



## **GO-Global *For Windows* Deployment Guide**

Considerations, Sizing and Scalability Guidelines

# Introduction

GraphOn's GO-Global software gives the Windows Server operating system the capability to serve 32-bit Windows® based applications to client desktops, terminals and web-browsers running on PC and non-PC desktops. The GO-Global Server environment is, by definition, a thin-client, 100 percent server-centric architecture. Because the GO-Global clients will be available for many different desktop platforms (Linux, Macintosh, and others), the GO-Global Server provides access to 32-bit Windows-based applications from virtually any desktop. The GO-Global Server allows you to roll out 32-bit Windows-based applications to a heterogeneous set of desktops while transitioning to a pure 32-bit desktop environment.

Unlike the traditional client/server environment, an application runs only on the server in the GO-Global Server environment. The GO-Global client performs virtually no local processing of applications. The GO-Global Server transmits only the application presentation—the Graphical User Interface, or GUI—down to the client. Each user logs on and perceives only his or her own *session*, which is transparently managed by the server operating system and is independent from any other client session.

From an application development perspective, one of the biggest benefits of the GO-Global Server is that the majority of existing 32-bit application programs run "as is"—no programming changes are required to run them on a GO-Global Server. This does not, however, mean that all existing 32-bit applications run equally well under GO-Global Server. As with the addition of any new operating system service, understanding how to design applications that take advantage of these new capabilities is important. It's also important to understand how this new service can magnify the negative results of bad programming habits, and to point out specific areas to watch for.

## Application Design Considerations

An important difference between the standard Windows server environment and the GO-Global Server environment is the design of applications that run in these environments. In a centralized host environment, applications are designed specifically to run in that environment. Utilizing GO-Global for Windows, applications designed for a Windows-based PC work without having to be explicitly developed for the GO-Global Server environment. Applications that run on Windows NT today will run on the GO-Global Server without modification.

This is important when you consider the implications of multiple users sharing a common Windows NT system concurrently. Instead of having different users running applications using their own local hardware resources (such as CPU, memory, and disk), and local software resources (such as registry settings, preference files, and dynamic link libraries), users in the GO-Global Server environment share the same hardware and software resources. For example, if two users run the same application in a GO-Global Server environment, two copies of that application are started in the same system, each operating under a different user context. All of this is managed transparently by the operating system.

Having multiple users access the same set of applications in a common system can create a number of areas of contention, however. These areas include:

**Users contend for CPU time.** In the GO-Global Server environment, each user has his or her own application environment and can run whatever applications are available to that user. However, all applications run by all users are contending for the central CPU resources available on the server machine. If one user runs a poorly written, CPU-intensive application, other users on that server may potentially experience a very visible loss of performance.

**Users contend for disk access.** This is analogous to how users contend for disk access using the traditional client/server network connections. In the GO-Global Server environment, the input/output demands are more intense because users not only contend for access to applications and related application files, but also for server operating system disk accesses. For example, multiple users may be making different dynamic link library (DLL) calls at the same time or swapping between real and virtual memory areas.

**Users contend for Random Access Memory (RAM).** Each user has an

independent session, which he or she can attempt to fill with all the memory-intensive applications at his or her disposal. Some users may try to open as many applications as they can, while others take a more conservative approach and run only those applications that they need. Nonetheless, the needs of all the users are satisfied from the same core set of server memory resources.

**Users contend for network access.** As in any distributed processing environment, the network provides the pipeline for communication between the desktop and the servers. In the GO-Global Server environment, the need for network access is more critical than in the traditional distributed client/server architecture because all desktop activity—graphical output and mouse/keyboard input—flows over the network links between the desktop and the server. Without a functioning network connection to the server, a GO-Global client cannot function at all.

**Users contend for server-side hardware components.** Hardware components such as CD-ROMs, disk drives, serial ports, and parallel ports are typically server-based, not client-based. Therefore, when a user goes to access server-side hardware components, he or she may be competing with other users making similar requests. Sharing these traditionally non-shared components creates new considerations for both the users and for the applications that access these components.

**Users contend for access to global Windows objects and resources.** In the GO-Global Server environment, users do not run individual copies of Windows NT—some of the core components are cloned, but the remaining components are shared among the users. Thus users are competing for access to the registry, the paging file, system services, and other global objects and resources.

