# Virtualizing IBM Lotus Domino 8.5.1 on VMware vSphere 4

***Functional and Performance Test Results***

February 2011

WHITE PAPER

# Table of Contents

# Executive Summary

This document presents the results of testing conducted jointly by VMware and IBM to characterize the performance and functionality of IBM Lotus Domino™ 8.5.1 running on virtual machines using VMware vSphere™ 4.0. The results of performance testing demonstrate that the latest enhancements in Domino 8.5.1 and VMware vSphere 4.0, combined with the latest x86 server processor technologies, deliver performance that is comparable to a native (non-virtualized) environment. Testing was also performed to validate performance with large mail files, Domino iNotes (HTTP) Server, and Domino clustering.

Results of functionality testing demonstrate that the VMware platform delivers a number of benefits to Domino servers that help provide a highly flexible platform at a much lower cost than a comparable physical environment. These benefits include reduced server hardware requirements, live migration with VMware vMotion™, and simplified high availability with VMware HA and VMware FT.

Finally, we present guidance on sizing and best practices, based on the lab environment created by VMware and IBM for this testing, that can be used as a reference for customers considering similar deployments.

# Introduction

IBM Lotus Notes® and Domino provide a complete collaboration solution to help companies work smarter and accelerate business processes. While customers have been successfully virtualizing Notes and Domino on VMware virtual infrastructure for many years, new product releases from both IBM and VMware, combined with the latest server processor technology from Intel and AMD, have led to significant advancements in performance and functionality. Domino's latest 64-bit enabled Windows release, when combined with vSphere 4 (especially with Intel EPT and AMD RVI), delivers substantially improved performance and greater efficiency. The upgrade cycle for Domino 8.5 provides an ideal opportunity for customers to adopt virtualization as their underlying platform and add VMware benefits such as consolidation, simplified DR, hardware independence, VMware vMotion, VMware DRS, and VMware HA. These technologies can provide a number of benefits to Lotus Domino environments including:

- **Reduced cost:** Advances in server processor technology (Intel EPT and AMD RVI) combined with significant performance enhancements in both VMware vSphere 4.0 and Domino 8.5.1 provide unprecedented opportunities for server consolidation, which can drive cost and complexity out of the Domino server infrastructure. Current testing has validated up to 4,000 Domino 8.5.1 users on a virtual machine with a single virtual CPU (vCPU), demonstrating that large environments can now be consolidated onto virtual machines with very efficient resource utilization.
- **Simplified management:** VMware virtual machines are hardware independent, which provides added flexibility and powerful new tools that the Domino Administrator can use to manage infrastructure. Utilizing hardware independent virtual machines to implement Domino server infrastructure makes it easier to perform functions such as server hardware refreshes, Domino server provisioning, testing, troubleshooting, and lifecycle management.
- **Improved service levels**: A virtualization-enabled platform for Domino provides administrators with new options for achieving high availability. VMware HA and VMware Fault Tolerance (FT) can provide out-of-the-box high availability for Domino servers without software clustering. VMware Site Recovery Manager can automate and orchestrate Domino disaster recovery and integrate the failover with other enterprise applications.

VMware and IBM have collaborated to test and validate performance and functionality of the latest Domino Server release running on VMware vSphere 4 to provide guidance to prospective customers. In addition to validating performance, this document introduces some of the key benefits, technical considerations, and resources available to customers who are considering a Domino deployment on VMware vSphere 4.

# VMware vSphere Overview

VMware's robust ESX® hypervisor has made inroads in the computer virtualization industry by allowing x86 servers to run multiple disparate operating systems (and their applications) in isolation on the same server. This allows x86 servers to be more fully utilized, while also providing isolation and encapsulation of separate workloads, which enables IT organizations to be more nimble in how they support and deploy IT infrastructure. VMware vSphere (founded on VMware ESX) is a robust platform that has taken the concepts and capabilities of isolation, encapsulation, and consolidation delivered by ESX to another level. Innovation in virtual deployment technologies has introduced "Zero Downtime" maintenance of enterprise workloads (using VMware vMotion), built-in "High Availability" for all workloads (with VMware HA), "Fault Tolerance" (with VMware FT) and unique "Disaster Recovery" capabilities (with VMware SRM) that have never been easier to achieve, on any platform. VMware and IBM Lotus have collaborated for the very first time around deploying various Lotus technologies with vSphere technology to ensure that everything works seamlessly together and enable Lotus customers to get all the best-of-breed capabilities offered by VMware technologies. For more information about VMware vSphere features, go to: http://www.vmware.com/virtualization/virtual-infrastructure.html.

## Performance Enhancements

There have been a number of technological advancements over the past several years that have dramatically improved the performance of Lotus Domino running in a virtual machine. These advancements have occurred in three primary areas:

1. **vSphere performance enhancements**: vSphere 4 has been optimized from the ground up to make it easier for customers to virtualize their most demanding workloads. These performance enhancements cover all areas of the platform including scalability enhancements, efficiency improvements, and resource management improvements. Customers looking to virtualize Domino can achieve the best possible performance by using the latest vSphere release. For more information go to: http://www.vmware.com/files/pdf/vsphere_performance_wp.pdf.

2. **Domino 8.5.1 performance enhancements**: The latest Domino release has also been optimized to deliver improved performance. Notable improvements[1] include:

   - Up to 38% reduction in CPU utilization
   - Up to 33% reduction in disk IOPS
   - Support for 64-bit platforms

   These enhancements lend themselves well to making Domino a much more efficient workload when running in a virtual machine.

3. **Server CPU enhancements**: For the majority of common workloads, performance in a virtualized environment is close to that in a native environment. The small difference in performance comes from the virtualization of the CPU, the MMU (Memory Management Unit), and I/O devices. In some of their more recent x86 processors, AMD and Intel have provided hardware extensions to help bridge this performance gap. In 2006, both vendors introduced their first-generation hardware support for x86 virtualization with AMD-Virtualization™ (AMD-V™) and Intel VT-x technologies.

   Recently, Intel introduced its second generation of hardware support that incorporates MMU virtualization, called Extended Page Tables (EPT). VMware studies have concluded that EPT-enabled systems can improve performance compared to using shadow paging for MMU virtualization. EPT provides performance gains of up to 48% for MMU-intensive benchmarks and up to 600% for MMU-intensive micro-benchmarks. We have also observed that although EPT increases memory access latencies for a few workloads, this cost can be reduced by effectively using large pages in the guest and the hypervisor.

[1] On physical servers, compared to previous versions of Domino (Domino 8 and earlier) running 4,000 users on Windows Server 2003. See performance results at: http://www.ibm.com/developerworks/lotus/library/domino85-performance/

Our testing with Domino has demonstrated notable improvements in performance when using these new technologies. More details are provided later in this document.

For more information on Intel EPT enhancements, go to:

http://www.vmware.com/pdf/Perf_ESX_Intel-EPT-eval.pdf

Similar performance results would be expected when using AMD RVI processors (as compared to Intel EPT processors); however, servers with these processors were not used during testing described in this document.

# IBM Lotus Domino Overview

IBM Lotus Notes and Domino are enterprise collaboration software used by more than half of the Fortune Global 100. Lotus Notes brings together email, collaboration tools, and business applications within a rich, integrated desktop experience. Lotus Domino provides a world class platform for building, deploying and administering critical business, collaboration, and messaging applications.

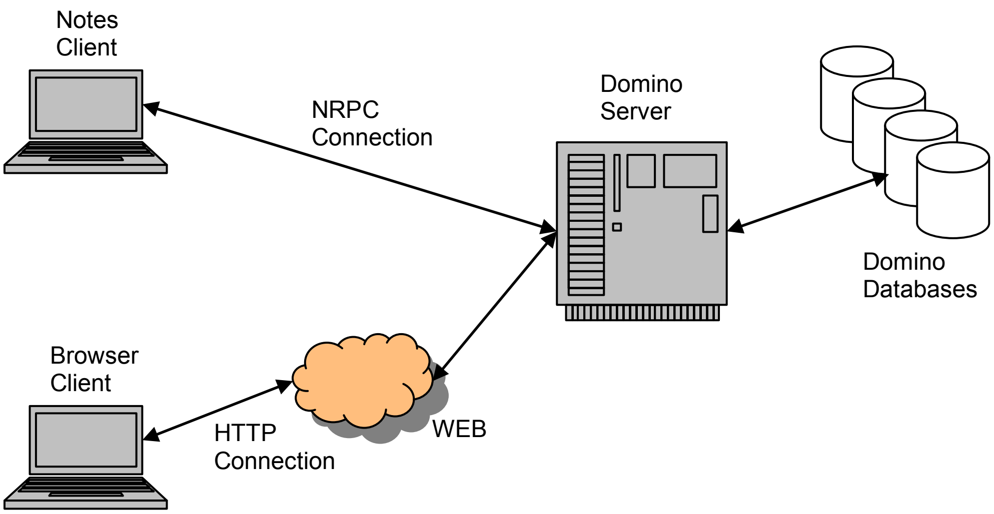


Figure 1. Lotus Domino System Architecture

The version of Domino used for testing and test results described in this paper was Domino 8.5.1, which was the latest version available at the time testing was performed.

# NetApp Storage Overview

All performance and functionality testing described in this document used storage that consisted of NetApp FAS3020 storage arrays running Data ONTAP storage management software. NetApp technologies enable companies to extend their virtual infrastructures to include the benefits of advanced storage virtualization.

The NetApp unified storage architecture provides customers with an agile and scalable storage platform. By using a combination of RAID-DP®, deduplication, FlexVol®, FlexClone®, thin provisioning, FlexShare®, Flash Cache, and Snapshot™ technologies, plus tools such as SnapManager® for Virtual infrastructure (SMVI), VMware vCenter plug-ins Virtual Storage Console (VSC), and the Rapid Cloning Utility (RCU), NetApp enables customers to achieve storage savings and operational efficiency in a virtual environment. For more information on NetApp storage solutions and best practices for storage virtualization with VMware vSphere, go to:

http://media.netapp.com/documents/tr-3749.pdf

# Validation and Performance Testing

The objectives of testing described in this section were to validate IBM Lotus Domino performance and functionality running on the VMware vSphere platform in different configurations and deployment scenarios. Testing was performed jointly by VMware and IBM using the latest versions of VMware vSphere and Lotus Domino, and the latest x86 server technologies. The test results included here supersede any previous performance studies and demonstrate that performance gains due to recent technology advancements have made IBM Lotus Domino an ideal candidate for virtualization on VMware vSphere.

## Lab Environment Overview

The lab environment used for testing was hosted at VMware facilities in Palo Alto, CA. Remote access to this lab was made available to the IBM Lotus Domino Performance Engineering team to assist with test setup, configuration, and analysis of results. The diagram in Figure 2 depicts the environment used for Domino testing:

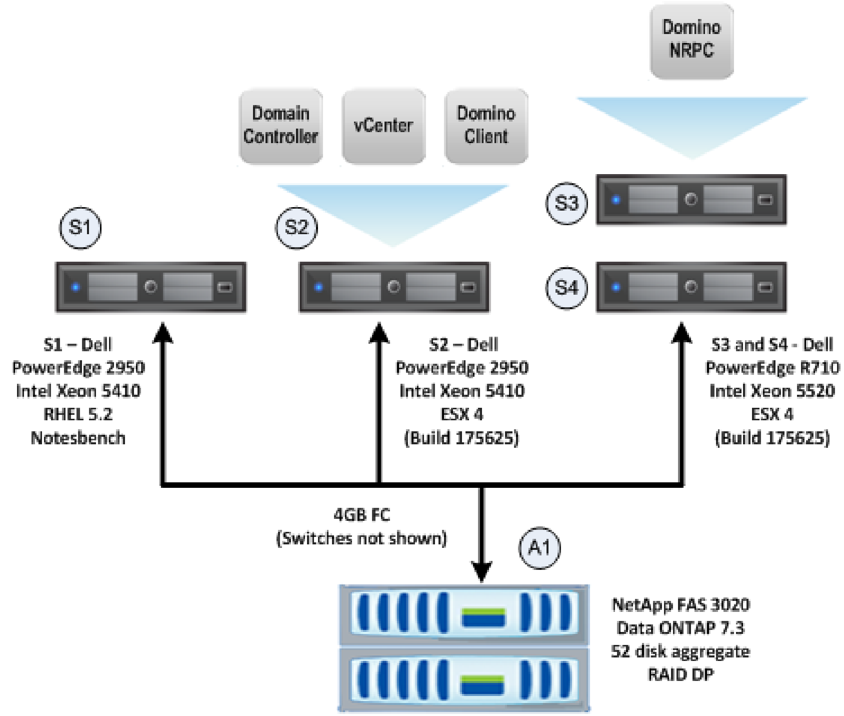


Figure 2. Physical Lab Setup Used for Domino Testing in Palo Alto, CA

### Lab Environment Servers

Figure 2 shows a total of four servers in the test harness configured at the primary site. The function of each server was as follows (starting from the far left):

- **NotesBench Client (S1):** This Dell PowerEdge 2950 system ran the NotesBench load generator and was running a native, non-virtualized installation of Red Hat Enterprise Linux 5.2. Although timing in a virtual machine is expected to be accurate, the load generator was run natively to eliminate any questions about timing.
- **Infrastructure Server (S2):** This Dell PowerEdge 2950 system ran VMware ESX 4.0 and contained virtual machines for the Active Directory domain controller, VMware vCenter™ Server, and a Windows XP virtual machine with the Domino client installed.

