

A020501

Monitoring IIS Performance

The best way to monitor the performance of your IIS installation is to use the Performance monitor and logs to record the long term statistics for the IIS server in question.

To use Performance in this way, you should leave it running for a number of hours or, better still, days to allow the information to collect over a reasonable time during typical usage.

Armed with this information, you should be able to identify any potential bottlenecks and from there make a reasoned decision about which part of the system needs to be upgraded or which parameters need to be modified to alleviate the problem.

If it's a hardware issue, it's usually a case that only the primary bottleneck gets displayed. That is, identifying a RAM bottleneck and resolving that issue might solve the problem, but it might also lead to you identifying a further disk issue and from there a CPU limitation.

With software or configuration issues, it's important to address a specific parameter one at one time so you can examine the effects of the modification. Change more than one, and it becomes impossible to identify which parameters has had the greatest effect.

Performance Counters

You probably already use Performance for monitoring other areas of your system, and you will already know that Performance works through a series of counters organized by the object on which they report.

For example, the Available Bytes counter within the Memory object reports the number of free bytes after system and user memory have been taken into consideration.

The counters you should monitor when checking the performance of IIS are listed in the following table.

Object	Counter	Notes
Memory	Available Bytes	Should be kept below about 20–30MB; lower than this indicates a lack of memory.
	Committed Bytes	This is the amount of memory already in use by the system and applications. In an ideal world, it shouldn't get above 75%.
	Pages/sec	Too high a value (more than 70–80), and you probably have too little RAM and the machine is having to swap too much out to disk all the time.
Processor	% Processor Time	If high but with low disk and RAM figures, it indicates a potential lack of CPU resources, so either upgrade to a faster CPU or add an additional CPU to share the load if you can.
Disk	% Disk Time	The amount of time spent working with the disk should be kept relatively low.
	Avg. Disk Bytes/Transfer	Should be relatively high because it indicates how much data is being transferred from the disk.
	Avg. Disk Queue Length	Above about 4 indicates that a high number of requests are being sent to the disk subsystem but not processed effectively.
Network	Bytes Total/sec	Should be high, but should also be within the limits of the network bandwidth for the interface on which the Web site or sites are running.
System	Context Switches/sec	Should be relatively low—it indicates how many times the system is switching between applications and services. Too high probably means you need additional CPU horsepower.
	Processor Queue Length	Should stay below 2 for each CPU in the system. Higher than this, and it indicates too many processes are running on the system. If you are using Web gardens to execute multiple instances, make sure you have not configured too many processes per CPU.
Paging File	% Usage	The usage of your page file should be within twice the amount of your physical RAM. Higher usage than this indicates a lack of physical memory.
WWW Service	Bytes Total/sec	The total number of bytes transferred by the WWW service. Should obviously be as high as possible, but you should really use it to make a comparison against other parameters to ensure that high CPU, RAM, or disk loading is related to the WWW service.

Object	Counter	Notes
	File Cache %	Shows the hit rate for items in the IIS cache. If it's a low value, consider disabling caching or changing the way your Web site or application operates to use cache more effectively.
	File Cache hits	Should be high if you have static content, but could be low if you have a heavy dynamic site.
	ISAPI Extension	Shows the rate of your ISAPI Extensions; a drop indicates an ISAPI failure or an application that is not terminating properly.