

System Monitoring

*System Monitoring - Seeing is Believing for
IBM System i Managers*


System Monitoring – Think Inside the Box

Capability breeds dependence. The more you can do, the more is asked of you. The more time you save, the more time you spend seeking out ways to save more time or money. Improved performance leads to higher expectations. If less really is more, how do you do more with less? Such are the challenges faced by many an IT Manager. If the odds are stacked against you despite your best efforts, perhaps it's time to redress the balance. Effective system monitoring could help you break free of the system shackles and use your System i (iSeries) network to save time and money...just like magic.

Money Trees and Time Machines - Myth Busting

'A powerful, robust network of servers – imagine all we could achieve with that. Think of the money we'll save; think of the time we'll save, think of that time as money and just imagine the savings...' The technologies that support businesses that dare to dream big are built on concepts such as this. The leap to seeing your servers as money trees and time machines isn't so large. This is the company's tool for modern day alchemy, isn't it?

The problem is not the server's ability to reach optimal performance so you can reap the rewards of increased productivity – it's making sure it maintains optimal performance on a 24/7 basis. Add to this the typical problems of a reduction in staff to manage the network and a growing user base that is calling in problems faster than you can identify or resolve them. It could feel as though the very technology that was designed to increase your business productivity takes more people (money) and time (money) to manage. How much is the added productivity really costing you and is increasing the size of your network only multiplying your existing, recurring and unresolved problems?



Calculating the Cost

An underperforming system or network is a wasteful extravagance. Most businesses simply can't afford to profitably operate at less than optimal levels. Problems occurring in the following areas typically take the most time and resource to investigate and resolve or, left unattended, are most likely to lead to extremely serious and costly problems such as a lack of availability on the system.

- Job Performance
- System Performance
- Memory Pool Performance
- System Configuration

Case Study

Analysis of problems encountered in the network environment over a three month period at Company X, a large food manufacturer, shows the extent that their underperforming network is costing them. The environment supports 7,000 users with annual revenues of \$1.4billion. The cost associated with each of these problems is only representative of the time it takes to investigate and resolve the problem, not the additional expenses of further impact such as a loss of availability or additional disk purchase etc.

Area	Problem	Reason	Consequence	Cost
Job Performance	There is no obvious reason as to why disk space is being rapidly consumed	A looping job is consuming vast amounts of temporary storage	Time consuming investigation and/or additional disk is required to avert system crash	\$1,760 ¹ Annual Est: \$7,040 ++
System Performance	CPU usage is consistently high throughout the day	Several QZDASOINIT take more CPU than they should	Time consuming, jobs all have same name, which ones are causing problem	\$1,210 ² Annual Est: \$4,840 ++
Memory Pool Performance	Certain jobs experience poor performance	There is a high rate of faulting in a specific pool indicating too many jobs in the pool or insufficient memory in the pool	Time consuming, which memory pool are jobs running in, then how many jobs are in pool etc.	\$79,965 ³ Annual Est: \$319,860++
System Configuration	Users complain of poor response times	1 IOP is being utilised more than the others	Time consuming investigation	\$1,440 ⁴ Annual Est: \$5,760++

Key:

- ¹ Two 4hr incidences of wasted profitability hrs for 2 people and 2 idle workers' time
- ² One 6hr incidence affecting a group of 10 users running at ½ productivity capacity
- ³ Three 3hr incidences affecting a group of 130 users running at ¼ productivity capacity
- ⁴ Two operators at weekend rates to investigate problem – still no conclusion after 2 days

Company X is planning on expanding their System i environment through the acquisition of a new food brand. The deal will see them inherit additional servers presently used by the brand. The benefit of additional processing power could be completely undermined by the problems they are experiencing if they are not addressed before the expansion. The central issue to all their problems is a lack of visibility. Without this, they struggle to deal with the evidence of problems without understanding how they occurred or why and stand little chance of increasing their productivity or profitable use of resources. If these problems sound familiar, ask yourself how long your organisation could reasonably sustain such a reactive approach to systems management?