- **Systems under test (S3 and S4):** Two Dell PowerEdge R710 servers were used as the primary systems for hosting the Domino mail servers subjected to benchmark testing. These servers had the following configuration:
  – 2 x quad-core Intel Xeon 5520 @ 2.27GHz
  – 32GB RAM
  – Dual-port 4GB fiber channel HBA connect to a NetApp storage array

  **NOTE:** These servers were configured with dual-boot capability between Windows Server 2003 Enterprise Edition SP2 (x64) and VMware ESX 4.0 (build 175625). This provided the ability to quickly switch between a virtualized system and a non-virtualized system so tests could be run on both platforms using identical hardware, for comparison purposes. When booted on ESX 4.0, the virtual machine running Domino 8.5.1 was also configured with Windows Server 2003 Enterprise Edition SP2 (x64).

## Lab Environment Storage and Disk Layout

All performance and functionality testing used storage (A1 in Figure 2) hosted at the lab test site. The storage configuration consisted of:

- A NetApp FAS3020 storage array running Data ONTAP version 7.3.
- The storage array included an aggregate total of 52 disks:
  – *24 x 146GB 10K RPM drives*
  – *28 x 72GB 10K RPM drives*

  All the vmdk files for virtual machine storage layouts were created on VMFS volumes on separate LUNs for each disk partition. Figure 3 below describes the disk layouts that were used in both native and virtual environments where G:\ stored Domino Data and H:\ held the transaction logs. The other four 350 GB LUNs (drives I through L) provided a home for the Domino mail databases, each LUN storing information for 1,000 mail users, for a total of 4,000 users.

**Hard Disk Drives**

| | | |
|---|---|---|
| Local Disk (C:) | Local Disk | 59.9 GB |
| Data (G:) | Local Disk | 28.9 GB |
| Trans (H:) | Local Disk | 58.9 GB |
| Mails1k (I:) | Local Disk | 348 GB |
| Mails2k (J:) | Local Disk | 348 GB |
| Mails3k (K:) | Local Disk | 348 GB |
| Mails4k (L:) | Local Disk | 348 GB |

Figure 3. Disk Layout for Domino Virtual Machine Storage

In addition, all ESX servers were connected to the storage array using dual-port 4GB fiber channel HBAs.

## Lab Network Configuration

Figure 4 shows the networking configuration for the Dell R710 server systems under test.

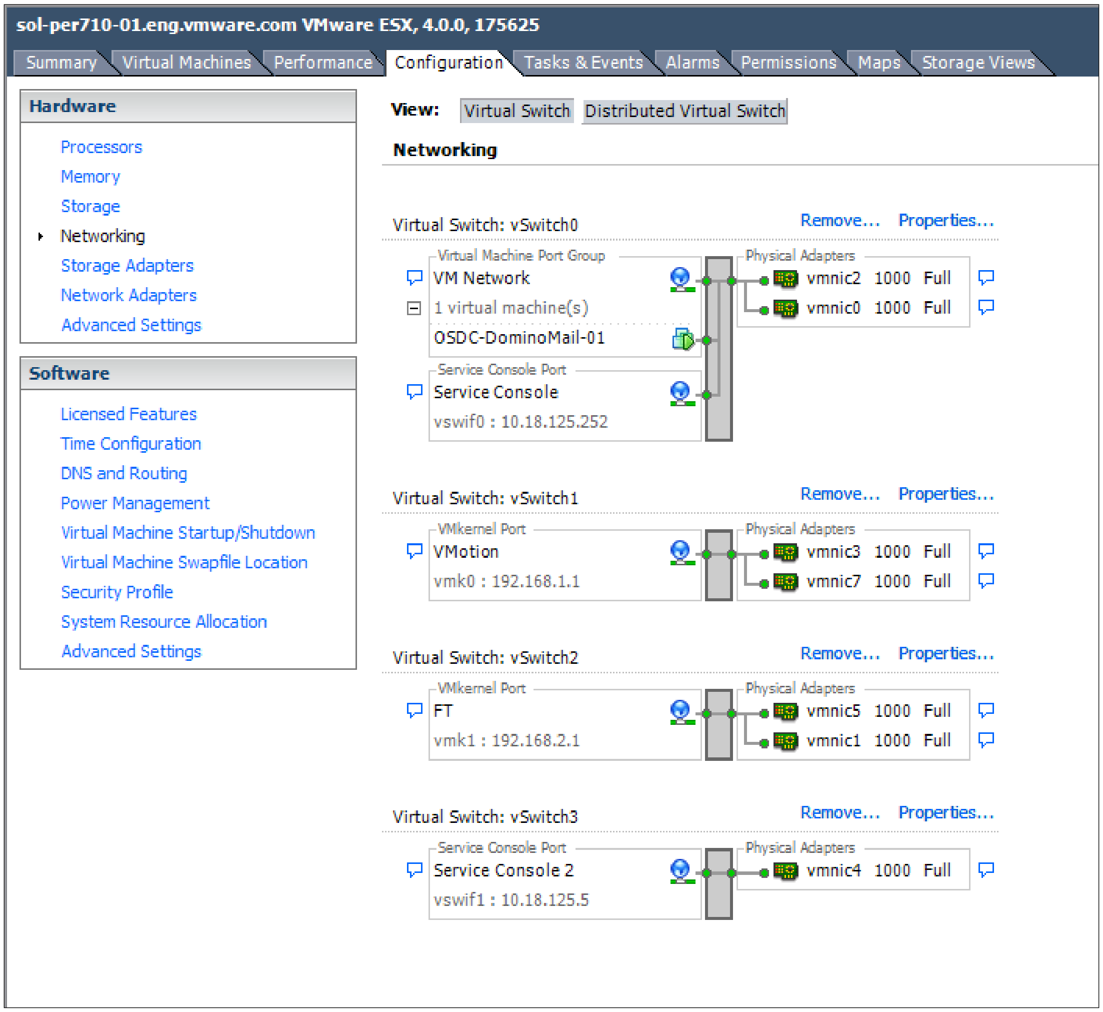


Figure 4. Network Configuration on ESX Servers Hosting Domino Mail Servers

Networking was set up using a total of seven physical NICs running at 1GB/s each:

- **vSwitch 0:** Redundant network team for virtual machines to access public network. Also provided primary service console access.
- **vSwitch 1:** Redundant network team for vMotion traffic. This was set up on a private VLAN. Note that redundant NICs are not required for vMotion.
- **vSwitch 2:** Redundant NICs for VMware Fault Tolerance logging. This was set up on a private VLAN. Note that redundant NICs are also not required for VMware FT.
- **vSwitch 3:** Secondary service console.

## Test Workload Generation

IBM's NotesBench benchmark application running the N85mail and DWA85mail workloads was used to provide loads on the Domino server for the tests described in this paper. (IBM provides a GUI-based test engine version of this benchmark tool called Server.Load that is available as an install option for the Notes Admin client.) Both N85mail and DWA85mail workloads simulate mail user activity by making the same API calls that a Notes client user would perform for the equivalent mail operations. Testing was performed with increments of 1,000 users added over several hours, ending with 4,000 concurrent users.

- The N85mail workload simulates a Notes client user performing mail operations. The DWA85mail workload simulates a browser-based user performing those same mail operations. This browser-based test was used to exercise the HTTP (TCP/IP port 80) capabilities of the Domino server and simulate users accessing their mail with a browser. For both these workloads, each user's mail database is configured as follows:

| | |
|---|---|
| Mail database size | Approximately 250 MB |
| Average document size | 100KB |
| Initial document count | 3,000 |

- The following list describes operations performed per user for both of these workloads over a six hour period:

| Workload actions | Action count |
|---|---|
| Refresh inbox | 24 |
| Read messages | 120 |
| Reply to all | 12 |
| Send message to one recipient | 24 – See note below |
| Send message to three recipients | 12 – See note below |
| Create appointment | 1 |
| Send invitation | 1 |
| Send RSVP | 1 |
| Move to another folder | 24 |
| Delete two documents | 24 |

**NOTE**: Messages varied in size between 500 and 150,000 bytes, with a random number of these messages containing attachments of 525 Kbytes.

# Domino Test Descriptions and Results

The following list describes tests performed to validate IBM Lotus Domino performance and functionality running on the VMware vSphere platform. Tests can be divided into two categories, baseline performance testing and testing of advanced features and functionality.

***Section 1. Baseline Performance Testing:***

- **Domino NRPC (Notes Remote Procedure Call) Physical to Virtual Comparison**: Provides a performance baseline comparing physical-to-virtual performance by examining CPU utilization and response times.
- **Domino NRPC Nehalem/EPT Performance Tests**: Examines performance gains of a virtualized Domino NRPC workload using the latest Intel EPT hardware assist on Xeon 5500 series processors. Compares results with those collected using previous generation processors.
- **Domino NRPC Server Scalability Tests for a Single Domino Virtual Machine**: Characterizes the performance of a virtualized Domino NRPC workload with various configurations of user counts, vCPUs, and memory.
- **Determining the Optimal "Building Block" Configuration for Scaling Out Domino NRPC Workloads**: Identifies an optimal configuration that can be used as a building block for scaling out Domino NRPC virtual machines.
- **Scale-Out Testing Using "Building Block" Configuration:** Characterizes the performance behavior of multiple Domino virtual machines running on a single ESX host using the optimal Domino NRPC "building block" configuration.

***Section 2. Advanced Features and Functionality Testing:***

- **Domino NRPC with vSphere Advanced Features Tests**: Characterizes the impact of deploying Domino NRPC workloads with advanced VMware technologies including VMware vMotion, VMware HA and VMware FT.
- **Domino NRPC Workload Test with Large NSF Files**: Determines the impact on performance of a virtualized Domino mail server when mail file sizes were increased from 250MB to 1GB.
- **Test of Domino NRPC Clustering on VMware vSphere**: Characterizes the behavior of deploying a Domino Cluster (two Domino virtual machines, each with their own Domino partition) across two ESX hosts.
- **Virtualizing Domino iNotes (HTTP) Server Tests**: Characterizes the behavior of running Domino iNotes (HTTP) Server in a virtual machine.

The following sections provide details of each test including detailed test objectives, a description of test setup, test results, observations, and analysis of results. Refer to previous sections for more information on the lab system configurations and workloads used in testing.

# Section 1. Baseline Performance Testing

## Domino NRPC Physical to Virtual Comparison

The main objectives of this test were to measure baseline performance of Domino NRPC workloads on a physical server compared to a virtualized environment and to verify that CPU utilization and response times running on VMware vSphere were comparable to Domino running on a physical server.

To obtain test results, individual tests were run on 1-vCPU, 2-vCPU, and 4-vCPU virtual machine configurations with 16GB RAM against identically-configured physical configurations. The tests executed a 4000-user NRPC workload (using Lotus NotesBench suite) to compare the response times and CPU utilization running natively (on a physical server) versus a VMware ESX 4.0 virtual machine.

Figure 5 provides a summary of the virtual to physical comparisons.

