

Network Monitoring

*Environmental Hazards to Peak Performance:
Why Network Monitoring is So Important for
System i Managers*

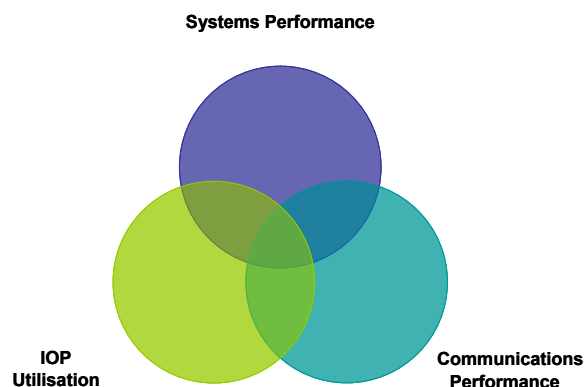
Network Monitoring – Blame the Environment

A professional athlete has many considerations in ensuring their personal best. The food they eat, the way they physically train and mentally prepare, as well as their strategy and technique, all play a factor in optimising their best result. Imagine all that had been attended to and when the athlete arrived for the race of their life, his performance was ruined as he was struck by lightning. Unlucky, some would say – or, it's unfortunate, but that's life. IT Managers are not in the business of luck. They need to prepare for every disaster scenario, every contingency and ensure that even when every object, job and process is running 100% perfectly on the system, some other unforeseen network element doesn't swoop in and ruin all their efforts, just because it resides outside the parameters of the System i. After all, in a situation like this, people remember the one thing that went wrong – not the millions of things that went right – it's unfortunate, but that's life.

Not Your Fault...Just Your Problem

Whilst no one can be aware of every potential problem that may or may not impact the system, IT Managers can reduce the incidence and impact of any potential problem in two ways. The first is to raise their awareness of the boundaries of their systems within the network and identify the areas of overlap where potential problems may occur. The second is to extend the proactive approach to systems management that keeps all their own elements running well, to these overlap areas, in an effort to protect the system from unforeseen problems and deal with them promptly and effectively when they do occur.

Frustrated users, who are delayed in their work, have no means of pinpointing the issue, so often any 'lag' to productivity is initially deemed an issue for Operators to resolve. In the case where these may actually be network issues, a bitter blame game can ensue with both Operations and Network teams pitted against one another as they each dig up 'proof' that the problem is not theirs. As this lengthy elimination process is carried out, the users are still waiting and productivity and profitability still suffers. As an Operator, it may not be your fault that the IOP's were incorrectly configured by an engineer, or that a user has sent a huge file that has had a knock-on effect on the network traffic, or the TCP Ping has fallen over and severed data communications between users and the server. None of it may be your fault – but with immediate visibility to these issues, solving these types of problems becomes a fast and pain free experience.



Calculating the Cost

Without sophisticated tracing tools to accurately pinpoint common 'network' issues that could be impacting users' productivity, it is very difficult to know where to start looking in identifying potential causes. What's more, as these issues have the potential to be either system or network related, quite often, members from both teams become involved in determining the root cause, effectively utilising the time and resources of twice as many people. The costs of this type of system/network detective work can soon add up to significant sums. Frequently, issues such as this often fall into two major categories and should be primary considerations for anyone considering this type of network monitoring:

- Network Bottlenecks and Errors
- IOP Utilisaton

Case Study

Company x is a large retail operation that was struggling with network issues that were not immediately visible on their centrally managed System i network. The network supports 10,500 users nationwide and the company generates \$2.1 billion in revenues annually. The IT Manager and Network Manager often wrestled with the same issues and decided to review the financial impact this had made in the past 12 months. The cost –impact results of the review were far higher than they anticipated, primarily due to delays and the associated cost of the time spent to identify the issue. The team was also acutely aware of the (as yet) unquantifiable fiscal damage to the brand reputation and loss of customer loyalty. The figures below only account for the immediate impact.

