



In association
with
Common France

COMMON EUROPE CONGRESS 2026

14 - 17 June
Lyon, France

The largest conference in Europe
for solutions around IBM Power (IBM I, AIX, Linux) & IBM Storage

common
EUROPE

www.comeur.org

common
FRANCE


LYON | CENTRE DE CONGRÈS
EVENTS DE LYON



**Welcome to Lyon, France
and the 2026 Common Europe Congress**

**Bienvenue à Lyon, en France,
et au Congrès de Common Europe 2026**

Building Resilient Architecture with PBHA and PBR on FlashSystem



Guillaume Legmar

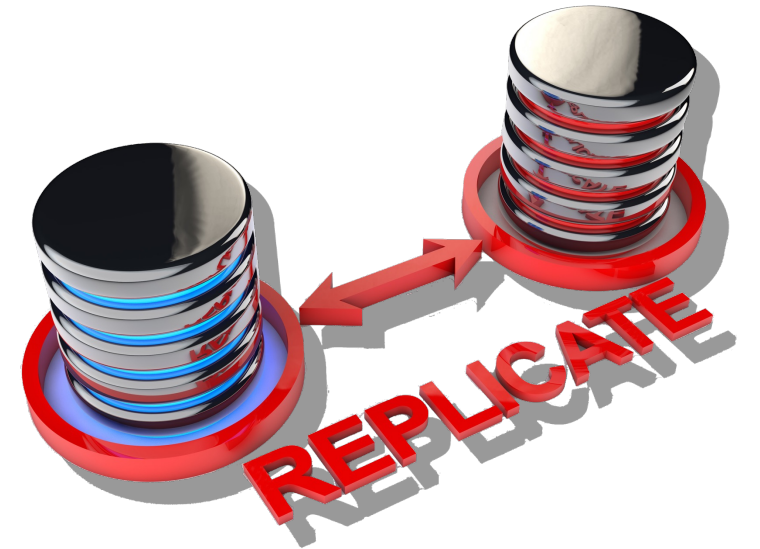
IBM Storage Flashsystem Specialist

Guillaume.legmar@fr.ibm.com

IBM Policy-Based Replication (PBR)

Simplicity, performance and efficiency

Policy-Based Replication



Policy-Based Replication (PBR)

Async DR

Introduced previously, Policy-Based Replication provides a simple to use, high performance DR solution

HA

New Policy-Based High Availability provides high throughput, low latency HA with no single point of failure

3-Sites

Policy-Based HA + DR for enterprise **three-site replication**



Clients are really impressed with PBR and PBHA

HA replication is so light and so easy to manage.

- Storage Administrator from Italy

And the truth is it is wonderful. It works great, and the replication task is much easier.

- Storage Administrator from Spain

We fully converted over to PBR yesterday. It worked exactly as advertised and is FAST!

- Major US Supermarket chain

We can testify to the improvement in performance and stability provided.

- Large Bank in France

>4X

Higher performance

32K

Replicated volumes

4PiB

Replicated capacity

The importance in using policies



Consistency

Policies define how replication must be configured, and how the system implements it, to ensure that every volume is configured correctly, using “hard-to-do-harm” principles.

Simplicity

All remote provisioning is managed by the system reducing the burden on storage admins, removes opportunities to make mistakes and all day-to-day mgmt. is from a single pane of glass.

Automated

Replication that auto-manages itself to adapt to the current conditions without causing performance problems, with reporting and alerting if Recovery Point Objectives (RPOs) are exceeded.

Configuring replication

1

Partner systems together and confirm the certificates are correct.

2

Follow the step-by-step walkthrough to define the policies for remote provisioning.

3

Create a replication policy, defining the RPO and the systems involved.