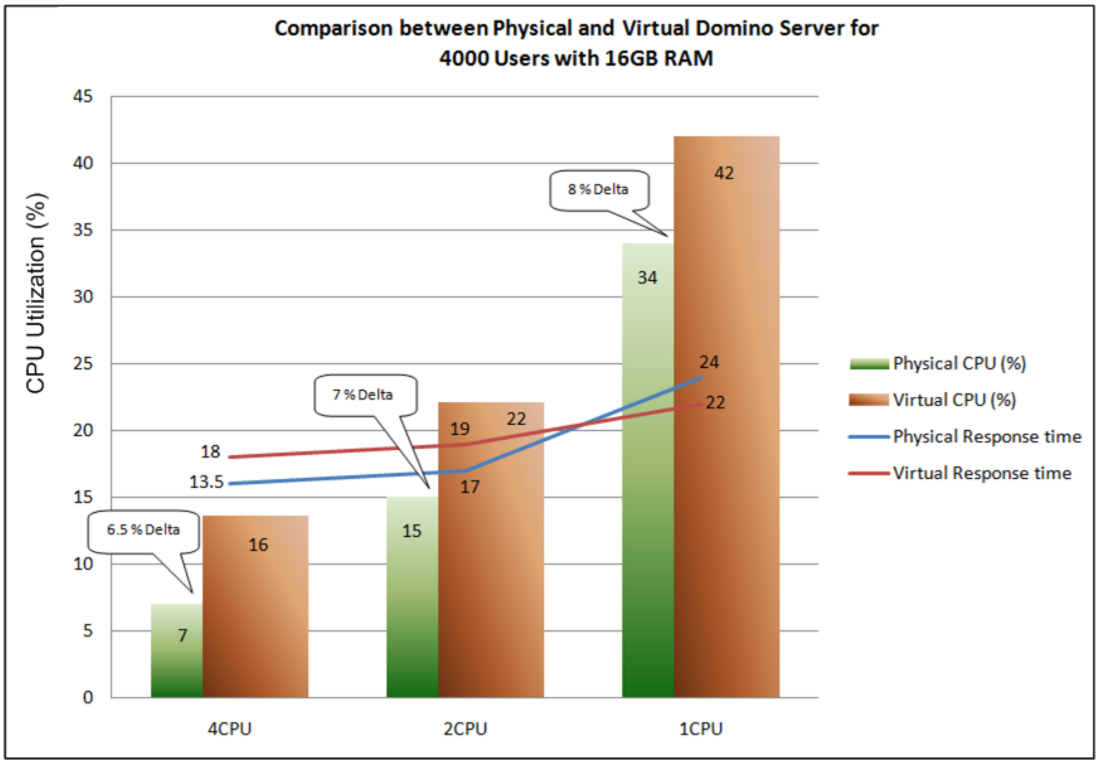


Figure 5. Comparison of Physical and Virtual Domino Server Performance

Here is a summary of observations from the virtual to physical comparison tests:

- Response times for all configurations were exceptional (less than 25ms in all cases) and in compliance with successful Notesbench NRPC test runs of under one second.
- CPU utilization remained well within limits across all configurations and scaled well in both physical and virtual configurations. The difference in CPU utilization between virtual and physical configurations was approximately 7%.

Based on test results, we can also conclude that a single 1-vCPU virtual machine can easily support a workload of 4,000 Domino NRPC users when running on an Intel Xeon 5500 series processor. However, customers will likely choose to allocate additional vCPUs depending on their specific user requirements and to take advantage of Domino's multi-thread design.

## Domino NRPC Nehalem/EPT Performance Tests

The purpose of these tests was to compare the performance of a virtualized Domino NRPC workload using the latest processor technologies (Nehalem and Nehalem with EPT) versus non-Nehalem/EPT processors. (AMD processors provide similar virtualization hardware-assist technology referred to as RVI.)

Historically, Domino NPC workloads have placed larger demands on virtualization (memory, I/O, and MMU utilization) resulting in higher than expected CPU utilization when virtualized. This test compares the CPU utilization across Nehalem and Non-Nehalem processors, with and without EPT, using a workload of 4,000 Domino users and virtual system resources of four vCPUs and 16GB RAM.

Figure 6 provides a summary of the Nehalem versus non-Nehalem test results.

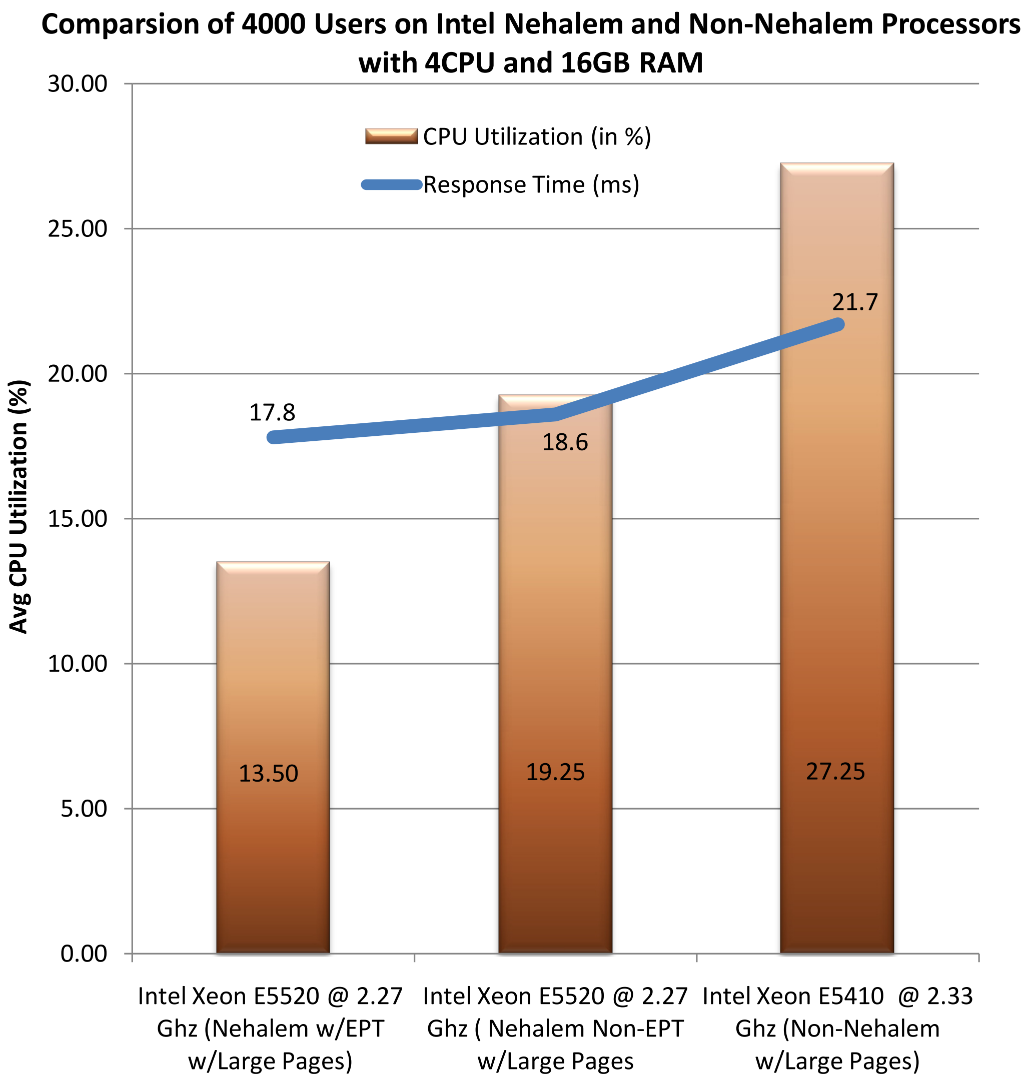


Figure 6. Comparison of Domino Performance on Nehalem Versus Non-Nehalem Processors

From the test results, we observed the following:

- Using Intel Xeon (Nehalem) processors with EPT enabled and large memory pages[2] provided the optimal configuration in terms of CPU utilization (13.5%) and response times (17.5ms).
- Using the same processor with EPT disabled increased CPU utilization by approximately 30%. While response times remained comparable, this demonstrates the benefits in CPU efficiency that can be obtained by using virtualization enhancements built into the latest Intel and AMD processors.
- Using earlier generation processors (Xeon 5400 series in this case) still provided comparable response times but resulted in even higher CPU consumption. Customers wishing to use these processors can expect acceptable response times, however, they should allow for more CPU headroom when sizing their solutions.

In conclusion, test results showed that Nehalem processors (with EPT and large memory pages enabled) provided the most optimal CPU configuration for deploying Domino NRPC workloads.

## Domino NRPC Server Scalability Tests for a Single Domino Virtual Machine

The purpose of this test was to characterize the performance of a virtualized Domino server with different-size Domino NRPC workloads and system resource configurations:

- 1000-user workload with one vCPU and 4GB RAM
- 2000-user workload with two vCPUs and 8GB RAM
- 4000-user workload with four vCPUs and 16GB RAM

Figure 7 provides a summary of the results for these tests.

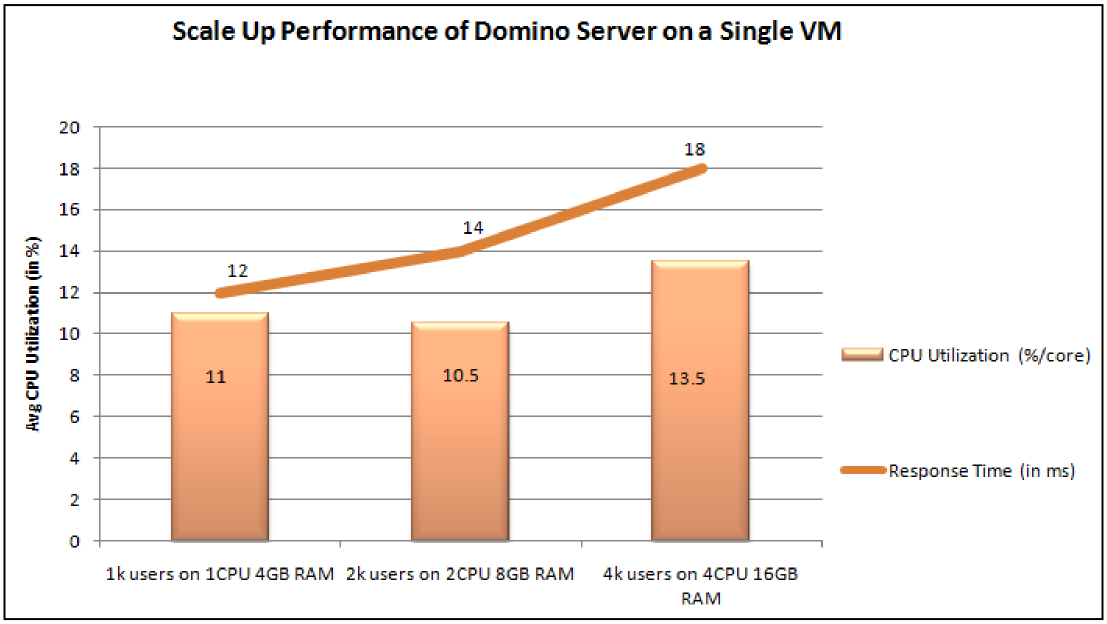


Figure 7. Scaling Up a Domino Virtual Machine with Different Workloads and System Resources

[2] To enable large page support in Windows, the system administrator must grant appropriate users the privilege to "Lock pages in memory" (in Windows Server 2003, the Control Panel > Administrative Tools > Local Security Policy > Local Policies > User Rights Assignment security policy setting). For more information, see the following document: http://www.vmware.com/files/pdf/large_pg_performance.pdf

From the test results, we can observe the following:

- For smaller environments, virtual machines can be configured with less RAM and lower vCPU counts.
- Virtual machines can be easily scaled as an environment grows by adding additional vCPUs and RAM to the virtual machine. This means customers can avoid over-provisioning of server CPU and RAM resources and scale resources only as the environment grows.
- For environments beyond 4,000 users, VMware and IBM recommend adding additional virtual machines in a scale-out model. Additional information on this approach is covered in the next sections.

## Determining the Optimal "Building Block" Configuration for Scaling Out Domino NRPC Workloads

The purpose of these tests was to identify the optimal "building block" configuration for scaling out larger Domino environments. Per IBM recommendations, 4,000 users is the upper end for a single Domino server to handle. Beyond 4,000 users, it is recommended that customers scale out by adding additional virtual machines. Thus, in this section, we wanted to identify the optimal CPU and RAM configuration for supporting a 4000-user "building block" virtual machine.

Figure 8 provides a summary of the results for virtual SMP testing on a 4000-user virtual machine.

