


Running Citrix Presentation Server 4.0 on X3 Architecture Servers

Discover the benefits of 64-bit computing with CPS 4.0 x64

Choose the best IBM System x server for the task

Read about best practices for Presentation Server



David Watts
Richard Bell
Raphael Buehner



International Technical Support Organization

**Running Citrix Presentation Server 4.0 on X3
Architecture Servers**

July 2006

Note: Before using this information and the product it supports, read the information in “Notices” on page v.

First Edition (July 2006)

This edition applies to Citrix Presentation Server 4.0 running on the IBM System x3950, System x3850 and System x3800.

© Copyright International Business Machines Corporation 2006. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	v
Trademarks	vi
Preface	vii
The team that wrote this Redpaper	vii
Become a published author	ix
Comments welcome	ix
Chapter 1. Selecting the right 64-bit server	1
1.1 What is 64-bit processing?	2
1.2 Comparing x86-64 to Itanium 2	3
1.3 Dual-core processors	4
1.4 The benefits of a 64-bit platform	5
1.4.1 Kernel memory	5
1.4.2 Scaling up	6
1.4.3 Easy migration path	6
1.5 X3 Architecture	6
1.5.1 Memory features	8
1.6 The X3 Architecture servers	10
1.7 The benefits of X3 in a Presentation Server x64 environment	11
1.8 Selecting a Presentation Server x64 platform	12
1.8.1 Storage in a Presentation Server x64 environment	12
1.8.2 Processor scalability in a Presentation Server x64 environment	15
1.8.3 How many processors?	15
1.8.4 NUMA in a Presentation Server x64 environment	16
1.8.5 Choosing a server	17
1.9 Virtualization using the System x3800	18
1.9.1 Presentation Server consolidation barriers	19
1.9.2 Virtualization Technology and Presentation Server x64	20
1.9.3 Characteristics of the System x3800	20
1.9.4 Virtualization platforms	22
Chapter 2. Introducing Citrix Presentation Server 4.0	27
2.1 What is new in Citrix Presentation Server 4.0	28
2.1.1 CPU utilization management	29
2.1.2 Virtual memory optimization	30
2.1.3 Application Isolation Environment	31
2.1.4 Improvements on printing	33
2.1.5 Virtual IP addressing	34
2.1.6 New licensing design	35
2.1.7 Support for Windows Server 2003 x64 Edition	35
2.2 Citrix Presentation Server 4.0 x64 Edition	36
2.3 Why upgrade to Presentation Server 4.0?	36
2.4 Benefits of Windows Server 2003 x64 Edition	37
2.5 Why upgrade to Presentation Server 4.0 x64 Edition?	38
Chapter 3. Moving to a 64-bit Citrix Environment	41
3.1 Upgrade planning	42
3.2 Supported upgrade paths	45

3.3 Migration strategies	46
3.3.1 In-place farm migration	46
3.3.2 Two farm aggregation	49
3.3.3 Multi-farm load sharing	52
3.4 Operating in a mixed 32/64-bit environment	53
3.5 Lab upgrade	54
Chapter 4. Best practices	59
4.1 Server layout	60
4.1.1 Hardware layout	60
4.2 Sizing for scalability, high availability, and fault tolerance	61
4.2.1 Scalability	62
4.2.2 High availability and fault tolerance	63
4.3 Configuring Load Manager	67
4.3.1 Publishing applications	68
4.4 Restricting user access	71
4.4.1 NTFS and registry security	72
4.4.2 User profiles	72
4.4.3 User Policies	77
4.5 Printing in a Presentation Server 4.0 environment	80
4.5.1 Session printing	80
4.5.2 Client printing with auto-created printers	81
4.5.3 Using the Universal Printer Driver	82
4.6 Performance monitoring with key performance counters	84
4.6.1 Citrix Resource Manager	84
4.6.2 Windows Performance Monitor	86
4.7 Performance tuning	88
4.7.1 Increase performance for Presentation Server 4.0	88
4.7.2 Increase performance for Windows Server 2003	90
4.7.3 Using third-party products	91
4.8 Adding and deploying new Presentation Server systems	92
4.8.1 Unattended installation	92
4.8.2 Cloning	92
Appendix A. CPS User Policies	95
Abbreviations and acronyms	99
Related publications	101
IBM Redbooks	101
Other publications	101
Online resources	102
How to get IBM Redbooks	103
Help from IBM	103

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law. INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.


COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

BladeCenter®
Chipkill™
DB2®
eServer™
IBM®
iSeries™
Lotus Notes®

Lotus®
Netfinity®
Notes®
OS/2®
Predictive Failure Analysis®
Redbooks™
Redbooks (logo) ™

ServeRAID™
SmartSuite®
System x™
Virtualization Engine™
xSeries®
zSeries®

The following terms are trademarks of other companies:

Java, JRE, Solaris, StarOffice, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Active Desktop, Active Directory, Internet Explorer, Microsoft, Outlook, Windows Server, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Itanium, Xeon, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Preface

Citrix Presentation Server 4.0 is the latest version of the popular server application for virtualizing large numbers of client systems on the one large server. This means that the user workstations can be lower function and that overall management can be centralized and streamlined. New to Citrix Presentation Server 4.0 is support for 64-bit EM64T-based servers and 64-bit operating systems, such as Windows® Server 2003 x64. 64-bit support means a larger number of users can be supported on the one server, as much as three times over the previous 32-bit version.

The X3 Architecture servers are a family of high performance IBM® System x™ servers, with four processor sockets and support for up to 64 GB of memory and six high-speed PCI-X slots. The X3 family comprises of the flagship x3950 server with its XpandOnDemand capabilities to attach up to eight systems together to form a single 32-way complex, the x3850, with its rack-dense 3U configuration, and the x3800, which supports up to 12 internal SAS hot-swap drives for up to 3.6 TB of internal storage. All three servers include the IBM X3 chipset and memory bandwidth of up to 21 GBps.

This IBM Redpaper describes what's new in Version 4.0 of Presentation Server and explains how the 64-bit combination of hardware, operating system, and Citrix application supports a greater number of users than the 32-bit equivalent. We also explain the migration steps that are available to move you to 64-bit processing, complete with migration strategies that eliminate user downtime. Finally, we provide details on best practices on using Presentation Server.

This Redpaper is aimed at the reader who is already familiar with earlier versions of Presentation Server, but are new to 64-bit computing and want to understand what makes a suitable platform to run Presentation Server 4.0 x64.

The team that wrote this Redpaper

This Redpaper was produced by a team of specialists from around the world working at the International Technical Support Organization, Poughkeepsie Center.

David Watts is a Consulting IT Specialist at the IBM ITSO Center in Raleigh. He manages residencies and produces Redbooks™ on hardware and software topics related to IBM System x servers and associated client platforms. He has authored over 40 Redbooks and Redpapers. He holds a Bachelor of Engineering degree from the University of Queensland (Australia) and has worked for IBM for over 15 years. He is an IBM Certified IT Specialist.

Richard Bell is an IT Specialist with IBM Australia. He has over 13 years of experience in network operating system and infrastructure sizing, design, and implementation. His areas of expertise include Microsoft® server products, Citrix Access Suite products, VMware server products, and IBM BladeCenter®, System x, and Storage products. Richard has many industry certifications, including Citrix Certified Integration Architect, Microsoft Certified Systems Engineer, IBM Professional Server Engineer, and Cisco Certified Network Associate.

Raphael Buehner is Team Leader of the System Integration Team for InLIS GmbH, an IBM Business Partner in Germany. He has more than seven years experience in IBM products, including System x, xSeries®, Netfinity® servers, BladeCenter, and Storage. Raphael studied Information Science and Technology at the University of Horb in Horb, Germany, and graduated as Diplom Engineer in August 1995. His areas of expertise include Microsoft Windows Network design and implementation, Citrix Access Infrastructure integration, VMware Virtual Infrastructure implementation, and xSeries and System x hardware. He has many industry certifications, including IBM eServer™ Certified xSeries Specialist, Citrix Certified Administrator, Microsoft Certified Systems Engineer, and Virtual Certified Professional. In the past, he held several Workshops in the area of VMware, Citrix, and Microsoft.



Figure 1 The team (l-r): Richard, Raphael, David

Thanks to the following people for their contributions to this project:

From the ITSO:

Cheryl Gera
Linda Robinson
Margaret Ticknor
Erica Wazewski

From IBM:

Jay Bretzmann
Marcus Kornegay
Heather Richardson
Rick Rudd
Andy Sylivant
Chuck Weber
Bob Zuber

From Citrix:

Barry Flanagan
Chris Mumford

Mark Snell
Ashley Willis

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners, and customers.

Your efforts will help increase product acceptance and client satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our papers to be as helpful as possible. Send us your comments about this Redpaper or other Redbooks in one of the following ways:

- ▶ Use the online **Contact us** review redbook form found at:

ibm.com/redbooks

- ▶ Send your comments in an e-mail to:

redbook@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400



Selecting the right 64-bit server

IBM X3 Architecture servers represent the most advanced and powerful Intel® based servers in the industry today. The IBM System x3800, x3850, and x3950 are the three IBM servers included in the X3 lineup. Together, they represent the ideal 64-bit computing platform where high processor, memory, and I/O throughput are required.

To take advantage of X3 Architecture and the 64-bit computing platform, Citrix has released a 64-bit version of their popular application presentation platform, Citrix Presentation Server 4.0. In this chapter, we guide you through the process of answering why you would choose a 64-bit Presentation Server environment, and which System x server is the best choice for your needs.

This chapter has the following topics:

- ▶ 1.1, “What is 64-bit processing?” on page 2
- ▶ 1.2, “Comparing x86-64 to Itanium 2” on page 3
- ▶ 1.3, “Dual-core processors” on page 4
- ▶ 1.4, “The benefits of a 64-bit platform” on page 5
- ▶ 1.5, “X3 Architecture” on page 6
- ▶ 1.6, “The X3 Architecture servers” on page 10
- ▶ 1.7, “The benefits of X3 in a Presentation Server x64 environment” on page 11
- ▶ 1.8, “Selecting a Presentation Server x64 platform” on page 12
- ▶ 1.9, “Virtualization using the System x3800” on page 18

1.1 What is 64-bit processing?

When discussing what constitutes a 64-bit platform, there are several areas that need to be addressed. They are:

- ▶ Operand size (integer and floating point)
- ▶ Register size
- ▶ Internal and external bus widths
- ▶ Physical and virtual address sizes

At times, some, or all of these factors, are used to distinguish processor sizes. For example, the original i8086 and i8088 processors are identical apart from the fact that the i8086 had a 16-bit data bus and the i8088 had an 8-bit data bus. The i8086 was considered a 16-bit processor, while the i8088 was considered to be an 8-bit processor.

The Intel Xeon® (*Prestonia* or *Gallatin*) supports 64-bit integer and floating point operands, 128-bit internal data paths, 64-bit external data paths, and 36-bit physical addressing. Yet it is considered a 32-bit processor because it has 32-bit General Purpose Registers (GPRs) and a 32-bit virtual address space.

The Intel Itanium® 2 supports 64-bit GP registers, 82-bit floating-point registers, 64-bit virtual addresses and 128-bit external data paths. So, why is it a 64-bit architecture? The operands size alone would not be considered 64-bit. In addition, it cannot be the data paths because none of those are 64-bit. Physical addresses are 50 bits, not 64. The one thing of significance here is that the virtual address space is 64 bits.

Definition: In terms of processor architecture, a 64-bit processor is a processor that is able to address 64 bits of virtual address space. A 64-bit processor can store data in 64-bit format and calculate arithmetic operations on 64-bit operands. In addition, a 64-bit processor has GPRs and Arithmetic Logic Units (ALUs) that are 64 bits wide.

Advanced Micro Devices (AMD) introduced an open-source extended instruction set for the IA-32 platform and named it AMD64. The instruction set was designed to enable an x86 compatible processor to overcome the memory limits of 32-bit virtual addressing (4 GB), yet provide backwards compatibility for 32-bit operating systems and applications. AMD initially named the extensions x86-IA64, but quickly renamed it AMD64. Intel followed with its own extended instruction set and named it Enhanced Memory 64 Technology (EM64T).

Both processors share many features and are code-compatible. However, the Intel EM64T does not provide support for the 3DNow instruction set from AMD. AMD64 does not support Hyper-Threading Technology from Intel. For the purpose of this Redpaper, the two processors will be deemed to be one and the same.

Note: In this redpaper, we will collectively refer to AMD64 and EM64T as *x86-64*.

The AMD64/EM64T(x86-64) processors have 64-bit GP registers, and they support 64-bit integer and floating-point operations. Data buses are all 64 bits wide or greater. In addition, 64-bit virtual addresses are supported. It can be stated that the x86-64 is a 64-bit processor.

Note: The address space of the x86-64 processors does support 64-bit addresses, but only the least significant 40 or 48 bits are used. The top 24 or 16 bits are required to be zero. This is a distinction that can be raised in the discussion of what constitutes a 64-bit processor, but it is only a theoretical discussion, because 48-bit addresses (2^{48}) is 256 TB and Windows Server® 2003 R2 Standard x64 Edition *only* supports up to 32 GB of physical RAM and 16 TB of Virtual Memory. Even Windows Server 2003 R2 Datacenter x64 Edition *only* supports up to 1 TB of physical RAM.

Having said that, instructions involving 64-bit addresses are given full support in the x86-64. The processor can operate in either 32-bit or 64-bit modes, and there are sufficient resources within the processor to support either mode. Both modes are supported for the sake of backwards compatibility, not because of any inherent restriction on 64-bit operation. The x86-64 is a 64-bit processor, but it is a processor that retains its compatibility with earlier generations of hardware.

In the discussion of what constitutes a 64-bit architecture, the use of *extensions* is one of the prime areas of confusion. AMD and Intel uses the term *extensions* to describe how the 32-bit instruction set is *extended* to provide greater flexibility in 64-bit support while maintaining backwards compatibility for existing software. It is the implementation of that instruction set that makes the difference in trying to give any sensible definition to what constitutes a true 64-bit architecture. The x86-64 core was enhanced where necessary to make sure it has full 64-bit support. So, in effect it is more accurate to describe them as 64-bit architectures that support backwards compatibility modes.

1.2 Comparing x86-64 to Itanium 2

The Itanium 2 processor from Intel is the company's highest performing computing platform. The processor has been designed for and marketed at high-end workstations and compute intensive applications.

A key difference between x86-64 and Itanium 2 is that the Itanium's Explicitly Parallel Instruction Computing (EPIC, as opposed to RISC or CISC) instruction set does not support both 32-bit and 64-bit modes. In fact, it will not run applications or operating systems compiled to run on an IA32 system. The Itanium 2 processor does not include nor support the x86-64 extensions.

x86-64 processors can run in three modes:

- ▶ 32-bit mode
- ▶ Compatibility mode
- ▶ Full 64-bit mode, "*Long Mode*" or "*IA-32e*"

In the near future, the most widely used mode will most likely be 32-bit mode. In this mode, you can install your 32-bit operating system, with its 32-bit drivers, and 32-bit applications. You will not be able to realize the advantages of the new features, such as the addressing above 4 GB or the additional registers. However, 32-bit applications will run somewhat faster due to improvements inherent in the new chips design.

x86-64 compatibility mode gives you the ability to run a 64-bit operating system and device drivers while still running 32-bit applications. Legacy applications will still only be able to address memory up to 4 GB, but the memory that they see will be dedicated to their own process. Previously, the system kernel and applications had to share physical memory. Compatibility mode has dropped support for virtual 8086 mode, as such real-mode existing

applications are not supported. 16-bit Protected-Mode applications are, however, still supported by the processor.

Note: Microsoft Windows Server 2003 x64 does not provide support for DOS, 16-bit, OS/2®, or POSIX applications.

The full functionality of x86-64 is implemented in full 64-bit mode. Applications that run in full 64-bit mode will get access to the entire 40 bits of memory address space, the new GPRs, and also the extended GPRs. This mode of operation not only requires a 64-bit operating system and 64-bit drivers, but also an application that has been compiled to take full advantage of the 64-bit architecture. Citrix Presentation Server for Windows Server 2003 x64 Edition (Presentation Server x64) is one such application.

The Intel Itanium 2 processor runs in one mode: 64-bit. The EPIC and x86-64 instructions sets are code incompatible. However, the Itanium 2 processor does provide support for 32-bit software by way of a hardware-based IA32 emulator built into the chip and a software based IA32 translator called the *IA32 Execution Layer* (IA32-EL). The performance of IA32 applications running on Itanium 2 is approximately half of what you would expect from today's AMD Opteron and Intel Xeon processors. This may be sufficient for some workloads, for example, you may want to run a small 32-bit Web server along side your large SQL Server 2005 Itanium Edition. However, for most 32-bit production workloads, running them on Itanium will result in unsatisfactory performance.

Note: While Microsoft and many other Independent Software Vendors have invested much time and money to produce operating systems and applications optimized to run on the Itanium 2 platform, Citrix Presentation Server x64 will not run on Itanium 2. In fact, Terminal Services (Application Services) is not an option of the Itanium release of Windows Server 2003. It is simply not available.

1.3 Dual-core processors

A dual-core processor is an integrated circuit in which two processor cores have been integrated to provide greater performance by way of more efficient simultaneous processing of multiple tasks. Having a dual-core processor is similar to having two discreet processors in the same server, but since the two processors are integrated on the same silicon die, the connection between them is much faster. Dual-core processors generally have slower clock speeds than their single-core counterparts. This is to keep power consumption down. However, the additional performance capacity means an overall increase in performance.

Some would speculate that a dual-core processor would be twice as powerful as a single-core processor in a Terminal Server environment; however, this is not correct. A dual-core processor does share some circuitry between the cores, such as the memory bus. As such, instead of a dual-core CPU being twice as powerful as two single cores, in reality, a dual-core processor is likely to be 120-150% faster than a single core processor with the same characteristics (front-side bus speed, cache size, and so on).

The block diagram in Figure 1-1 on page 5 shows the logical differences between a single-core Xeon MP *Cranford* processor and a dual-core Xeon 7020/7040 *Paxville* processor.

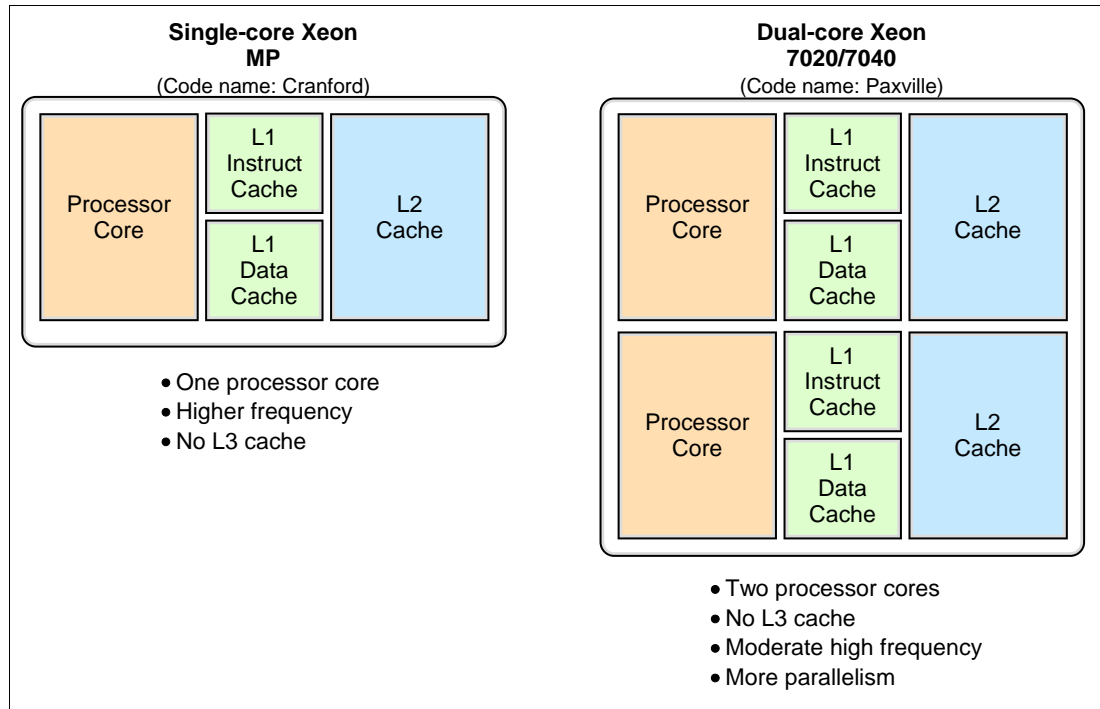


Figure 1-1 Features of single core and dual core processors

As the current line of single-core processors approach the limitations of speed and complexity, dual-core processors will become more mainstream rather than high end.

1.4 The benefits of a 64-bit platform

The most obvious benefit of running Presentation Server x64 is the large increase in the number of user sessions that can be supported on a single 64-bit server. However, there are other less obvious benefits to running Presentation Server x64.

1.4.1 Kernel memory

In a Citrix Presentation Server environment, the benefits of a 64-bit platform are typically quantified in terms of user scalability. Any advancement in server hardware usually results in an increase in the number of users you can get onto a Presentation Server with a given workload.

As advancement in hardware design has outpaced the capacity of 32-bit Windows operating systems, limitations in kernel memory have limited end-user scalability in servers with four or more processors. Windows Server 2003 Standard Edition is a 32-bit OS that has a memory limitation of 4 GB (2^{32}). This 4 GB is divided into 2 GB for user-mode processes and 2 GB for kernel-mode processes. The /3gb switch can be used in the boot.ini file to increase the amount of memory available to user-mode processes, but it is only beneficial in situations where the Presentation Server server is not kernel memory bound and the applications have been programmed to take advantage of the Address Windowing Extensions (AWE) application programming interfaces included in Windows Server 2003.

Enabling the /3gb switch effectively halves the amount of memory allocated to kernel mode processes. Enabling the /3gb switch in Presentation Servers that are kernel memory bound is not recommended.

With the introduction of the X3 Architecture servers with their x86-64 processors, Windows Server 2003 x64 Edition, and Citrix Presentation Server x64, previous scalability limits no longer exist. Windows Server 2003 x64 can address up to 16 TB of virtual address space. There is 8 TB for kernel-mode processes and 8 TB for user-mode processes. Presentation Server x64 can fully utilize a 64-bit quad processor server before kernel-memory becomes a bottleneck.

1.4.2 Scaling up

Getting more concurrent users onto each Presentation Server results in cost savings;

- ▶ Fewer administrators required to maintain fewer physical servers
- ▶ Reduced overhead required to manage patches, service packs, and hotfixes
- ▶ Less power, cooling, and space requirements in computer rooms
- ▶ Lower server-based licensing costs

1.4.3 Easy migration path

Many of the benefits of 64-bit computing are fully realized in an optimized 64-bit environment. The benefits are centered mainly around the integer operation size and the memory addressing components of an operating system or application.

However, Windows Server 2003 x64 is able to run 32-bit and 64-bit applications side-by-side by isolating the 32-bit applications inside an execution layer called Windows on Windows 64 (WoW64). WoW64 provides an isolation environment for 32-bit applications while providing interoperability between 64-bit applications and other 32-bit applications through COM and Remote Procedure Calls (RPCs).

While WoW64 supports 32-bit applications at the operating system level, x86-64 processors provide support for 32-bit apps by their compatibility mode function. This is good in a Presentation Server environment where many desktop applications have not yet been ported to a 64-bit architecture. As always though, not all 32-bit applications can make use of memory above 2 GB. Many applications may not have the necessary APIs to take advantage of memory above 2 GB. Running 32-bit applications in the WoW64 environment can result in a slight degradation of 32-bit performance, but this is more than offset by the increases in I/O performance that you would typically see in the IBM System x servers with X3 Architecture.

1.5 X3 Architecture

In 2005, IBM introduced the IBM X3 Architecture, the culmination of a three year, \$100 million development effort, to bring mainframe-inspired capabilities and sophisticated high-end technology to the company's next-generation 64-bit Intel Xeon processor-based System x servers. The new X3 Architecture provides up to 38% higher 4-way performance than the previous generation of Intel Xeon processor-based systems. It enables businesses to run 32-bit and 64-bit applications simultaneously and to process massive amounts of data more rapidly.

The third generation of Enterprise X-Architecture design, X3 Architecture, is optimized for superior server consolidation and enterprise application software, and uses the IBM high-end

Virtualization Engine™ technology to enhance virtualization capabilities. The X3 Architecture offers investment protection for IBM customers with its family of pay-as-you-grow Intel-based servers by supporting both 32-bit and 64-bit applications on the same platform so clients can migrate to 64-bit as needed.

XA-64e, the third-generation Enterprise X-Architecture chipset (codename *Hurricane*), is the heart of the X3 Architecture, providing an integrated processor and memory controller that significantly reduces memory latency, improving response times and overall system performance. The XA-64e chipset was specifically designed for System x servers by a cross-platform architectural team.

With support for up to 32 Intel Xeon MP processors and over 20 Gbps of memory bandwidth per 64 GB of RAM (up to a maximum of 512 GB), the System x servers that are based on the X3 Architecture offer maximum performance and broad scale-up capabilities. X3 offers break-through performance and reliability by leveraging off technologies commonly seen in big-iron p, i, and zSeries® systems. IBM has incorporated eDRAM technology, a Distinct Filtering Directory, and a Snoop Filter, right into the chipset.

Figure 1-2 shows the main components and pathways of the X3 Architecture chipset.

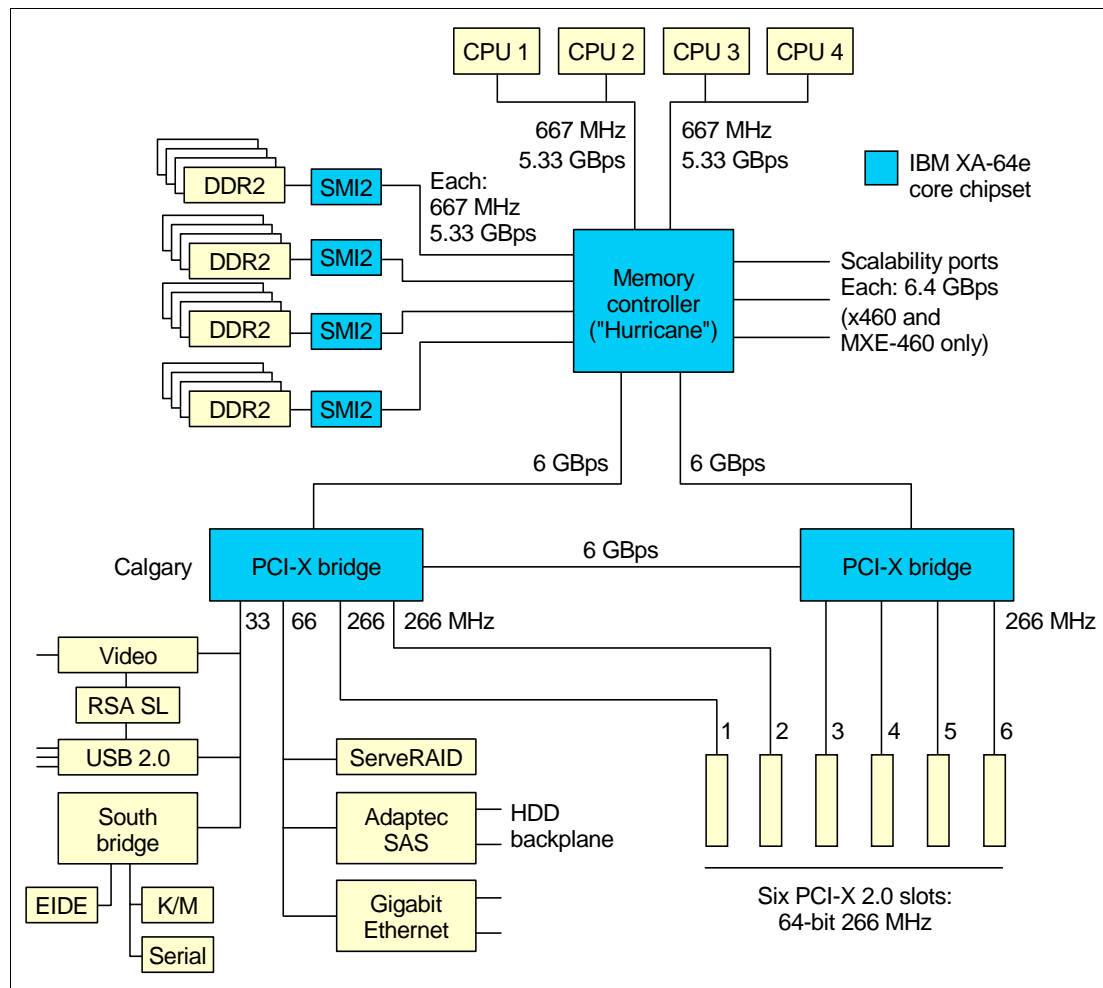


Figure 1-2 X3 Architecture block diagram

Some of the features of the X3 Architecture include:

- ▶ Active Memory with Memory ProteXion, and ChipKill.
- ▶ XceL4v Cache
- ▶ Support for dual-core Xeon 7020 and 7040 (*Paxville*) processors
- ▶ Supports PC2-3200 ECC DDR2 DIMMs

1.5.1 Memory features

Memory in the X3 Architecture servers is PC2-3200 ECC DDR2 DIMMs. Standard memory is as follows:

- ▶ IBM System x3800 (x260): 1 GB or 2 GB of RAM standard (using 512 MB DIMMs)
- ▶ IBM System x3850 (x366): 2 GB standard (using 512 MB DIMMs)
- ▶ IBM System x3950 (x460): 2 GB standard (using 512 MB DIMMs)

Memory is implemented in the servers using memory cards. Each server supports up to four memory cards, and each card has four DIMM sockets. Memory cards are standard as follows:

- ▶ IBM System x3800: One card standard
- ▶ IBM System x3850: One card standard
- ▶ IBM System x3950: Two cards standard
- ▶ IBM System x3950 E (MXE-460): Two cards standard

Using 4 GB DIMMs in every socket, a total of 16 DIMMs, the servers can hold 64 GB of RAM.

The memory is two-way interleaved, meaning that memory DIMMs are installed in pairs. There are four ports to memory, with each supporting up to 5.33 GBps data transfers.

The DIMMs operate at 333 MHz instead of 400 MHz, as per the PC2-3200 spec, so that throughput is 2.67 GBps, or 333 MHz x 8 bytes. At 2.67 GBps and two-way interleaving, the throughput matches that of the front-side bus at 5.3 GBps so that bus transfers remain in sync.

There are a number of advanced features implemented in the X3 Architecture memory subsystem, collectively known as *Active Memory*:

- ▶ Memory ProteXion

The Memory ProteXion feature (also known as *redundant bit steering*) provides the equivalent of a hot-spare drive in a RAID array. It is based in the memory controller, and it enables the server to sense when a chip on a DIMM has failed and to route the data around the failed chip.

Normally, 128 bits out of every 144 are used for data and the remaining 16 bits are used for ECC functions. However, X3 Architecture needs only 12 bits to perform the same ECC functions, thus leaving four bits free. These four bits are equivalent to an x4 memory chip on the DIMM that Memory ProteXion uses. In the event that a chip failure on the DIMM is detected by memory scrubbing, the memory controller can reroute data around that failed chip through these spare bits.

It can do this automatically without issuing a Predictive Failure Analysis® (PFA) or light path diagnostics alert to the administrator, although an event is logged to the service processor log. After the second DIMM failure, PFA and light path diagnostics alerts would occur on that DIMM as normal.

