



TiGi Performance Analyzer Report

Sample Report XYZ Corporation

Executive Summary

TiGi Corporation has compiled the performance counter data captured from [Company XYZ's] Citrix MetaFrame production server environment. The analysis of this data, along with our corresponding recommendations, can be found in this document which is intended as a baseline for examining the performance and scalability of your Citrix MetaFrame environment.

Based on the data provided, your current hardware platform cannot accommodate the volume of requests at a rate which allows for predictable and acceptable performance levels. The root cause of these disk requests include the applications you are running, their utilization of virtual memory (the pagefile), their "temp" directories or cache, the base operating system itself in registry reads and writes, printing, and the general random nature of adding more simultaneous users to the same server.

Expected Impact Analysis

RECOMMENDATION	EXPECTATION
1) This is one	Performance:
2) This is two	Scalability:

Detailed Analysis

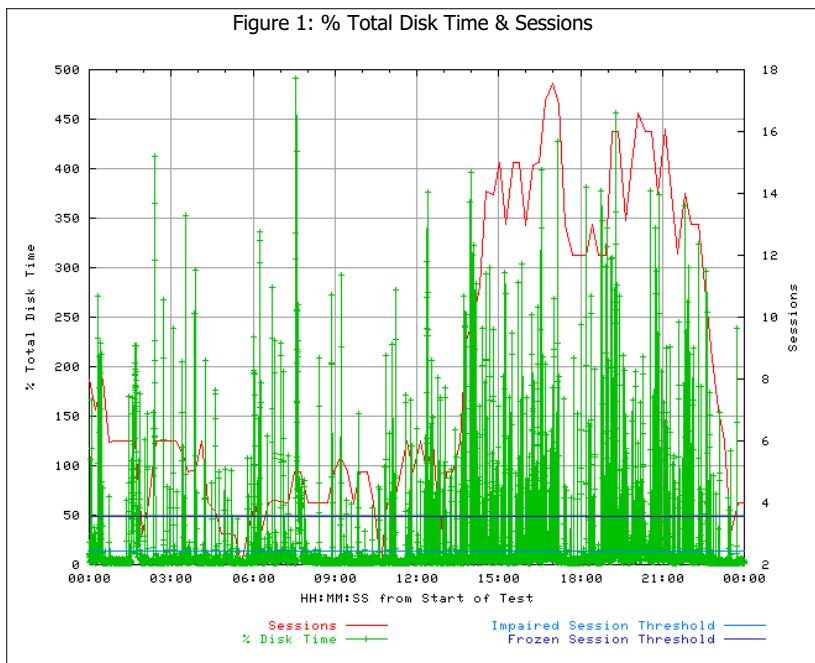


Figure 2: % Total Disk Time & Sessions

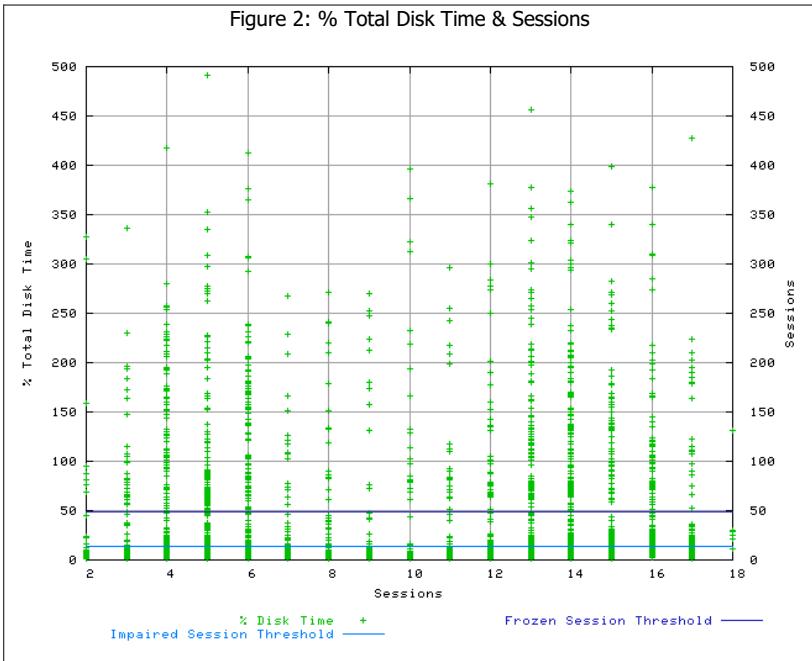


Figure 2 shows the distribution of Total % of Disk Time and how it is an issue that gets worse with each additional users login. **Most multi-user environments, such as Citrix MetaFrame, have impaired performance in sessions at % of Disk Time levels as low as 10% through 45% consistently. Any samples above the "blue target zone" at right usually represent user sessions which are consistently impaired, suffer from periodic freezing, or end users encounter regular "white-screens."**

% Total Disk Time: % Disk Time is the percentage of elapsed time that the selected disk drive is busy servicing read or write requests.

