

# PowerHA SystemMirror for i

## Ready to switch

Jenny Dervin

April 12, 2011

PowerHA SystemMirror for i is low maintenance, easy to manage, and always ready to switch. Read about how PowerHA SystemMirror has a replication technology to fit every IBM i environment.

Are you looking for a high availability solution for your business which is always ready to switch? Look no further than IBM and the PowerHA SystemMirror for i product. PowerHA SystemMirror can be configured to minimize both planned and unplanned outages and can also be used for offline backups. Setting up your PowerHA environment is not difficult, and once it is set up, it requires minimal administration. PowerHA SystemMirror is deeply integrated into the operating system and monitors for and reacts to both hardware and software failures.

First, we'll explain the wide range of replication technologies available within PowerHA SystemMirror, and then we'll cover its reliability and ease of use advantages.

## Replication technologies

Most IBM i high availability products monitor for changes to an object, replicate those changes to the backup node, and then apply the changes to the object on the backup node. This type of replication is commonly called logical replication. PowerHA SystemMirror replicates at the disk IO level, which is referred to as hardware replication.

PowerHA SystemMirror has no knowledge of objects. Instead, every time a disk write is issued, that write is done on the primary node and sent over to the backup node and written to disk there, as well. Both synchronous and asynchronous replication is available. With synchronous replication, the write on the production node does not complete until it has been received on the target node, so you are guaranteed identical copies on each side. With asynchronous replication, order is preserved, but the write is allowed to complete without waiting for an acknowledgement from the target node that the write has been received there.

Since PowerHA SystemMirror does not operate at an object level, the system administrator specifies which objects must be replicated by creating an independent auxiliary storage pool (IASP) and placing the production data and possibly also the production applications in the IASP.

Everything within the IASP is replicated. There are some object types which cannot be placed in an IASP. Many of those object types specific to the production environment can be synchronized via another PowerHA technology called Administrative Domain. We cover that technology in more detail a bit later in this article.

Typically, migrating applications into an IASP environment does not require changes to the application itself but only work management related changes, such as making sure that users and jobs are connected to the correct IASP. Many of our more popular ISV applications are now enabled for IASPs, with more declaring support every day.

Hardware replication is a broad category, and within it, there are several different replication technologies available as part of PowerHA SystemMirror.

## Geographic mirroring

The first type of hardware replication made available was geographic mirroring. The IBM i operating system handles the replication at a disk page level. Geographic mirroring can be used with any type of disk, internal or external, and is also agnostic to whether the disk is natively or virtually attached. As with the other hardware replication technologies, although the IASP is not accessible on the backup node while its actively replicating, the node itself is available to do other work. Switchover or failover involves varying off the IASP on the production node and varying it on the backup node.

With synchronous geographic mirroring, the backup copy of the IASP is guaranteed to be equivalent to the production copy. The disk write on the production system does not complete until the page has reached the backup system. Typically, synchronous geographic mirroring is used when the production and backup systems are within a short distance, such as within the same city. Sufficient communications bandwidth will result in the best response times on the production system.

Starting in version 7.1, asynchronous geographic mirroring is also available. This allows for a greater distance between the production and backup systems without affecting the local response time. The environment must still be properly sized and provide sufficient bandwidth to keep up with the average production transaction rate, but response time is not be negatively affected if replication falls behind during a peak period.

## Metro Mirror and Global Mirror

For customers with access to a DS8000 storage system, Metro Mirror and Global Mirror are alternative PowerHA SystemMirror replication technologies. Metro Mirror and Global Mirror make use of the copy services or peer-to-peer remote copy (PPRC) technology available on the DS8000. Instead of IBM i doing the replication, the replication is done within the DS8000. Any performance overhead of the high availability environment is therefore moved to the DS8000.

As with geographic mirroring, there are both synchronous and asynchronous options. Metro Mirror is the synchronous replication option. So again, when the disk write completes, it is guaranteed to be both on the production and the backup copy. This technology is most commonly used when the

production and backup systems are relatively close to each other. Global Mirror is asynchronous and can be used when the production and backup systems are a greater distance apart.

By integrating the Metro Mirror and Global Mirror technologies into PowerHA SystemMirror, the entire high availability environment can be managed from one interface on your IBM i system. PowerHA SystemMirror does the work of managing the storage server copy services so a seamless switchover or failover can be achieved. While it is possible to use the DS8000 copy services technology to replicate the full system, IASP replication within PowerHA is preferred by most customers. Full system replication solutions don't integrate the IBM i and DS8000 functionality, so a seamless switch is hard to achieve and requires an IPL. With full system replication, the backup server is not accessible, whereas with PowerHA IASP replication, the system can be used for other work which can be offloaded or stopped if production is switched to the server. Finally, with full system replication, you can't take advantage of the high availability environment to reduce OS upgrade outages, or for less disruptive hardware maintenance.