Note: In BIOS, there is a Memory Array setting in Advanced Settings. When you set it to **High Performance Memory Array**, BIOS reconfigures the server for maximum performance at the expense of some fault tolerance features, including Memory ProteXion. For a production environment, we recommend that you not select this setting, thereby keeping your system protected from memory failures with Memory ProteXion.

► **Memory scrubbing**

Memory scrubbing is an automatic daily test of all the system memory that detects and reports memory errors that might be developing before they cause a server outage.

Memory scrubbing and Memory ProteXion work in conjunction with each other and do not require memory mirroring to be enabled to work properly.

When a bit error is detected, memory scrubbing determines if the error is recoverable or not. If it is recoverable, Memory ProteXion is enabled and the data that was stored in the damaged locations is rewritten to a new location. The error is then reported so that preventative maintenance can be performed. As long as there are enough good locations to allow the proper operation of the server, no further action is taken other than recording the error in the error logs.

If the error is not recoverable, then memory scrubbing sends an error message to the light path diagnostics, which then turns on the proper lights and LEDs to guide you to the damaged DIMM. If memory mirroring is enabled, then the mirrored copy of the data from the damaged DIMM is used until the system is powered down and the DIMM replaced.

► **Memory mirroring**

Memory mirroring is roughly equivalent to RAID 1 in disk arrays, in that usable memory is halved and a second copy of data is written to the other half. If eight GB is installed, then the operating system sees four GB once memory mirroring is enabled. It is disabled in the BIOS by default. Because all mirroring activities are handled by the hardware, memory mirroring is operating system independent.

When memory mirroring is enabled, certain restrictions exist with respect to placement and size of memory DIMMs and the placement and removal of memory cards.

► **Chipkill™ memory**

Chipkill is integrated into the XA-64e chipset, so it does not require special Chipkill DIMMs and is transparent to the operating system. When combining Chipkill with Memory ProteXion and Active Memory, X3 Architecture provides very high reliability in the memory subsystem.

When a memory chip failure occurs, Memory ProteXion transparently handles the rerouting of data around the failed component, as described above. However, if a further failure occurs, the Chipkill component in the memory controller reroutes data. The memory controller provides memory protection similar in concept to disk array striping with parity, writing the memory bits across multiple memory chips on the DIMM. The controller is able to reconstruct the missing bit from the failed chip and continue working as usual. One of these additional failures can be handled for each memory port for a total of four Chipkill recoveries.

► **Hot-Swap memory**

The X3 Architecture servers support the replacing of failed DIMMs while the server is still running. This hot-swap support works in conjunction with memory mirroring.

- ▶ Hot-Add memory

The server also supports adding additional memory while the server is running. Adding memory requires operating system support.

Hot-Swap and Hot-Add are features that are mutually exclusive. Hot-add requires that memory mirroring be disabled and hot-swap requires that memory mirroring be enabled.

In addition, to maintain the highest levels of system availability, if a memory error is detected during POST or memory configuration, the server can automatically disable the failing memory bank and continue operating with reduced memory capacity. You can manually re-enable the memory bank after the problem is corrected by using the Setup menu in the BIOS.

Memory mirroring, Chipkill, and Memory ProteXion provide multiple levels of redundancy to the memory subsystem. Combining Chipkill with Memory ProteXion allows up to two memory chip failures for each memory port on the system, for a total of eight failures sustained. For example:

1. The first failure detected by the Chipkill algorithm on each port does not generate a light path diagnostics error because Memory ProteXion recovers from the problem automatically.
2. Each memory port could then sustain a second chip failure without shutting down.
3. Provided that memory mirroring is enabled, the third chip failure on that port would send the alert and take the DIMM offline, but keep the system running out of the redundant memory bank.

1.6 The X3 Architecture servers

The three servers based on the X3 Architecture are the IBM System x3800, x3850, and x3950. They have a common set of technical specifications and characteristics, but the key differences are:

The System x3950 (formerly the xSeries 460) is the flagship server with the following characteristics:

- ▶ Each chassis occupies 3U of rack space and supports four CPUs and 64 GB of RAM.
- ▶ Six hot-swap drive bays and six 266 MHz PCI-X 2.0 hot-swap slots.
- ▶ Targeted at 8-way and above configurations where effective scale-up options are essential.
- ▶ Up to eight systems can be connected together to form one single 32-way complex with up to 512 GB RAM.

The System x3850 (formerly the xSeries 366) is a high performance 4-way server with the following characteristics:

- ▶ The same mechanical design as the x3950, offering up to 64 GB of RAM and up to four Xeon MP processors.
- ▶ Targeted at 2-way and 4-way high performance commercial computing, such as database, e-mail, and e-commerce, and Terminal Server applications.

The System x3800 (formerly the xSeries 260) is also a high performance 4-way server. The x3800 has the following characteristics:

- ▶ Same central electronics as the x3850 and x3950.

- ▶ Larger 7U chassis to house up to 12 hot-swap disk drives and a full-height internal tape drive.
- ▶ Targeted at two-way and four-way high performance commercial computing applications where more internal disk storage is required.

Key features

The x3800, x3850, and x3950 have a number of common features:

- ▶ IBM eServer X3 Architecture, featuring the XA-64e third-generation chipset.
- ▶ Common system boards: the CPU/memory board, the I/O board, and the PCI-X board.
- ▶ Up to four Intel Xeon processors. These Processors support 64-bit addressing with the Intel Extended Memory 64 Technology (EM64T) architecture.
- ▶ Support for the Intel Xeon 7020/7040 dual-core processors (some models have dual-core processors standard).
- ▶ Up to 64 GB of RAM, using high performance PC2-3200 ECC DDR2 DIMMs.
- ▶ Active Memory with Memory ProteXion, memory mirroring, memory hot-swap and hot-add, and ChipKill.
- ▶ Six full-length 64-bit 266 MHz PCI-X 2.0 Active PCI slots.
- ▶ Integrated Adaptec AIC-9410 serial-attached SCSI (SAS) controller.
- ▶ Support for internal RAID arrays using an optional ServeRAID™-8i adapter. ServeRAID-6M is also supported for external SCSI storage with the EXP400 enclosure.
- ▶ Integrated dual-port Broadcom 5704 PCI-X Gigabit Ethernet.
- ▶ Integrated Baseboard Management Controller. Remote Supervisor Adapter II SlimLine adapter standard (x3950) or optional (x3850 and x3800).
- ▶ Support for the IBM Integrated xSeries Adapter for iSeries™ (IXA) for a direct high speed link to an iSeries server (x3950 and x3850 only).
- ▶ Hot-swap fans and power supply.
- ▶ Light path diagnostics to identify any failed components.
- ▶ Three-year warranty on-site, nine hours per day, five days per week, with a next business day response.

For more details about the X3 Architecture servers, see the redbook *Planning and Installing the IBM @server X3 Architecture Servers*, SG24-6797.

1.7 The benefits of X3 in a Presentation Server x64 environment

Advancements in technology by made Intel, Microsoft, and Citrix enable an x86-64 server to host literally hundreds of user sessions per server at excellent speed and responsiveness. An organization is able to consolidate its Terminal Servers from many 1U rack servers to a couple or a few large 64-bit servers. Therefore, businesses can reap the benefits of the savings associated with reduced power consumption, reduced administrative overhead, reduced cooling, and reduced server-based licensing costs.

Of course, the downside of the upside is that server consolidation amplifies the risk of server hardware failure affecting potentially hundreds of users. IBM has recognized this risk and has responded by incorporating many mainframe-inspired features into the X3 Chipset.

Memory ProteXion, Memory Scrubbing, Chipkill, hot-swappable and redundant components such as power supplies, fans, and hard disk drives enable a level of availability that is equal to, or better than, a large farm of smaller, lower specification servers.

It is also important to consider that, with fewer servers doing the same or more work, organizations should implement stringent change and configuration management systems to protect production servers from inadvertent (or otherwise) corruption, and clumsy patch implementation.

1.8 Selecting a Presentation Server x64 platform

So far, you have been shown some of the key features and benefits of the EM64T processors running on X3 Architecture. You have seen an outline of the three X3 Architecture servers - the x3800, the x3850, and the flagship x3950. Now, let us determine which is the right X3 server on which to deploy Citrix Presentation Server x64.

Note: This is not a server sizing discussion. That is, “How many CPUs and how much RAM do I need to run this many users on this type of load?”. There have been numerous articles published on this subject.

For sizing information, review the *IBM BladeCenter and Citrix MetaFrame Presentation Server sizing* white paper, available from:

<http://www.pc.ibm.com/support?page=MIGR-52686>

While the above paper does not directly address Presentation Server x64 running on X3 servers (it is based on MetaFrame XP and HS20/HS40 Blades), there is some excellent information presented that will help you perform your own server sizing computations.

The three servers in the X3 Architecture lineup have identical processor, memory, and bus architectures. Where they differ is in the areas of storage and processor scalability. The focus of this section will be around identifying where storage, and the number of processors, affect a Presentation Server environment, then choosing a configuration that is a best-fit for the job.

Both Microsoft and Citrix have conducted their own performance tests of Windows Terminal Server x64 and Citrix Presentation Server x64.

The Microsoft paper, *Terminal Services Scaling and Performance on x64-Based Versions of Windows Server 2003* can be found at

<http://www.microsoft.com/windowsserver2003/technologies/terminalservices>

The Citrix paper entitled *Citrix Presentation Server for Microsoft Windows Server 2003 x64 Edition - Performance and Scaling Capabilities* can be found at:

<http://support.citrix.com/article/CTX107341>

1.8.1 Storage in a Presentation Server x64 environment

First, the storage requirements in a Presentation Server environment encompass:

- ▶ Operating system files:
 - Page file
 - Application executables and dynamic linked libraries (DLLs)
- ▶ User profile storage

- ▶ Application data storage:
 - Temporary data
 - User session data
 - Shared data
- ▶ Users' home drives
- ▶ Farm data store
- ▶ Print spooler files

Careful consideration should be given to storage when designing any Presentation Server farm. When a Presentation Server x64 server is being implemented, any mistakes or oversights made during the analysis of requirements will be amplified by the number of users operating on one Presentation Server x64 server, which would potentially be hundreds.

- ▶ Operating system files:

There is not much to consider here. Unless the server has been configured to boot from a SAN, the operating system will be installed on a local hard disk drive.

- Page file

For performance, the Page file should be configured to reside on a different spindle, or set of spindles, than the operating system files. However, this only becomes an issue if the Presentation Server is running short on memory and the operating system is swapping user session data (usually kernel mode data) between RAM and the page file. The server should be configured with enough RAM (within the realms of budget realities) to avoid this situation.

- Application executables and DLLs

As with the operating system files, unless the server has been configured to boot from a SAN over a Fibre Channel adapter, all application executables and DLLs will be installed locally.

Application Isolation Environment uses a larger amount of local storage than would otherwise be normally used. Take this into account if you plan to use AIE in your environment.

- ▶ User profile storage:

When designing a Presentation Server farm, there are many factors to consider in how to handle user profiles. A full discussion of the implementation of the different types of user profiles is out of the scope of this Redpaper. For the purposes of the storage discussion, we will take a quick look at each profile type. There are several different types of user profiles that can be deployed. There are others, but the main types of profiles are:

- Local profiles

User profile data is kept locally on the Presentation server, usually under `C:\Documents and Settings\%username%`.

- Roaming profiles

User profile data is located somewhere on a file server share.

- Mandatory profiles

User profile data can be located locally or in the network. The important thing here is that the profile is locked-down so that the user cannot make permanent changes to his or her operating environment. One copy of the profile is shared among many users.

– Hybrid profiles

A Hybrid profile uses a mandatory profile with certain areas of the profile being saved for concurrency between sessions.

As an example, a user is running Microsoft Word as a published application and changes their Toolbar settings. Using a local or roaming profile, these changes are saved to NORMAL.DOT and the changes are constant between sessions. If the user has been configured with a mandatory profile, the toolbar customization is discarded as soon as the user logs off. When the user logs back on, the mandated settings are reloaded. A hybrid profile is a mandatory profile with redirected user data folders, such as the %appdata% folder. The hybrid profile model offers great flexibility and security in a Presentation Server environment because IT administrators have the ability to lock down and secure a users' environment without restricting the functionality of the applications that are being accessed.

The Flex Profile Kit by Jeroen van de Kamp puts a new twist on the hybrid profile. See: <http://www.brianmadden.com/people/person.asp?id=13>

The difference with the Flex Profile Kit is that it allows the Administrator to customize which areas of a mandatory profile are to be saved. This is particularly good for applications that do not store user settings under the %appdata% (C:\Documents and Settings\%username%\Application Data) area.

Depending on the type of user profile deployed, profile storage space may grow to many gigabytes per user. Of course, this depends on the applications in use and how they store their data.

▶ Application data storage

As mentioned previously in the discussion about user profiles, the %appdata% folder can be redirected to a network share, or the users' home drive. Applications such as Internet Explorer® make use of a temporary storage area to make retrieval of cached data faster. Applications produce data to be shared with others in the organization.

▶ Users' home drives

Users' home drives may contain Microsoft Outlook® 2003 offline storage files (.ost), redirected application data files, a redirected My Documents folder, and, of course, a user's personal files.

▶ Print spooler files

The print spooler files are normally generated and filed under C:\WINDOWS\system32\spool. It is common practice to relocate the spool folder to a separate spindle from the operating system and page file if the server will be processing a lot of print jobs. Consideration should be given to this in the design of your Presentation Server systems.

▶ Farm data store

The data store is a central database that holds records for information about the farm that remains relatively static. Information such as the configuration of the servers in the farm, details of published applications, policies, printer, printer drivers, and isolation environments are all kept in the data store. Information that changes regularly, such as server load metrics and performance data, is not kept in the farm data store.

The data store may be housed on one of the servers in the farm. This is acceptable for a small implementation. For medium and large implementations, we recommend configuring a dedicated server running a Relational Database Management System (RDBMS) such as IBM DB2®, Microsoft SQL Server, or Oracle.

User profiles, application data, and users' home drives can require large amounts of storage space. Added to that, in an environment where there can be up to 600 users logged on per server, these folders also generate a significant amount of disk activity. Best practice for Presentation Server 4 dictates that a Presentation Server should not be used as a file server. Indeed, a Presentation Server should be dedicated to the tasks involved directly with delivering applications to user desktops. Mixing server roles, particularly in a Presentation Server x64 environment where you have many concurrent users, could result in server slowness due to disk I/O bottlenecks and cause user sessions to hang.

User and application data should be stored in a SAN, NAS, or dedicated file server.

1.8.2 Processor scalability in a Presentation Server x64 environment

In previous versions of Presentation Server, the most common deployments were on dual processor servers. Implementing Presentation Server on quad-processor servers only yielded a minor increase in performance. A minor increase in performance did not justify the extra cost for the additional processors.

In a MetaFrame XP Presentation Server environment, it was fairly common to see a situation where a dual-processor server would support a maximum of 30-35 concurrent users before the machine started to slow down and cause user session hangs. Adding more processors to the server did not have much effect because the server would bottleneck in the memory subsystem before the processors were even marginally utilized.

Adding memory only alleviated part of the problem, because 32-bit Windows can only address up to 4 GB of RAM. Even on a quad processor server with 4 GB of RAM, the operating system would run out of kernel memory or page table entries before the hardware was fully utilized.

With IBM X3 Architecture servers, Windows Server 2003 x64, and Citrix Presentation Server x64, previous scalability limits around memory have been removed. It is now feasible to consider Presentation Server deployments on four-core and eight-core systems. The following section looks at the questions surrounding processor density in a Presentation Server x64 environment.

1.8.3 How many processors?

Both the x3800 and the x3850 support up to four dual-core processors. The x3950 can scale up to 32 dual-core processors using the scalability ports.

Note: Performance and Scalability testing of Presentation Server x64 has not been included as part of this Redpaper. Refer to the Citrix white paper *Citrix Presentation Server for Microsoft Windows Server 2003 x64 Edition - Performance and Scaling Capabilities* at <http://support.citrix.com/article/CTX107341> and the Microsoft white paper *Terminal Services Scaling and Performance on x64-Based Versions of Windows Server 2003* at <http://www.microsoft.com/windowsserver2003/technologies/terminalservices/default.aspx> for detailed analyses of Presentation Server x64 and Windows Terminal Server x64 performance and scalability.

While it would certainly be possible to deploy Presentation Server x64 on an x460 with 64 cores, would it be feasible? Performance metrics obtained by Citrix indicate that, in a 64-bit environment, the performance increase from a 1-way server to a 2-way server is a 116% gain in user density. Going from a 2-way server to a 4-way server results in a 60% gain in user density. Table 1-1 shows these results.

Table 1-1 Multi-processor scalability in a Presentation Server x64 environment

Processors	Number of users	User density gain
1-way server	74 +/- 1	Baseline
2-way server	160 +/- 1	+116%
4-way server	255 +/- 1	+60%

The above figures were obtained using a “light user” load on AMD Opteron processors. Performance results on x86-64 processors have not been published by either Microsoft or Citrix. At this point in time, scaling Opteron processors above four CPUs yields a performance decrease due to the L3 cache coherency architecture inherent in the processors themselves. Processor configurations above four processors in an IBM X3 Architecture server would produce better scalability due to the low latency design of processor L3 cache sub-system of the X3 chipset.

Note: These figures are referenced for the purpose of demonstrating how increasing processor density, even in a Presentation Server x64 environment, does not yield a linear growth pattern. Your results may vary.

IBM Global Services and its partners have extensive knowledge and expertise in the area of Presentation Server sizing and deployment. Please see your IBM representative to organize a detailed consultation.

Given the above, it is likely that in an environment where 32-bit operating system kernel memory limitations are not a limiting factor, and processor and memory architecture are not limiting in any way, Presentation Server x64 (itself) can in fact scale beyond four processors. However, due to the lack of 64-bit applications that run on top of Presentation Server x64, linear scalability from four to eight processors and beyond is highly unlikely due to the way in which 32-bit applications run in Compatibility Mode on an x86-64 architecture.

1.8.4 NUMA in a Presentation Server x64 environment

Non-Uniform Memory Access (NUMA) architecture is a way of building very large multi-processor systems without jeopardizing hardware scalability. The name NUMA is not completely correct, since not only memory can be accessed in a non-uniform manner, but also I/O resources, such as PCI-X devices.

NUMA effectively means that every processor or group of processors has a certain amount of memory local to it. Multiple processors or multiple groups of processors are then connected together using special bus systems (for example, HyperTransport or the scalability ports of a System x3950) to provide processor data coherency. The essence of the NUMA architecture is the existence of multiple memory subsystems, as opposed to a single one on a SMP system.

The so-called “local” or “near” memory has the very same characteristics as the memory subsystem in a SMP system. But by limiting the number of processors that directly access that memory, performance is improved, because of the much shorter queue of requests.

Since each group of processors has its local memory, memory on another group of processors would be considered to be remote to the local processor. This remote memory can be accessed but at a longer latency than local memory. All requests between local and remote memory flow over the inter-processor connection (HyperTransport or scalability ports).

Operating system design needs to take into account NUMA architecture. Operating systems need the appropriate APIs to read an area of firmware called the *Static Resource Allocation Table* (SRAT). All editions of Windows Server 2003 x64 and Linux® kernels above Version 2.6 support NUMA. Applications such as Microsoft SQL Server 2005 have NUMA support built-in.

NUMA awareness is not a characteristic of Presentation Server x64. Citrix relies on the capabilities built into Terminal Services for this function. However, as the function of Terminal Services is merely to provide a multi-user, multi-session environment for Microsoft Windows, it can be said that the responsibility for NUMA awareness is passed on to the application layer. Microsoft provides a set of APIs so that developers can take advantage of the NUMA support built into the operating system. Since most applications hosted in any Terminal Services environment are 32-bit, non-scalable, single-threaded, non-NUMA aware applications, it can be surmised that Presentation Server x64 will not scale well in a multi-node configuration.

1.8.5 Choosing a server

So far, it has been proposed that in large scale Presentation Server x64 environment, lots of local storage should not be required on the Presentation Server x64 server. With that being said, a relatively fast storage subsystem should be configured to accommodate an increase in disk reads due to larger numbers of users accessing applications local to the server.

It has also been stated that NUMA awareness is not a characteristic of Presentation Server x64 or Terminal Services. Due to the nature of the applications being delivered to end-users, Presentation Server x64 will not scale well on a multi-node server configured with more than four processors.

Taking these parameters into consideration, the System x3850 is an ideal choice for your Presentation Server x64 implementation.

Characteristics of the System x3850

Delivering an industry-leading, 64-bit framework for application presentation, the IBM System x3850 is built on the power of the third generation of IBM Enterprise X-Architecture technology. X3 drives the x366 to deliver the performance, availability, and manageability required for large scale Presentation Server x64 deployments.

Four-socket performance and 64-bit memory addressability provide an optimized platform for the application-serving tier. At the crossroads of a major industry transition to mainstream 64-bit applications, X3 delivers a formidable combination of 64-bit performance, availability, and investment protection.

With its extensive chipset development experience and industry-leading performance and availability breakthroughs, IBM is uniquely positioned to propose a robust and powerful server, offering innovation that delivers real business and IT results.

Figure 1-3 shows the IBM System x3850.



Figure 1-3 The IBM System x3850

The following are the key features of the System x3850:

- ▶ 4-way capable server in a rack-dense 3U form factor.
- ▶ Models with single-core or dual-core processors, and single-core models are upgradeable.
- ▶ XA-64e third-generation chipset.
- ▶ Models with one Intel Xeon MP processor, up to 3.66 GHz and 1 MB L2 cache, which can be upgraded to 4-way. Processors support 64-bit addressing with the Intel Extended Memory 64 Technology (EM64T) architecture.
- ▶ 2 GB memory standard expandable to 64 GB (using 4 GB DIMM slots), using high performance PC2-3200 ECC DDR2 DIMMs.
- ▶ Active Memory with Memory ProteXion, memory mirroring, memory hot-swap and hot-add, and ChipKill.
- ▶ Six full-length 64-bit 266 MHz PCI-X 2.0 Active PCI slots.
- ▶ Integrated Adaptec AIC-9410 serial-attached SCSI (SAS) controller. Support for internal RAID arrays using an optional ServeRAID-8i adapter. ServeRAID-6M is also supported for external SCSI storage with the EXP400 enclosure.
- ▶ Six internal hot-swap drive bays for up to 440 GB of internal storage (using 73.4 GB disks).
- ▶ Integrated Dual-port Broadcom 5704 PCI-X Gigabit Ethernet.
- ▶ Baseboard Management Controller standard with optional Remote Supervisor Adapter II SlimLine adapter.
- ▶ Supports the IBM Integrated xSeries Adapter for iSeries (IXA) for a direct high speed link to an iSeries server.
- ▶ Three-year warranty, on-site, nine hours per day, five days per week, with a next business day response.

1.9 Virtualization using the System x3800

In many situations, it may not be logical or technically feasible to consolidate all of your users onto one or a few, large scale Presentation Servers in a single data centre. You may have remote offices that require Presentation Servers to be close to users because of bandwidth

limitations. You may have applications that, even with the advances made by Citrix with the Application Isolation Environment and CPU sharing, are still unable to coexist with other applications on the same server. Or, you may run multiple, geographically dispersed data centers to spread the risks associated with natural disasters such as hurricanes, cyclones, tornadoes, earthquakes, and floods.

In any case, it may be advantageous to consider virtualizing some or all of your Presentation Server x64 servers (and other servers) onto a single hardware platform. In this section, we will take a look at the barriers to consolidation for Presentation Server x64. Then we will look at how, using virtualization technology, you can overcome these barriers. We will also look at some of the technology available today that will enable you to accomplish these tasks.

1.9.1 Presentation Server consolidation barriers

As already stated, running Presentation Server x64 on IBM X3 Architecture servers enables businesses to realize TCO savings related to data center power, cooling, and administration. However, some barriers still exist that prevent IT Administrators from replacing many Presentation Server 4.0 systems with a few Presentation Server x64 systems.

- ▶ Application isolation

Although Citrix have made many advances with the Application Isolation Environment and CPU utilization management, there have still been reported instances where customers have had to deploy multiple Presentation Server systems to satisfy the needs of individual applications.

Applications that do not work well under the Application Isolation Environment include applications that install device or kernel drivers, applications that install a service, and applications that require a reboot in the middle of the installation.

- ▶ Remote offices/branches

Many remote offices require at least two to six servers in order to support populations of thirty to fifty users. All but the smallest remote office requires an Active Directory® Domain Controller. They may require a local file server, local mail servers (Lotus® Notes® or Microsoft Exchange), and local application and database servers as well. In recent times, there has been a push to centralize all or most of these services and provide access to them over the WAN via the ICA protocol in Citrix Presentation Server. This model has been and continues to be very successful for many organizations. However, there are times when this model does not fit. A few common barriers always seem to present themselves.

The cost of bandwidth is usually at the top of the list as a major roadblock.

A distributed IT management model is another. An organization may have multiple departments or divisions, each with its own IT staff. While this situation may not be ideal in a business sense, in reality these situations do exist.

On the other hand, you may be a small organization that does not have the resources required to run a multi-server environment, yet the business requires multiple servers to be competitive.

1.9.2 Virtualization Technology and Presentation Server x64

The many benefits of the scalability features incorporated in the IBM X3 Architecture and Presentation Server x64 can still be realized by virtualizing the Presentation Server x64 environment on a suitable hardware platform.

The following points illustrate how the barriers to consolidation of your Presentation Server environment can be overcome by virtualizing Presentation Server x64 servers using VMware Server and the IBM System x3800.

- ▶ Application Isolation

Utilizing virtualization technology, applications are isolated from each other running on separate Presentation Server x64 systems. Using VMware Server, they run isolated from one another just as they would using discreet hardware. Traditional load managed groups can be deployed in this manner.

- ▶ Remote offices/branches

Using the System x3800, remote office server requirements can be hosted on one machine saving the business from having to deploy individual servers for individual roles.

The System x3800 offers a complete package being able to house a half-height tape drive and up to 3.6 TB of storage. This means that Active Directory Domain Controllers, e-mail servers, database servers, and file servers can all be hosted on a single high speed platform.

Additionally, access to applications can be provided by integrating Presentation Server x64 as one or more virtual machines. Having a consolidated infrastructure does not necessarily mean having to beef up WAN links.

Departmental IT Management can deploy dedicated Presentation Server x64 farms on the one server. This saves the business money in an otherwise uneconomical structure where IT resources are duplicated throughout the organization.

Small businesses have the ability to deploy more servers as they grow, simply by creating more virtual machines. Once again, this saves businesses money by utilizing existing hardware resources and not purchasing more servers.

1.9.3 Characteristics of the System x3800

In 1.6, “The X3 Architecture servers” on page 10, we present each of the X3 Architecture servers. The System x3800 is an ideal solution for Presentation Server x64 deployment utilizing a virtualization platform, such as VMware Server or Microsoft Virtual Server.

The System x3800 has all the benefits associated with running Presentation Server x64 on the X3 Architecture in a 64-bit environment:

- ▶ Supports more users through the removal of kernel memory limitations found in IA32 operating systems
- ▶ Increased memory bandwidth reduces bottlenecks in the memory subsystem for memory hungry applications
- ▶ Faster processor to processor communication using the Cache Coherency Filter, Snoopy Filter, and cache directory results in faster response times for CPU hungry applications
- ▶ Active Memory featuring Memory ProteXion, Memory Scrubbing, Chipkill, hot-swappable and redundant components, such as power supplies, fans and hard disk drives, enable a level of availability that is equal to, or better than a large farm of smaller, lower specification servers

- ▶ Plus, with up to 3.6 TB of internal storage, it can easily accommodate the storage requirements of multiple virtualized machines

Figure 1-4 shows the IBM System x3800.



Figure 1-4 The IBM System x3800

The features of the System x3800 include:

- ▶ IBM X3 Architecture feature the XA-64e third-generation chipset.
- ▶ One standard Intel Xeon MP processor, up to 3.66 GHz and 1 MB L2 cache, upgradeable to 4-way. These processors support 64-bit addressing with the Intel EM64T architecture.
- ▶ Xcel4v Dynamic Server Cache.
- ▶ Memory: 1 GB or 2 GB standard expandable to 64 GB (using 4 GB DIMMs), using high performance PC2-3200 ECC DDR2 DIMMs.
- ▶ Active Memory with Memory ProteXion, memory mirroring, memory hot-swap hot-add, and ChipKill.
- ▶ Six full length 64-bit 266 MHz PCI-X 2.0 Active PCI slots.
- ▶ Integrated Adaptec AIC-9410 serial attached SCSI (SAS) controller.
- ▶ Support for internal RAID arrays using and optional ServeRAID 8i controller. The ServeRAID 6M is also supported for external SCSI storage with the EXP400 enclosure.
- ▶ Six internal hot-swap drive bays standard, with an additional six bays optional for up to 3.6 TB of internal storage (using 73.4 GB drives).
- ▶ Integrated dual-port Broadcom 5704 PCI-X gigabit Ethernet.
- ▶ Integrated Baseboard Management Controller and optional Remote Supervisor Adapter II SlimLine adapter.
- ▶ Three year warranty and onsite 9 hours x 5 days a week with next business day response (upgradeable to 24 x 7 same business day service).

1.9.4 Virtualization platforms

The choice of which virtualization platform is the best fit for a virtualized Presentation Server x64 environment really comes down to multi-processor support and 64-bit guest operating system support. Not all versions of the most popular virtualization platforms support 64-bit guest operating systems. Below is a quick review of the most popular virtualization platforms at the time of writing.

Microsoft Virtual Server 2005 R2 Enterprise Edition

In April 2006, Microsoft announced that their popular virtualization product would be made freely available as a download or CD order.

Virtual Server 2005 R2 (VSR2) is available in a 32-bit version and a 64-bit version. VSR2 is an ideal way to consolidate multiple workloads onto a physical server, allowing organizations to make better use of their hardware.

Table 1-2 shows some of the key areas of support as they relate to a Presentation Server x64 environment on an IBM System x3800 server.

Table 1-2 Microsoft Virtual Server 2005 R2

Requirement	Response
Host OS Support	VSR2 is supported on the following host operating systems: <ul style="list-style-type: none"> ▶ Windows Server 2003 Standard, Enterprise, and Datacenter editions (including Service Pack 1) ▶ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition ▶ Windows XP Pro SP2, Pro x64 Edition (for non-production use) ▶ Windows Small Business Server 2003 Standard Edition or Premium Edition
Host CPU Support	Supports up to 32 physical processors. Supports dual-core processors. Supports AMD64 and EM64T processors.
Host RAM support	Minimum of 256 MB. Will scale to 64 GB RAM on a running host. Windows Server 2003 x64 or a running host. Windows Server 2003 (IA32) with PAE enabled.
Guest OS Support	Runs most major x86-IA32 operating systems, including support for Linux-based Oses. VSR2 takes advantage of the broad hardware support of the host operating system. VSR2 does not support x64 operating systems. Supports up to a maximum of 64 virtual machines (depending on the limits of the resources of the host machine).
Guest CPU Support	VSR2 supports a single virtual processor in a guest virtual machine. The processor ID is based on the ID of the real processor.
Guest RAM Support	Guest virtual machines can be allocated up to 3.6 GB RAM. The host memory cannot be over allocated to virtual machines.

VMware Server

VMware Server is a free for download virtualization platform from VMware. The product was announced at a time when VMware had discontinued the GSX version of their virtualization platform. So it could be argued that VMware Server is the free version of GSX server. This is not entirely true, because VMware Server supports a Linux host OS while GSX server did not. There are also other changes between the products that we will not list here. However, virtual machines created under VMware GSX Server will run unaltered in a VMware Server environment. Alternatively, virtual machines created under VMware Server will run unaltered in a VMware ESX Server environment. This allows organizations to step up to the full functionality of VMware ESX Server and Virtual Infrastructure as they grow.