All Seeing; All Knowing; All Powerful; All the Time

If system expectations are set at a level that allows for a zero margin of error, IT Managers will have to look to their resources to achieve the seemingly impossible. No system is immune to problems so the solution must lie in the approach to dealing with those problems in the most effective way. Increased visibility will help identify problems but no organisation can afford to assign the man power to watch every component at all times and even if they could, it's likely that a few would still slip through the net – we're only human after all.

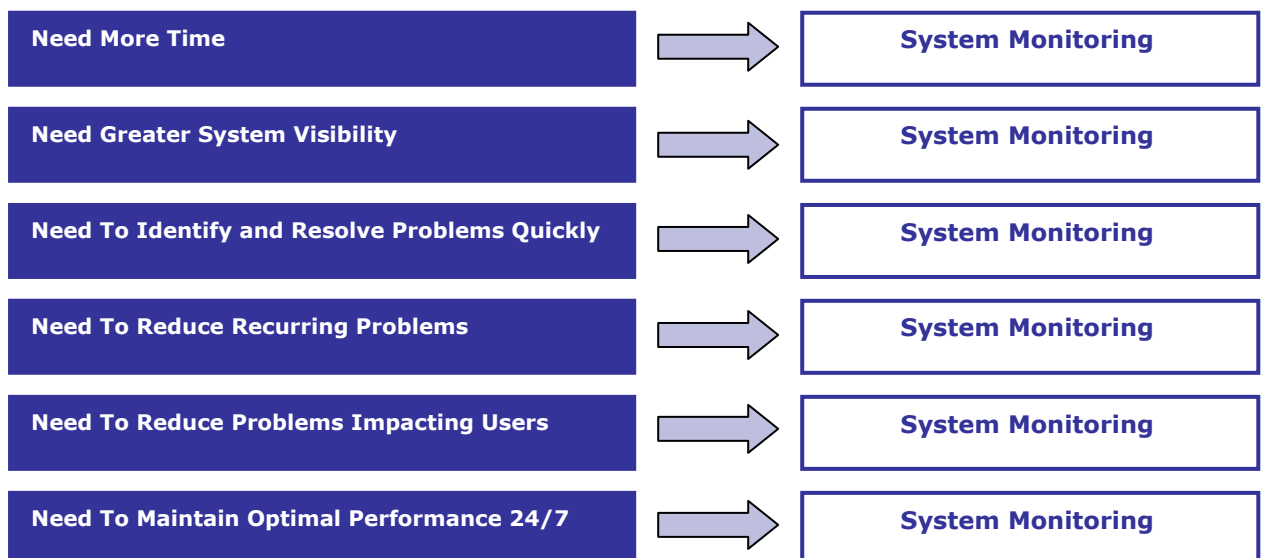
Many automated monitoring solutions exist that can help alleviate the burden of system monitoring but their effectiveness should be judged on the degree of visibility they can offer to potential problems, the time they save and the flexibility they offer in meeting the specific views and requirements of your existing and developing network. In short, they should accommodate the typical and likely problems that can be anticipated based on your system demands, applications, throughput and configuration. In addition, they need to provide fast access to information that will pinpoint the cause of unforeseen issues.

Creating Hindsight

If only Company X knew that a looping job was the cause of their disk space problem... 'If only' scenarios torment us with facts that, had they been known in advance, could have made a positive difference to the outcome. Whilst no solution can turn back the hands of time, it is possible to anticipate problems before they escalate in severity and impact users' productivity. By monitoring all the critical elements of the system in real time and only being alerted to exceptional conditions – i.e. those that breach their optimum boundaries, managers can use their problem solving time to best effect.

Evaluate Your Needs

A quick evaluation of your existing problems could reveal your most urgent managerial issues can be resolved with a new approach to system monitoring.



