Autodesk

Autodesk
AutoCAD Map 3D 2009

Citrix XenApp 4.5
Performance Analysis

Notice
The information in this publication is subject to change without notice.

THIS PUBLICATION IS PROVIDED “AS IS” WITHOUT WARRANTIES OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. CITRIX SYSTEMS, INC. (“CITRIX”), SHALL NOT BE LIABLE FOR TECHNICAL OR EDITORIAL ERRORS OR OMISSIONS CONTAINED HEREIN, NOR FOR DIRECT, INCIDENTAL, CONSEQUENTIAL OR ANY OTHER DAMAGES RESULTING FROM THE FURNISHING, PERFORMANCE, OR USE OF THIS PUBLICATION, EVEN IF CITRIX HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES IN ADVANCE.

The exclusive warranty for Citrix products, if any, is stated in the product documentation accompanying such products. Citrix does not warrant products other than its own.

Product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

Table of Contents

1. Executive Summary

This section provides an overview of the Scalability Analysis project including key findings.

1.1 Project Overview

The objective of this project is to assess the scalability of Autodesk AutoCAD Map 3D 2009 software on a Citrix XenApp 4.5 platform. This analysis is an important step in the partnership of Autodesk and Citrix, providing insight and support for a growing segment of mutual customers.

The analysis took place at the LPS Integration Citrix Testing Facility at LPS Integration headquarters in Nashville, TN. LPS Integration is the premier provider of Information Technology solutions in the Southeast. Using a boutique approach to best of breed technologies, LPS Integration leads the way in providing technology and highly qualified engineers for your technology needs. LPS Integration is a Citrix Platinum Partner. http://www.lpsintegration.com

Citrix EdgeSight for Load Testing, formerly known as TLoad was used as the primary tool for automation scripting and load testing by simulating concurrent users on an isolated XenApp environment. Virtual user actions were based on workflows confirmed by Autodesk to ensure simulations apply to projected real world use cases.

1.2 Key Findings

1.2.1 Test Overview

This project employed a 32-bit platform to deliver AutoCAD Map 3D 2009 software via Citrix XenApp 4.5. The test operated on a dual Intel Quad Core platform with 16 GB of RAM with a load of up to 46 users. AutoCAD Map 3D 2009 software is not supported on 64-bit platforms. As such, no testing was carried out on a 64-bit platform for XenApp.

Citrix EdgeSight for Load Testing was used to simulate user activity and gradually increase the load on the test servers by adding virtual users over a prescribed interval. Virtual users used AutoCAD scripting commands and followed a workflow provided by Autodesk to simulate the activity of live users.

For each test run, server performance was monitored in real time via the Citrix EdgeSight for Load Testing Controller, and more detailed data was collected in Microsoft Performance Monitor logs. Upon the completion of the test, graphical analysis was performed to obtain insight into server resource utilization with increasing concurrency.

1.3 Scalability Analysis

The chart below outlines the results obtained during testing.

Citrix XenApp 4.5 32-bit Hardware Platform		Analysis
Test Platform	**Server Model**: HP DL360 G5 **OS**: Windows 2003 R2 SP1 **CPU**: 2 Quad Core Intel 2.5 GHz **Memory**: 16 GB RAM	**1-30 Users** Processor Utilization Reached a maximum utilization of 50%. Memory Utilization Physical memory was never maximized, paging is minimal. Bandwidth Utilization Total bandwidth consumed was up to 1Mb at client resolution of 1280x1024. Conclusion User experience was excellent as defined in Section 2.2.2. **21-46 Users** Processor Utilization Reached a maximum utilization of 60%. Memory Utilization Physical memory was never maximized; however at 45 users 14Gb of memory was consumed. Bandwidth Utilization Total bandwidth consumed was up to 2Mb at client resolution of 1280x1024. Conclusion User experience was good as defined in Section 2.2.2.

Test Results

Citrix recommends that Citrix XenApp administrators evaluate their environment and test data to determine an acceptable threshold for average resource utilization. The threshold accounts for the periodic CPU spike experienced during normal system usage as well as providing a buffer for unexpected or scheduled server operations. In some instances, organizations may choose to increase or decrease the acceptable threshold based on application characteristics and redundancy requirements.

