

# **Exploding the 6 Common Myths of iSeries High Availability**

An Executive Overview



[www.iterainc.com](http://www.iterainc.com)

## 6 Common Myths of iSeries High Availability:

### Myth #1

Only larger enterprises can afford the total cost of ownership of high availability software

### Myth #2

A dedicated person is usually required to monitor and manage high availability software

### Myth #3

The high availability role-swap process (switchover) is typically cumbersome, and often unreliable

### Myth #4

Myth # 4. No real difference exists between the two methods of HA data replication: remote journaling and proprietary harvest-and-send

### Myth #5

The longer a high availability vendor has been in business, the more functional and reliable it is

### Myth # 6

High availability software only benefits companies that have a second iSeries (or plan to obtain one)

## The iSeries high availability landscape began to change about three years ago when two significant events occurred...

economically develop an entirely new generation of iSeries high availability software built natively on the game-changing technology of IBM's remote journaling. These two events have changed the HA market by lowering the cost of high availability software, and improving reliability and ease-of-use. Because of this, what used to be fact only a few years ago is simply no longer true.

Given that you picked up this document, it is likely that you are trying to determine if your company can benefit from high availability. Or perhaps you believe your company would benefit from HA, but think solutions are either too expensive, too complex, or too time consuming to operate (or all of the above). It could also be that you are frustrated with your current HA software and are looking for an alternative. Regardless, let's look at and dispel the more common myths surrounding iSeries high availability:

### Myth #1—Only larger enterprises can afford the total cost of ownership of high availability software.

A few years ago, it was common for a high availability software license to cost many tens of thousands of dollars for even the smallest iSeries/400 model—and this was only a portion of the total cost of ownership. Other factors included the cost of a second iSeries, installation and training, annual software maintenance fees, and the ongoing costs of managing the system (typically requiring 20-40 hours a week). When all costs were considered, only the largest enterprises could cost justify a solution. That's why HA software vendors mostly focused on Global 2000 enterprises. All other companies simply had to cope with the downtime resulting from system maintenance tasks, and had to accept significant data loss and recovery time in the event of a system failure or disaster. But the story is very different today.

The total cost of ownership (TCO) of a robust high availability solution is just a fraction of what it was a few years ago—especially if you shop wisely. Many factors have contributed to this:

- *Decreasing cost of hardware:* It is no secret that the price/performance of iSeries machines has improved significantly in recent years. Purchasing a second iSeries (AS/400) capable of acting as a suitable backup machine is only a fraction of what it used to cost. In addition, recent IBM announcements of new iSeries models configured and priced especially for high availability make buying a second machine much more affordable when it is purchased with high availability software.

First was IBM's introduction of remote journaling in release V4R2 of the OS/400 operating system. Journaling is needed in any high availability (HA) solution to track changes to data files and other journalable objects, and remote journaling made it possible to rapidly extend copies of journal entries to a remotely connected iSeries. The advent of remote journaling triggered the second significant event: an opportunity for an agile company like iTera to

**More about the affordability of iSeries high availability solutions in the iTera white paper:**

*iSeries High Availability: It's More Affordable Than You Think.*

Available at:  
[www.iterainc.com](http://www.iterainc.com)

high availability software.

- *Decreasing cost of HA software:* Full-featured high availability software is available today for a fraction of its former price. When iTera's Echo<sup>2</sup> High Availability software came on the market in 2000 it cost significantly less than other products, yet the product delivered levels of reliability and functionality that inspired scores of companies to switch.

An important note about price: Many HA vendors have recently cut functionality to offer "lite" or scaled-down versions of their software, mostly because they recognized the success iTera is having among the small and medium sized businesses (SMB) segment. But iTera has never cut its features and functionality to bring down the cost. Actually, the opposite has occurred. iTera continues to add functionality that makes Echo<sup>2</sup> even more powerful and easier to use, yet Echo<sup>2</sup> still costs the same or less than the scaled down offerings of other HA vendors. As with any software purchase, it is critical to closely compare features and talk to multiple customers of each vendor.

