


## Monitoring NT Performance

Jerry L. Rosenberg  
SRM Associates, Ltd.  
CMG 2002

e-mail: [Jerry@SRMAssoc.com](mailto:Jerry@SRMAssoc.com)  
<http://www.srmassoc.com>


NT Performance Notebook 1



## Capacity and Performance

- Performance tuning
  - Identify bottlenecks and take immediate corrective action
- Capacity planning
  - Anticipate future bottlenecks and avoid them
    - Understand current usage
    - Project for expected changes


©2002 SRM Associates, Ltd.



## Common Quote

For the most part, Windows 2000 is a "self-tuning" operating system. This means that in most cases, Windows 2000 automatically adapts to perform optimally "right out of the box" depending on the environment in which it's running—assuming that the hardware is properly configured.


©2002 SRM Associates, Ltd.



## Rest of Quote

However, as with any operating system, performance depends on many outside factors such as hardware, device drivers, applications, workload, the network, and so forth. In addition, there are certain practices and tuning guidelines that can be followed to optimize the performance of Windows 2000 in certain environments.


©2002 SRM Associates, Ltd.



## NT Data

- There is a very rich set of data available in NT from performance monitor
- It is critical to avoid data overload
  - Particularly in time critical performance situations, it is vital to pinpoint the problem and correct it rapidly

©2002 SRM Associates, Ltd.



## INFORMATION IS NOT KNOWLEDGE!!

©2002 SRM Associates, Ltd.

## A server can only be engaged in three specific types of activity:

- It can be acquiring or archiving information to be processed. This can be either data or instructions. This is I/O.
- It can be holding this data in its internal storage in some manner. This is memory.
- It can be performing operations (instructions) on the stored information (data). This is CPU or Processor.

©2002 SRM Associates, Ltd.

## Still Only Categories of Data

- Processor
- I/O
- Memory
  
- What about Network ?

©2002 SRM Associates, Ltd.

## Performance Standard

- If the processor is not excessively busy and work is completing on schedule, all is well.
- If work is not meeting service levels and the processor use is excessive, begin by examining the processor.

©2002 SRM Associates, Ltd.

## Performance Standard

- If work is not meeting service levels and the processor use is not excessive, look to I/O or memory as the problem.
  
- Beware: it is not quite that simple.

©2002 SRM Associates, Ltd.

## NT Differences

- **NT servers will run short of memory before any other resource. Watch this carefully. NT seems to consume memory.**
- **Even when it looks like there is another resource problem, check memory as well.**
- **When memory problems have been ruled out, disk and network should be checked next.**
- **The best performance gains in NT will come from tuning memory, disk and network subsystems, in that order.**
- **In case it hasn't become obvious, NT exhibits a change from previous platforms in that CPU is the least important resource to consider**

©2002 SRM Associates, Ltd.

## Performance Analysis Process

1. Start monitoring your system now and develop a baseline.
2. Monitor and review your system regularly to anticipate problem.
3. Identify the problem resource using the program outlined above; Using the metrics defined below.
4. Make sure you have a tested backup of the system, files and applications before making any changes.
5. Whenever possible, change one variable at a time and document everything.
6. Benchmark after the change to quantify the effects and to ensure system stability.
7. Go back to step 2.

©2002 SRM Associates, Ltd.

## Prerequisites

1. If you are interested in network data, make sure to add the SNMP service to collect data on the network interface object so that you will get info concerning the network interface card.
  - Select start/settings/control panel/networks/services/add and add the SNMP service. You will need to reboot to activate.
2. You should also add the network tools and agent to collect network segment data.
  - Select start/settings/control panel/networks/services/add and add network tools and agent. You will need to reboot to activate. Note: network monitoring under NT 4.0 will put the selected NIC in promiscuous mode, which will add overhead. Not so in W2K.
3. I strongly urge you to collect disk performance statistics by turning on diskperf. The rumors of high overhead are unfounded on today's processors. Just enter diskperf - ye from the command prompt and reboot. Full options for the diskperf command can be obtained by typing diskperf /? at the command line. This facility will stay on through subsequent reboots until it is disabled.

©2002 SRM Associates, Ltd.

## Available Performance Tools

©2002 SRM Associates, Ltd.

## Starting Perfmon or Sysmon

- Perfmon (NT) or Sysmon (W2K) can both be started by START/RUN and entering perfmon.exe.
- NT Perfmon can also be started by Settings/Control Panel/Administrative Tools/Performance.
- W2K Sysmon can also be started by Start/Programs/Administrative Tools/Performance
- Chart Mode:
  - Great for looking at current activity or for reviewing captured log files.

