Microsoft Exchange Server 2003 Deployment Considerations

for Small and Medium Businesses

A Dell™ PowerEdge™ 2800 server can provide an effective platform for Microsoft® Exchange Server 2003. A team of Dell engineers tested Exchange Server 2003 performance and scalability for varying numbers of users on a PowerEdge 2800 server platform, in scenarios typical of those in a small or medium-sized organization.

BY SUMAN KUMAR SINGH, ABHIJIT CHATTOPADHYAY, AND BHARATH VASUDEVAN

Related Categories:

Dell PowerEdge servers

E-mail technology

Microsoft Active Directory

Microsoft Exchange

Microsoft Windows Server 2003

Performance

Visit www.dell.com/powersolutions for the complete category index.

Today, small and medium businesses (SMBs) increasingly rely on messaging solutions. E-mail has become a mission-critical application. Whether in the office or on the road, communicating internally or with customers, enterprises depend heavily upon e-mail availability. Performance and availability of the messaging infrastructure is critical. Therefore, enterprise IT organizations must consider hardware sizing and mailbox design before the initial deployment or redeployment after hardware or software upgrades are performed.

This article examines how to select the appropriate platform for deploying Microsoft Exchange Server 2003 in SMB environments[1] and why this platform can be well suited for messaging solutions for the SMB segment. Dell engineers conducted testing in November 2005 with the intent of answering the following frequently asked questions:

- How many mailboxes can be hosted on one server?
- How much memory is needed for the Exchange server?
- If server processor utilization is very low, why is e-mail access slow?
- Is storage performance too slow?

Selecting appropriate hardware for Microsoft Exchange Server 2003

Dell provides several server platforms that can be used to deploy Exchange Server 2003 for SMBs up to approximately 900 employees. Different types of workload can be imposed by clients—consuming a combination of memory, processor, I/O, and network resources on the server. Typical client workloads for Microsoft Exchange tend to stress the memory and I/O subsystems, while exercising the processor and network subsystems to a lesser extent. Single-socket servers may not allow future scalability because they are limited in their expansion capabilities. Microsoft Exchange Server 2003 is a 32-bit application, which can only address 4 GB of virtual memory. The current version of Exchange Server 2003 does not support 64-bit versions of the Microsoft Windows® OS. Given the 4 GB memory limit and the lighter processor workload,

[1] For guidance on sizing large Microsoft Exchange installations, refer to the Dell Exchange Advisor Tool at www.dell.com/exchange.

Server	Dell PowerEdge 2800
CPU	Two Intel® Xeon® processors at 2.8 GHz with 1 MB level 2 cache
Memory	4 GB of double data rate 2 (DDR2) error-correcting code (ECC) of RAM
NIC	Two dual-port Intel 8254NXX Gigabit Ethernet* adapters
RAID controller	PowerEdge RAID Controller 4, Extended Dual Channel integrated
Internal disks	Ten 73 GB, 15,000 rpm SCSI drives
OS	Microsoft Windows Server™ 2003, Enterprise Edition, with Service Pack 1 (SP1)
Messaging application	Microsoft Exchange Server 2003, Enterprise Edition, with SP1

**This term does not connote an actual operating speed of 1 Gbps. For high-speed transmission, connection to a Gigabit Ethernet server and network infrastructure is required.*

Figure 1. Server configuration for Microsoft Exchange test environment

the processing power of a quad-socket server may not be required for hosting Exchange alone.

Ideally, a server should support 4 GB of memory and have scalable I/O options while providing the processor scalability and network bandwidth required to host Microsoft Exchange Server 2003. Dual-socket Dell PowerEdge servers, available in both tower and rack optimized form factors, can help meet such requirements. These general-purpose servers can provide the performance characteristics for Exchange as well as a large amount of internal storage. For future I/O scalability, using external storage is recommended. However, this article focuses on internal server storage only, which is typically used in small Exchange environments.

Configuring the Microsoft Exchange test environment

The Dell test team performed sizing and performance tests at the Dell Enterprise Solutions Engineering Labs. The purpose of the tests was to determine the processor, memory, network, and disk utilization while hosting between 600 to 1,050 Exchange mailboxes. Figure 1 describes the hardware components of the test environment. For more information about the test bed setup, the test tools, and the setup procedures, see the supplemental online section of this article at www.dell.com/powersolutions.