Sessions: Number of active Terminal Services sessions or ICA users.

Impaired Session Threshold shows the peak level for optimal server performance. Samples above this line can result in negative impacts on end user session performance.

Frozen Session Threshold shows the level where end user session impact can result in frozen sessions and significant application performance degradation.

Figure 3: Individual Drive % Disk Time & Sessions

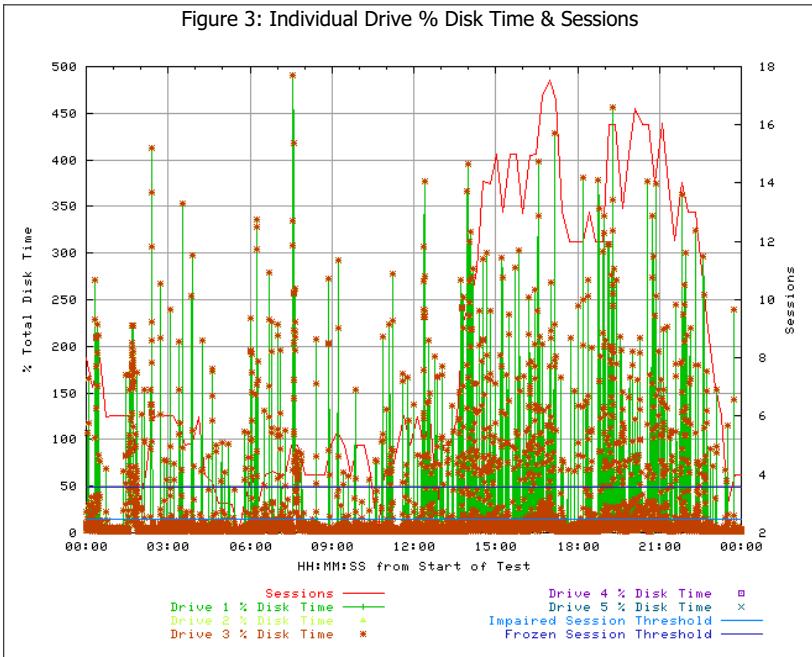


Figure 3 compares the % of Disk Time for each of the logical or physical drives. This graph is useful in showing which drives may be the largest consumer of available system IO. Each drive letter is represented by a separate color. Please look in Appendix A for determining which drive letters were captured.

% Total Disk Time: % Disk Time is the percentage of elapsed time that the selected disk drive is busy servicing read or write requests.

Sessions: Number of active Terminal Services sessions or ICA users.

Impaired Session Threshold shows the peak level for optimal server performance. Samples above this line can result in negative impacts on end user session performance.

Frozen Session Threshold shows the level where end user session impact can result in frozen sessions and significant application performance degradation.