As with Company X, you will already be aware of the most common problems that have occurred on your system or network. Accounting for unknown future problems presents a bigger challenge but one that can be met by considering the types of demands that are placed on your system and determining the impact your particular combination of applications, jobs and configuration plays in your daily operations. To find a solution that meets both your present and future problems, it's important to ensure it meets a checklist of requirements that can accommodate both present and future needs. Even if you have not experienced the problems associated with the areas below, ask yourself how likely they would be to occur in your environment and what extent their impact would have. Would you know about them immediately or only after system performance was compromised?

Job and System Performance Checklist for your System i Network:

Job Performance	System Performance
▪ Response Time	▪ Active Batch Jobs
▪ Transactions per Hour	▪ Active Interactive Jobs
▪ CPU per Transactions	▪ CFInt Overhead
▪ CPU Usage	▪ CPW Available
▪ Interactive CPU Usage	▪ CPW Percent of Available
▪ Database CPU Usage	▪ CPW Percent of Configured
▪ Faults per Second	▪ CPW Used
▪ Lock Wait Time	▪ Interactive Capacity Used
▪ Temporary Storage	▪ Microcode Tasks
▪ Thread Count	▪ System Overhead
▪ Disk I/O	▪ W/I over A/W

Most Common Associated Problems Relating to Job and System Performance:

- Users are complaining that the system is slow but the user or job responsible for the problem is not immediately obvious – painstaking investigation is required
- You need to plan for increased capacity but there is no quick way to determine how many transactions per hour are going through the system
- The Accounts department is running a group of critical jobs including year-end processing but it's not clear how much CPU these are using individually or collectively and at what expense to other departments that are also running critical jobs
- You need to evaluate your LPAR box to determine if an upgrade is necessary. There is no simple way to determine how much of the potential processing power of the partition/system is actually being used
- System performance is being impaired but the users are not responsible. The amount of resource being used by the system indicates a problem in the original set up. You need a fast view to quantify this resource use independently of user activities.

Memory Pool Performance, LPAR and System Configuration Checklist:

Memory Pool Performance	System Configuration
▪ Active to Ineligible Transitions	▪ UPS Battery Status
▪ Active to Wait Transitions	▪ Running on UPS
▪ Activity Level	▪ Unallocated Disk Units
▪ Database Page Faults	▪ Damaged Main Storage Unit
▪ Database Pages	▪ Front Panel Battery Status
▪ Non-Database Page Faults	▪ Processors Not Working
▪ Non-Database Pages	▪ Processors Partly OK
▪ Reserved Size	▪ Front Panel Status
▪ Current Size	▪ Last Termination Cause
▪ Wait to Ineligible Transitions	▪ Remote Power on Enabled
	▪ Auto Restore When Power Restored
LPAR	▪ Next IPL Type
▪ Current Processors Per Partition	▪ Key Position
▪ Current Memory Per Partition	▪ UPS Bypass
▪ Current Interactive Per Partition	▪ PTF Status

