High Availability and Disaster Recovery Planning: Next-Generation Solutions for Multiserver IBM Power Systems Environments

Redguides

for Business Leaders

Dino Quintero
Steven Finnes
Mike Herrera
Ravi Shankar
Jonathan Sigel

- Understand business continuity challenges, with a specific look at outages
- Learn how to build an effective high availability and disaster recovery solution
- Explore end-to-end solutions with IBM PowerHA SystemMirror

Redbooks

Overview

In a volatile and technology-dependent business climate, downtime, whether it is planned or unplanned, is costly. According to a recent IDC report,[1] the cost of IT downtime to organizations can range from thousands to millions of dollars per hour. A 2009 study published by the Information Availability Institute effectively makes the point that measurable objectives should be made based on "the needs of the business, not the capabilities of current technologies and procedures."[2] High availability and data protection are not new concepts to IT professionals. What is new are the expanding options and capabilities around deployment of highly available environments that offer varying levels of data, application and infrastructure resilience.

Many tiers of a high availability disaster recovery (HADR) solution are possible. The right HADR configuration is a balance between recovery time requirements and cost. Critical decisions must be made to determine which services within the business must remain online in order to continue operations.

IBM® Power Systems customers who use AIX® (UNIX®) and IBM i operating systems already benefit from the industry-leading resiliency capabilities (reliability, availability, and serviceability (RAS)) of the platform. These customers can further benefit from the continuous availability offered by the IBM PowerHA™ SystemMirror solution. Power Systems servers have become the platform of choice for clients who are running their business-critical workloads. Vital transactions, such as bank deposits, medical claims, logistics of critical goods shipments and benefit payments, are processed every day in IBM AIX and IBM i environments all over the world.

In this IBM Redguide™ publication, we define the fundamental concepts for high availability, present considerations for building a strategy, and make suggestions for choosing an effective high availability and disaster recovery strategy depending on your environment. We offer details about how IBM PowerHA SystemMirror tightly integrates with the Power Systems servers, IBM operating systems, and storage to help meet business continuity requirements. We also discuss the challenges posed by a myriad of environmental risks, changing business environments, and compliance issues.

---

1 Source: IDC report "Leveraging Clustered File Systems to Achieve Superior Application Availability," document 219198, published July 2009; available at (IDC login required):
`http://www.idc.com/getdoc.jsp?sessionId=&containerId=219198&sessionId=530BF8F472193221DE4563C85617FD4F`

2 Source: Information Availability Institute study "The State of Resilience on IBM® Power Systems™," published in 2009.

This document is a reference guide, with a definitive collection of approaches, that outlines high availability offerings from IBM and includes characteristics and respective requirements. The intended audience for this paper is both customers and consultants who are looking for an overview of high availability solution approaches specifically for Power Systems environments.

This guide includes the following topics:

What is high availability?

High availability is a generic term used by the IT industry to describe the accessibility and uptime of critical IT application environments. From a technology perspective, high availability solutions always involve redundancy. For example, RAID 5 enables highly available disk arrays through a form of data redundancy.

No one can expect to achieve the popularly coined “five nines of availability” (shown in Table 1) by using technology alone. A highly available environment is a combination of technology, change control, skills and overall IT operational discipline.

Table 1 The five nines of availability

Uptime	Uptime	Maximum downtime per year
Five nines	99.999%	5 minutes 35 seconds
Four nines	99.99%	52 minutes 33 seconds
Three nines	99.9%	8 hours 46 minutes
Two nines	99.0%	87 hours 36 minutes
One nine	90.0%	36 days 12 hours

The need for high availability and disaster recovery solutions

High availability is a key component of business resiliency. It is widely documented that outages increase the total cost of IT ownership and cause potential damage to client relationships and loss of revenue. Although hardware has become highly reliable, research shows that unplanned outages occur and typically result from operator error, software bugs, environmental conditions and other non-hardware related factors—problems that reliable hardware cannot prevent. Planned outages for application and system maintenance also impact business performance. Businesses are aggressively shrinking the time allotted for these types of activity.

