DEPLOYMENT GUIDE

INDUSTRY: HEALTHCARE

FUJIFILM Synapse® PACS Deployment Guide for VMware Infrastructure

FUJIFILM Medical Systems USA, Inc.
Thomas E. Riddle and Ray Milot

vmware FUJIFILM

Table of Contents

1.0 Introduction

This document provides direction to those interested in running the Fujifilm Synapse™ picture archiving and communication system (PACS) on VMware® Infrastructure. It provides a description of the architecture of Synapse, as well as the value of utilizing the VMware platform. It also includes results of recent testing done jointly by VMware and FUJIFILM Medical Systems USA (Fujifilm), which characterized the performance and functionality of Synapse on VMware infrastructure version 3.5 update 3. Finally, the document outlines best practices for using the two product sets together in your datacenter.

2.0 VMware Infrastructure

VMware infrastructure is a leading virtualization solution that provides multiple benefits to IT administrators and users. VMware Infrastructure provides a layer of abstraction between the resources required by an application and the operating system, and the underlying hardware that provides those resources. The value of this abstraction layer includes the following:

- Consolidation: VMware solutions allow multiple physical servers to be consolidated into one server, with little or no decrease in overall performance.
- Ease of Provisioning: VMware Infrastructure allows encapsulation of an application into an image that can then be duplicated or moved, greatly reducing the cost of application provisioning and deployment.
- Manageability: Virtual machines may be moved from server to server with no down time using VMware® Vmotion™, which simplifies common operations like hardware maintenance and planned downtime.
- Availability: Unplanned downtime can be reduced and higher service levels may be provided to an application. VMware High Availability (HA) ensures that in the case of an unplanned hardware failure, any affected virtual machines will be restarted on another host in a VMware cluster.

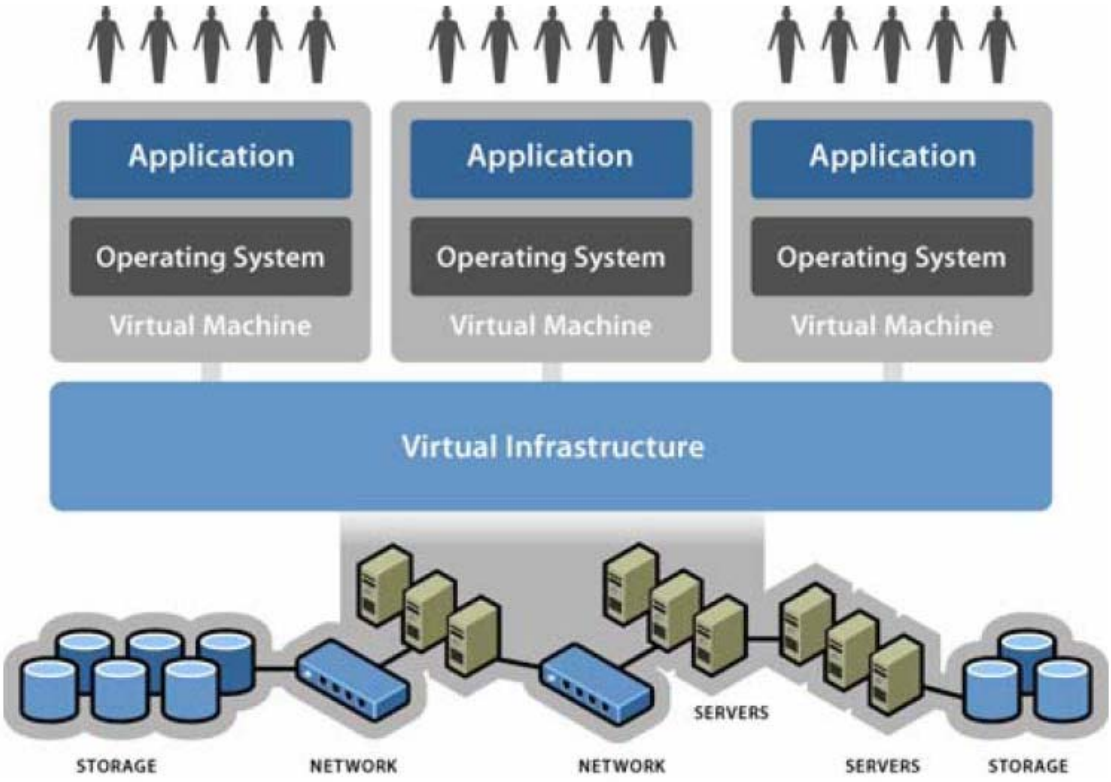

VMware Infrastructure includes the following components:

VMware ESX™ (ESX): A robust, production-proven virtualization layer, run on physical servers, that abstracts processor, memory, storage, and networking resources into multiple virtual machines.

VMware vCenter™ Server (vCenter): A central point for configuring, provisioning, and managing virtualized IT environments.

VMware Virtual Infrastructure Client (VI Client): A graphical interface that allows users to connect remotely to the vCenter Server or to individual ESX instances from any Windows PC.

VMware Virtual Infrastructure Web Access (VI Web Access): A Web interface that allows virtual machine management and access to remote consoles.

VMware Virtual Machine File System (VMFS): A high-performance clustered file system for ESX virtual machines.

VMware Virtual Symmetric Multi-Processing (SMP): Allows a single virtual machine to use multiple physical processors simultaneously.

VMware VMotion™ (VMotion): Enables the live migration of running virtual machines from one physical server to another with zero down time, continuous service availability, and complete transaction integrity.

VMware High Availability (HA): Provides easy-to-use, cost-effective high availability for applications running in virtual machines. In the event of server failure, affected virtual machines are automatically restarted on other production servers that have spare capacity.

VMware Distributed Resource Scheduler (DRS): Allocates and balances computing capacity dynamically across collections of hardware resources for virtual machines.

VMware Consolidated Backup (Consolidated Backup): Provides a centralized facility for agent-free backup of virtual machines. It simplifies backup administration and reduces the load on the ESX Servers.

VMware Infrastructure SDK (SDK): supports a standard programming interface for VMware and third-party solutions that want to access the VMware Infrastructure.

