

# COMMUNICATING WITH SAM

## High Availability Advances: Network Friendly Architecture



Sam Johnston

### Question:

We have investigated high availability solutions for iSeries for several years however we have never been able to justify the high initial cost for iSeries hardware and HA software as well as the ongoing operating expense associated with the network these solutions require. How do we minimize initial investment and minimize the number of lost transactions without severely impacting our network and system performance? What recommendations do you have?

### Answer:

There have been several significant developments recently that can address each of your concerns. In February this year IBM announced special high availability pricing for model 810s allowing smaller sized companies to more easily justify the initial expense of the iSeries hardware.

IBM has also recently expanded the number of qualifying HA Business Partners to include more cost effective solutions. These new HA solution providers have built their application around the latest remote journaling technology in OS/400. The combination of new applications and remote journaling options can have a significant impact on reducing HA costs along with reduced workload on the source machine and the latency between source and target machines. This can significantly improve the ROI for High Availability compared to the legacy HA solutions built on harvesting local journals that have been around for many years.

The remote journaling function of OS/400 offers a reliable and efficient method to transfer journal entries from a source to a target machine. Since this is a built in function of the base OS/400 code and not a separate product or feature it is implemented in the licensed internal code layer. The main benefits of this are that it eliminates the need to create copies and store harvested journal information on the source machine. It also reduces the CPU workload on the source by eliminating the processing required to harvest the journal and puts it where it belongs on the target machine. This approach can lower CPU and DASD workload on the source machine often more than 10%. This means you probably won't have to upgrade your source production system to accommodate HA workload.

In order to have a workable HA solution for today's iSeries you must consider replicating more than OS/400 database and access path information. In V5R1

remote journaling was enhanced to broaden the types of objects that can be journaled to include data queues, data areas and IFS files. As the list of object increases so does the amount of journal traffic so you want to ensure the underlying communications method is as efficient as possible. Legacy HA solution providers that harvest local journals on the source machine had to build their own proprietary transfer function into their product before remote journaling became available. This adds to the cost of legacy HA solutions, and results in lower overall performance capability. The next two figures show the difference remote journaling can make to the source machine workload.

### Legacy Approach (Figure 1)

Notice how many times the information must cross the machine interface to the HA software on the source machine in **Figure 1** and compare that to **Figure 2**. ▶

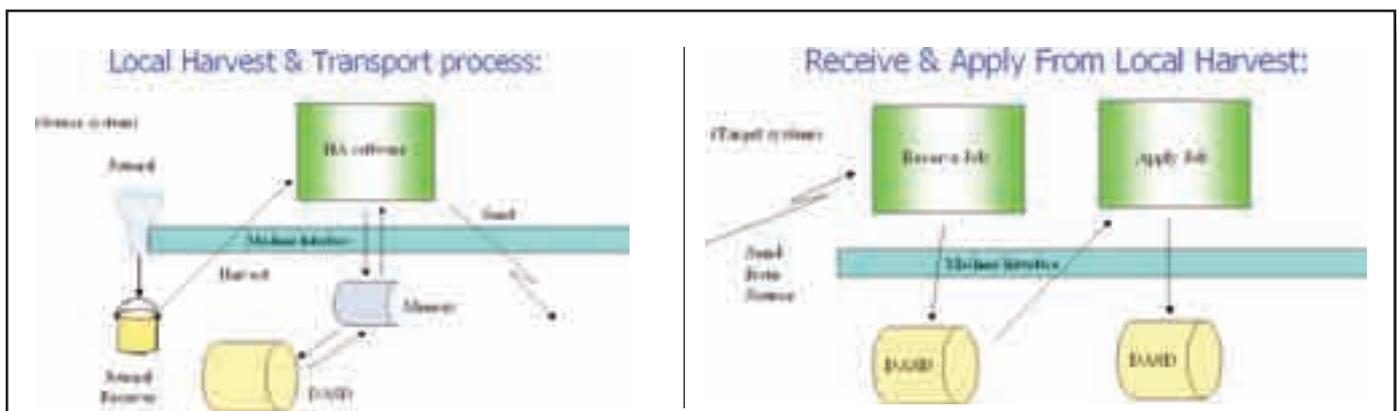


Figure 1

## Figure 2: Remote Journal Based Solution

Once we have accepted the superiority of remote journaling based HA applications we need to decide how we want to maintain the journals. There are two options here for consideration. Like everything else in life they are trade-offs so you need to determine which meets the need of your business.