- *Decreasing cost of system monitoring/management:* iTera has dramatically reduced the amount of time needed to manage and maintain high availability software. Operators that used to spend 20-40 hours per week managing other HA products, now spend less than 30 minutes a day managing Echo<sup>2</sup>. In fact, most Echo<sup>2</sup> customers spend only 10-20 minutes. This directly translates into a large cost saving as staff is freer to focus on other IT priorities.

iTera has reduced the time-to-manage for HA in three significant ways: 1) the product is simply more reliable than most other HA products; therefore, less problems occur in the first place; 2) many features automate functions of the software that would otherwise need to be handled manually; and 3) screens are designed for easy access to critical information. More about these points in Myth #2.

- *Decreasing cost of communication bandwidth:* Many companies wisely choose to locate the backup machine offsite—from a few miles to several hundred miles away—to improve the chances of continuous availability in the unlikely event of a site disaster at the location of the production machine. The good news is the cost of the communications bandwidth needed to remotely connect two machines has plummeted in recent years. Currently, each MB of communication bandwidth has an average cost of about \$50 per month, while four years ago each MB had an average cost of about \$800 to \$1,000 per month.

**Myth #2—A dedicated person is usually required to monitor and manage high availability software.**

Many HA products require from 20 to 40 hours per week to monitor and manage. But this does not hold true for all HA products. iTera customers will attest that it rarely takes over 30 minutes a day to monitor and manage the product; most customers are able to do this in 10-20 minutes. The primary reason for this difference is the improved reliability, but equally important is the ease-of-use that can be attributed to the very rich autonomic layer that iTera created between the technology and the user. The following details many of the ways that iTera reduced the time needed to manage high

**Speak with customers of all HA solutions, including iTera's and learn for yourself how much time each spends monitoring and managing the different HA products**

**Smooth, reliable role swapping is one of the most critical capabilities of high availability software**

availability:

- Echo<sup>2</sup> High Availability prevents problems from occurring in the first place. For instance, if an object is moved to another library or renamed, or if a new object is added in a designated library, Echo<sup>2</sup> automatically detects this and makes sure the object is properly replicated. This is one example of how the product proactively heads-off problems and reduces vulnerability.
- If Echo<sup>2</sup> detects an out-of-sync condition with an object, it is often resynchronized automatically—usually before the operator detects the situation on the system monitor. In addition, Echo<sup>2</sup> normally does this without interruption of the replication process, requiring restricted states, or suspending the replication of all other objects. Self-healing functions are one of the many autonomic capabilities included in Echo<sup>2</sup>.
- The Echo<sup>2</sup> system monitor was created to make it very simple for users to see where a problem or potential problem might be developing in the replication process. Most critical indicators are on a single screen, and all others can be seen with the push of a function key. If a problem appears on the monitor, message logs and other resources make it easy to spot the source of the problem.
- Echo<sup>2</sup> significantly automates the “role-swap” process (the backup system taking on the role of the production system during a planned or unplanned downtime event). Automation makes this critical process easier to use, less time consuming, and safer to execute. Most importantly, having confidence in the role swap encourages routine testing of this critical process to ensure that it executes smoothly when it is needed. More about this in Myth #3.