Disaster recovery solutions are an extension of high availability solutions with the added capability of providing resiliency with geographic dispersion. More IT shops are moving away from outsourced disaster recovery operations to insourced disaster recovery operations based on continuous data replication between geographically dispersed locations. Modern disaster recovery solutions require both geographic dispersion and recovery point objectives as close to zero as possible. Increasingly, IT shops are being asked to prove periodically that they can recover operations at a remote facility. The simple fact is that owning your own disaster recovery solution is both economically sensible and gives you greater control over your environment.

Business Continuity Resiliency Services: IBM can also provide remote facilities worldwide through Business Continuity Resiliency Services (BCRS). You can learn more about BCRS at the following Web address:

http://www.ibm.com/services/continuity

Establishing common requirements

The following requirements are the most common IT considerations for establishing an HADR solution:

Recovery time objective (RTO)
The time as measured from the time of application unavailability to the time of recovery (resuming business operations).

Recovery point objective (RPO)
The last data point to which production is recovered upon a failure. Ideally, customers want the RPO to be zero lost data. Practically speaking, we tend to accept a recovery point associated with a particular application state.

Planned downtime In normal day-to-day operations, the largest share of time that an application is rendered unavailable because of planned maintenance procedures, such as system saves or operating system upgrades and the like. In an ideal environment, a redundant resource is used to carry the production workload so that the primary environment can undergo maintenance. The faster that you can switch between the primary and secondary nodes in a cluster, the less impact there is to the production environment.

Geographic dispersion
In the context of a multisystem HADR solution topology, the capability to recover operations at a remote location. This requirement is increasingly driven by compliance regulations, dispersion of the data, and the growing importance to have a complete disaster recovery solution.

Ease of management
The degree of automation that an HADR solution offers to an IT operations staff. Consider both the degree of skill specialization required to manage the solution and its practical capability when applied to various resiliency operations such as planned failover or role swap operations.

Ease of deployment Clients ultimately want an HADR solution that is simple to configure. Through the use of node discovery functions, Smart Assists and configuration wizards for AIX deployments, and independent auxiliary storage pools (IASPs) for IBM i environments, IBM clustering solutions can reduce the amount of time required to deploy an environment.

Integration and support

When up and running in production, the degree of integration with the operating system influences the robustness of the solution, the types of skills required to manage the solution, and the types of support that might be involved in the event of a problem.

Outage types to consider

This guide focuses specifically on high availability solutions that involve multiple independent systems referred to as *nodes*, which are incorporated into a cluster. To achieve high availability, the solution must be designed with consideration of all aspects of the environment with the infrastructure being a critical building block.

A high availability solution should address the following types of outages:

- Scheduled outages
- Outages in the hardware subsystems (central processor complex (CPC), I/O, disk)
- Outages in the application, operating system, or both
- Operator error outages

For any high availability solution approach, evaluate it in the context of the previously mentioned outage types. Considerations include whether the solution covers each of these outages or some of these outages, and whether the coverage is adequate for the stated IT resiliency objectives and requirements. Customers must identify the outage types that require coverage and evaluate the solution options to determine which option best fits their requirements.

Generally speaking, disaster recovery solutions provide protection against natural disasters (earthquake, floods, or fire, which can lead to extended site power outages). Increasingly, corporate and governmental regulations are driving focus in this area. As you evaluate your solution options, consider whether a given solution can be effectively deployed in a geographically dispersed topology to meet your corporate compliance objectives.

Customer infrastructure can fail because of various outage causes such as those that we previously mentioned. Hardware failures represent only a small percentage of the total failures. Nearly 50% of outages are the result of software and operator errors. These are all instances in which a high availability solution can help shorten outage windows and can provide a reliable mechanism to move critical resources between highly available servers. In a multisite disaster recovery implementation, a clustered solution can further extend its reach and manage the replication of the data between the sites. Such a solution can automatically reverse the roles in the event of a local site failure.

Table 2 outlines the groups of outage types to consider when evaluating a high availability solution.