©2002 SRM Associates, Ltd.

## Starting Log Mode

- Start Sysmon. Expand the Performance Logs and Alerts icon in the left pane. In the right pane, right-click and select New Log Settings.
- At this point, enter a name for the logging session – make it descriptive. You will then see the Properties tabs of the log set you are defining. In the General tab, add those objects/counters that you wish to log.
- Next, select the Schedule tab and select the appropriate check boxes to identify the log performance data.

©2002 SRM Associates, Ltd.

## Performance

- Concentrate on the area of failure
- Limit the metrics analyzed to key counters
- Have historical data
- Build reports that allow for quick, high-level review of systems performance

©2002 SRM Associates, Ltd.

## Capacity

- Build an historical data base
- Establish baseline values
- Review the report periodically for trend changes
- Model possible effects

©2002 SRM Associates, Ltd.

## General System Hardware Tuning

- Make sure that you have updated BIOS and drivers on all server hardware
- Ensure that BIOS is configured for maximum performance
- Ensure that the BIOS settings match the installed hardware.
- Verify the actual delivered hardware.

©2002 SRM Associates, Ltd.

## Tuning Benchmarks

- Make one change at a time, such as adding hardware resources or changing an operating system setting.
- Repeat monitoring after every change..
- In addition to monitoring, review event logs, because some performance problems generate output you can display in the Event Viewer.

©2002 SRM Associates, Ltd.

## Tuning Benchmarks

1. Put a repeatable load on a server.
2. Measure how system resources perform.
3. Analyze those measurements.
4. Modify server hardware and software to eliminate the bottlenecks.
5. Start the process over until you have reached the ultimate performance bottleneck.

©2002 SRM Associates, Ltd.

## CPU Tuning

- Make sure that the root cause is not within another resource area.
- Remove CPU overhead
  - wasteful hardware components
  - don't implement compression
  - offload CPU intensive work.
- Remove faulty hardware.
- NT service packs
- Last resort – upgrade the CPU.

©2002 SRM Associates, Ltd.

## CPU




- Switch to CPUs with a larger L2 cache.
- Switch to faster CPUs.
- For multiprocessor systems, add more CPUs.
- On multiprocessor computers, manage the processor affinity with respect to process threads and interrupts.
- After you have done everything to eliminate a CPU bottleneck, you can still improve performance by clustering the server with another.

©2002 SRM Associates, Ltd.

## CPU RED FLAG

System % Total Processor Time—Performance Monitor object consistently near 100%.

©2002 SRM Associates, Ltd.



## Memory Tuning

1. Select the appropriate NT Memory strategy from Control Panel/ Network/Services/Server/Properties or
2. System/Advanced/Performance Options/Virtual Memory
3. Optimize virtual memory and the paging file system
4. Remove unnecessary processes from the server
5. Schedule memory intensive jobs to off hours.
6. Last resort – Add RAM

©2002 SRM Associates, Ltd.

## MEMORY

- Increase physical memory above the minimum required.
- Create multiple paging files.
- Determine the correct size for the paging file.

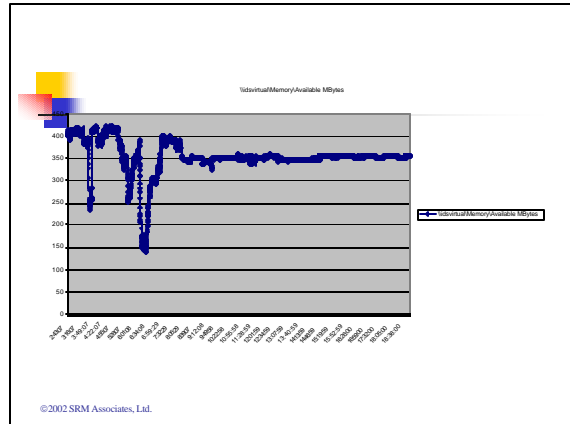
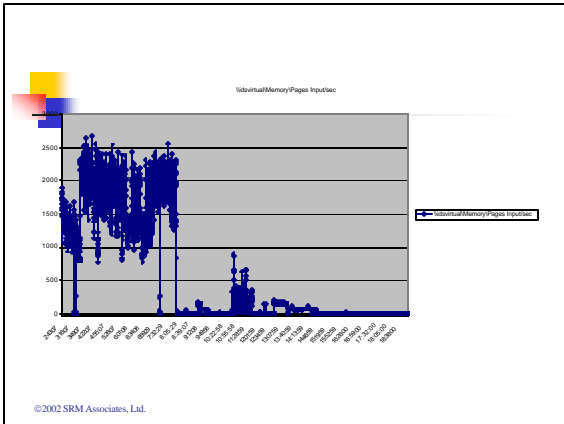
Ensure that memory settings are properly configured.