2. Test Methodology

This section provides a brief overview of LPS Integration Test Methodology.

2.1 Scripted Test

For this method, a standard set of scripts are leveraged to control the actions of test users that are similar to typical production users. These scripts are developed to simulate a desired set of predefined actions (workflows), which are based on the user's role and applications used during a typical user session. Each workflow may contain sub-workflows that dictate the multiple paths users take to complete these daily tasks. These sub-workflows will be the basis for scripts that are generated. Initiation of script execution would be at set intervals to ensure that steps taken while working in an application are not repeated simultaneously for all virtual users during the test. These intervals ensure more accurate results since the application is able to respond in a more realistic manner.

2.1.1 Scripted Test Method Summary

The table below summarizes the advantages and disadvantages of the scripted test method described above.

Testing Method	Advantages:	Disadvantages:
Scripted Test	• Completely controlled with no variables. • Identical tests can be repeated as many times as needed. • User time is not required to do test. • Tests can be re-run as environment grows. • Citrix recommends this approach at the beginning of the scalability testing effort.	• Takes significant time, tools and expertise to create automated test scripts. • User skill levels not incorporated in test.

Scripted Test Summary

2.2 Performance Monitoring

Throughout the scalability testing cycles, server performance activity should be closely monitored and recorded using a monitoring tool like Citrix EdgeSight, Resource Manager or Microsoft Performance Monitor. In addition, Citrix recommends that from time to time throughout a testing cycle, a real user or administrator should launch an ICA connection to the test XenApps to examine the application's performance from an end user perspective. The shadowing feature included with XenApp can also be leveraged to ensure application response time is acceptable during a Scripted Test cycle. With Citrix EdgeSight for Load Testing, sessions can be viewed directly from their respective Launchers. No shadowing is required.

2.2.1 Performance Objects and Counters

The table below contains the most critical performance counters that should be recorded and monitored during test execution and then later analyzed after test execution. These performance objects and counters can be utilized from within Resource Manager, Microsoft Performance Monitor, or Citrix EdgeSight. For more information on the Performance Objects and Counters, please see Appendix C.

Objects:	Counters
Cache	Copy Read Hits %
Logical Disk	% Disk Time, % Free Space
Memory	Available Bytes, Free System Page Table Entries, Page Reads/Sec, Page Writes/Sec, Pages Input/Sec, Pages Output/Sec, Pool Nonpaged Bytes, Pool Paged Bytes
Network Interface	Bytes Total/Sec
Paging File	% Usage
Physical Disk	Current Disk Queue Length
Process	Working Set
Processor	% Interrupt Time, % Processor Time
Redirector	Current Commands
Server	Pool Nonpaged Failures, Pool Paged Failures
System	Context Switches, Processor Queue Length
Terminal Services	Active Sessions

Recommended Performance Objects and Counters

2.2.2 Real-User Experience

The real-user experience is a key value used to establish thresholds for the scalability testing, as application responsiveness exposes the net effect of all stress testing on the test servers. Unlike other data, the real user experience is a subjective measure that cannot be found in any monitoring tool.

In scalability testing, the real user experience is determined by launching a published application session to the test XenApp and executing predefined tasks during the test to validate the "user experience." The real user experience has been defined by LPS Integration, Inc in four levels or categories detailed below. The categories are determined by login times, the browsing menus experience, executing queries and screen refresh speed.

	Application Responsiveness
Excellent	Equivalent or better than local PC performance.
Good	Screen updates are fluid and there is minimal effect on user's workflow.
Acceptable	Slower screen updates are somewhat noticeable and latency is increased; however the user is still able to function in a productive manner.
Poor	The session becomes unbearably slow, frozen or disconnected. Therefore, the user cannot continue his/her tasks in a productive manner.

User Experience Measurements

LPS Integration recommends that administrators explain these user experience measurements to any "live users" before they take part in the scalability testing effort. Then after test execution completes, the end-users will be able to accurately relay to administrators the type of user experience while using the application(s). As mentioned earlier, LPS Integration believes it is critical to understand and determine what an acceptable level of response time is when the system is considered "fully loaded."