Table 2 Outages to consider

Group	Possible outage types that an HADR solution might cover
Group 1	CPC (hardware: CPU and memory)
Group 2	Network or storage adapter failures, cable disconnects and so on External errors: storage errors, switch errors and so on
Group 3	Critical operating system resource: volume, file system, IP and so on
Group 4	Application, middleware, and operator actions

Group	Possible outage types that an HADR solution might cover
Group 5	Site outages

The IBM PowerHA SystemMirror Enterprise Edition solution reliably orchestrates the acquisition and release of cluster resources locally or from one site to another in the event of an outage or natural disaster. The solution incorporates end-to-end components as described in the following sections.

End-to-end solution components

A comprehensive HADR solution has the following basic components:

- Application data resiliency

  Applications require access to data or copies of the data to perform business-critical operations. Therefore, data resiliency is the base or foundational element for a high availability and disaster recovery solution deployment.

- Application infrastructure resiliency

  Infrastructure resiliency provides the overall environment that is required to resume full production at a standby node. This environment includes the entire list of resources that the application requires upon failover for the operations to resume automatically.

- Application state resiliency

  Application state resiliency is characterized by the application recovery point as described when the production environment resumes on a secondary node in the cluster. Ideally the application resumes on an alternate node at the last state where the application was on the primary system when a failure occurred. Practically speaking, the characteristic of the application to resume varies by application design and customer requirements.

A complete end-to-end solution incorporates all three elements into one integrated environment that addresses one or all of the outage types as described previously in this paper. The behavior of a solution to a customer depends upon the inclusion and incorporation of these basic elements into the clustering configuration. For example, you can have a solution based purely upon data resiliency and leave the application resiliency aspects of the final recovery process to IT operational procedures. Alternatively you can incorporate the data resiliency into the overall clustering topology enabling automated recovery processing.

Application data resiliency: Methods and characteristics

Basic technologies are employed by HADR solution implementations to provide application data resiliency. Each one has its particular characteristics and applications. Generally speaking, there are two distinct groups: storage-based resiliency and log-based replication.

Storage-based resiliency

Data resiliency across multiple nodes in a cluster is the foundation for building an effective high availability solution. Storage replication is the most commonly used technique for deploying cluster-wide data resiliency. There are two general categories for storage-based resiliency: shared-disk topology and shared-everything topology.

In this context, the following critical storage-related high availability criteria can be considered:

- Active storage sharing across the cluster (concurrent access)
- Shared-disk configuration (active-passive)
- Multisite replicated storage

> **Local clustering versus multisite replication:** Active storage sharing across the cluster and shared disk configuration are both specific to local clustering. Storage replication can provide multisite replication for the environment.

Active storage sharing across the cluster (concurrent access)

Often referred to as *shared-everything storage*, active storage sharing across the cluster is an active-active ownership arrangement. In this arrangement, nodes in the cluster have simultaneous read/write access to the shared data. The cluster management technology performs locking operations to ensure that only one node can perform an update or write operation at a time.

The benefit of this approach is that no switching operation is associated with storage resiliency because the nodes simultaneously own the shared resource. If a node outage occurs, another node in the cluster resumes production through a reassignment process. The technology used to resume the application can be based on either journals or memory replication. Another consideration is the degree to which the entire application infrastructure is monitored and recovered. However, this type of sharing requires the individual software subsystem to be aware of concurrent disk access to avoid data corruption.

Shared-disk configuration (active-passive)

Shared-disk configuration is an active-passive shared ownership arrangement between nodes in the cluster. One node in the cluster performs read/write operations to the disks. Ownership of those resources can be passed to other nodes in the cluster as part of a failover (or rollover) operation. The operating system, application and data are all switched between nodes. The recovery point is established by applying the journals that are in the shared disk resources. The recovery time is associated with the time it takes to apply the journal.