©2002 SRM Associates, Ltd.

## MEMORY RED FLAGS

- Memory Available bytes consistently less than 4MB (Intel) or less than 8MB (Alpha).
- Memory Available bytes decreasing over time.
- Paging File: % Usage, % Usage Peak is near 100%.
- Memory Committed Bytes is less than RAM.
- If Memory Committed Bytes approaches Memory Commit Limit, and if the page file size has already reached the maximum size as defined in Control Panel, System, there are simply no more pages available in memory or in the page file.

©2002 SRM Associates, Ltd.



## Disk Tuning

- Disk Performance vs. Disk Capacity
- File System Cache Tuning
- Network Shared Subsystems
- Distribute file system activity
- One logical disk per physical disk
- Group similar disk work.
- Use RAID appropriately.
- Avoid compression and encryption.
- Last resort – Add hardware.

©2002 SRM Associates, Ltd.

## DISK

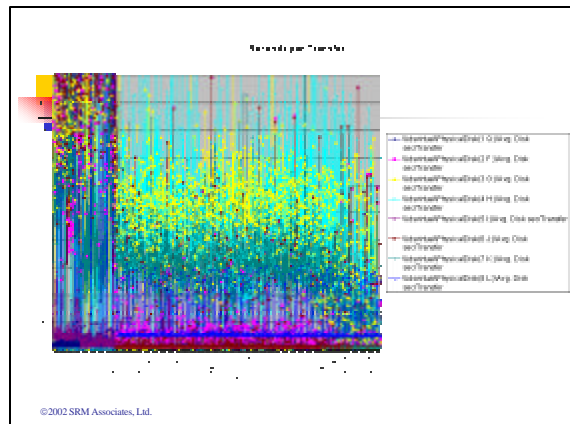
- If the server uses a RAID, add more disk drives, using faster drives if possible, and increase the memory cache size on the controller.
- If the server does not use a RAID, switch to higher-speed disk drives.
- Use Disk Defragmenter to optimize disk space.

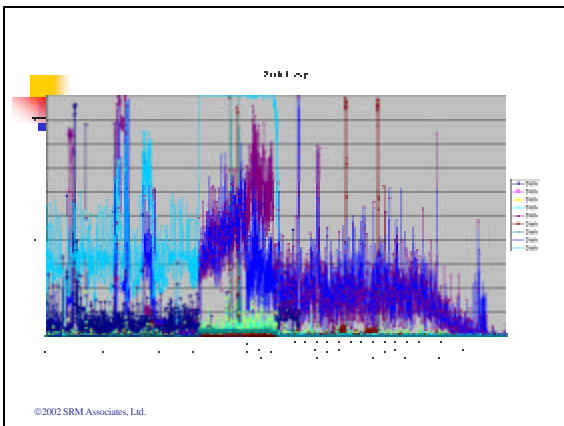
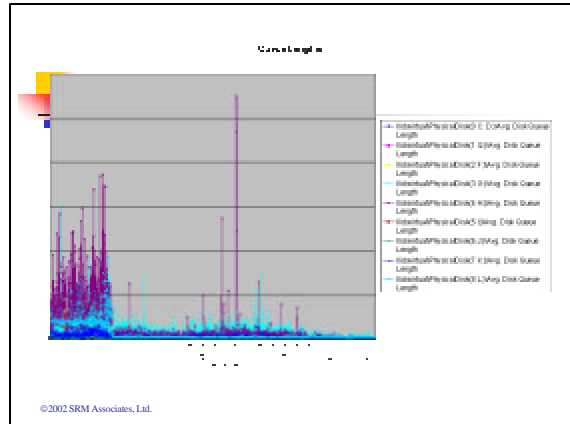

©2002 SRM Associates, Ltd.

## DISK RED FLAGS

- Physical Disk % Disk Time counter consistently at or near 67%.
- Physical Disk Queue Length > 2.

©2002 SRM Associates, Ltd.