Area	Problem	Consequence	Cost
Network Bottlenecks	A user sent out a huge file in the form of a promotion to the customer base causing severe delays in transmitting critical stock replenishment data for 'just in time processing'.	Shelves which should have been filled with the best selling items on the busiest shopping days of the year were empty.	\$600,720 ¹ Annual Est: \$1,201,440
Network Errors	A comms line became inactive and the CEO did not receive the authorisation note to sign off his top directors' payroll just prior to the Christmas break.	The Directors did not receive their annual Christmas bonus as contractually promised. The monthly accounts were delayed as a result.	\$3,300 ² Annual Est: \$3,300
IOP Utilisation	A large group of users is experiencing poor response time, slowing their use of critical applications	Operators cannot see an obvious reason as to why as system looks ok. Network team says it must be a system issue – deadlock ensues, users impacted until proof is found.	\$158,887 ³ Annual Est: \$635,550
Total Annual Cost of Job Performance and Status issues			\$1,840,290
Key:			
¹ Two 6hr incidences of problem identification for 4 IT people and total loss of 1% of annual profitability			
² One 16hr incidence of problem identification for 3 IT people, 20 hrs overtime for 2 Accounts people, 7 hours overtime for one legal counsel person to ensure the breach of contracts could be resolved without legal action			
³ Four 5hr incidences affecting a group of 500 users running at ½ productivity capacity requiring 8 Ops/Ntw members to resolve it (IT time + Users' time + loss of profits @ \$96.15 p.h.p.p)			

See No Evil; Encounter no Evil

When users report slow response times to critical applications it can be hard to determine where to begin checking to isolate the issue. This is especially true when dealing with a distributed network environment. For example, users in Belgium may be experiencing poor response times but between the data centre in the US and the users in Belgium, there are numerous components and possibilities as to the cause of the lag on the network.

In other situations, there may be 'clues', such as high CPU usage but again this only serves to highlight an additional consequence, rather than the problem itself. Without visibility as to the primary cause, operators and network managers face a long elimination process that could see them involving additional resources in the form of IBM server engineers to try to eliminate the possible reasons for the lag. A lack of visibility inevitably translates to a negative impact on the user community and as such, must be the primary goal of any network monitoring solution.

Anticipating Communications Issues

Instilling a pro-active approach to monitoring communications elements on the network can pay dividends for operators and Managers. Whilst the communications themselves may not fall directly into the path of responsibility for IT/Ops Managers, real-time knowledge of when a line is down can help identify and resolve the issue before important jobs are delayed. To be able to produce an easy to understand graph showing the status of a particular communications line not only gives proof of 'availability', which may form part of a service level agreement, but also can be used as an additional security alert.

In this case, it may be important to generate a report showing exactly the times when an important line becomes active. Managers could employ real-time alerts to be notified of when a line is active as this may constitute a security breach. This can also give senior management an idea of how often third-parties, who may have a legitimate reason and full authorisation, are logging into the network to download support applications, for example. Creating such a process for monitoring lines will not only give management exact figures of when the line was opened, but also by whom. Using this type of pin-point analysis leaves no room for misunderstandings as the audit trail report will also give the profile name. Similarly, a third-party support team may claim that the line was inactive and they were unable to upload an important application 'fix' – the status report will provide the objective view on line availability at any given time so there is no dispute.

Beyond communication line status, operators should also be aware of throughput levels. By applying a define threshold to the line, operators can be alerted to any breach of that threshold, giving them advance warning that action must be taken, by them or by the network team.

Exposing Network Bottlenecks

Network bottlenecks have the potential to cause serious delays to users who may be prevented from running their critical applications. In time-sensitive industries such as banking, trading, logistics and even retail, as was shown in our case study, the impact on the organisation can be extremely costly. Network monitoring that offers an initial 'heads up' on conditions that have changes on the system dramatically reduce investigation time so operators and network teams know where to start looking to resolve the problem. What's more, operators can produce immediate verification on important surrounding factors that can help to determine the cause. Reports showing users activity, communications lines performance, utilisation and IOP utilisation eliminate the need for guesswork.

Creating a 'view' on IOP utilisation can be extremely useful as it could indicate an issue with the original configuration – a problem so 'hidden' that unless all other options have been ruled out, it is unlikely to be investigated. However seeing that utilisation is unusually high on a particular processor combined with the knowledge that users are experiencing very poor response when all other system elements are normal, would immediately highlight that processor for operators to carry out a more detailed examination rather than wasting their time investigating numerous other possibilities.