Live Partition Mobility and shared-disk technology

Virtualizing physical resources is becoming prevalent on Power Systems servers. Virtualized storage area network (SAN)-attached volumes can be combined with POWER6® and POWER7™ hardware for clients to use extended features. For example, with Live Partition Mobility, logical partitions (LPARs) can be dynamically moved between servers with minimal impact to the application. Live Partition Mobility functions, in combination with PowerHA SystemMirror cluster solutions, can compliment the environment by providing non-disruptive planned maintenance while protecting against unexpected outages.

HyperSwap configurations

Shared-disk configurations can be extended by using a HyperSwap® configuration where the disks in a shared storage pool are mirrored between separate storage servers enabling resiliency between two storage servers. Continuous availability of the solution is assured even if one of the storage servers fails.

With this type of configuration, high availability is realized against server-based errors, and continuous availability is realized against storage errors. The nodes in the cluster have access to the mirrored data across two separate storage servers, giving protection against both a primary server outage and a storage server outage.

Multisite replicated storage

The shared-disk topology can further be extended for geographic dispersion by using either host-based or storage server-based replication technology. Data in a storage pool is replicated in a synchronous manner for zero loss implementations and in an asynchronous manner for geographically dispersed sites where latency might impact operations.

Host-based replication

Host-based replication implies that the host mirroring technology is doing the work. Hosts on the two sides coordinate and replicate the data across the cluster. The major advantage of host-based replication is that it can work for any storage, irrespective of whether it supports mirroring.

Storage server-based replication

Storage server-based replication means that the storage server is performing the data replication on behalf of the primary node. Perhaps more importantly, this type of replication can provide continuous data replication in the event of a production node outage and provides a common replication mechanism across various platforms.

Synchronous replication

Synchronous replication means that the application state is directly tied to the completion of the write operation to both the local node and remote node. Synchronous replication provides a mechanism to ensure that no data (which is written to disk) is lost in relationship to the application in the event of an unplanned outage. It also means that the distance between the primary and secondary nodes have a direct impact on the application response time.

Asynchronous replication

Asynchronous replication means that the operations are synchronous and the storage subsystem manages the synchronization between the primary and the secondary nodes. This technology enables the application to continue operations without waiting for the remote storage operation to complete minimizing a performance impact. Data is typically replicated in groups of volumes, and consistency is established within group boundaries. In the event of an unplanned outage, the recovery point on the remote server will be based on the last consistency boundary that is replicated to the remote storage subsystem.

Log-based replication

Log-based replication is a form of resiliency primarily associated with databases. Typically, database logs are used to monitor changes that are then replicated to a second system where those changes are applied. IBM i solutions that are based on this technology are referred to as *logical replication*.

Application infrastructure resiliency

Application infrastructure resiliency has two aspects. First, it provides the application with all the resources that it requires to resume operations at an alternate node in the cluster. Second, it provides for cluster integrity by using monitoring and verification.

For an automated or semi-automated failover operation to work, all of the resources that the application requires to function on the primary node must also be present on the secondary node. These resources include items such as dependent hardware, middleware, IP connectivity, configuration files, attached devices (printers), security profiles, application specific custom resources (crypto card) and the application data itself. These redundant resources are typically referred to as *cluster resources* and are managed by forming cluster resource groups. Dependencies established between the application and these resource groups form the key control mechanism, which ensures that the resources are prepared and available before the application resumes operations on the standby node.

During day-to-day operations, the cluster monitors the resilient infrastructure resources for changes that indicate a failure, a pending failure or a possible configuration change that might cause a cluster operation (such as a failover) or an operator to take corrective action. Monitoring is primarily about the resources that must be tracked by the high availability solution and the internal solution notifications that can trigger an action as defined by the policy for high availability management. An aspect of monitoring is performing periodic verification checks where specific or custom scans of the cluster resources are conducted to assess status against the intended configuration.

These operations are performed in addition to real-time monitoring as an integrity check that supersedes the real-time monitoring function. A modern high availability solution automatically identifies changes and addresses them through auto-corrective features or notifications methods. In addition to the monitoring and verification capabilities of a modern resilient infrastructure solution, cluster-wide management functions should be available that enable the operator to perform various operations on behalf of the application and operating system to maintain or update the resilient infrastructure.

