

## Introduction

Disasters happen. Your organization is not immune. They don't happen often. In fact, they are extremely rare, but all companies must prepare for the worst even while they hope for the best.

Yet disasters aren't the only availability threats. Hardware and software failures also cause outages. And essential server maintenance—including nightly backups, hardware and software upgrades, and more—occurs much more frequently. What's more, unlike disasters, good fortune won't save you from the need for planned maintenance events. Regular maintenance is essential to continue to serve evolving business needs and deal with technology issues.

Protecting your business from downtime, whether planned or unplanned, requires making one or more copies of the data and maintaining this copy or copies at a distance sufficient to account for the threat that you want to countervail. This copy may either be maintained in real time on a recovery server that is always ready to assume responsibility for production operations or it may require an action to restore the environment to the original.

There are three basic technologies employed to provide high availability or disaster recovery protection for a business IT environment. The first is periodic tape saves, which in today's fast-paced world is barely enough protection to qualify for disaster recovery because too much data can be lost between periodic saves. This leaves a choice of two techniques that provide real-time replication: either replication of changed hardware sectors in a storage subsystem, with recovery of the server from the copy of the storage subsystem, or replication of data changes to a recovery server, allowing a quick switch to the waiting server.

IBM® PowerHA® SystemMirror® for i addresses the need for high availability by replicating hardware sectors from a source storage subsystem to a matching sector address on a target storage subsystem. At the point of a failover from the production server to a recovery server, the target storage subsystem is attached to the recovery server and the business environment is recovered.

MIMIX® Availability™ is a logical replication solution that delivers high availability by replicating data changes at the conclusion of each transaction and immediately applying them to a recovery server. At the point of a failover, the recovery server is available without further action or recovery steps.

This paper explores storage subsystem-based replication technology and explains how MIMIX for PowerHA uses MIMIX technology to augment and complete the protection provided by IBM PowerHA solutions.

## Glossary

Following is a glossary of terms relating to IBM storage subsystem high availability and the IBM PowerHA for i software which adapts the storage based solution to the IBM i platform.

<b>SAN (Storage Area Network)</b>	Storage hardware that can be shared among servers.
<b>DS8000 SAN / IBM® TotalStorage® DS8000®</b>	A family of IBM SAN hardware that is designed to support the sector format used with IBM i. It attaches only to IBM Power servers running IBM i.
<b>V7000 SAN / IBM® Storwize® V7000</b>	A family of IBM SAN hardware that attaches to IBM Power servers running IBM i. It provides virtualization of the sector format which is common to other platforms so that IBM i OS can use it.
<b>Metro Mirror</b>	A licensed HA feature of IBM SANs that provides synchronous replication of data from a source SAN to a target SAN, triggered by a source sector write operation
<b>Global Mirror</b>	A licensed HA feature of IBM SANs that provides asynchronous replication of data from a source SAN to a target SAN, triggered by a source sector write operation.
<b>IBM PowerHA SystemMirror for i</b>	An IBM i licensed program product designed to implement HA using switchable IASPs in conjunction with IBM SAN HA features such as Metro Mirror and Global Mirror.
<b>ASP (Auxiliary Storage Pool)</b>	A software definition of a pool of disk units allocated to an IBM i server to which data can be assigned. Typically, a user-generated ASP corresponds to a physical arrangement of disks.
<b>IASP (Independent Auxiliary Storage Pool)</b>	An ASP with all of the necessary system information associated with the data it contains. When the IBM i server is active, the IASP can be varied off and varied on.
<b>Switchable IASP</b>	An IASP that corresponds to a physical arrangement of disks that can be accessed by two servers and can be varied off one server and switched to the other server where it can be varied on.
<b>SYSBAS</b>	The term used for the set of system ASP and basic user ASPs. It represents all of the data not in an IASP.
<b>Cluster Resource Services</b>	Functionality within IBM i that ties together the addressing of storage subsystems between multiple servers in a cluster, allowing for movement of data between servers to achieve high availability.
<b>Journaling</b>	Use of an IBM i journal construct that logs sequential object changes for use in system recovery.

## SAN-based High Availability

SAN-based HA functionality copies data sector-by-sector from a source SAN to a target SAN. The SAN, not the server, performs the replication, thereby off-loading the replication processing from the server.