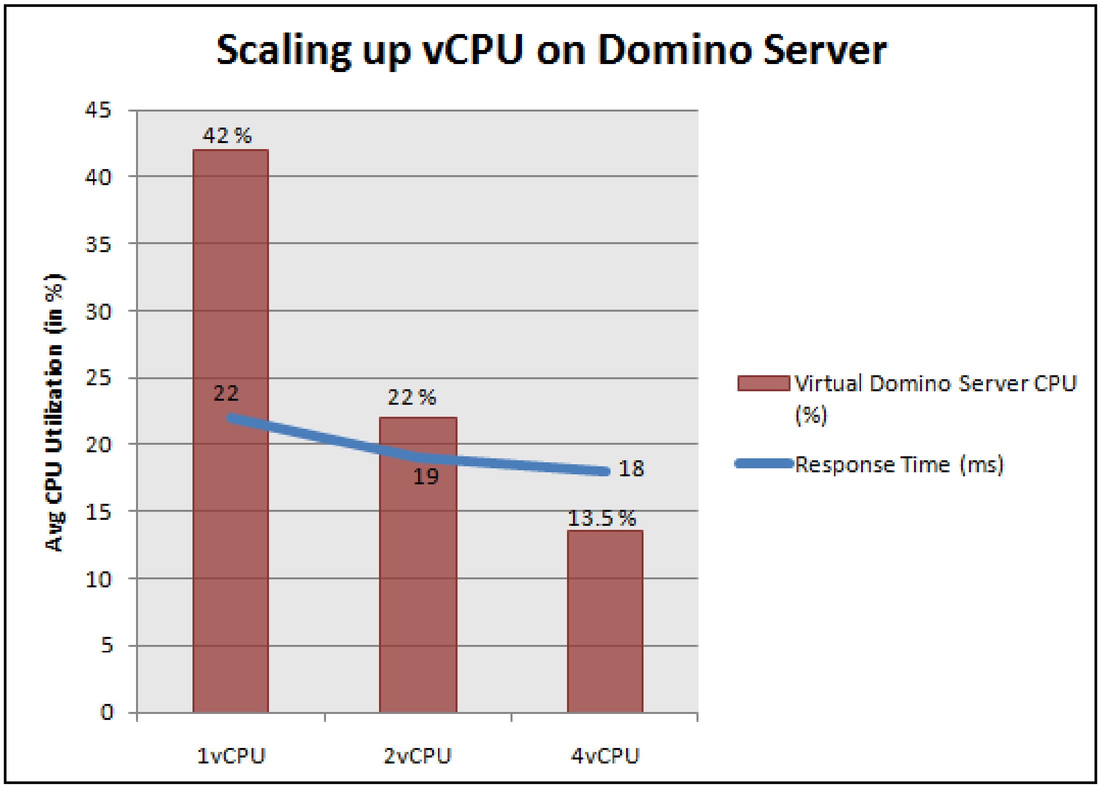


Figure 8. Domino Server Performance with vCPU Scaling

Keeping the N85mail workload (4,000 users) and RAM (16GB) constant and changing only the number of vCPUs, test results showed the following:

- The 1-vCPU configuration resulted in 42% CPU utilization, which is on the high side.
- The 4-vCPU configuration had the lowest CPU utilization, but did not make the most efficient use of all four processors. Thus, the ideal configuration for 4000-user workloads was two vCPUs, which provided the best balance between overall performance and CPU efficiency.

Figure 9 show the impact of RAM on disk IOPS and response times for a 4000-user virtual machine.

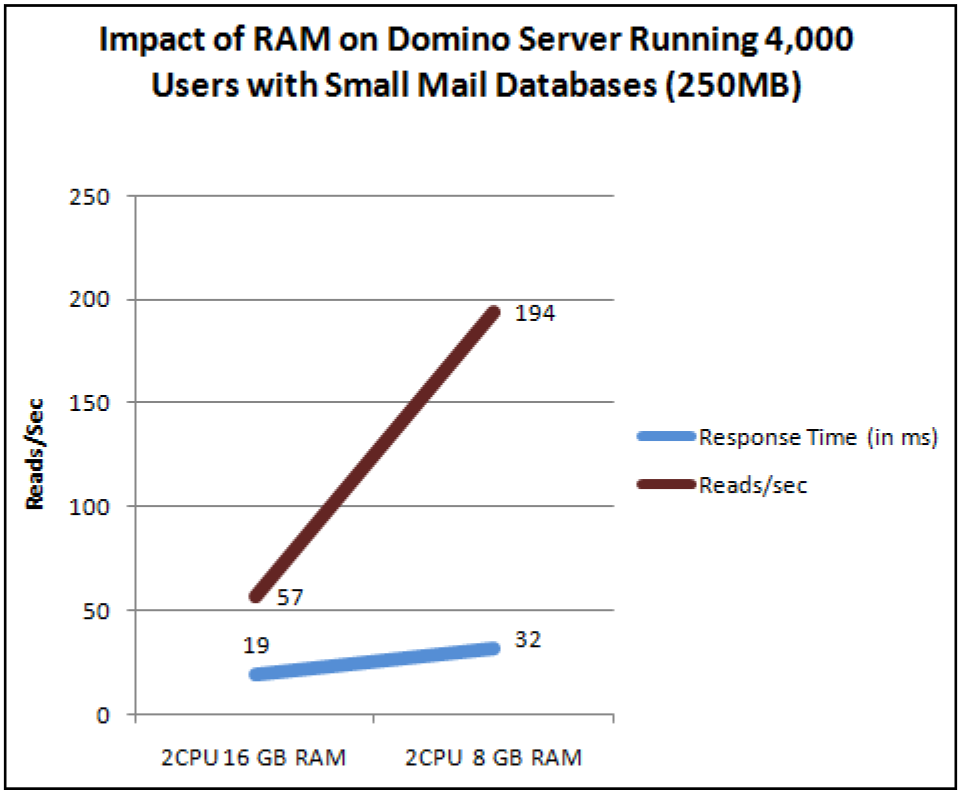


Figure 9. Impact of Memory on Disk IOPS, and Response Times (Small Mail Databases)

From the test results, we observed the following:

- By allocating more memory to the virtual machine, Domino can cache additional mail database data, thus reducing IOPS and improving response times. The graph in Figure 9 shows that as we reduce the memory from 16GB to 8GB, the number of read I/O's increased by nearly 3x, impacting the response time, which increased from 19 to 32ms.
- As long as the storage subsystem is properly configured, 8GB RAM yielded the best results for both price and performance, and thus is the recommended 4000-user "building block" configuration.
- We also noted that if the disk sub-system is under-performing, we could allocate more memory to offset I/O demand and achieve acceptable response times.

The tests provided a good learning tool to help understand the effects of CPU resources and memory size on Domino performance and response times. Based on the results above, we can conclude that a Domino virtual machine with two vCPUs and 8GB RAM represents a viable configuration for scaling out environments in 4000-user virtual machine increments. In the next section we observe the impact of scaling out two building block virtual machine configurations on a single ESX server.

## Scale-Out Testing Using "Building Block" Configuration

Based on the optimal building block reference configuration we designated above, additional testing was performed to scale out to two virtual machines on a single ESX host. Lotus customers want assurance that they can scale out Domino servers on a single VMware host and achieve acceptable scalability. For this test, VMware housed two separate Domino mail servers running in two separate virtual machines on the same ESX host server.

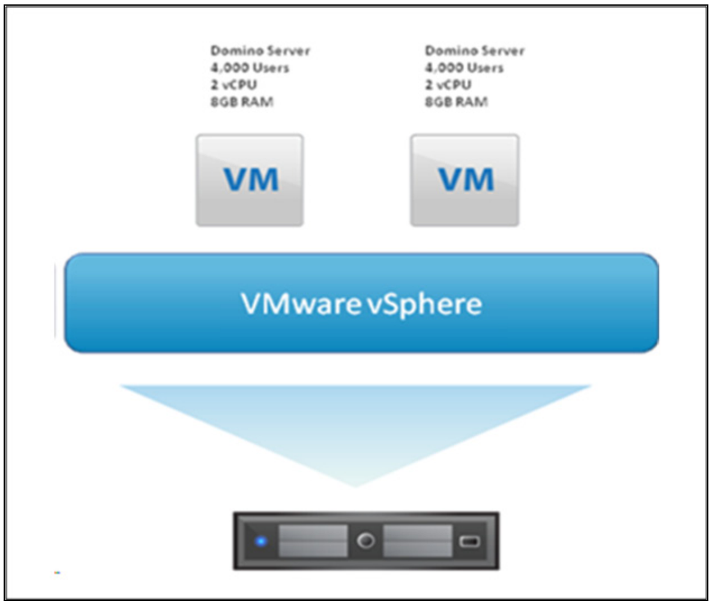


Figure 10. Scale-Out Testing Using Optimal Building-Block Configuration

In this test, each virtual machine running Domino was pushed to a load of 4,000 simulated NRPC mail users, or a total of 8,000 simulated users on the ESX host.

To create the desired test scenario, the second Domino virtual machine was created by cloning the first Domino server instance. Both Domino servers were running with two vCPUs and 8GB each, the optimal building block reference configuration chosen earlier. The disk layout remained the same throughout (as described earlier in the "Lab Environment Storage and Disk Layout" section).

The Domino servers were set up as members of the same Domino domain, but not cluster members. Each Domino server was set up with 4,000 mail users and their databases. The Domino directory on both of the servers was a replica copy, since they were in the same Domino domain. Therefore, it supported a total workload of 8,000 users, with 4,000 users assigned to one server and another 4,000 users assigned to the other server.

Table 1 shows the resource utilization for the two virtual machines running on a single ESX host.

Table 1. Virtual Machine Resource Profile for Domino Server Scale-Out Test

| | CPU | | | MEMORY | | NETWORK | | DISK | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | (%) | | | (MB) | | PACKETS/SEC | | I/O SEC | | LATENCY(MS) | |
| VIRTUAL MACHINE | AVG | MAX | READY (AVG/MAX) | AVG | MAX | RECD | TXMIT | AVG | MAX | AVG | MAX |
| OSDC-DominoMail-02 | 28.22 | 66.59 | 0.16/2.795 | 3338.66 | 4014.08 | 1961.98 | 1050.34 | 743.97 | 1993.08 | 4.6 | 10.08 |
| OSDC-DominoMail-01 | 29.26 | 50.385 | 0.15/2.775 | 3276.57 | 4014.08 | 1964.28 | 1047.64 | 698.68 | 1744.92 | 4.64 | 10.77 |

As we scale out to two virtual machines, each supporting a workload of 4,000 users, all system resources on the virtual machines remained consistent. The workload run against the servers was the n85mail workload, ramping up to 4,000 users on each server, while allowing users to send mail to any of the total 8,000 users. To emulate a more realistic customer environment, a small percentage of mail messages were routed between the two systems. The workload was ramped up to 4,000 users on one virtual machine in 1000-user increments, followed by ramp-up to 4,000 users on the second virtual machine.

Looking at the graph in Figure 11, you can see the steps of the ramp-up for the first 4,000 users is larger in %CPU gain, transaction rate, etc. than the ramp-up from 4,000 to 8,000 users. This occurs because in this test the first 4,000 users are only on the first server (and this behavior matches that of a regular 4000-user NRPC test with this configuration). The second 4000-user increment, 4,000 to 8,000 users, reflects the impact of the extra mail being delivered to the users on the first server from the second test server as it is ramping up to its 4,000 users. (Design of the test allowed simulated users to send mail to any of the 8,000 users on either of the two servers.)

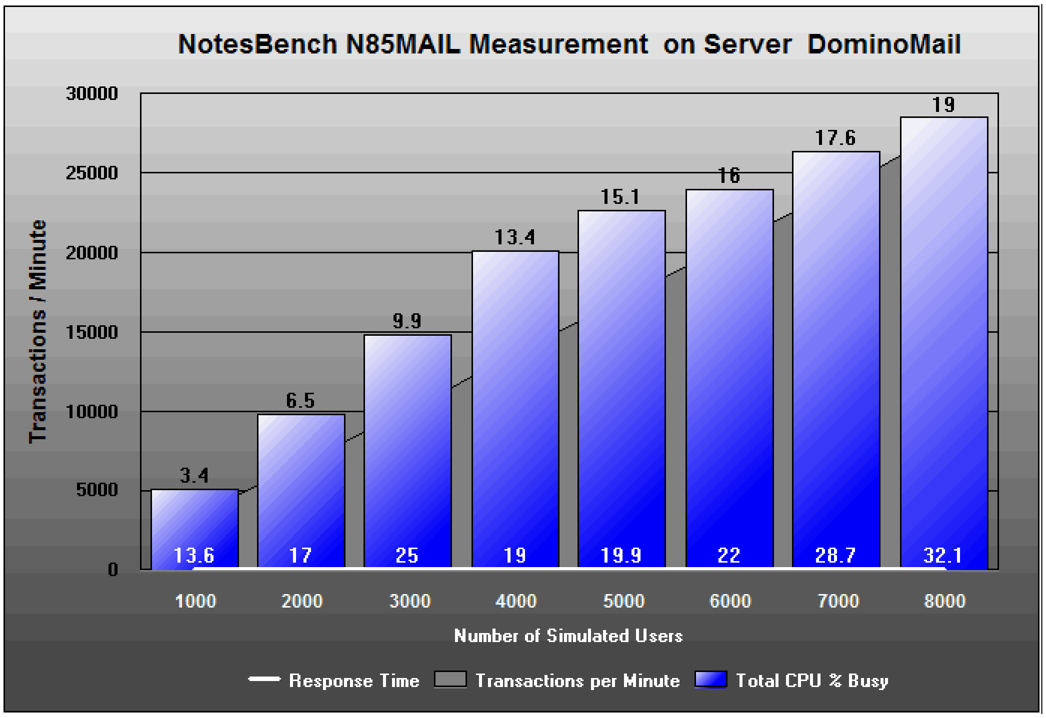


Figure 11. Response Time and Transaction Rates for Domino Server Scale-Out Test

From the results above, we can observe that use of CPU resources remain efficient and response times[3] remain under 33ms–demonstrating excellent performance as we scale to 8,000 users on a single ESX server. Depending on the number of CPU cores and RAM on a given server, customers can easily scale out Domino virtual machines to accommodate large numbers on users with limited server hardware requirements.