Many of these points of contention can be mitigated by sizing the GO-Global Server system with sufficient CPU, memory, and disk resources to handle the client demand. For example, a multiple processor configuration can maximize CPU availability. Memory availability can be maximized by installing extra physical memory. Finally, disk access performance can be maximized by configuring multiple SCSI channels and distributing your operating system and application loads across different physical drives. Properly configuring a GO-Global Server is a critical element of perceived application performance

## **Significant Performance Factors**

### *Know Your Application*

Knowing your application may require asking the following questions: How graphics-intensive is the application? How much memory does the application require? How much of the application's memory is shareable between users? How does the

application refresh the screen? Does the application require a lot of typing? What is acceptable performance?

### *Know Your Users*

Knowing your user may require asking the following questions: How fast do your users type? Do your users leave the applications open? When do users log on? Do the users stay logged on throughout the day? Are logons evenly distributed throughout the day or are most logons at a specific time of the day?

### *Know Your Network*

Because GO-Global provides application sessions to the user, network communications are very important. Even infrequent network slowdowns can provide unacceptable performance to your users. Knowing your network may include the following questions: How fast do your users type? Are the applications graphics-intensive? What is the typical display resolution of your client workstations? What is the network bandwidth required for a user running the application?

## **GO-Global Server Hardware Requirements**

The following guidelines are based on the minimal requirements of the operating systems on which you run GO-Global for Windows.

### **Windows NT 4.0 Server and Workstation** with Service Pack 6.

Japanese Windows NT 4.0 with Service Pack 6

Simplified Chinese Windows NT 4.0 with Service Pack 6

GraphOn recommends a dual 500 MHz Pentium or better microprocessor, 128 MB of RAM, and a hard disk with at least 100 MB of free space.

### **Windows 2000 Server and Advanced Server** with Service Pack 3.

Japanese Windows 2000 with Service Pack 2

German Windows 2000 with Service Pack 2

French Windows 2000 with Service Pack 2

Traditional Chinese Windows 2000 with Service Pack 2

Simplified Chinese Windows 2000 with Service Pack 2

GraphOn recommends a dual 500 MHz 500 or faster Pentium-compatible processor, 128MB of RAM, and a hard disk with at least 100 MB of free space.

### **Disk and Memory Requirements**

In addition to the Windows operating system requirements for your server, the following are required for GO-Global for Windows:

- 48 MB disk space for standard GO-Global for Windows installation, including the GO-Global Cluster Manager
- 5 MB disk space for installing all GO-Global client software
- 10 MB for web related HTML files

### **Networks and Modems**

In addition to connections over LAN or WAN based networks, GO-Global for Windows supports dial-up networks utilizing modem as well as VPN connections. TCP/IP connectivity is a minimum requirement.

### **Processor, Bus, and Memory**

The processor and bus architecture are fundamental to GO-Global for Windows server performance.

The ISA (AT bus) architecture is low-bandwidth and is not recommended for GO-Global for Windows servers. Use a higher-performance bus, such as EISA or PCI, for best performance. These buses support the sustained data transfer rates of a GO-Global for Windows server.

The memory (RAM) requirement for GO-Global for Windows is 16 MB plus 4 MB for each typical user or 8 MB for each power user. In many cases, adding RAM has a larger effect on system performance than upgrading to a faster processor. In general, processor and RAM requirements for GO-Global for Windows scale linearly. You can roughly double the number of users supported on a multiprocessor-capable system by doubling the number of processors and doubling the amount of RAM. By purchasing multi-processor capable systems (even if you initially purchase only one processor), you provide for convenient system scaling as your requirements grow.

### **Win16 Applications**

Windows NT and Windows 2000 are Win32 (32-bit) environments. Windows 3.x and DOS are part of a Win16 (16-bit) environment. GO-Global for Windows only supports 32 bit applications.

### **Hard Disks**

The hard disk subsystem in a server is an important factor in system throughput. Small Computer System Interface (SCSI) disk drives and adapters, especially Fast Narrow SCSI (SCSI-2), Fast Wide SCSI, Wide Ultra SCSI, and Wide Ultra2 SCSI devices, have significantly better throughput than ST-506, Integrated Device Electronics (IDE), or Enhanced Small Device Interface (ESDI) disk drives and adapters.

For the highest disk performance, consider using a SCSI-based Redundant Array of Independent Disks (RAID) controller. RAID controllers automatically place data on multiple disk drives and can increase disk performance and improve data reliability. Use NTFS for all disk partitions on your GO-Global for Windows servers. NTFS allows security configuration, better performance, and more fault tolerance.