At the time of writing, VMware Server is in Beta 2 with an expected GA date in the second half of 2006.

VMware Server is a robust yet easy to use product that brings the powerful benefits of virtualization to every server. It is based on VMware's proven technology, which has been used by thousands of customers for more than six years

Table 1-3 shows some of the key areas of support as they relate to a Presentation Server x64 environment on an IBM System x3800 server.

Table 1-3 VMware Server

Requirement	Response
Host OS Support	VMware Server is supported on the following host operating systems: <ul style="list-style-type: none"> ▶ Linux Server ▶ Microsoft Windows 2000 Server and Advanced Server ▶ Windows Server 2003 Web, Standard, Enterprise, and x64 Editions
Host CPU Support	Supports up to 16 physical processors. Supports dual-core and AMD/EM64T processors. Experimental support for the Intel EM64T virtualization Technology extensions.
Host RAM support	Minimum 512 MB. Will scale to 64 GB RAM on Windows or Linux hosts with large memory support (x64) or that are PAE enabled.
Guest OS Support	<ul style="list-style-type: none"> ▶ Supports most major x86-IA32 operating systems ▶ Supports 64-bit guests including Windows Server 2003 Enterprise, Standard, and Web editions, Windows XP Pro x64, and 64-bit Linux distros <p>VMware Server performs an internal check to verify that the host processor is a supported 64-bit platform. If the host CPU is not a supported 64-bit platform, and an error message is displayed indicating that the hardware of your host machine is not compatible with 64-bit guest operating systems.</p>

Requirement	Response
Guest CPU Support	VMware Server supports 2-way virtual SMP guests on SMP hosts. The processor ID in the guest depends on the processor ID of the host.
Guest RAM Support	Guest virtual machines can be allocated up to 3.6 GB RAM. The number of concurrent virtual machines and the amount of RAM contained in the host should be considered when allocating RAM to virtual machines.

VMware ESX Server

ESX Server is VMware's enterprise-class virtualization platform for consolidating and managing servers in mission critical environments.

Unlike VMware Server, which installs upon a host operating system and is reliant on the compatibility and stability of the underlying OS, VMware ESX Server is a bare metal install. It contains it's own highly optimized Linux-based operating system.

Using VMware ESX Server and VMware VirtualCenter organizations can deploy multiple hosts and manage them as a single pool of resources. Workloads can be distributed evenly across multiple hosts, thereby maximizing resource usage. Virtual machines can also be moved between VMware ESX Server hosts while running, enabling dynamic load distribution and downtime management.

Table 1-4 VMware ESX Server

Requirement	Response
Host OS Support	ESX Server is a bare-metal install. It contains its own OS. ESX Server will not install as an application under Windows or Linux OSes.
Host CPU Support	2-16 Intel Xeon or AMD Opteron (32-bit mode) processors.
Host RAM support	512 MB - 64 GB.
Guest OS Support	Runs most major x86-IA32 operating systems, including support for Linux-based OSes. ESX Server 2.5 does not include support for 64-bit OSes.
Guest CPU Support	Virtual Intel or AMD x64 uniprocessor system (depending on the PID of the host). Virtual Intel or AMD 2-way SMP system when the VMware Virtual SMP add-on is included.
Guest RAM Support	Guest virtual machines can be allocated up to 3.6 GB RAM. The number of concurrent virtual machines and the amount of RAM contained in the host should be considered when allocating RAM to virtual machines.

Using the above tables as a guide, it can be seen that VMware Server is the only platform available that will support 64-bit guest operating systems. This is a critical area of support for virtualizing a Presentation Server x64 environment on the IBM System x3800.

It is highly likely that further support for 64-bit guests will be announced in the near future. Both VMware and Microsoft have the next generation of their enterprise-class virtualization platforms currently in development. Microsoft is set to include *Hypervisor* virtualization technology in the next release of the Windows Server system codenamed “*Longhorn*”. VMware will be releasing Version 3 of its enterprise-class virtualization platform, ESX Server, in the later half of 2006.

Intel is supporting the drive towards virtualization by developing a set of extensions for EM64T processors called *Virtualization Technology*. VMware and Microsoft are supporting Intel Virtualization Technology in the upcoming releases of their next products. Intel Virtualization Technology is available in selected Intel Xeon processors at the time of writing.



Introducing Citrix Presentation Server 4.0

This chapter introduces what is new in Citrix Presentation Server 4.0 in comparison to earlier versions, specifically Citrix MetaFrame XP and Citrix MetaFrame Presentation Server 3.0. We discuss the benefits of upgrading to Citrix Presentation Server 4.0 x64 Version running on Windows Server 2003 x64 Edition.

This chapter has the following topics:

- ▶ 2.1, “What is new in Citrix Presentation Server 4.0” on page 28
- ▶ 2.2, “Citrix Presentation Server 4.0 x64 Edition” on page 36
- ▶ 2.3, “Why upgrade to Presentation Server 4.0?” on page 36
- ▶ 2.4, “Benefits of Windows Server 2003 x64 Edition” on page 37
- ▶ 2.5, “Why upgrade to Presentation Server 4.0 x64 Edition?” on page 38

2.1 What is new in Citrix Presentation Server 4.0

In this section, we show the new features of Citrix Presentation Server 4.0 in comparison to Citrix MetaFrame XP and Citrix MetaFrame Presentation Server 3.0.

Citrix Presentation Server 4.0 adds significant new features to an existing MetaFrame XP or Presentation Server 3.0 system. There are many new features and we will pick out some highlights of these. Some of the shown features are also implemented in Presentation Server 3.0.

There are three versions of Presentation Server 4.0: Standard, Advanced, and Enterprise. The features depend on the version you intend to use, as shown in Table 2-1.

Table 2-1 Packaging and features

Feature	CPS 4 Standard	CPS 4 Advanced	CPS 4 Enterprise
Operating system support			
Windows 2000	Supported	Supported	Supported
Windows Server 2003	Supported	Supported	Supported
Windows Server 2003 x64	Supported	Supported	Supported
Solaris™	No	No	Supported
IBM-AIX	No	No	Supported
HP-UX	No	No	Supported
Functions			
Load Manager	No	Supported	Supported
4.0 Printing	Supported	Supported	Supported
Access Client	Supported	Supported	Supported
USB Support	No	Supported	Supported
Application Isolation	No	No	Supported
virtual IP addressing	No	Supported	Supported
CPU Management	No	No	Supported
Memory Management	No	No	Supported
Resource Manager	No	No	Supported
Network Manager	No	No	Supported
Installation Manager	No	No	Supported
MOM support	No	No	Supported
Components			
ICA Clients 9.x	Supported	Supported	Supported
Web Interface 4.0	Supported	Supported	Supported
Secure Gateway	Supported	Supported	Supported
Conferencing Manager	Supported	Supported	Supported

We describe the following features in more detail so you can gain the most improvements with these features and the X3 architecture servers:

- ▶ 2.1.1, “CPU utilization management” on page 29
- ▶ 2.1.2, “Virtual memory optimization” on page 30
- ▶ 2.1.3, “Application Isolation Environment” on page 31
- ▶ 2.1.4, “Improvements on printing” on page 33
- ▶ 2.1.5, “Virtual IP addressing” on page 34
- ▶ 2.1.6, “New licensing design” on page 35
- ▶ 2.1.7, “Support for Windows Server 2003 x64 Edition” on page 35

For a complete list of new features in comparison to older versions down to Citrix MetaFrame 1.8, refer to the Citrix Migration white paper *Maximizing the Value of Your Citrix Investment - Why Upgrade to Citrix Presentation Server 4.0* at:

http://www.citrix.com/English/ps2/products/documents_onecat.asp?contentid=186&cid=White+Papers

2.1.1 CPU utilization management

In CPU utilization management, two new features are implemented to improve user density:

- ▶ CPU reservation

This feature reserves an amount of CPU time that always can be used if needed. If a user is not using all his CPU time, it can be used by another user and vice versa. By default, 20% of the CPU is reserved.

- ▶ CPU fair share

This is a relative amount of CPU time based on the number of other users that have shares. So, if there are four users using the CPU with every user having an equal share, every user has a guaranteed CPU time of 25%. If one of the users does not need all of his CPU time, another user can use it.

CPU Fair Share is shown in Figure 2-1. Using this feature is discussed in 4.7, “Performance tuning” on page 88.

With the CPU Utilization Management feature, it is possible to keep applications or processes from using 100% of the CPU time and to increase the density of users on the server.

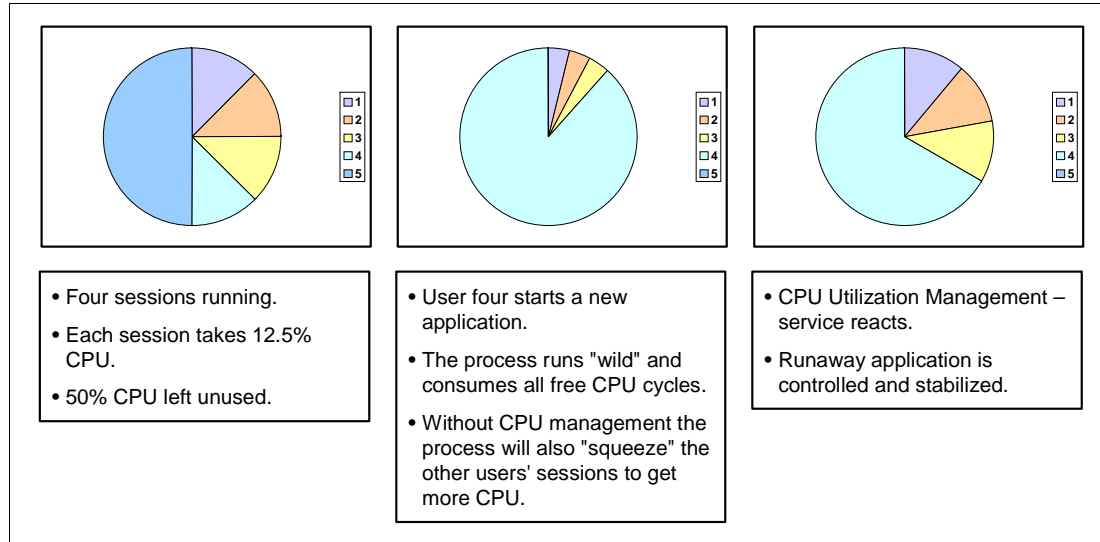


Figure 2-1 CPU Fair Share

The CPU utilization management normalizes the usage of each user on a Presentation Server by smoothing out the normal CPU peaks that most applications have. It also limits very CPU intensive applications to a fair-share of the CPU so that one user does not inadvertently have an inappropriate adverse affect on the other users. As a result, the CPU efficiency is increased and more users can be bound on one CPU.

Note: The CPU Utilization Management feature is not activated by default.

2.1.2 Virtual memory optimization

The virtual memory optimization feature is intended to improve performance and user density. This technology reduces the amount of used memory by rebasing DLLs to an optimized virtual address. It also prepares the import section of the DLL with the virtual address of all imported symbols before loading the DLL. So, it also improves the initialization time. We can increase the amount of users per server when different users load the same DLL by rebasing and decrease the application loading time. How the virtual memory optimization feature works is shown in Figure 2-2 on page 31.

This feature is not activated by default. You have to pay attention using this feature because it is using approximately 20% of CPU processing power for automatic optimization. So we recommend running the automatic optimization once a day during an offpeak time if this is possible.

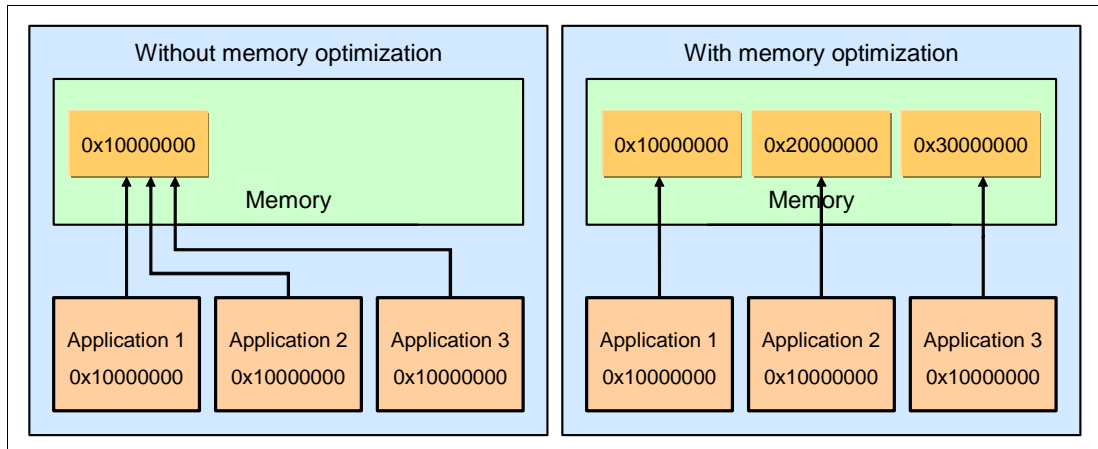


Figure 2-2 Virtual memory management

Referring to Figure 2-2:

- ▶ Without memory optimization configured, every application attempts to write to the same base memory address space. The first application is successful and occupies that space, but all others must get reallocated.
- ▶ With memory optimization configured, applications are pre-assigned an alternate memory address space and can be loaded into that address space efficiently.

For more information about this feature, refer to the Citrix Knowledge base Article CTX106023 at:

<http://support.citrix.com/article/CTX106023>

2.1.3 Application Isolation Environment

Application Isolation is a technology to make applications in a Citrix Presentation Server environment more compatible and work better together. On a Terminal Server, the applications have to share system components and resources. Under some circumstances, you can be faced with the following problems:

- ▶ Multiple instances of an application cannot be opened.
- ▶ Different versions of the same application cannot run on a single server.
- ▶ Applications cannot share system resources properly.
- ▶ Applications are not Terminal Services aware.
- ▶ Applications use hardcoded file paths or settings.

All of the above issues are solved with Application Isolation. Application Isolation redirects specific system resources like parts of the registry and the file system. The installed application gets a redirected view of the system resources and is running in an isolated sandbox. All special calls that the application will make are automatically redirected. For example, a call to a DLL in C:\Windows\System32 is automatically redirected to c:\Program Files\Citrix\AIE\AIE001\c\Windows\System32. This is shown in Figure 2-3.

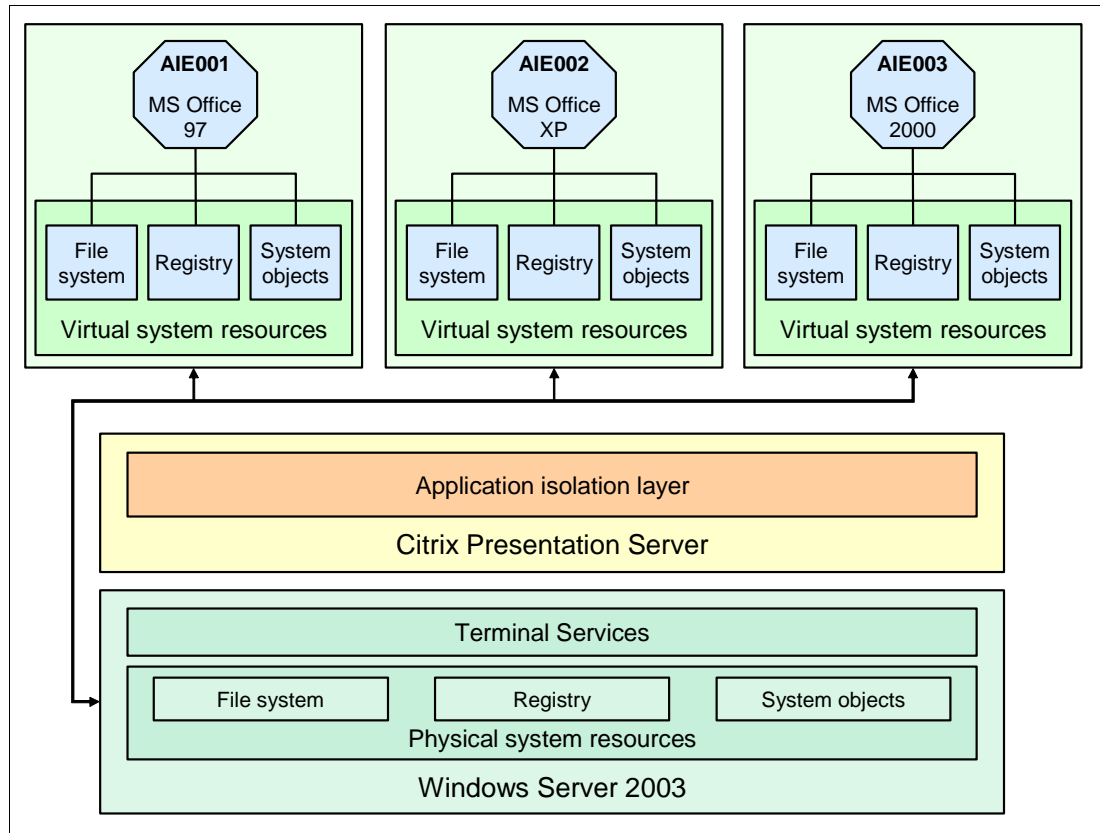


Figure 2-3 Application Isolation Environment (AIE)

The application isolation layer has control of the different application isolation environments. It maps the physical resources of the operating system to the virtualized resources of the application isolation environment.

If you want to use Application Isolation, you should consider the following:

- ▶ Application Isolation can significantly increase the hard drive and registry space needed.
- ▶ Possible increase of execution overhead depending on the applications usage of the registry and file system.
- ▶ Decreasing the number of users per server.

For more information about this feature, refer to the Citrix Knowledge base Article *CTX106868-Isolation Environments on CPS 4* at:

<http://support.citrix.com/article/CTX106868>

However, not all applications are Application Isolation aware. See the following list for tested applications. If you have an application that is not listed below, you should test it before installing it on a production Presentation Server.

- ▶ Microsoft Office 97, 2000, XP, 2003
- ▶ Oracle Client 8i and 9i
- ▶ SAP GUI front end 4.6 (Compilation 4), 6.2 (Compilation 1), and 6.2 (Compilation 5)
- ▶ StarOffice™ 5.2 and 7.0
- ▶ WinZip 9.0
- ▶ Acrobat Reader 4.05, 5.1, 6.0 and 7.0
- ▶ Crystal Reports 7.5 and 8.0
- ▶ FoxPro 2.6 for DOS
- ▶ JRE™ 1.3 and 1.4
- ▶ Mozilla Firefox 1.0
- ▶ Internet Explorer 5.5 and 6.0
- ▶ Lotus SmartSuite® 97 and 2000
- ▶ MDAC 2.8
- ▶ Project 98
- ▶ WordPerfect Office 11

The latest list of supported applications is listed in Citrix Knowledge base Document *CTX106100 - Application Isolation Environment Supported Applications*, at:

<http://support.citrix.com/article/CTX106100>

For more information concerning common limitations and compatibility issues of Presentation Server 4 Application Isolation Environment, see:

<http://support.citrix.com/article/CTX109254>

2.1.4 Improvements on printing

Citrix Presentation Server 4.0 includes a complete new and redesigned printing subsystem. It has many improvements in usability, security, and performance. The three highlights are:

- ▶ New printer subsystem
All session printer management is now handled by a new user-mode system service called Citrix Print Manager Service. All printer management is centralized in a secured service process for creating and deleting printers.
- ▶ New universal printer driver 3 (UPD)
The new universal printer driver 3 leverages the Enhanced Metafile (EMF) file format to increase printing performance up to four times more than the earlier versions. This is because the spool files are 80-90% smaller through the EMF format.
- ▶ Simplified printer management
Printer connections are managed through policies.
Auto-created printers can be managed by the user and the printing on auto-created printers is spooled directly on the users client to save resources on the server and improve performance.
Printer naming of auto-created printers is changed so it is more transparent to the user.

For more information of the new printing features see the *Presentation Server 4.0 Printing Enhancements* White paper, at:

http://www.citrix.com/English/ps2/products/documents_onecat.asp?contentid=186&cid=White+Papers

2.1.5 Virtual IP addressing

Virtual IP addressing means that we now can use separate IP addresses for a user's session. So it is possible to use more applications that depend on different IP addresses per user. Examples are:

- ▶ CTI (Computer Telephony Integration)
- ▶ Siebel
- ▶ FileNet and Symitar

An example is shown in Figure 2-4 and Figure 2-5.

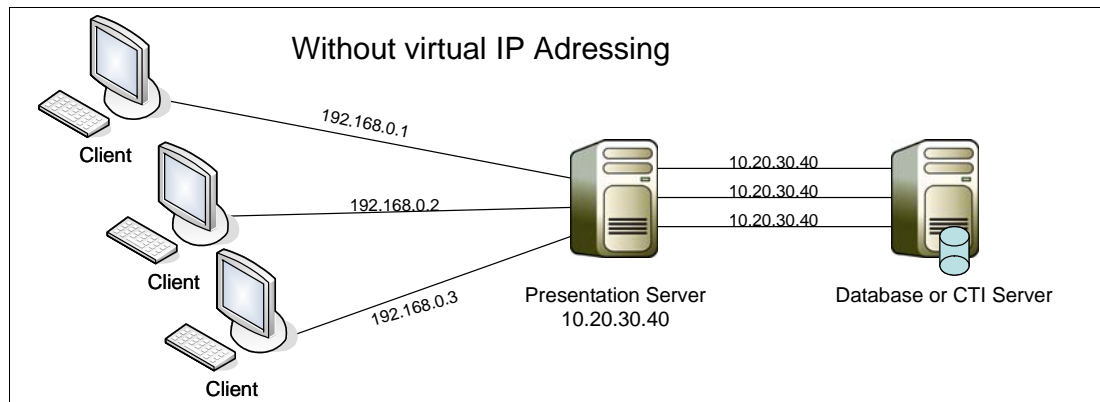


Figure 2-4 Without virtual IP addressing

Without virtual IP addressing configured, each session would not be distinguished by the back-end database, since the IP address of the Presentation Server would be used for communication with the application.

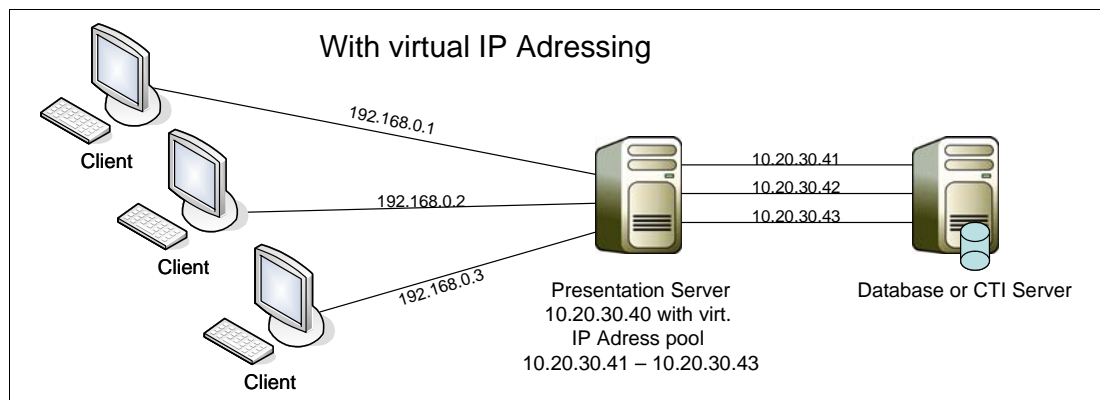


Figure 2-5 With virtual IP addressing

With virtual IP addressing configured, the IP addresses that each ICA session uses to communicate with the back-end database is unique.

Virtual IP addressing is managed through the Presentation Server Console. However, it is only possible to define a IP address range that can be used. It is not possible to assign a specific IP address for a specific user or client. If a static IP address in the session for the client is required, the IP address of the client has to be passed through the ICA session. For more information about how to configure virtual IP addressing, refer to the *Citrix Presentation Server Administrator's Guide* at:

<http://support.citrix.com/article/CTX106319>

2.1.6 New licensing design

Starting with Citrix Presentation Server 3.0, there is a new licensing design implemented. The licenses are no longer in the Citrix Database and they are no longer seen in the Presentation Server Console. With Presentation Server 4.0, a separate Licensing Server is required.

The License Server is based on Internet Information Server 5 or 6 or Apache HTTP Server 2.0.49. We recommend not installing the Citrix License Server on one of the Citrix Presentation Servers — it should be another dedicated server.

The licenses are now activated and applied through a license file that is downloaded from the personal Citrix Web site.

Note: Citrix Presentation Server for UNIX® licenses are implemented into the licenses of Citrix Presentation Server 4.0 Enterprise Edition and extra licenses are no longer available. So the only way to use Citrix Presentation Server for UNIX is to buy Citrix Presentation Server 4.0 Enterprise Edition or Citrix Access Suite 4.0.

An example of license usage in a mixed UNIX and Windows environment is shown in Figure 2-6.

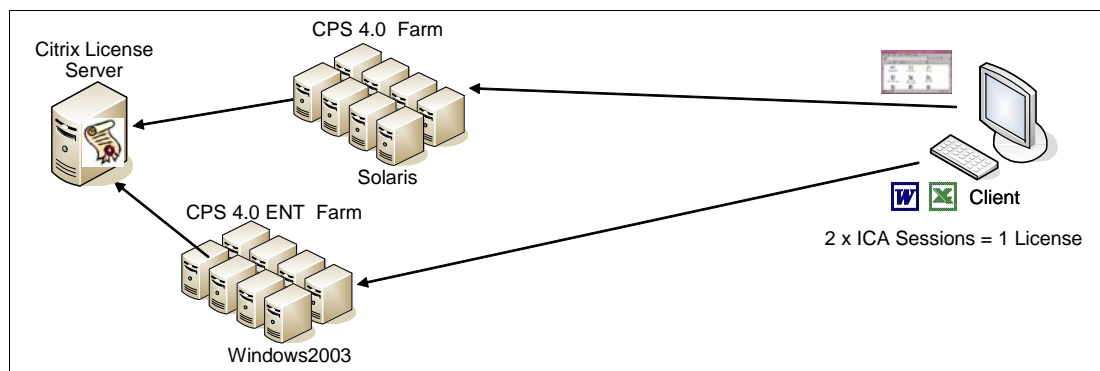


Figure 2-6 Citrix Presentation Server 4.0 Licensing.

For more information about installing, deploying, and managing the new Citrix License Server, refer to the Citrix Knowledge base Article CTX106282 *Citrix Access Suite Licensing Guide* at:

<http://support.citrix.com/article/CTX106282>

2.1.7 Support for Windows Server 2003 x64 Edition

The 64-bit processors from AMD and Intel have brought 64-bit capabilities to the mainstream. The 64-bit capable processors are supported with the release of Windows Server 2003 x64 Edition. Presentation Server 4.0 x64 Edition is available to support Windows Server 2003 x64 Edition to take full advantage of the scalability and performance of 64-bit processors, such as those in the IBM System x servers with X3 Architecture. We discuss this more in the next section.

2.2 Citrix Presentation Server 4.0 x64 Edition

In this section, we describe the differences between the 32-bit and x64 editions of Presentation Server 4.0.

Most of the features from the Presentation Server 4.0 32-bit Edition were ported to the x64 Edition. However, there are some functions that are not supported in Presentation Server 4.0 x64 Edition. These are:

- ▶ Novell NDS (no 64-bit Novell Client)
- ▶ Smartcards
- ▶ Secure ICA
- ▶ .NET 2.0
- ▶ Client database- and Client-Assistant
- ▶ Oracle
- ▶ Port sharing between IIS and XML-Service
- ▶ TWAIN (only 32-bit applications supported)

Presentation Server 4.0 x64 Edition can only run on Windows Server 2003 x64 Edition. It is not possible to install it on a 32-bit Windows 2000 Server or Windows Server 2003.

Note: Citrix License Server is not supported running on a Windows Server 2003 x64 Edition system. The Citrix Web interface is also not supported on the 64-bit operating system.

For more information of the advantages and benefits of upgrading to Citrix Presentation Server 4.0, especially the x64 Edition, refer to the following two sections:

- ▶ 2.3, “Why upgrade to Presentation Server 4.0?” on page 36
- ▶ 2.5, “Why upgrade to Presentation Server 4.0 x64 Edition?” on page 38.

2.3 Why upgrade to Presentation Server 4.0?

In this section, we describe some of the benefits of why we recommend migrating to Presentation Server 4.0, regardless of whether you use the 32-bit or the 64-bit edition.

As we mentioned in 2.1, “What is new in Citrix Presentation Server 4.0” on page 28, there are many new features that are good reasons for upgrading to Citrix Presentation Server 4.0. The best reason for upgrading to Presentation Server 4.0 is the performance improvements. The user density on the same hardware can be significantly increased and this is why upgrading has a positive impact on scale-up scenarios. We also recommend upgrading because user satisfaction about visibility and performance will increase. The key benefits are listed below:

- ▶ Broader application support

Many applications are not designed to run in a multi-user environment. Some applications may conflict with others when installed on the same server. With the new features of Application Isolation and virtual IP Addressing, it is now possible to overcome this problem. So more applications can be published on fewer servers.

- ▶ Increased user density

With the new features of CPU utilization management and virtual memory management, more applications can be installed on one server and the number of users can be increased up to 25% per server.

- ▶ Improved printing

The new Universal Printer Driver makes printing up to four times faster while reducing bandwidth and memory. Almost any printer can be used, including advanced printer functions like trays and stapling. It is easier to print to the closest printer of the user while roaming from device to device.

2.4 Benefits of Windows Server 2003 x64 Edition

In this section, we take a short look at the benefits of the Windows Server 2003 x64 Edition and what advantages this operating system has to achieve your business goals.

Windows Server 2003 x64 Edition supports both the AMD AMD64 and the Intel EM64T processors. One of the advantages of using Windows Server 2003 x64 Edition is the increased virtual address space. A 32-bit operating system has a limited address space of 32 bits, so a maximum of 4 GB of memory can be addressed.

This 4 GB of memory in a 32-bit address space is divided into two equal parts:

- ▶ 2 GB are allocated to the processes.
- ▶ 2 GB are allocated to the operating system that is used for kernel mode memory, system cache, and drivers.

The only way to go beyond this limitation is to use the /PAE (Physical address Extensions) switch in the BOOT.INI of the server. The memory allocated with the /PAE switch is allocated for the user space, so applications that are programmed to use the Address Windows Extensions (AWE) API can take advantage of the memory. However, the kernel space is still limited to 2 GB.

Windows Server 2003 x64 Edition can now address up to 16 TB of address space. This is also divided into two equal parts. So it is possible to get 8 TB of kernel space and 8 TB of user space.

The memory limits are shown in Table 2-2.

Table 2-2 Memory limitations comparison of Windows Server 2003

Description	32-bit	64-bit (x64)
Total virtual address space	4 GB	16 TB
Virtual address space per 32-bit application	2 GB (Note 1)	2 GB (Note 2)
Virtual address space per 64-bit process	Not applicable	8 TB
Virtual address space for operating system kernel	2 GB (Note 1)	8 TB
Paged Pool	470 MB	128 GB
Non-paged pool	256 MB	128 GB
System cache	1 TB	1 TB
Notes:		
1. 3 GB for the application and 1 GB for the kernel if the system is booted with the /3GB switch		
2. 4 GB if the 32-bit application has the LARGEADDRESSAWARE Flag set (LAA)		

Windows Server 2003 x64 Edition can run 64-bit and 32-bit applications; however, all drivers must be 64-bit versions. For more information about Windows Server 2003 x64 Edition, refer to:

<http://www.microsoft.com/windowsserver2003/64bit/x64>

For additional information, see the IBM Redpaper *Introducing Windows Server x64 on IBM @server xSeries Servers*, REDP-3982, at:

<http://www.redbooks.ibm.com/abstracts/redp3982.html>

2.5 Why upgrade to Presentation Server 4.0 x64 Edition?

In 2.3, “Why upgrade to Presentation Server 4.0?” on page 36, we described the benefits to upgrading to Presentation Server 4.0, either 32-bit or 64-bit. In this section, we discuss the advantages and benefits of upgrading to the x64 Edition of Presentation Server 4.0.