### **Myth #3—The HA role-swap process (switchover) is typically cumbersome and often unreliable.**

When the role-swap process executes during a planned or unplanned downtime event, the backup system assumes the role of the production system (and vice-versa). **The entire purpose of the role swap is to quickly move necessary users, devices and processes to a fully synchronized backup environment during system maintenance, failure or disaster.**

Like so many of the other HA myths, this one began legitimately from actual experience by many companies who have used other HA products. The primary reason that the role swap is often a pain is because when operators prepare for role swap they usually find objects that are not synchronized between the two systems. Even if full synchronization is actually verified, the role swap can be fraught with problems because several actions need to execute seamlessly:

- Production system put in a restricted state
- A variety of different HA jobs ended and restarted on each system
- Journaling processes on each system ended and restarted
- IP addresses on each system reassigned
- User profiles on each system disabled and enabled

## Isn't it ironic?

**Many shops that use other brands of HA software don't role swap their systems to prevent downtime. That's because they are afraid of incurring even more downtime than they were trying to prevent in the first place!**

- Production devices pointed to the backup system
- Interfaces switched to the backup system
- System prepared for reverse of the role swap
- And many more

The more processes that need to be done manually, the more time-consuming and error-prone the role swap. There are many iSeries shops that have had HA software installed for months or even years yet have never tried to perform a role swap. This is because of a lack of trust in the accuracy of replicated data and/or a lack of confidence in the reliability of the role-swap. IT people understandably are not willing to unnecessarily risk additional downtime or data corruption. Are you beginning to see the irony here? **Many shops don't use HA to prevent downtime because they are afraid of having downtime last longer than it otherwise would!**

A crucial portion of the Echo<sup>2</sup> installation process is a very thorough role swap training and testing procedure in the customer's environment. This is done with one goal in mind: achieve a high level of confidence in the role swap so that it will be tested regularly. iTera recommends that customers test the role swap on a quarterly basis (monthly is better), which is in line with IBM's own recommendations.<sup>1</sup>

Once Echo<sup>2</sup> is properly configured and users are trained, the entire role-swap process typically executes in less than five minutes. One Echo<sup>2</sup> customer clocked a role swap at 24 seconds from the time a function key was pressed to the time that both systems assumed their new roles and replication resumed. Keep in mind that these times don't include the time to end subsystems and physically move devices and interfaces.

Something else to keep in mind about the role-swap: if the process works smoothly and is tested regularly, it will reduce the amount of downtime when the system has to be brought down for maintenance, and it is far more likely to work quickly and smoothly in the event of a system failure or disaster ("failover").

References from customers that regularly and successfully test the role swap are absolutely critical when comparing the reliability of this process in various HA software products. When you speak to references, be sure to have them describe the tasks they need to perform during the role swap process and how much time the entire process takes.

## **Myth # 4—No real difference exists between the two methods of HA data replication: remote journaling and proprietary harvest-and-send.**

The foundation of iTera's Echo<sup>2</sup> High Availability software has always been IBM's remote journaling data transport engine, which is part of the OS/400 operating system. The other types of data transport engines in high availability products are proprietary processes that were developed by HA vendors.

Before going into the differences between remote journaling and the proprietary methods, it is important to emphasize that all HA solutions require standard journaling functions to track data changes that are made to database files and other journalable objects selected to be mirrored by the HA software. As changes are made

**What differentiates the efficiency and speed of HA data transport engines is the machine that the process uses to "harvest" the data from journal entries.**

to these files on the live production environment, the HA software “harvests” the data-change information from the journal entries and “applies” this information to copies of the objects that exist on the backup environment.

What differentiates the efficiency of HA transport engines is the machine where the data is harvested from journal entries. HA software that uses a proprietary process must harvest information from journal entries on the production (source) system. That’s because the journal entries only exist there. However, HA software that uses remote journaling is able to harvest information from journal entries on the backup (target) because “carbon copies” of the journal entries are rapidly distributed to the backup by the remote journaling process. See the following illustrations:

The red circles in the illustration on the right indicate points of work that are caused by activity above the machine interface level