Flexible Monitoring

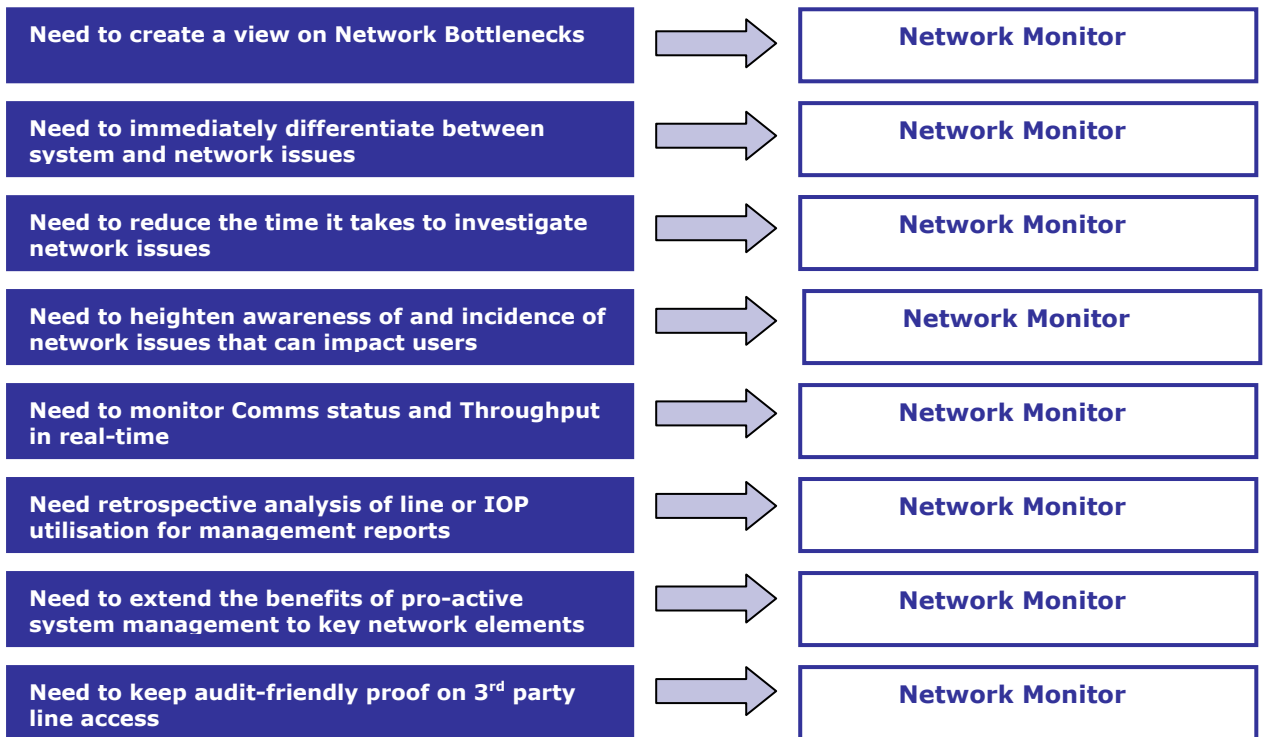
If your system i environment is already benefiting from a pro-active approach to monitoring, it's likely you'll want to extend this to areas of the network that, left unmonitored could impact a number of crucial system performance elements. Your choice of monitor should have the reach to accommodate this and allow you to keep a watchful eye on the status, and in some cases, the utilisation levels of, virtually any network element or device with an IP address. This can include any or all of the following:

- Other system i machines
- Specified PC's of individuals (e.g. the CEO)
- Mainframes
- Other Servers
- Routers

You could also use flexible 'network' monitoring to ensure that critical elements are active 24/7 – such as the organisation's firewall. In this scenario, Managers would be notified in real-time if the firewall became inactive and immediate steps could be taken to limit any resulting breach of security to ensure data was not compromised.

Evaluate Your Needs

Use the chart below to determine if your job issues could be resolved with a flexible 'network monitor'.