# Section 2. Advanced Features and Functionality Testing

## Domino NRPC with vSphere Advanced Features Test

Once an optimal "building block" virtual machine configuration was identified, this configuration was used to validate the functionality and behavior of VMware vMotion, VMware HA and VMware FT when Domino virtual machine are placed under heavy stress from a NotesBench workload.

### *VMware vMotion Test*

VMware vMotion provides the ability to move a live, running virtual machine from one ESX host to another, with no downtime.  Because virtual machines running on vSphere are abstracted from the underlying hardware, vMotion can even move virtual machines across hardware from different vendors and between physical machines or servers having different hardware configurations as long as the CPUs meet compatibility requirements.[4]

VMware vMotion can be an invaluable tool for the Domino administrator as it presents capabilities such as the following:

- **Avoid planned downtime:** Move Domino virtual machines off of an ESX host that requires downtime (hardware replacement/upgrade, firmware upgrade, etc.) with no loss of service.
- **Simplify server refresh cycles:** Server refresh cycles can be challenging as the application and operating system typically need to be re-installed. With vMotion, moving a Domino server onto new hardware can be done in minutes, with no downtime.
- **Troubleshooting:** Moving Domino virtual machines onto a different ESX host can be an effective tool for troubleshooting suspected issues with underlying hardware.

Figure 12 shows the behavior of the Domino virtual machine when subjected to a total of eight vMotion operations during the test run. Since the Notesbench workload gradually scales up the number of users during the test run, we ran vMotion tests at multiple points:

- Two vMotion operations at 1,000 users
- Two vMotion operations at 2,000 users
- Two vMotion operations at 3,000 users
- Two vMotion operations at 4,000 users

Objectives of the tests were to validate we could maintain sub-second response times without disconnects.

---

[3] Response times (in milliseconds) are the white numbers displayed just above the x-axis along the bottom of the graph.

[4] The only requirement for vMotion migration is that all servers in the ESX cluster must meet certain CPU compatibility requirements.  See the following web location for more information:

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1003212

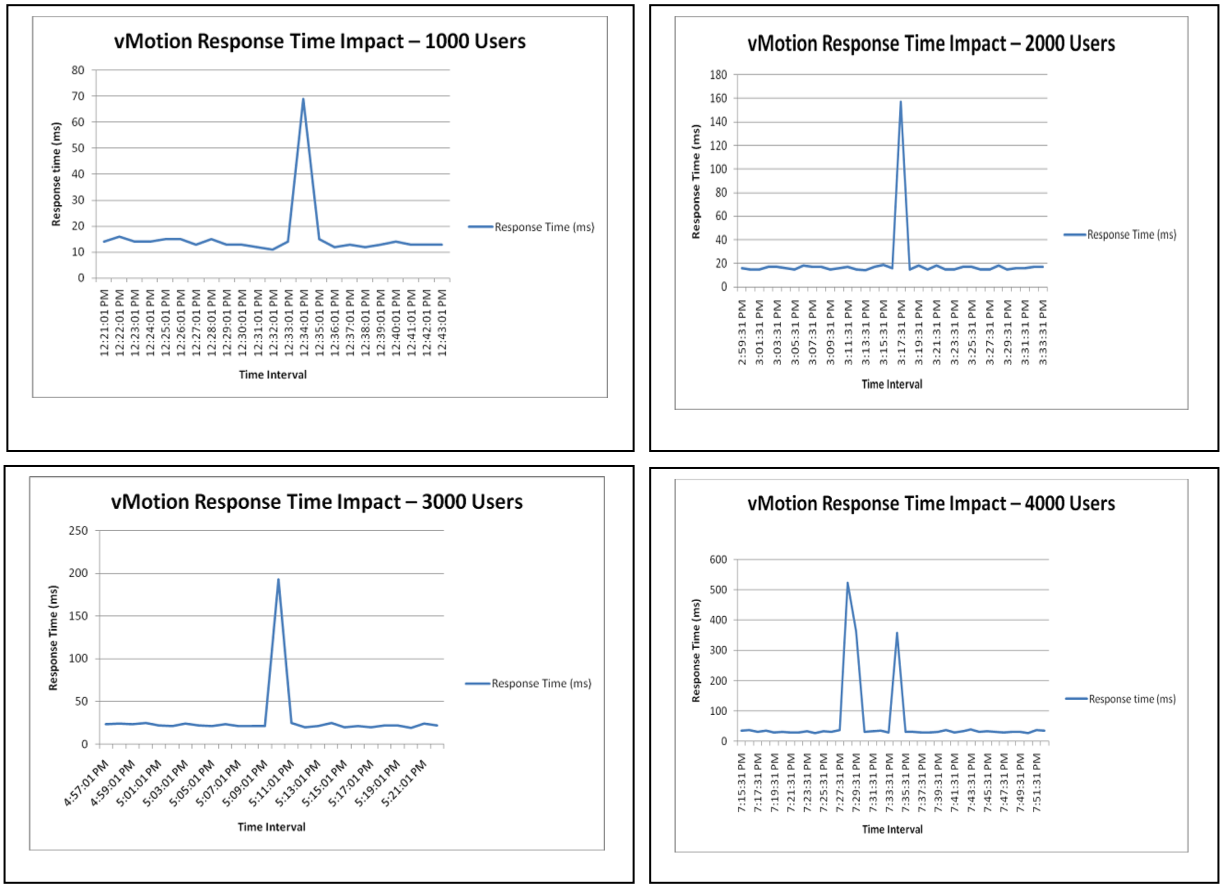


Figure 12. vMotion Results for Domino Mail Server with Workloads of 1,000, 2,000, 3,000, and 4,000 Users

For all vMotion tests:

- Response times remained “sub-second” during migration of workloads.
- There were no disconnects of the Domino NotesBench sessions during migrations.
- Loads continued to run seamlessly across migrations.

In all test cases, the vMotion migrations did not affect the NotesBench workload and sub-second response times were maintained. These results validate vMotion as an effective tool for migrating Domino servers across physical ESX hosts for various reasons: load balancing, avoiding planned downtime, and server refresh cycles.

### *VMware HA Test*

VMware HA provides an automatic rapid restart of all virtual machines running on a given ESX host in the event of server hardware failure. This provides a simple and effective high availability solution for any virtual machine, regardless of the application or operating system, and without requiring implementation of any software clustering solutions.

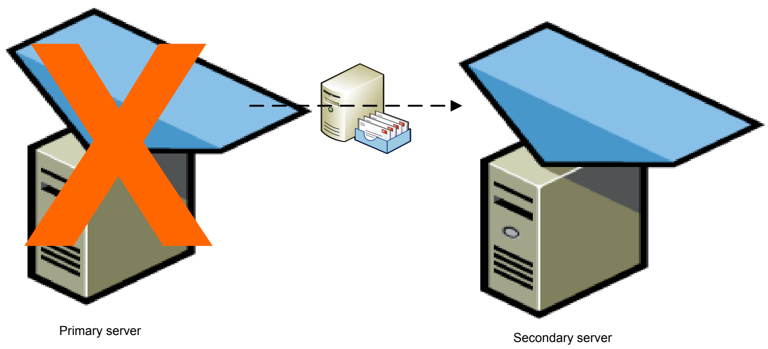


Figure 13. Domino Virtual Machine Restart on Another ESX Host after Server Hardware Failure

In the Domino environment, VMware HA can be used on the mail servers for customers who may prefer using a stand-alone Domino server, instead of using Domino clustering. Since VMware HA reboots the virtual machines on a surviving host during a failure, downtime is usually measured in minutes while the guest OS restarts and Domino services come back online. Alternatively, VMware HA can be combined with Domino clustering solutions to provide even higher levels of availability. Domino clustering provides high availability at the application layer while VMware HA provides high availability at the server or host layer. Thus, using this combination can eliminate downtime and vulnerability completely during server hardware failure.

Finally, VMware HA can also be used to protect all the other servers in a Domino environment, providing them with simple, low-cost high availability. These systems include all the other application servers in the Domino environment, for example, BlackBerry Enterprise Servers, Active Directory servers, and others. Protecting these systems with VMware HA can eliminate many single points of failure in the Domino environment and improve overall availability and service levels.

Figure 14 shows the behavior of VMware HA during a server hardware failure. To simulate a failure, one server hosting the Domino virtual machine was powered off while the NotesBench workload was driving a 4000-user workload.

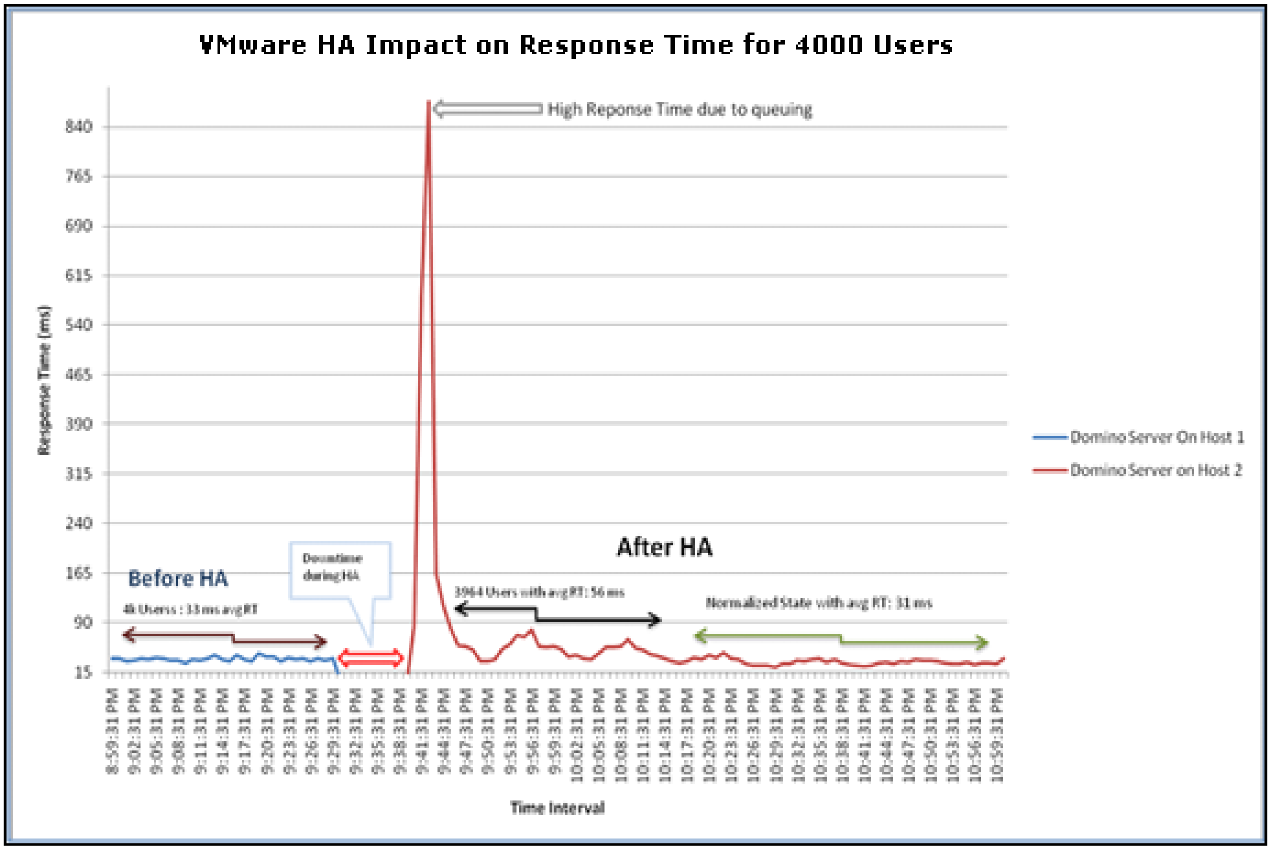


Figure 14. Effects of VMware HA on a 4000-User Domino Virtual Machine during Server Failure

The graph in Figure 14 shows the workload driving 4,000 users and maintaining an average response time of approximately 33ms. At 9:30 PM, the server is powered off to simulate an outage due to hardware failure. Transactions begin to queue on the Notesbench client, which is waiting for the Domino Server and virtual machine to come back online and begin processing messages. At 9:39 PM, the virtual machine has restarted on the surviving ESX host, all Domino services have started, and the server has begun processing messages again. Response times spike up as high as 850ms as the server works through user authentication and the high volume of messages from the Notesbench client. By 9:47 PM, response times are back under 50ms and by 10:14 PM have worked their way back down to the 31ms average observed before the failure.