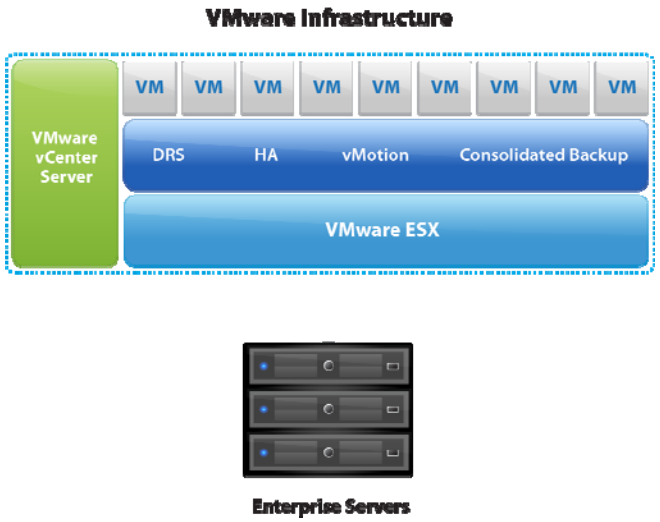

Figure 1: VMware Infrastructure Architecture

3.0 Synapse Architecture and Deployment Strategy

Synapse was designed to provide on-demand image access, a scalable architecture, and utilize integrated Web technology. The Oracle Database and Synapse components are designed to be "hardware agnostic" – whether the hardware is physical or virtual. Synapse is a software package designed to fully utilize the Microsoft Windows Workstation and Server Operating Systems. Synapse takes full advantage of Microsoft Windows Active Directory security to authenticate all users, allowing secure access into Synapse.

Fujifilm Synapse® Overview

Fujifilm Synapse™ PACS is a collection of software modules built on the Microsoft Windows Server platform, which together provide the core software functionality for the product. The server software communication with the workstation is entirely web based and uses Microsoft Internet Information Server (IIS) to provide image access to client workstations.

Synapse consists of the following server modules:

Database Server
The Synapse database operationally tracks all aspects of the PACS. The database supports the system folder structure, which organizes the patients and studies. All workstations communicate with the database through Hyper Text Transfer Protocol (HTTP) communication. Synapse uses Oracle Database 10g as its database foundation.

Web Server
Synapse uses Windows Internet Information Server (IIS) as its core Web server. Synapse is a Web-based system – all data going to and from the Synapse workstation goes through the Web server(s). All images, information and user authentication are sent over standard Web ports – port 80 for standard communication and port 443 for SSL are typical in most installations.

Storage Server
Synapse storage servers are Windows servers that are used for the storage and distribution of images, documents and other Synapse file objects. Storage directories are then presented as Universal Naming Convention (UNC) paths to the Web servers, which wrap the file content for Web-based distribution to the Synapse workstation.

DICOMServer
Synapse DICOMServer software receives studies directly from DICOM modalities without the need for modality interface gateways or interface units. All modalities are direct TCP/IP connections to the network. Synapse DICOMServer software also provides direct, brokerless DICOM Modality Worklist Management (DMWL) to any modality supporting this functionality and responds to all query/retrieve, modality performed procedure step and storage commitment requests.

HIS Server
The HIS (Hospital Information System) Synapse RIS Interface Software is integrated as an HL-7 interface engine, which provides direct brokerless connections to any HL-7 information system. It supports patient, order and report related information. Admit/Discharge/Transfer (ADT) related information could originate from the RIS or HIS.

Synapse Server
Figure 2 depicts the architecture of the Synapse servers.

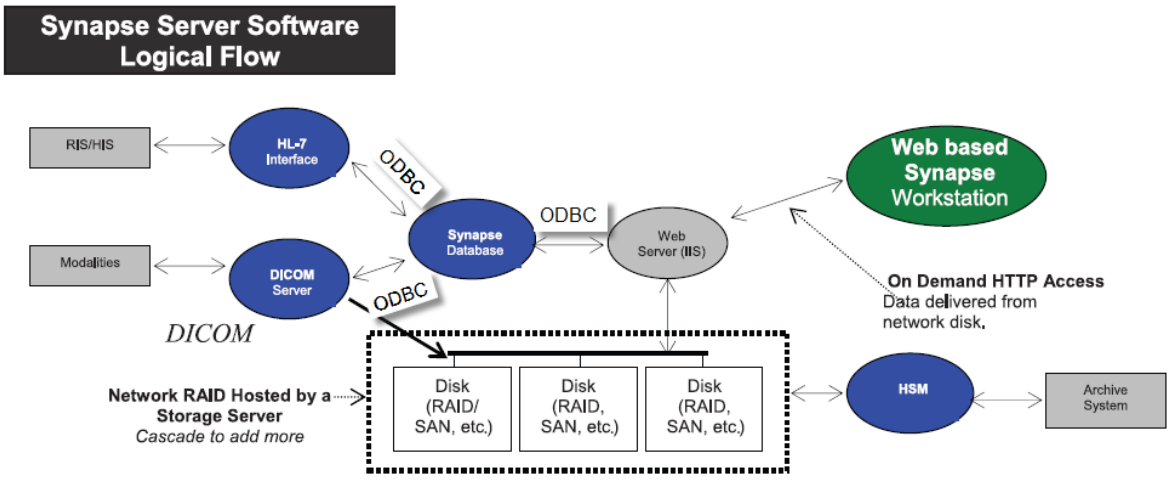

Figure 2: Synapse Server Architecture

4.0 Testing Process and Results

To characterize the performance of Synapse on VMware Infrastructure, VMware and Fujifilm jointly carried out performance tests. The tables below contain information about the test configuration.

Hardware Configuration:

Hardware	Configuration
2 machines of HP Proliant BL 460c G1	3 GHz dual quad-core Intel Xeon E5450
	16GB RAM
Storage: Netapp FAS 3020	Aggregate of 10 disks
	RAID DP
	4 Gbps fibre connectivity
	10K RPM disks
	FC protocol

Virtual Machine Configuration:

VMs	Configuration
DICOM and Web server	2 vCPUs
	2GB memory (reservation= 2 GB)
	1 Ethernet card
	20GB storage (1 x 10 GB, 1 x 10GB)
	(vmdk only)
Synapse Storage server	1 vCPUs
	4GB memory (reservation= 2 GB)
	1 Ethernet card
	110GB storage
	(1 x 10 GB, 1 x 100GB)
	(100GB = RDM)
Database server	2 vCPU
	2GB memory (reservation= 2 GB)
	130GB storage (O: Drive)
	(3 x 10GB, 1 x 40GB)
	(40GB = RDM)
Modality simulator	2 vCPU
	2GB memory
	40 GB storage

Workload Used

Fujifilm provided a modality simulator provided that was used as the load driver. It was configured to send a variable number of studies from several modalities as described in the table below.