### **Network Interfaces**

The RXP protocol is highly compressed and causes negligible loading on a network, but because the GO-Global for Windows server handles all network requests, a high performance network interface card (NIC) is recommended.

### **Using Performance Monitoring Tools**

GraphOn recommends that you use performance monitoring tools to get accurate

accounts of system performance and the effects of configuration changes on system throughput. The most important measurements for performance monitoring are the percentage of total processor time, memory pages per second, percentage of network utilization, and hard disk I/O rates.

A good way to estimate how many users a server can support is to measure system performance with two to five users on the system and then scale the results. This method has been found to yield reliable results.

### **GO-Global Cluster Manager**

GraphOn's Cluster Manager is a centralized management and monitoring utility. It is used by system administrators to monitor or manage user sessions, server processes, and to register and manage applications on the GO-Global server. The Cluster Manager allows for remote management of all GO-Global servers from a central location.

## Sizing GO-Global for Windows Servers (basic)

GO-Global for Windows supports multiple users on Windows NT 4.0 Server, and Windows 2000 Servers. A multi-user system requires more system resources than a single-user system. This section provides some system sizing guidelines that can help you decide on a hardware configuration that will support your users with optimal performance. Most companies find that their users can be categorized as typical users or power users.

*Typical user:* Generally uses one or two applications but normally only one at a time. Little actual program data is transferred between the client and server, and the users rarely use Object Linking and Embedding (OLE).

*Power user:* A more sophisticated user who uses three or more applications, often with several active at the same time. Data is often cut and pasted between local and remote applications, and OLE is used heavily. Power users consume more resources than typical users. A good rule of thumb is that one power user is equivalent to two typical users in processor utilization and RAM requirements.

### Memory Calculation Example

A basic GO-Global server allocates memory as follows:

Service	Memory allocation	Program
Application Publishing Service (one instance per server)	8 MB	aps.exe
Program Window (one instance per user)	8 MB	pw.exe
License Service (one instance per server)	4 MB	lmgrd.exe
License interface (one instance per user)	3 MB	blm.exe

**Note:** It is important to realize that any subsequent application memory is in addition to the basic GO-Global server configuration.

### Basic allocation example

In the following example we calculate the total GO-Global server consumption plus any subsequent user allocation based upon such things as log-on services and licensing. A sample application such as Microsoft Word is factored in to the final equation. If we utilize an application such as Word, the additional allocated RAM for an average document will be ~12 MB per user.

<b>Program</b>	<b>MB</b>		
Aps.exe	8	→	
Lmgrd.exe	4		
<b>Total server consumption</b>	<b>12</b>		<b>12 MB</b>
Pw.exe	8	* <b>Number of users</b> (15)	
Blm.exe	3		
Winword.exe	12		
<b>Total client consumption</b>	<b>23</b>		<b>345 MB</b>

The example above shows that 345 MB of RAM would be required to support 15 users utilizing Word for Windows. These numbers are approximate and are subject to change depending on specific system and hardware configuration. Note that these numbers reflect only the GO-Global system services and do not take in to account other applications or system services, such as a Web server. Standard system overhead is in addition to any GO-Global services.

## Sizing GO-Global for Windows Servers *(detailed)*

**Issue:** How to determine the maximum number of users that can be supported on a GO-Global Server

The number of users that can be supported from a GO-Global Server is influenced by many factors, including:

- The speed and number of processors
- The amount of physical memory
- Memory requirements of the application(s)
- Processor utilization requirements of the application(s)
- Kernel memory requirements of hardware devices (for example, RAID hard disk controllers, video cards, etc.)

A combination of these factors will generally determine the number of users that can effectively run from a GO-Global Server. This article describes how to evaluate the scalability impact of these factors and estimate the maximum number of users that can be supported from a GO-Global Server. Administrators can limit the number of users that can log on to a server through the Cluster Manager. (Tools | Sessions | Options...).

### Resolution

The number of users that can be supported on a given GO-Global Server is generally limited by one or more of the following system resources:

- **Processor (CPU) Utilization:** The percentage of time that the server's processors are executing non-idle threads.
- **Physical Memory:** The amount of physical memory used by processes running on the computer.
- **Committed Memory:** The amount of memory for which space has been reserved in the disk paging file.
- **Paged Pool Memory:** The amount of memory allocated within the Paged Pool, an area of kernel-mode memory used to store data for the operating system and other kernel-mode modules.
- **System Page Table Entries (PTEs):** The amount of space consumed within the System Page Table Entry (PTE) area, an area of kernel-mode memory that is used by the operating system to store kernel stacks, and by device drivers (e.g., RAID hard disk controllers, video cards, etc.) to map views of device-specific data into the kernel address space.