User density on a server in a Presentation Server environment is highly dependent on four key factors:

- ▶ Operating system
- ▶ Presentation Server version
- ▶ Server hardware
- ▶ Application

A Presentation Server environment is very memory intensive. Because of the improved memory capabilities of Windows Server 2003 x64 Edition, Presentation Server 4.0 x64 Edition does benefit from this enhancement.

Figure 2-7 shows a comparison of the user density on a 4-way server between Windows 2000 Server, Windows Server 2003, and Windows Server 2003 x64 Edition, all running Presentation Server 4.0 x64 Edition.

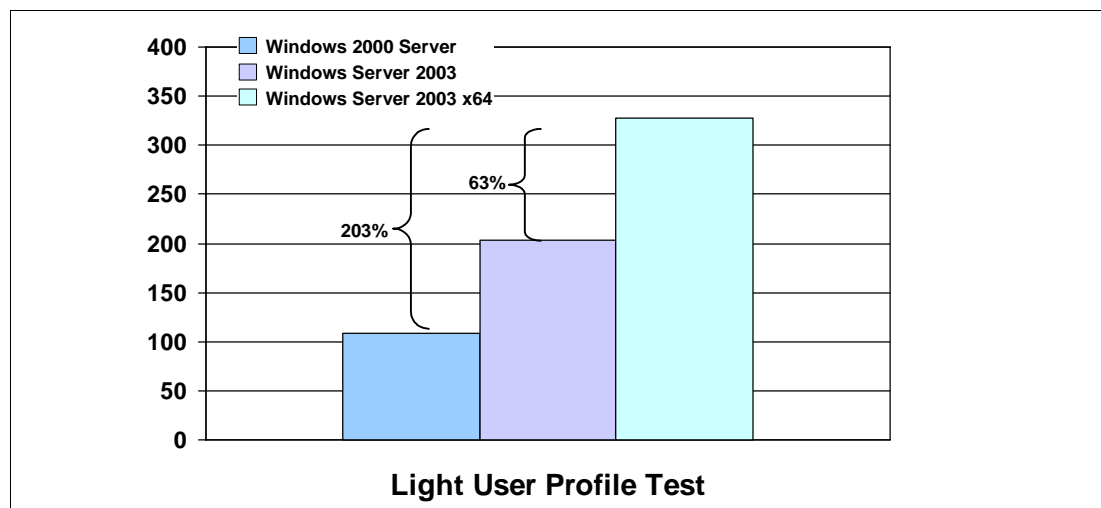


Figure 2-7 Max. number of users comparison on a quad processor hardware

In a Presentation Server 32-bit environment, the server is bound by the kernel memory limits (paged pool, non-paged pool, and system cache) of the operating system, as shown in Table 2-2 on page 37 and discussed in 1.4.1, “Kernel memory” on page 5. The more users that log onto the server, the more kernel and user memory is needed.

When upgrading to four or more processors, the processors may only use 20% of their capacity, but you do not get any more users on the server, because the operating system is running out of kernel memory. As a result, there is no improvement in user density when upgrading from a two-way system.

With the x64 environment, these kernel memory limits are obsolete, so scaling up to four or more processors achieves a high improvement in user density. Now there will be other limits of the hardware before running out of kernel memory.

For more information on performance testing, see the Citrix White paper *Citrix Presentation Server for Microsoft Windows Server 2003 x64 Edition - Performance and scaling capabilities*, at:

http://www.citrix.com/English/ps2/products/documents_onecat.asp?contentid=186&cid=White+Papers



Moving to a 64-bit Citrix Environment

In this chapter, we recommend the areas of concern that should be addressed before upgrading your farm to Presentation Server x64. We present the upgrade paths that are supported by Citrix. We also give some real-world migration scenarios that can be used to help in the design phase of a migration project. We present a few strategies around how to best implement load balancing in a mixed 32-bit/64-bit Presentation Server environment.

This chapter discusses the following topics:

- ▶ 3.1, “Upgrade planning” on page 42
- ▶ 3.2, “Supported upgrade paths” on page 45
- ▶ 3.3, “Migration strategies” on page 46
- ▶ 3.4, “Operating in a mixed 32/64-bit environment” on page 53
- ▶ 3.5, “Lab upgrade” on page 54

3.1 Upgrade planning

As previously mentioned, there are some compelling reasons to consider migrating or upgrading to Presentation Server 4.0 or Presentation Server x64.

- ▶ Increased user density
 - CPU utilization management
 - Virtual memory optimization
- ▶ Broader application support
 - Application Isolation Environment
 - Virtual IP Addressing
 - USB support for PDAs, scanners, and cameras
- ▶ Improved printing
 - Up to four times faster printing using less bandwidth and memory
 - Support for advanced printer features, like multiple trays and output bins, staplers, and so on.
 - Session printer policies for printing to the nearest printer

Presentation Server x64 combined with the IBM X3 Architecture servers dramatically increases user density, which presents a total cost of ownership case that is hard to overlook.

- ▶ Fewer administrators required to maintain fewer servers
- ▶ Reduced overhead required to manage patches, service packs, and hotfixes
- ▶ Less power, cooling, and space requirements in computer rooms
- ▶ Lower server-based licensing costs
- ▶ 32-bit applications run seamlessly

However, as with any upgrade or migration, careful consideration should be given to the health and suitability of your system as a whole before you begin. Citrix recommends performing an infrastructure analysis prior to undertaking any project to upgrade your environment. Some of the key evaluation areas to examine include:

- ▶ Server operating systems

Presentation Server 4.0 comes in different packages for Windows 2000 Server and Windows Server 2003. The separate packages will only install on the platform for which they were intended. Similarly, Presentation Server x64 will only install on Windows Server 2003 x64 Edition.

If you will be upgrading your operating system from Windows 2000 Server to Windows Server 2003, we recommend that you perform a complete rebuild of the entire server rather than an in-place upgrade. Citrix does not support an in-place upgrade of the operating system of an existing Presentation Server.

It is not possible to upgrade to Windows Server 2003 x64 Edition from either Windows 2000 Server or Windows Server 2003. If your goal is to upgrade from a 32-bit environment of Windows Server and Presentation Server to a 64-bit environment incorporating the 64-bit versions of both products, it will be necessary to perform a fresh installation of the operating system and build the server from scratch.

- ▶ Server hardware

It may be time to consider replacing some of your older farm servers. Many older servers do not support the Intel EM64T 64-bit extensions. If you plan to be moving to a Presentation Server x64 environment, it is important to ensure that the hardware is capable of running a 64-bit operating system.

Additionally, Presentation Server 4.0 runs much better in an environment that is not constrained by a lack of server RAM. Consider upgrading the amount of RAM in existing servers.

► Farm design

Since many Presentation Server implementations start out small and end up as a sprawling, out-of-control mess, now is the time to consider whether the needs of your network, users, and applications have changed since the farm was first installed:

- Network segments may have been added.
- New sites may have been implemented.
- Remote access needs may have grown.

These are factors that should be considered when the time comes to upgrade to Presentation Server 4.0 or Presentation Server x64. It may be better to consider implementing a new farm rather than building upon an old farm that may need major redesign.

► Application environment

All existing and future applications should be tested in an isolated lab that reflects your projected environment.

► Printing

Citrix has re-designed the entire printing sub-system between MetaFrame Presentation Server 3 and Presentation Server 4.0.

Your printing architecture should also be redesigned to make use of the new functionality found in the Citrix Universal Printer driver (UPD).

At the very least, an analysis exercise should be undertaken to determine the UPD's usability within your environment.

► Network architecture

It is prudent to review the network at this stage before an upgrade to Presentation Server 4.0 is undertaken.

It may be necessary to consider implementing quality of service (QoS) or a packet shaping/prioritization device in the network to guarantee service level agreements to the business.

It is also a good time to produce a baseline of network performance before the upgrade takes place. That way, further baselining and analysis can give you an exact measurement of the effect (if any) on the network the upgrade to Presentation Server 4.0 has had.

Having baseline figures available is also very handy when tendering for a new network provider.

► Lab environment

A dedicated, isolated lab environment is one of the most crucial pieces of trouble-shooting equipment that any organization can possess. A lab environment should reside on a completely isolated test network. If need be, you can connect it to your main Internet connection via your production network, but only if it is firewalled away from your production environment.

The test environment should reflect the production environment in every way - from the IP addressing scheme, to device names, to Active Directory OU structure, and everything in between. This is one of the reasons why it should be on an isolated LAN.

A dedicated test bed makes patch management, script development, application testing, upgrade testing, and documentation so much easier and risk free.

A test lab can be created on virtual machines using a minimal amount of hardware. There is no need to set aside five of the latest servers, routers, and switches in order to create a lab environment.

- ▶ Change control

An effective change control regime is a central to providing a stable, consistent environment to the business. A stable environment is also crucial for a successful upgrade to Presentation Server 4.0. That is, if it is a stable environment to begin with, the upgrade process should be able to be executed and re-executed in the lab with consistent, reproducible outcomes every time.

Change control is the key to providing a stable platform from which to upgrade.

- ▶ Systems management

A decent systems management system should be able to provide the business with server performance metrics, asset data, event notification, alerts, and software inventory.

Just as it is important to baseline the network before any work is undertaken, it is just as important to have a baseline of the servers as well. A baselining exercise will give you valuable information about the relative health of your servers and farm. Any problem areas can be addressed before any attempt is made to upgrade or migrate the servers.

Systems management software comes in many shapes and sizes. Some systems may cost many thousands of dollars and do things that save a business many thousands of dollars. You may only require a system that performs certain tasks, and you will be pleasantly surprised that you can pick it up for free.

- ▶ Support structure

The help desk will need to be informed of any impending upgrade to the infrastructure. Are they equipped for a potentially greater number of calls? Hopefully, you will not be expecting very many. The help desk may require additional training, or at least notification of the upgrade.

- ▶ User access needs during the upgrade

Take into consideration the time frame in which the upgrade is to take place. Is the upgrade to be completed in a few days, maybe over a weekend? Is it necessary for users to be able to access their applications during the upgrade? What would be the impact to the business if published applications were not available during the upgrade?

Web Interface or the Program Neighborhood Agent are the two easiest methods with which to provide access to applications. This is because they can be centrally controlled and configured. Citrix recommends the use of these two technologies during an upgrade to limit the impact to end users and the business. In 3.3, "Migration strategies" on page 46, we will show you how you can configure Web Interface during the upgrade period to reduce the impact of outages to the business.

- ▶ Citrix licensing during an upgrade

If you are upgrading from MetaFrame Presentation Server XP or MetaFrame 1.8, during the upgrade, you will need to provide connection licenses for both the previous platform and Presentation Server 4.0 or Presentation Server x64. This is because the licensing mechanisms changed between MetaFrame XP Feature Release 3 and MetaFrame Presentation Server 3. The same issue is not encountered when upgrading from Presentation Server 3 to Presentation Server 4.0 (or x64).

One advantage of subscribing to the Subscription Advantage licensing scheme is that users maintain connectivity to MetaFrame XP and previous MetaFrame farms during the migration period.

- ▶ Break the upgrade into manageable chunks

If you have large or multiple farms, consider breaking the upgrade into manageable chunks. Every organization is unique and the upgrade plan you employ will be influenced by the needs of the business.

- ▶ Project communication

It is very beneficial to the success of a project to communicate any outages or changes to the end-user population. The help desk would like a bit of forewarning as well. That way, if there are delays or unscheduled interruptions, users can prepare other arrangements to continue working.

- ▶ Disaster Recovery

Do you have a backout plan just in case the upgrade goes wrong?

You should have. If not, create one, and test it in your isolated test lab.

Citrix encourages engaging the services of a well qualified consulting firm that has experience in similar migrations or upgrades. Employing an external consulting firm will save an organization time and money.

3.2 Supported upgrade paths

Table 3-1 shows the supported migration paths from earlier versions of Presentation Server.

Citrix does not support performing an in-place upgrade of the server's operating system. If you upgrade the operating system of your server, you must install Presentation Server 4.0 as a new installation.

Note: Citrix makes available separate versions of Presentation Server 4.0 for Windows 2000 Server and Windows Server 2003. They do not differ in functionality, options, or support. Citrix makes no differentiation between the two versions when discussing them in their documentation.

A direct upgrade from a 32-bit version of the Windows operating system to a 64-bit version of the Windows operating system is not possible. If your requirement is to upgrade from a 32-bit version of Presentation Server (from MetaFrame XP FR3 or MetaFrame Presentation Server 3) to Presentation Server x64, a complete rebuild of the server is required. Further information regarding migration and co-existence is contained in the preceding sections.

Table 3-1 Supported upgrade paths

Existing installation	Supported upgrade path
MetaFrame XP FR3 for Windows 2000 Server	Presentation Server 4.0 for Windows 2000 Server with SP4
MetaFrame XP FR3 for Windows Server 2003	Presentation Server 4.0 for Windows Server 2003
MetaFrame Presentation Server 3.0 for Windows 2000 Server	Presentation Server 4.0 for Windows 2000 Server with SP4
MetaFrame Presentation Server 3.0 for Windows Server 2003	Presentation Server 4.0 for Windows Server 2003

You cannot upgrade to Presentation Server 4.0 from Presentation Server versions prior to MetaFrame XP Presentation Server Feature Release 3. We recommend upgrading these servers to MetaFrame Presentation Server 3 first, then upgrading to Presentation Server 4.0. Citrix Subscription Advantage subscribers have the advantage of being able to upgrade their product and connection licenses without having to purchase extra licenses for each step of the upgrade from MetaFrame 1.8 to Presentation Server 4.0. Customers without Subscription

Advantage will need to purchase product and connection licenses for each step of the upgrade. This could be a very expensive exercise in some circumstances. If this is the case, it may be more economical to install an entirely new Presentation Server x64 farm and migrate your users into it.

Refer to the *Presentation Server 4.0 Administrator's Guide* for more details at:

<http://support.citrix.com/article/CTX106319>

Important: Installing Service Pack 2005.04 on Presentation Server 3.0 effectively upgrades its binaries to Presentation Server 4.0 levels. This service pack is not considered a “feature release” or an “upgrade”. It will not activate any of the new features found in Presentation Server 4.0, such as AIE and CPU sharing. It is important to note, however, that the printing sub-system is not license dependent, and that by applying Service Pack 2005.04, the new printing architecture introduced in Presentation Server 4.0 will be used.

If you plan to install Service Pack 2005.04, be aware that some farm servers can stay at their existing version level while the database connection servers, farm metric servers, and data collectors should be upgraded to Presentation Server 4.0 prior to other servers. The complete order of operation is shown in 3.3.1, “In-place farm migration” on page 46.

Please see Citrix knowledge base article CTX106831, *Service Pack 2005.04 - FAQ* for more information:

<http://support.citrix.com/article/CTX106831>

Citrix recommends that a mixed environment should only be maintained for the duration of the upgrade. Some servers, such as data collectors, farm metric servers, and database connection servers, should be upgraded prior to other servers in the farm. In the next section, we will review these upgrade steps.

3.3 Migration strategies

In this section, we present three migration strategies for upgrading a MetaFrame XP FR3 or Presentation Server 3 farm to a Presentation Server 4.0 farm. While taking into consideration the limitations surrounding upgrading the Windows Server 2003 x64 Edition, the same migration strategies can be applied for migrating to Presentation Server x64. The migration strategies are;

- ▶ In-place farm migration
- ▶ Two farm aggregation
- ▶ Multi-farm load sharing

We discuss these in the following sections.

3.3.1 In-place farm migration

An in-place migration is best suited to a farm that is built on a solid foundation. There should be no performance issues or bottlenecks on your servers. The servers should be stable with all the latest and relevant service packs and hotfixes installed.

Performing an in-place upgrade simply entails installing Presentation Server 4.0 over the top of the current version of Presentation Server.

An in-place migration can only be performed on the Presentation Server component(s) of the server. Once again, Citrix does not support performing an in-place upgrade of the operating system of the server. Added to that, an upgrade from any version of the Windows operating system to Windows Server 2003 x64 Edition is not possible. Since Presentation Server x64 does not install on a 32-bit operating system, it is not possible to perform an in-place upgrade to Presentation Server x64.

In-place upgrade advantages:

- ▶ Low cost
- ▶ Simplified administration and troubleshooting
- ▶ Load balancing hardware not required

In-place upgrade disadvantages:

- ▶ No redundancy
- ▶ Limited rollback strategy
- ▶ Complexity of multiple license scenarios

Figure 3-1 is a representation of a Presentation Server 3 farm in its simplest form, and is a typical candidate for an in-place upgrade.




Figure 3-1 Presentation Server 3 farm

The following steps describe how to do an in-place migration from a farm consisting of MetaFrame XP FR3 or MetaFrame Presentation Server 3 servers to Presentation Server 4.0:

1. Install the Licensing Server.

If you are upgrading from MetaFrame XP Presentation Server, you will need to install a Licensing server. If you are upgrading from Presentation Server 3.0, you should already have a license server. Be sure to download and install the current licenses

2. Upgrade the Conferencing Manager Server from Version 2.0 to Version 4.0.
3. If you are running in Interoperability Mode, turn it off at this point.
4. If you are running the Enterprise version of Presentation Server, upgrade the primary and backup Farm Metric Servers before any other servers.

Resource Manager, a component of the enterprise version of Presentation Server products, uses the farm metric server(s) to interpret information collected from other servers. Farm metric servers running earlier versions of Presentation Server than other servers in the farm may cause inconsistencies.

5. Upgrade the Presentation Server Console, Access Suite Console, and Web Interface.
 To automatically upgrade the consoles and the Web Interface and preserve custom configuration settings, accept the default settings of the Install MetaFrame Presentation Server and its components option in Autorun. Accepting the default settings automatically upgrades the consoles and the Web Interface before upgrading the server and preserves custom configuration settings.
6. Upgrade the Zone Data Collectors and Database Connection Server (where applicable).

Figure 3-2 shows the in-place upgrade at mid-point. After upgrading the critical components of the farm infrastructure, you will have a mixed environment with Presentation Server 4.0 systems and systems with earlier versions. No changes need to be made to the Web Interface server and the application servers are upgraded in-place to Presentation Server 4.0.

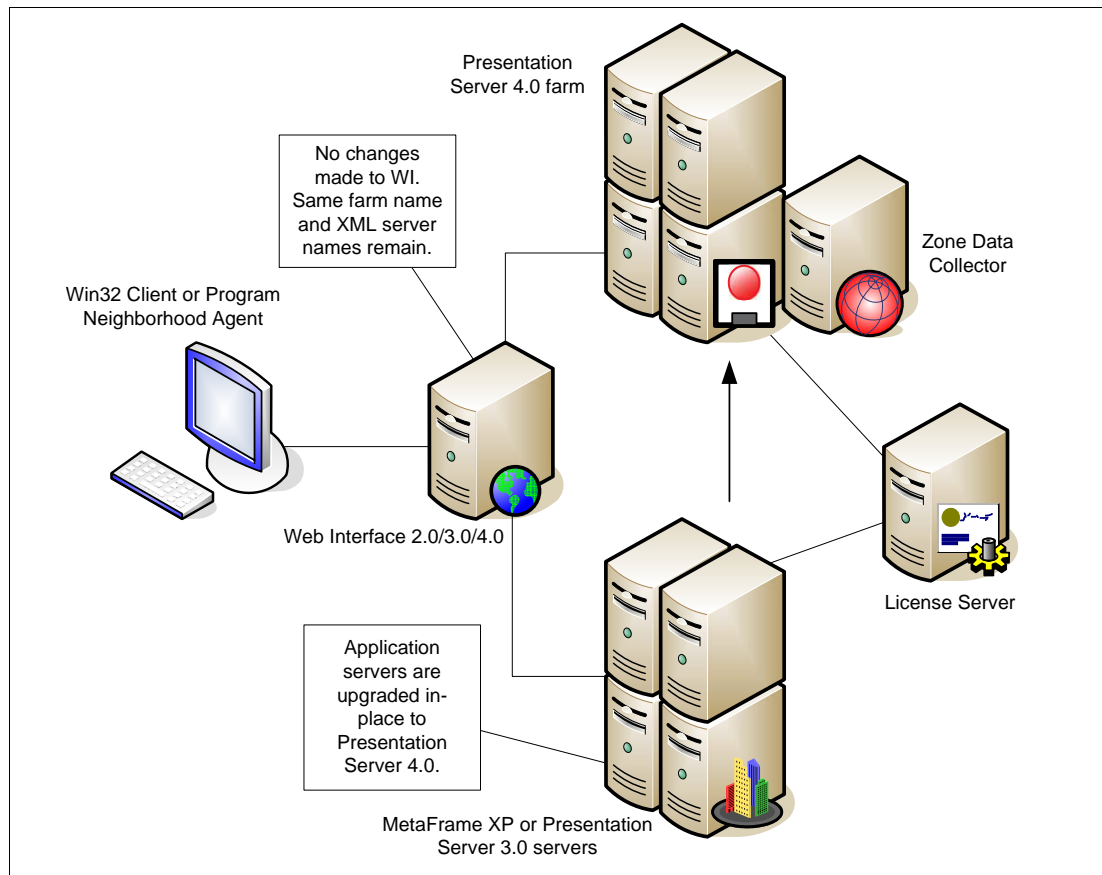


Figure 3-2 Mixed environment

7. Upgrade other farm servers.

Figure 3-3 shows the end result of an in-place upgrade to Presentation Server 4.0. If the foundation for the upgrade is a solid farm with no known issues, the end result will be a Presentation Server 4.0 farm that is stable, solid, and inherently more manageable and scalable.

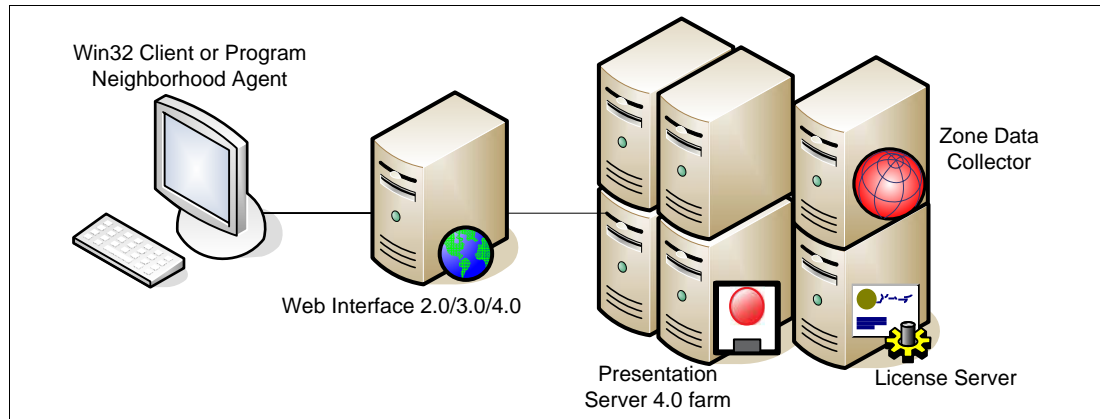


Figure 3-3 Presentation Server 4.0 farm

3.3.2 Two farm aggregation

A two farm aggregation upgrade is a farm installation that is performed side-by-side with an existing farm. While the two farms run in parallel, application aggregation is performed by Web Interface, MetaFrame Secure Access Manager, or Program Neighborhood Agent.

Two farm aggregation advantages:

- ▶ Simplified administration and troubleshooting
- ▶ Load balancing hardware not required

Two farm aggregation disadvantages

- ▶ Complexity of multiple license scenarios

The migration can be performed using the following process:

1. Install a new Presentation Server 4.0 farm.

This will include installing the first License Server if one does not already exist. It also includes creating and configuring Zones and Zone Data Collectors. A full discussion of the best practices for installing and configuring a Presentation Server 4.0 farm is out of the scope of this Redpaper.

Refer to the following for a complete discussion regarding best practices in farm design:

- *Presentation Server 4 Administrators Guide*
<http://support.citrix.com/article/CTX106319>
- *Presentation Server 4 Advanced Concepts Guide*
<http://support.citrix.com/article/CTX107059>

The Presentation Server 4.0 infrastructure runs alongside the MPS3 or MFXP FR3 infrastructure

Figure 3-4 shows the two farms running side-by-side with a single Web Interface server. The most important part of this migration scenario is the configuration of Web Interface. It should be emphasized that while the second farm is added to the Web Interface configuration, no load sharing should occur between the farms. Using this method, the applications are cut-over to the new farm one by one using the Presentation Server Console.

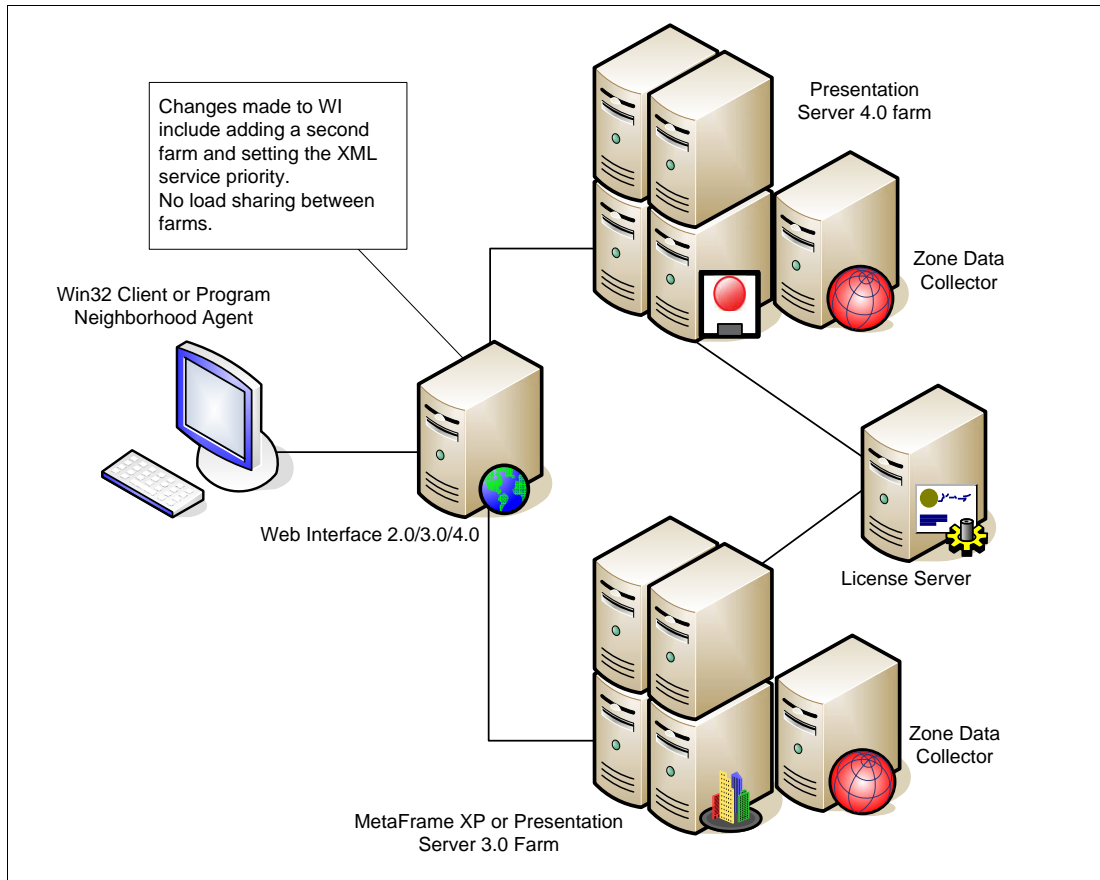


Figure 3-4 Two farm aggregation upgrade method

2. Add the new farm to the existing Web Interface server.

Figure 3-5 on page 51 and Figure 3-6 on page 51 shows the process of adding the new farm to the Web Interface server. In our case, we configured our Web Interface 4.0 server via the Access Suite Console. However, the same configuration options apply to any version of Web Interface. The existing farm name is DE AUS. The new farm is called DE AUS_X64.

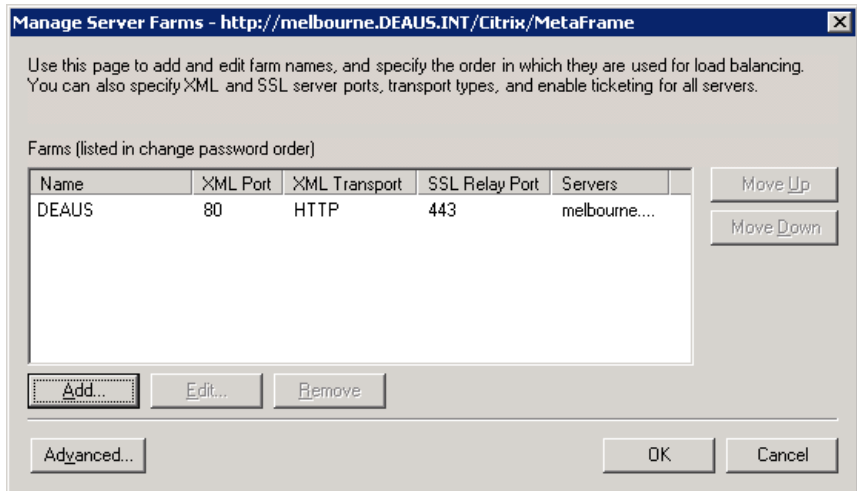


Figure 3-5 The existing farm - click Add to add a new farm

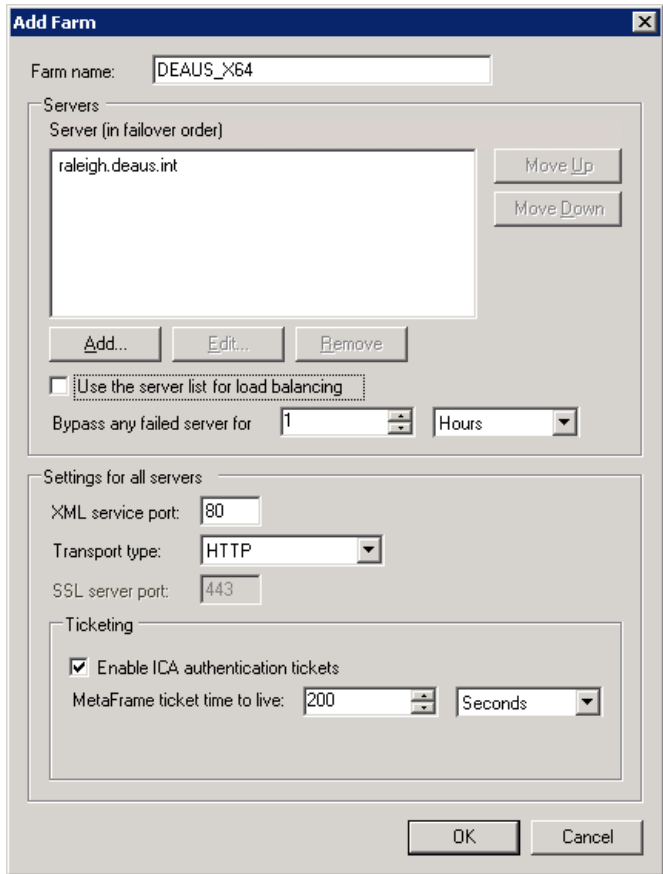


Figure 3-6 Adding the new farm to Web Interface using the Access Suite Console for Web Interface 4.0

3. Published applications can be gradually moved over to the new farm by changing their properties within the management console. As the published applications are migrated, be sure to remove access to the applications in the old farm.
4. Once all published applications have been migrated to the new farm, decommission the old farm and update Web Interface to remove any references to the old farm.