Description of load			
Modality Type	# of Images Per Study	Size of Image in KB	Size of Study in MB
CT	500	515	251
US	40	301	12
MG	4	34406	134
CR	30	7352	215
MR	400	151	59
XA	10	2171	21

Testers injected the studies into the DICOM server and retrieved them at the workstation. Testers monitored the image retrieval time on the workstation as well as the resource utilization on the DICOM, storage and DB server. They tried different configurations for the DICOM server based on its resource utilization.

Testers also executed the Zeus bench test to test the ability to handle simultaneous Web connections.

Finally, they tested the horizontal scalability of the DICOM server by adding a second DICOM server of the same configuration.

Results

Figure 3 shows the results of testing on three different virtual machine configurations.

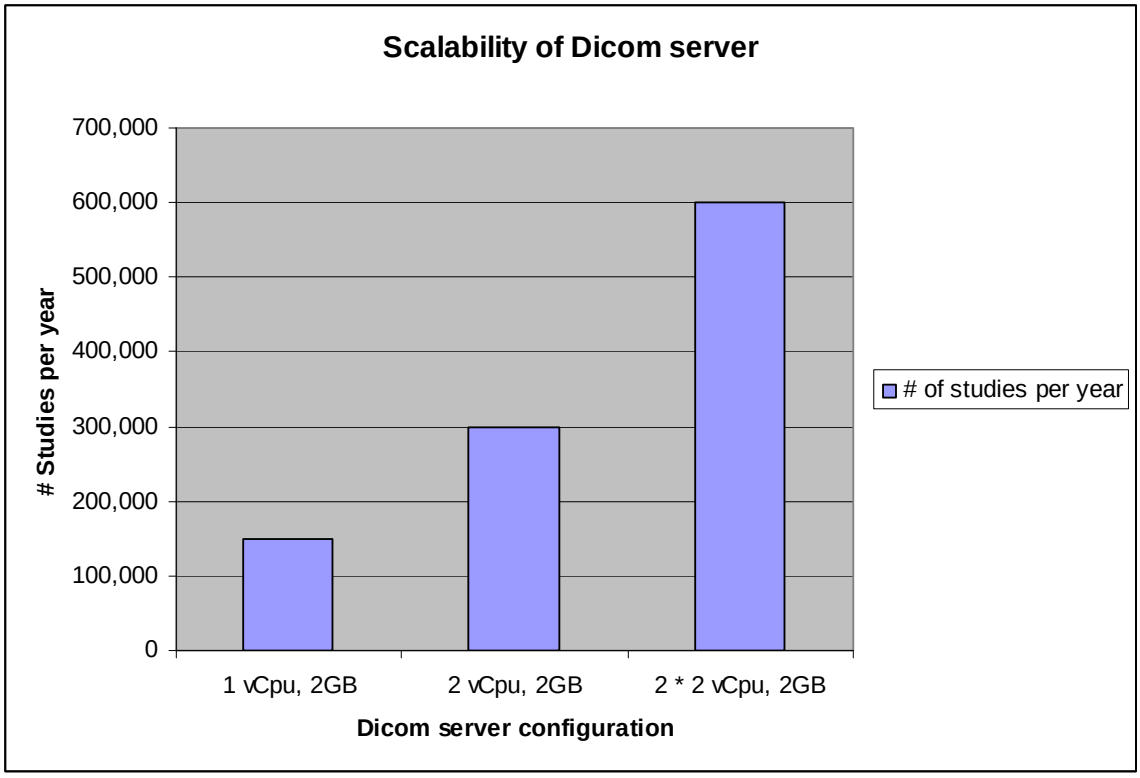

Figure 3: Scalability of DICOMServer running in VMware virtual machines

The DICOM server scales linearly when scaled both vertically and horizontally. Doubling the capacity by increasing virtual CPU resources or adding additional DICOM servers doubles the number of studies that can be performed per year. This table summarizes test results:

VM spec			Load driver	Performance	Scale
DICOM	DB	storage	# of modalities sending images in parallel	Image retrieval time sec	# of studies per year
2 vCpu, 4GB	2 vCpu, 4GB	1 vCpu, 4GB	5 (1500 studies)	< 1	150,000
2 vCpu, 4GB	2 vCpu, 4GB	1 vCpu, 4GB	12 (1000 studies)	< 1	300,000
2 vCpu, 4GB	2 vCpu, 4GB	1 vCpu, 4GB	10 simultaneous web connection pulling 7 MB image (zeus bench)	< 1	300,000
1 vCpu, 2GB	2 vCpu, 2GB	1 vCpu, 2GB	6(500 studies)	2-5	150,000
2 vCpu, 2GB	2 vCpu, 2GB	1 vCpu, 2GB	12 (1000 studies)	< 1	300,000
4 vCpu, 2GB	2 vCpu, 2GB	1 vCpu, 2GB	12 (1000 studies)	< 1	300,000
2 * 2 vCpu, 2GB	2 vCpu, 2GB	1 vCpu, 2GB	24 (2000 studies) horizontal scalability with 2 DICOM servers	< 1	600,000
2 vCpu, 2GB	2 vCpu, 2GB	1 vCpu, 2GB	Vmotion	< 1	300,000

The following tables summarizes resource utilization on the DICOM server for different virtual machine configurations as the number of modalities sending images increases:

Load driver	VM spec	Resource utilization - DICOM		
# of modalities sending images in parallel	DICOM	Cpu % (avg/max)	Memory MB (avg/max)	Disk - latency ms (avg/max)
5 (1500 studies)	2 vCpu, 4GB	35/82	202/573	2/32
12 (1000 studies)	2 vCpu, 4GB	75/100	820/1147	1/61
10 simultaneous web connection pulling 7 MB image (zeus bench)	2 vCpu, 4GB	72/100	772/1147	1/83
6(500 studies)	1 vCpu, 2GB	87/100	1553/1905	2/8
12 (1000 studies)	2 vCpu, 2GB	80/99	717/819	2/35
12 (1000 studies)	4 vCpu, 2GB	43/83	1389/1905	2/10
24 (2000 studies)	2 * 2 vCpu, 2GB	Utilization similar to 1 Vm case		
Time required for migration of virtual machines with VMotion	2 vCpu, 2GB	VMs migrated in 19 to 42 seconds		

The next two tables summarize resource utilization on the database server and the storage server:

Load driver	VM spec	Resource utilization - DB		
# of modalities sending images in parallel	DB	Cpu % (avg/max)	Memory MB (avg/max)	Disk - latency ms (avg/max)
5 (1500 studies)	2 vCpu, 4GB	4/17	116/205	2/315
12 (1000 studies)	2 vCpu, 4GB	10/26	138/369	1/10
10 simultaneous web connection pulling 7 MB image (zeus bench)	2 vCpu, 4GB	9/26	156/328	1/13
6(500 studies)	2 vCpu, 2GB	9/15	126/143	2/5
12 (1000 studies)	2 vCpu, 2GB	13/22	175/246	2/5
12 (1000 studies)	2 vCpu, 2GB	8/18	152/246	1/10
24 (2000 studies)	2 vCpu, 2GB	Utilization similar to 1 Vm case		
Time required for migration of virtual machines with VMotion	2 vCpu, 2GB	VMs migrated in 19 to 42 seconds		

Load driver	VM spec	Resource utilization - storage		
# of modalities sending images in parallel	storage	Cpu % (avg/max)	Memory MB (avg/max)	Disk - latency ms (avg/max)
5 (1500 studies)	1 vCpu, 4GB	7/19	661/2458	1/69
12 (1000 studies)	1 vCpu, 4GB	14/29	1400/1966	2/9
10 simultaneous web connection pulling 7 MB image (zeusbench)	1 vCpu, 4GB	13/28	1029/1843	2/17
6(500 studies)	1 vCpu, 2GB	10/19	516/1024	4/11
12 (1000 studies)	1 vCpu, 2GB	14/22	1143/1741	5/11
12 (1000 studies)	1 vCpu, 2GB	11/24	893/1372	3/9
24 (2000 studies)	1 vCpu, 2GB	Utilization similar to 1 Vm case		
Time required for migration of virtual machines with VMotion	1 vCpu, 2GB	VMs migrated in 19 to 42 seconds		

Observations:

- The CPU utilization on the DICOM server depends on the number of modalities running in parallel. For a virtual machine with one vCPU, six modalities can drive CPU utilization to an average of 87 percent and maximum of 100 percent. Similarly, for a two vCPU virtual machine, 12 modalities can utilize an average of 80 percent of CPU resources, and for a 2 * 2 vCPU virtual machine, 24 modalities can utilize on average 80 percent of CPU resources.

- The image retrieval time on the workstation in most cases is less than one second when observed through http analyzer.

Additional Testing

VMotion technology was tested by migrating the three virtual machines -- DICOMServer, database and storage -- one at a time from one physical host to another. The modality simulator was engaged to send 1000 studies from 12 modalities on the DICOM server with a two vCPU 2GB configuration. During the migration, no server took more than 42 seconds to migrate and image retrieval was always less than one second.

5.0 Deployment Best Practices

The best practices for deploying Synapse are explained here:

1. When configuring virtual machines, remember that ESX itself has some overhead and allow for the CPU overhead required by virtualization. If the virtual machines are highly utilized (80 to 90 percent), leave one core on the host for the hypervisor's processing and don't allocate that to any virtual machine. If the virtual machines are lightly utilized (40 percent or less) you may over-commit in terms of vCPUs. In the case of the testing described in this document, the system was loaded with a large number of studies to see the performance under high cpu utilization (80 percent or greater), so one core was left uncommitted so it could be used for hypervisor processing.

2. Carefully select the amount of virtual memory you allocate to your virtual machines to allow enough memory to hold the working set of applications you will run in the virtual machine. DICOM and storage server use less than 2GB, so 2GB is enough for the two virtual machines. The storage server uses around 2GB, so allocate a maximum of 4GB to the storage server. Allocating more resources than are required may lead to performance degradation.

3. ESX requires some memory for its own processing. The ESX service console typically uses 272MB and the VMkernel uses 273MB. Overhead memory includes space reserved for the virtual machine frame buffer and various virtualization data structures. Overhead memory depends on the number of virtual CPUs, the configured memory for the guest operating system, and on whether you are using a 32-bit or 64-bit guest operating system. For Synapse virtual machines the overhead is about 150MB per virtual machine, which can be arrived at from the VMware reference document "http://www.vmware.com/pdf/vi_performance_tuning.pdf".

4. Use separate vSwitches (and consequently separate physical network adapters) to avoid contention between the service console, VMkernel, and virtual machines, and between virtual machines running heavy networking workloads. Service console is on vSwitch0, VMkernel is on vSwitch1, DICOM server and storage server are on vSwitch2.

5. To establish a network between two virtual machines that reside on the same ESX host, connect both virtual machines to the same virtual switch. If the virtual machines are connected to different virtual switches, traffic will go through wire and incur unnecessary CPU and network overhead. For example, keep the DICOM and storage servers vSwitch 2.

6. Ensure that heavily-used virtual machines do not all access the same Virtual Machine File System (VMFS) volume concurrently and that they are spread across multiple VMFS volumes (typically a VMFS volume will span a single LUN, but this is not always true). Heavy SAN I/O when a large number of virtual machines access the same VMFS volume concurrently will cause poor disk performance. In the test configuration, three separate luns were created for the three virtual machines, and they were not kept in the same datastore.

7. Avoid running a four-processor virtual machine on a dual-processor host system, even if the dual-processor host has hyper-threading (that is, four logical CPUs). Testers disabled hyperthreading on the host during the tests described in this document.