There are no fundamental architectural concerns when implementing this functionality on SANs attached to Windows, Linux and AIX servers. The same is not true, however, for IBM i servers. In IBM i environments, issues associated with IBM i's unique Single Level Store architecture must be addressed when using SAN-based HA.

Two distinctive aspects of Single Level Store make it difficult to recover data right up to the point of failure and restore operations quickly, if at all, when using a SAN-based HA option.

1. Applications are unaware of sector addresses. Consequently, the cross-reference table that is maintained in memory must be rebuilt or identical addressing must be maintained on the recovery server for use when the production server crashes.
2. IBM i uses memory as a very large cache. This allows for very high performance, but it also means that the most recent data changes may still reside in memory and, therefore, may not be written to disk before the production server crashes. This lost data must be accounted for in some manner.

When relying on SAN-based HA to maintain a real-time recovery server in an IBM i environment you have two solutions to choose from:

1. Reconstruct the cross-reference table as part of the recovery process. This is sometimes referred to as an abnormal IPL because of the undetermined amount of time required to bring the server back to where the Single Level Store is operational again.
2. Go beyond a strictly SAN-based HA solution and employ a method for keeping the addresses on the production and recovery servers in sync. This is one of the functions of IBM PowerHA for i, which is described below

Recognizing IBM i's use of memory as cache, it is highly recommended to implement journaling to replicate the additional data required to recover data up to the point of failure, protecting the environment from data corruption. Configuring IBM i journal functionality accomplishes this task at the cost of additional resources and time for failover recovery.


## Full System Copy

Full System Copy is the name used to define the standard SAN-based HA solution used by other platforms—Windows, Linux and AIX—when used to provide HA on IBM i. It does not take into account the unique IBM i server architecture and, instead, treats HA as a solution restricted to the storage subsystem.

In a Full System Copy SAN HA implementation, the recovery server remains powered off until production is switched to that recovery server. At that point, the copy of the entire storage subsystem is used to power up the server, resulting in an identical server, including system name, but with a different serial number. The unique IBM i issues are then dealt with as part of the abnormal IPL.


Customers who choose this approach benefit from simplicity and commonality with other platforms, but they need to deal with three significant issues: lost data that did not get out of memory before a crash, the real possibility of the need to repair damaged data created at the time of the server crash, and the time required to recover from the damage.

Because the server is powered off until the point of switch, there is nothing for an IBM i high availability solution to do, whether it is a licensed program product from IBM or from a third-party high availability provider. It is entirely a SAN solution using the SAN interface and the basic server power-on functions.



## PowerHA-based High Availability

IBM PowerHA for i was designed to overcome the obstacles that are inherent in a SAN-based high availability solution. The two major functions on which PowerHA for i is based are Cluster Resource Services and switchable Independent Auxiliary Storage Pools (IASP). IBM PowerHA for i software uses these functions to replicate the data in an IASP to a target SAN. IBM PowerHA does not replicate data and programs that cannot or do not reside in an IASP. Nor does it protect applications that are not certified to run in an IASP.




IBM developed IBM PowerHA for i to provide high availability by combining clustering and switchable IASPs with SAN HA technology. The Metro Mirror and Global Mirror functions of the SAN provide the actual replication of changed sectors as defined in the IASP. It is most successful if all customer data can reside in an IASP and not in SYSBAS. However, in many organizations, that is simply not possible.

# Extend and Enhance Availability Protection with MIMIX for PowerHA

MIMIX for PowerHA complements IBM PowerHA to fill in the gaps in its availability protection. These gaps are the protection of data and applications in SYSBAS, support for additional nodes for local or remote high availability and disaster recovery, and the detection of events that may slow or halt a switch. MIMIX for PowerHA leverages the replication technology in MIMIX Availability, a solution that is the result of more than 20 years of partnering with IBM to develop a world-class cluster architecture that meets the HA/DR needs of IBM i customers.

## Protection for SYSBAS Data

MIMIX for PowerHA protects the data that IBM PowerHA doesn't protect, namely data that does not reside in an IASP. In doing so, it also safeguards against downtime for applications that have not been certified to run in an IASP.



## Administrative Domain

The Administrative Domain functionality found in PowerHA for IBM i provides synchronization of 19 objects found in SYSBAS and which may be important for some applications. These include environmental variables, device & job descriptions, user profiles, authorization lists and some other non-database objects. Objects are registered and then any change to the registered object on any of the servers in the high availability environment will trigger an update of that object on the other server or servers.

