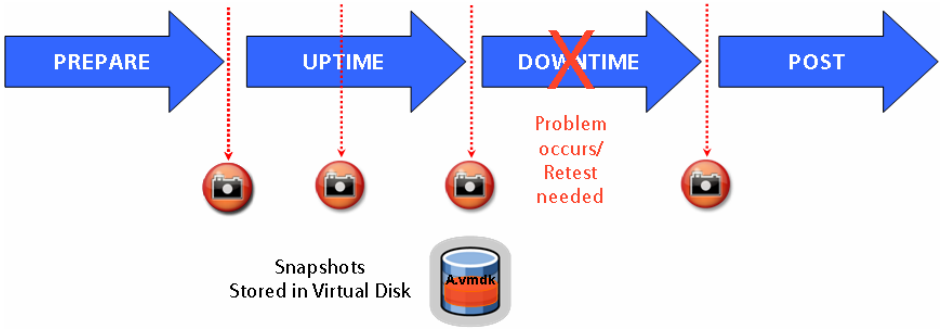

Figure 8-1. Using VMware Snapshots to Test SAP Upgrade

8.2 VMware VMotion, VMware DRS and VMware HA

VMotion technology enables the migration of running virtual machines from one physical server to another without service interruption This migration allows you to move SAP virtual machines from a heavily loaded server to a lightly loaded, one or to offload them to another server to allow for hardware maintenance without any down time.

VMware Distributed Resource Scheduler (DRS) takes the VMotion capability one 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. VMware DRS dynamically monitors the workload of the running virtual machines and the resource utilization of the physical servers within a cluster.

With VMware High Availability (HA), SAP virtual machines on a failed ESX Server host can be restarted on another ESX Server host. This feature provides a cost-effective failover alternative to expensive third-party clustering and replication solutions.

VMotion and DRS perform best if the following elements are in place:

- The source and target ESX Server hosts must be connected to the same Gigabit network and the same shared storage.
- A dedicated Gigabit network for VMotion is recommended.
- The destination host must have enough resources.
- The virtual machine must not use physical devices like CD_ROM, Floppy.
- The source and destination hosts must have compatible CPU models; otherwise migration with VMotion will fail. For a listing of servers with compatible CPUs consult VMotion compatibility guides from the hardware vendors.
- To minimize network traffic it is ideal to keep virtual machines that communicate with each other together (e.g. SAP dialog and database virtual machines) on the same host machine.
- Virtual machines with smaller memory sizes are better candidates for migration than larger ones.

If you use VMware HA, be aware that:

- VMware HA handles ESX Server host hardware failure and does not monitor the status of the SAP services such as the SAP enqueue service or the database; these must be monitored by separate third party clustering software.
- Virtual machines are automatically restarted on the failover ESX Server hosts but the SAP application on the virtual machine is not. Separate startup scripts or services in the guest operating system can enable automatic start up of the SAP applications.

- Proper DNS hostname resolution is required on each ESX Server host in a VMware HA cluster.
- VMware HA heartbeat is sent via the ESX Server service console network so redundancy in this network is recommended (see Part 6 Network Design and Configuration Guidelines).

8.3 General Recommendations

Performance Tuning

For general performance tuning ESX Server best practices consult the VMware whitepaper *Performance Tuning Best Practices for ESX Server 3* : http://www.vmware.com/resources/techresources/707.

Performance Monitoring

Use the VI Client, or esxtop to measure resource utilization. The measuring tools in the guest operating system may produce different results, for example: the memory of the guest operating system as measured within the guest itself is not necessarily the actual memory currently allocated by ESX Server; available memory within the guest may come from swap on the ESX Server host.

VMware® Tools

VMware Tools consists of a set of kernel modules used by guest operating systems running inside of a VMware virtual machine. The modules facilitate interaction between the guest operating system and the ESX Server vmkernel. The tools provide key components for optimal performance including: the balloon and optimized network/storage drivers. The latest version of VMware Tools should be installed in each guest operating system.

9 Best Practice Guidelines Summary

This section summarizes the guidelines included in this document.

General

- Work with your hardware vendors to get detailed sizing for your solution and follow the normal SAP Quick Sizer[4] sizing process.
- To address sizing requirements, scale the SAP application tier horizontally across multiple virtual machines as is typically done on physical servers in the native world.
- A viable architecture for large SAP implementations includes deploying SAP application instances in virtual machines, but leaving the database and CI on native hardware.
- Consult SAP note 1056052 (*VMware ESX Server 3.0 Configuration Guideline*) for additional guidelines.

CPU

- Setting the number of vCPUs: start with smaller sized virtual machines and add vCPUs as required. More vCPUs in virtual machines can increase scheduling overhead. For example, while management tasks increase with more virtual machines, you will get better throughput with a larger number of 2 x vCPU virtual machines as opposed to having fewer 4 x vCPU virtual machines.
- To maximize performance of production-critical systems and prepare for the worst case scenario of the CPU load of each virtual machine peaking simultaneously, size the total number of vCPUs of all the virtual machines to be equal to or less than the total number of cores on the ESX Server host machine.