For example, if there is a need to expand the storage capabilities of the cluster, then an operator should be able to add disks and include them as a cluster-wide resource by using the high availability-solution-provided central interface. The modern resilient clustering infrastructure should be able to monitor and manage the critical resources from a central point of control versus performing these cluster-wide operations on each node individually.

In regard to effective health monitoring and cluster-wide system management, the high availability solution must be closely integrated with the host operating system. This integration provides for synergetic health management, application management, and system state management.

Cluster-aware operating systems

Health management, if done externally to the operating system, can be prone to errors and scheduling issues. It can also cause ongoing management challenges. The health management process must be done from within the operating system or operating system kernel to make it highly reliable and less reliant on user monitoring and intervention.

A cluster-aware operating system naturally leads to the exploitation of the hardware resources within the infrastructure (network adapters, SAN connections and so on) with minimal inputs from the user while providing for multiple redundant communication links between the nodes in the cluster. This discovery-based configuration capability reduces the monitoring and configuration burden on the user who is responsible for maintaining the health management infrastructure.

Enabling cluster awareness in the operating system enables operating-system-based operations to be in harmony with the cluster-wide high availability solution. In particular, it ensures that operating-system-based operations do not accidentally disturb the cluster. A high availability solution implemented with a cluster-aware operating system, exploits the operating system features, extending them across the entire cluster, and enables centralized cluster management of the infrastructure.

The IBM PowerHA SystemMirror strategy is rooted in deep integration with the operating system. IBM’s intention and strategy in an upcoming release of the PowerHA solution is to provide a deeper integrated cluster aware operating-system-based high availability solution. For the next generation of the PowerHA SystemMirror solution on AIX, IBM aims to exploit AIX cluster-aware operating system capabilities, thus providing a more robust high availability solution. The IBM i operating system is cluster aware and is exploited by PowerHA for IBM i 6.1 and PowerHA SystemMirror for i 7.1.

Application-state resiliency

Assume for a moment that you have deployed a clustering infrastructure. You have implemented your data resiliency based on a shared configuration, and your application infrastructure resiliency is in place so that all of your application resources are available on a secondary node. You can failover (or roll over) to the alternate node in your cluster at will.

The question you need to ask yourself is: Where will the application recovery point be with respect to the last application transaction? If your application is designed with commit boundaries and the outage is an unplanned failover, then the recovery point in the application will be to that last commit boundary. If you are conducting a planned outage role swap, then the application is quiesced so that memory can be flushed to the shared-disk resource and the data and application are subsequently varied on to the secondary node.

Many factors related to the environment will drive design decisions that will help achieve a set of resiliency objectives. While RPO, RTO and network recovery objective (NRO) play a critical role in establishing objectives, middleware and the application recovery characteristics will also play an important role. For example, the database might take much longer time to do the recovery processing for 1 TB of data as compared to 100 GB of data.

While application state resiliency depends on many characteristics of the environment, a high availability solution should aid in health monitoring of the application stack. This solution should also provide for corrective actions in the cluster to reduce failover times and to accelerate recovery times. For example, some middleware in a high availability configuration might create a cache of application state information on another node in the cluster apart from the active node, thus enabling a quicker failover.

Application friendly cluster infrastructure

The IBM PowerHA SystemMirror for AIX solution provides high availability management agents called *Smart Assists* for key middleware management and application deployment. Smart Assists help customers define high availability policies that can be rapidly integrated for critical workloads. They help in discovering the complex software that is deployed in the cluster. This discovery-based information is presented to the customer and aids in defining the high availability policy for the site. After the high availability policy is set, the Smart Assists provide health monitoring methods and periodically checks the health of the middleware. Upon failure detection, middleware and its resource dependencies are restarted on the node specified by the policy. Smart Assists thus provide for end-to-end high availability management of the middleware and application stack.