Analyzing the test results

Of the various factors affecting the performance of an Exchange server, the test team studied the effect of increasing the number of mailboxes on processor, memory, and network utilization as well as on disk response times. Using these results, the team analyzed the potential bottlenecks that can be encountered in an Exchange deployment. The tests were conducted for 600 to 1,050 simulated Exchange users.

Processor performance

To observe the effect of increasing the number of mailboxes on processor utilization, the test team installed LoadSim on each of the clients—which were Dell PowerEdge 750 servers—to simulate an Exchange workload. The test was run several times for different user counts. Figures 2 and 3 show the results obtained for processor measurement counters.

For 600 to 1,050 users, the average processor utilization ranged from 5 to 10 percent, while the maximum processor utilization ranged from 17 to 27 percent. Processor privileged time and user time followed the same trend as processor utilization. As shown in Figure 2, the average processor utilization reached only about 10 percent for 1,050 Messaging Application Programming Interface (MAPI) Messaging Benchmark 3 (MMB3) users. Typically, if processor utilization is consistently greater than 75 percent it may be considered a bottleneck. In this case, an average processor utilization of 10 percent shows that the PowerEdge 2800 had ample computing power available to handle 1,050 or more mailboxes. Thus, processing power should not be expected to become a bottleneck for this size of Exchange deployment. The spare computing power

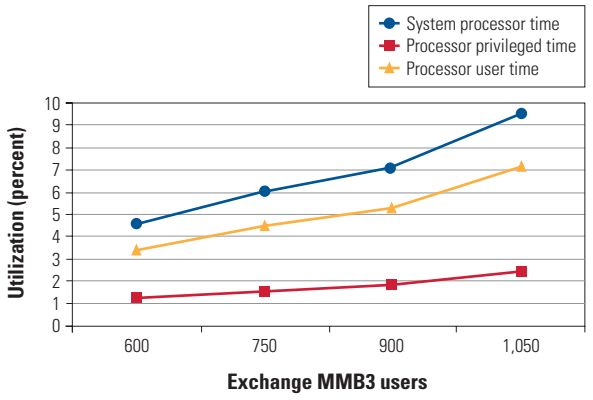

Figure 2. Average processor utilization for test environment

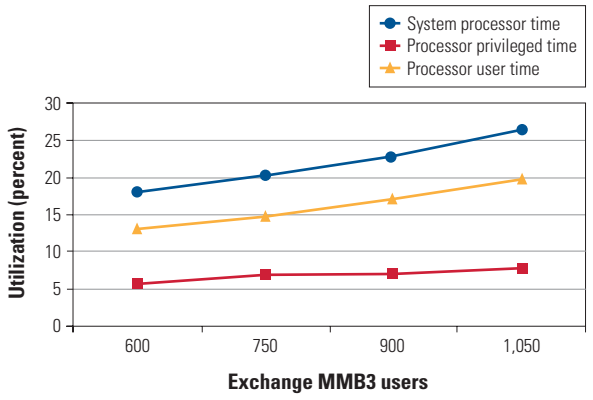

Figure 3. Maximum processor utilization for test environment

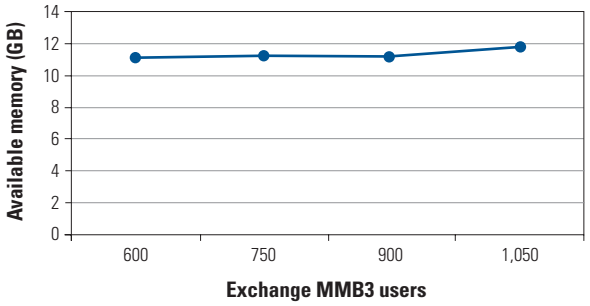

Figure 4. Memory utilization for test environment: Available memory

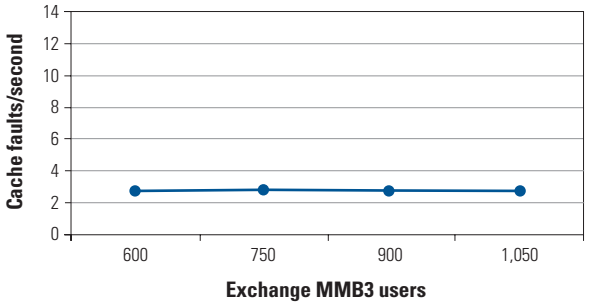

Figure 5. Memory utilization for test environment: Cache faults

available may be used by other applications such as antivirus, antispam, or backup software.