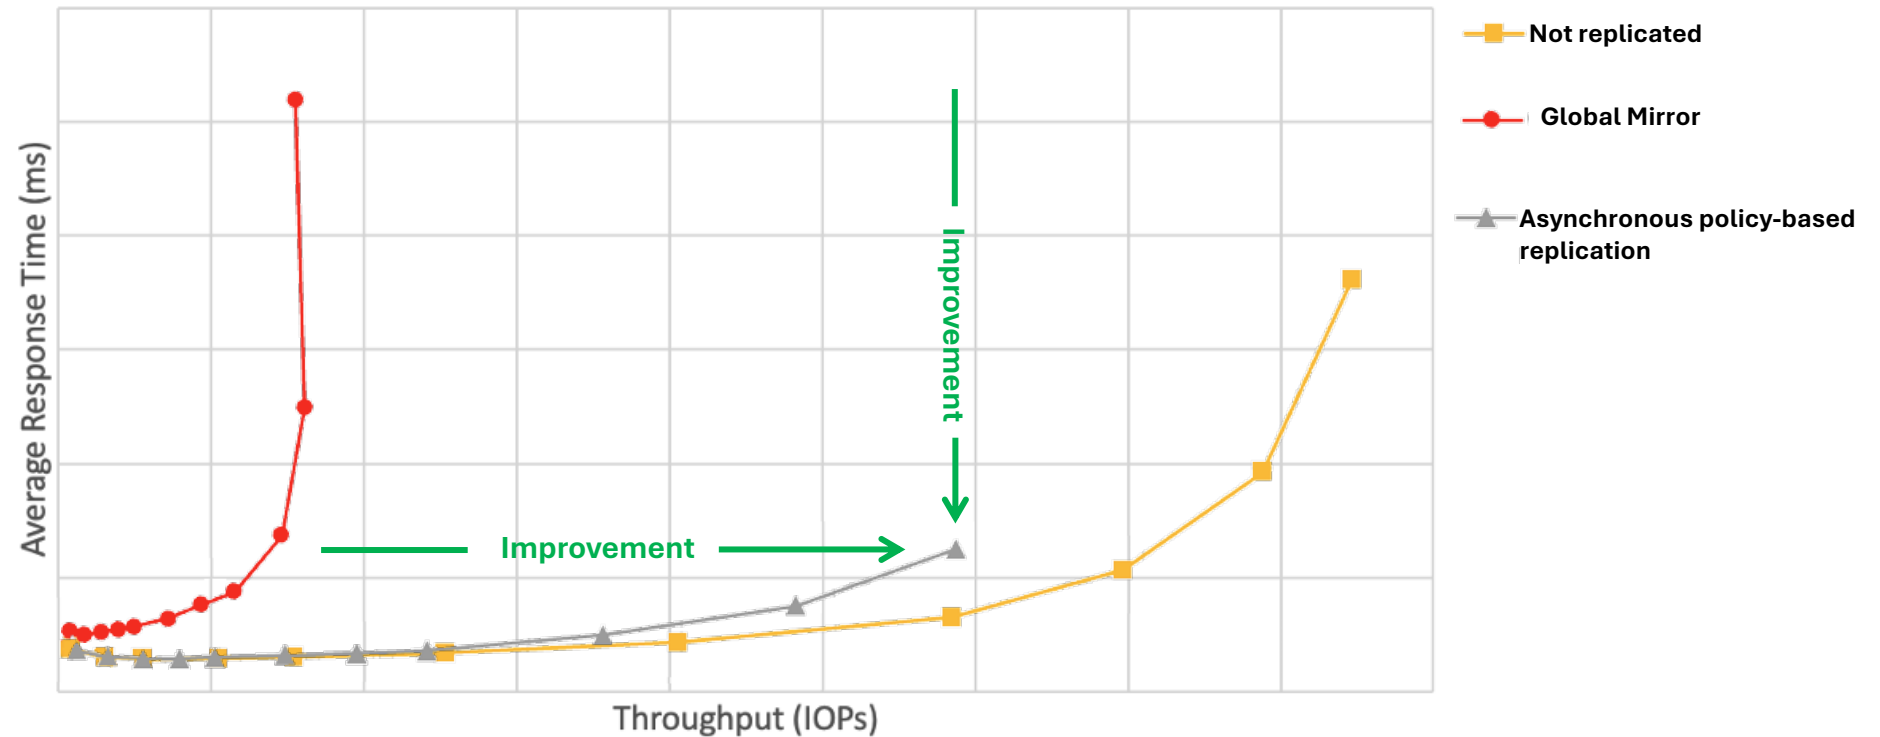
4

Assign replication policy to a volume group and the system will automatically configure and start replication.

Performance gains from ground-up redesign

The entire replication path was re-written from scratch, with emphasis on throughput, response times and scalability.

Resulting in storage systems that gained massive increases in throughput, with significantly lower response times.

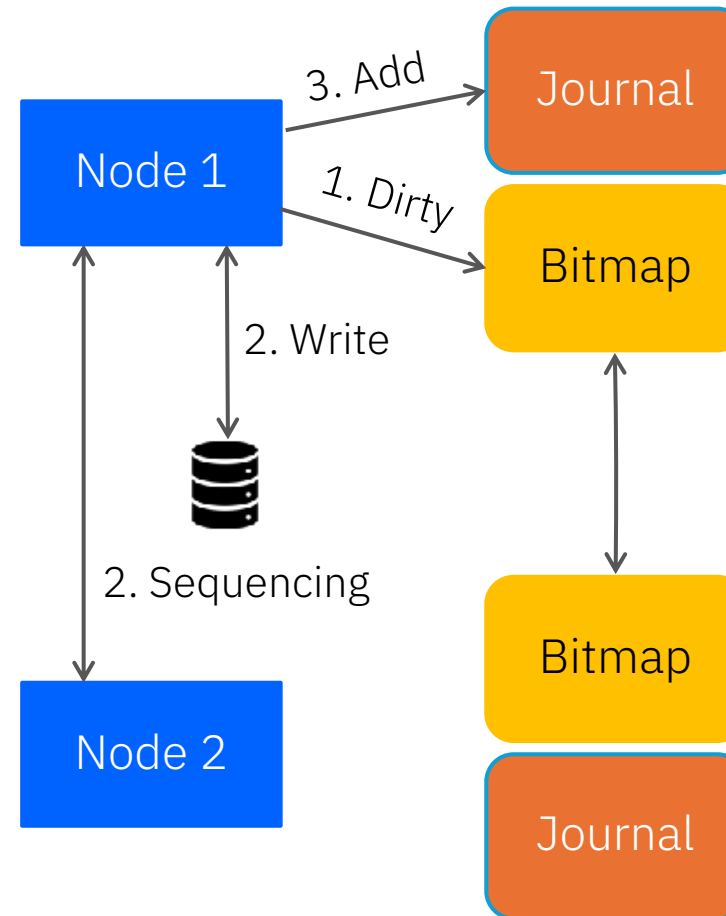


2x FS9500 connected via FC (16K/70/30/50)

Two methods to provide Asynchronous Replication

#1 Journaling Mode

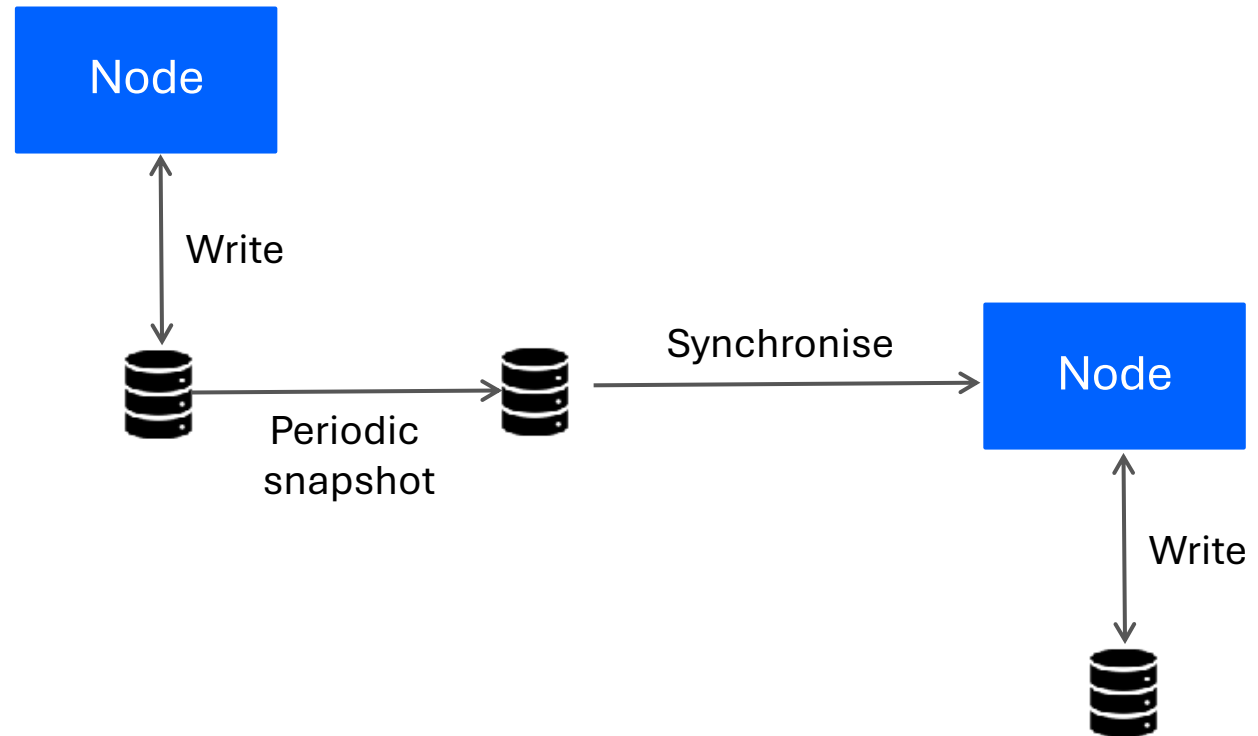
- Writes are added to a non-mirrored, volatile, in-memory journal.
- Allows for buffering of many seconds (or even minutes!) of writes to absorb replication performance problems without host impact.
- Backed up by a non-volatile bitmap.
- Journal sizes vary from 8 GiB to 32 GiB per node.
- GlobalMirror used this method.



