

Disk Monitoring

*Pay Now or Pay Later? Why Real-Time Disk
Monitoring is Money in the Bank for IBM
System i Managers*

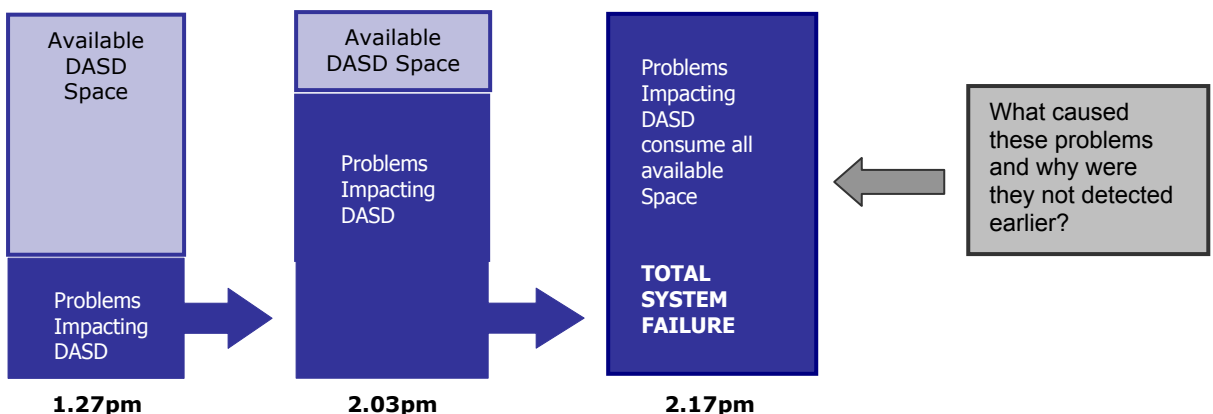
Disk Monitoring – Calculating the Real Cost of Disk

Many IT Managers will be familiar with the 'disk or die' ultimatums that a disgruntled system can issue. If the system demands feeding and threatens to shut down if it doesn't get more disk to chew, it's time to address the issue of disk once and for all. The expense of using disk to accommodate problems in the hope of buying extra investigation time is a quick fix that few organisations can sustain. If you've been held hostage to the system's insatiable appetite and are looking for alternative solutions to soothe the savage beast, look no further than your approach to disk monitoring.

Hungry Systems with Expensive Tastes

Whilst disk problems are not unique and could strike at any moment, they don't strike without good reason. Without knowing what the reasons are, Managers are forced to make difficult and expensive disk decisions that could create a big dent in their annual budget at the expense of other critical projects. What's more, if problems that affect disk are left unattended, they can escalate quickly and impact many other areas of system performance so Managers are under even more pressure to now solve a number of issues requiring immediate attention. With no solution to pinpoint the cause, Operators may not know where to begin investigations as there is no apparent chain of events with a root cause evident, only the combined results of multiple problems.

The goal is not to make less use of DASD, after all, it exists to be used, but rather to ensure that *all use is best use*. If any unnecessary or wasteful use is suspected, the organisation responsible must view this as an unnecessary expense. Catching instances of runaway DASD immediately and being able to isolate the problem that caused it for fast resolution are the main objectives of effective disk monitoring. The company that can do this can put an end to their system's expensive taste for disk.



Calculating the Cost

Potential problems with any system object be it a file, data queue, ASP or spooled file could trigger a situation that will impact DASD resource. The most common are listed below and left unattended, could have the greatest impact on the system. Any solution

for monitoring disk should pay particular attention to these areas and alert Managers to any change in their individual size, collective quantity, percentage statistics or status in real-time if they are to minimise impact on disk usage. Doing so will not only put Operators in a pro-active position to manage and respond to DASD threats generally but will also safeguard against the unnecessary expense of avoidable DASD usage.