Figure 4: Pages/Sec, % Disk Time, & Sessions

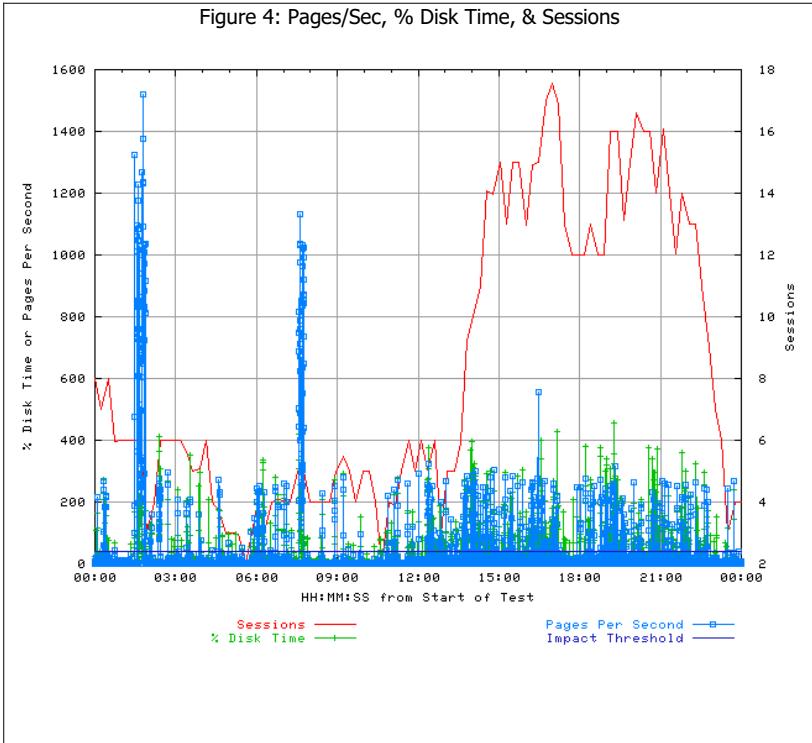


Figure 4 compares Pages per Second and % Total Disk Time against Sessions. This graph correlates the impact of virtual memory paging on available system I/O. This graph is capped at 100% Total Disk Time and 100 Pages per Second to show the necessary granularity. Pages per Second rates of 40 or greater can cause system delays. **In multi-user environment, such as Citrix MetaFrame, this number can be as low as 10 when its impact occurs.** This is because of the random nature of access by many users on the same system can read or write data to varying locations within the pagefile as well as the random nature of application access. **On some servers moving the pagefile to a TiGiJet can allow the same system to accommodate up to 4 times the number of users.**

Pages/sec: Pages/sec is the number of pages read from or written to disk to resolve hard page faults. (Hard page faults occur when a process requires code or data that is not in its working set or elsewhere in physical memory, and must be retrieved from disk, virtual memory). This counter was designed as a primary indicator of the kinds of faults that cause system-wide delays.

% Total Disk Time: % Disk Time is the percentage of elapsed time that the selected disk drive is busy servicing read or write requests.

Sessions: Number of active Terminal Services sessions or ICA users.

Impact Threshold shows the peak level of Pages Per Second for optimal server performance.

Figure 5: % Total Processor & Sessions

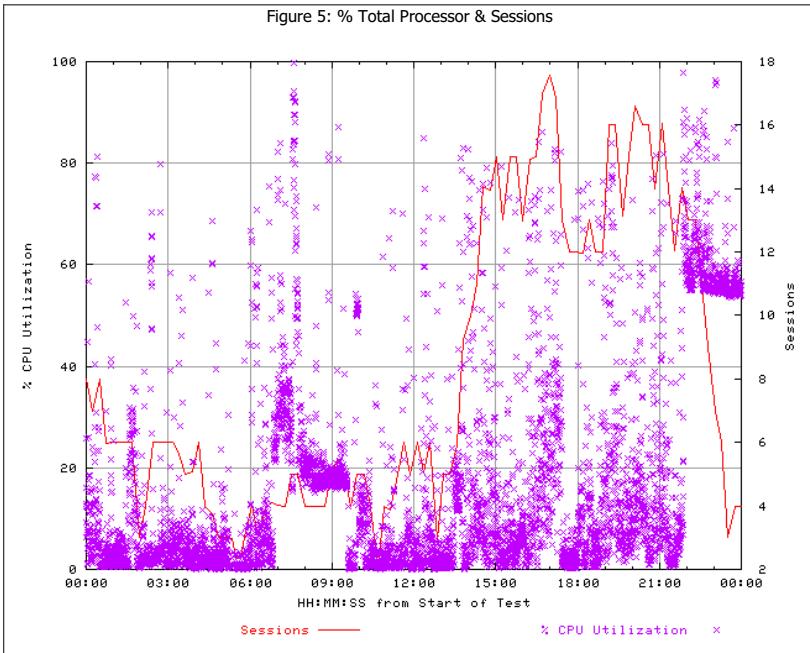


Figure 5 compares % Total Processor Time against Sessions. This graph demonstrates the impact of the number of sessions on CPU utilization. This graph also shows the impact that logon and logoff's may have on a system's CPU. **Most multi-user environments, such as Citrix MetaFrame, try to maintain an efficient % of Total Processor Time less than 80% to accommodate for peaks in application requirements or users.**

% Total Processor Time: % Processor Time is the percentage of time that the processor is executing a non-Idle thread. This counter was designed as a primary indicator of processor activity. It is calculated by measuring the time that the processor spends executing the thread of the Idle process in each sample interval, and subtracting that value from 100%. (Each processor has an Idle thread, which consumes cycles when no other threads are ready to run)..