Memory performance

The tests used to study memory utilization were similar to those used for evaluating processor utilization. These tests measured available memory (see Figure 4) and cache faults (see Figure 5). The results show that available memory slightly declined and cache faults per second slightly increased as the number of mailboxes increased.

With the rapid increase in the speed of processors, memory, and system buses, the disk has become the slowest component of a server, making the disk seek process a very expensive operation in terms of performance. Exchange Server 2003 is an I/O- and memory-intensive application, so cache faults should be minimized as much as possible to limit data seek from disks and thus optimize performance. However, because Exchange is a 32-bit application, 4 GB of RAM is the maximum amount of memory that an Exchange server can efficiently use. Best practices recommend installing the maximum 4 GB of physical RAM.

Exchange memory optimization. By default, a 32-bit Windows OS allocates 2 GB of virtual address space for kernel mode (OS) and 2 GB for user mode (applications).[2] But 2 GB of virtual address space may not be enough for memory-intensive applications like Exchange. To solve this problem and make more memory space available for user mode, administrators can set the `/3GB` switch in the boot.ini file. The `/3GB` switch makes 3 GB (3,072 MB) of memory space available for user mode, and therefore more memory is available for Exchange services. Administrators should also set the `/USERVA=3030` parameter in the boot.ini file to allocate 3,030 MB of the 3 GB to Exchange and keep 42 MB available for dynamically allocating memory back to kernel mode if needed.

However, all client activity generated by a MAPI client causes updates to the Exchange Jet database, which produces random I/O activity. Installing the recommended 4 GB of RAM does not necessarily solve all memory issues nor guarantee that most I/Os will be cached. Using this switch may only help ensure that the Exchange components do not lack memory resources. For optimized performance, administrators should allocate an appropriate value of virtual memory to different Exchange components as well as to the OS. If the server is not a dedicated Exchange server and it hosts other applications such as antivirus, antispam, or backup software, additional RAM may be required.

As shown in Figures 4 and 5, increasing the number of users does not drastically increase the cache faults per second; neither does it significantly reduce the available memory. These results indicate that the Exchange components have sufficient memory and the deployment may be performing satisfactorily.

Network performance

Figure 6 shows network performance (kilobytes of data sent and received) for a varying number of users. These test results show that as the number of mailboxes increased, network traffic increased. Both the incoming and outgoing network traffic increased proportionally with the increasing number of mailboxes.

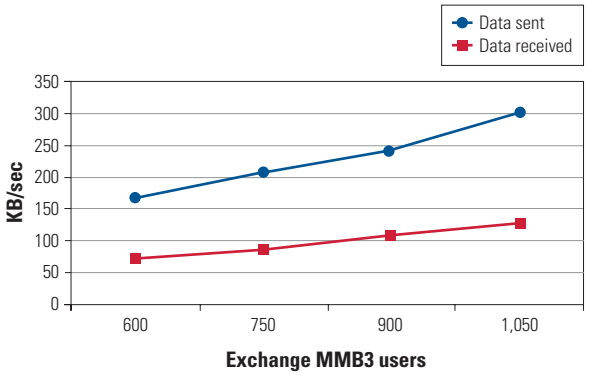

Figure 6. Network utilization for Microsoft Exchange test environment

[2] Microsoft Windows 2000 Advanced Server and 32-bit Windows Server 2003, Enterprise Edition, can support more than 4 GB of RAM by using Physical Address Extension (PAE). For more information about PAE and large memory support in Windows 2000 and Windows Server 2003, visit www.microsoft.com/whdc/system/platform/server/PAE/PAEdrv.mspx and support.microsoft.com/?kbid=283037.