#### To determine the number of users that can run effectively from a given server

1. Determine how much of each of the above resources is required for an average user.
2. Determine how much of each resource is initially available on the server computer.

3. Calculate the user limit using the information from steps 1 and 2.
4. Verify the calculated result in a production environment.

The details of this process are described below:

## 1. Determining User Resource Requirements

User resource requirements are generally determined by the applications that users will run on the server. If an application supplier has tested an application with GO-Global, they may be able to provide per-user resource requirements for the application. Otherwise, the best way to determine user requirements is to run a multi-user test that simulates the operations that typical users will perform on the server.

### To determine the per-user resource requirements for a given set of applications

1. Restart the server computer.
2. Begin logging system resources using the Windows Performance Monitor.
  - a. On a separate Windows 2000 computer, start the Windows Performance Monitor.
  - b. Right-click **\Performance Logs and Alerts\Counter Logs**.
  - c. Choose **New Log Settings**.
  - d. Type a name for the log.
  - e. Click **Add**.
  - f. Enter the name of the server computer in the **Select counters from computer** field.
  - g. Add the following counters:
    - i. \Processor\% Processor Time
    - ii. \Memory\Available Bytes
    - iii. \Memory\Committed Bytes
    - iv. \Memory\Pool Paged Bytes
    - v. \Memory\Free System Page Table Entries
  - h. Set the **Sample Interval** to 10 seconds.
  - i. Select the **Log Files** tab.
  - j. Change the **Log file type** to Text File – CSV.
  - k. Click **OK**.
  - l. Right-click the icon for the new log and choose **Start**.
3. Log 10 users on to the GO-Global Server. Each user should be running on a separate client computer.
4. In each session, start all the programs that you expect users to run during a typical session. In each session, simulate typical user interaction with the applications. This may be done with live users driving each session, or

with a script that is recorded and played back on each client using a third-party automated testing tool.

5. Run the test for 15-30 minutes.
6. Stop the performance monitor log.
  - a. Right-click the icon for the log and choose **Stop**.
7. Open the log file in a spreadsheet program such as Microsoft Excel and optionally chart the results.
8. Determine the per-user resource requirements from the data as follows:

- a. The average processor utilization per user:

$$\text{AverageProcessorUtilizationPerUser} = \frac{(\text{AveragePercentProcessorTime} * \text{NumberOfProcessors})}{10}$$

where *AveragePercentProcessorTime* is the average value of the \Processor\% Processor Time \\_Total counter during the period when typical user interaction was simulated, and *NumberOfProcessors* is the number of processors installed in the computer.

- b. The physical memory required per user:

$$\text{PhysicalMemoryPerUser} = \frac{(\text{InitialAvailableBytes} - \text{MinimumAvailableBytes})}{10}$$

where *InitialAvailableBytes* and *MinimumAvailableBytes* are respectively the initial and minimum values of the \Memory\Available Bytes counter.

- c. The committed memory required per user:

$$\text{CommittedMemoryPerUser} = \frac{(\text{MaximumCommittedBytes} - \text{InitialCommittedBytes})}{10}$$

where *MaximumCommittedBytes* and *InitialCommittedBytes* are respectively the maximum and initial values of the \Memory\Committed Bytes counter.

- d. The Paged Pool memory required per user:

$$\text{PagedPoolMemoryPerUser} = \frac{(\text{MaximumPagedPoolBytes} - \text{InitialPagedPoolBytes})}{10}$$

where *MaximumPagedPoolBytes* and *InitialPagedPoolBytes* are respectively the maximum and initial values of the \Memory\Pool Paged Bytes counter.

- e. The number of System PTEs required per user:

$$\text{SystemPTEsPerUser} = \frac{(\text{InitialFreeSystemPTEs} - \text{MinimumFreeSystemPTEs})}{10}$$

where *InitialFreeSystemPTEs* and *MinimumFreeSystemPTEs* are respectively the initial and minimum values of the \Memory\Free System Page Table Entries counter.