3.3.3 Multi-farm load sharing

A multi-farm load sharing migration strategy is more complex than other migration strategies because it requires the use of multiple Web Interface servers and a hardware load balancing device or Windows Network Load Balancing.

Multi-farm load sharing advantages:

- ▶ Load sharing
- ▶ Robustness
- ▶ Farm affinity
- ▶ Redundancy during migration

Multi-farm load sharing disadvantages:

- ▶ Load balancing equipment required
- ▶ Complex administration and troubleshooting
- ▶ Extra hardware required for redundant environment

The following steps show the high level tasks involved in a multi-farm load sharing migration strategy:

1. Follow the standard steps to install a new Presentation Server 4.0 farm.
 - a. Install License Server, licenses, and a new farm.
 - b. Install and publish business applications.
2. Deploy two or more Web Interface servers.
3. Configure each Web Interface server so that they both have access to each farm.
4. Deploy a load balancing device in front of the WI servers, ensuring session affinity is enabled.

Figure 3-7 gives an overview of this scenario.

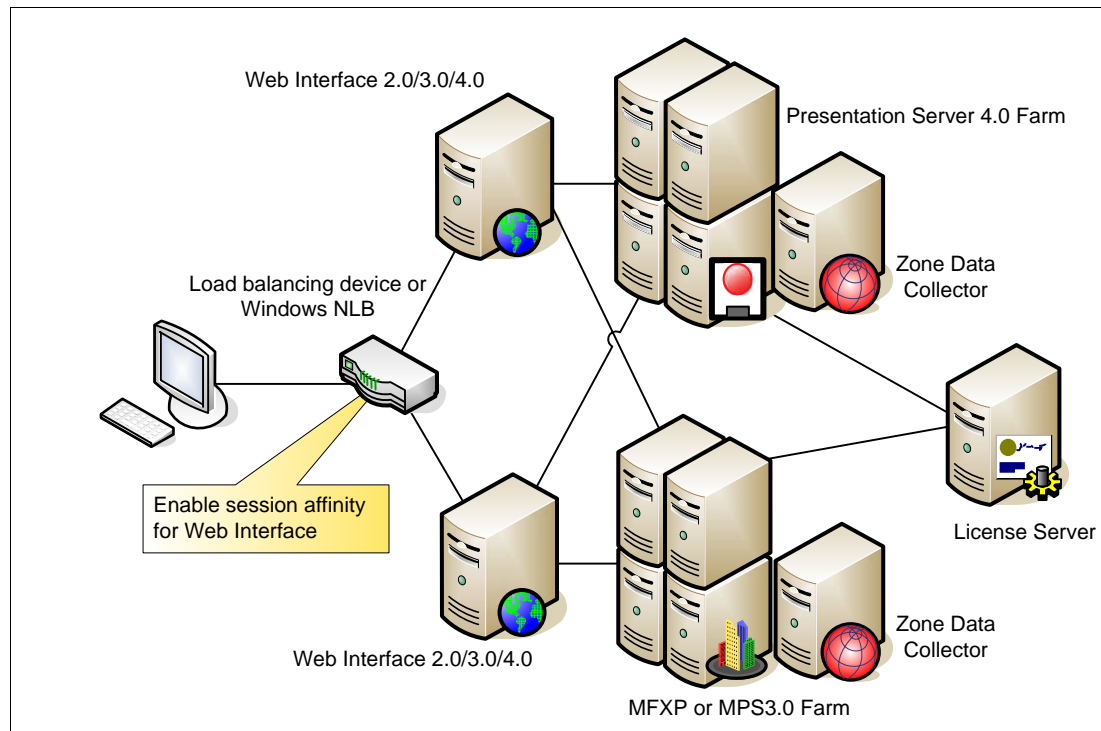


Figure 3-7 Multi-farm load sharing using load balancing device

5. Migrate or upgrade MFXP and MPS3.0 servers using the in-place migration method in 3.3.1, “In-place farm migration” on page 46. Bear in mind that it makes more sense to rebuild the servers entirely rather than simply install Presentation Server 4.0 over the top of what is already there. Ideally, you will end up with a pristine Presentation Server 4.0 environment.

3.4 Operating in a mixed 32/64-bit environment

It is important to note that many deployments of the Citrix Access Suite will include Presentation Server x64 and Presentation Server 4.0 servers. One of the most important things to consider when running in a mixed 32-bit and 64-bit environment is load management.

Presentation Server Advanced and Enterprise Editions include a sub-component called Load Manager. Using Load Manager, you can optimize server resources throughout your Presentation Server farm. Load Manager uses *load evaluators* to balance load across the farm by setting limits or thresholds based on certain evaluation criteria. Citrix provides two load evaluators by default that work well in most situations.

- ▶ The *Default* load evaluator

As its name suggests, this is the evaluator that is attached to every server in the farm by default. It contains one rule: when 100 users are logged on to the server, that server reports full load.

- ▶ The *Advanced* load evaluator

This evaluator is somewhat more advanced in that it contains rules based on CPU utilization, memory usage, and page swaps.

Load evaluators can be assigned to servers or published applications. Different load evaluators can be assigned to different servers and different published applications. A load evaluator can contain any mix of rules and thresholds. This makes load evaluators very flexible when it comes to calculating loads across platforms with differing specifications.

Note: The Citrix provided load evaluators cannot be deleted or changed. If you want to change rules or settings, you can copy the rules to a new load evaluator and then customize them as needed.

How load management works

Each sever in the farm maintains a table in the memory of its own performance metrics. This data is periodically reported back to a particular member of the farm called the Zone Data Collector (ZDC). The Load Manager subsystem uses the performance data from every server in the zone to formulate a load figure based on the rules of the load evaluator assigned to that server. The ZDC then distributes the load figure to other ZDCs in other zones of the farm. The Load Manager subsystem maintains an up to date record of the load of every server in the farm.

When a client launches a published application, it contacts the ZDC to obtain the address of a server hosting the published application. Load Manager selects the server that hosts the requested published application and then supplies its address to the client. The client then initiates a session on that server and launches the published application.

Example situation

Consider a scenario where the environment consists of a Presentation Server 4 farm that has a mixed hardware platform:

- ▶ Five HS20 Blade servers running Windows Server 2003 Standard Edition and Presentation Server 4.0
- ▶ One System x3850 (x366) running Windows Server 2003 Standard Edition x64 and Presentation Server x64

The HS20 blades have been sized to support 100 users each. The x366 with X3 Architecture, 4 Intel dual-core Xeon *Paxville* processors, and 64 GB RAM has been sized to support 400 users with the same workload. The Default load evaluator has been assigned to all servers.

With 100 users logged on to the blades, they have a CPU utilization of 60%, and memory usage of 80%. At this point, the blades are pretty well utilized. Maybe a few more users on each blade would not hurt, but you do not want to negatively affect the server responsiveness for all your users. At 100 users, Load Manager reports that the blade servers are full. No more connections will be routed to those blades.

With 100 users logged on to the x366, Load Manager reports a CPU utilization of only 15%, and memory usage of only 20%. The x366 is not fully utilized at all — it has been sized to take another 300 users — however, with the Default load evaluator assigned to the server, Load Manager reports that the server is fully loaded. No more user connections will be routed to this server. In this case, it would be advantageous to assign the Advanced load evaluator to the x366. The advanced load evaluator reports a full load when CPU utilization is greater than 90%, memory usage is greater than 90%, and page swaps are greater than 100 per second.

These figures have been provided to illustrate a point. A proper server sizing exercise should be conducted as part of the design phase of a Citrix implementation project.

3.5 Lab upgrade

Our lab for this Redpaper consisted of two xSeries 236 (8841-41U) servers and a System x3800 (8865-1RZ). The layout of the lab is shown in Figure 3-8 on page 55.




Figure 3-8 Lab configuration

The lab migration was performed as an in-place upgrade. The environment that was migrated from was a MetaFrame XP FR3 environment. The following steps were taken to migrate it to a Presentation Server 4.0 mixed 32-bit/64-bit environment.

1. A Licensing Server was installed on the domain controller “Rottweil”, and evaluation licenses for Presentation Server 4.0 were installed
2. On the MFXP FR3 server “Melbourne”, the Presentation Server 4.0 installation was started by letting the Presentation Server 4.0 CD autoplay and choosing **Product installations and updates** → **Install MetaFrame Presentation Server 4 and its components**.
3. Accept the license agreement.
4. Take note of the Prerequisites Installation window, particularly the note regarding Web Interface and ASP.NET, as shown in Figure 3-9 on page 56.

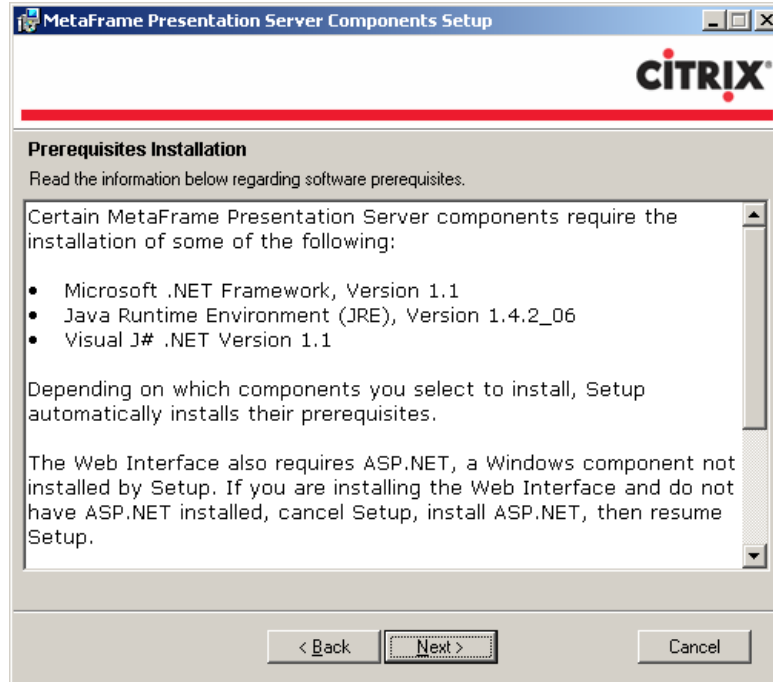


Figure 3-9 MetaFrame Presentation Server Components Setup - Prerequisites Installation

5. On the Component Selection window, leave the default choices. A license server was already installed in step 1 (see Figure 3-10).

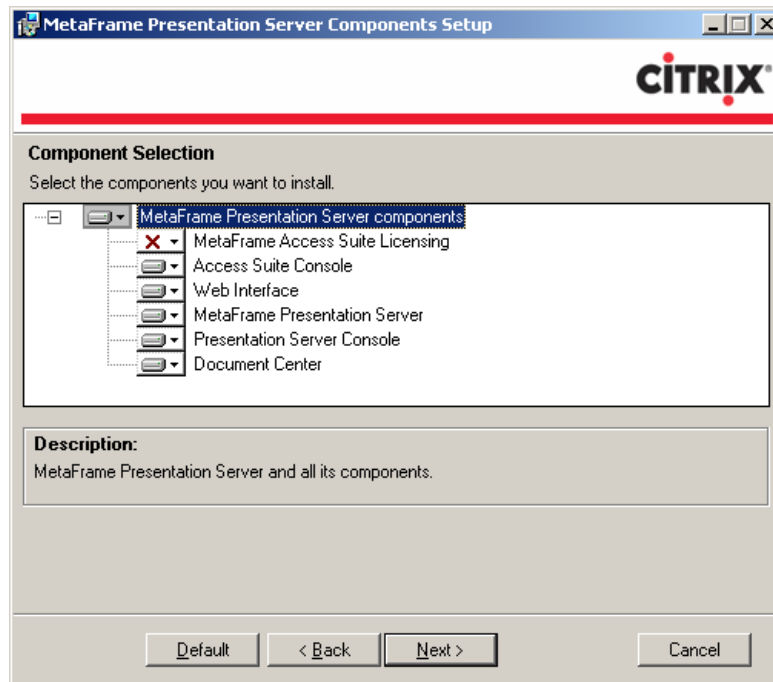


Figure 3-10 MetaFrame Presentation Server Components Setup - Component Selection

Note: Running the setup from the autoplay screen allows the Presentation Server 4.0 installer to launch the installation screens for all components that are to be installed. The installer will launch the installation screens for the administration snap-in, Visual J#.NET1.1, and Web Interface.

6. After installing the required sub-components of Presentation Server, the installer will display the warning shown in Figure 3-11.

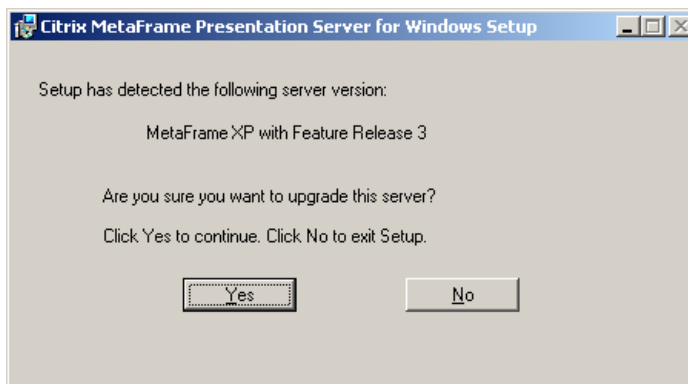


Figure 3-11 Upgrading to Presentation Server 4.0 - warning dialogue box

7. To continue the upgrade, click **Yes**.
8. Enter the name of the licensing server previously installed.

Once the Presentation Server has been upgraded, the Administration Console and Document Center components are installed.

Presentation Server x64 installs in exactly the same fashion as Presentation Server 4.0. In the test lab, Presentation Server x64 was installed on the System x3800 and brought into the DE AUS farm.

Note: 32-bit applications installed on Windows Server 2003 x64 are installed under C:\Program Files (x86). In a Presentation Server farm that includes 64-bit and 32-bit servers load balancing the same published application, the path to the executable file will possibly need to be manually configured to take into account the different paths.

Important: Although a direct upgrade from MetaFrame XP FR3 is possible and is supported, Citrix highly recommends that production-level servers are rebuilt from scratch using a reproducible, scripted installation, including all service packs, hotfixes, and published applications.

9. With the farm functioning properly under Presentation Server 4.0, the role of Data Store server was moved to the System x3800. The x236 was then rebuilt with a fresh install of Windows Server 2003 and Presentation Server 4.0.



Best practices

This chapter provides information about how to get the best performance and availability from a Citrix Presentation Server 4.0 environment. Furthermore, we will discuss some techniques on the best ways to secure Presentation Server and what is recommended in different printing environments. First, we have a brief look at a recommended server layout.

This chapter covers the following topics:

- ▶ 4.1, “Server layout” on page 60
- ▶ 4.2, “Sizing for scalability, high availability, and fault tolerance” on page 61
- ▶ 4.3, “Configuring Load Manager” on page 67
- ▶ 4.4, “Restricting user access” on page 71
- ▶ 4.5, “Printing in a Presentation Server 4.0 environment” on page 80
- ▶ 4.6, “Performance monitoring with key performance counters” on page 84
- ▶ 4.7, “Performance tuning” on page 88
- ▶ 4.8, “Adding and deploying new Presentation Server systems” on page 92

4.1 Server layout

In this section, we discuss considerations for the server layout from the view of the hardware and the software. This section is not a complete sizing topic, but it contains some helpful suggestions for choosing the hardware. We also discuss some considerations about the software layout of a Presentation Server system.

4.1.1 Hardware layout

Before implementing Presentation Server, we have to consider the hardware we want to use. The best hardware architecture to use is discussed in 1.8, “Selecting a Presentation Server x64 platform” on page 12. Furthermore, we have to discuss the layout of the selected hardware.

It is a best practice to run a Presentation Server environment on multiple dedicated servers; however, in branch offices or small-medium businesses, it may be more appropriate to use a single high-capacity (CPU, memory, or disk) server like the IBM System x3800 and run different operating systems through virtualization. This is discussed in 1.9, “Virtualization using the System x3800” on page 18.

When determining the hardware layout for a Presentation Server solution, consider the following aspects to the configuration:

- ▶ Hard drive RAID level

RAID 1 is using fully redundant mirroring for the disks, so we have a complete copy of one drive on the other drive. It provides high fault tolerance and can improve read performance. However, RAID 1 write performance is lower compared to RAID 5. Furthermore, this type of redundancy needs twice the disk space.

RAID 5 is a striped disk array. All data written to disk is divided into several data blocks and distributed on the disks. Also, parity data is written so fault tolerance is included. The writing data performance on a RAID 5 array is better than in RAID 1. Reading data is similar to RAID 1. For a RAID 5 array, a minimum of three disks is required.

We recommend using a two-drive RAID 1 array. The biggest benefit of this array is that all the data, including the operating system, remains on both disks. So it is possible to keep one copy of the complete system before something is installed only through unplugging one of the disks. For this task, the disks should be hot pluggable, so there is no need to power down the system and the RAID 1 array must only consist of two disks. The degraded write performance we mentioned has nearly no impact on the performance of the system because most data accesses on a Terminal Server are read accesses. Most writings to the disk happen when the operating system is writing to the page file.

- ▶ Processor, bus, and memory

The layout for the processor, bus, and memory highly depends on the number of users, the applications, and the type of users that are connected to the server. Because Presentation Server is memory and I/O intensive, we recommend a minimum of two processors and high scalable hardware. We also recommend calculating a minimum of 25 MB RAM for each user so that a good performance for typical and power users is ensured. Adding more memory has often a larger effect on the system than adding faster processors. In addition, the minimum requirements for the operating system must be achieved. In a 64-bit environment, more memory is needed for the operating system and the applications. This is because the pointers (and things that act like pointers, such as handles and size of variables) double in size. Thus, the amount of memory needed in fact does not double but it can increase.

For more information about performance considerations for xSeries server, refer to the IBM Redbook *Tuning IBM eServer xSeries Servers for Performance*, SG24-5287.

► Software considerations

Before installing Presentation Server 4.0, Terminal Services has to be installed and active.

As a best practice, we recommend not installing any other service on the server running Presentation Server. Presentation Server is an application server that is used to provide applications to users from any place with any client. If you install other services, the system may become slow and unstable. No data should be saved on a Terminal Server.

Furthermore, with Citrix Presentation Server 4.0, it is no longer possible to install the server as a Active Directory domain controller. So the server should be a standard member of a Active Directory domain. We recommend that all Citrix Presentation Servers in a farm belong to the same Active Directory forest. Citrix Presentation Server does not support UPD logons if a server farm spans multiple Active Directory forests.

If your requirements are to minimize the number of physical servers, consider implementing virtualization software such as VMware Server (formerly GSX Server) or VMware ESX Server on a server with sufficient processing, memory, and storage capacity, such as the System x3800.

Before implementing Citrix Presentation Server 4.0 in a production environment, we always recommend performing a pilot installation. For a pilot installation, we recommend choosing different user groups to receive the best result of the tests. Then performance should be measured with the key performance counters shown in 4.6, “Performance monitoring with key performance counters” on page 84. With this method, it is possible to make the decision for the best hardware layout.

For more information about sizing the hardware and designing a Citrix Presentation Server 4.0 server farm, refer to:

- *MetaFrame Presentation Server Administrator's Guide*, found at:

<http://support.citrix.com/article/CTX106319>

- *Advanced Concepts Guide Citrix Presentation Server for Windows 4.0*, found at:

<http://support.citrix.com/article/CTX107059>

- Redpaper *Implementing Windows Terminal Server and Citrix MetaFrame on IBM eServer xSeries Servers*, REDP-3629, found at:

<http://www.redbooks.ibm.com/abstracts/redp3629.html>

4.2 Sizing for scalability, high availability, and fault tolerance

A Citrix Presentation Server environment must always be up and running because sometimes thousands of users depend on it. So, we always have to consider how to make such an environment more available and fault tolerant. We have to include the scalability of the systems so you can respond fast on new business requirements. This is discussed in the following sections.

4.2.1 Scalability

When we think of scalability, we have to examine two ways of scalability. These are scale up or scale out.

- ▶ Scale up

Increasing the capacity of a single system by adding processors, memory, and disk. The advantage of scaling up is that it is relatively easy, because it only requires hardware and software that can take advantage of adding more hardware.

- ▶ Scale out

Adding systems that can be managed and run together. With scaling out, discrete servers are added to the environment to gain more process power. Load balancing and clustering are the main goals of scaling out. Due to these goals, scaling out could be more complex than scaling up.

In the past, scale out was mostly recommended for Presentation Server implementations. Today, as times change, with the new features of Presentation Server 4.0 and the possibility to run it on a Windows Server 2003 x64, we recommend striking a balance between scale up and scale out.

So, when using an IBM X3 Architecture server, as described in 1.8, “Selecting a Presentation Server x64 platform” on page 12, this is the happy medium for scaling up and scaling out, as shown in Figure 4-1. With scaling up the server, the user density on one server can be increased. With scaling out in a Presentation Server environment, we gain the advantage of higher availability and fault tolerance, as described in 4.2.2, “High availability and fault tolerance” on page 63.

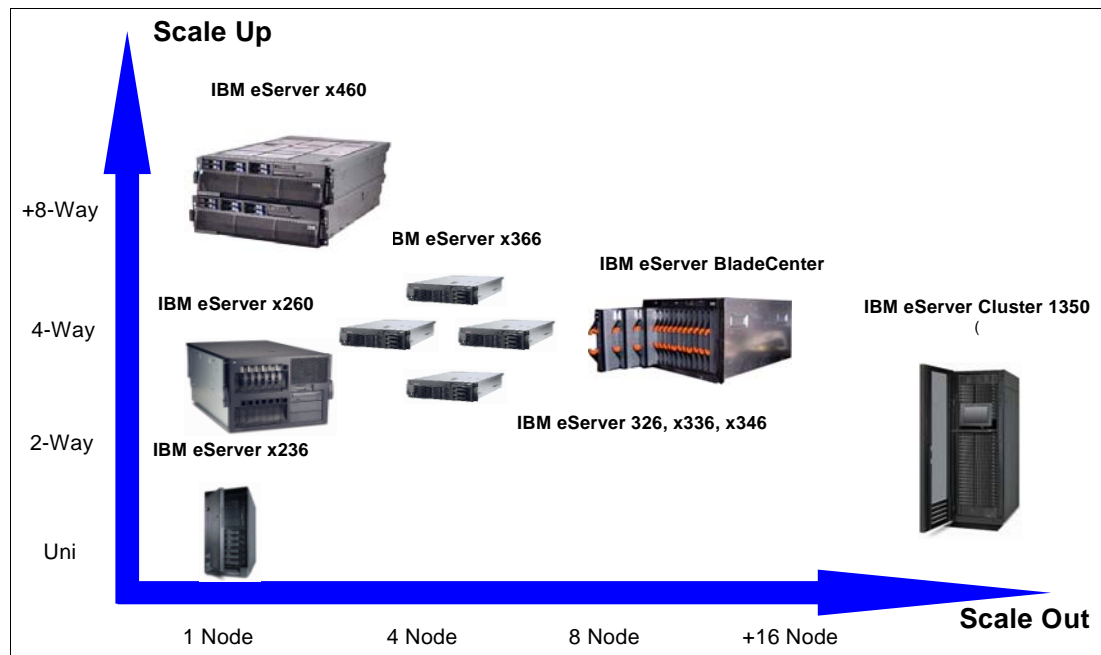


Figure 4-1 Scale Up versus Scale Out

The advantages of scaling up in a Presentation Server environment are:

- ▶ Fewer licenses required
- ▶ More economies of scale
- ▶ Fewer servers to manage
- ▶ Better scalability, especially with the x64 Edition

The disadvantages are:

- ▶ Single or fewer points of failure.
- ▶ Multiple applications have to be installed on the same server, so the testing of the collaboration of the applications will increase (recommend having a test server). This disadvantage can be eliminated with the use of virtualization with VMware or the Application Isolation Environment feature of Citrix Presentation Server 4.0 Enterprise Edition.

4.2.2 High availability and fault tolerance

To implement high availability and fault tolerance, two major aspects have to be considered: the hardware and the software. In this section, we discuss how to increase availability and fault tolerance on the application view of the Presentation Server. For more information regarding other services, such as the Citrix License Server, Citrix Data Store, Citrix Metric Server, and so on, refer to the Citrix white paper *Citrix Access Suite 4.0 - Disaster Recovery Planning*, found at:

http://www.citrix.com/English/ps2/products/documents_onecat.asp?contentid=186&cid=White+Papers

First of all, the server hardware itself has to be fault tolerant. This means we always use redundant power supplies and fans and redundant and hot swappable hard disks configured in an array. Furthermore, we recommend using a minimum of two network cards and teaming them. Also, memory needs to be examined and it would be best to use the memory fault tolerance feature of the X3 Architecture servers.

All of these hardware requirements are met with the IBM System x high-end servers. Furthermore, when sizing the hardware, it should be more than one server, so that there is no single point of failure. For example, if your server sizing says that you will need five servers to connect all your users, add about 25% of the needed resources, so you would install six servers to have improved fault tolerance.

Note: The opinion of what is high availability and fault tolerance may differ from administrator to administrator. It mostly depends on the following questions:

- ▶ What is the longest acceptable downtime a user can experience?
- ▶ How often are changes made to the environment?
- ▶ During changes, how much time can you afford to have the server offline?

When installing the Presentation Server farm, consider installing the servers on different sites, so that the users can be routed to another site when one site is complete down because of disasters like fire, flood, and power loss.

The second aspect is the software. Presentation Server 4.0 does not support Microsoft clustering. This is because Terminal Services does not run in application server mode when running on a cluster node. Only remote administration mode is supported. Only the Citrix Licensing Server and the data store can be clustered because they do not need to run on a Terminal Server in application server mode.

For more information about Clustering with Citrix components, refer to:

- ▶ Citrix knowledge base article *Microsoft Server Clustering and Terminal Services Limitations*, found at:

<http://support.citrix.com/article/CTX999186>

- ▶ Microsoft Knowledge base Article Q294926:
<http://support.microsoft.com/support/kb/Q294/9/26.asp>

So, for Presentation Server, we need other ways to come up to the requirement of high availability and fault tolerance. There are two main methods to accomplish this:

- ▶ Citrix Presentation Server Load Manager
- ▶ Virtualization with VMware

These two methods are discussed in the following sections.

Load balancing with Load Manager

Load Manager is for Citrix Presentation Server 4.0 (Advanced and Enterprise Edition).

With Load Manager it is possible to publish applications on multiple servers. When a user starts an application, Load Manager is responsible for determining from which server the application will start, depending on the load rules you have defined. Load Manager enables you to:

- ▶ Optimize resource usage in a server farm. When starting an application, the users ICA connection is routed to the server with the lightest load in the farm, so the efficiency of the servers in the server farm is maximized.
- ▶ Set the date and time for load management. You can schedule the availability of a published application based on date and time.
- ▶ Monitor application loads across server farms. You can chart the loads on servers that host published applications.
- ▶ Check disk data. Load Manager can check processes like disk operations on servers that host published applications.
- ▶ Span subnets. Load Manager can be used in networks that contain multiple server farms within a single network or multiple subnets.

The rules you define for Load Manager can be based on published applications or servers. You can apply thresholds for the following components for Load Manager:

- ▶ Content switches: A range of context switches per second for a selected server. A context switch occurs when the operating system switches from one process to another.
- ▶ Application user load: Limits the number of users allowed to connect to a selected published application
- ▶ Server user load: Limits the number of users allowed on a selected server.
- ▶ Disk operations: A range of disk operation, in read/write cycles per second, for a selected server.
- ▶ IP address range: A range of allowed or denied Client IP addresses for a published application. This rule can only be used in conjunction with another rule.
- ▶ Disk data I/O: A range of data throughput, in kilobytes, for a selected server.
- ▶ Scheduling: Schedules the availability of selected published applications.
- ▶ Page swaps: A range of page swaps per second for a selected server. A page swap occurs when the operating system moves data between the physical memory and the disk.
- ▶ Page faults: A range of page faults per second for a selected server. A page fault occurs when the operating system tries to access data that has been moved from physical memory to disk.

- ▶ CPU Utilization: A range of processor utilization, as a percentage, for a selected server.
- ▶ Memory Usage: A range of memory usage for a selected server.

With the Load Manager Monitor, you can display statistics monitored by the load evaluators after they are attached to a server or a published application, so it is always possible to get real-time statistics of the actual load. It is also possible to use Load Manager in a mixed environment with MetaFrame 1.8 and MetaFrame Presentation Server 3.0 or earlier when the farm is running in interoperability mode. However, load evaluators that you create with Load Manager for Citrix Presentation Server cannot be used with MetaFrame 1.8 servers.

There are two built-in load evaluators that come with Citrix Presentation Server:

- ▶ Default: This load evaluator is attached to every server by default. It contains only one rule, the server user load, that reports a full load when 100 users log on to the server.
- ▶ Advanced: This load evaluator contains the CPU utilization, memory usage, and page swap rules. This load evaluator is not attached to any server by default.

The built-in load evaluators cannot be deleted or modified, but the default load evaluator can be detached from the servers. To define your own load evaluator, copy one of the built-in ones and change the rules and settings. Figure 4-2 shows how to configure a new load evaluator.

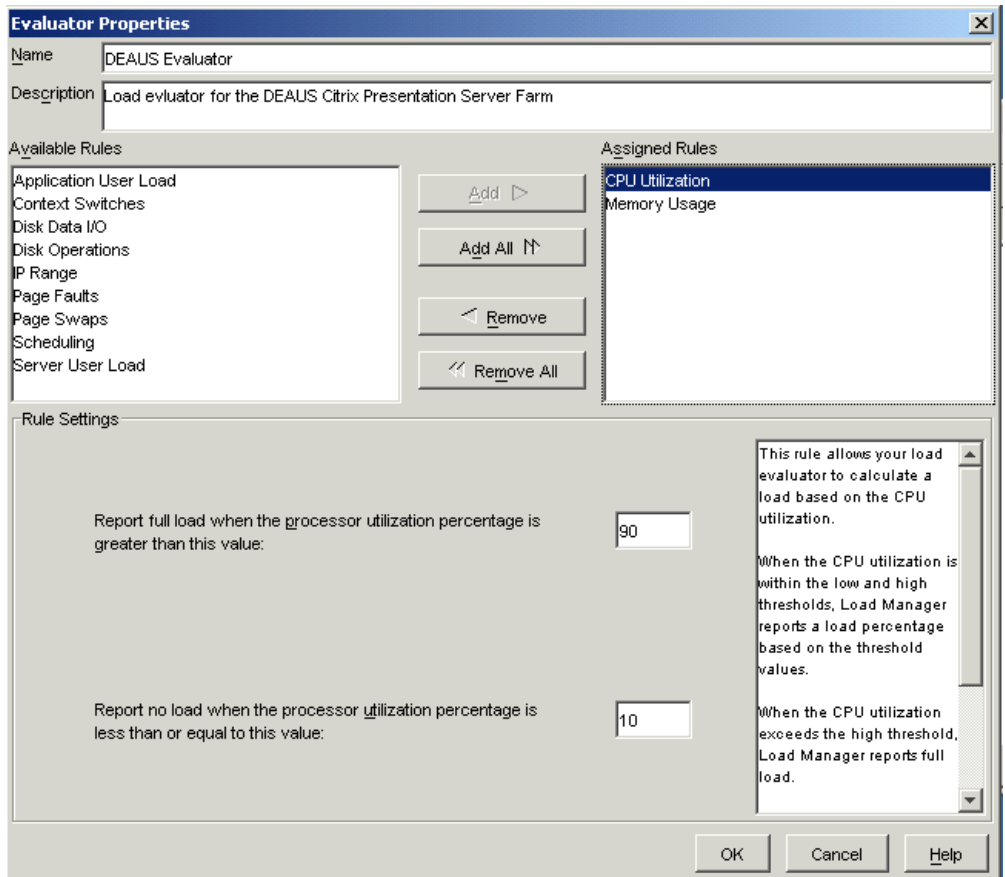


Figure 4-2 Configuring a new load evaluator

Before using Load Manager, consider the following:

- ▶ A published application that is unique for only a small user group does not need to be load managed unless the application has to be highly available.
- ▶ Published applications that require significant resources from a server can be configured to report full loads at a lower threshold than the actual server limit.
- ▶ Do not only configure the user load rule on a server because the overall load not only depends on the number of users on the server; it depends more on the applications the users start and how the users work.
- ▶ To get the best results with Load Manager, measure the performance of an application that you want to load manage. With the results of this measurement, it is best to define the rules for the load evaluator and to say when the server should report full load.