3. Test Environment

The purpose of this section is to document the test environment built by Citrix Consulting for the scalability testing effort.

3.1 High Level Diagram

The following diagram provides a high-level overview of the scalability testing environment.

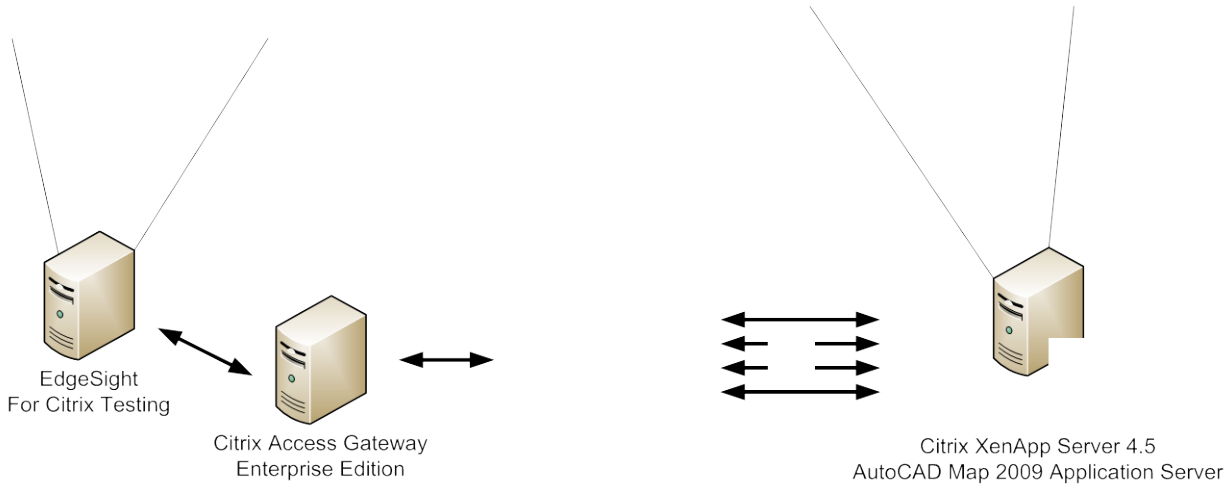

Scalability Test Environment

The diagram above provides a high-level overview of the scalability testing environment. As illustrated, the scalability test environment consists of the following components:

- **Citrix EdgeSight for Load Testing Controller** – The Controller is used to manage and run the load test. It is the management console for the recording, editing, and playback of scripts as well as the central monitoring and reporting console.
- **Citrix Access Gateway Enterprise Edition** – The Launcher is used by the Controller to launch virtual users.
- **Citrix XenApp** – Allows multi-user access to the published resources on a Microsoft Terminal Server.

3.2 Citrix XenApp Specifications

3.2.1 Hardware Specifications

The table below outlines the hardware specifications for the Citrix XenApp used in the scalability testing.

Servers	AutoCAD Map 3D
Make, Model, Type	HP DL360 G5
Processor Type/Speed	Intel 2.5 GHz
# of Processors	2 Quad Core
Memory (GB)	16 GB
Disk Capacity	72 GB
Disk Speed	15K
Array Controller	SmartArray 6i
RAID	Not Enabled
Cache Ratio (Read/Write)	100 % Read
Page File Size	4 GB min/4GB max
NIC card	HP NC7782 Gigabit Server Adapter
Operating System	Windows 2003 R2 Enterprise, SP1
XenApp	XenApp 4.5

Hardware Specifications

4. Scalability Analysis and Results

This section is broken into two separate parts, showing separate analysis and results for the two tests. The details of the scalability results are discussed with graphical analysis to provide additional insight on the impact of users on server resources. For more information on each platform, please refer to Section 3.2, Citrix Server Specifications .