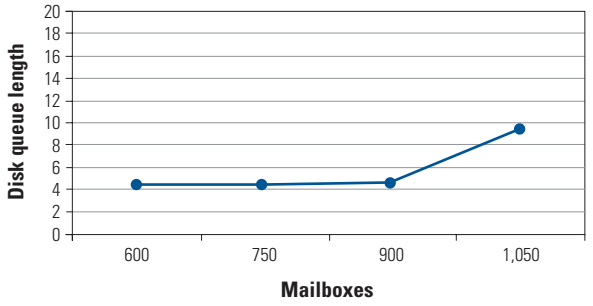

Figure 7. Storage performance for test environment: Disk queue length

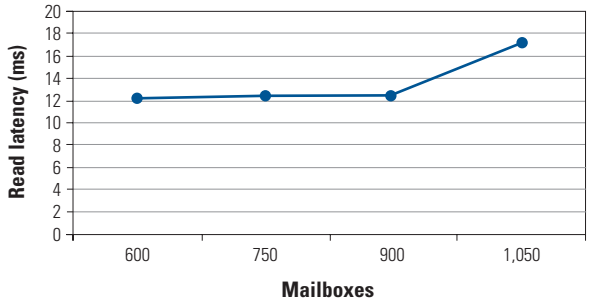

Figure 8. Storage performance for test environment: Read latency

From the magnitude of network traffic shown in Figure 6, the full-duplex 100 Mbps LAN connection appeared to perform sufficiently. However, these results do not include backup traffic or any other application traffic. If enterprise IT organizations plan to use the network for backups and restores or if other applications generate substantial network traffic, they should consider using a Gigabit Ethernet network.

Exchange uses the Microsoft Active Directory® directory service. Every Exchange mailbox must be an Active Directory user. Exchange servers and messaging clients access Active Directory in various situations, such as when logging on to the network and connecting to a mailbox or accessing server-based address lists. These activities create heavy network traffic between the servers, so there should be sufficient network bandwidth available between Exchange and the Active Directory server.

Storage performance

To observe the effect of increasing the number of mailboxes on storage performance, the test team used the Microsoft JetStress tool to simulate an Exchange I/O workload. The test was run several times for different mailbox counts with 50 MB mailbox size and 0.67 I/Os per second (IOPS) per user. Figures 7, 8, and 9 show the test results.

As shown in Figures 7, 8, and 9, disk queue length, read latency, and disk transfers per second rose steeply for more than 900 mailboxes. The increase in read latency was proportional to the increase in disk queue length. For 900 mailboxes and fewer, the read latency and disk queue length remained nearly constant.

Exchange is an I/O-intensive application and all client activity causes updates to the Exchange database, which produces I/O operations to disk. The disk subsystem should be able to meet these demands, and thus administrators should size the disks for performance and not just for capacity of the mailboxes.

For high performance, best practices recommend that the average read or write latency not exceed 20 milliseconds (ms) and the average disk queue length per spindle be less than 2. A disk queue length of less than 2 per spindle suggests that, while one I/O operation is being processed by the disk, another I/O is waiting in the queue. A value greater than 2 per spindle suggests that I/O requests are coming in at a higher rate than they can be processed by the disk. This may ultimately result in high disk latencies and therefore is not recommended. Best practices also recommend using fast drives rather than high-capacity drives. Fast drives can process more I/Os per second with low latency. If considering using SCSI drives, enterprise IT organizations should deploy 15,000 rpm Ulra320 drives. To further improve disk performance, IT administrators may also consider using the Microsoft DiskPar utility to verify that the disk tracks are sector aligned.

As shown in Figures 7, 8, and 9, the disk queue length, read latency, and disk transfers per second remained relatively consistent up to 900 mailboxes, even though the workload was increasing. However, these values increased sharply for 1,050 mailboxes. For example, the disk queue length for 1,050 mailboxes increased to 10. The database RAID group in this test environment comprised six spindles, so the disk queue length per spindle for 1,050 mailboxes was 1.67 (10 divided by 6), which is close to the maximum recommended limit of 2. Also, the read latency increased to 17 ms, which is approaching the limit of 20 ms.

These test results show that the disk performance parameters approached their maximum recommended limit for 1,050 mailboxes. To allow for spikes and periods of heavy load, best practices recommend no more than 900 mailboxes on the Exchange server.

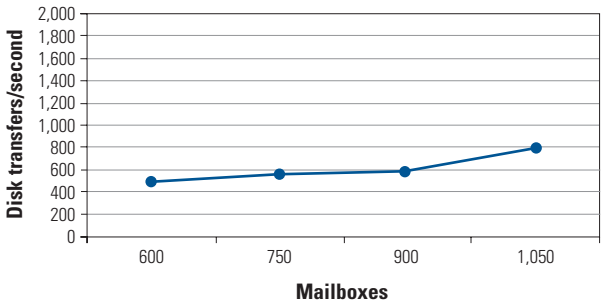

Figure 9. Storage performance for test environment: Disk transfers per second

To accommodate future growth, administrators should consider further reducing the number of hosted mailboxes. For example, to accommodate 25 percent future growth up to a maximum of 900 mailboxes, no more than 675 mailboxes should initially reside on the Exchange server.[3]

Available disk bandwidth can become a bottleneck as the number of mailboxes increases. To overcome this limitation and host large numbers of mailboxes, administrators may consider moving storage to external storage devices including direct attach storage (DAS) or storage area networks (SANs).