Two methods to provide Asynchronous Replication

#2 Cycling Mode

- Uses change volumes to periodically copy data from the production system to the recovery system.
- Cycling frequency is based on the time until the RPO is exceeded.
- Change volumes protect consistent images at the recovery location.
- Number of volumes affect the pause time while taking snapshots.
- GlobalMirror with Change Volumes used this method.



Policy-based Replication provides Adaptive Replication to switch between Journaling and Cycling modes automatically

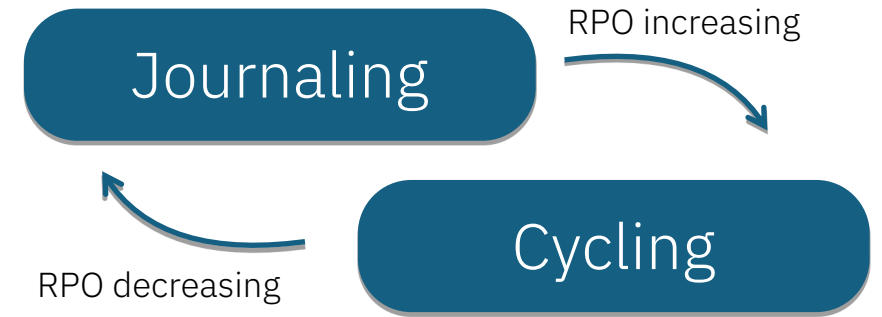
Volume Groups:

- Prefer to use journaling mode
- May switch to cycling mode (if the host throughput is greater than the replication throughput)
- With higher RPO alerts will switch to cycling mode first. (if all RPO are the same, those doing highest write bandwidth will switch first)

The system always aims to meet the RPO for all volume groups, regardless of the mode.

The current RPO of a volume group is reported.

There is no user control or visibility over the mode (if there are zero tunable settings).



Monitoring

Replication status can be seen under

Volumes > Volume Groups

Selecting a volume group and selecting the Policies:

- Displays detailed replication status for the group
- Provides links for the actions for replication that can be performed on the group


The screenshot shows the IBM FlashSystem 9100 GUI interface. The browser address bar indicates the URL is <https://zooty/gui#volumes-volumegroups>. The user is logged in as 'superuser' with the role of 'Security Administrator'. The page title is 'myAppGroup1' and it shows '40.00 GiB Total Group Capacity'. The navigation menu on the left includes options for Home, Replication Policy, Volumes (4), Policies (1), and Snapshots (0). The main content area displays the 'Replication' status for the 'GoldTierPolicy' with a topology of '2 Site, Asynchronous' and an RPO Alert of '1 minute'. A diagram shows data flowing from the 'zooty' production copy to the 'snazzy' recovery copy. The recovery copy status is 'Recovery point within policy' with '0 seconds behind the production copy'. A button 'Assign internal snapshot policy' is highlighted in a blue box. The bottom status bar shows performance metrics: Latency (144 µs), Bandwidth (426 MBps), and IOPS (9,995).

Metric	Value	Read	Write
Latency	144 µs	111 µs	244 µs
Bandwidth	426 MBps	319 MBps	105 MBps
IOPS	9,995	7,489	2,506

Policy Based Replication - Performance

- Significant improvement in maximum throughput compared to Global Mirror - compare with (yellow) non-replicated!
- Lower latency compared to Global Mirror.
- Performance scales throughout the portfolio.
- Cycling mode roughly equivalent to the performance of taking snapshots.

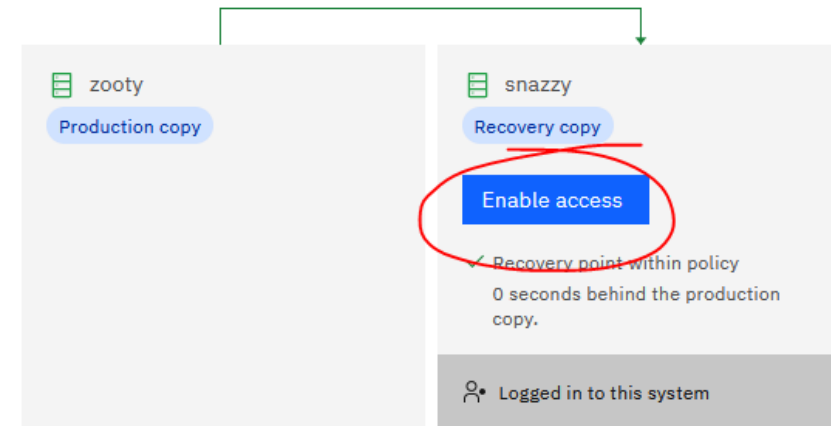
FlashSystem 9500 running 8.5.2 in journaling mode. 16KiB 70/30/50%



Policy Based Replication - Failover

- A failover is performed by enabling access to the recovery copy of the volume group.
- This action must be taken on the system with the recovery copy.
- The volume group comes online with the volumes and data at the point in time represented by the recovery point.
- The production and recovery volume groups each become 'independent'.
- Host I/O and configuration changes are permitted on both copies.
- Any configuration changes made are not replicated (until replication is restarted).

Must be performed on the recovery system.
Using the GUI



Using the CLI

```
lsvolumegroupreplication  
<volumegroup_id/name>
```

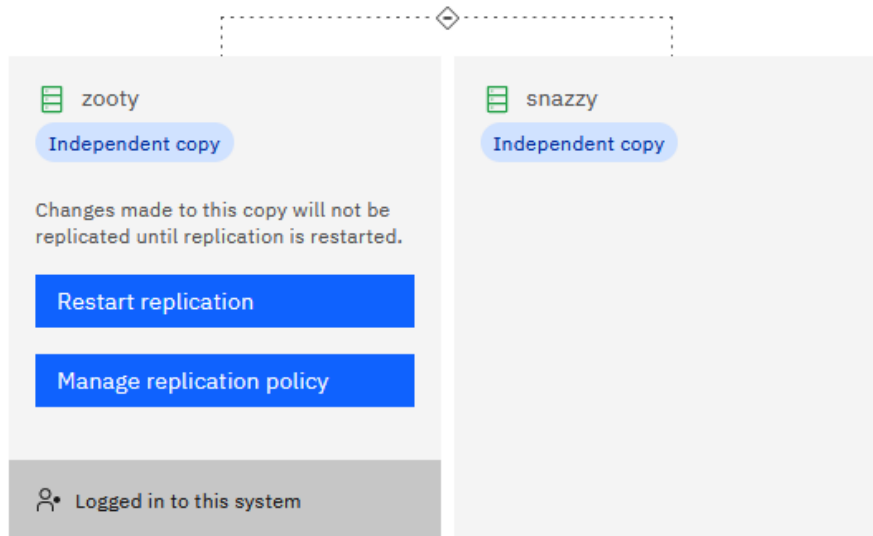
```
chvolumegroupreplication -mode  
independent
```

```
<volumegroup_id/name>
```