Sessions: Number of active Terminal Services sessions or ICA users.

Figure 6: % Total Processor, Privileged Time, & Sessions

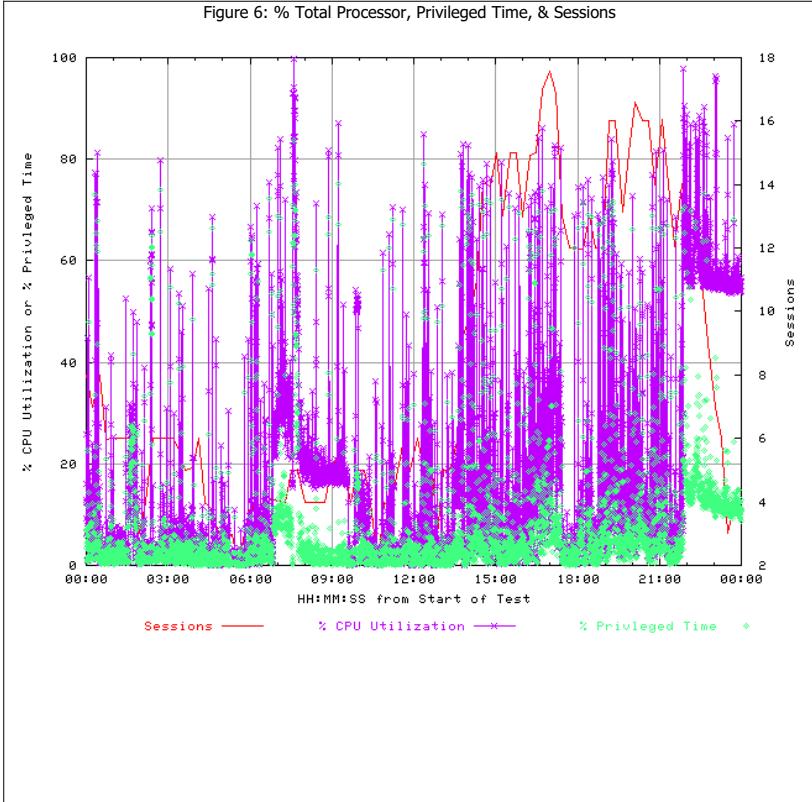


Figure 6 examines CPU utilization one step further by comparing % Total Processor Time and % Total Privileged Time against Sessions. This graph shows the percentage of CPU resources which are being "wasted" doing I/O management and could be used for end-user application functions. **Most multi-user environments are at their peak efficiency when % Total Privileged time is less than 10% of the % Total Processor Time.** (Please refer to Appendix A for the exact comparison). With the TiGiJet, the system will spend fewer CPU cycles in privileged time freeing up additional CPU resources to accommodate more users.

% Total Processor Time: % Processor Time is the percentage of time that the processor is executing a non-Idle thread. This counter was designed as a primary indicator of processor activity. It is calculated by measuring the time that the processor spends executing the thread of the Idle process in each sample interval, and subtracting that value from 100%. (Each processor has an Idle thread, which consumes cycles when no other threads are ready to run)..

% Total Privileged Time: % Total Privileged Time shows the percentage of non-idle processor time spent executing code in privileged mode. Privileged mode is a processing mode designed for operating system components and hardware-manipulating drivers. It allows direct access to hardware and memory. The alternative, user mode, is a restricted processing mode designed for applications, environmental subsystems, and integral subsystems. The operating system switches application threads to privileged mode to access operating system services. Typically the largest consumer of privileged time is the IO subsystem for IO wait state management.

Sessions: Number of active Terminal Services sessions or ICA users.

Appendix A: Summary of Captured Counters

Statistical Measurements

The following statistical breakdowns are provided in Appendix A for all of the captured counters;

- Min – The minimum measurement over the entire sample period.
- Max – The maximum measurement over the entire sample period.
- Sum – The sum of all of the samples. This measurement is useful with counters such as Current Disk Queue. The sum will show how many total Disk Queues occurred over the sample period.
- Avg – Average is the statistical Mean of the entire sample period.
- 75th – The 75th percentile is the smallest number that is greater than 75% of the numbers in a given set. The reason this statistic is so useful in measuring data throughput is that it gives a comparative picture of the specific counters against the Avg and 95th percentile.
- 95th – The 95th percentile is the smallest number that is greater than 95% of the numbers in a given set. The reason this statistic is so useful in measuring data throughput is that it gives a very accurate picture of the specific counter's impact on system performance by eliminating data peaks and offering a point of reference for managing a system's maximum capacity to. Another important reason for using 95th percentile is to compare against the Max measurement to identify whether the Max entry was an anomaly or is descriptive of actual system performance.