**NOTE:** The impact of the virtual machine reboot on users can be minimized by configuring Lotus Notes users with local mail files on their clients. In this situation, users will still be able to access their mail database and even send messages during a VMware HA failover. Outbound messages will be queued in the clients local mail.box file and sent out once the server comes back online. In this situation, the failover will not impact users other than introducing a slight delay in receiving mail while the virtual machine fails over.

To summarize observations about the graph for the VMware HA test shown in Figure 14:

- Response times increased at time of failure for approximately 5 minutes before VMware HA restarted the Domino Server on a second ESX host.
- Once the Domino virtual machine rebooted and Domino Server was brought back online, the NotesBench load reconnected and ramped back up to nearly peak load.
- Response times peaked initially while Domino went through its logs as part of recovery, but the response time eventually settled back down to lower levels.

Note that VMware HA is not application-aware and it cannot be used as a failover mechanism to provide capabilities such as no-downtime OS and application patching. It can, however, provide a mechanism to protect against server hardware failure where minimal downtime might be acceptable. VMware HA can also detect failures within the operating system through VMware Tools and take corrective action if the operating system becomes unresponsive. However, it does not have any view into what is happening with the Domino application itself. For customers who require application-level failover capability (for example, to patch software with no downtime), it is recommended that you use Domino clustering. Testing with Domino clustering in virtual machines is described later in this document.

### *VMware FT Test*

Similar to VMware HA, VMware Fault Tolerance (FT) can also protect Domino servers from a failure of the underlying ESX host, but VMware FT uses a mechanism for protecting virtual machines that does not result in any downtime during failover. Figure 15 shows how VMware FT protects virtual machines by creating a shadow image of the protected virtual machine on a separate ESX server.

**NOTE:** VMware FT is currently only supported on virtual machines with a single virtual CPU.

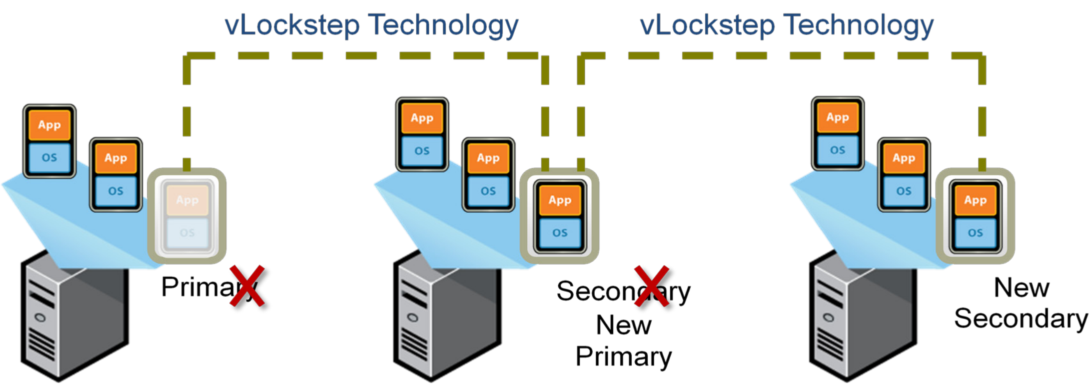


Figure 15. VMware FT Provides Zero Downtime and Zero Data Loss for Mission-Critical Applications

When you enable VMware FT on a virtual machine in an HA-enabled cluster, it creates a duplicate, secondary copy of the virtual machine on a different host. Record/Replay technology then records all operations executed on the primary virtual machine and replays them on the secondary instance. vLockstep technology ensures the two copies stay synchronized, which allows the workload to run on two different ESX/ESXi hosts simultaneously. To the external world, the two virtual machines appear as a single virtual machine. That is, they have one IP address, one MAC address, and you need only manage the primary virtual machine.

Heartbeats and replay information allow the virtual machines to communicate continuously with each other to monitor each other's status. If a failure is detected, VMware FT creates a new copy of the virtual machine on another host in the cluster. If the failed virtual machine is currently the active primary machine, the secondary machine takes over and a new secondary virtual machine is established. If the secondary virtual machine fails, another secondary machine is created to replace the one that was lost.

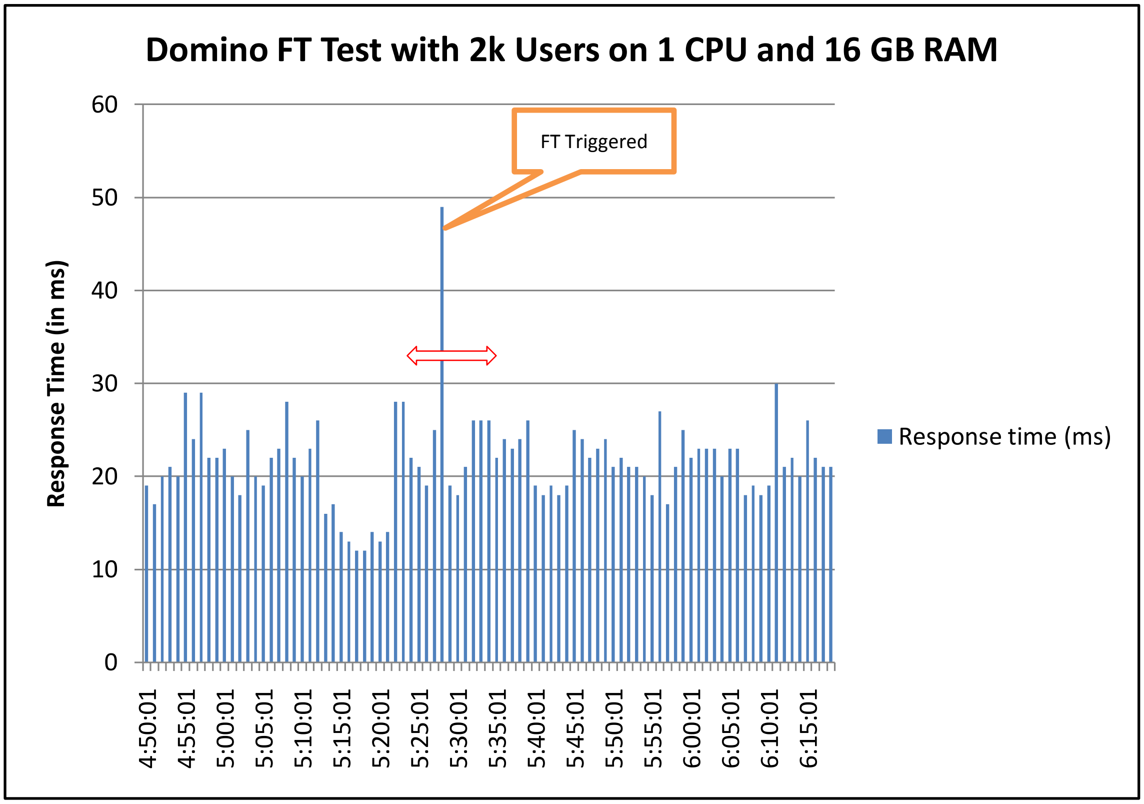


Figure 16. Response Times for the Domino VMware FT Test

From the graph above, you can see that VMware FT provides extremely rapid failover of the Domino virtual machine during hardware failure. Actual failover time was measured in seconds and was completely seamless to the NotesBench client.

To summarize the behavior shown in Figure 16, VMware FT worked as intended. When the primary server that the Domino virtual machine was running on failed, from the client side, there was only a slight spike in response time during the recovery to the shadow or secondary virtual machine. Thus, the results show that VMware FT can provide a simple solution for any single-vCPU server in a Domino environment to provide high availability at the server level with zero downtime and zero data loss.

VMware FT is not application-aware, so it cannot be used to provide failover for operations such as no-downtime OS and application patching. It can, however, provide a mechanism to protect against server hardware failure where minimal downtime might be acceptable. For customers who require application-level failover capability (for example, to patch software with no downtime), it is recommended you use Domino clustering. Testing with Domino clustering in virtual machines is covered later in this document.

**NOTE:** In our lab environment, we had a redundant 1GB connection between the ESX hosts to handle the VMware FT logging traffic. Due to the heavy CPU, memory, network, and disk activity generated by NotesBench simulating a workload of 4,000 active users, we approached the saturation point of the 1GB VMware FT logging network. In some instances, VMware FT was not able to spawn the secondary virtual machine under this heavy load. For this reason, we recommend limiting workloads to 2,000 users for VMware FT when using a 1GB connection for VMware FT logging. Alternatively, you can use a 10GB connection for VMware FT logging, which will provide scalability to support a higher number of user counts. (This configuration was not tested as part of this project.)

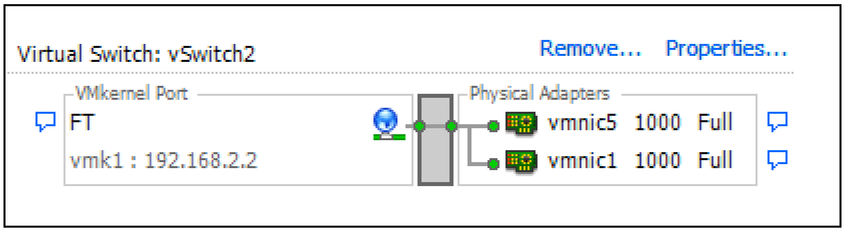


Figure 17. Lab environment showing 1GB connection for FT logging

The second NIC is provided for redundancy purposes only. VMware FT does not support teaming multiple NICs for increased bandwidth.

## Domino NRPC Workload Testing with Large NSF Files

Lotus Domino Mail Server has the ability to support thousands of Notes client users with varying mail database sizes and content. The purpose of this test was to determine performance behavior when deploying Domino NRPC with larger user mail databases. All previous tests described in this document used 250MB NSF mail databases. The goal of this test was to increase mail database size to 1GB and ensure that overall performance and response times were not impacted.

Table 2 below shows the CPU utilization and response times for Domino NRPC workloads using 1 GB NSF mail databases and compares results to the same tests with 250MB NSF databases.

Table 2. CPU Utilization and Response Times for Domino Workloads using 250MB and 1 GB NSF Mail Files for Domino Virtual Machine with 2 vCPU and 8GB RAM

| | CPU UTILIZATION (AVERAGE) | RESPONSE TIME (MS) | READS/SEC |
|---|---|---|---|
| Domino Mail Server (250MB NSF files) | 24 | 32 | 194 |
| Domino Mail Server (1GB NSF files) | 25 | 54.7 | 364 |

From the results shown in the table, we could observe the following:

- CPU utilization and response times were well within acceptable ranges for both the smaller NSF mail file tests and the larger 1GB NSF file tests.
- The use of larger mail files resulted in increased disk activity as reads per second increased from 194 to 364. Customers should size storage accordingly or consider adding additional RAM to the virtual machine. (See Figure 18 below.)

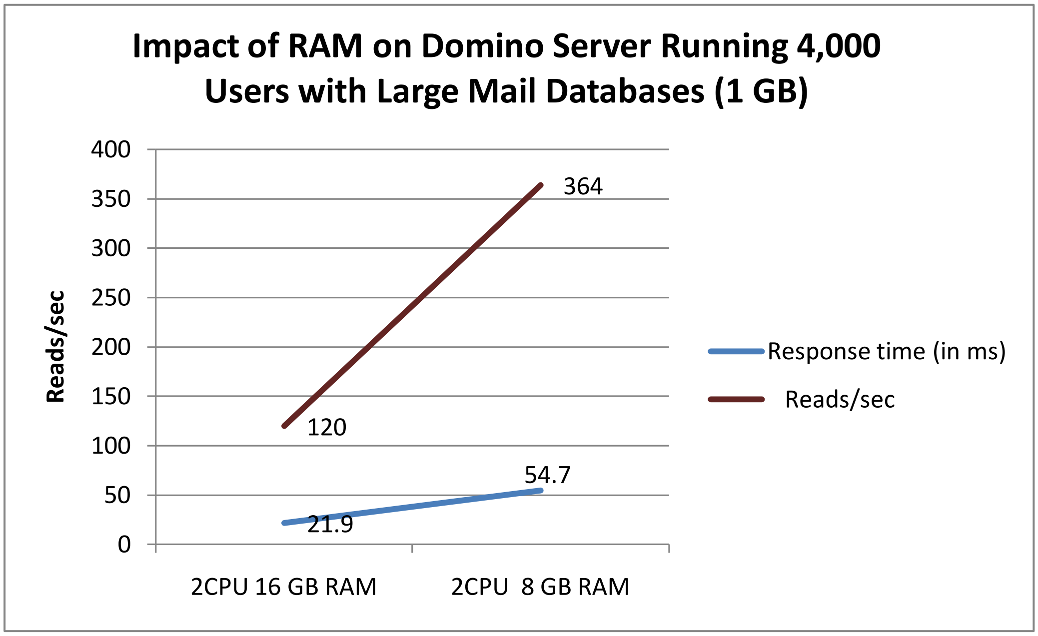


Figure 18. Impact of Adding Additional RAM with Large NSF Files

As shown in the graph, when we increased RAM from 8GB to 16GB for the large mail file tests, we were able to offset the increase in storage IOPS and reduce response times considerably.