PowerHA SystemMirror for IBM i clustering technology is an extension of the IBM i operating system. The IBM i operating system is cluster aware, and PowerHA SystemMirror exploits the integrated clustering technology, thus enabling a complete end-to-end solution for HADR. The application friendly cluster infrastructure is primarily centered on exploiting the IASP data structure. A commercial application that has been set up for IASPs is readily deployed into a PowerHA SystemMirror HADR cluster. Deploying an IASP data structure is generally straightforward. Many IBM i customers with home-grown applications have implemented PowerHA for IBM i 6.1. Several of the major IBM commercial application providers have already implemented PowerHA support, and many more are coming on line.

Solution options for Power Systems clients

In combination with leading RAS features of IBM Power Systems servers, IBM PowerHA clustering software effectively helps detect any component failure and reacts accordingly, shortening the amount of time to recovery or masking a failure altogether.

The IBM PowerHA SystemMirror Enterprise Edition solution for Power Systems can provide a valuable proposition for reliably orchestrating the acquisition and release of cluster resources locally or from one site to another in the event of an outage or natural disaster.

Now that we have reviewed the dimensions involved in the establishment of an optimal solution for your HADR operations, you can see the options that are available to Power Systems clients in Table 3 and Table 4 on page 11. The event groups are defined in Table 2 on page 4.

Table 3 High availability solutions offerings for Power Systems customers

Solution	Group coverage	Data resilience	Replication	Infrastructure resilience	Integration level
PowerHA SystemMirror for AIX	Groups 1, 2, 3, 4 and 5	Shared-disk, replication, local HyperSwap	► Geographical Logical Volume Manager (GLVM) ► Metro Mirror ► Global Mirror ► EMC Symmetrix Remote Data Facility (SRDF)	Cluster-wide monitoring for groups 1, 2, 3, 4 and 5	AIX cluster aware
VMControl, Live Partition Mobility or both	Groups 1 and 2	Virtualized storage	NA	Central processor complex	Hypervisor
IBM DB2® pureScale	Groups 1, 2, 3, 4 and database layer	Shared everything storage	NA	Cluster-wide monitoring (database layer) for groups 1, 2, 3 and 4	DB2
DB2 HADR	Groups 1, 2 and 5 (DB2 data)	DB2 log-based replication	DB2 log-based replication	NA	DB2
PowerHA SystemMirror for i	Groups 1, 2, 3, 4 and 5	Shared-disk, replication	► Geomirror ► Metro Mirror ► Global Mirror ► FlashCopy®	Cluster-wide monitoring for groups 1, 2, 3, 4 and 5	IBM i cluster aware (high availability shared resource (HASR))
IBM iCluster®	Groups 1, 2, 3, 4 and 5	Log-based replication	Log-based replication	Log monitoring	Remote journaling

Table 4 Power Systems HADR solution coverage matrix

Solution	Topologies supported	Resilient data	Coverage	Redundant resource utilization	Additional capabilities
PowerHA SystemMirror for AIX	Multinode	Multiple copies	► Data center ► Multisite	Dynamic LPAR	► Cluster-wide administration ► Cluster-wide file synchronization
VMControl, Live Partition Mobility, or both	Two-node	Single copy	Data center	NA	NA
DB2 pureScale	Multinode	Single copy	Data center	Active-active	NA
DB2 HADR	Two-node	Two copies	Data center	NA	NA
PowerHA SystemMirror for i	Multinode	Multiple copies	► Data center ► Multisite	► CBU ► Flash Copy	Integrated HADR and IBM Storage Server Integration
iCluster	Multinode	Multiple copies	Multisite	Read access	Active-active configurations

Summary

You must consider several factors prior to choosing a solution for high availability or disaster recovery. You need to establish comprehensive requirements that meet your business and detailed plans for change management and back up operations. Consider the type of outages you want to cover, the degrees of automation and the scope of coverage. Evaluate each solution approach. For example, is it an integrated extension of the host operating system technology, or is it an add-on solution? How long does it take, and how difficult is a planned or unplanned operation? A resilient IT infrastructure with 7x365 application availability is becoming the norm. An optimal resiliency solution for your environment is one that addresses all facets of your requirements including the day-to-day IT operations aspect.