### Proprietary Source-Harvest & Send

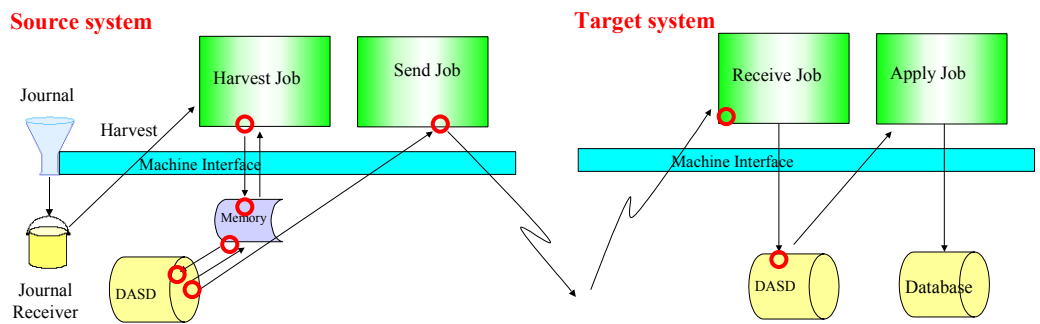


Figure 1

### Remote Journaling Send & Target Harvest

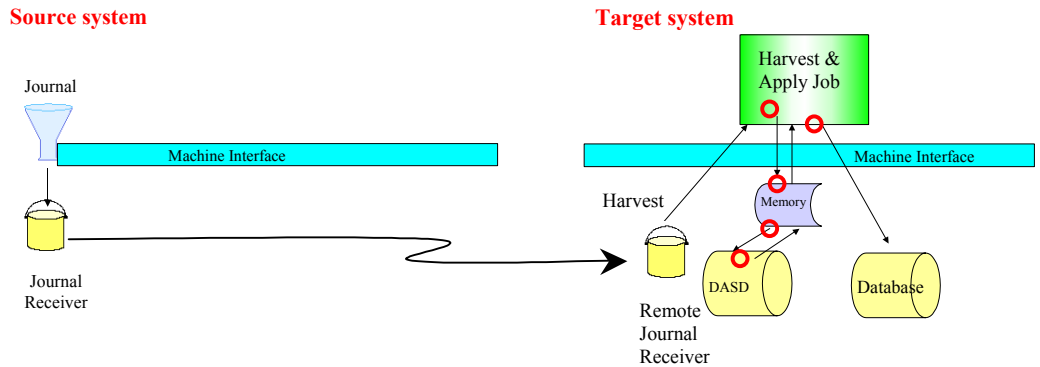


Figure 2

Figure 1 illustrates a proprietary “source harvest & send” process. As you can see, most of the work (filtering, validating, sending, etc.) occurs on the live production environment and above the machine interface (outside of the operating system). All of this adds a great deal of system overhead to the production environment, and also makes data more vulnerable. More about this shortly.

Figure 2 shows a remote journaling replication process. Here, all of the filtering and validating is moved to the backup machine, and the send process is handled entirely by remote journaling, all within the OS/400 operating system.

**Remote journaling provides several proven advantages to iSeries high availability performance and reliability****A detailed analysis of the advantages of remote journaling in high availability solutions can be found in the iTera white paper:**

*[The Benefits of Remote Journaling in iSeries High Availability](#)*

Available at:  
[www.iterainc.com](http://www.iterainc.com).

Remote journaling provides several distinct advantages:

1. Little overhead is incurred on the production system by the HA processes because most of the work created by harvesting data from journal entries is moved to the backup environment as it can be done on copies of journal entries created by the remote journaling process. IBM tests have shown that this can result in a performance hit of up to 10% on the production environment.<sup>2</sup> However, many iTera customers that have switched from another HA product (that previously had a proprietary harvest-and-send process), have told us that their previous product added 20-30% additional overhead.
2. Even though remote journaling moves the harvesting workload to the backup system, it doesn't mean the same amount of overhead then occurs on the backup. First, there is no separate "send" process on the production system as this is handled by remote journaling on the production system. Second, there is no "receive" process above the machine interface on the backup as this is also handled by remote journaling on the backup.
3. Since IBM's remote journaling sends and receives data entirely beneath the machine interface, the time needed to place data on the communication wire is often substantially reduced. In tests by IBM, it never took more than 5 milliseconds to place a journal entry on the wire, even at extremely high transaction levels.<sup>3</sup> That's because this occurs entirely within the operating system; in other words, the process has a very high priority for system processor time because it is part of the OS. With remote journaling, it is extremely rare to have any backlog of journal entries waiting to be sent to the production machine.
4. Without remote journaling, there can be significant data vulnerability because backlogs of data are often waiting in the queue to be sent to the backup. That's because data-change information needs to be harvested from journal entries on the production system, then queued for sending to the backup, then transmitted to the backup by a proprietary "send" job. All of this happens above the machine interface, therefore it has a lower priority for processor time. In addition, other processor demands on the production system can slow the process even further, which can cause backlogs, especially when there are peaks in transaction volumes and/or high priority workloads.
5. Because latency is rare with remote journaling, high availability products that use remote journaling are better able to handle high transaction volume environments. iTera customers with these type of environments that had used another HA product prior to Echo<sup>2</sup> attest that they were often hours behind in their replication process before switching to Echo<sup>2</sup>; after switching to Echo<sup>2</sup>, this backlog or latency had disappeared altogether.
6. Remote journaling has its own built-in auditing functions that keep track of the journal entries sent to the backup machine. If acknowledgement is not received that a sent journal entry was delivered, the process automatically resends the journal entry. This process guarantees that journaled information reaches the backup.

**The older the software, the more likely it gets entrenched in its original design unless it is periodically redeveloped to incorporate new core technologies**

### **Myth #5—The longer a vendor’s HA software has been on the market, the more functional and reliable it is**

Some of iTera’s competitors have been selling HA software since the mid 1980s, while iTera brought Echo<sup>2</sup> to market in 2000. It is common to believe that if a product has been on the market longer, it is better. Our competitors will quickly tell you that this longevity and accompanying experience automatically make their product superior. But there is another side to this argument (which applies to software of all types): the older the software, the more likely it still has its original design. In other words, new technologies are either not incorporated or they are retrofit into the original architecture, which can make for a less elegant use of the technology, and often, a less efficient product.

During the last few years, two technologies have emerged from IBM’s iSeries division that have been shown to improve the efficiency, intelligence, and reliability of iSeries high availability software. One is remote journaling; the other is clustering.

The benefits of remote journaling in iSeries high availability have just been detailed in Myth #4. Not long after remote journaling was introduced by IBM, iTera’s Echo<sup>2</sup> High Availability software was created, and from the very beginning it incorporated this cutting-edge technology. Because of the market success of iTera, other HA vendors were forced to acknowledge the benefit of remote journaling, and have finally begun to offer this technology. However, instead of completely redesigning and rewriting their solution based on remote journaling, most simply retrofit the technology into an older architecture.

Clustering is another technology that recently appeared in the OS/400 operating system. When a cluster is defined within cluster resource services of OS/400, some of the benefits include being able to distribute workloads among all resources defined in the cluster (i.e. multiple iSeries systems), and managing all resources from a single workstation. Integrating this technology with HA software brings the advantage of automatically detecting node failures and executing self-correcting actions—particularly the rapid initiation of a “failover” to another designated node.

A few years after Echo<sup>2</sup> was introduced, iTera invested in a full redesign of the software to incorporate the many benefits of clustering. As with remote journaling, most other HA vendors only retrofit clustering functions into existing HA technology rather than rewrite software from scratch to incorporate these. An add-on benefit of iTera embedding the many functions of clustering in Echo<sup>2</sup> is that several self-correcting and other autonomic benefits were added—many that benefit all customers whether clustering is enabled in OS/400 or not.