Considering other Microsoft Exchange configuration options

IT administrators deploying Microsoft Exchange in SMB environments may consider configurations such as high-availability (HA) clusters and front-end Microsoft Exchange servers.

High-availability clustering

The basic goal of HA clustering is to make sure that the physical server hosting an application is not a single point of failure by providing the ability for that application to be restarted on one of multiple servers in a cluster. If the server running the application fails, another designated server takes over the responsibility of running that application. Dell HA cluster implementations employ Microsoft Cluster Service (MSCS), and are designed and tested to help make sure that no single point of failure exists.[4] HA clustering using MSCS requires shared storage, because every node in the cluster needs access to the Exchange data. Therefore, to implement this type of configuration, administrators should consider moving storage to external devices such as DAS or SANs.

Front-end servers

Microsoft Exchange Server 2003 supports an Exchange architecture consisting of front-end and back-end servers. The front-end server accepts requests from clients and proxies them to the appropriate back-end server for processing. This architecture is usually recommended when the Exchange environment has multiple back-end servers. However, administrators may deploy a front-end server in an environment with a single back-end server because this configuration can provide the following benefits:

- **Single namespace:** Allows single namespace for all users; same URL for Microsoft Outlook® Web Access, Post Office Protocol 3 (POP3), and Internet Message Access Protocol 4 (IMAP4) clients; and consistent server name (does not change even if mailbox is moved or new servers are added)
- **Offload processing:** Uses front-end server to manage all encryption and decryption processing
- **Strong security:** Uses front-end server as a single point of access either on or behind a firewall; provides an additional layer of security for mailboxes; and does not require Remote Procedure Call (RPC) ports to be opened from perimeter network to internal network
- **Scalability:** Enables the number of front-end or back-end servers to be increased or decreased without disrupting users

Deploying Microsoft Exchange in an SMB environment

SMBs can range in size from 50 to 2,000 employees. A dual-socket Dell PowerEdge 2800 server can provide a suitable platform for SMBs to deploy messaging solutions. The results presented in this article show that, if used as a dedicated Microsoft Exchange server, the internal drives of a PowerEdge 2800 server can host up to 900 mailboxes.

Although the PowerEdge 2800 has enough processing power, memory, and network capacity to support much more than 900 mailboxes, the available disk bandwidth using the on-board hard drives can become a limiting scaling factor. To overcome this limitation and to host more than 900 mailboxes on a single server, enterprise IT administrators can move storage to external SCSI devices or host the Exchange mailboxes on a Dell/EMC Fibre Channel SAN. The best practices discussed in this article also can help administrators identify the bottlenecks in their deployments and take the appropriate actions to correct them.

Suman Kumar Singh is a systems engineer in the High-Availability Systems Group at Dell. He specializes in messaging systems architecture and sizing. His other interests include SANs, virtualization, and security. Suman has published several papers and presented at virtualization-related industry conferences.

Abhijit Chattopadhyay is a senior engineer in the High-Availability Solutions Engineering Group at the Dell Bangalore Development Center. Abhijit has a degree in Electronics and Communication Engineering from K.L.E. Society's College of Engineering and Technology in Karnataka, India. His current interests include high-availability clustering, SANs, and Microsoft Exchange.

Bharath Vasudevan currently manages the High-Availability Cluster Group at Dell. He has previously designed server hardware and served as a lead for multiple cluster releases. His current interests include application performance characterization and storage technologies. He has a master's degree in Electrical and Computer Engineering from Carnegie Mellon University.

[3] These storage recommendations are for Exchange users each generating 0.67 IOPS. Performance may vary depending on the nature of the Exchange users.

[4] For more information about deploying Microsoft Exchange Server 2003 on a Dell HA cluster, see "Microsoft Exchange Server 2003 Scale-Out Performance on a Dell PowerEdge High-Availability Cluster" by Arrian Mehis, Ananda Sankaran, and Scott Stanford in *Dell Power Solutions,* February 2005; www.dell.com/downloads/global/power/ps1q05-20040216-Stanford.pdf.