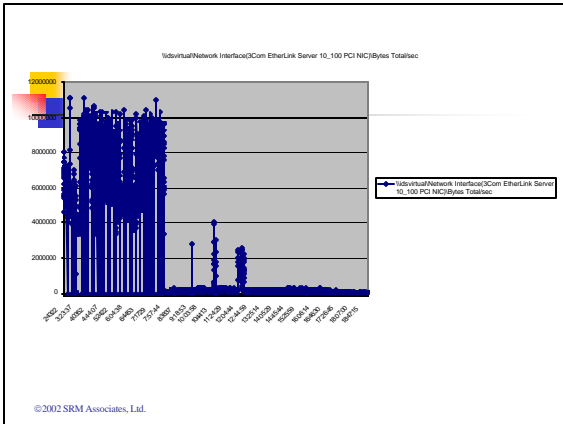
- ## Network Tuning
- Remove unnecessary protocols and services.
  - Network binding search order.
  - Periodically run Network Monitor.
  - Balance network loads
  - NIC settings
  - Device drivers and BIOS levels
  - Last resort – Faster NICs
- ©2002 SRM Associates, Ltd.

## NETWORK

Network monitoring typically consists of observing server resource usage and measuring overall network traffic. With System Monitor you can handle both of these activities, although for in-depth traffic analysis, you should use Network Monitor.

©2002 SRM Associates, Ltd.

- ## NETWORK
- Partition the network load so that it is balanced among all network adapters in the server. This is most easily done using VLANs in a switch connected to the server.
  - Use smart network adapters to take advantage of the advanced offloading features in Windows 2000.
  - Add more network adapters to increase the available bandwidth, although doing so will increase the interrupt load on the CPU if you use "dumb" network adapters.
  - Change to a higher-speed network, for example, upgrade to gigabit networking from 100Base-TX.
  - Unbind infrequently used network adapters.
- ©2002 SRM Associates, Ltd.



## Sizing Strategies

- Define your objectives.
  - Identify and understand the business and technical requirements to support the objectives.
  - Quantify the loading characteristics.
  - Determine the performance objectives.
  - Anticipate future business requirements.
  - Understand future system architecture.
- Configure the system.
- Stress test to validate the configuration.
- Utilize the tuning methodology during testing.
- Implement the system in production.
- Continue to utilize the tuning methodology.

©2002 SRM Associates, Ltd.

## TUNING LAWS

- Thou shalt not change more than one system parameter at a time.
- There is always a bottleneck; tuning just minimizes it and moves it around.
- One bottleneck may mask another.
- The Heisenberg Uncertainty Principle also applies to performance monitoring.

• Use the Task Manager.

©2002 SRM Associates, Ltd.

## Tuning Advice

- Limit the metrics that you collect and analyze.
- Build an historical database of performance metrics for the key systems in your environment.
- Develop reports and graphs that allow for a quick high-level review of system performance.
- Conduct this high level review on a regular basis.
- Through this review process, develop an understanding of the baseline values for normal, acceptable behavior of these systems.

©2002 SRM Associates, Ltd.

## Tuning Advice

- Be on the lookout for any changes in this expected behavior pattern.
- Use the metrics to isolate the resource that is the bottleneck (or potential bottleneck).
- Take corrective action.
- Go back and make sure that the correction was effective.
- Start regular monitoring again.

©2002 SRM Associates, Ltd.

## Performance Counters

Windows 2000 Performance Guide  
Mark Friedman & Odysseas Pentakalos

Tuning and Sizing NT Server  
Chris Aubley

Windows NT Performance  
Mark Edmead & Paul Hinsberg

©2002 SRM Associates, Ltd.



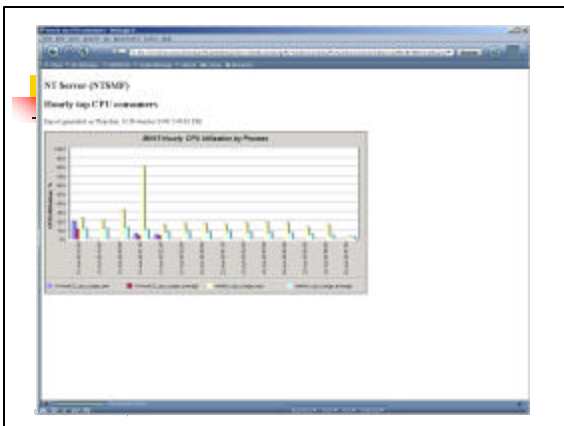
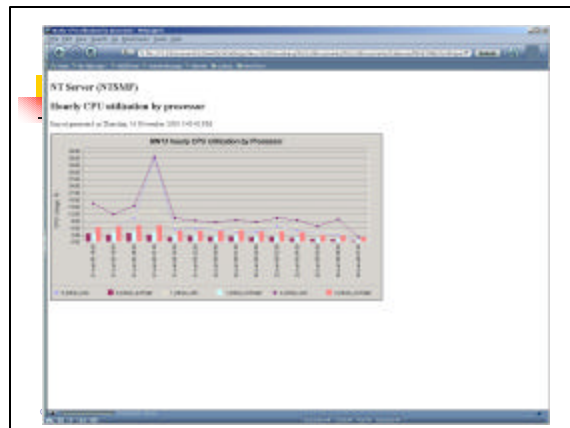