Note: With Load Manager, it is not possible to migrate a user automatically or manually from one server to another if one server is going down. Users that are logged on to the server that is going down have to log in again and are then routed to another server.

For more information about configuring an using Load Manager, refer to Citrix Knowledge base Article CTX106452, *Load Manager Administrator's Guide*, found at:

<http://support.citrix.com/article/CTX106452>

Virtualization with VMware

Over the last few years, the use of virtual machines on Intel based servers has increased in popularity through the use of VMware's virtual machine products. VMware ESX Server is the most common product when consolidating servers in an Intel based environment through virtualization.

One of the key features of VMware's virtualization products is running more than one operating system on one server, and therefore using more of the server's capacity. VMware products, such as Virtual Center and VMotion, can also be considered for additional availability and fault tolerance in a Citrix Presentation Server environment. Presentation Server is fully supported for running on a guest instance of Microsoft Windows Servers hosted by VMware ESX Server and VMware Server (formerly GSX Server).

For details of support, see the Citrix knowledge base entry *Support for MetaFrame with VMware ESX Server or Microsoft Virtual Server*, found at:

<http://support.citrix.com/article/CTX997956>

When implementing a Presentation Server environment as guest instances of two or more VMware ESX Server systems, it is possible to gain high availability and fault tolerance on a higher level than using only Load Manager. With the additional VMware products VirtualCenter and VMotion, you can gain the advantage of zero downtime in your Citrix environment even in the failure of a complete server hardware. Such a scenario is shown in Figure 4-3 on page 67.

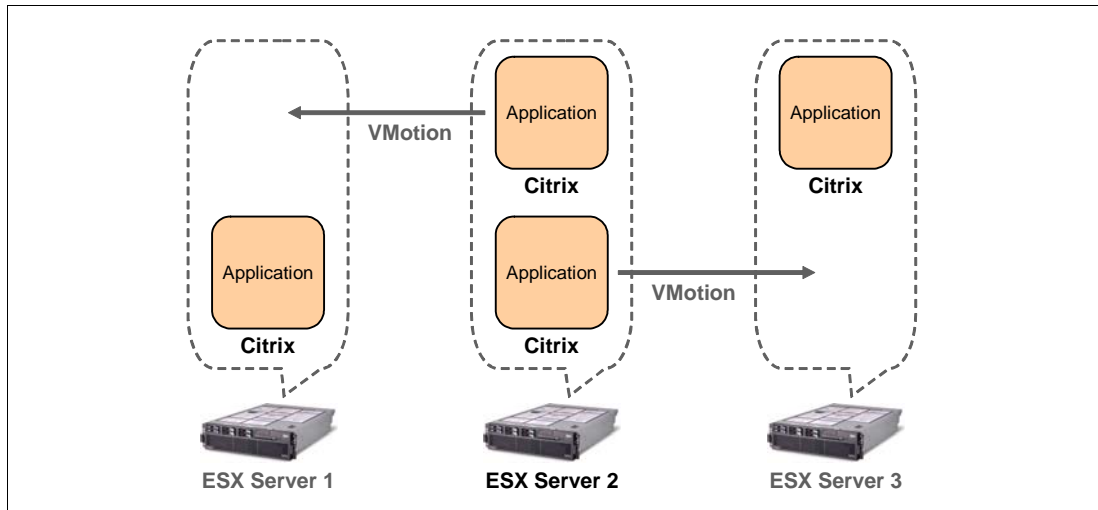


Figure 4-3 Using VMware ESX Server with VMotion

You can manually move the two Presentation Server instances on Server 2 to the other servers so that Server 2 can be powered down for maintenance. Also, it is possible to move the instances automatically, depending on the defined thresholds.

In addition, a new installation of Presentation Server instances can be deployed through the different VMware tools and techniques.

For more information about implementing and using VMware, refer to the IBM Redpaper *Server Consolidation with VMware ESX Server*, REDP-3939.

Attention: You still need a disaster recovery plan where you define how to recover from any possible disaster. Also, keep in mind that redundancy and disaster recovery plans for all other network infrastructure services should be available, but these are not in the scope of this Redpaper.

In conclusion, there are many possibilities to gain higher availability and fault tolerance. In our opinion, a mixture of using VMware ESX Server together with VMotion and the Load Balancing feature in Citrix Presentation Server could be a very reliable solution. How to configure higher availability with the Citrix Load Manager is discussed in the next section.

4.3 Configuring Load Manager

In this section, we describe how to configure high availability and fault tolerance for a Presentation Server environment by configuring the Citrix Load Manager. We show how to publish applications and what are the requirements for publishing and managing applications.

To use Load Manager for application fault tolerance, you need a minimum of two Presentation Server installations. These can be two separate servers, but it can also be two virtual machines with Presentation Server 4.0 installed. Furthermore, the Presentation Server package has to be at least the Advanced Edition. Then you have to publish the applications you want to load manage because only published applications can be load managed. Similarly, if you want to load manage the desktop, you must also publish the desktop.

4.3.1 Publishing applications

First, consider which servers you want to load manage and what applications you want to publish. The servers and applications depend on how critical the performance for the servers and applications is that are installed in your server farm. Before publishing the applications, consider the following:

- ▶ How and where are the applications installed in the server file system?
- ▶ How will the applications be published?
- ▶ Which user gets which application?

To publish an application on more than one server, the application has to be installed in the same directory structure for all servers. In addition, the same features for the application should be installed on all servers and they should be configured the same way. This is because when publishing an application, the properties for the application are configured one time for all servers. This is shown in Figure 4-4.

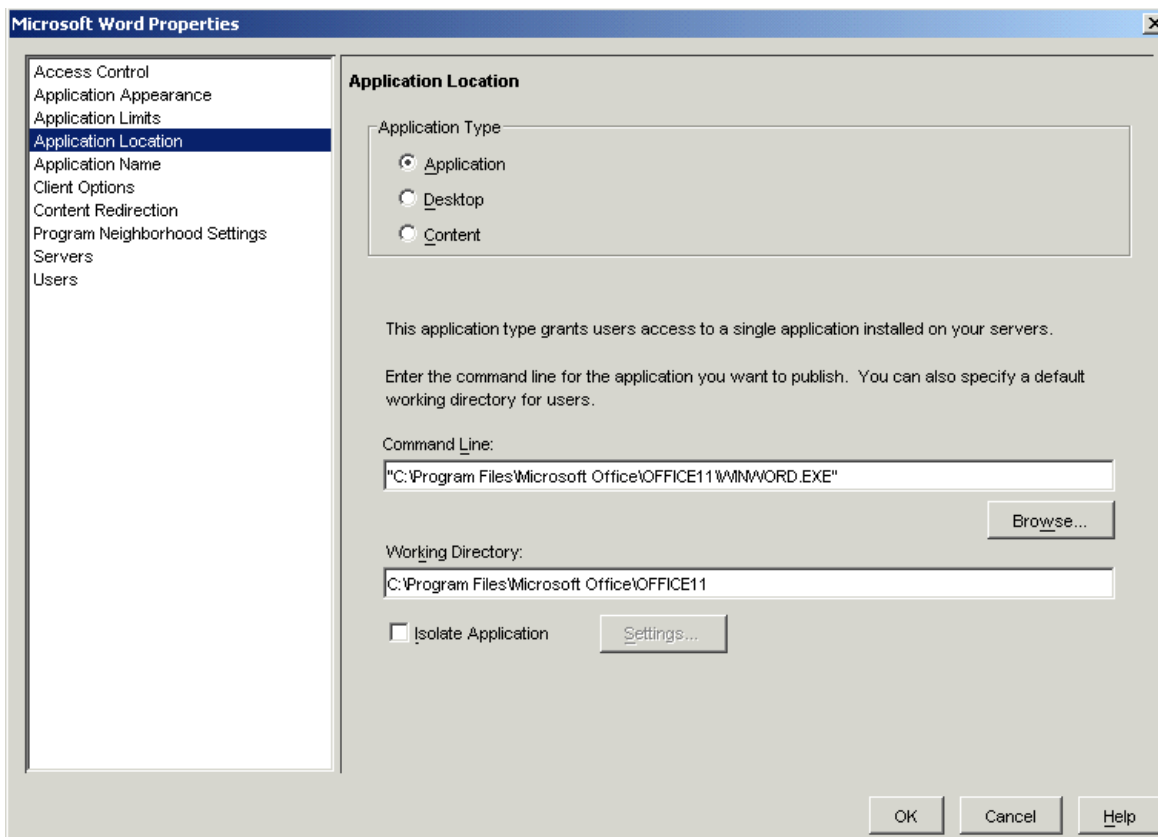


Figure 4-4 Published application starting properties

In the example of Figure 4-4, you see that the path for the application is hardcoded, as is the working directory. So we need to install the applications being published all the same way on the different servers. If you are publishing the desktop rather than individual applications, you do not need to worry about keeping the application settings the same across servers.

While publishing the application, we also define how the application is published. This includes the following:

- ▶ Program Neighborhood settings

There you define the Program Neighborhood directory for the user where the application will appear if you use Program Neighborhood. You also define whether an icon should be on the user's desktop or if it should be integrated into the user's Start menu as well.

- ▶ Application appearance

Here you configure the window size and color depth for the published application. You can also define if the applications task bar is hidden and if the application starts maximized.

- ▶ Client options

Here existing audio for the application is defined as well as the encryption type.

- ▶ Content redirection

Content redirection means that when a user opens a document from his local workstation and the appropriate application is not installed on the local workstation, the file is redirected to the corresponding published application and is opened through the published application on the Presentation Server.

The use of content redirection is not possible with Program Neighborhood; you must use the Program Neighborhood Agent instead.

After publishing the application, establish the servers where the application is published, as shown in Figure 4-5.

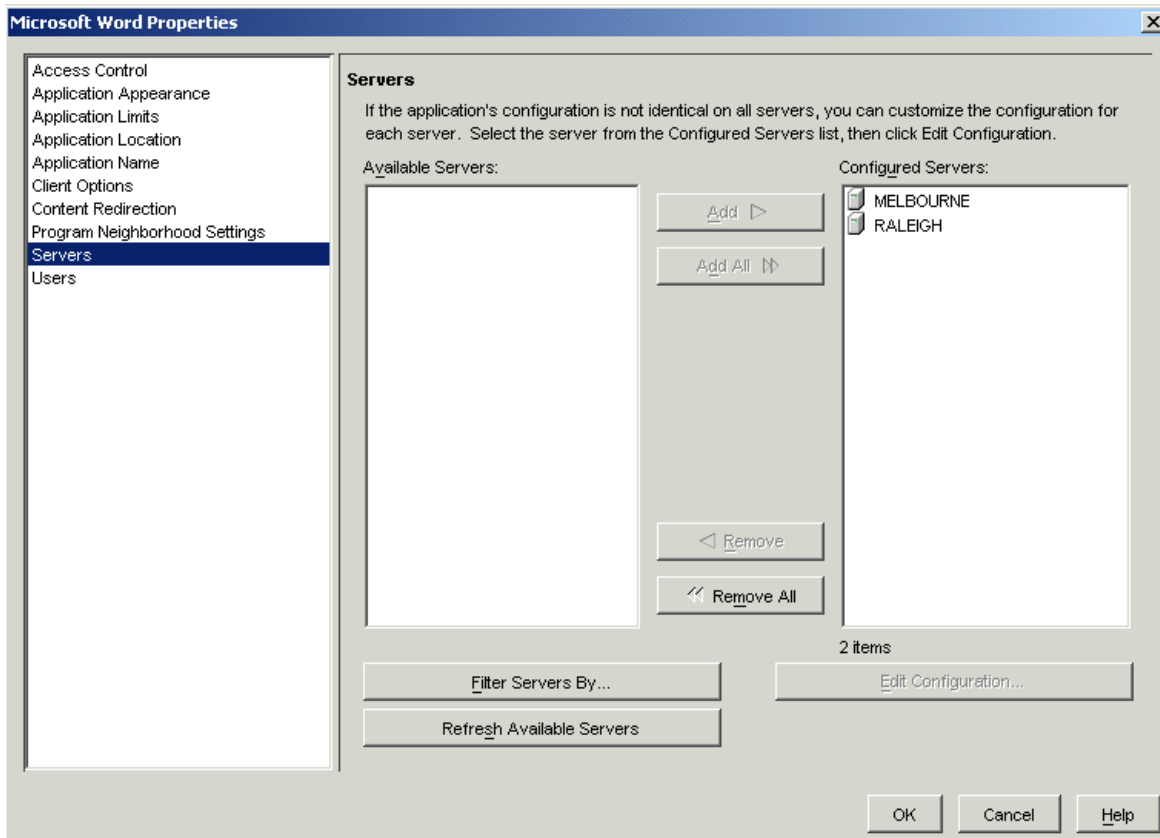


Figure 4-5 Connecting published application to the servers

Define the users or user groups that will be assigned to the published application. We recommend using groups for the applications so that any users added to the group at a later date automatically get access to the necessary applications. It is easier to manage the published applications through Active Directory. This is shown in Figure 4-6.

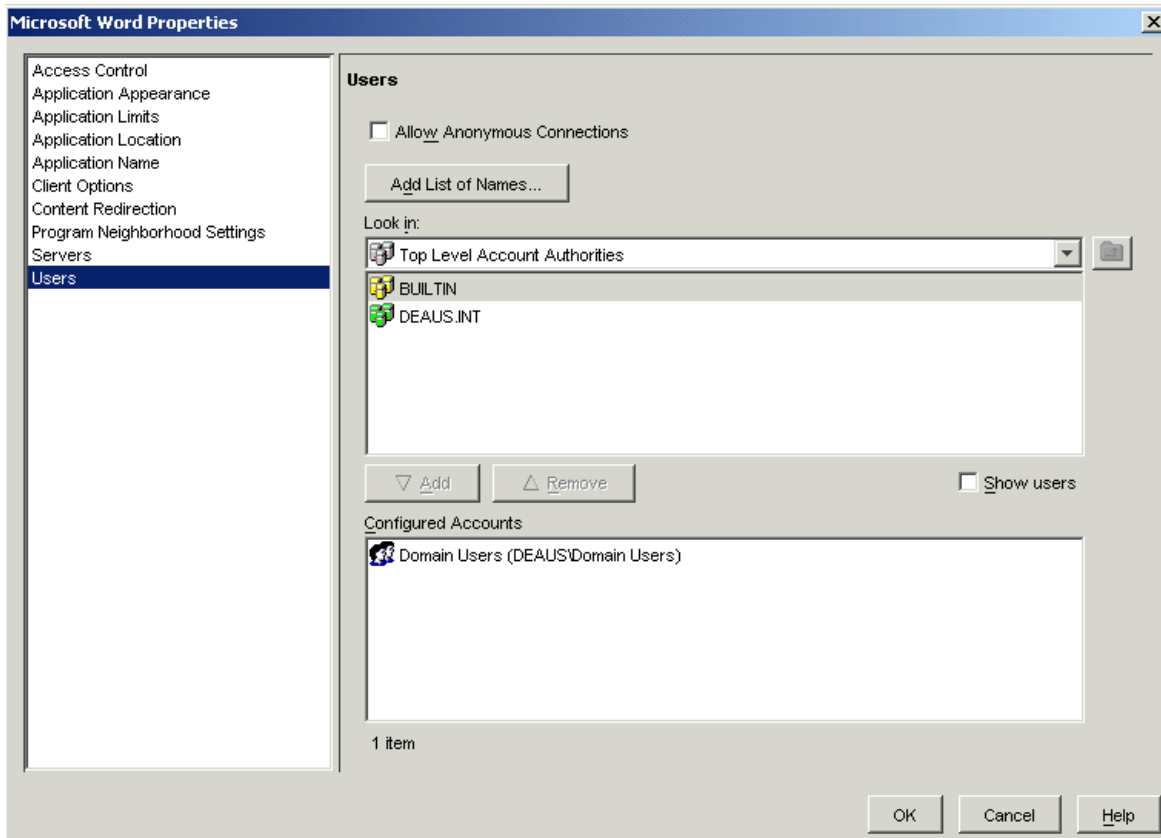


Figure 4-6 Defining the users for the published application

All users and groups that are defined for the application get the icon for it at the next refresh or restart of the Citrix Neighborhood Programs or the Citrix Neighborhood Agent.

Now define the load evaluator for the servers or the published application, as shown in Figure 4-2 on page 65. The metrics and values that are defined for the load evaluator depend on your experiences with the application. The load evaluator can always be changed without downtime, so we recommend watching the arrangement of the applications and the performance of the servers to get the most efficient utilization of Presentation Server.

Figure 4-7 on page 71 shows how to bind the defined load evaluator to a published application. Binding a load evaluator to a server independent of the application works the same way.

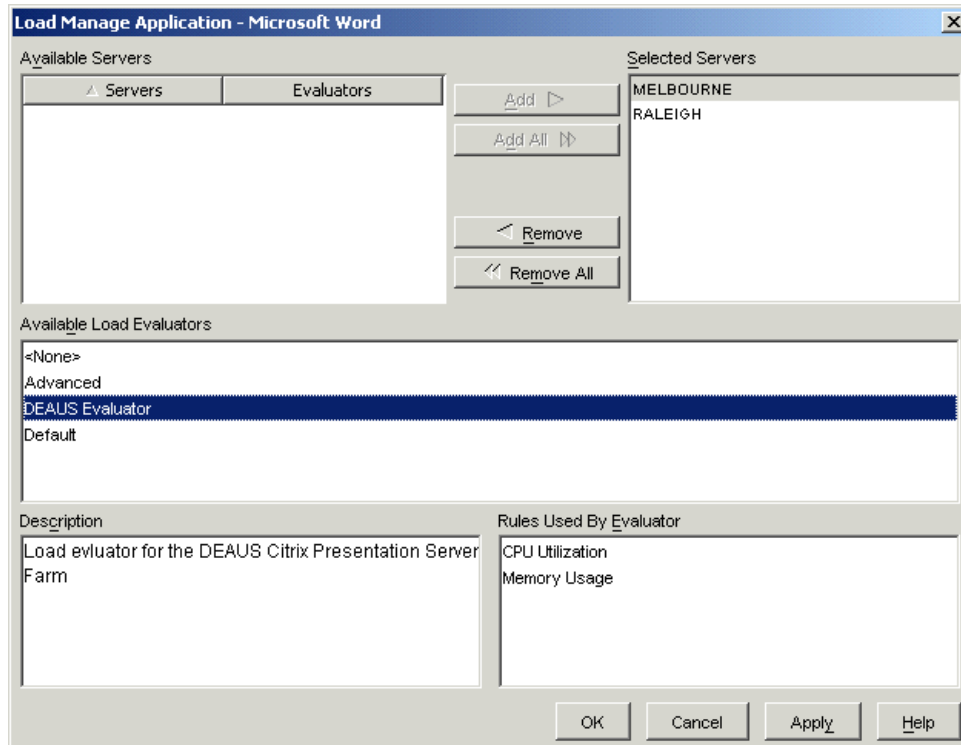


Figure 4-7 Bind load evaluator to a published application

When the user starts the application, Load Manager checks all the servers and routes the user to the server that is up and has the lightest load. So, if one of the servers goes down or has to be updated, the users can still log on to their applications by routing them to another server.

4.4 Restricting user access

In a Presentation Server environment, it is important to implement high security and to restrict user access (also called “locking down the users”). Restricting user access in this sense means that you restrict the user’s environment so the desktop or the application configuration is fully managed by the administrator and the user cannot do anything else than he is allowed to do.

When a user logs on to Presentation Server, he gets only the display, mouse, and keyboard information and he is logged on locally to the server. So the user could also see local drives of the server and could delete files or change system configurations. In this section, we will discuss how to lock down the users. There are also many tools on the Web to reach this target, but in this Redpaper, we discuss the standard possibilities without third-party software. In this Redpaper, we do also not discuss any further security features like encryption, certificates, and so on.

The following are aspects of security that you should review:

- ▶ NTFS and registry security
- ▶ User profiles
- ▶ User policies (Windows and Citrix)

These bullets are discussed in the following sections.

4.4.1 NTFS and registry security

You should implement the NTFS file system for all drives as well as secure the registry so that users logging on to the server may neither delete operating system files, the Program Files directory, or edit the registry.

When installing Windows Terminal Services, you are asked to select the default permissions for application compatibility, as shown in Figure 4-8. Using Full Security will prevent standard users from making changes to the registry.

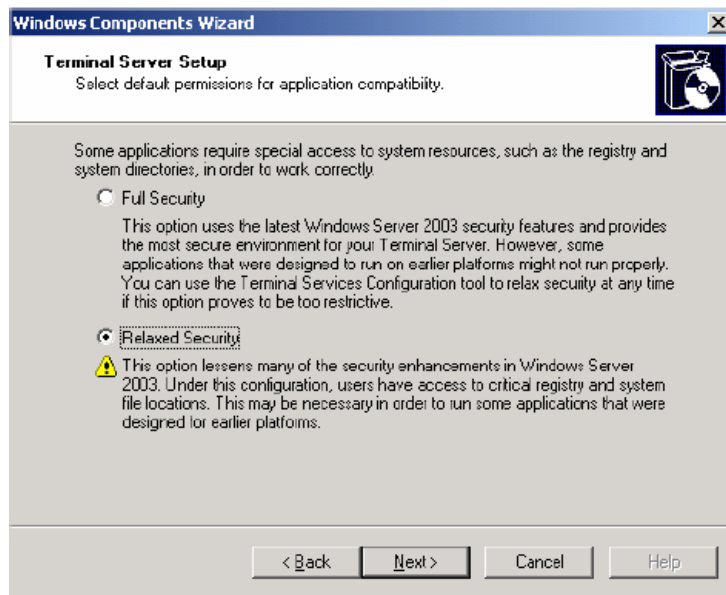


Figure 4-8 Selecting Terminal Server default permissions for application compatibility

We recommend using Full Security. You can change this setting at any time through the Windows Terminal Services Configuration Tool.

All users should be configured as standard domain users. Do not give the users Administrator or Power User permissions.

Windows Server 2003 installs by default with the system drive sufficiently secure, with the Everyone group not having Full Control to the system drive. The authenticated user group has Read and Execute access to the system drive so they can use the files they need, but cannot delete anything.

4.4.2 User profiles

As mentioned in 4.3, “Configuring Load Manager” on page 67, you use more than one server to load manage the users and applications and to get higher availability. When a user or application is load managed, the administrator is not able to know where the user is logging on.

As discussed in 1.8.1, “Storage in a Presentation Server x64 environment” on page 12, there are four kinds of user profiles: local, roaming, mandatory, and hybrid. You would not use local profiles in an environment with Load Manager because users will always want to have the same desktop each time, regardless of the server they are connecting to.

The best way of locking down user profiles is the use of hybrid profiles. You can use mandatory roaming profiles for the users so they always get the same look and feel. And to

allow users to make changes to application settings (such as the configuration of the Office Toolbar), you can use hybrid roaming profiles and specify what configuration changes will be saved.

By default, user profiles contain many files and folders. Every time the user logs on to Presentation Server, his roaming profile is copied from the master network share to the local profile directory of the server. By default, this is C:\Documents and Settings\Profiles\%USERNAME%. At logoff, the profile is copied back to the network share. It is therefore important to keep the roaming profile as small as possible to avoid slow logons.

There are a few strategies that you can use to limit the size of roaming profiles:

- ▶ Redirect certain folders to network locations outside of the user's profile.
- ▶ Exclude certain default folders from being copied.
- ▶ Apply a size limit that will not allow the profile to exceed a certain size.

To implement these strategies, we also use Windows Group Policies, as discussed in 4.4.3, "User Policies" on page 77.

To create a standardized user profile, create a user template and give this template the configuration needed. A roaming profile is defined, as shown in Figure 4-9. The user gets only a roaming profile when logging on to the Presentation Server. His workstation profile can be configured separately.




Figure 4-9 Creating roaming user profile

To create the template, do the following:

1. Log on to the Presentation Server with the template user to create the profile.
2. Make every setting you want a user to get when he first logs in to the Terminal Server.
3. Log off again and the profile is written back to the file share that you defined for the template user.

4. Rename the nuser.dat file to nuser.man to create a mandatory profile.
5. To allow the user making changes to his application configuration like MS Office Toolbar or Internet Explorer Favorites, create a Group Policy to redirect folders from the profile to the users Home drive or Root drive where the changes can be saved.

Table 4-1 lists our recommended folder redirections.

Table 4-1 Folder redirection

Folder	Redirected to
Application Data	R:\AppData
Temporary Internet Files	R:\TemplnetFiles
Cookies	Not redirected
Desktop	H:\Desktop
Favorites	H:\Favorites
History	Not redirected
Local AppData	Not redirected
Local Settings	Not redirected
My Pictures	H:\My Documents\My Pictures
Nethood	Not redirected
Personal	H:\My Documents
Printhood	Not redirected
Programs	Not redirected
Recent	R:\Recent
SentTo	Not redirected
Start Menu	Not redirected
Startup	Not redirected
Templates	H:\Templates

You can use any drive letter for folder redirection. In Table 4-1, the drive letters H:\ and R:\ were used. The *Terminal Server Home Drive*, also known as *Root drive*, is mapped to drive R:\. We recommend that you hide this Root drive in Explorer, since it only contains configuration data. This network drive is mapped through the domain level user attributes.

The normal home drive of the user H:\ is mapped through a logon script and contains the user's private data, so it should be visible.

When setting up folder redirection through Windows Group Policies, it is only possible to redirect to an UNC path. Often, it is better to redirect to a drive letter as shown in Table 4-1. In fact, some applications work better when using a drive letter for folder redirection. Doing this also prevents open or save dialog boxes from showing long UNC paths to the user. To do this folder redirection with drive letters, you can use the True Control policy templates. You can download these templates from:

<http://portal.loginconsultants.nl/forum/index.php?board=16>

Figure 4-10 shows how to add a group policy template.

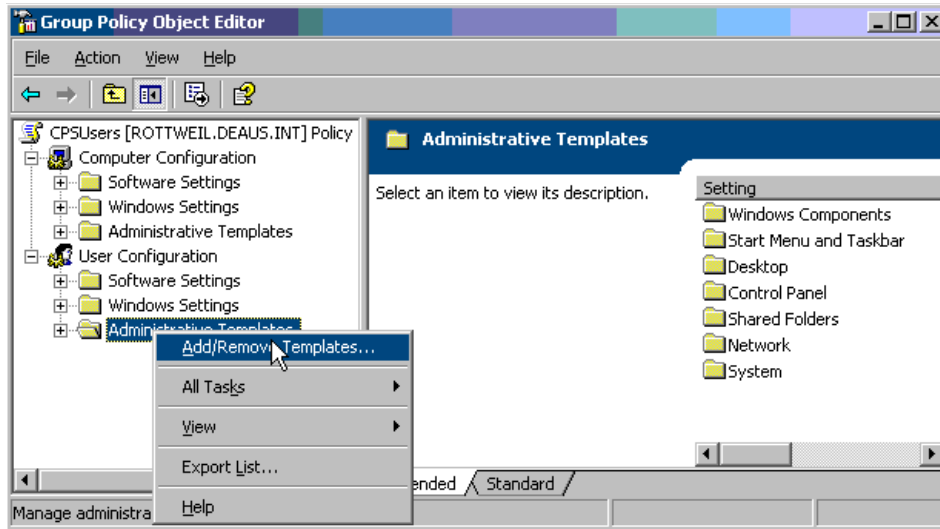


Figure 4-10 Adding a policy template

After adding the template, you have to set the filtering options to see all policies instead of only fully configurable policies. This is shown in Figure 4-11.

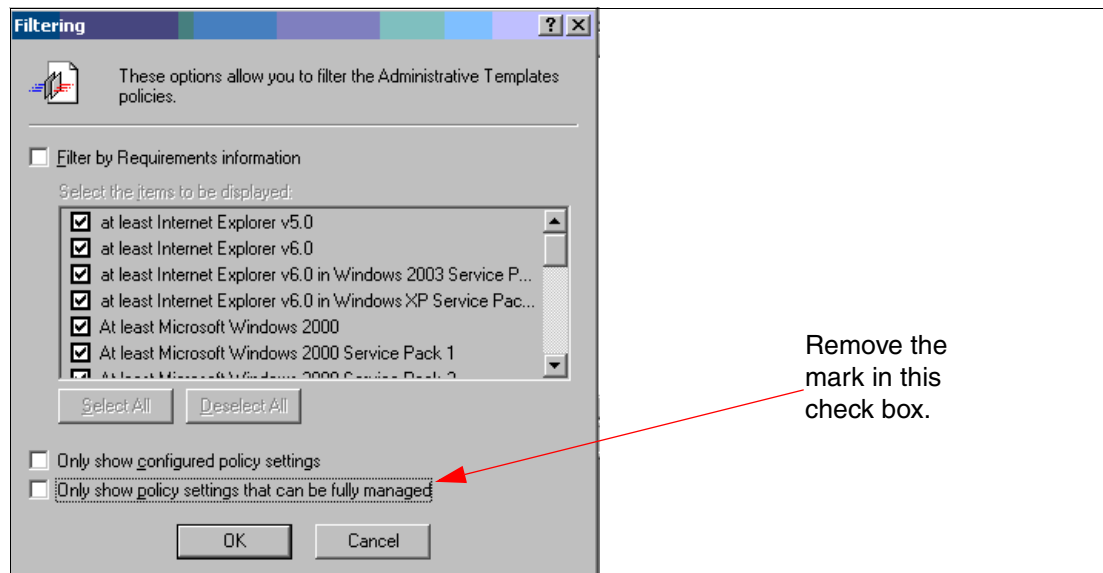


Figure 4-11 Switch to only show policy settings that can be fully managed

Now you can set the folder redirection to drive letters instead of UNC path names, as shown in Figure 4-12.