- Journal Receivers
- Temporary Storage
- ASP
- Important Files
- Looping Jobs

Case Study

Company X is a large Facilities Management organisation. They are servicing three important, different customers on an i5 550 that are experiencing a number of problems that are impacting disk. The cost of purchasing more disk is an option but would greatly impact their profit for these customer accounts and would not resolve the problems long-term. Their Service Level Agreements are below target for two consecutive months on all three accounts and they face stiff financial penalties as a result. They are struggling to resolve the underlying issues as they are only aware of the result of the problems in the form of escalating DASD use. The table below reveals the true cost that a lack of effective disk monitoring is costing them.

Problem Area	Impact	Frequency	Cost
Journal Receivers	Average Avoidable DASD Usage: 15GB*	Customer A: x 1 Customer B: x 1 Customer C: x 2	<ul style="list-style-type: none"> • Total Penalty charges for sub-standard SLA terms = \$1,000 • Proportional DASD Costs = \$858 • Staff Resource to investigate issue = \$360
Temporary Storage	Average Avoidable DASD Usage: 9GB*	Customer A: x 5 Customer B: x 3 Customer C: x 2	<ul style="list-style-type: none"> • Total Penalty charges for sub-standard SLA terms = \$2,500 • Proportional DASD Costs = \$1,287 • Staff Resource to investigate issue = \$1,200
ASP	Average Avoidable DASD Usage: 5GB*	Customer A: x 1 Customer B: x 1 Customer C: x 1	<ul style="list-style-type: none"> • Total Penalty charges for sub-standard SLA terms = \$750 • Proportional DASD Costs = \$215 • Staff Resource to investigate issue = \$225
Total Cost			US\$ 8,395

* Market rate of US\$1,000 per 70GB; Agreed penalty of \$250 per incidence; Av. \$30ph staff time

The profit margins of Company X greatly depend on their ability to deliver agreed levels of service to customers 24/7 without resorting to expensive disk additions to combat unforeseen issues. Company X might just as easily be any large organisation

managing resource for specific groups within the company. In this case, justifying additional resource for these problems becomes a tense negotiation of budgets at best and at worst, a daunting trial of accountability for delayed processing, user inconvenience/unproductive time and other consequential problems that come with a heavy price tag, such as a lack of availability. Effective disk monitoring can provide companies struggling with issues that impact DASD the means to eradicate unnecessary use and solve the root cause of the problems.

Real-Time Insight into Disk, Temporary Storage & ASP's

Knowing about problems that could impact DASD is the first step to achieving effective disk monitoring. Sudden, unexplained temporary storage increases are a good example of how difficult it is to identify the root cause of problems without a specialist solution in place. Typically, Managers would notice a leap in temporary storage and be left with little option but to carry out a thorough investigation of each job in each sub-system, submitted by each user. The checking of hundreds and hundreds of jobs to identify which one was problematic is a marathon task that few Managers have the time to carry out once, let alone on a regular basis.