Jerry Rosenberg  
 Strategic Resource Management  
 Associates, Ltd.  
 www.srmassoc.com  
 516 433-1817  
 jerry@srmassoc.com

©2002 SRM Associates, Ltd.

## Processor

- Processor: % Processor Time
  - < 50%      80% = Danger
- System: % Total Processor Time
- Processor: Interrupts/sec
  - < 3500 for Pentium
- System: Processor Queue Length
  - < 2

©2002 SRM Associates, Ltd.



## Processor – Other Metrics

- System: Context Switches/sec
- Process: % Processor Time, % Priv. Time, % User Time, Priority Base
- Thread: % Processor Time, Priority Current

©2002 SRM Associates, Ltd.

## Memory

- Soft Faults
  - Page/faults/sec
- Hard Faults
  - Pages/sec

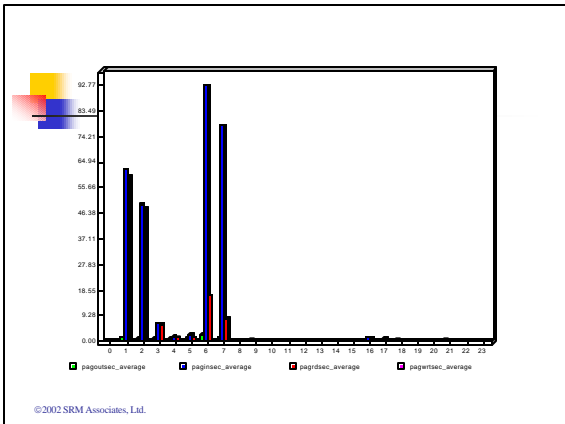

Note: Hard fault = I/O; More RAM yields Soft Faults

©2002 SRM Associates, Ltd.

## Memory

- Examine both memory and paging
- Hard faults s/b < 20
- Hard Faults/Total Faults should not exceed 5-10%

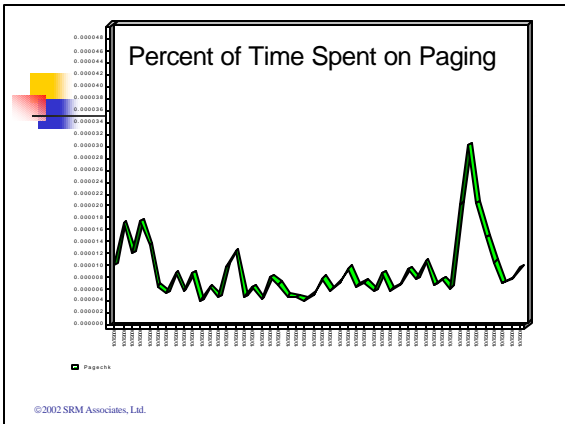
©2002 SRM Associates, Ltd.



## Paging Standard

- Memory:  $\text{Pages/sec} * \text{PhysicalDisk: Average Disk sec/Transfer} = \text{amount of disk time spent on paging activity during the interval.}$
- Should not exceed 10%
- 20% = thrashing

©2002 SRM Associates, Ltd.



## Paging

- Server: Pool Non-paged Failures.
  - Indicates that physical memory is too small
- Server: Pool Non-paged peak
  - Reasonable indicator of how much physical memory you need

©2002 SRM Associates, Ltd.