Administrative Domain is not designed to protect the customer's data-related objects such as database, data areas, data queues and IFS. The assumption is that all of the customer's data and applications are running in the IASP. It is the environmental variables and other system objects that are not changed very often that are being protected. Of these, the most important may be user profiles which must remain in SYSBAS.

MIMIX for PowerHA replicates 73 object types, essentially covering all of the data and applications that must reside in SYSBAS. There is a small overlap between objects replicated by MIMIX for PowerHA and the 19 objects handled by Administrative Domain, most notably user profiles. However MIMIX for PowerHA is complementary and Administrative Domain can be integrated with MIMIX for PowerHA to maintain synchronization of the objects not supported by MIMIX for PowerHA.

## Confidence in a Speedy, Accurate Recovery

An HA/DR solution is only a solution if it works when you need it, i.e., when your production server is down or some of your data has been destroyed. If that should happen, you need to know that your recovery server is ready to assume the production role quickly. Yet, a number of problems might arise that can slow an IBM PowerHA for i switch operation or even halt it.

MIMIX for PowerHA continuously checks your environment for these issues. If a problem should occur, it allows you to correct the problem before it turns into a major downtime event. The audits that MIMIX for PowerHA performs include the following:

- **Duplicate library name detection.** When you vary on an IASP, the system performs a check to determine if a duplicate library name exists in SYSBAS. If so, the vary on will fail. If this occurs during a switch from a production to a recovery server, the switch will fail. To avoid these problems and protect the viability of switches, MIMIX for PowerHA monitors library names and notifies you whenever a library name in an IASP on the production server is the same as a library name in SYSBAS on the recovery server. Consequently, you can take action immediately, rather than having to correct the problem only after you discover it during a switch.
- **User profile configuration audit.** If user profiles differ between the production and recovery servers, the OS must update the profiles before the IASP can be used. This could significantly increase recovery times. MIMIX for PowerHA audits for out-of-sync user profiles in real time so you can update them to maintain an optimum environment long before you need to switch. Recovery time will be minimized when you do switch.
- **Library ratio audit.** Maintaining an optimum ratio between libraries in SYSBAS and libraries in the IASP reduces vary on times, and, thus, recovery times. MIMIX for PowerHA monitors this ratio so you can keep your environment optimized.

## Added High Availability and Disaster Recovery Protection

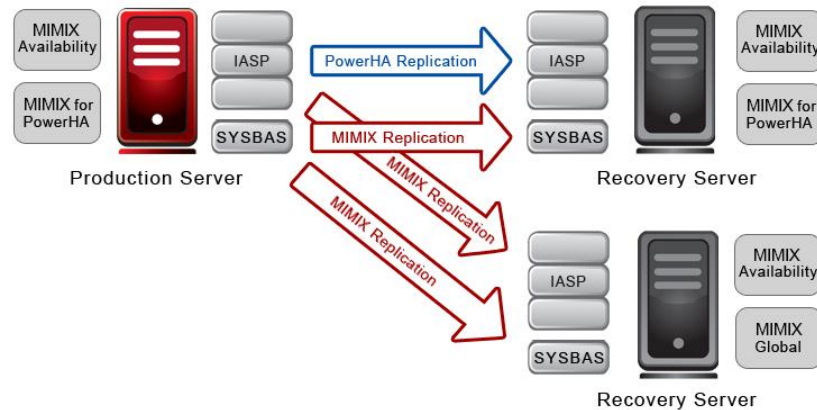
In addition to providing replication of SYSBAS data, MIMIX for PowerHA can be used in conjunction with MIMIX Availability to enhance the protection that IBM PowerHA provides by replicating data and applications to additional local and remote servers. This allows you to maintain a local recovery server so you can resume operations quickly after a highly localized problem, such as a hardware failure. You might, at the same time, also replicate to a remote recovery server to ensure that your organization's business operations will incur minimal downtime even if a natural disaster knocks your entire primary data center offline.

Duplicate recovery servers are essential for complete availability protection. Consider your exposure during planned maintenance. One of the benefits of an HA environment is that you can switch production to a recovery server while you perform maintenance on the production server. However, if you have only the one recovery server, your data and applications will be vulnerable during maintenance operations because replication will cease and, therefore, your data and applications will not be backed up during that time.