8. Enable transmit coalescing and set it to 128 in host->configuration->advanced setting ->net ->Net.vmxnetThroughputWeight. This favors throughput over latency and is suitable for applications with high throughput requirements.

9. Follow other general best practices mentioned in "http://www.vmware.com/pdf/vi_performance_tuning.pdf" such as installing the latest version of VMware Tools in the guest operating system and disconnecting unused, unnecessary devices on both the guest and on the host such as USB adapters.

6.0 Reference Configurations

Based on the testing the following recommendations are made for sizing virtual machines for Synapse deployments:

Small deployment:
When the site has five or fewer modalities sending images simultaneously and the total studies per year is less than 125,000, it can be considered as a small size deployment and the following configuration can be used:
1 DICOM/web server virtual machine: 1 vCpu , 2GB memory
1 DB server virtual machine: 2 vCpu , 2GB memory
1 Storage server virtual machine: 1 vCpu , 4GB memory

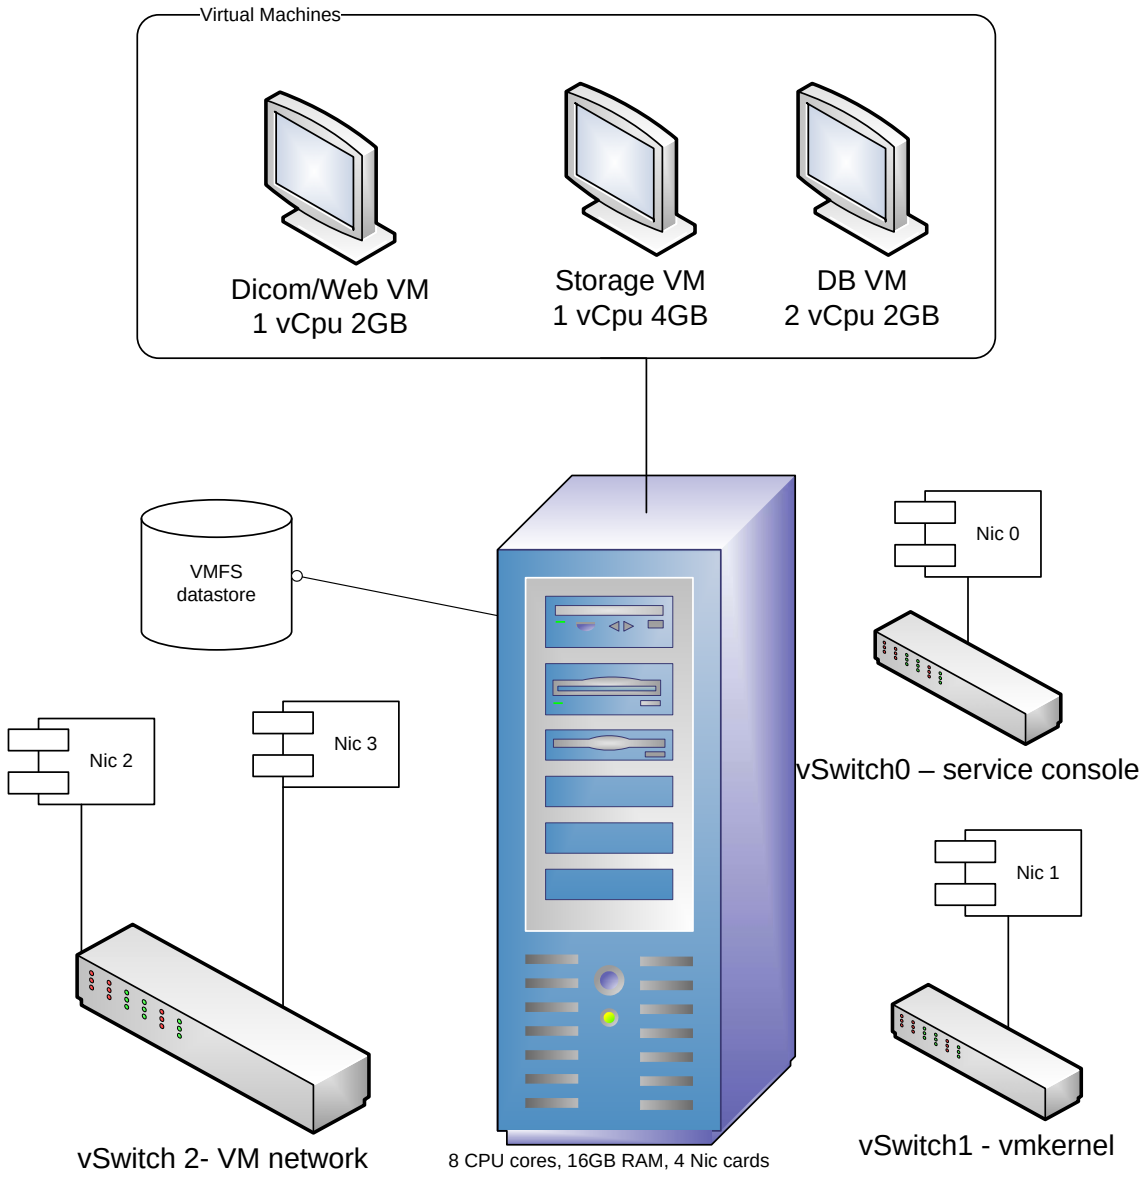

Figure 4: Recommended configuration for small deployments

Medium deployment:
When the site has six to 12 modalities sending images simultaneously and the total studies per year are between 125,000 and 300,000, it can be considered a medium size deployment and the following configuration can be used:
1 DICOM/web server virtual machine: 2 vCpu , 2GB memory
1 DB server virtual machine: 2 vCpu , 2GB memory
1 Storage server virtual machine: 1 vCpu , 4GB memory