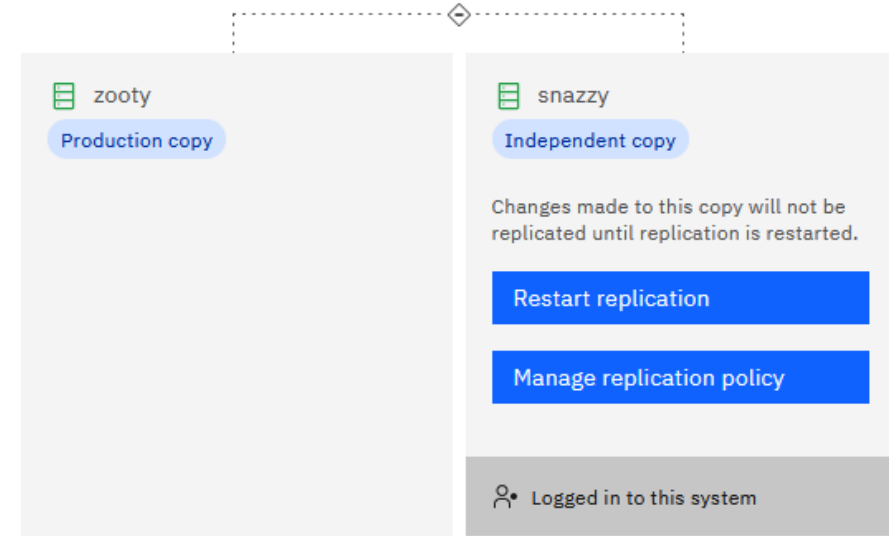
Failover

Here is what the resulting independent state looks like on each system (when connected):

Logged in to the system at production

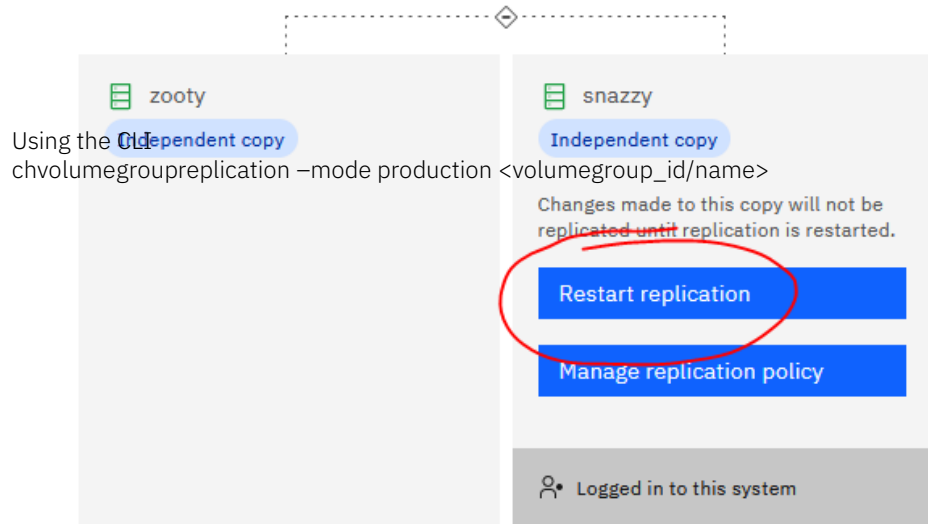


Logged in to the system at recovery

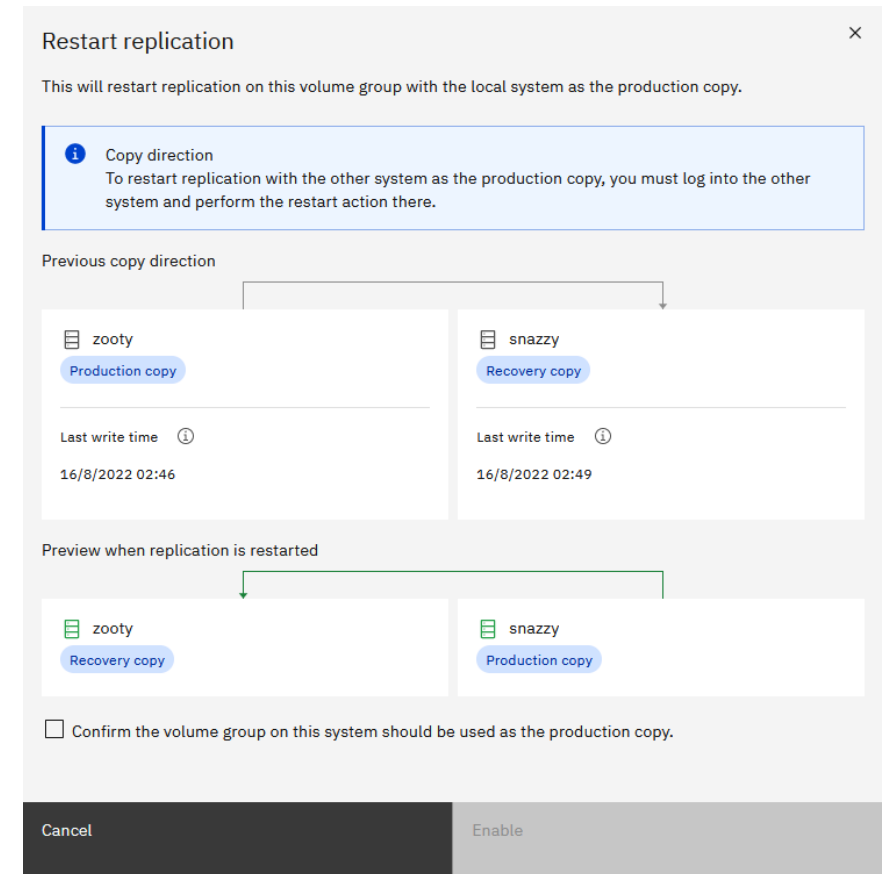


Failback/restart

A failback or resuming replication is performed by restarting replication from the system that is to continue as the production copy of the volume group.



The CLI and GUI both indicate the previous direction of replication and the last time of any host writes performed to volumes in the volume group to help ensure replication is being restarted in the appropriate direction.



Recapping the benefits of policy-based replication



>4X

More than four times the host throughput compared to Global Mirror and HyperSwap.

32K

Replicate up to 32,500 volumes from a single system (versus 2,500 for GMCV and 2,000 for HyperSwap).

4PiB

Replicate up to 4 pebibyte per system, which reduces the number of systems and resources required.

IBM Policy-Based High Availability (PBHA)

Simplicity, performance and efficiency

Policy-Based High Availability



Policy-based high availability (PBHA)

- Active/active mirroring between independent systems
- FlashSystem 5045 through FlashSystem 9500
- No single point of failure
- Up to 4X higher performance
- Partitions are a key component to IBM's new FlashSystem grid architecture

IBM SAN Volume Controller Cluster_9.71.19.88

Partition01 Overview

Replication Policy

Policy name	Topology
RepPol_01	2 Site High Availability

High Availability Status

Two 'bath' Production Systems are shown, connected via a green line, indicating active/active mirroring. Each system is labeled 'bath' and 'Production System'. Management is through this system, and it is the preferred management system. A 'Go to system' button is available for each system.