## 2. Determining Available Resources

Once the user requirements for the desired set of applications is known, the next step is to measure the available resources of the computer(s) that will be used to run the applications. This may be done as follows:

1. If it is not already installed, install GO-Global on the server computer.
2. Restart the computer.
3. Obtain the initial values of the following performance counters using the Windows Performance Monitor:
  - a. \Memory\Available Bytes
  - b. \Memory\Committed Bytes
  - c. \Memory\Commit Limit
  - d. \Memory\Pool Paged Bytes
  - e. \Memory\Free System Page Table Entries

4. Use this data to calculate the following available resources for the computer:

- a. The available processing capacity:

$$\text{AvailableProcessorUtilization} = \frac{(\text{NumberOfProcessors} * \text{ProcessorSpeed} * \text{MaximumPercentUtilizationAllowed})}{\text{BaseProcessorSpeed}}$$

where *NumberOfProcessors* is the number of processors installed in the server, *ProcessorSpeed* is the speed in MHz of the server's processors, and *BaseProcessorSpeed* is the speed of the processor used to determine the *AverageProcessorUtilizationPerUser* in step 1.1.a above. *MaximumPercentUtilization* is the maximum percentage of processor utilization that should occur during normal operation. The recommended value for *MaximumPercentUtilization* is 0.80 (80%).

- b. The available physical memory:

$$\text{AvailablePhysicalMemory} = \text{InitialAvailableBytes} - \text{MinimumAvailableBytesAllowed}$$

where *InitialAvailableBytes* is the initial value of the \Memory\Available Bytes counter, and *MinimumAvailableBytesAllowed* is the minimum number of bytes of physical memory that should be available during normal operation. The recommended value for *MinimumAvailableBytesAllowed* is 41943040 bytes (40 MB).

- c. The available committed memory (i.e., available space in the disk paging file):

$AvailableCommittedMemory = CommitLimit - InitialCommittedBytes -$

$MinimumAvailableCommittedBytesAllowed$

where *CommitLimit* and *InitialCommittedBytes* are respectively the values of the \Memory\Commit Limit and \Memory\Committed Bytes counters, and *MinimumAvailableCommittedBytesAllowed* is the minimum number of bytes of committed memory that should be available during normal operation. The recommended value for *MinimumAvailableCommittedBytesAllowed* is 0.20 times the Commit Limit.

- d. The available Paged Pool memory:

$AvailablePagedPoolMemory = SizeOfPagedPool - InitialPagedPoolBytes -$

$MinimumAvailablePagedPoolBytesAllowed$

where *SizeOfPagedPool* is the size of the paged pool in bytes (generally 0x0A000000 (167772160) bytes on Windows 2000 Servers) and *InitialPagedPoolBytes* is the initial value of the \Memory\Pool Paged Bytes counter.

*MinimumAvailablePagedPoolBytesAllowed* is the minimum number of paged pool bytes that should be available during normal operation. The recommended value for *MinimumAvailablePagedPoolBytesAllowed* is 1048576 bytes (10 MB).

The actual size of the paged pool can be determined using a kernel debugger as follows:

**Note:**

1. Break into the debugger
2. Type !vm, and then press ENTER.
3. Information about the computer's virtual memory will be displayed. The size of the Paged Pool will be listed as *PagedPool Maximum*.

For information on how to use a kernel debugger, refer to article 151981 in the Microsoft Knowledge Base.

- e. The available System PTEs:

$AvailableSystemPTEs = InitialFreeSystemPTEs -$

$MinimumAvailableSystemPTEsAllowed$

where *InitialFreeSystemPTEs* is the value of the \Memory\Free System Page Table Entries counter, and *MinimumAvailableSystemPTEsAllowed* is the minimum number of SystemPTEs that should be available during normal operation. The recommended value for *MinimumAvailableSystemPTEsAllowed* is 4000.