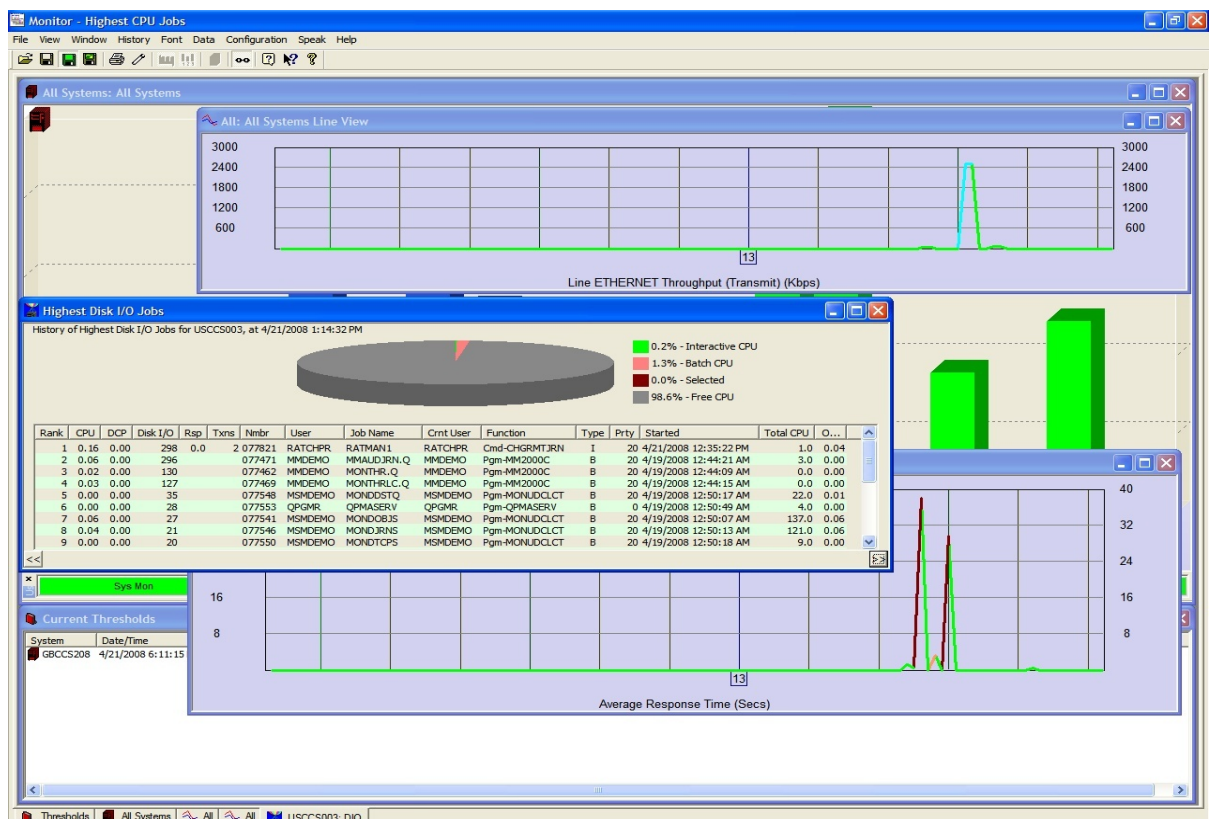
Looking at the range of network related issues experienced by Company x from our case study, it's plain to see that any potential monitor must have the flexibility to accommodate a virtually limitless choice of elements to monitor. The checklist below offers some key areas to consider when weighing up the functionality of any system monitoring solution that extends to network elements:

Network Monitoring Checklist for your System i Network:

IOP Utilisation	Communications Performance
<ul style="list-style-type: none"> ▪ IOP Utilisation 	<ul style="list-style-type: none"> ▪ APPC Ping
<ul style="list-style-type: none"> ▪ IOP Utilisation (Comms Only) 	<ul style="list-style-type: none"> ▪ Comms Ethernet >16 Retries
<ul style="list-style-type: none"> ▪ IOP Utilisation (Disk Only) 	<ul style="list-style-type: none"> ▪ Comms Ethernet > 1 Retry
	<ul style="list-style-type: none"> ▪ Comms Ethernet 1 Retry
	<ul style="list-style-type: none"> ▪ Comms Errors
	<ul style="list-style-type: none"> ▪ Comms Throughput Receive
	<ul style="list-style-type: none"> ▪ Comms Throughput Transmit
	<ul style="list-style-type: none"> ▪ Comms Utilisation
	<ul style="list-style-type: none"> ▪ TCP Ping

Network Monitor in Action

Below is a good example of how a network monitor can assist in identifying a resource problem on the System i. In this case, an earlier remote journal monitor highlighted the fact that a remote journal was in a status of "Inactive". The user issued the command CHGRMTJRN to activate the remote journal which resulted in unprocessed journal entries being sent from one system to another. The result is that the response time is affected (as shown by the peaks in the lower window), however the operator can see that the transmit figure for the Ethernet line during the same time period was also "peaking". By drilling down the operator is able to identify that the cause of the poor "response" times is in fact due to the user running the CHGRMTJRN command.



Perhaps you're tired of everyone blaming the system environment for network issues and need independent proof on your side, or perhaps you just need a better view on critical network elements because your team and time is already stretched, either way, if you've experienced some of the issues outlined in this white paper (or are hoping to avoid them in future), contact the systems management experts, CCSS, to discuss how we can help you to extend a pro-active approach to your network monitoring.

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.

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, Switzerland and Sweden.

www.ccsltd.com