Connectivity

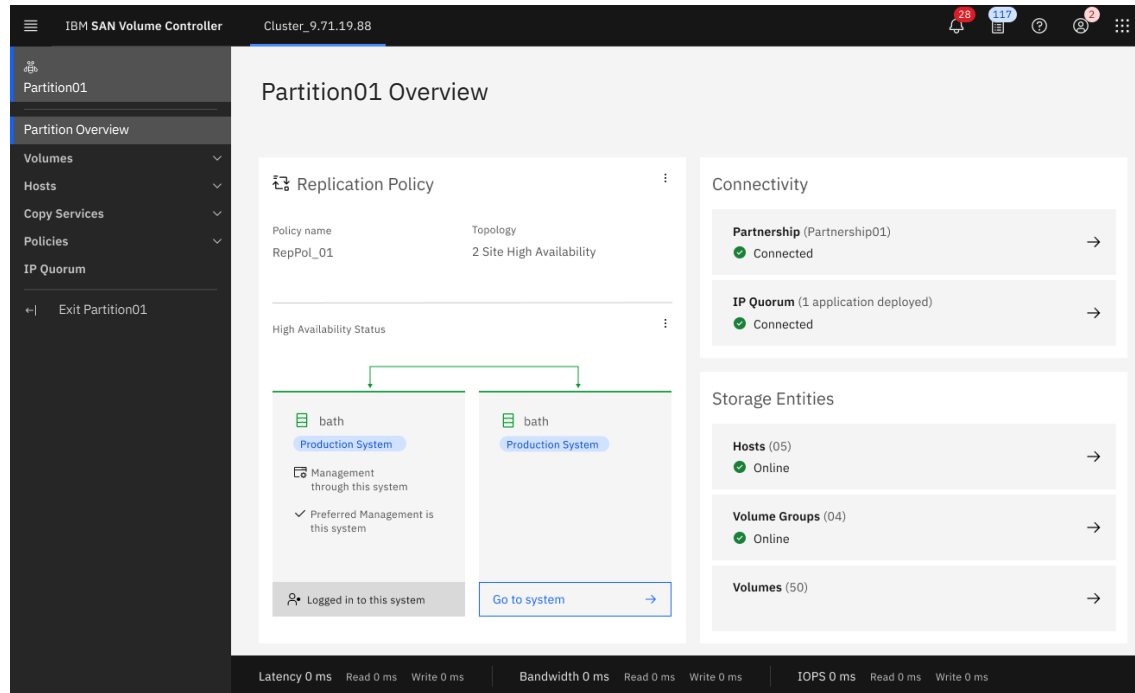
- Partnership (Partnership01) - Connected
- IP Quorum (1 application deployed) - Connected

Storage Entities

- Hosts (05) - Online
- Volume Groups (04) - Online
- Volumes (50)

Latency 0 ms | Read 0 ms | Write 0 ms | Bandwidth 0 ms | Read 0 ms | Write 0 ms | IOPS 0 ms | Read 0 ms | Write 0 ms

Storage partitions simplify the configuration, management and monitoring, with a single point of control for HA



Create and populate

Create a storage partition in a few clicks, and create or add any number of hosts, volume groups and volumes.

Policy-based

Add a HA replication policy to the partition, and everything within it will automatically be configured to be highly available, with all remote provisioning handled automatically.

Highly available

In a disaster, the storage partition automatically manages its availability ensuring that applications remain accessible, with per-partition control over which system is preferred in the event of a loss-of-connectivity.

Scoped

HA problems have zero impact on any non-HA volumes: HA and non-HA storage can happily co-exist on the same system.

Policy-based High Availability

Fixing the HyperSwap limitations



Simplifies the configuration, management and monitoring of HA through the use of Storage partitions, policies and volume groups.



Addresses the performance, usability, and scalability challenges with HyperSwap.

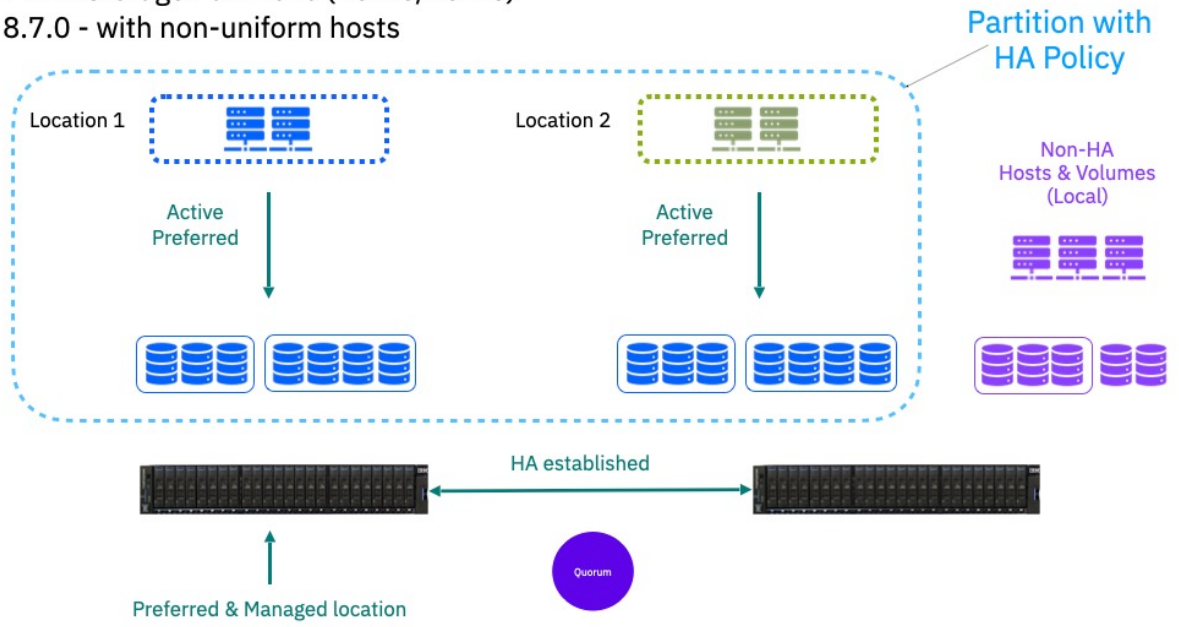


Complete ground-up re-design and implementation of the entire HA software within Storage Virtualize.



8.7.0 provides high-availability over Fibre Channel and short-distance NVMe Ethernet partnerships.