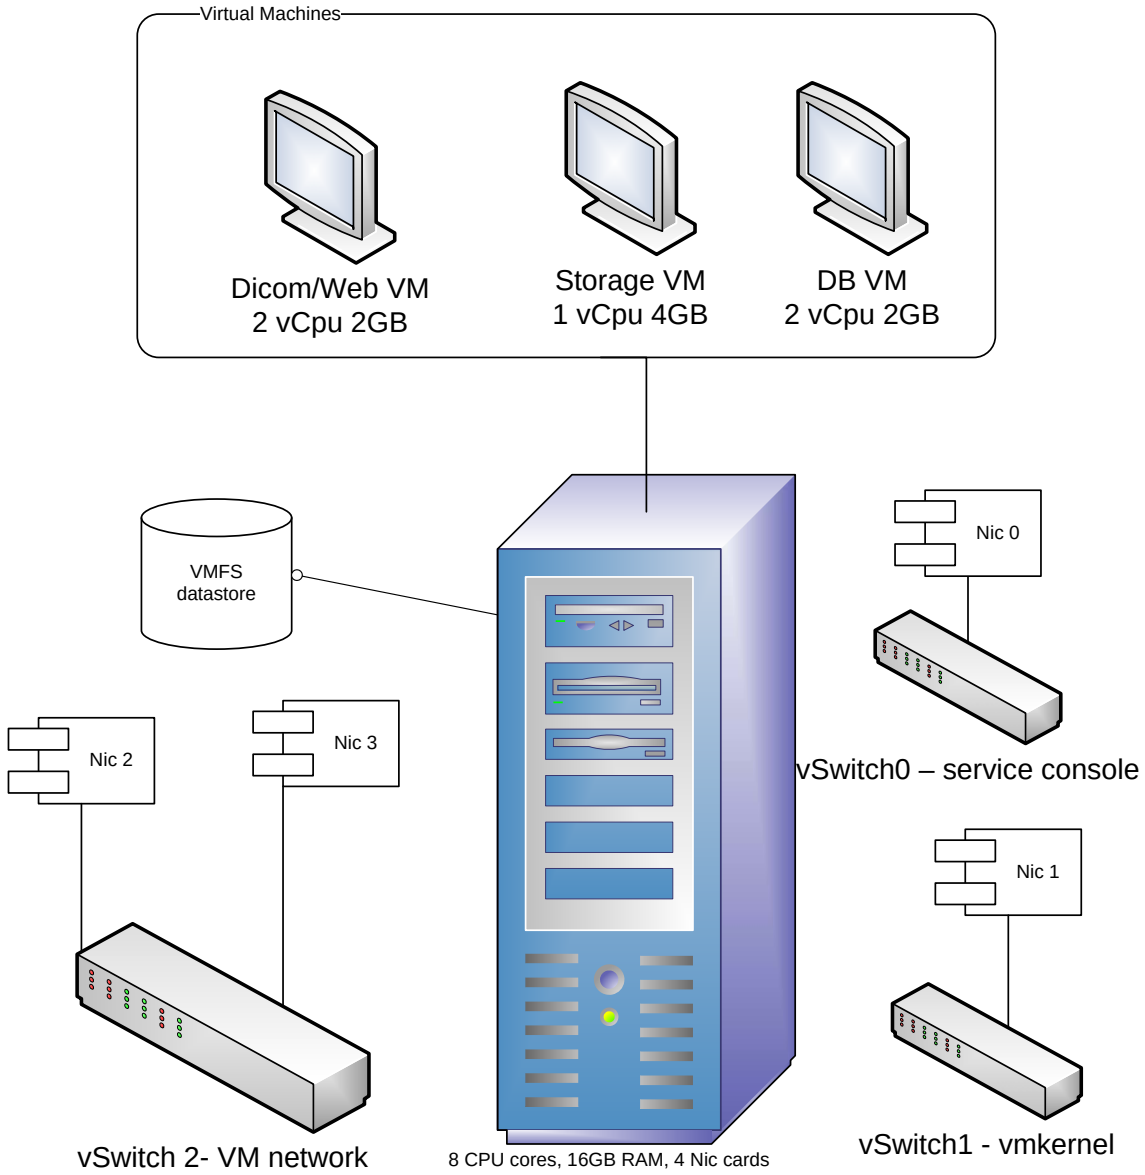

Figure 5: Recommended configuration for medium deployments

Large deployment:
When the site has 13 to 24 modalities sending images simultaneously and the total studies per year are between 300,000 and 600,000, it can be considered a large size deployment and the following configuration can be used:
2 DICOM/web server virtual machine: 2 vCpu , 2GB memory
1 DB server virtual machine: 2 vCpu , 2GB memory
1 Storage server virtual machine: 1 vCpu , 4GB memory