## Switched disk and LUN-level switching

Switched disk and LUN-level switching are technologies usually used in conjunction with another replication technology. Switched disk and LUN-level switching use only one copy of the IASP but can switch the connection of that IASP from one system to another. Switched disk can be used with any type of storage and switches the IASP between partitions on the same CEC. LUN-level switching is only available when using the DS8000 storage server and can switch the IASP between partitions or between systems. Prior to POWER7, there was a third variation of this technology, utilizing a switchable storage tower on an HSL loop, but that option is no longer available on POWER7 hardware.

Since these technologies do only utilize one copy of the data, the disk would still be a single point of failure. Typically, clients use switched disk or LUN-level switching within a data center to provide local protection from planned and unplanned outages not involving the disk subsystem. They can also use a replication technology such as geographic mirroring, Metro Mirror, or Global Mirror for another level of high availability or for disaster recovery.

## FlashCopy

For customers using a DS8000 storage system, FlashCopy has also been integrated into PowerHA SystemMirror. FlashCopy is a point in time copy instead of live replication. A FlashCopy can be generated very quickly and accessed immediately from a separate partition or system. There are options for creating a FlashCopy, depending on its intended use. A full copy FlashCopy will copy the entire IASP in the background and is commonly used for a development or test environment or for queries or business intelligence. A no copy FlashCopy will only copy original pages before they change on the production system and is typically used for backups.

## Advanced Copy Services

Advanced Copy Services for PowerHA is an IBM Lab Services offering. Advanced Copy Services (ACS) uses PowerHA technology and provides more customized and complex solutions for customers. One example is Metro Global Mirror. Metro Global Mirror is a combination of Metro Mirror and Global Mirror and provides two live replicas of the data. Metro Mirror is used to replicate

to the first backup, and then Global Mirror is used to replicate from the first backup to a second backup. Many of our PowerHA customers also use Advanced Copy Services to customize and automate their high availability environment.

## **Administrative domain**

As stated before, administrative domain is another PowerHA SystemMirror technology which can be used to synchronize those parts of the production environment which cannot be stored in an IASP. Most of the common object types used in production environments (such as user profiles, system values, environment variables, network attributes, and device descriptions) are supported within administrative domain. An administrative domain is simple to set up. Add the nodes which you would like to keep synchronized and then add the resources to be monitored. Once an object is added to the administrative domain, changes to that object on any node within the administrative domain will be synchronized across all the nodes. Administrative domain also provides flexibility in which parts of an object are synchronized. In the case of user profiles, for example, all attributes can be synchronized, or individual attributes can be selected.

## **Reliability and ease of use**

PowerHA SystemMirror is deeply integrated into the IBM i operating system and is built upon clustering technology. The clustering technology provides a low-level messaging structure which keeps all the nodes in the cluster synchronized. Low-level constructs monitor for hardware, OS, and network failures, and when a failure is detected, notifications are sent to the other cluster nodes. PowerHA SystemMirror can be configured to fail over automatically to a backup node when a failure is detected on the production node, minimizing the need for system operator intervention. If automatic failover is not desired, it is also possible to configure PowerHA to wait for direction from an system administrator.

Since everything in the IASP is guaranteed to be replicated, the administration and maintenance of your high availability environment is minimal. You won't need to spend time verifying that new objects are being replicated, as long as the library or directory is in the IASP. This in turn makes switchovers more reliable, with less chance that changes on the production system have been left out of the high availability environment.

PowerHA SystemMirror provides either a command-based interface to control your environment or a GUI interface. When using DS8000 copy services, PowerHA integrates the commands controlling replication on the DS8000, so the environment can be completely controlled using IBM i interfaces. For planned outages, switching to the backup system can be as easy as one command or one action in the GUI.

Once again, Advanced Copy Services for PowerHA provides added integration of the DS8000 into your IBM i environment, providing commands to configure and manage storage from the IBM i command line.

## **Summary**

With the various replication technologies available within PowerHA SystemMirror, maintenance of your high availability environment is pain-free. Many customers using PowerHA SystemMirror

switch from production to backup regularly, giving them added confidence in their disaster recovery plans. If your business is looking for a high availability solution, PowerHA SystemMirror for i is a great choice to make.

© Copyright IBM Corporation 2011

([www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml))

Trademarks

([www.ibm.com/developerworks/ibm/trademarks/](http://www.ibm.com/developerworks/ibm/trademarks/))