PBHA Storage Partitions (Active/Active) 8.7.0 - with non-uniform hosts



SP 1 Overview

Replication Policy

Policy name: RepPoL_01 | Topology: 2 Site High Availability

High Availability Status: [Status]

Connectivity:

- Partnership (Partnership01) - Connected
- IP Quorum (1 application deployed) - Connected

Storage Entities:

- Hosts (05) - Online
- Volume Groups (04) - Online
- Volumes (50) - Online

Performance Metrics:

Latency 0 ms | Read 0 ms | Write 0 ms | Bandwidth 0 ms | Read 0 ms | Write 0 ms | IOPS 0 ms | Read 0 ms | Write 0 ms


8.6.1 - Policy-based High Availability

- Limitations of HyperSwap today :
 - Requires a single clustered system that is split across sites - all or nothing configuration- i.e. all volumes must be in HyperSwap configuration
 - Split brain requires half the cluster to stop
 - Performance limits
 - Hardware at each site must be the same
 - Software the same and upgraded at the same time at both sites

All of these limitations are addressed by the new Policy-based HA functionality introduced in 8.6.1

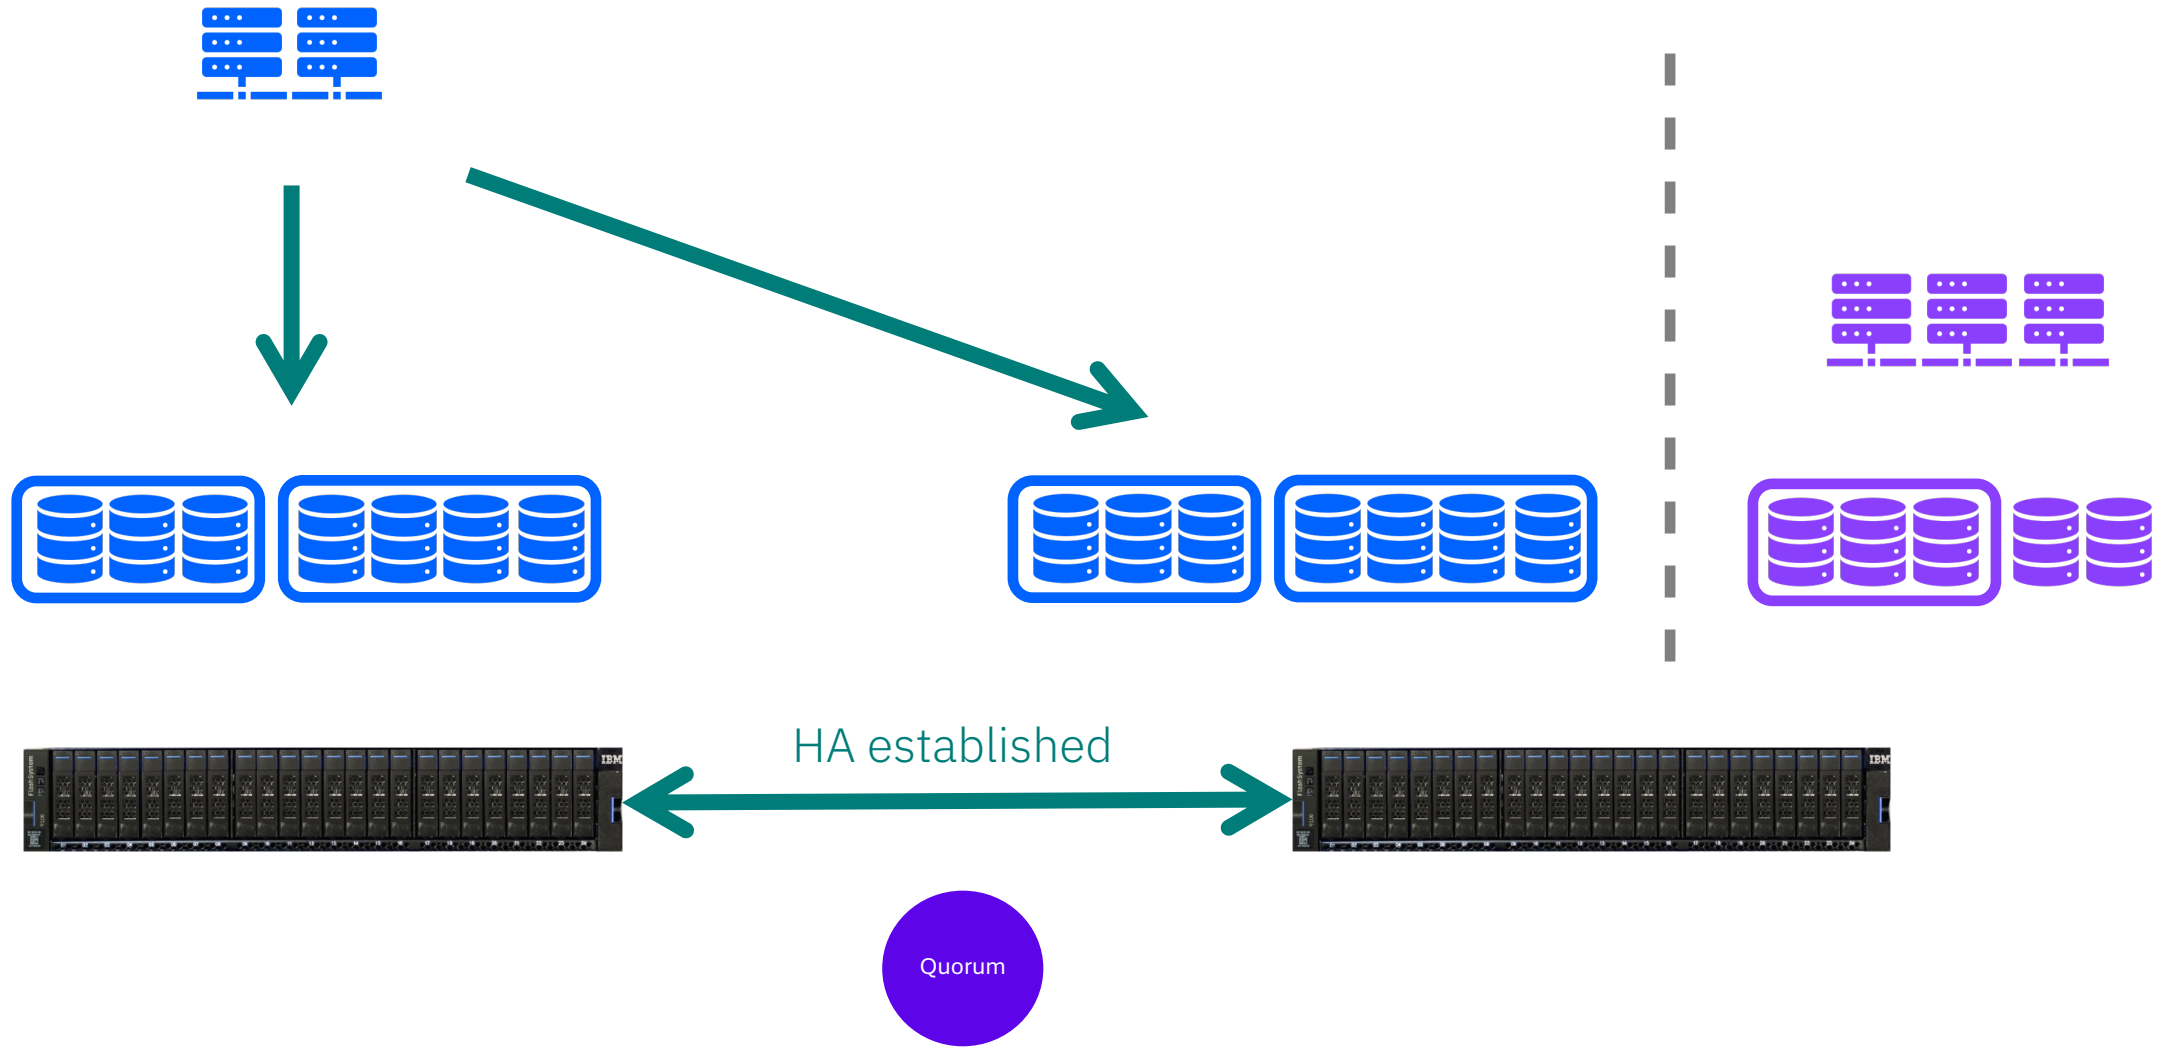
The first release is targeted at new installations and proof of concepts, not existing systems
Support for adding HA to an existing non-HA system is planned to be included in the next LTS*
There will be functionality missing from the first release, but many enhancements are already being planned for upcoming releases