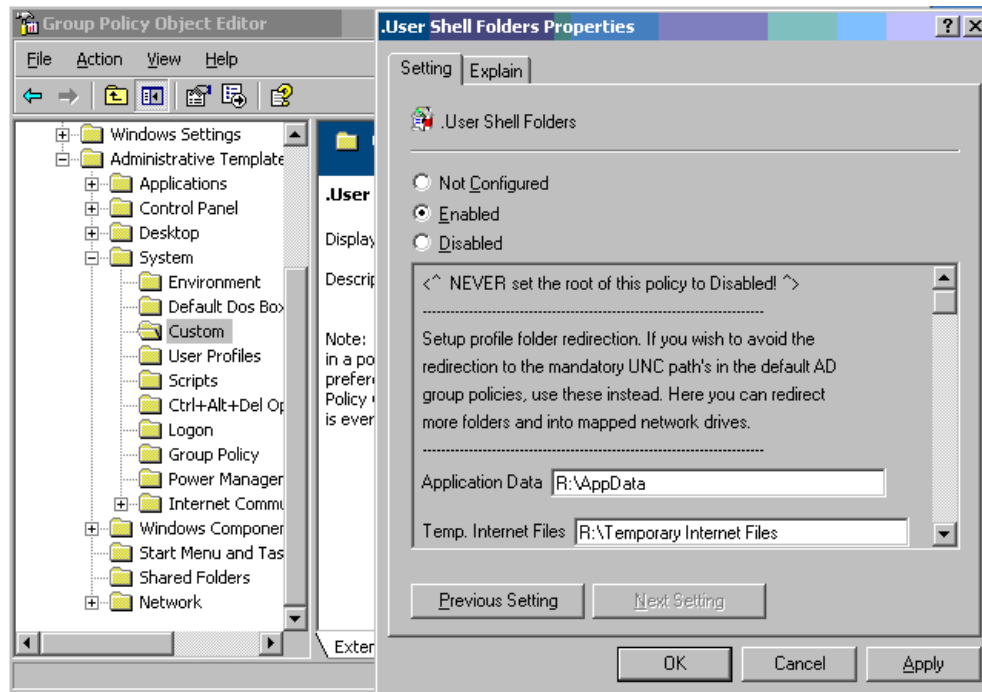


Figure 4-12 Folder redirection with drive letters

Now you created a user profile template, redirected folders, and excluded them from being copied with the profile at logon or logoff. The next step is to create a policy to limit the size of the profile and warn the user when this size is exceeded, as shown in Figure 4-13.

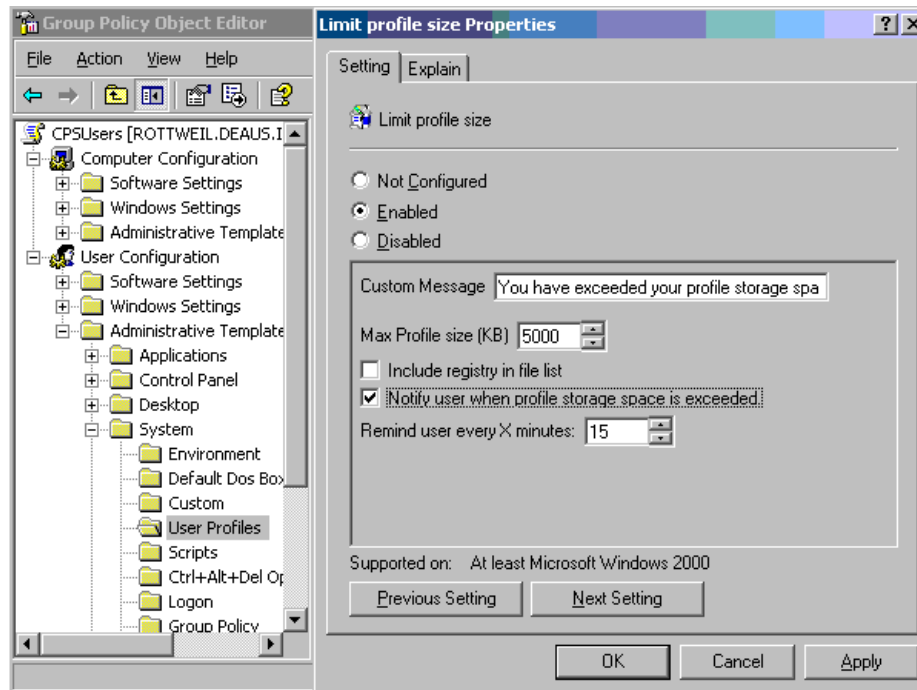


Figure 4-13 Limit User Profile size

Now you have finished creating the profile and the two most important policy settings for the profile. To make this template available for every new user, copy the template profile from the template user profile path to the local default profile of all your Presentation Server systems.

Note: Only new users get this new created profile. If you want to give it to users who still have logged on to the terminal server, you have to copy the template user profile to the network share of the involved user.

Tip: Install the User Profile Hive Cleanup Service from:

<http://www.microsoft.com/downloads/details.aspx?familyid=1b286e6d-8912-4e18-b570-42470e2f3582&displaylang=en>

This helps you ensure user sessions are completely terminated when a user logs off. Without this service, user profiles are often not unloaded successfully, which causes the copy to the roaming profile location to fail.

Some other useful policies and how to assign these policies to different Presentation Server users is discussed in the next section.

4.4.3 User Policies

Since the release of Windows 2000 Server and Active Directory, policies for user restriction have become a powerful and easy-to-manage tool. With the release of Windows Server 2003, the templates for the policies have been enhanced so that nearly every requirement to restrict users can be achieved. We discuss the use of these policies in this section.

Many administrators do not implement any policy for users that use a local PC or they implement only very light policies to give the users the freedom to manage their PC themselves. In a Presentation Server environment, you would want to restrict most of the desktop or applications that the user will get so as to avoid any failures on the server due to changes made by the user. It is therefore helpful to configure very strong policies on a Presentation Server system that will affect the user only when he logs on to a Presentation Server and leave the light policies for his local PC environment (in the case, we are using PCs and not Thin Clients). In this section, we show how to achieve this goal.

We perform the following tasks:

- ▶ Create an Organizational Unit (OU) especially for the Presentation Servers and put all Presentation Server systems in this OU, as shown in Figure 4-14 on page 78.
- ▶ Create a policy for the computer on this OU to turn on loopback processing of Group Policy Objects (GPOs).
- ▶ Create useful user policies and apply them to the user group.

We described creating the policy in 4.4.2, “User profiles” on page 72. This policy is now assigned to the new created OU and all Presentation Servers are put in this OU.

Figure 4-14 shows the Create OU for the Terminal Servers.

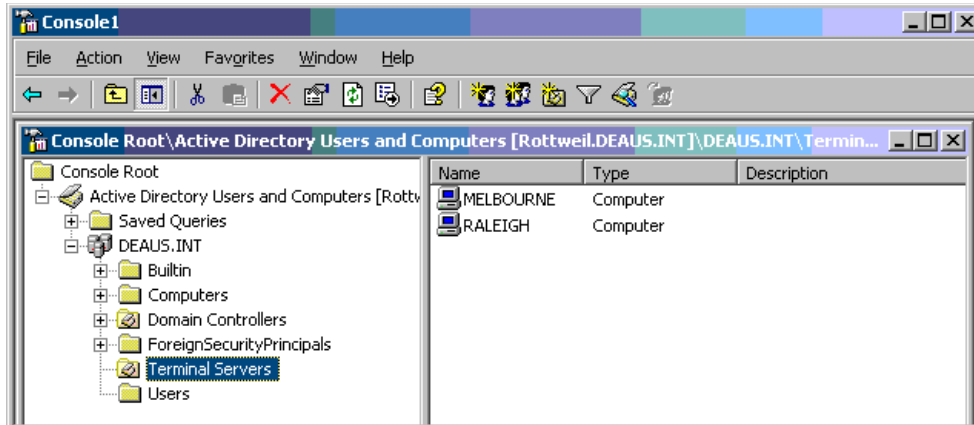


Figure 4-14 Create OU for the Terminal Servers

When a user logs on to a computer that is part of an Active Directory domain, the GPO that is applied is determined by the location of the account object of the user in Active Directory. If a user object is located in one OU and the Presentation Server system is in another OU, the Group Policy applied to the user is a result of the OU that the user's object resides in and not the OU the Presentation Server system resides in. This could be a problem for administrators when they attempt to lock down a Presentation Server system without affecting users when they log on to their normal PC.

To remedy this, you can configure a GPO to merge the settings from the Group Policy of the OU where the user's object resides in and the Group Policy of the Presentation Server system's OU. Alternatively, you can configure the Group Policy applied to the Presentation Server system's OU so that it completely replaces the user's policy. So it is possible to completely control the policy of the Presentation Server system without affecting the original policy applied to the user. To do this, enable the loopback processing in the GPO that is applied to the Presentation Server system, as shown in Figure 4-15 on page 79.

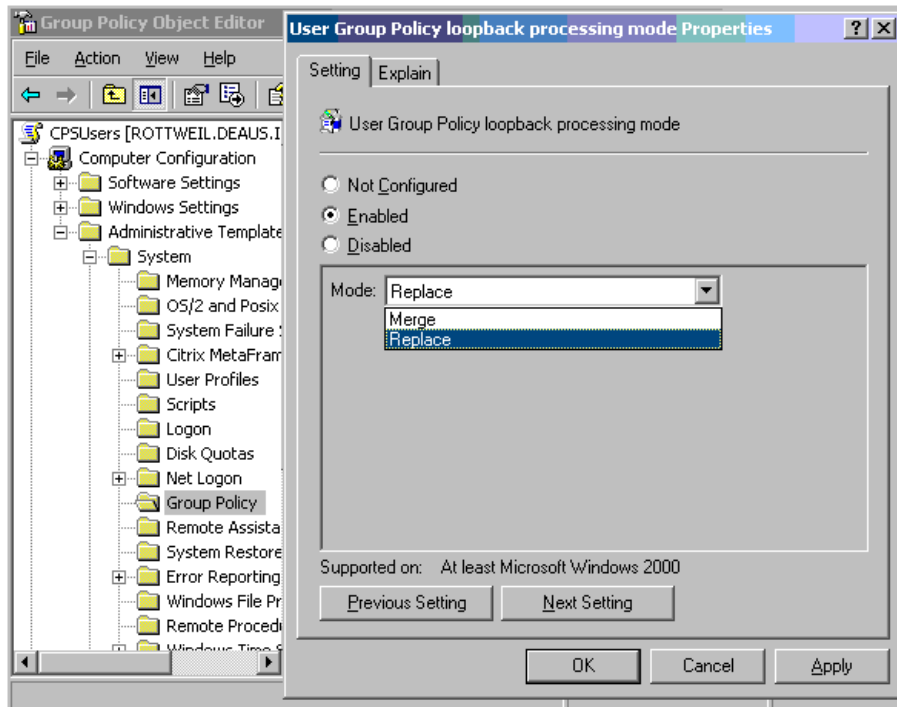


Figure 4-15 Enable loopback processing

There are two modes to enable this policy, as shown in Figure 4-15:

- ▶ Replace specifies that the user settings normally applied by the OU of the user's object are ignored and replaced with the settings configured for the OU of the Presentation Server system.
- ▶ Merge specifies that the user settings normally applied by the OU of the user's object are applied along with the settings configured for the OU of the Presentation Server system. If any two settings conflict with each other, the Presentation Server system's settings take precedence.

Note: Loopback processing is nearly always necessary to be applied in a Presentation Server environment due to the need for much more restrictive settings than in a desktop environment.

Now you can enable the restrictions that are useful for a Presentation Server system. You can find a report of the most useful settings in Appendix A, "CPS User Policies" on page 95. These settings are shown with the Group Policy Management Console. We recommend using this console to simplify the management of Group Policies for Windows Servers and Workstations. However, this Management Console does not run on x64 Editions of Windows.

There are many more policies that can be configured. Moreover, you can find many templates to plug into the GPO Manager, for example, for different versions of Microsoft Office. These can be found in the Microsoft Office Resource Kit. For more information about locking down a Windows Server 2003, refer to:

<http://www.microsoft.com/windowsserver2003/techinfo/overview/lockdown.mspx>

Also, many new Citrix Presentation Server Policies are added to Presentation Server 4.0 to manage the user sessions for a Presentation Server environment with the Citrix Management Console. One of the new policies is that control of session printers and auto-created printers is now migrated into the policies of Presentation Server. This is discussed in the next topic.

4.5 Printing in a Presentation Server 4.0 environment

Printing in a Presentation Server environment has always been a key topic. In Presentation Server 4.0, printing becomes easy to configure and many problems from older versions are solved. In this section, we take a short look at the printing enhancements and how to meet your requirements.

Printing can be divided into two main topologies:

- ▶ Session printing
- ▶ Client printing with auto-created printers

These two topologies are discussed in the following two sections. Both are shown in Figure 4-16.

4.5.1 Session printing

With session printing, no printers or printer drivers are installed on the client device; the Presentation Server has direct access to the print queue. In most environments, these are standard network printers that are accessible via a \\servername\printername share. It can also include printers where the print queue is located locally on a Presentation Server system where the printers are network-attached or direct-attached printers. Session printers are installed on the server.




Figure 4-16 Session and client printers

The session printers are those printers a user gets mapped when he logs on to Presentation Server. The rules that specify which user gets which session printer and which is the default printer are configured through Citrix Presentation Server policies. These policies can be configured for user name, user group, client IP address, or client name. How to configure a policy for session printer connection is shown in Figure 4-17.

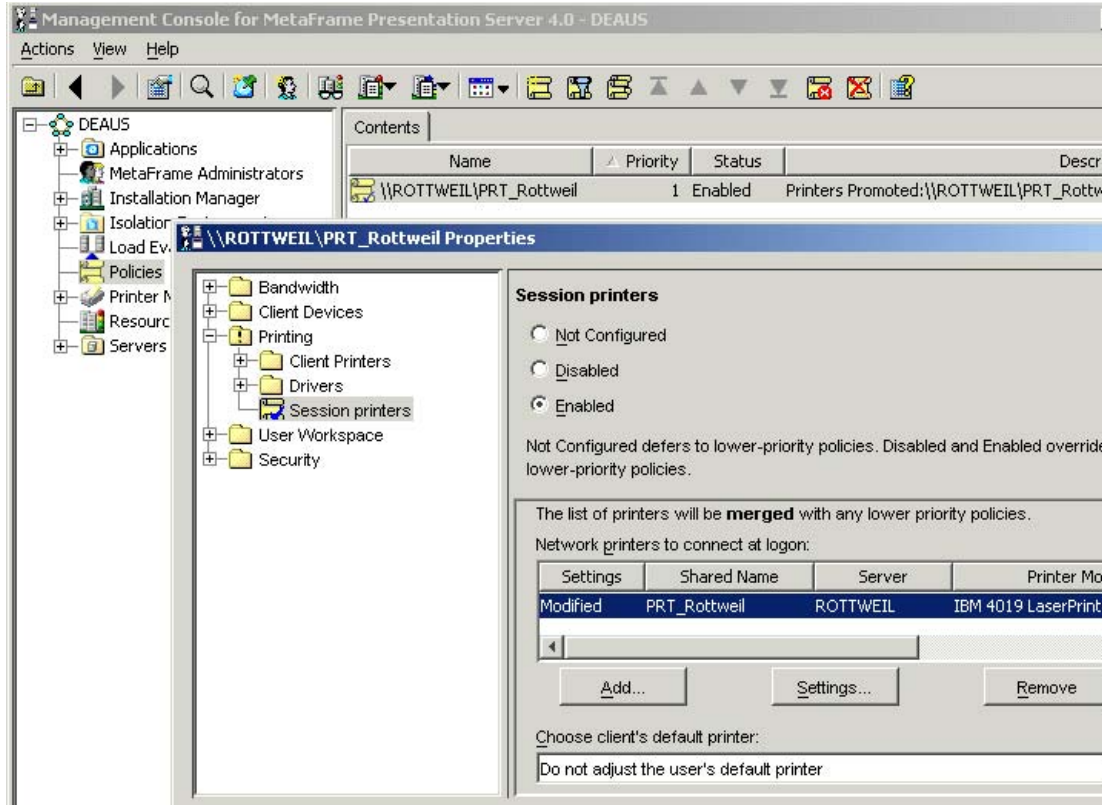


Figure 4-17 Policy for session printer

Note: The printer driver still has to be installed on the Presentation Server system to connect the defined printer for the user in his session, unless you are using the Universal Printer Driver. Using the UPD is recommended to narrow the variety of printer drivers in a heterogeneous environment down to a minimum.

With session printing, we gain the following advantages:

- ▶ Reliable printing
- ▶ Very good performance when Presentation Server and Print Server are on the same LAN
- ▶ Users always get the same printers no matter where they log on

4.5.2 Client printing with auto-created printers

With client printing, the printers are available on a user's client device. This can include printers physically attached to the client device or printers mapped through the network. Client printers are installed locally on the client.

Auto-creation of client printers is controlled through Presentation Server policies like the session printers. A policy for auto-created printers is shown in Figure 4-18.

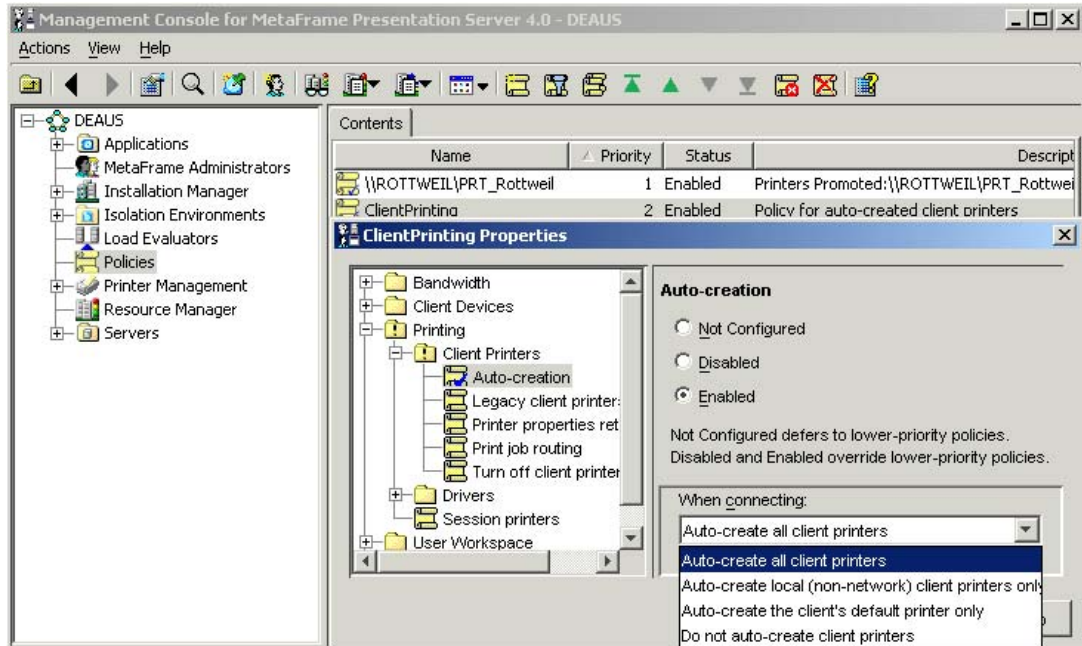


Figure 4-18 Client printing policies

The policies shown in Table 4-2 can be configured for client printing.

Table 4-2 Citrix policies for client printing

Policy rule name	Settings controlled
Auto-creation	Which client printers (if any) should be auto- created
Legacy printers	Enable fall back to Presentation Server 3.0 rules constructing client printer and port names
Printer properties retention	Where the system should retain modified printer settings
Print job routing	Whether to connect directly to the printer server for network attached client printers or always print indirectly through the client
Turn off client printer mapping	Disables client printer mapping

4.5.3 Using the Universal Printer Driver

The new Universal Printer Driver Version 3 (UPD) acts like a proxy for printer drivers operating on the Windows client. For example, actual printer capabilities and validated printer document settings are retrieved from the client and then transferred to the server so that the application on the server can see the correct settings of the printer attached to the client.

In many cases, it is useful to use the UPD instead of auto-created printers to keep Presentation Server clean from printer drivers that are not validated by the administrator. The administrator often wants to use only tested and validated printer drivers because some drivers are not Terminal Server-aware and so they can cause performance problems, hanging printer queues, or, at worst, blue screens on the server. For example, for a home office user, the administrator does not know what printer the user has attached to his workstation, so the administrator creates a policy that says that all printer drivers that are not known by

Presentation Server must use the UPD for printing. However, the user can still edit the printer and document settings as he would use his normal printer driver. This printer interface is shown in Figure 4-19.

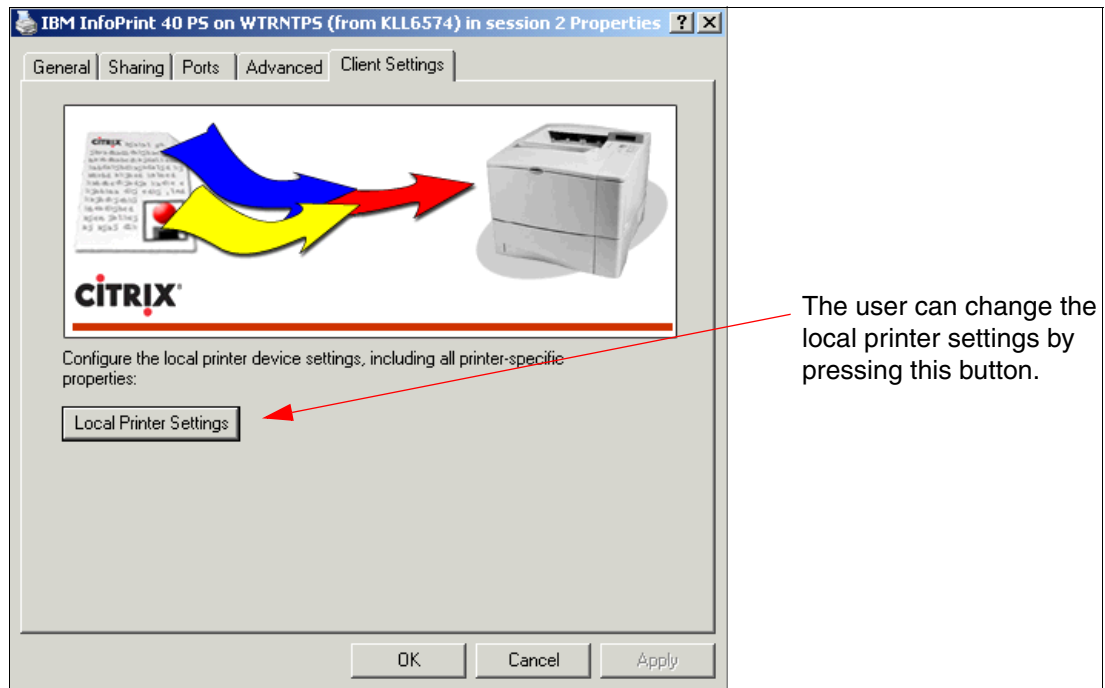


Figure 4-19 UPD user interface

Figure 4-20 shows how to activate generic UPD printing.

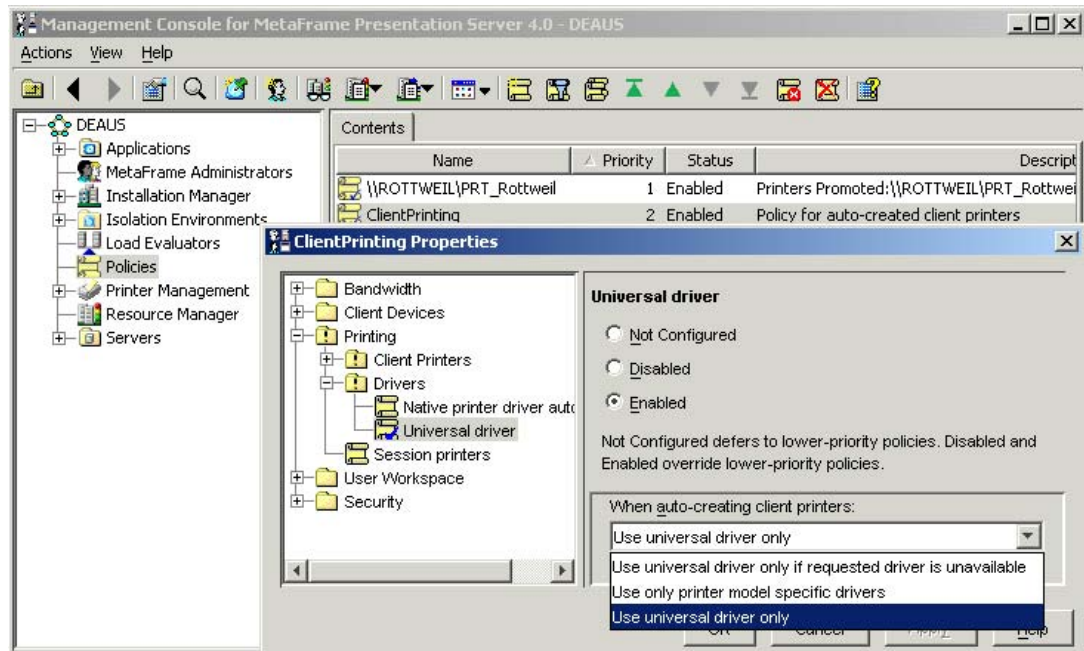


Figure 4-20 Activate generic UPD printing

Note: Citrix Presentation Server Policies use prioritization to check their validity for a user. You should consider the priority of the different policies you define. The policy with the highest priority is always applied first.

So when using the UPD 3 for printing, you have the advantage of minimizing or even eliminating native printer drivers on the Terminal Server. So it is possible to secure the Terminal Server printing environment from drivers that may cause problems on the server. If UPD is not used, the only way to print to a client attached printer is to install the printer driver onto the Terminal Server. This can happen automatically when a user is connecting to the server and the installed client printers are auto-created. So, without using the UPD 3, it is more difficult for you to keep the Terminal Server free from unwanted printer drivers.

4.6 Performance monitoring with key performance counters

In this section, we discuss the standard tools for performance monitoring and the key counters that are important for control and monitoring of performance of Presentation Server systems.

Performance monitoring is essential in a Presentation Server environment, as minor changes can directly affect many users. With the standard given tools, there are three ways to monitor server performance:

- ▶ Citrix Resource Manager
- ▶ Windows Performance Monitor
- ▶ Windows Task Manager

We recommend that you use all three tools. Windows Performance Monitor is a good tool to monitor the key counters in a Presentation Server environment. With Citrix Resource Manager, it is possible to set thresholds and send messages to the administrator if problems occur. Windows Task Manager can provide a quick view of the actual resources of the server.

4.6.1 Citrix Resource Manager

The Citrix Resource Manager comes with the Citrix Presentation Server 4.0 Enterprise Edition. It is a tool to monitor performance of Presentation Server systems, send notifications to administrators, generate reports from the servers, and even provide billing reports for application or server usage.

The default counters and thresholds that are defined in Resource Manager are shown in Figure 4-21 on page 85.

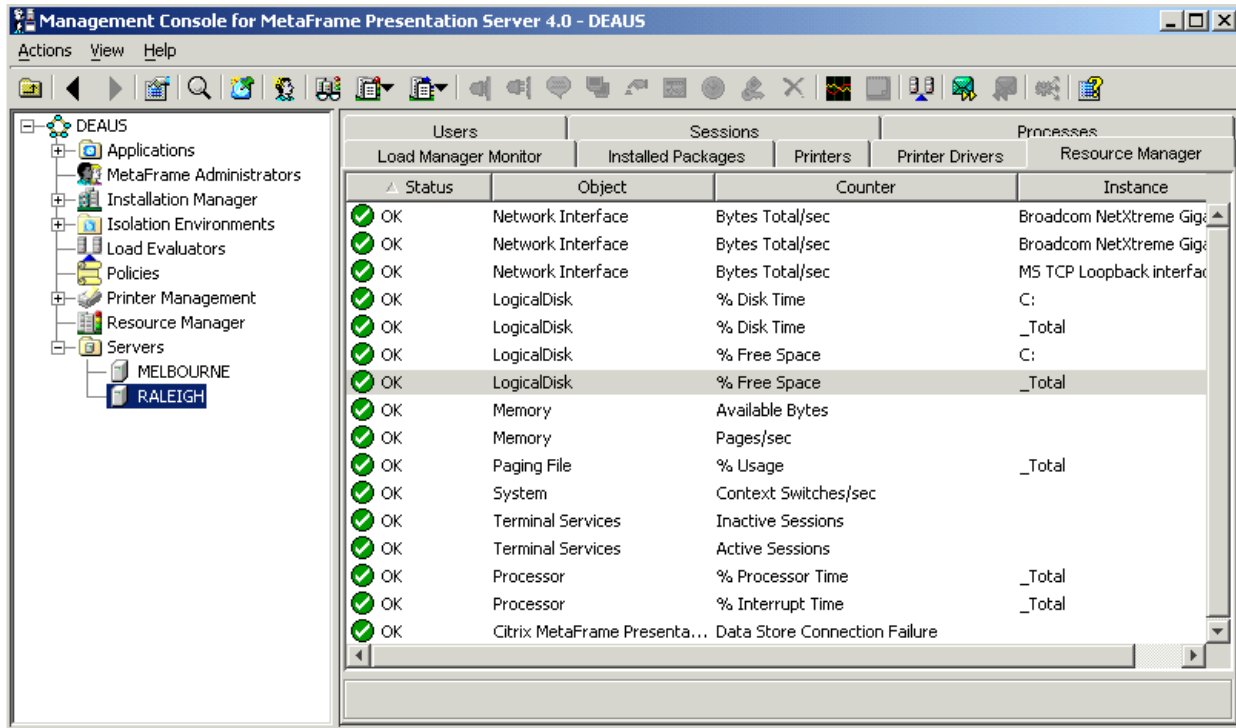


Figure 4-21 Default Resource Manager counters

In addition to these default metrics, there are many other counters you can select, as shown in Figure 4-22.

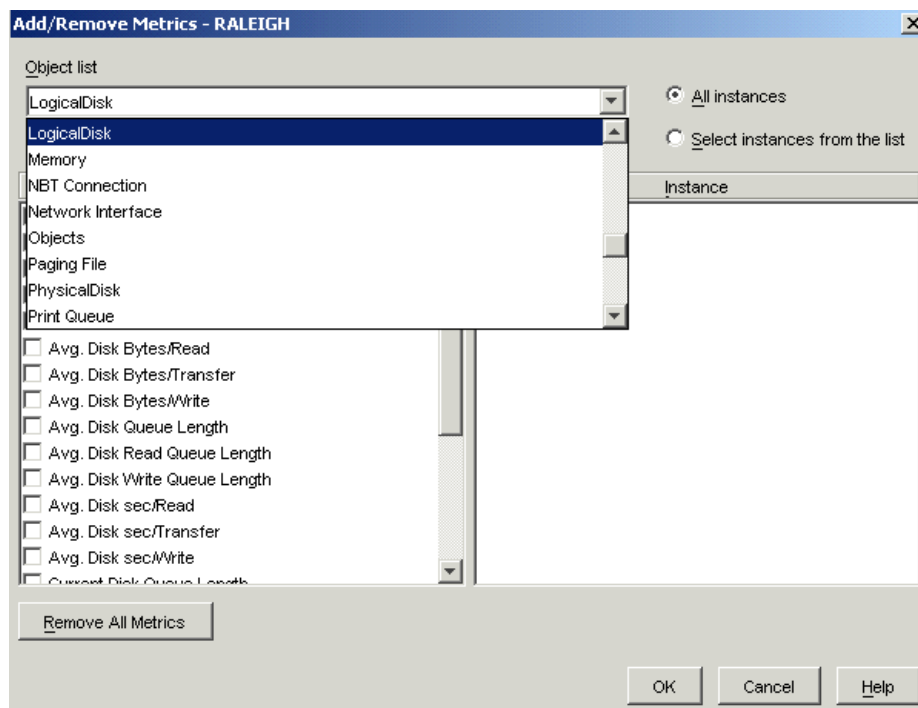


Figure 4-22 Defining Resource Manager metrics

After you have created your metrics, you can create a real time graph, as shown in Figure 4-23 on page 86.




Figure 4-23 Real-time graph

For more information about Citrix Resource Manager, see the *Resource Manager Administrator's Guide*, which can be found in the Citrix Knowledge base Article CTX106471 at:

<http://support.citrix.com/CTX106471>

4.6.2 Windows Performance Monitor

With the Windows Performance Monitor, nearly every component of the server can be monitored. We confine ourselves to the key counters for monitoring a Presentation Server environment. These counters can also be monitored with the Citrix Resource Manager.