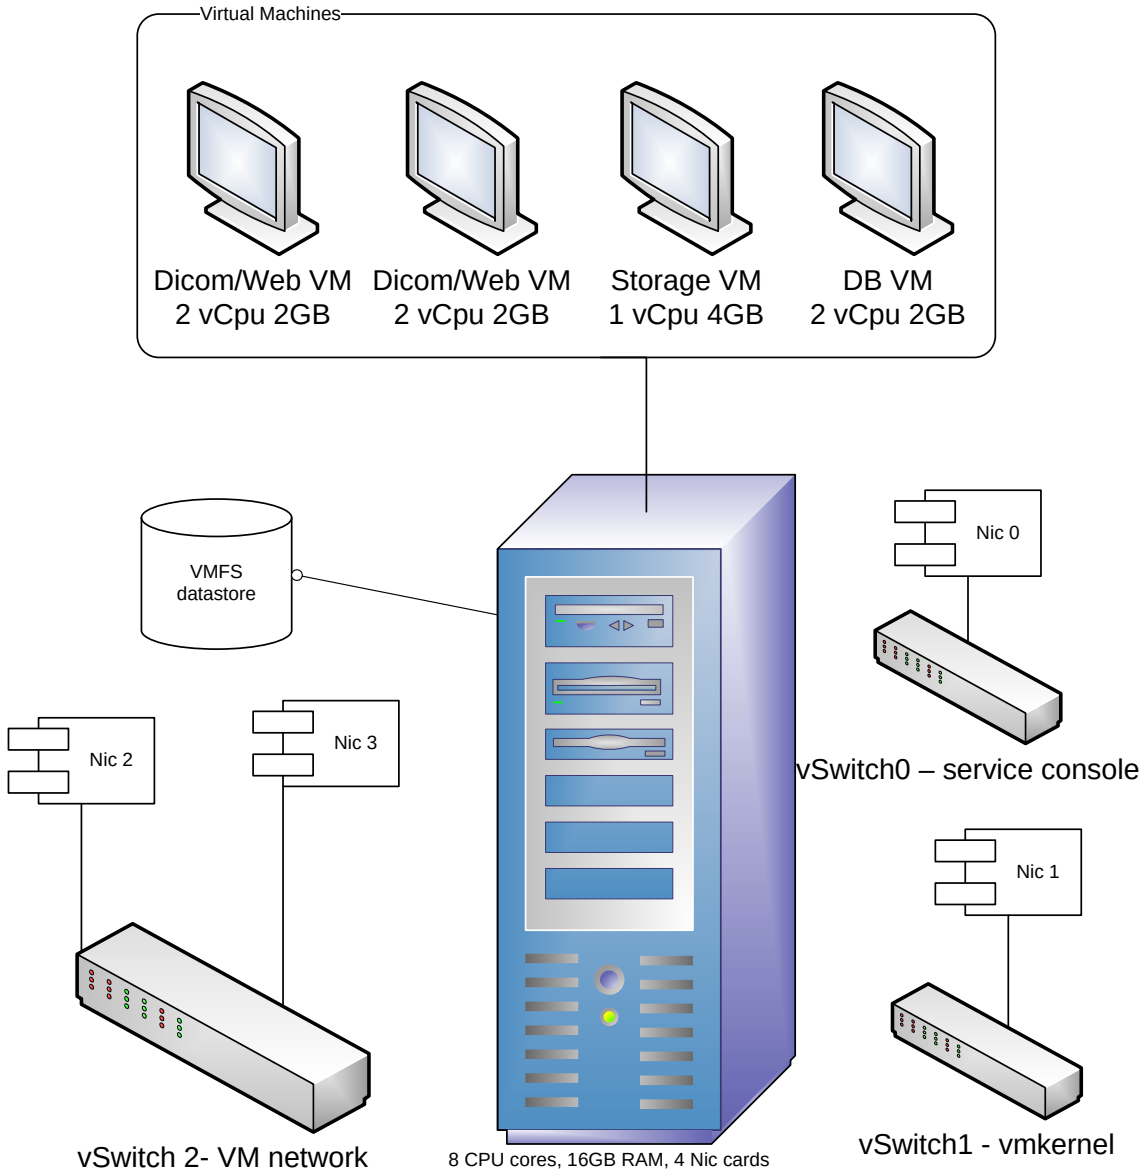

Figure 6: Recommended configuration for large deployments

You can add additional DICOM virtual machines of the same spec to scale DICOM server to handle more studies per year, with each additional virtual machine capable of handling about 150,000 studies per year.

7.0 Conclusions

Overall, the testing results show that Synapse performs well on VMware Infrastructure. Furthermore, virtualization has the potential to reduce costs, increase service levels, and simplify the manageability of the application.

- The performance of the Synapse application on VMware Infrastructure 3.5 is comparable to a physical deployment, with image retrieval time of less than one second.
- The application scales linearly as the virtual machines are scaled out. One DICOM server virtual machine can handle 300,000 studies per year, and two DICOM server virtual machines can handle 600,000 studies per year.
- Synapse is able to leverage advanced VMware features such as VMotion and DRS. Virtual machines migrate from one host to another in 40 seconds without any impact on performance.

8.0 Resources

You can find more information about VMware products and the complete Synapse healthcare informatics solution portfolio at the following locations:

VMware Web site: http://www.vmware.com/

FUJIFILM Medical Systems Web site: http://www.Fujimed.com/

VMware Infrastructure 3 product Web site:
http://www.vmware.com/products/data_center.html

VMware Infrastructure 3 download Web site:
https://www.vmware.com/download/vi/

VMware support Web site:
http://www.vmware.com/support/

VMware Performance Tuning Paper:
http://www.vmware.com/pdf/vi_performance_tuning.pdf

System Compatibility Guide for a complete list of compatible hardware:
http://www.vmware.com/pdf/vi35_systems_guide.pdf

Storage/SAN Compatibility Guide for a complete list of compatible storage devices:
http://www.vmware.com/pdf/vi35_san_guide.pdf

I/O Compatibility Guide for a complete list of compatible networking devices:
http://www.vmware.com/pdf/vi35_io_guide.pdf