### 3. Calculating the User Limit

Using the values obtained in steps 1 and 2, a user limit for the computer may be estimated as follows:

1. Calculate the maximum number of users that can be supported by each of the computer's resources:

- a. Calculate the maximum number of users that can be supported by the computer's processors:

$$\text{ProcessorUserLimit} =$$

$$\frac{\text{AvailableProcessorUtilization}}{\text{AverageProcessorUtilizationPerUser}}$$

- b. Calculate the maximum number of users that can be supported by the computer's available physical memory:

$$\text{PhysicalMemoryUserLimit} = \frac{\text{AvailablePhysicalMemory}}{\text{PhysicalMemoryPerUser}}$$

- c. Calculate the maximum number of users that can be supported by the computer's available committed memory:

$$\text{CommittedMemoryUserLimit} =$$

$$\frac{\text{AvailableCommittedMemory}}{\text{CommittedMemoryPerUser}}$$

- d. Calculate the maximum number of users that can be supported by the computer's available paged pool memory:

$$\text{PagePoolMemoryUserLimit} = \frac{\text{AvailablePagedPoolMemory}}{\text{PagedPoolMemoryPerUser}}$$

- e. Calculate the maximum number of users that can be supported by the computer's available System PTEs:

$$\text{SystemPTEUserLimit} = \frac{\text{AvailableSystemPTEs}}{\text{SystemPTEsPerUser}}$$

2. Calculate the maximum number of users for the computer using the following formula:

$$\text{MaximumNumberOfUsers} = \text{Minimum} ( \text{ProcessorUserLimit}, \\ \text{PhysicalMemoryUserLimit}, \\ \text{CommittedMemoryUserLimit}, \\ \text{PagePoolMemoryUserLimit}, \\ \text{SystemPTEUserLimit} )$$

3. If the maximum number of users is lower than required, attempt to modify the configuration of the computer to increase the availability of the limiting resource.

- a. If processor utilization is the limiting resource, increase the number of servers, or use servers with faster and/or a larger number of processors.
- b. If physical memory is the limiting resource, install additional memory in the computer. The maximum amount of memory that may be installed in a Windows 2000 Server is 4 GB.
- c. If committed memory is the limiting resource, increase the size of the disk paging file:
  - i. Right click on **My Computer** and choose **Properties**.
  - ii. Choose **Advanced | Performance Options | Change**.
  - iii. Increase the maximum size of the paging file.
  - iv. Choose **Set**.

**Warnings:**

? The recommended size of the disk paging file will generally be less than or equal to twice the amount of physical memory installed in the machine. Use care when increasing the size of the paging file above the recommended size. If the amount of committed memory is allowed to rise far above the amount of physical memory, the system may begin performing an excessive number of disk I/O operations to retrieve data from paging file. This can seriously degrade system performance.

? The following information involves opening and manipulating the Windows Registry. Carrying out operations other than those described here may cause configuration errors, possibly rendering your system unusable. Please use extreme caution any time you work in the Registry.

- d. If Paged Pool memory is the limiting resource, increase the size of the Paged Pool:
  - i. Choose **Start | Run**.
  - ii. Type **regedit, ENTER**.
  - iii. Select the following registry value:  
**\\HKEY\_LOCAL\_MACHINE\\SYSTEM\\CurrentControlSet  
\\Control  
\\Session Manager\\Memory  
Management\\PagedPoolSize**
  - iv. Choose **Edit | Modify**.
  - v. Choose **Hexidecimal**
  - vi. Type **FFFFFFFF** in the **Value data** field.
  - vii. Click **OK**.
  - viii. Restart the computer.

**Warning:**

Paged Pool memory and System PTEs share an area of virtual kernel-mode memory, and increasing the amount of Paged Pool Memory will decrease the number of System PTEs. The above change will add approximately 200 MB of memory to the Paged Pool, but it will reduce the number of Free System Page Table Entries by 50,000.

- e. If the number of System PTEs is the limiting resource:
  - i. Ensure that the system is configured to maximize the number of System PTEs:
    1. Double-click the following registry value:  
**\\HKEY\_LOCAL\_MACHINE\\SYSTEM\\CurrentControlSet\\Control\\Session Manager\\Memory Management\\SystemPages**
    2. Choose **Hexadecimal**
    3. Check the value. It should be FFFFFFFF. The GO-Global Setup Program sets this field to this value when the product is installed. If the value has been subsequently modified, reset it to FFFFFFFF.
    4. Check the value of the PagedPoolSize value (also located under the Memory Management key). It should be zero (0).
  - ii. If the system is configured to maximize the number of System PTEs, review the requirements of hardware devices installed in the computer. Some devices, such as RAID disk controllers and video cards can consume large numbers of System PTEs. If such a device is installed in the server, ensure that it is required, and, if it is not, uninstall the device.
4. If any configuration modifications are made in step 3.3 above, recalculate the affected available resource, the resource user limit, and the maximum user limit for the server computer.