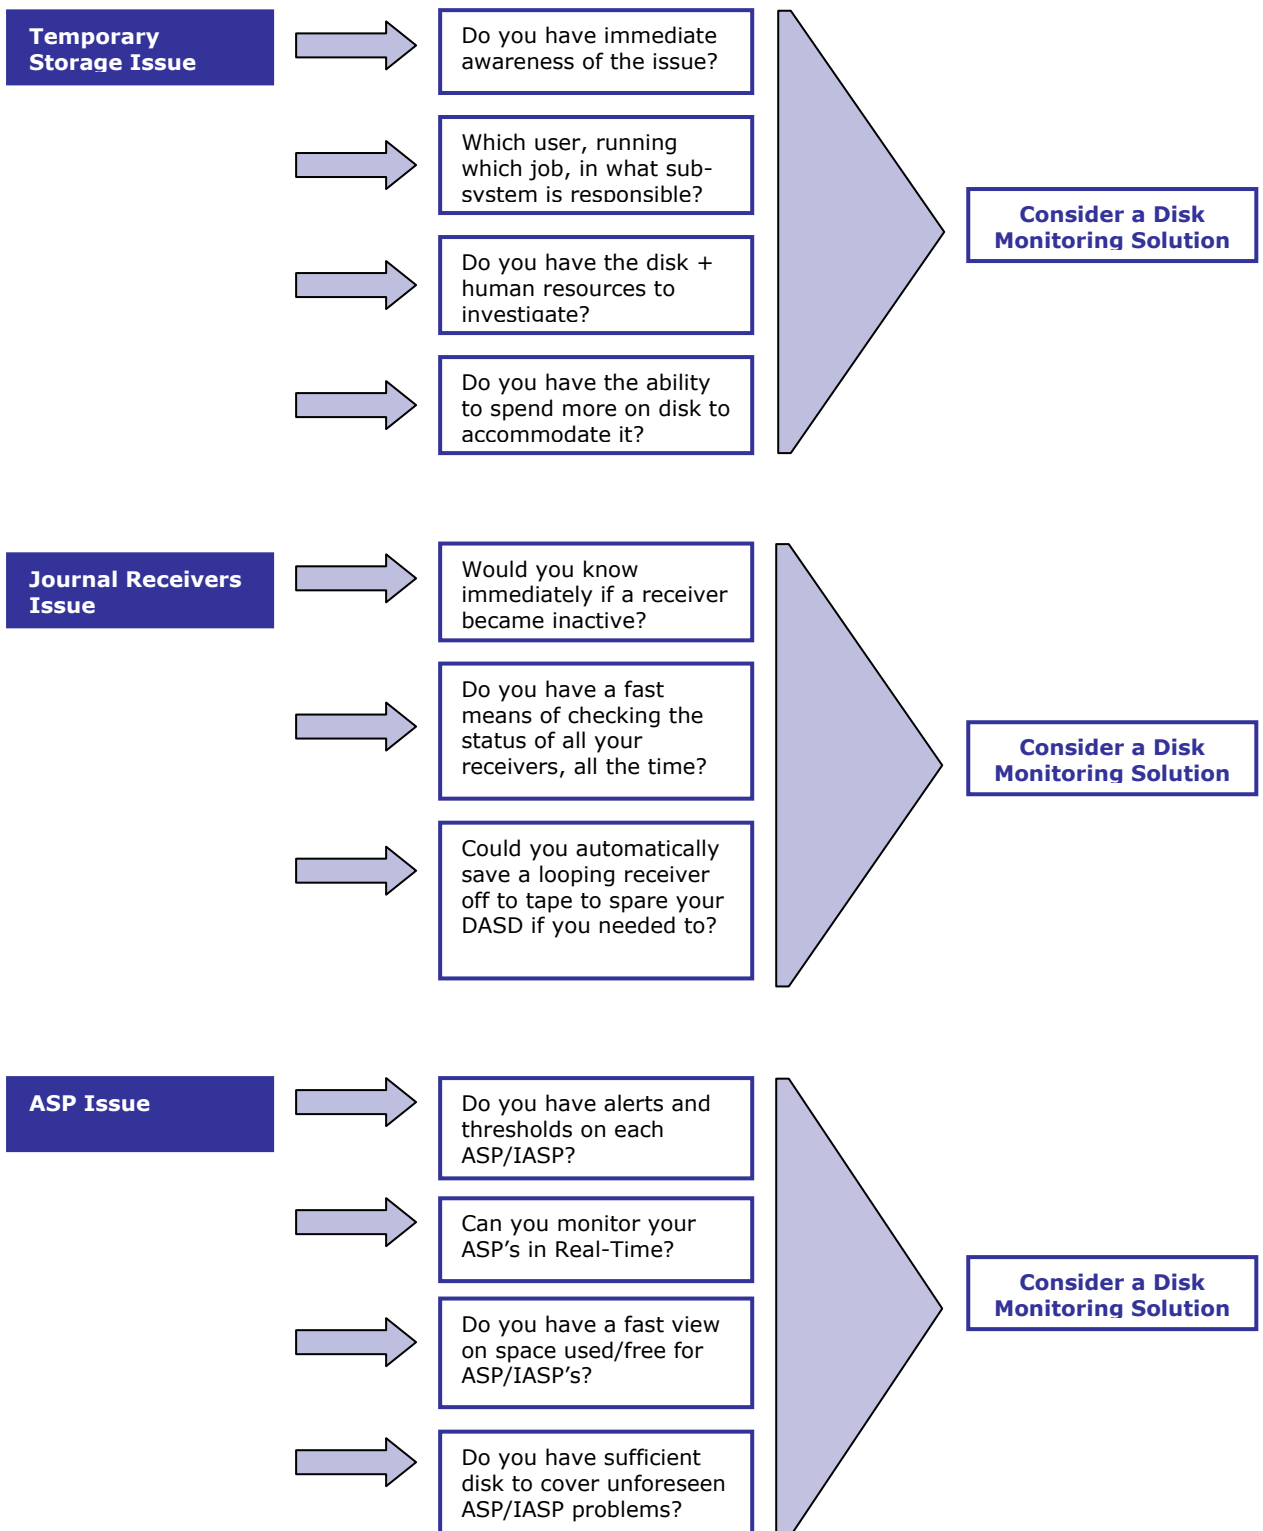
A similar issue exists for unchecked ASP's. Overflow of a user ASP can leave data at risk of being lost altogether should the system crash. Independent ASP's, employed in clustering and Global Mirroring environments are unable to overflow and as such require regular monitoring to prevent potential loss of data.