## Test of Domino NRPC Clustering on VMware vSphere

The purpose of this test was to characterize the behavior of deploying a Domino Cluster across two ESX hosts. For the test setup, two identical physical servers were configured with VMware ESX and a virtual machine was installed on each ESX host. The virtual machines were configured to run Windows 2003, 64-bit with two vCPUs and 16GB of RAM. Domino was installed on both virtual machines and configured to work as Domino cluster members. We pushed the cluster workload beyond normal recommendations for the Domino configuration to increase our confidence that the VMware virtual infrastructure could handle normal cluster loads.

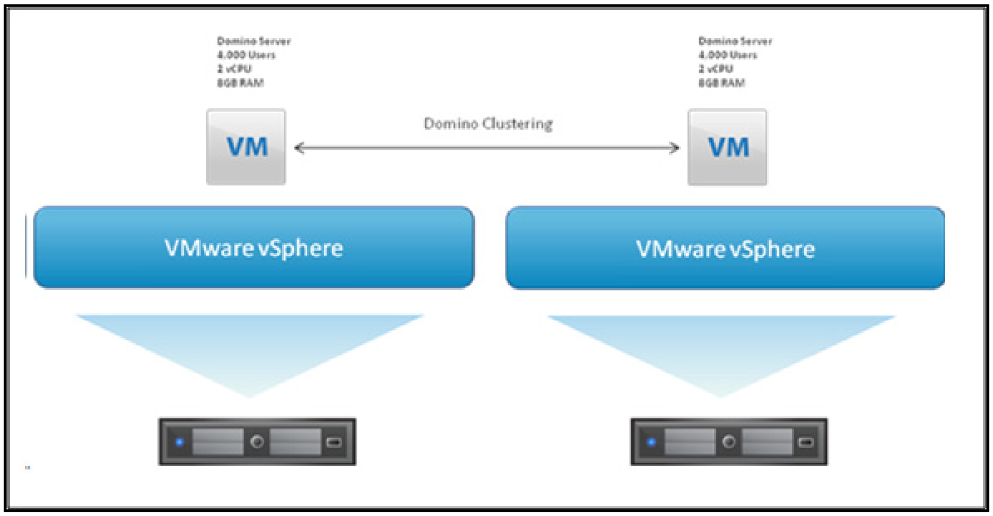


Figure 19. Domino Clustering across Two Virtual Machines Running on Separate ESX Servers

The test environment we used consisted of 8,000 mail users and 8,000 mail databases on each server. Cluster replication for all 8,000 mail databases was enabled as well as mail routing between both of the servers. The n85mail workload was run against the two servers, ramping up to 4,000 users on each server in 1000-user steps. This provided a more intense test and workload than would normally be seen in a typical Domino cluster configuration of 2,000 active users (with 4,000 maximum in failover mode), but we wanted to prove there were no limitations caused by VMware ESX. Table 3 lists key performance metrics recorded at the maximum load levels applied during the Domino Clustering test.

Table 3. Key Performance Metrics for Domino Clustering Tests

| CPU AVG (PERFMON) | AVG. RESPONSE TIME | DISK IO/SEC | NETWORK BYTES/SEC |
|---|---|---|---|
| 32% | 61ms | 1130 | 6MB/sec |

As you can see from the table, the system performed well with response times significantly below the 1-second limit. The results provided here may help you determine sizing requirements and determine how many virtual machines with Domino clusters a configuration can handle. Table 4 below should help you further with system sizing and provide more information on the resource impact of a cluster environment over a non-clustered Domino server. The information should also help you estimate the resources required to go from 2,000 active users to 4,000 active users.

Table 4. Key Sizing Guidelines for Domino Clusters with 2,000 and 4,000 Users

**2,000 Users:**

| | VCPUS | RAM | CPU AVG | RESPONSE TIME | DISK IO/SEC | NETWORK BYTES/SEC |
|---|---|---|---|---|---|---|
| non cluster | 2 | 16GB | 6% | 12ms | 298 | 1MB |
| cluster | 2 | 16GB | 9% | 16ms | 324 | 2MB |

**4,000 Users:**

| | VCPUS | RAM | CPU AVG | RESPONSE TIME | DISK IO/SEC | NETWORK BYTES/SEC |
|---|---|---|---|---|---|---|
| non cluster | 2 | 16GB | 14% | 19ms | 605 | 2MB |
| cluster | 2 | 16GB | 16% | 21ms | 624 | 4MB |

**NOTE**: The hardware and operating system configurations used to capture test results listed in these tables were identical and the same as those used in previous tests described in this document. All tests used the NRPC mail workload.

From the results shown, we could observe the following:

- Using Domino clustering on vSphere has minimal impact on overall performance and response times.
- Domino clustering with vSphere can provide an excellent solution for achieving maximum levels of availability and gaining additional application failover capabilities.
- The cluster environment adds slightly to the CPU utilization percentage and disk access rates (IO/sec), but it has more impact on network performance. Therefore, proper planning of the use of network cards is essential.

It should be noted that, when planning resources for a Domino cluster environment, resources should be calculated on the total failover load and sized such that the system can easily handle that load. For example, if a cluster has two members with 2,000 active users on each member, each of the cluster members should be sized to handle a 4000-active user load. Sticking to these guidelines is equally important when running in a virtual environment. Looking at a host that has cluster members, it may look like you have extra resources that can be allocated to additional virtual machines but, in reality, these resources may be needed in a cluster failover.

## Virtualizing Domino iNotes (HTTP) Server Tests

The HTTP service in Domino is its web engine, which renders database documents to be displayed in a user's web browser. To exercise this Domino feature, we used the DWA85mail workload that simulates web-based mail users (iNotes users).

Testing was done using the 4000-user optimal building block configuration chosen earlier, with two vCPUs and 8GB of memory, followed by a second test where the memory was increased to 16GB. Table 5 shows the results of the test with two vCPUs and 8GB RAM and Table 6 shows results with two vCPUs and 16GB RAM.

Table 5. Virtual Machine Profile for Domino HTTP Test with Two vCPUs and 8GB RAM

| | CPU | | | MEMORY | | NETWORK | | DISK | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | (%) | | | (MB) | | PACKETS/SEC | | I/O SEC | | LATENCY(MS) | |
| VIRTUAL MACHINE | AVG | MAX | READY (AVG/ MAX) | AVG | MAX | RECD | TXMIT | AVG | MAX | AVG | MAX |
| OSDC-DominoMail-01 | 67.5 | 100 | 0/0.00 | 8108 | 8110.08 | 1261 | 679 | 683 | 1367.78 | 4 | 7.79 |

Table 6. Virtual Machine Profile for Domino HTTP Test with Two vCPUs and 16GB RAM

| | CPU | | | MEMORY | | NETWORK | | DISK | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | (%) | | | (MB) | | PACKETS/SEC | | I/O SEC | | LATENCY(MS) | |
| VM | AVG | MAX | READY (AVG/ MAX) | AVG | MAX | RECD | TXMIT | AVG | MAX | AVG | MAX |
| OSDC-DominoMail-01 | 66 | 100 | 0/0.09 | 15917 | 16220.16 | 1260 | 677 | 511 | 1300.81 | 2 | 4.15 |

Both sets of tests were able to attain a steady state for 4,000 concurrent, simulated users with reasonable CPU utilization percentage and user response times. This set of tests provided good insight into understanding the effect of memory and number of vCPUs selected for a guest system configuration. The graph in Figure 20 shows testing results with two CPUs and 8GB of memory. The second graph in Figure 21 shows test results for two vCPUs and 16GB of memory.

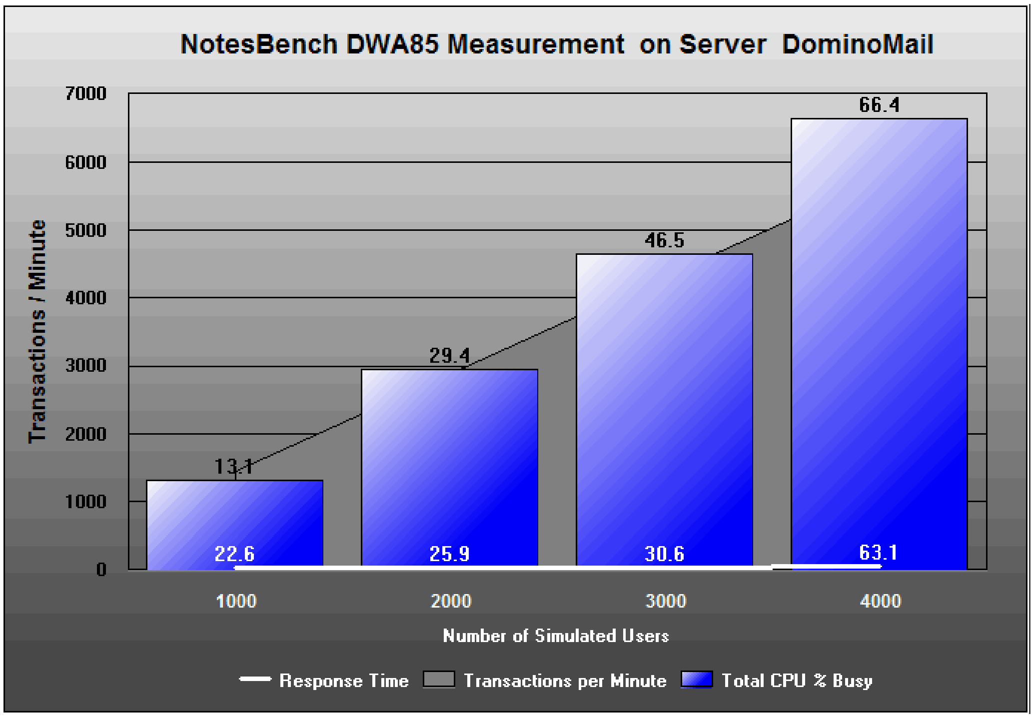


Figure 20. Average Response Times & CPU Utilization for Domino HTTP Test with 2 vCPUs and 8 GB Memory

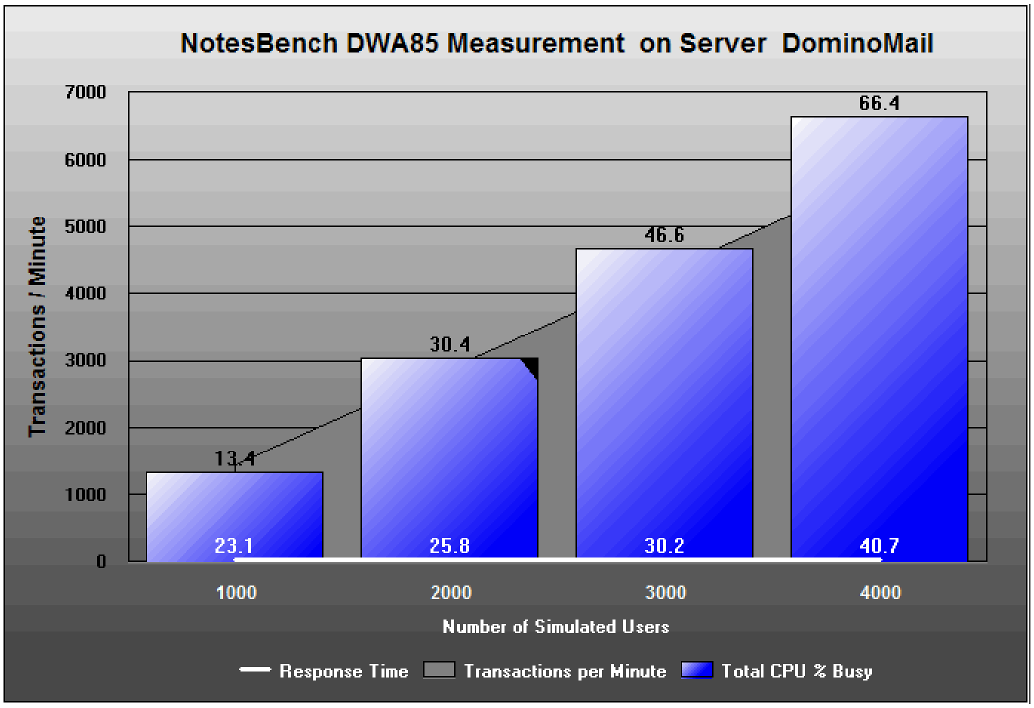


Figure 21. Average Response Times & CPU Utilization for Domino HTTP Test with 2 vCPUs and 16 GB Memory

Note the average response time[5] at the 4000-user mark in each graph, 63.1MS versus 40.7MS. The next set of Disk I/O graphs, in Figure 21 and Figure 22 respectively, help in understanding the difference in response times.