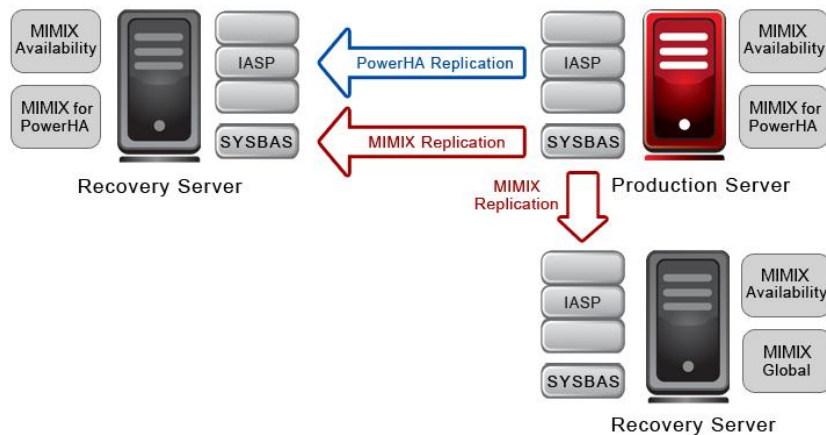
Simultaneous disasters are exceptionally rare, but not non-existent. In the case of simultaneous disasters that bring down both the production server and one of the recovery servers, another recovery server will keep your data and applications available. This is especially important if the first recovery server is in the same city or, particularly, in the same

building as the production server. In that case, even comparatively minor events might knock both servers offline for an unacceptable period.

Using MIMIX for PowerHA with MIMIX Availability to maintain multiple recovery servers avoids this vulnerability. When you switch the production role to the first recovery server, MIMIX for PowerHA ensures that MIMIX Availability will continue to replicate from the new production server to the additional recovery server(s), keeping your data and applications fully protected even while your production server is unavailable.



Following a switch of production to the recovery server, two copies are still maintained. If this were a failover of the production server instead, there is still a copy of the data protecting you from further outages.



## Greater Access to Data in Real Time

MIMIX for PowerHA protects against downtime and data loss, but real-time replication to multiple servers can help you realize even more value from your data by making data available where it is needed and when it is needed for other business purposes. MIMIX for PowerHA combined with MIMIX Availability does that.

One or more of the additional real-time replica servers that the combined MIMIX solutions maintain can be used for purposes other than HA or DR. Because a replica contains a continuously up-to-date copy of your production data, you could, for example, use the server as a:

- **Data warehouse or query server.** By taking the query load off the production server, you can allow users to mine their data more fully, while also improving transaction-processing performance on the production server.
- **Reporting server.** Likewise, by running reports on a replicate server you can remove that load from your production server, while opening up new possibilities for more complex reporting.
- **Tape-based backup source.** Even after implementing an HA/DR solution, many organizations still create offline backups, typically on tape, for archival purposes and as the last line of defense against data loss. However, when run on a production server, backup jobs require that the application be stopped until the backup operation has completed.

Because replicas maintained by MIMIX are always current, you can create backups on a replica server by stopping the application of replicated changes to the database at a specific point in time while the tape operation is conducted, confident that the tape backup will be as up-to-date as it would have been if you created it on the production server. However, the production server will experience none of the downtime that it would otherwise have incurred during backup operations because changes are being buffered on the recovery server and will be applied to the database after completion of the tape operation.

## Trusted Technology Built on Partnership

Vision Solutions has served the needs of the IBM midrange availability market for more than 25 years. At the onset of our involvement in this industry, we pioneered infrastructure technologies that optimize performance and ensure efficient, reliable replication.

For the past two and a half decades, we have leveraged our close relationship with IBM to optimize our high availability and disaster recovery solutions for IBM i. In addition, we've used—and continue to use—the insights we have gathered from working with thousands of customers worldwide to continually enhance our offerings.

MIMIX Availability—our flagship solution for IBM i—is an enterprise-class solution that provides faster recovery, stronger data protection and a smarter way to manage high availability for IBM i.

And Vision Solutions is more than just technology. We are the global availability experts. We have a worldwide network of certified partners and a team of veteran industry professionals with unmatched expertise to help you maximize the return on your availability investment. In addition, our CustomerCare team delivers around-the-clock, around-the-globe support with the most advanced and accessible tools.