For System i shops running High Availability software that rely on Journal Receivers, such as the Audit Journal, issues can arise when a looping job causes the receivers to continually fill up. As the availability and audit-ready status of these files are crucial to the overall compliance of Sarbanes-Oxley, Managers may be obliged to make DASD available to them in such a circumstance. Given the correct monitoring solution and the same circumstance, Managers would have the far more economical option of running the receivers to tape whilst the problem was resolved. Commonly, issues impacting disk can form an epicentre of trouble that radiates out to other areas of the system. This is why it is not sufficient to monitor the disk usage in isolation, but rather all the elements that could impact disk. For example, a Journal Receiver may become inactive which can not only degrade system performance for users but also, should the system fail and 'switch over' as per High Availability protocol, data could be permanently lost.

In each of these three instances, Managers that practice a less reactive approach to DASD monitoring can save considerable time and money. Any solution that tackles these issues must have the flexibility to both identify problems in real-time and then help Managers to bypass the list interrogation method and pin-point them for fast resolution.

Need to Know Basis – Roadmap your Requirements for Disk Monitoring

The table below offers Managers a guide to help them evaluate their present disk monitoring requirements. If you are experiencing issues similar to the ones below and find your answers fall into either the 'no' or 'don't know' categories, don't wait until it's too late to find out if a disk monitoring solution could have prevented wasted DASD.



Other Threats to Disk

Whilst the problem areas above make up some of the most typical that can impact disk, it is important to remember that virtually *any* object on the system that encounters a problem and is left unattended, could potentially impinge on DASD resource. One such example of this is the heavily utilised QGPL (General Purpose Library). This library typically acts as a repository for save files that are stored there. Unless this is monitored frequently, the sheer number of files can quickly escalate, especially during periods of heavy processing and DASD can take the brunt of the burden as a result. Similarly, certain individual files such as history files and PM400 collection services files are prone to quick growth and should be monitored in real-time.

Checklist Your Disk Monitoring Flexibility

Any solution that you may decide to use must incorporate the flexibility to monitor a number of different attributes of multiple system elements in real-time. Furthermore, thresholds that alert Managers to impending breach levels can offer 'built in' time safeguards so adequate time can be given to investigations before any impact on users is felt. The table below gives examples of the types of monitoring metrics that can be used to ward off disk dramas. These monitors can be applied across the system to protect DASD and associated cost of 'unnecessary' disk usage. ASP Busy monitors as shown below can also keep a check on associated disk issues that impact performance rather than usage alone. Automated tasks can then be put in place to run routine operations that keep DASD resource usage at acceptable levels, such as deleting certain files after they reach a predetermined number, percentage or size.