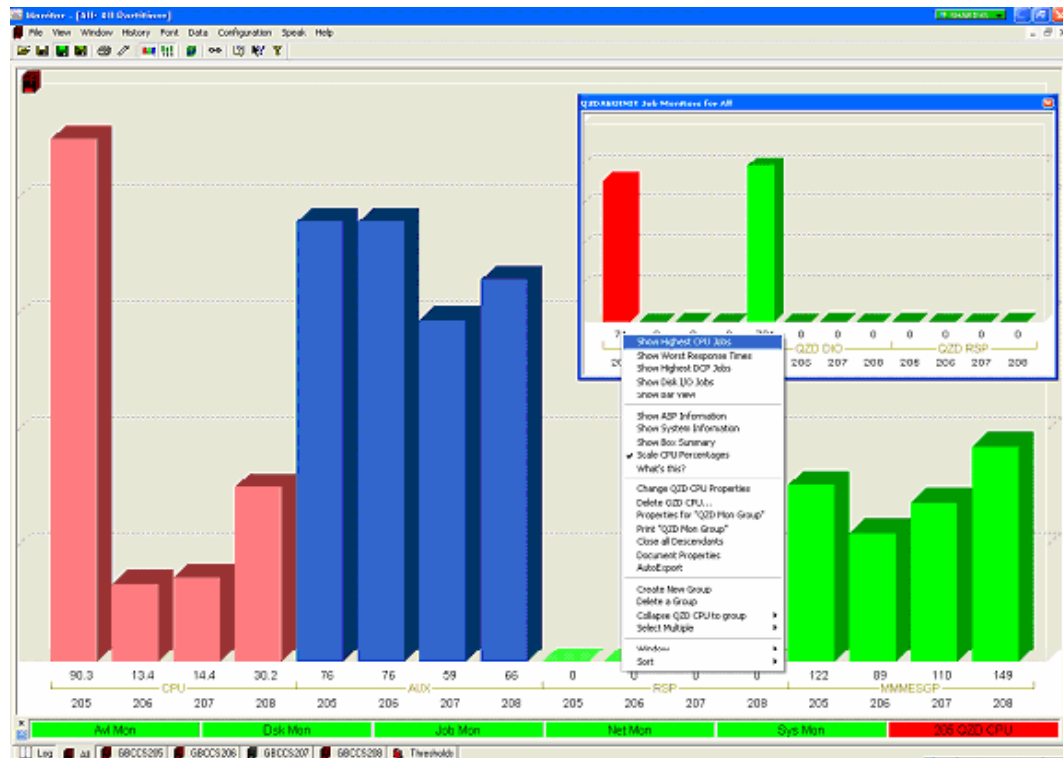
Most Common Associated Problems Relating to Memory Pool Performance, System Configuration and LPAR:

- Multiple jobs are sharing insufficient main memory increasing I/O and CPU usage. Monitoring the number of non-database page faults could alert operators to the situation so they can assign more memory to the jobs to decrease unproductive use of disk and CPU resources and avert the associated problems of diminished throughput and longer response times
- Immediate summary of LPAR system information is required in problem analysis without a lengthy interrogation of OS for each partition
- Immediate visual confirmation of key system information is needed to immediately rule out or identify serious issues relating to status of elements in problematic scenarios such as power outage or for a quick reference to important information such as the number of unallocated disk units should a fast reallocation of resources be required

What's the Problem?

Accountability for system resource and associated performance, its problems and cost justifications is becoming an increasingly important consideration for System Managers forced to play detective in a continuous 'who done it' analysis of jobs gone awry, temporary storage missing in action and users beating a well worn path to the helpdesk, yelling 'what's the problem?' Effective system monitoring eliminates the problems that lead to these types of situations by imposing the visibility to identify exceptional metric or status conditions before users are affected and giving immediate access to the information to resolve them.

The Screenshot below is an example of the real-time visibility that can be achieved through system monitoring with products available on the market today. Below: High CPU is a cause for concern. Real-time verification shows a group of QZDASOINIT jobs are responsible, so Managers can quickly isolate the particular jobs within that group and resolve the situation.



If your network could benefit from the type of foresight that comes with efficient system monitoring and you're ready to take advantage of doing more with less on your system i, contact the systems management experts, CCSS, to discuss how we can assist you in seeing the full potential of your system. After all, seeing is believing.

About CCSS

CCSS develops, supports and markets IBM System i performance monitoring and reporting, message management and remote management solutions. An Advanced IBM Business Partner, CCSS develops powerful solutions to support some of the world's most demanding System i environments across many industries including insurance, banking, pharmaceutical and manufacturing. All CCSS solutions are IBM ServerProven and qualify for IBM's ServerProven rebate program.

Existing customers that rely on CCSS's feature-rich solutions include leading organisations such as Volvo, Mattel, Newell-Rubbermaid, The Royal Bank of Scotland, Siemens Medical, RWE npower and Waterstone's. CCSS is headquartered in Gillingham, Kent, UK with key regional headquarters in Raleigh, North Carolina, USA; Bonn, Germany and Makati City, Philippines together with a global agent network spanning Portugal, Brazil, the Netherlands and Sweden.

www.ccssltd.com