Finally, the strategy for IBM Power Systems is to deliver more advanced functional capabilities for business resiliency. The strategy is also to enhance product usability and robustness through deep integration with AIX, IBM i and affiliated software stack and storage technologies. PowerHA SystemMirror is architected, developed, integrated, tested and supported by IBM from top to bottom.

Additional resources

For more information, consult the following resources:

- IBM Redbooks® Web site:

  http://www.ibm.com/redbooks

  – *IBM i 6.1 Independent ASPs: A Guide to Quick Implementation of Independent ASPs*, SG24-7811

  – *Implementing PowerHA for IBM i*, SG24-7405

  – *Implementing SAP Applications on the IBM System i Platform with IBM i5/OS*, SG24-7166

  – *PowerHA for AIX Cookbook*, SG24-7739

- IBM PowerHA SystemMirror for AIX

  http://www.ibm.com/systems/power/software/availability/aix/index.html

- IBM PowerHA High Availability wiki

  http://www.ibm.com/developerworks/wikis/display/WikiPtype/High%20Availability

- Implementation Services for Power Systems for PowerHA for AIX

  http://www-935.ibm.com/services/us/index.wss/offering/its/a1000032

- IBM Systems Lab Services and Training

  http://www-03.ibm.com/systems/services/labservices/

- IBM System Storage™ Virtualization

  http://www-03.ibm.com/systems/storage/virtualization/

- Ready for IBM Power Systems Software™

  http://www-304.ibm.com/isv/tech/validation/power/index.html

For education and workshops, register for the following learning opportunities at:

http://ibm.com/services/learning

- “PowerHA for AIX Implementation, Configuration and Administration,” AN410
- “IBM Power HA for i, Clustering, and Independent Disk Pools Implementation,” AS541 or OV541

Finally, you can build a roadmap to continuous availability with the IBM Availability Factory. To learn more, contact your IBM representative or an IBM Business Partner.

The team who wrote this guide

This guide was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO).

Dino Quintero is a Project Leader and an IT generalist with the ITSO in Poughkeepsie, NY. His areas of expertise include enterprise continuous availability planning and implementation, enterprise systems management, virtualization, and clustering solutions. He is currently an Open Group Master Certified IT Specialist - Server Systems. Dino holds a Master of Computing Information Systems degree and a Bachelor of Science degree in Computer Science from Marist College.

Steven Finnes is Offering Manager for IBM Power Systems High Availability solutions. He is responsible for PowerHA SystemMirror for AIX and IBM i and for Power Systems Capacity Backup (CBU) offerings. Steve has worked for IBM for 29 years. As the chief engineering manager for I/O and storage, he and his team brought the first RAID 5 product offering to market in the early 1990s. He has been involved in both research and development and marketing of high availability offerings for over 10 years. Steve has authored numerous articles for IT publications on business resiliency and disaster recovery strategies.

Mike Herrera is a Certified IT Specialist in the Advanced Technical Skills team out of Dallas, TX. He has worked for IBM for 11 years and was formerly a support escalation contact in the IBM AIX Support Line. He specializes in AIX and PowerHA and SAN environments and is certified by IBM as an Advanced Technical Expert. He has co-authored two IBM Redbooks publications. Mike holds a bachelor degree in Management Information Systems from the

University of Connecticut and an Master of Business Administration (MBA) from the University of Dallas.

Ravi Shankar is an Architect in Operating System and High Availability Areas. He has worked in the IT industry for nearly 20 years and has specialized in a wide set of technologies in the areas of embedded systems, security, and business resiliency. Ravi has in-depth knowledge about real-time programming, operating system internals, high availability and disaster recovery.

Jonathan Sigel is the Market Segment Manager in North America for IBM's high availability solutions for Power Systems. He has been with IBM for six years in roles spanning offering management, market analysis and business development. Jonathan holds an MBA in Marketing and Finance from City University of New York, and a Bachelor of Arts degree from Harvard University.

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/pages/IBM-Redbooks/178023492563?ref=ts

- Follow us on twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

This document, REDP-4669-00, was created or updated on May 5, 2010.

Trademarks