Object and ASP Monitoring Checklist list for your System i Network:

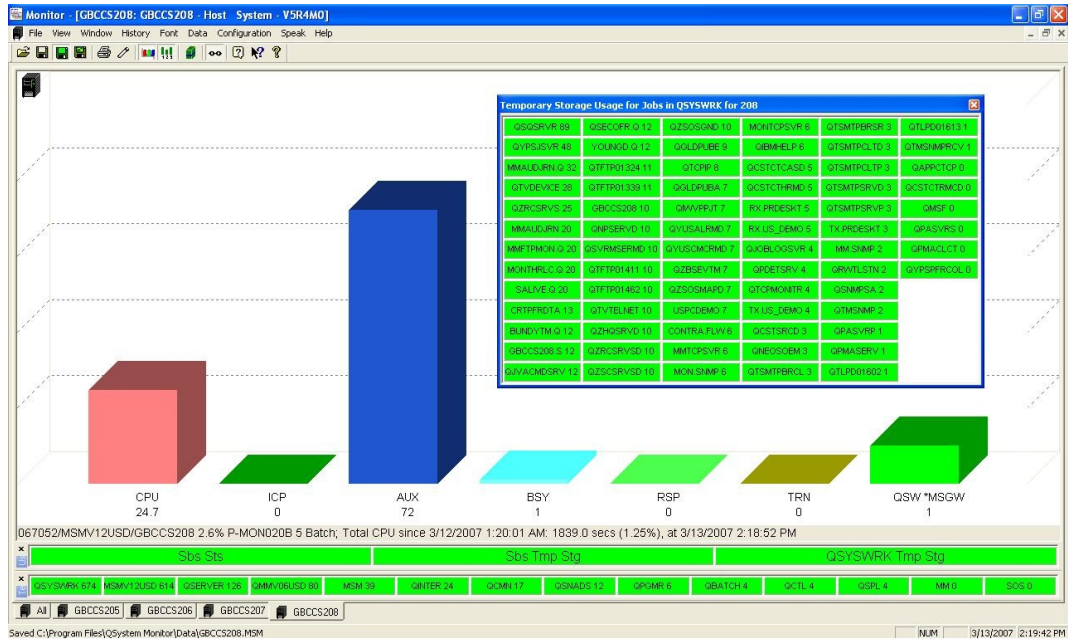
Objects	ASP Information
▪ Data Queue Entries	▪ ASP Busy
▪ Data Queue Percentage	▪ ASP Busy Unit #
▪ Distribution Queue Entries	▪ ASP Busy Unit %
▪ File Record Count	▪ ASP Space Free
▪ File Deleted Record Count	▪ ASP Space Used
▪ File Deleted Record Percentage	▪ ASP Status
▪ Journal Receiver Size	▪ ASP Usage
▪ Network File Count	
▪ Object Size	
▪ Object Count	
▪ Spooled File Count	

Why has the Disk increased in the Last 2 hours?

Implementing the right approach to disk monitoring means Managers and Operators facing on-the-spot disk dilemmas have the resources at hand to not only answer formerly impossible questions about the where and why of sudden usage but also to take appropriate action. Keeping them pro-active to potential threats to DASD, spikes

in file sizes, counts and object growth and the status of ASP or IASP could be all that is required to keep the system robust, budgets in check and time well spent.

The Screenshot below is an example of how Company X from our case study, might have avoided their Temporary Storage issues with the benefit of real-time disk monitoring. Alerted to a sudden surge in temporary storage, they could have immediately identified the sub-system and job responsible without ever having to waste valuable time interrogating the system.



Ongoing problems that impact DASD are uncompromising in their ability to put the pressure on any network so whether you pay now for more disk or pay later in inconvenience, associated problems and financial consequence, your disk issues demand attention. If the issues covered in this white paper sound familiar or you're struggling with other DASD related issues, contact the systems management experts, CCSS, to discuss how we can assist you.

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 and Sweden. www.ccsltd.com