4.1 Scalability Analysis

4.1.1 Server Resource Analysis

The graph below shows the results of processor performance testing for 46 users. While load increased as users increased, the processor performance was within acceptable limits. A consistent increase in processor utilization and process queue length is apparent. Scripted user testing revealed no noticeable performance impacts for this level of users. Based on previous tests, CPU optimization, a feature available with Citrix XenApp 4.0 and later, will not provide significant gains for this application.

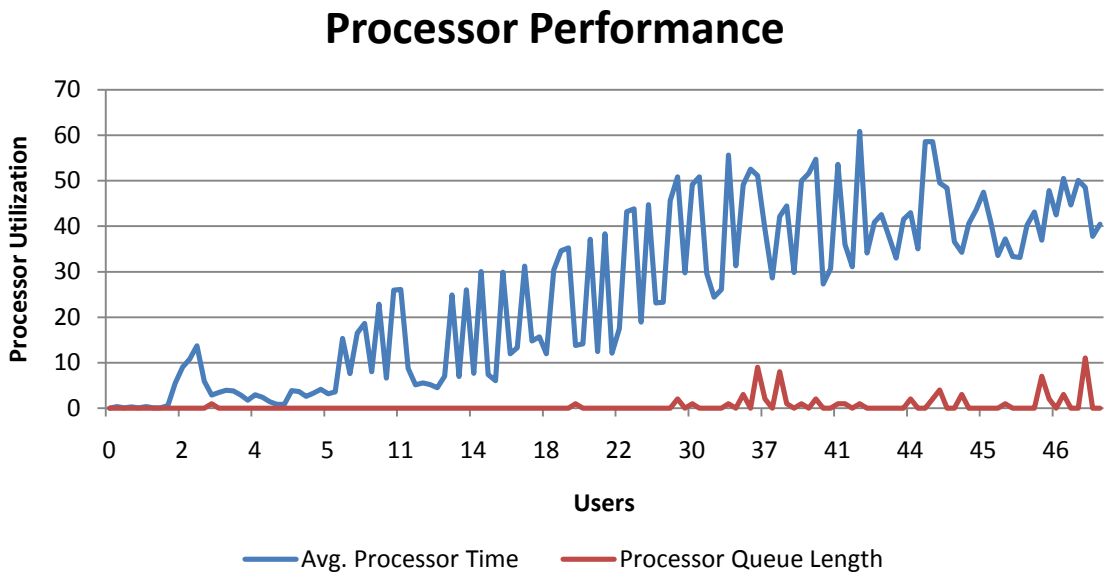

Further analysis reveals a steady increase in the amount of memory consumed by running processes, but the amount of available memory is still within operational limits. Because physical memory was never maximized, paging is minimal. Based on the results shown below, we believe that additional users could be added to the test server.