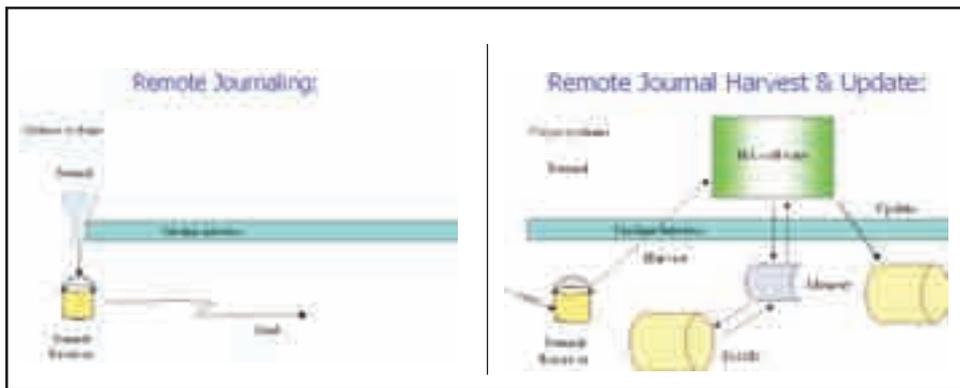


Figure 2

The first option is **Synchronous Delivery Mode**. This method guarantees the journal entries are in the main storage of the target system before the control is returned to the application on the source system. This is really the only way to guarantee the recovery point is the last transaction. This may be critical to some organizations. However this may also have a negative effect on application performance from an end user point of view if your systems are widely geographically spread out and communications are over a WAN (wide area network). The end users must wait for the synchronous acknowledgement response from the target system before completing the transaction.

The second option is to use **Asynchronous Delivery Mode**. Sending entries this way means that they are sent to the target machine sometime after the control is released back to the end user application. This method would not guarantee the recovery point is the last journal entry however the other positive business benefits are likely to mitigate the cost of a small number of lost transactions in the event of a failure. In testing asynchronous delivery IBM

found it to be extremely fast and efficient. “In every instance we never saw more than 5ms transpire between the time we produced the journal entry on the source system and the time it was sent down the communication wire” – IBM Striving for Optimal Journal Performance. Provided that your communication infrastructure is sized to handle the traffic volume asynchronous delivery is quite capable of providing a reasonable recovery point

for most business needs. If your plan is to keep the transport overhead to a minimum then asynchronous is the way to go.

The potential for a communication bottleneck exists in any situation where information must be exchanged between two systems regardless of which solution you choose. The performance of the end user application can also be significantly affected, or the latency between the systems can increase beyond desired levels, if the network infrastructure and the available bandwidth are not sufficient. The good news is that the cost of high-speed services has reduced significantly over the past few years, so the communication infrastructure needed to link the source and target machine now represents a smaller portion of the overall HA investment. While communication costs have generally decreased low latency network solutions are still expensive and having adequate bandwidth is still critical to the success of the HA solution. That being said, the new generation HA solutions are more flexible in that they can be supported with lower bandwidth, with the cost being the increased risk of losing a transaction in the event of system failure during the

delivery lag. This means that the new solutions offer greater granularity in the additional operating expenses needed to support the solution to allow customers to more easily calculate the breakeven point between the cost of increasing bandwidth and the cost of recreating a lost transaction.

Now that we have discussed a few ways the new HA solutions and offerings from IBM can lower the cost of HA we should take a look at the causes and increasing cost of downtime. Most downtime is scheduled and planned, and only a small percentage is unplanned in disaster situations. The planned causes include daily backups, reorganization of files to reclaim disk space, application software updates, IBM OS releases and PTFs, hardware upgrades.

Many business pressures are demanding continuous operations from our mission critical system. Globalization and centralization, the economic pressures of demanding more from less, and WEB based services means that the window for these scheduled activities is shrinking to zero. The cost of downtime is also increasing. The number of employees that are idle during an outage due to their dependence on production systems is increasing, and the work they do is increasing in importance. Their wages alone can be a significant cost. The potential lost customer business is also a factor that adds to the cost of an outage. New HA solutions (by decreasing the operating expenses needed to support the solution, and improving the business case) make it easier to deliver continuous availability and to accommodate the need to complete system maintenance tasks.



**Sam Johnston** is a partner and Chief Technology Officer of Intesys Network Communications Ltd., providing value-added networking and e-commerce solutions to the iSeries community. He can be reached at (416) 438-0002 or via e-mail at [sjohnston@intesys-ncl.com](mailto:sjohnston@intesys-ncl.com). Any TUG member wishing to submit a question to Sam can forward their typewritten material to the TUG office, or to Intesys. The deadline for our next issue is Friday June 11, 2004.