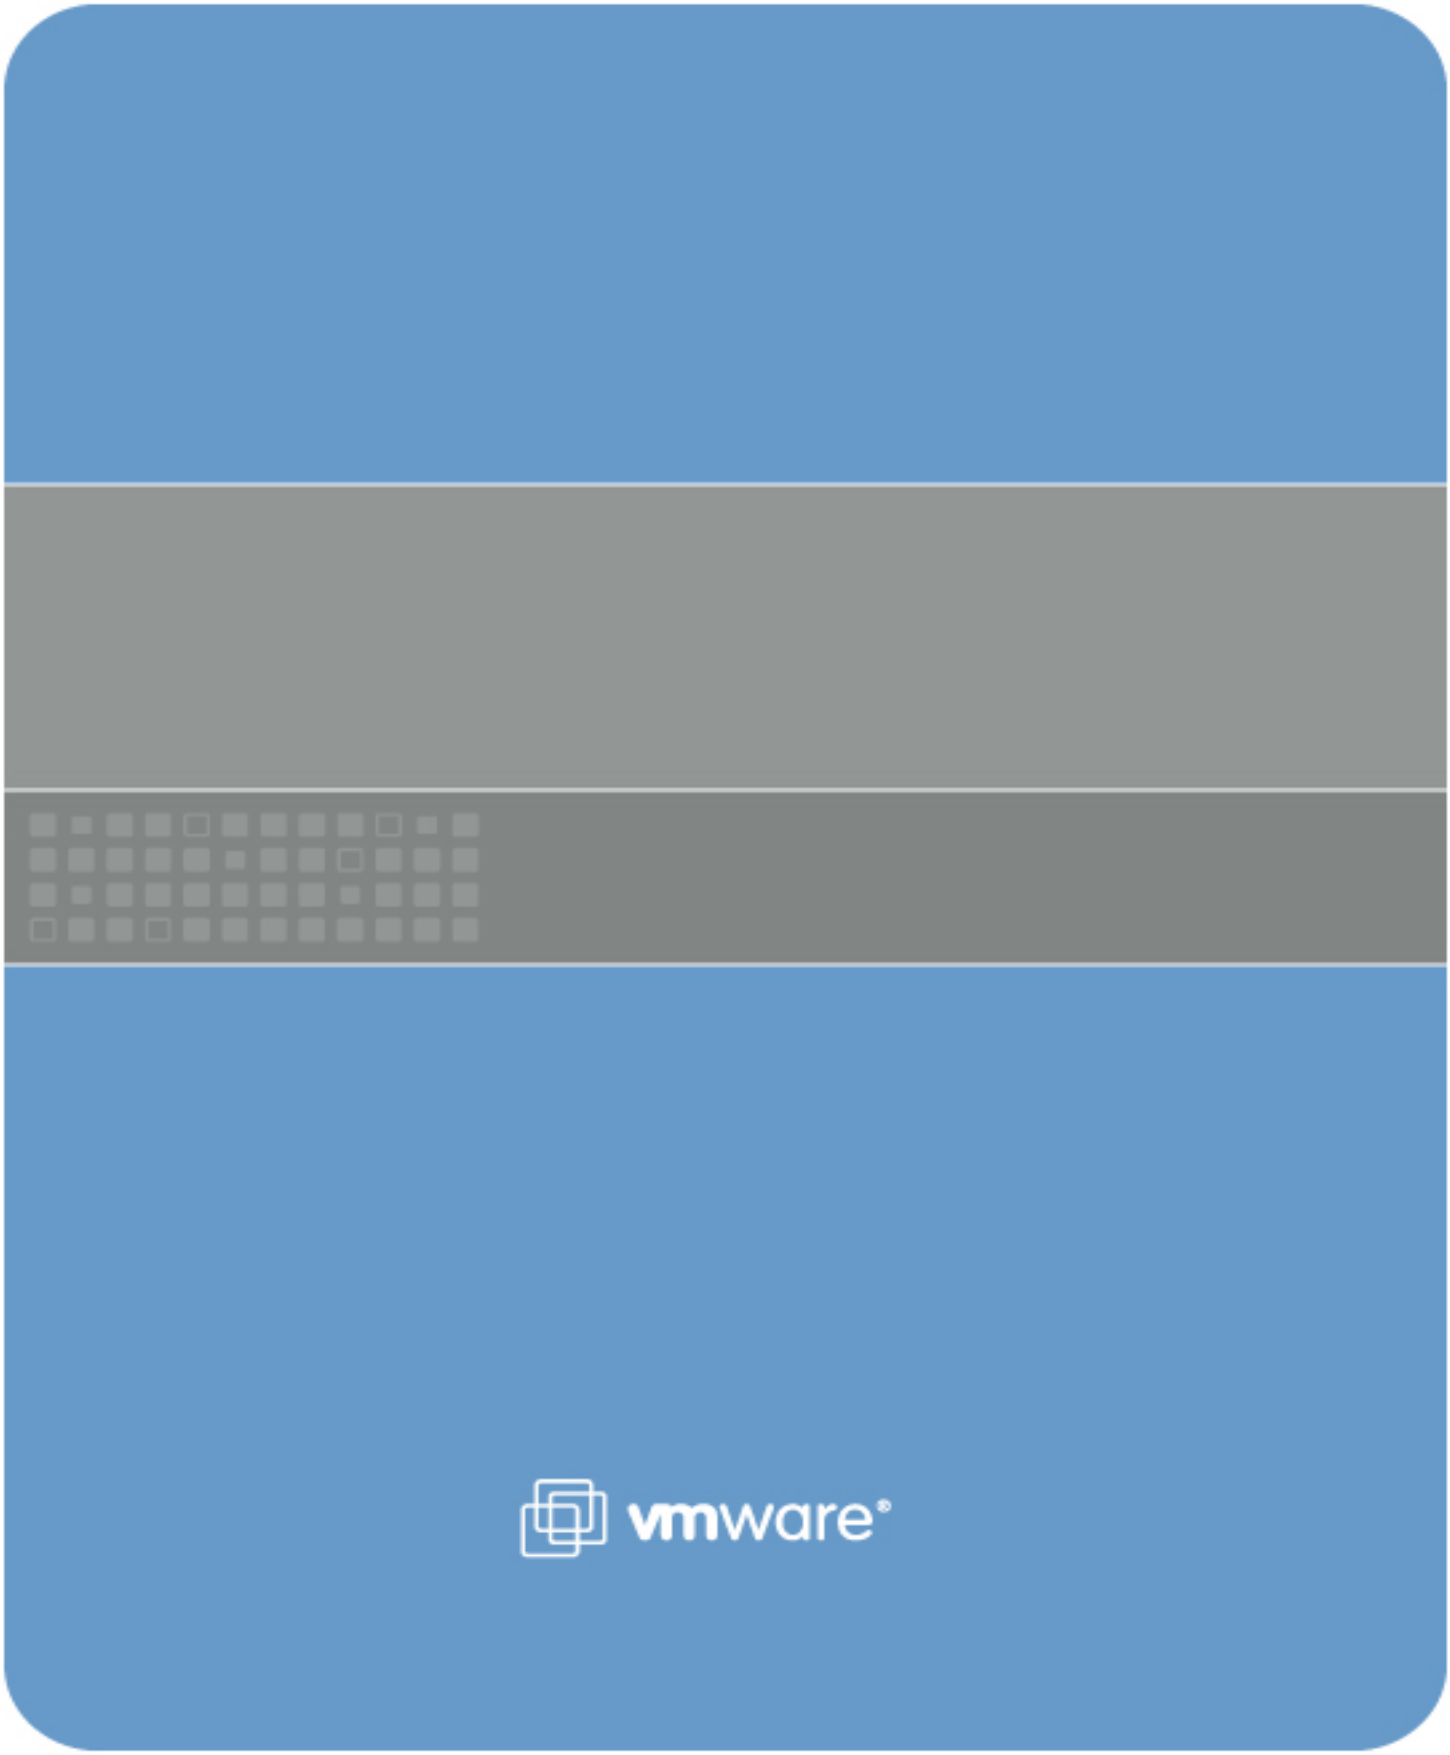

VMware, Inc. 3401 Hillview Ave Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

vmware®