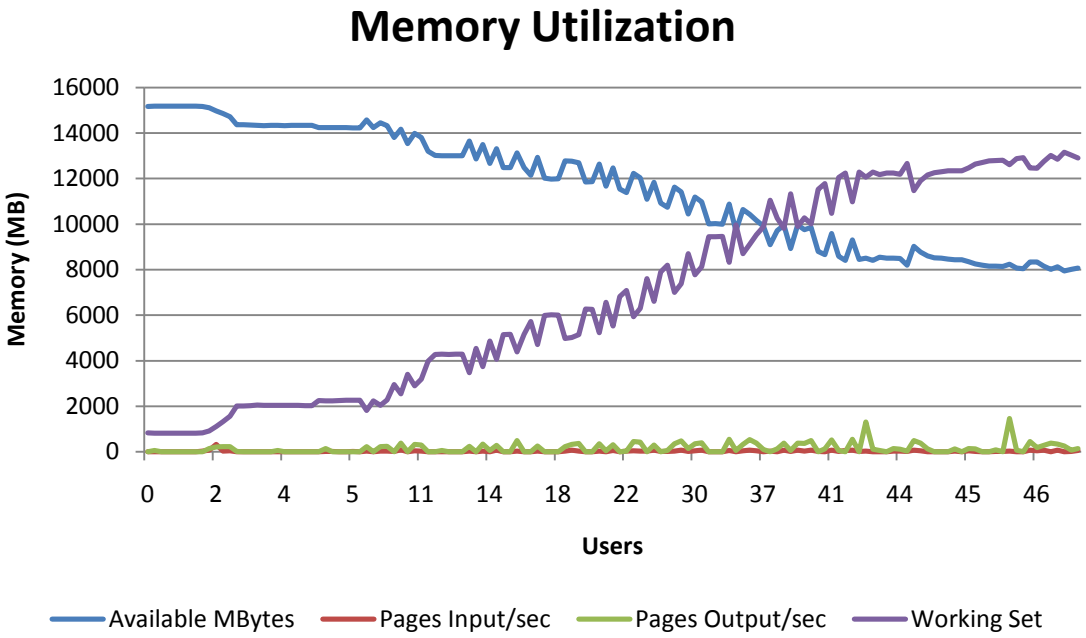

4.1.2 Anticipated User Experience

Anticipated usability thresholds are listed in the chart below. For an explanation of the rating system, please refer to Section 2.2.2, Real User Experience.

Rating	User Load
Excellent	0 – 30 Users
Good	31 – 46 Users
Acceptable	47 – 66 Users (projected)
Poor	> 67 Users (projected)

Performance also varies based on specific tasks performed in the workflows. Performance Indicators show acceptable results for all functionality for 46 users or less.

4.2 Bandwidth Analysis

4.2.1 Bandwidth Utilization Results

Based on the performance in the test, with a projected load of 111 users the bandwidth utilization is displayed below. As you can see total bandwidth approaches 50 percent of the available bandwidth. Based on these results, many more users could be added to the server before bandwidth would reach its maximum. These results bode well for using the application across WAN links that have limited bandwidth availability such as T1's and T3's.

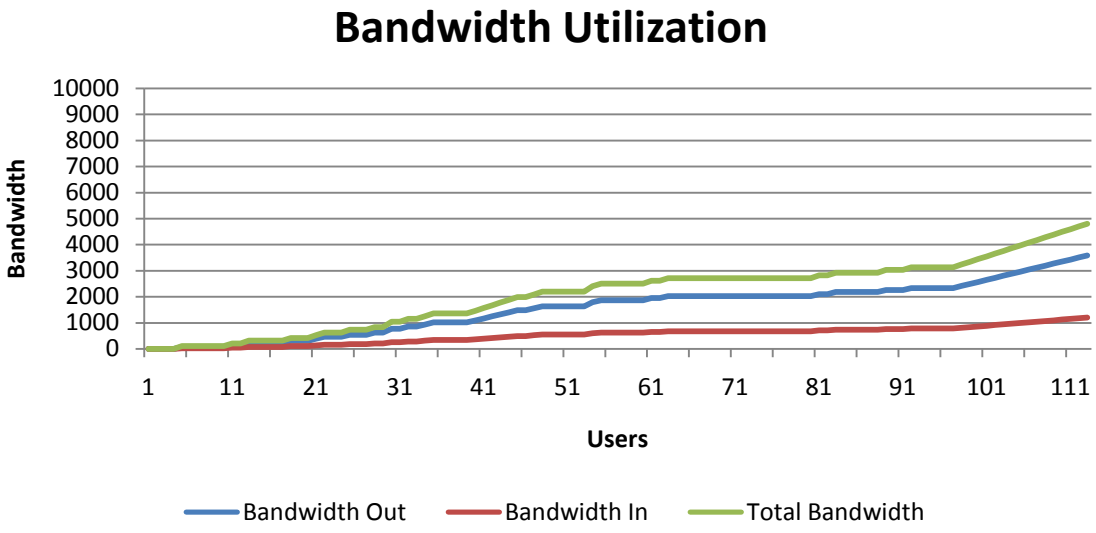

5. Appendix A: Application Workflows

This section details the application workflow that was approved by Autodesk and used in the scalability testing.