There are only a few counters to watch when monitoring or sizing Presentation Server systems. At a minimum, the following performance counters should be monitored:

- ▶ **Memory: Pages/sec**

The counter shows the number of times per second that the server looked for something in physical memory, but did not find it there, so it needed to load it from the page file on the hard drive. Ideally, this number stays around zero. Any value up around 20 or 30 indicates that you might need more memory or fewer users.

- ▶ **Processor: % Processor Time: _Total**

This counter shows how busy the processors are. If the processor is too busy, do not automatically think that you need more processor power. The processor might also be busy when running out of memory and the processor is spending a lot of time writing and reading from the page file. If the processor utilization is high for a longer time, it might be useful to monitor the Processor Queue Length as well. This counter shows how many requests have to wait for the processor to get processing time from it. So it is possible to see how busy the processor is. The Processor Queue Length should be around zero.

- ▶ **Physical Disk: % Disk Time**

This counter shows how busy the server's hard drives are. A value around 100% indicates that the disks are very busy so you might need faster disks, more memory, or fewer users for that server. If this counter stays around 100%, monitor the Current Disk Queue Length.

This counter shows you, like the Processor Queue Length, how many requests are waiting for disk time.

- ▶ Network Interface: Bytes Total/sec

This counter shows how busy your network interface card is. Keep in mind that the network interface speed is given in megabits per second and this counter tracks bytes. So when monitoring a 100 Mbps network interface card, this counter shows 12.5 MB when the network card is fully loaded.

Most Presentation Server systems are connected to 100 Mbps networks or faster, so other hardware components are often fully loaded long before the network interface.

- ▶ Terminal Services: Active Sessions

This performance counter provides a total number of ICA or RDP sessions that are active on your server. This counter is very valuable if you have performance problems at special times of the day, so you can see how many users were on the server when the problem came up.

These performance counters are available on every Windows Server installed. In addition to these counters, there are seven more performance objects that are Presentation Server-specific or Terminal Server-specific. They are:

- ▶ Citrix IMA Networking counters

These counters monitor the IMA background communication traffic.

- ▶ ICA session counters

These counters monitor the details of specific ICA session bandwidth, including printing, compression, and the various virtual channels.

- ▶ Citrix MetaFrame Presentation Server counters

These counters monitor Presentation Server traffic, such as local host cache, zone election, IMA data store communications, and application enumerations.

- ▶ Citrix Secure Ticket Authority counters

These counters monitor the Citrix Secure Gateway.

- ▶ Citrix CPU Utilization management User counters

These counters monitor CPU entitlement, CPU reservation, CPU shares, and CPU usage on a per user base.

- ▶ Terminal Services counters

These counters monitor the total, active, and inactive Terminal Services and Presentation Server sessions.

- ▶ Terminal Services Session counters

These counters monitor the detail of specific sessions from the perspective of Terminal Services. These counters are similar to the ICA session counters except that they monitor the Terminal Server specific elements of each session. These include the percentage of processor time that a session uses, the number of handles, and the threads a session has open. Also, session input and output bytes, frames, errors, and cache performance of a session are monitored.

4.7 Performance tuning

In this section, we discuss some basic topics on how to get the best performance out of your Presentation Server environment. We give tips on how to configure Presentation Server and its applications.

Performance tuning in a x64 environment is nearly the same as in a 32-bit environment, so most of the following section can also be used in a 32-bit environment.

First of all, before tuning, we recommend monitoring the performance counters, as discussed in 4.6, “Performance monitoring with key performance counters” on page 84.

We divide this topic into the following three sections:

- ▶ Increase performance for Presentation Server 4.0
- ▶ Increase performance for Windows Server 2003
- ▶ Using third-party products

4.7.1 Increase performance for Presentation Server 4.0

Activating the CPU utilization management and the memory management features is always recommended, as long you are using Presentation Server 4.0 Enterprise Edition. These two features are enabled as shown in Figure 4-24 and Figure 4-25 on page 89.

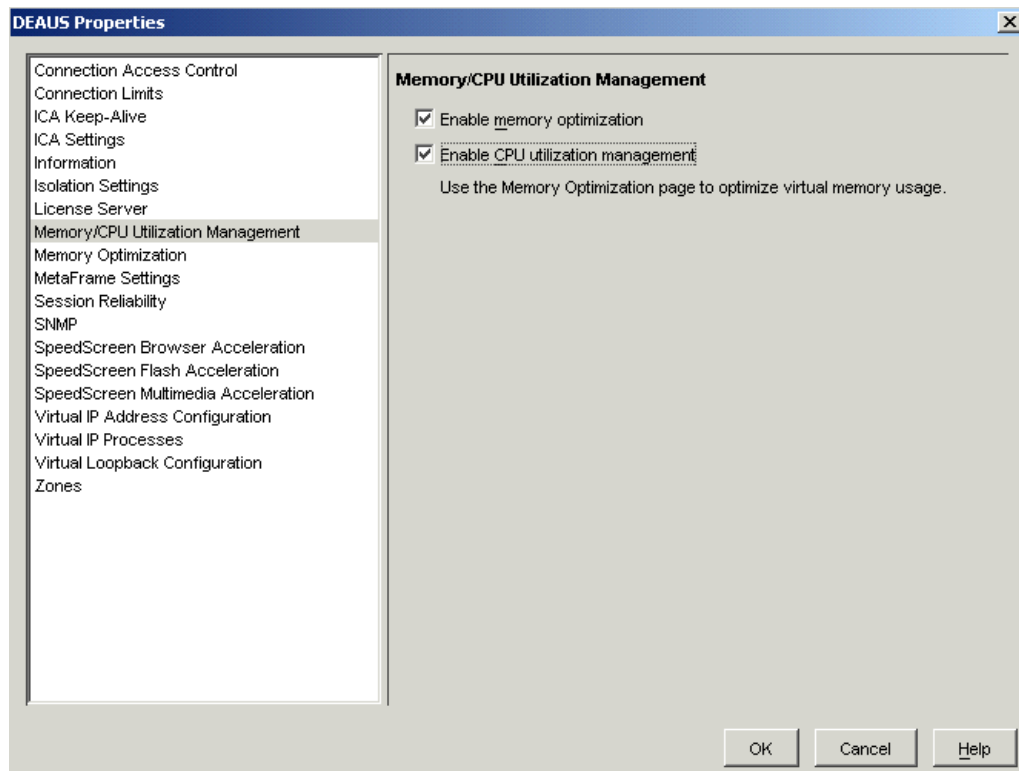


Figure 4-24 Enabling CPU Utilization Management on a Citrix Presentation Server Farm base

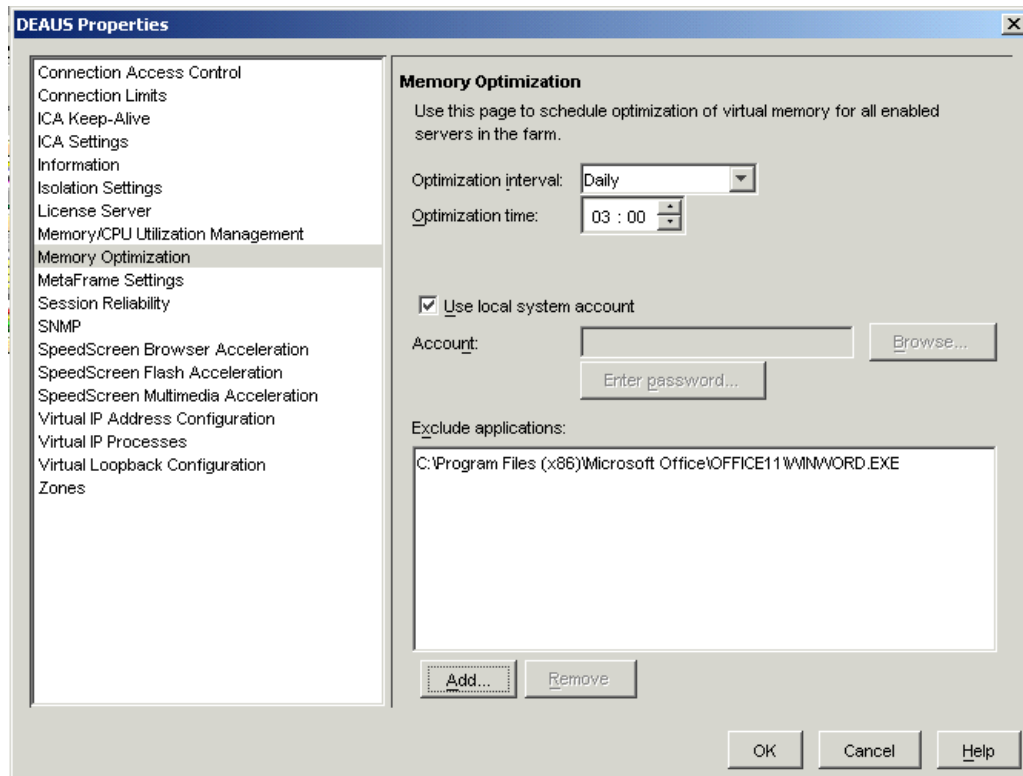


Figure 4-25 Enabling Memory management on a Citrix Presentation Server Farm base

When Memory Management is enabled, every application's base address is logged when starting. An optimization will run on the server on a defined schedule and optimize the applications that were logged. There are no known applications that do not work with one of these two features.

With the Access Suite Console, you can then create a Virtual Memory optimization report. This report shows details of how much virtual memory has been saved by using this feature.

We recommend configuring the ICA connections and the published applications to be as lightweight as possible. This includes the following:

- ▶ **Minimum color depth for your needs.**
Always configure the minimum color depth for the published application. If the application does not need more than 256 colors, configure it to use 256 colors. The fewer colors you configure, the less overhead you have in your network.
- ▶ **Disable audio**
Most applications do not need audio, so disable it centrally on the server.
- ▶ **Disable client drive mapping, COM port mapping, and LPT port mapping.**
This reduces the logon time and the time for opening the open/save dialog box on the server. Disable the mapping for all users and enable it for special users who need the mappings with a policy.
- ▶ **Install Presentation Server updates.**
Always be up-to-date with the Presentation Server updates and hotfixes. Also, deploy the latest Citrix client for the workstations. You gain the best performance by using the latest ICA client for your Presentation Server version. Furthermore, some features do not work when an older ICA client is installed on the user's workstation.

- ▶ CPU priority for the published application.

It is possible to set the CPU priority for a published application. These are divided into five settings: low, below normal, normal, above normal, and high. Use this feature as a last resort, because you have to be very sensitive with these settings. Setting an application to high priority may very quickly use the full capacity of the server.

- ▶ Set the **Allow only one instance of application for each user** check box.

For normal or light users, they mostly only need one instance for an application, so you can force this for an application. This can improve performance in particular for applications that consume a great deal of memory.

- ▶ Automatic reboot cycles.

Set the server to automatic reboot at a minimum of once a week. Many applications do not free all the memory they occupied when they are closed. So, there could be a great deal of unused memory and the performance of the system can decrease. Rebooting the server from time to time alleviates this problem.

4.7.2 Increase performance for Windows Server 2003

In addition to the recommendations in 4.7.1, "Increase performance for Presentation Server 4.0" on page 88, there are also some tips for tuning the Windows Server 2003 itself. The following methods can improve performance:

- ▶ Page file size.

Most recommendations say to configure a page file size to 1.5 times to 2.5 times the total amount of physical RAM. This does not need to be strictly followed. When determining the page file size, look at the types and numbers of applications that the users are using.

Also, consider the amount of physical memory in the system. If you have a server with 1024 MB, the 1.5 times page file is adequate. If you have 8 GB of physical memory, you can probably configure a smaller page file. Try a 4 GB page file first and increase it when it is necessary. Also, set the page file to a fixed size. Set the starting and ending size to the same amount.

- ▶ Page file location.

Put the page file on a separate fast disk drive, for example, a RAID-0 disk array or a fast flash RAM hard drive.

- ▶ Spool file location.

Move the spool file to a separate partition or disk drive.

- ▶ Disable unneeded overhead from the operating system.

Disable the use of wallpaper, menu animations, desktop themes, and active desktop. Also, disable all power consuming screen savers and set one light screen saver for all users centrally. Consider also disabling all animated icons in the system tray, including the network indicator and the system clock.

- ▶ User profiles.

Keep the user profiles small. In most cases, we use roaming user profiles for Presentation Server. When the profile becomes too big, slow logons can occur because of loading the roaming profile from the network share. Limit the profile size and remove all unneeded overhead.

- ▶ Virus scanning.

Disable real-time scan of a virus scanner. Scan once a night in a unproductive time because real-time scanning can extremely decrease performance. Most of the user files

are away from the server on a network file share of another server, so you can enable real-time scanning on that server where the files are located.

- ▶ Windows performance options.

Set the Windows performance options, as shown in Figure 4-26.

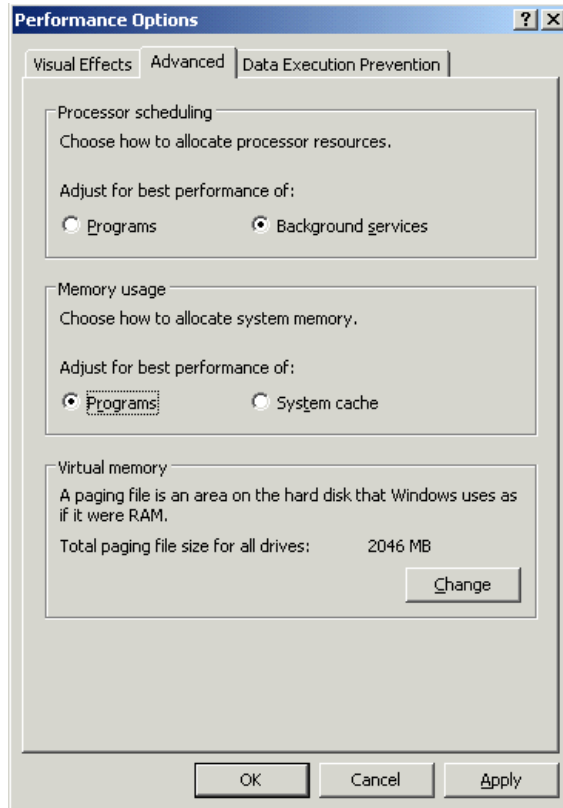


Figure 4-26 Windows performance options

- ▶ Operating system updates.

Keep the operating system up-to-date with the most recent service packs and hotfixes. Also, keep the operating systems drivers actual.

For more information about performance tuning a Windows Server 2003 refer to the Redpaper *Tuning Windows Server 2003 on IBM @server xSeries Servers*, REDP-3943. This Redpaper contains excellent information about configuring a Windows Server 2003 for best performance.

4.7.3 Using third-party products

The best way to get the highest performance out of a Presentation Server is to use a add-on product for Presentation Server. Today, there are many products on the market for Terminal Server and Presentation Server. Our recommended products are:

- ▶ Appsense Management Suite for Terminal Server

With this suite, you can manage the security and performance of Presentation Servers. You can dynamically control processors, memory, and system resources to gain the best performance. For more information, see <http://www.appsense.com>.

- ▶ TScale from RTO software

TScale is a pure software to optimize performance in a Terminal Server environment. It optimizes the use of registry, memory, and processes. Furthermore, it offers a dynamic CPU management and supports Windows Server 2003 x64 Edition. For more information, see <http://rtosoft.com>.

4.8 Adding and deploying new Presentation Server systems

In a bigger Presentation Server environment, where you have more than one or two servers and maybe thousands of users, you need methods to deploy Presentation Server systems and add them to a farm. This is discussed in the following sections.

The most common methods are:

- ▶ Unattended installation
- ▶ Cloning

4.8.1 Unattended installation

With an unattended installation, a new server is deployed through a completely new automated installation of the operating system, Presentation Server, and all applications and configurations. To complete this task, we recommend using a deployment tool like IBM Director with Remote Deployment Manager to deploy the created software images to the servers. For an unattended installation, you need to know how to create answer files and script files for a silent installation of the operating system, Presentation Server, and the applications you intend to install.

The advantages of unattended installations are:

- ▶ Unattended installations can be deployed to many different types of hardware
- ▶ Clean installation

The disadvantages are:

- ▶ Strong scripting knowledge needed
- ▶ Application deployment needs to be scripted separately or distributed via packages

For detailed information about automating deployment of Citrix Presentation Servers with IBM Director and Remote Deployment Manager, refer to the IBM Redpaper *Automating the Deployment and Image Management of a Citrix Hosted Client Environment*, REDP-3981.

4.8.2 Cloning

Server cloning can provide more rapid deployment than an unattended and scripted installation. There are a few steps to prepare a Presentation Server system for cloning. The steps to perform depend on the type of data store that is used in the farm. There are two types of data store that are possible in your Citrix Server farm. These are:

- ▶ Microsoft Access or MSDE
- ▶ Oracle, SQL Server, or DB2

If the Citrix farms data store is based on Access or MSDE, the first installed Citrix Server holds the data store by default, unless it is moved to another Citrix Server. This Citrix Server cannot be cloned. All other Servers in the farm are possible to clone. This is by design,

because only one server in the Citrix farm can hold the data store if you use MS Access or MSDE database. Furthermore, Citrix must be uninstalled to clone the server and reinstalled after cloning. When using an Access or MSDE database as a data store, all the published application information and other centrally configured Citrix specific information would be lost.

If you use an Oracle, SQL Server, or DB2 database server for the Citrix data store, no Citrix specific information is saved on any Citrix Server, so all Servers can be cloned.

So, depending on your environment, there are the following major steps for cloning a Citrix Server:

- ▶ Using MS Access or MSDE data store and still having a farm deployed.
 - a. Take any Server from the Citrix farm except the server holding the data store and uninstall Citrix Presentation Server 4.0.
 - b. Clone this server as often as you want.
 - c. Reinstall Citrix Presentation Server 4.0 on the cloned servers and put them into your server farm.
- ▶ Using Oracle, SQL Server, or DB2 data store and still having a farm deployed.
 - a. Take any Server from the Citrix farm and uninstall Citrix Presentation Server 4.0.
 - b. Clone this server as often as you want.
 - c. Reinstall Citrix Presentation Server 4.0 on the cloned servers and put them into your server farm.
- ▶ Deploying a new farm with a data store that depends on MSDE or Access.
 - a. Install the first server with Windows 2003 Server x64 and Terminal Services.
 - b. Install and configure all required applications.
 - c. Clone the server as often as you want.
 - d. Install the first server with Citrix Presentation Server 4.0 and the data store on it and create the farm.
 - e. Install Citrix Presentation Server 4.0 on all other servers and put them into the created farm.

So, for installing and deploying the first Citrix Server farm, it is a fast and easy way to clone the prepared servers. Furthermore, it is possible to keep one clone so that you have a “Master-Copy” of your servers for fast deployment if one server in the farm is going down or you need one more server.

The advantages of server cloning are:

- ▶ Applications can be imaged in addition to the operating system.
- ▶ Does not need strong scripting knowledge.
- ▶ All servers are identical.

The disadvantages are:

- ▶ All server hardware should be the same.
- ▶ Depends on the applications installed whether they can be cloned.


For more information about cloning a Presentation Server and the detailed steps, refer to:

- ▶ Citrix Knowledge base Article CTX107406 *Server Cloning with Presentation Server 4.0*, found at:

<http://support.citrix.com/article/CTX107406>

- ▶ Pages 56-59 of *Advanced Concepts Guide*, found at:

<http://support.citrix.com/article/CTX107059>



CPS User Policies

The information in this User Policy document is described in 4.4.3, “User Policies” on page 77.

Table 4-3 User Policies

Group	Policy	Value
Computer Configuration		
Windows Settings > Security Settings > Local Policies > User Rights Assignment	Allow log on through Terminal Services.	DEAUS\Domain Users
Windows Settings > Security Settings > Local Policies/Security Options > Interactive Logon	Interactive logon: Do not display last user name.	Enabled
Administrative Templates > System/Group Policy	User Group Policy loopback processing mode.	Enabled Mode: Merge
Administrative Templates > System > User Profiles	Delete cached copies of roaming profiles.	Enabled
Administrative Templates > Windows Components > Terminal Services	Allow users to connect remotely using Terminal Services.	Enabled
User Configuration > Administrative Templates		
Control Panel	Prohibit access to the Control Panel.	Enabled
Desktop	Disable Windows and Menu Animated Effects.	Enabled
Desktop	Disable Windows Contents While Dragging.	Enabled
Desktop	Prohibit user from changing My Documents path.	Enabled
Desktop	Remove Properties from the My Computer context menu.	Enabled
Desktop > Active Desktop®	Disable Active Desktop.	Enabled

Group	Policy	Value
Network > Offline Files	Prohibit user configuration of Offline Files.	Enabled
Start Menu and Taskbar	Add Logoff to the Start Menu.	Enabled
Start Menu and Taskbar	Prevent changes to Taskbar and Start Menu Settings.	Enabled
Start Menu and Taskbar	Remove and prevent access to the Shut Down command.	Enabled
Start Menu and Taskbar	Remove Clock from the system notification area.	Enabled
Start Menu and Taskbar	Remove common program groups from Start Menu.	Enabled
Start Menu and Taskbar	Remove links and access to Windows Update.	Enabled
Start Menu and Taskbar	Remove Run menu from Start Menu.	Enabled
Start Menu and Taskbar	Turn off notification area cleanup.	Enabled
System	Do not display the Getting Started welcome screen at logon.	Enabled
System	Prevent access to registry editing tools.	Enabled
System	Disable regedit from running silently?	Yes
System	Prevent access to the command prompt.	Enabled
System	Disable the command prompt script processing also?	No
System	Windows Automatic Updates.	Enabled
System > Custom	<p>User Shell Folders</p> <p>Application Data: R:\Application Data Temp. Internet Files: R:\Temporary Internet Files Cookies: %USERPROFILE%\Cookies Desktop: H:\Desktop Favorites: H:\Favorites History: %USERPROFILE%\local settings\history Local AppData: %USERPROFILE%\Local Settings\Application Data Local settings: %USERPROFILE%\Local Settings My Pictures: H:\My Documents\My Pictures Nethood: %USERPROFILE%\Nethood Personal / My Docs: H:\My Documents Printhood: %USERPROFILE%\Printhood Programs: %USERPROFILE%\Start Menu\Programs Recent: R:\Recent SendTo: %USERPROFILE%\SendTo Start Menu: %USERPROFILE%\Start Menu Startup: %USERPROFILE%\Start Menu\Programs\Startup Templates: H:\Templates</p>	Enabled

Group	Policy	Value
System > User Profiles	Limit profile size Custom Message: You have exceeded your profile storage space. Before you can log off, you need to move some items from your profile to network or local storage. Max Profile size (KB): 5000 Include registry in file list: Disabled Notify user when profile storage space is exceeded: Enabled Remind user every X minutes: 15	Enabled
Windows Components > Windows Explorer	Hide these specified drives in My Computer.	Enabled
Windows Components > Windows Explorer	Pick one of the following combinations.	Restrict A, B, C and D drives only
Windows Components > Windows Explorer	Hides the Manage item on the Windows Explorer context menu.	Enabled
Windows Components > Windows Explorer	Prevent access to drives from My Computer.	Enabled
Windows Components > Windows Explorer	Pick one of the following combinations.	Restrict A, B, C and D drives only
Windows Components > Windows Explorer	Remove CD Burning features.	Enabled
Windows Components > Windows Explorer	Remove Hardware tab.	Enabled
Windows Components > Windows Explorer	Remove Security tab.	Enabled
Windows Components > Windows Messenger	Do not allow Windows Messenger to be run.	Enabled
Windows Components > Windows Movie Maker	Do not allow Windows Movie Maker to run.	Enabled
Windows Components > Windows Update	Remove access to use all Windows Update features.	Enabled

Abbreviations and acronyms

AIE	Application Isolation Environment	MB	megabyte
AMD	Advanced Micro Devices	MDAC	Microsoft Data Access Components
API	application programming interface	MFXP	MetaFrame XP
AWE	Address Windowing Extensions	MP	multiprocessor
BIOS	basic input output system	MS	Microsoft
CD	compact disk	MSDE	Microsoft Data Engine
CISC	complex instruction set computing	NAS	network addressable storage
COM	Component Object Model	NDS	NetWare Directory Services
CPS	Citrix Presentation Server	NTFS	NT file system
CPU	central processing unit	NUMA	Non-Uniform Memory Access
CTI	Computer Telephony Integration	OS	operating system
DIMM	dual inline memory module	OU	Organizational Unit
DLL	dynamic link library	PC	personal computer
DOS	disk operating system	PCI	Peripheral Component Interconnect
ECC	error checking and correcting	PFA	Predictive Failure Analysis
EMF	Enhanced Metafile	POST	power on self test
EPIC	Explicitly Parallel Instruction Computing	RAID	redundant array of independent disks
GA	general availability	RAM	random access memory
GB	gigabyte	RDBMS	relational database management system
GP	general purpose	RDP	Remote Desktop Protocol
GPO	Group Policy Object	RISC	reduced instruction set computer
GPR	general purpose register	SAN	storage area network
GUI	graphical user interface	SAS	Serial Attached SCSI
HTTP	Hypertext Transfer Protocol	SCSI	small computer system interface
I/O	input/output	SMP	symmetric multiprocessing
IBM	International Business Machines Corporation	SQL	structured query language
ICA	Independent Computing Architecture	SRAT	Static Resource Allocation Table
IIS	Internet Information Server	TB	terabyte
IMA	Independent Management Architecture	TCO	Total Cost of Ownership
IP	Internet Protocol	UNC	universal naming convention
IT	information technology	UPD	universal printer driver
ITSO	International Technical Support Organization	URL	Uniform Resource Locator
IXA	Integrated xSeries Adapter	USB	universal serial bus
JRE	Java™ Runtime Environment	WAN	wide area network
LAN	local area network	WI	Web interface
LED	light emitting diode	ZDC	Zone Data Collector
LPT	line printer		

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this Redpaper.

IBM Redbooks

For information about ordering these publications, see “How to get IBM Redbooks” on page 103. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *Automating the Deployment and Image Management of a Citrix Hosted Client Environment*, REDP-3981
- ▶ *Implementing Windows Terminal Server and Citrix MetaFrame on IBM @server xSeries Servers*, REDP-3629
- ▶ *Introducing Windows Server x64 on IBM @server xSeries Servers*, REDP-3982
- ▶ *Planning and Installing the IBM @server X3 Architecture Servers*, SG24-6797
- ▶ *Server Consolidation with VMware ESX Server*, REDP-3939
- ▶ *Tuning Windows Server 2003 on IBM @server xSeries Servers*, REDP-3943

Other publications

These publications are also relevant as further information sources:

- ▶ *Advanced Concepts Guide*, found at:
<http://support.citrix.com/article/CTX107059>
- ▶ *Application Isolation Environments - Compatibility Issues*, found at:
<http://support.citrix.com/article/CTX109254>
- ▶ *Application Isolation Environments Explained*, found at:
<http://support.citrix.com/article/CTX106868>
- ▶ *Application Isolation Environment Supported Applications*, found at:
<http://support.citrix.com/article/CTX106100>
- ▶ *Citrix Access Suite Licensing Guide*, found at:
<http://support.citrix.com/article/CTX106282>
- ▶ *Citrix Presentation Server for Windows Server 2003 x64 Editions - Performance and Scaling Capability*, found at:
<http://support.citrix.com/article/CTX107341>
- ▶ Citrix White Papers
http://www.citrix.com/English/ps2/products/documents_onecat.asp?contentid=186&id=White+Papers
- ▶ *Considerations of Terminal Services on a Server Cluster*, found at:
<http://support.microsoft.com/support/kb/Q294/9/26.asp>

- ▶ *IBM eServer BladeCenter and Citrix MetaFrame Presentation Server Sizing* white paper, found at:
<http://www.pc.ibm.com/support?page=MIGR-52686>
- ▶ *Load Manager Administrator's Guide*, found at:
<http://support.citrix.com/article/CTX106452>
- ▶ *Locking Down Windows Server 2003 Terminal Server Sessions*, found at:
<http://www.microsoft.com/windowsserver2003/techinfo/overview/lockdown.aspx>
- ▶ *MetaFrame Presentation Server 4.0 Administrator's Guide*, found at:
<http://support.citrix.com/article/CTX106319>
- ▶ *Microsoft Server Clustering and Terminal Services Limitations*, found at:
<http://support.citrix.com/article/CTX999186>
- ▶ *Resource Manager Administrator's Guide*, found at:
<http://support.citrix.com/article/CTX106471>
- ▶ *Server Cloning with Presentation Server 4.0*, found at:
<http://support.citrix.com/article/CTX107406>
- ▶ *Service Pack 2005.04 - FAQ*, found at:
<http://support.citrix.com/article/CTX106831>
- ▶ *Support for MetaFrame with VMware ESX Server or Microsoft Virtual Server*, found at:
<http://support.citrix.com/article/CTX997956>
- ▶ *Virtual Memory Optimization Feature*, found at:
<http://support.citrix.com/article/CTX106023>

Online resources

These Web sites and URLs are also relevant as further information sources:

- ▶ AppSense
<http://www.appsense.com>
- ▶ Jeroen van de Kamp
<http://www.brianmadden.com/people/person.asp?id=13>
- ▶ RTO Software
<http://rtosoft.com>
- ▶ User Profile Hive Cleanup Service
<http://www.microsoft.com/downloads/details.aspx?familyid=1b286e6d-8912-4e18-b570-42470e2f3582&displaylang=en>
- ▶ Windows Server 2003 x64
<http://www.microsoft.com/windowsserver2003/64bit/x64>
- ▶ Windows Server 2003 Terminal Services
<http://www.microsoft.com/windowsserver2003/technologies/terminalservices>

How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



Running Citrix Presentation Server 4.0 on X3 Architecture Servers



Discover the benefits of 64-bit computing with CPS 4.0 x64

Choose the best IBM System x server for the task

Read about best practices for Presentation Server

Citrix Presentation Server 4.0 is the latest version of the popular server application for virtualizing large numbers of client systems on the one large server. This means that the user workstations can be lower function and that overall management can be centralized and streamlined. New to Citrix Presentation Server 4.0 is support for 64-bit EM64T-based servers and 64-bit operating systems, such as Windows Server 2003 x64. 64-bit support means a larger number of users can be supported on the one server, as much as three times over the previous 32-bit version.

The X3 Architecture servers are a family of high performance IBM System x servers, with four processor sockets and support for up to 64 GB of memory and six high-speed PCI-X slots. The X3 family comprises of the flagship x3950 server with its XpandOnDemand capabilities to attach up to eight systems together to form a single 32-way complex, the x3850, with its rack-dense 3U configuration, and the x3800, which supports up to 12 internal SAS hot-swap drives for up to 3.6 TB of internal storage. All three servers include the IBM X3 chipset and memory bandwidth of up to 21 GBps.

This IBM Redpaper describes what's new in Version 4.0 of Presentation Server and explains how the 64-bit combination of hardware, operating system, and Citrix application supports a greater number of users than the 32-bit equivalent. We also explain the migration steps that are available to move you to 64-bit processing, complete with migration strategies that eliminate user downtime. Finally, we provide details on best practices on using Presentation Server.

This Redpaper is aimed at the reader who is already familiar with earlier versions of Presentation Server, but are new to 64-bit computing and want to understand what makes a suitable platform to run Presentation Server 4.0 x64.

**INTERNATIONAL
TECHNICAL
SUPPORT
ORGANIZATION**

**BUILDING TECHNICAL
INFORMATION BASED ON
PRACTICAL EXPERIENCE**

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:
ibm.com/redbooks**