The bottom line is this: iTera’s agility allowed the company to quickly and fully incorporate critical changes in technology, making a significant contribution to a more reliable, versatile, and easy to use product. Fully a third of iTera’s customers have abandoned the investment already made in another brand of HA software to switch to Echo<sup>2</sup>. It is clear that they recognized a significant difference between the technologies offered by iTera and other HA vendors.

**A fast ROI can be achieved with HA software without necessarily needing a second iSeries**

### **Myth # 6—High availability software only benefits companies that have a second iSeries (or plan to obtain one)**

Some shops believe that it only makes sense to evaluate high availability software if there is another iSeries available in the computing environment that can act as a backup, or there is an intention to purchase another machine for this purpose. But many system managers aren't aware that they can achieve a good return on investment (ROI) on HA software without having a second iSeries.

One of the most fundamental benefits of HA software is the elimination of downtime for daily and system tape backups, and this can easily be accomplished using a single-system replication topology. In this topology, essential data is replicated into another OS/400 environment defined in a separate logical partition (LPAR) on the same machine. This allows tape saves to be done over the mirrored data rather than over live data on the production environment. This function alone can eliminate hundreds of hours of downtime each year, which can quickly pay for a HA solution.

Single-System Replication




Figure 2 – Single-system replication environment

Keep in mind that it is possible to role swap to the backup LPAR environment in order to eliminate other kinds of downtime, such as operating system upgrades.

An additional, often overlooked benefit of single-system replication is that reporting and data warehousing functions can be performed on mirrored data in the backup environment. If the backup environment is configured properly, this can noticeably improve system performance by reducing disk-arm contention in the production environment. In addition, moving reporting and data warehousing to the backup environment can provide more options for enhanced security settings.

### **Conclusion**

Because changes in technology from IBM allowed the entry of more agile companies like iTera to enter the market, the cost of high availability software has dropped, and a significantly more automated, more efficient, and more user-friendly product has emerged. In fact, the innovations that have been brought to market by iTera have actually changed the iSeries HA landscape in many ways:

- Competitors are finally starting to include remote journaling in their offerings—a critical technology that iTera embraced from the very beginning
- Competitors are starting to offer “lite” and otherwise scaled-down versions of their HA products in an attempt to compete at the same low price points that Echo<sup>2</sup> has enjoyed for years. But there is nothing “lite” or scaled down about Echo<sup>2</sup>

**What used to be true about iSeries high availability has clearly changed**

**For more information about iSeries data availability strategies and iSeries high availability software, please obtain iTera's other white papers at [www.iterainc.com](http://www.iterainc.com).**

- Hundreds of SMB-level companies are reaping the benefits of high availability software as they realize that it doesn't take the enormous financial and personnel investment it once did.
- Scores of frustrated companies that already have HA are switching to iTera and finding dramatic HA improvements in reliability, performance and ease of use.

As a result, things that were considered fact about high availability a few years ago have clearly changed. This is very good news for IT managers and operators seeking to improve the data availability of their iSeries.



iTera, Inc.  
801-799-0300  
[info@iterainc.com](mailto:info@iterainc.com)  
[www.iterainc.com](http://www.iterainc.com)

© Copyright 2004, iTera, Inc. No portion of this document may be reproduced without permission from iTera, Inc. IBM, eServer, and iSeries are trademarks of International Business Machines Corporation. All other trademarks are property of their respective companies. Rev 05/04

Notes:

1- IBM - Independent Software Vendors (ISVs) Company Criteria for VAE Selection ([http://www-1.ibm.com/servers/eserver/iseries/ha/pdf/ISV\\_criteria.pdf](http://www-1.ibm.com/servers/eserver/iseries/ha/pdf/ISV_criteria.pdf)). Page 3.

2 - IBM Redbook: Striving for Optimal Journal Performance on DB2 Universal Database for iSeries, Chapter 6.5.6. The entire Redbook can be found at [www.ibm.com/redbooks](http://www.ibm.com/redbooks). Search under book ID#: SG24-6286-00

3 - *ibid*, Chapter 6.2.1