5.1 AutoCAD Map 3D 2009 Workflow

Scripted Functions

1. Startup…
2. Add data
    a. Parcels
    b. ODBC – parcel information
    c. Roads
3. Calc select zoom
    a. Add Calculated field for parcel area
    b. Round calculation to .00
    c. Add buffer1.sdf (buffer was previously created as it could not be done multiple times)
    d. Zoom to extents of buffer and make transparent
    e. Add labels to parcels for prop value
        i. Zoom threshold of 20000
    f. Add labels for Road names
        i. Zoom threshold 40000
    g. Zoom in and out and pan
    h. Zoom t o parcel extents
    i. Zoom window and Pan
    j. Zoom to buffer extents
4. Query
    a. Query parcels using buffer polygon
    b. Edit query to add Prop value constraint
    c. Theme remaining parcels
5. Plot
    a. Initialize plot to a DWF
    b. Preview
    c. Process

More Complex Functions

1. Startup…
2. Add data
    a. Parcels
    b. ODBC – parcel information
    c. Roads
3. Calc select zoom

    a. Add Calculated field for parcel area
    b. Round calculation to .00
    c. Add buffer1.sdf (buffer was previously created as it could not be done multiple times)
    d. Zoom to extents of buffer and make transparent
    e. Add labels to parcels for prop value
        i. Zoom threshold of 20000
    f. Add labels for Road names
        i. Zoom threshold 40000
    g. Zoom in and out and pan
    h. Zoom t o parcel extents
    i. Zoom window and Pan
    j. Zoom to buffer extents
4. Query
    a. Query parcels using buffer polygon
    b. Edit query to add Prop value constraint
    c. Theme remaining parcels
5. Plot
    a. Initialize plot to a DWF
    b. Preview
    c. Process

6. Appendix B: Performance Objects and Counters

This section details the counters used to monitor Citrix XenApp during the scalability testing effort.

Counter	Description
Cache: Copy Read Hits %	The percentage of cache copy read requests that hit the cache, that is, they did not require a disk read in order to provide access to the page in the cache. A copy read is a file read operation that is satisfied by a memory copy from a page in the cache to the application's buffer.
LogicalDisk: % Disk Time	The average number of read and write requests that were queued for all logical disks. Sustained value of 2-3 or greater indicates disk speed may become a bottleneck, and typically increases processor activity. If hard disk performance becomes a bottleneck, a hardware disk controller that includes both read and write cache can improve disk performance.
LogicalDisk: % Free Space	% Free Space is the percentage of total usable space on the selected logical disk drive that was free.
Memory: Available Bytes	Amount of physical memory available to processes, measured in MB. Paging should be monitored if less than 25% of physical memory is available, as excessive paging may occur.
Memory: Free System Page Table Entries	The number of page table entries not currently in used by the system. This counter displays the last observed value only; it is not an average.
Memory: Pages Reads/sec	The rate at which the disk was read to resolve hard page faults. It shows the number of reads operations, without regard to the number of pages retrieved in each operation. Hard page faults occur when a process references a page in virtual memory that is not in working set or elsewhere in physical memory, and must be retrieved from disk. This counter is a primary indicator of the kinds of faults that cause system-wide delays. It includes read operations to satisfy faults in the file system cache (usually requested by applications) and in non-cached mapped memory files. Compare the value of Memory\\Pages Reads/sec to the value of Memory\\Pages Input/sec to determine the average number of pages read during each operation.
Memory: Pages Writes/sec	The rate at which pages are written to disk to free up space in physical memory. Pages are written to disk only if they are changed while in physical memory, so they are likely to hold data, not code. This counter shows write operations, without regard to the number of pages written in each operation. This counter displays the difference
Memory: Pages Input/sec	The rate at which pages are read from disk to resolve hard page faults. Hard page faults occur when a process refers to a page in virtual memory that is not in its working set or elsewhere in physical memory, and must be retrieved from disk. When a page is faulted, the system tries to read multiple contiguous pages into memory to maximize the benefit of the read operation. Compare the value of Memory\\Pages Input/sec to the value of Memory\\Page Reads/sec to determine the average number of pages read into memory during each read operation.
Memory: Pages Output/sec	The rate at which pages are written to disk to free up space in physical memory. Pages are written back to disk only if they are changed in physical memory, so they are likely to hold data, not code. A high rate of pages output might indicate a memory shortage. Windows writes more pages back to disk to free up space when physical memory is in short supply. This counter shows the number of pages, and can be compared to other counts of pages, without conversion.