#### 4. Verifying the Result

The result of the procedure outlined in steps 1–3 is an estimate of the maximum number of GO-Global users that can be supported on a given server. Several assumptions are made to arrive at this result. For example, it is assumed that the operations performed during the multi-user resource requirements test are representative of the operations that real users will perform when using the product in a production environment. It is also assumed that the number of users will be limited by one of the system resources measured in the test. To ensure the accuracy of these assumptions and the final result, the server configuration should be tested in a production environment, generally as part of a pilot deployment to the number of users calculated above. During this deployment, system resources should be monitored using the Windows Performance Monitor, and users should be asked to report the nature and time of any problems that they observe. In

addition to the performance counters for the resources mentioned in this article, the following performance counters should also be monitored:

- **\System\Processor Queue Length:** The Processor Queue Length is the number of active threads that are available but unable to execute due to another active thread. One of the clearest indicators of a processor bottleneck is a sustained processor queue length greater than 1-3 threads per processor. High processor queue lengths generally occur when processor utilization is high, but they can develop when processor utilization is well below 90%.
- **\Memory\Pages/sec:** The Pages/sec performance counter measures the number of pages read from or written to disk to resolve hard page faults. Hard page faults occur when memory must be read from or written to the disk paging file. System-wide performance delays can occur when the number of hard page faults is high, even when CPU usage is low.

If problems are reported:

1. Attempt to identify the resource that caused the problem by reviewing the data from the Windows Performance log at the time the problem occurred:
  - a. If the user reported sluggish performance, check:
    - i. **% Processor Utilization:** If greater than 90%, the computer's processors may be the bottleneck.
    - ii. **Processor Queue Length:** If greater than 1-3 threads per processor for a sustained length of time, the computer's processors may be the bottleneck.
    - iii. **Available Bytes:** If this value is low (e.g., less than 10% of the installed physical memory), check the Pages/sec counter. If the Pages/sec value is significantly higher than normal, the system is frequently swapping memory to and from disk. This problem is generally caused by having insufficient physical memory.
  - b. If the user reported an inability to connect to the server or reported errors while running an application, check:
    - i. **Available Bytes:** GO-Global prevents new users from connecting to the server if Available Physical Memory is less than 20 MB.
    - ii. **Committed Bytes:** If the value of this counter is near the Commit Limit, random errors may occur in all sessions. Also, high values for Committed Memory often coincide with high disk paging (Pages/second) values and sluggish performance.
    - iii. **Paged Pool Bytes:** If near the maximum value (*0x0A000000 (167772160)* bytes on Windows 2000 Servers), the system may be running out of paged pool memory. If this is the case, random errors will occur in all running sessions.

- iv. **Free System Page Table Entries:** GO-Global prevents new users from connecting to the server if the number of free System PTEs is less than 2000. If the number of free System PTEs drops below 1000, random errors can occur in all running sessions.
2. Once you have identified the resource that caused the problem:
- a. Attempt to increase the availability of the resource as described in step 3.3,  
or
  - b. Reduce the maximum number of users for the specified resource to a safe value.

# GO-Global Server Deployment Checklist

- \_\_\_ Configure hardware
- \_\_\_ Configure Operating System
- \_\_\_ Add international keyboard support (as needed)
- \_\_\_ Create/modify NT user & server policies
- \_\_\_ Attach server to the network
- \_\_\_ Add DNS records for the server
- \_\_\_ Install Web server
- \_\_\_ Install application(s) (reboot typically required)
- \_\_\_ Configure application(s) (as needed)
- \_\_\_ Install GO-Global for Windows software (reboot required)
- \_\_\_ Login with a user account that has administrative privileges
- \_\_\_ Execute Cluster Manager from the desktop icon
- \_\_\_ Add applications to Cluster Manager
- \_\_\_ Add server options to Cluster Manager (as needed)
  - \_\_\_ Printer drivers
  - \_\_\_ Data Encryption
  - \_\_\_ Client Drives
  - \_\_\_ Client Sound
  - \_\_\_ Client Clipboard
  
- \_\_\_ Secure the server and applications
- \_\_\_ Change server performance options to Applications (Windows 2000 only)
- \_\_\_ Deploy client software

## **Resources and Additional Information**

For additional information on server sizing, please refer to the following:

- Thin Client Server Computing Sizer from HP (Compaq)  
<http://activeanswers.compaq.com/sizers/TSESizer/TSESizer.html>

Some information contained in this document is courtesy of Microsoft and their guidelines as it related to the Windows® Operating System