## Memory – Other Metrics

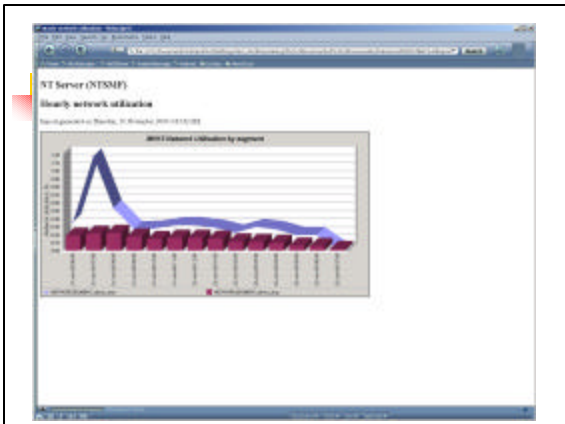
- Memory: Available Bytes
  - Never below 4Mb; Ideally 10 Mb
- Process: Page Faults/sec, Working Set
- Page Input/Sec: # pages read from disk
- Page Reads: # reads from disk
- Memory: Committed Bytes
  - Amount of memory in use; current memory footprint

©2002 SRM Associates, Ltd.

## Memory – Other Metrics

- Memory: % committed Bytes in Use
  - Ratio of committed bytes to commit limit
  - Hit commit limit -> resize paging file -> frag
- RAM size S/B 1.5-2 X Avg. Committed bytes
- Also check Paging File: % usage and % usage peak

©2002 SRM Associates, Ltd.



## I/O

- Physical
  - Actual Device
- Logical
  - Partition
- Need diskperf -y or diskperf -ye

©2002 SRM Associates, Ltd.

## I/O

- % Disk Time
  - < 55%
- Current Disk Queue Length
  - < 2 (QL - # spindles)
- Avg. Disk Bytes/Transfer
  - Big numbers are good

©2002 SRM Associates, Ltd.

## I/O – Other Metrics


- Memory: Pages/sec
- Physical Disk: Average Disk Bytes Read, Disk Bytes/sec
- Processor: % Processor Time, Interrupts/sec

©2002 SRM Associates, Ltd.

## I/O – Other Metrics

- Physical Disk: Avg. Disk Sec/Transfer
  - S/B < 0.03
- Avg. Queue time = Disk Queue Length \* Avg. Disk Sec/Transfer
  - Compare over time and with other disks

©2002 SRM Associates, Ltd.



## Network

- Need to examine:
  - Amount of network Activity
  - Throughput

©2002 SRM Associates, Ltd.

## Network

- Complex and configuration dependent
  - Protocols
  - NICs
  - Network Applications
  - Topology
- Maybe multiple protocol stacks
- Hint: monitor Redirector object

©2002 SRM Associates, Ltd.

## Network

- Redirector:
  - Bytes Total/Sec
  - Current Commands
    - Will increase when there is a delay in placing frames into the network. Should not get much longer than number of NICs in the box.
  - Network Errors/Sec
  - Reads and Writes Denied/Sec
  - Raw Reads and Writes Rejected/Sec

©2002 SRM Associates, Ltd.

## Network

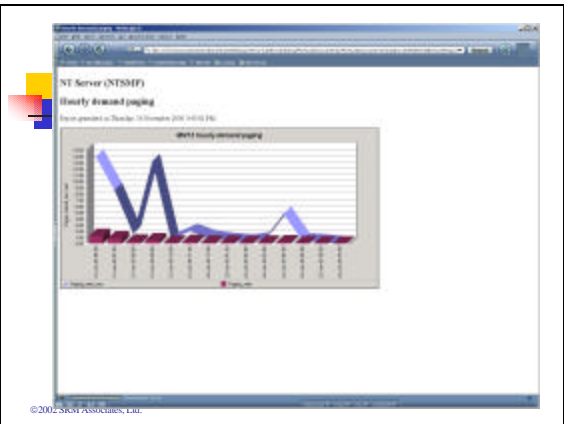
- If the sum of Server:Bytes Total/sec for all servers approaches the maximum transfer rate of your network, you are approaching saturation and should segment.

©2002 SRM Associates, Ltd.

## NIC example

- NIC:
  - Output Queue Length
    - < 2
  - Bytes Total/sec
    - Should not be high if OQL < 2
- If collisions greater than 10%, Network issue

©2002 SRM Associates, Ltd.



## Network – Other Metrics

- Network Segment: % Net Util., Total Bytes Received/sec
- Memory: Pages/sec
- Logical Disk: % Free Space
- Paging File: % Peak Usage
- Physical Disk: % Disk Time, Avg. Q Len.
- Processor: % Processor Time

©2002 SRM Associates, Ltd.