Summary Data

Appendix B: Definition of Counters

Counter Name	Counter Description
Processor: % Processor Time_Total Instance	% Processor Time is the percentage of time that the processor is executing a non-idle thread. This counter was designed as a primary indicator of processor activity. It is calculated by measuring the time that the processor spends executing the thread of the Idle process in each sample interval, and subtracting that value from 100%. (Each processor has an Idle thread, which consumes cycles when no other threads are ready to run).
System: Processor Queue Length	Processor Queue Length is the number of threads in the processor queue. There is a single queue for processor time even on computers with multiple processors. Unlike the disk counters, this counter counts ready threads only.
System: % Total Privileged Time	% Total Privileged Time shows the percentage of non-idle processor time spent executing code in privileged mode. Privileged mode is a processing mode designed for operating system components and hardware-manipulating drivers. It allows direct access to hardware and memory. The alternative, user mode, is a restricted processing mode designed for applications, environmental subsystems, and integral subsystems. The operating system switches application threads to privileged mode to access operating system services.
Terminal Services: Active Sessions	Number of active Terminal Services sessions.
Memory: Available Bytes	Available Bytes is the amount of physical memory available to processes running on the computer, in bytes.
Memory: Pages/Sec	Pages/sec is the number of pages read from or written to disk to resolve hard page faults. (Hard page faults occur when a process requires code or data that is not in its working set or elsewhere in physical memory, and must be retrieved from disk). This counter was designed as a primary indicator of the kinds of faults that cause system-wide delays.
Memory: Pages Reads/sec	Page Reads/sec is the number of times the disk was read to resolve hard page faults. (Hard page faults occur when a process requires code or data that is not in its working set or elsewhere in physical memory, and must be retrieved from disk).
Memory: Page Writes/sec	Page Writes/sec is the number of times pages were written to disk to free up space in physical memory.
Memory: Paging File %Usage	Percentage of page file utilized by computer.
Memory: Pool Paged Bytes	Pool Paged Bytes is the number of bytes in 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 Nonpaged Bytes	Pool Nonpaged Bytes is the number of bytes in 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.
Physical Disk: % Disk Read Time	The % Disk Read Time counter is the percentage of elapsed time that the selected disk drive is busy servicing read requests.
Physical Disk: % Disk Time	% Disk Time is the percentage of elapsed time that the selected disk drive is busy servicing read or write requests.
Physical Disk:	% Disk Write Time is the percentage of elapsed time that the selected

% Disk Write Time	disk drive is busy servicing write requests.
Physical Disk: % Idle Time	% Idle Time reports the percentage of time during the sample interval that the disk was idle.
Physical Disk: Avg. Disk Bytes/Read	Avg. Disk Bytes/Read is the average number of bytes transferred from the disk during read operations.
Physical Disk: Avg. Disk Bytes/Transfer	Avg. Disk Bytes/Transfer is the average number of bytes transferred to or from the disk during write or read operations.
Physical Disk: Avg. Disk Bytes/Write	Avg. Disk Bytes/Write is the average number of bytes transferred to the disk during write operations.
Physical Disk: Avg. Disk Queue Length	Avg. Disk Queue Length is the average number of both read and write requests that were queued for the selected disk during the sample interval.
Physical Disk: Avg. Disk Read Queue Length	Avg. Disk Read Queue Length is the average number of read requests that were queued for the selected disk during the sample interval.
Physical Disk: Avg. Disk sec/Read	Avg. Disk sec/Read is the average time in seconds of a read of data from the disk.
Physical Disk: Avg. Disk sec/Transfer	Avg. Disk sec/Transfer is the time in seconds of the average disk transfer.
Physical Disk: Avg. Disk sec/Write	Avg. Disk sec/Write is the average time in seconds of a write of data to the disk.
Physical Disk: Avg. Disk Write Queue Length	Avg. Disk Write Queue Length is the average numbers of write requests that were queued for the selected disk during the sample interval.

Appendix C: Links for Additional Data

- Performance Counters Reference for Windows Server 2003;

http://www.microsoft.com/technet/prodtechnol/windowsserver2003/proddocs/deployguide/counters_overview.asp