Memory: Pages/sec	The number of memory pages read from or written to disk to resolve memory references that was not in memory at the time of reference. A value greater than 100 is not a problem unless it is accompanied by low Available Bytes or high Disk Transfers/sec
Memory: Pool Nonpaged Bytes	The size, in bytes, of the nonpaged pool, an area of system memory (physical memory used by the operating system) for objects that cannot be written to disk, but must remain in physical memory as long as they are allocated. Memory\\Pool Nonpaged Bytes is calculated differently than Process\\Pool Nonpaged Bytes, so it might not equal Process\\Pool Nonpaged Bytes_Total. This counter displays the last observed value only; it is not an average.
Memory: Pool Paged Bytes	Pool Paged Bytes is the size, in bytes, of the paged pool, an area of system memory (physical memory used by the operating system) for objects that can be written to disk when they are not being used. Memory\\Pool Paged Bytes is calculated differently than Process\\Pool Paged Bytes, so it might not equal Process\\Pool Paged Bytes_Total. This counter displays the last observed value only; it is not an average.
Network Interface: Bytes Total/sec	The rate at which bytes are sent and received over each network adapter, including framing characters. Network Interface\\Bytes Received/sec is a sum of Network Interface\\Bytes Received/sec and Network Interface\\Bytes Sent/sec.
Paging File: % Usage	Percentage of page file in use. If greater than 75% of the page file is in use, physical memory (RAM) should be increased.
PhysicalDisk(_Total): Current Disk Queue Length	Current Disk Queue Length is the number of requests outstanding on the disk at the time the performance data is collected. It also includes requests in service at the time of the collection. This is an instantaneous snapshot, not an average over the time interval. Multi-spindle disk devices can have multiple requests that are active at one time, but other concurrent requests are awaiting service. This counter might reflect a transitory high or low queue length, but if there is a sustained load on the disk drive, it is likely that this will be consistently high. Requests experience delays proportional to the length of this queue minus the number of spindles on the disks. For good performance, this difference should average less than two.
Process: Working Set (_Total)	Working Set is the current size, in bytes, of the Working Set of this process. The Working Set is the set of memory pages touched recently by the threads in the process. If free memory in the computer is above a threshold, pages are left in the Working Set of a process even if they are not in use. When free memory falls below a threshold, pages are trimmed from Working Sets. If they are needed they will then be soft-faulted back into the Working Set before leaving main memory.
Processor: % Interrupt Time	Percentage of total usable space on the selected logical disk drive that was free.
Processor: % Processor Time	Percentage of elapsed time a CPU is busy executing a non-idle thread. High value is a concern only if accompanied by a Processor Queue Length sum greater than <12 x # of CPU's> or growing with % Processor Time greater than 80-90%.
Server: Pool Nonpaged Failures	The number of times allocations from nonpaged pool have failed. Indicates that the computer's physical memory is too small.
Server: Pool Paged Failures	The number of times allocations from paged pool have failed. Indicates that the computer's physical memory or paging file are too small.

System: Context Switches/sec	Combined rate at which all CPU's are switched from one thread to the other. This occurs when a running thread voluntarily relinquishes the CPU, is preempted by a higher-priority thread, or switches between user mode and privileged mode to use an executive or subsystem service. A baseline should be established to determine if excessive context switching is occurring. For example, some systems have been observed to behave just fine with context switches between 50,000 – 60,000, but on other systems values this high negatively impact performance.
System: Processor Queue Length	Number of threads in the processor queue; for ready threads only, not threads that are running. Greater than *<12 x # of CPU's>* for 5-10 minutes or with %Total Processor Time of 80%-90%.
Terminal Services: Active Sessions	Number of active Terminal Server sessions.

Performance Objects and Counters

851 West Cypress Creek Road Fort Lauderdale, FL 33309 954-267-3000 http://www.citrix.com