Memory

- To maximize performance of production-critical systems and prepare for a worst case scenario of memory utilization of each virtual machine peaking simultaneously, define virtual machine memory as follows:
    - Do not over commit memory – the configured memory size of all the virtual machines should be equal to the physical memory on the ESX Server host minus the system and virtual machine memory overhead.
    - Set the reservation of the virtual machine to the same value as the configured size. This prevents virtual machine memory being reclaimed by ESX Server for use by other guests.
    - Do not disable the balloon driver.
- Follow the existing SAP recommendations, documentation, and notes to configure memory and swap/page file within the virtual machine.
- For NUMA based systems:
    - Set the vCPU and memory of virtual machines so that they fit into a NUMA node.
    - Do not set CPU or NUMA node affinities.

CPU and Memory Over-commitment

[4] The Quick Sizer is a Web-based tool designed to make the sizing of SAP solutions easier and faster. It has been developed by SAP in close cooperation with all platform partners and is free of cost.

The recommendation not to over-commit CPU and memory for SAP virtual machines for which performance requirements are critical (e.g. production systems) is a conservative guideline. Consider it a starting point for sizing SAP virtual machines across ESX Server hosts to maximize performance and to handle a worst case scenario where peak resource usage of each virtual machine occurs simultaneously. You can assign higher levels of over-commitment in the following situations:

- Non-production SAP systems subject to less stringent performance requirements.
- Monitoring of actual resource usage of CPU and memory indicate unused capacity.
- ESX Server memory page sharing generates appreciable memory savings.

Network

- Set up larger implementations of SAP virtual machines on two separate networks :
  - Private: traffic between SAP virtual machines.
  - Public: traffic between SAP virtual machines and end users.
- Use a minimum of four NIC ports per ESX Server host to separate VMotion, ESX Server console and virtual machine traffic.
- There are three different configuration modes for VLAN tagging - Virtual Switch Tagging (VST) mode is recommended.
- Use NIC teaming for load balancing and passive failover.

Storage

- A minimum of two HBA ports per ESX Server host is required. Configure multipathing to allow for four paths to storage. This configuration is similar to native setups for SAN redundancy and I/O performance.
- Generally a mixed practice of RDM and VMFS use is practical for SAP installations:
  - RDM is required for third party clustering software and is compatible with third party storage management solutions.
  - VMFS increases storage utilization and facilitates easier administration via the VI client for features like VMware template cloning and snapshots.
- Carefully consider the number of SAP virtual machines allocated per LUN as this can impact I/O performance:
  - Performance-critical production virtual machines will benefit from being spread across multiple LUNs.
  - Storage layout of data files for performance-critical production databases should follow the guidelines for native environments: separate LUNs recommended in native architectures will translate to separate VMFS volumes (or RDM disks) in the virtualized architecture.
  - Consolidating multiple non-production SAP virtual machines into fewer LUNs is viable as performance is less critical.

Deployment / Operations

- Use of template cloning and snapshots can increase productivity of system administration tasks and testing of SAP code changes and upgrades.
- VMotion and VMware DRS can balance resource load and maximize performance across ESX Server hosts.
- VMware HA offers cost-effective out-of-the-box high availability for VMware Infrastructure. However, it does not monitor the health of individual SAP services such as enqueue and database. Use third party clustering software to monitor these single points of failure.
- As the VMware HA heartbeat uses the console service network, set this network up in a redundant configuration..

- Install the latest version of VMware Tools in the guest operating system for optimal performance.
- Use VI client and esxtop to measure resource utilization of the ESX Server host, as resource utilization measured within the guest may be different from what is measured on the ESX Server host.

Appendix 1 – Additional Resources

SAP notes:

- 1056052 - VMware ESX Server 3.0 Configuration Guideline
- 1122388 - Linux: VMware ESX Server 3 configuration guideline
- 1084815 - SAP support in virtualized Solaris environments
- 895807 - SAP support in virtualized Linux environments
- 1104578 - Virtualization in Windows: Enhanced monitoring
- 674851 - Virtualization on Windows
- 989963 – Linux: VMware timing problem

SAP benchmark certifications:
http://www.sap.com/solutions/benchmark/sd2tier.epx

VMworld 2007 Breakout Sessions (Useful sessions on performance best practices are available here. You must be a member of VMworld.com to access these presentations.):
http://www.vmworld.com/vmworld/sessions.jspa

VMware White Papers
Performance Tuning Best Practices for ESX Server 3:
http://www.vmware.com/resources/techresources/707

VirtualCenter Monitoring and Performance Statistics:
http://www.vmware.com/resources/techresources/771

Resource Management Guide ESX Server 3.0.1 and VirtualCenter 2.0.1:
http://www.vmware.com/pdf/vi3_301_201_resource_mgmt.pdf

VMware Virtual Networking Concepts:
http://www.vmware.com/resources/techresources/412

VMware ESX Server 3 802.1Q VLAN Solutions:
http://www.vmware.com/resources/techresources/412

Server Configuration Guide (networking and storage section):
http://www.vmware.com/pdf/vi3_301_201_server_config.pdf

VMware SAN System Design and Deployment Guide:
http://www.vmware.com/resources/techresources/772

VMware Virtual Machine File System: Technical Overview and Best Practices:
http://www.VMware.com/pdf/vmfs-best-practices-wp.pdf

VMware, Inc3401 Hillview Ave Palo Alto, CA 94304 USA
Tel 650-427-5000 Fax 650-427-5001 www.vmware.com