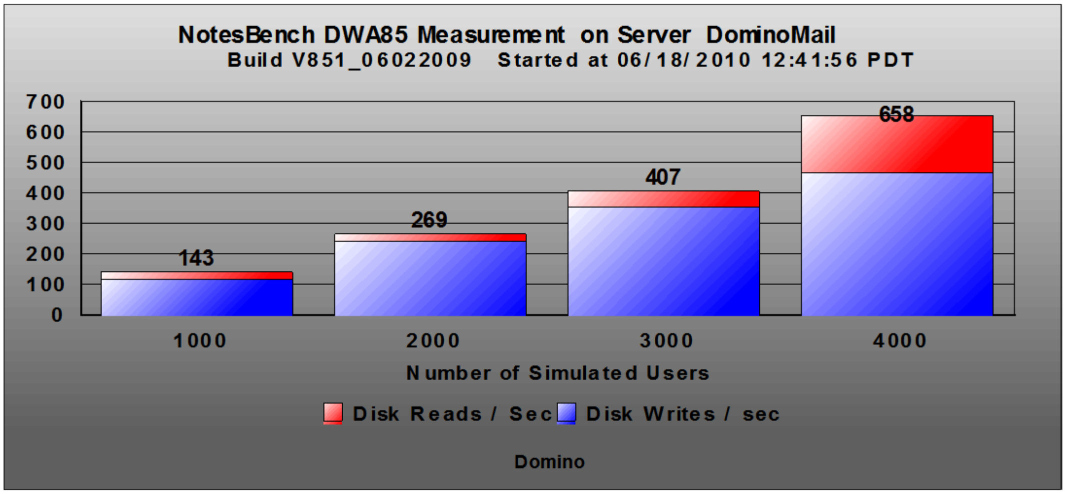


Figure 22. Disk Rates for Domino HTTP Test with Two vCPUs and 8 GB Memory

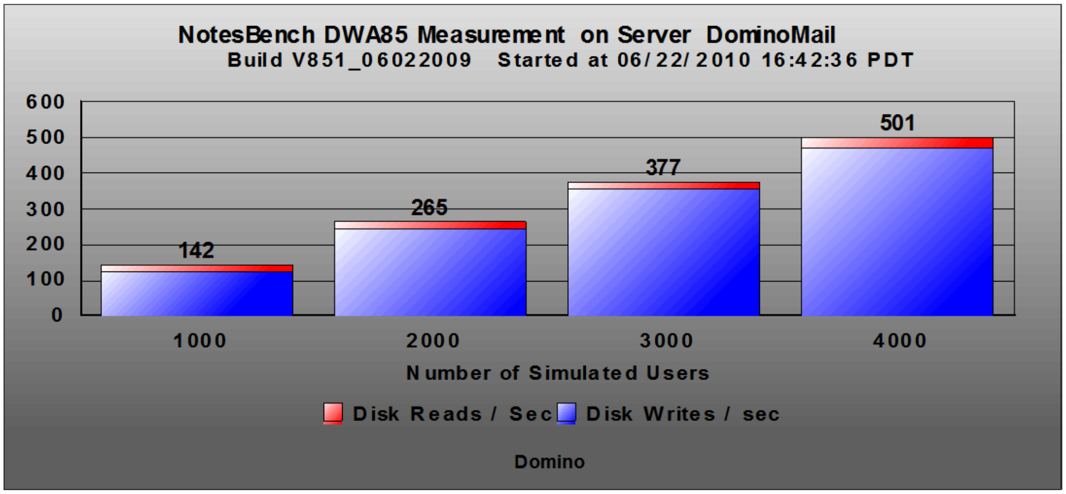


Figure 23. Disk Rates for Domino HTTP Test with Two vCPUs and 16 GB Memory

From the results shown, we can conclude that:

- Figure 22 and Figure 23 show that there is a substantial increase in the number of reads to disk at the 4000-user mark with 8GB selected (versus 16GB). The reason for this is that HTTP tasks utilize memory, so there is not enough extra memory with 8GB of RAM to cache all the databases. This is not necessarily a bad thing, since the response time with 8GB is well within expected service levels.
- The 66% CPU utilization in both the 8GB and 16GB tests could be considered to be on the high side so, depending upon the applications planned for this server, it may be good to increase the number of vCPUs for Domino HTTP environments.
- Looking at memory and CPU utilization provides a lesson in what needs to be considered when sizing Domino virtual machine system resources. As long as there is enough I/O bandwidth, response times should meet expectations.

[5] Response times (in milliseconds) are the white numbers displayed just above the x-axis along the bottom of the graph.

# Lessons Learned and Best Practices

The following table provides a summary of best practices, gleaned from the results of all testing, for virtualizing Lotus Domino using VMware vSphere.

Table 7. Summary of Best Practices for Lotus Domino Virtualization

| CONSIDERATION | BEST PRACTIVE DESCRIPTION |
|---|---|
| BIOS Settings | Enable Virtualization Technology to run the 64 bit guest OS. |
| | Enable Turbo Mode for balanced workload over unused cores. |
| | Enable VT-x for hardware based virtualization support. |
| | Enable High Performance mode instead of power saving mode to enhance performance. |
| CPU Considerations | Enable Intel VT-x technology that ensures hardware virtualization support instead of software Binary Translation (BT). |
| | Over-Commit is fine, but ensure that the host is not over-utilized. (Host itself required a minimum of two cores for service console, vmkernel, and other components.) |
| | Disable IRQ's like Floppy/CD/DVD ROM, USB Controllers and other video/audio programs. |
| Memory Considerations | Enable Extended Page Support (EPT) support for MMU overheads. |
| Network Considerations | Use VMXNET 3 Paravirtualized Network Adapters. |
| | Ensure that the virtual machine, Service Console, and vmkernel are on separate physical adapters to avoid any sort of contention. |
| | Ensured that all the virtual machines are on the same virtual switch as having virtual machines on separate switch will cause transmission of traffic over different wires, resulting in additional CPU and network overhead. |
| General Storage Considerations | Create an aggregate across disks to ensure more bandwidth and high IOPS, thereby reducing the latency. Separate Data stores for GOS and Domino Mail Databases. Recommendations on designing a scalable storage solution can be found in the vSphere Performance Best Practices whitepaper listed in the Additional Resources section below. |
| | 4 GB FC HBA. |
| | For active-active storage arrays, the policy for path selection is recommended to be Fixed (Preferred Path). Doing so allows for the best possible utilization of your bandwidth to the disk array, whereas, in case of active-passive storage arrays, the recommended path is MRU (Most Recently Used). |
| | All disks eager zeroed – An eager-zeroed thick disk has all space allocated and zeroed out at creation time. Longer creation time, but better performance and security. |
| | Ensure that best practices have been followed for proper disk partition alignment. For additional information, see: http://www.vmware.com/pdf/esx3_partition_align.pdf Also you may want to consult with your storage vendor for guidance on any array-specific alignment recommendations. |
| Others | Create a golden image of the virtual machine so it is easier to deploy; deploying virtual machines from this image will also save significant time. |

# Summary and Conclusions

VMware and IBM Lotus have collaborated for well over four years in virtualizing Domino on VMware's ESX and vSphere technologies. In this paper, we've walked through the process and methodologies we used to obtain performance baselines and provided performance scalability characteristics of different configurations that customers can use as guidance in sizing their own virtual environments. With the introduction of the latest technologies from VMware, IBM Lotus, and the processor companies AMD and Intel (Intel Nehalem with EPT and AMD with RVI), testing described in this paper has also shown CPU utilization and other performance metrics to be stellar for virtualized applications such as Domino, on par or within single percentage points of physical system performance.

While Domino NRPC scalability has been the primary target of previous VMware and IBM Lotus collaboration, customers have repeatedly asked us for best practices and what they should expect when virtualizing Domino NRPC Clusters with VMware vSphere and other VMware technologies. This paper should provide the reassurance customers need to feel comfortable virtualizing Domino Clusters and the best practices to follow when doing so. We also tested Domino compatibility with VMware vSphere advanced features and demonstrated that IBM Lotus Domino is fully compatible with tools such as VMware HA, vMotion, and VMware Fault Tolerance (FT).

The same conclusions about excellent performance of virtualized system proved true with workload testing of Lotus Domino iNotes and the results provided in this paper. This testing clearly showed favorable scaling of Domino HTTP in a VMware virtual infrastructure environment and also provided information helpful in looking at memory and CPU requirements of virtualized applications and determining what best meets your performance needs. Another important subject we addressed was the impact that large size mail databases would have on Domino virtualized to run in the VMware environment. As the results indicated, VMware vSphere and ESX continued to perform well, very close to performance levels provided by physical server systems. Using large mail files is not a concern when virtualizing Domino systems, provided the environment is sized properly.

In conclusion, the Domino server provides a robust platform for building and deploying applications to suit any environment, including many “custom” designed applications. Many VMware customers have asked us how their “custom” applications might perform with VMware virtualization. While, the test cases described in this paper don't address all varieties of “custom” Domino workloads, we expect you will likely be pleased with your results─ if you deploy your Lotus Domino “Custom” applications in accordance with best practices provided in this paper.

# Additional Resources

***VMware References***

- Customers can post questions in the VMTN Community specifically for messaging and collaboration products:

  http://communities.vmware.com/community/vmtn/general/emailapps

  This forum is monitored by VMware subject matter experts (SMEs) as well as other customers with similar interests.
- VMworld 2009 Session – Virtualizing IBM Lotus Domino and Sametime:

  http://www.vmworld.com/sessions/EA2649.html
- VMware Documentation:

  http://www.vmware.com/support/pubs/
- VMware Licensing:

  http://www.vmware.com/support/licensing/
- Application Performance Troubleshooting:

  http://www.vmware.com/resources/techresources/10066
- VMware vSphere 4.1 Performance Best Practices:

  http://www.vmware.com/resources/techresources/10161
- VMware Proven Best Practices (VIOPS):

  http://communities.vmware.com/community/viops
- VMware Solutions/Product/Partner Podcasts:

  http://www.vmware.com/technical-resources/podcasts/
- VMware Global Support KB:

  http://kb.vmware.com
- VMware Global Support Videos:

  http://blogs.vmware.com/kbtv/
- VMware TV:

  http://www.youtube.com/user/vmwaretv
- VMworld TV:

  http://www.youtube.com/user/VMworldTV
- VMware KB TV:

  http://www.youtube.com/user/VMwareKB

### IBM References

- IBM Lotus Domino website:
  http://www.ibm.com/software/lotus/products/domino/
- IBM Lotus Domino workloads:
  http://www.ibm.com/developerworks/lotus/library/notes8-workloads/
- General software support for IBM SWG products in a VMware environment:
  http://www-01.ibm.com/support/docview.wss?&uid=wws1e333ce0912f7b152852571f60074d175
- VMware product support information for IBM Lotus Domino-based server products:
  http://www-01.ibm.com/support/docview.wss?rs=203&uid=swg21106182
- Support statement for release 8.0.1 of Notes client, Domino server, and Domino Web Access:
  http://www-01.ibm.com/support/docview.wss?rs=899&uid=swg21252786
- Pricing policies for IBM Software on virtualization, including VMware:
  http://www01.ibm.com/software/lotus/passportadvantage/Counting_Software_licenses_using_specific_virtualization_technologies.html
- Pricing policies for IBM support:
  http://www-01.ibm.com/software/lotus/passportadvantage/pvu_licensing_for_customers.html

### NetApp References

- NetApp official website:
  www.netapp.com
- TR-3428: NetApp and VMware Virtual Infrastructure 3 Storage Best Practices:
  http://media.netapp.com/documents/tr-3428.pdf
- TR-3505: NetApp Deduplication for FAS and V-Series Deployment and Implementation Guide:
  http://media.netapp.com/documents/tr-3505.pdf
- TR-3705: NetApp and VMware View Solution Guide
  http://media.netapp.com/documents/tr-3705.pdf
- TR-3737: SnapManager 2.0 for Virtual Infrastructure Best Practices:
  http://media.netapp.com/documents/tr-3737.pdf
- TR-3749: NetApp and VMware vSphere Storage Best Practices:
  http://media.netapp.com/documents/tr-3749.pdf
- TR-3671: VMware vCenter Site Recovery Manager in a NetApp Environment:
  http://media.netapp.com/documents/tr-3671.pdf



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