PBHA - Introducing Storage partitions




- A Storage partition contains a set of hosts, volume groups, and volumes.
- Storage partitions can be configured to be highly available, automatically configuring every host and volume within them to be highly available.
- The management of hosts, volume groups, and volumes automatically swaps between systems to provide highly available management access.
- Highly available storage partitions have an attribute that defines which system is preferred for management access when HA is established.
- Max 4 Storage partitions per system

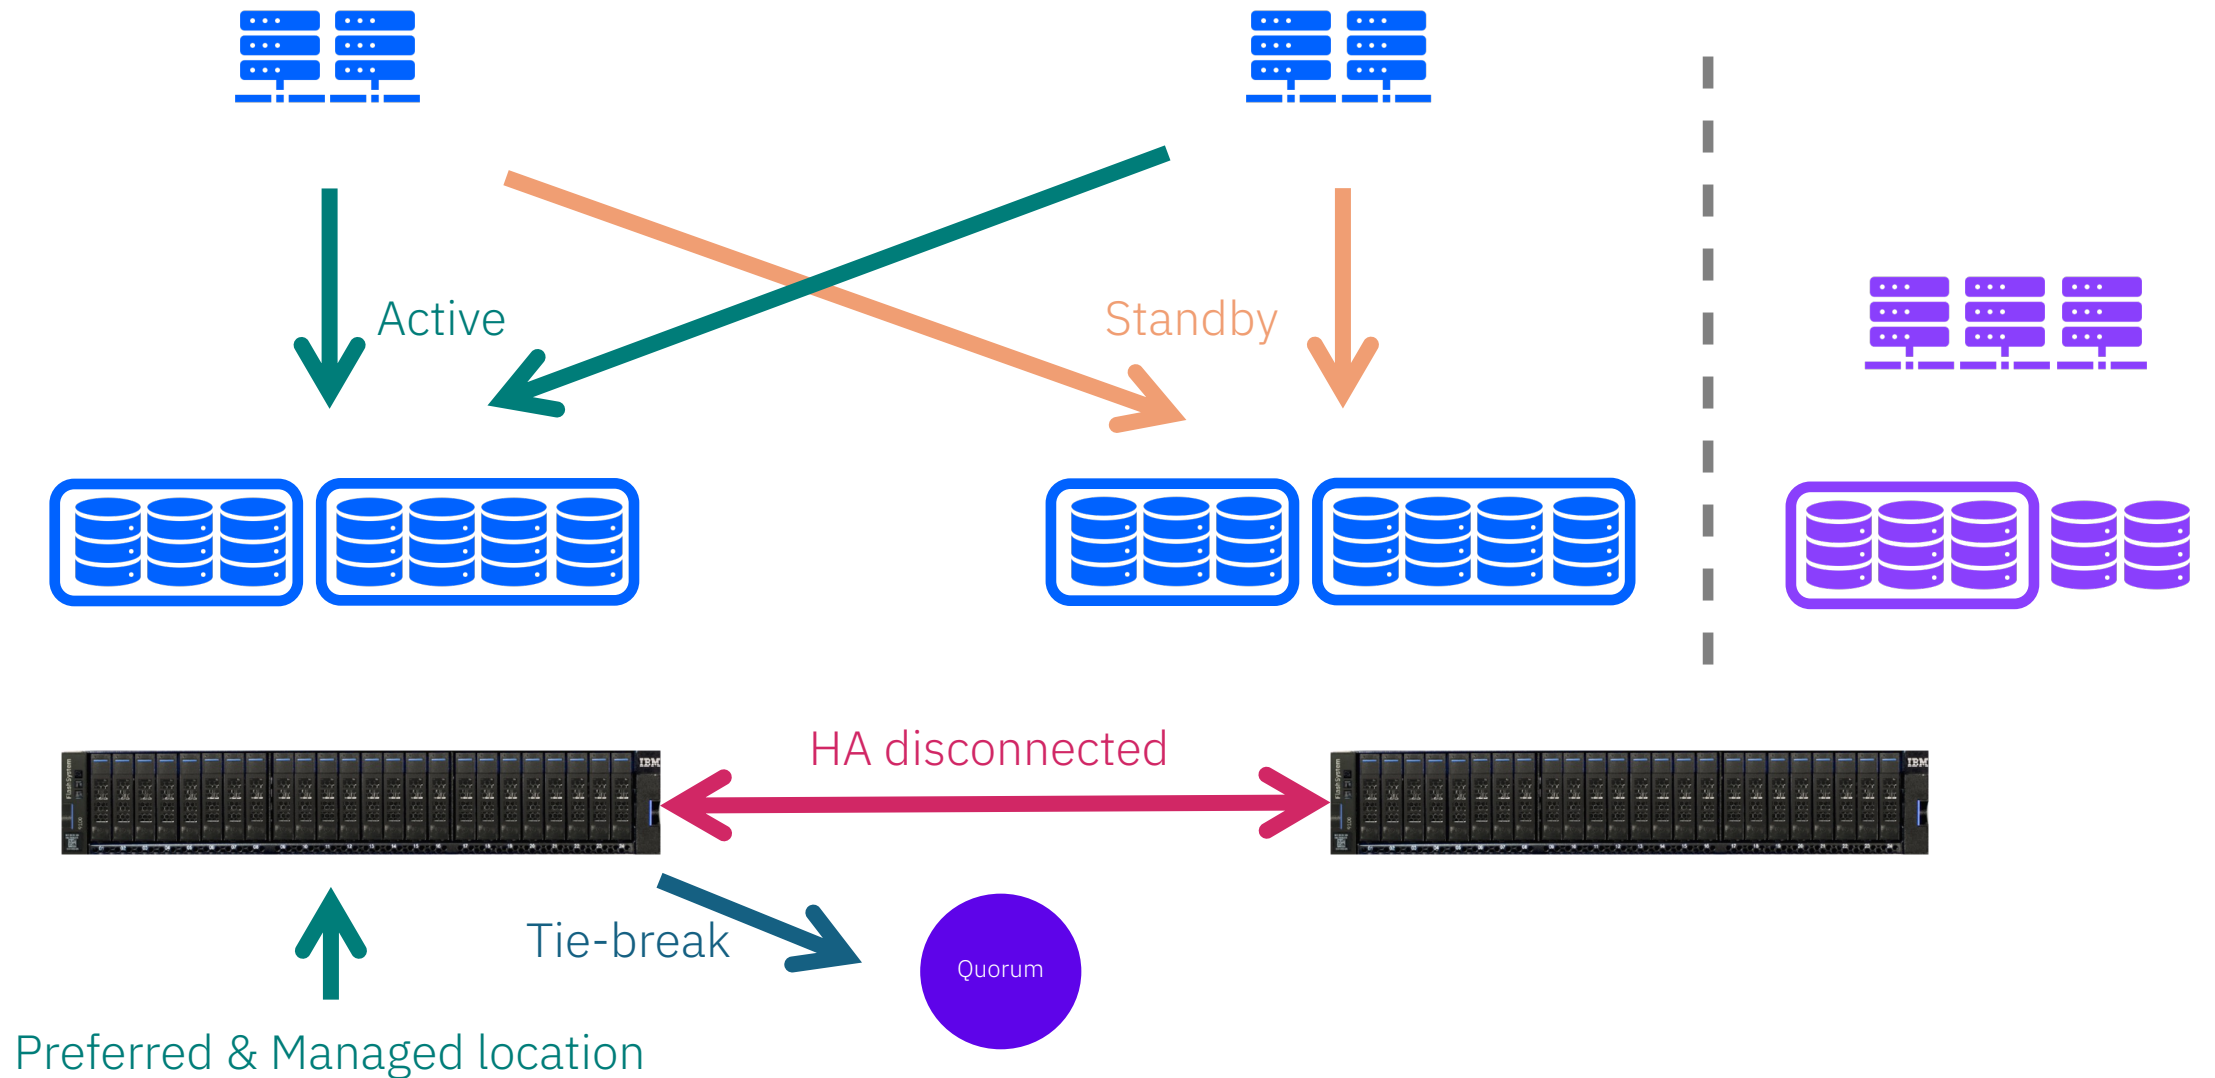
HA storage partitions - normal running



HA storage partitions - preferred management site



HA storage partitions - split brain scenario



Recapping the benefits of policy-based high availability



>4X

More than four times the host throughput compared to Global Mirror and HyperSwap.

32K

Replicate up to 32,500 volumes from a single system (versus 2,500 for GMCV and 2,000 for HyperSwap).


4PiB

Replicate up to 4 pebibyte per system, which reduces the number of systems and resources required.

Policy-based three-site replication (HA + DR)

Simplicity, performance and efficiency

Policy-Based 3-Site Replication



Simple policy-based management

Available for new configs and to upgrade existing:

- Add DR to existing 2-site HA configuration
- Add HA to existing 2-site Async replication

Supported on NVMe FlashSystems and SVC


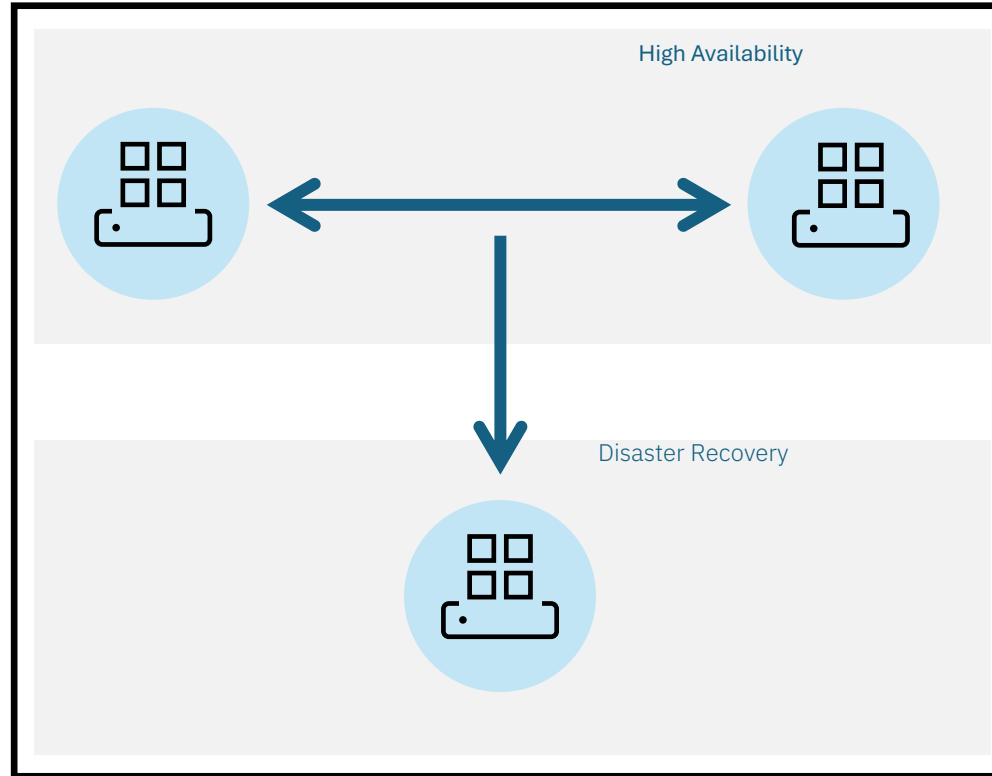
The screenshot displays the IBM SAN Volume Controller (SVC) management interface for 'System A'. The main view is titled 'My_salesforce_partition' and shows a 'Replication overview' diagram. This diagram illustrates a 'High availability' configuration between 'System A' and 'System B', with 'System A' being the 'Logged in here' system. Below this, a 'Disaster recovery' link points to 'My_async_partition', which is associated with 'System C'. The right-hand side of the interface provides details on 'Storage objects', including 'Hosts (2)' (Available), 'Volume groups (06)' (Healthy), and 'Connectivity' (High availability partnership: Online, IP Quorum: Connected, Disaster recovery partnership: Online). A 'Manage partition' button is visible in the top right. The bottom status bar shows performance metrics: Latency 0 ms, Read 0 ms, Write 0 ms, Bandwidth 0 MBps, Read 0 MBps, Write 0 MBps, and IOPS 0.

Enterprise-grade business continuity

Active/Active high-availability
across metro distances

+

High performance, low Recovery Point Objective (RPO) replication to
DR site



Enterprise-grade business continuity

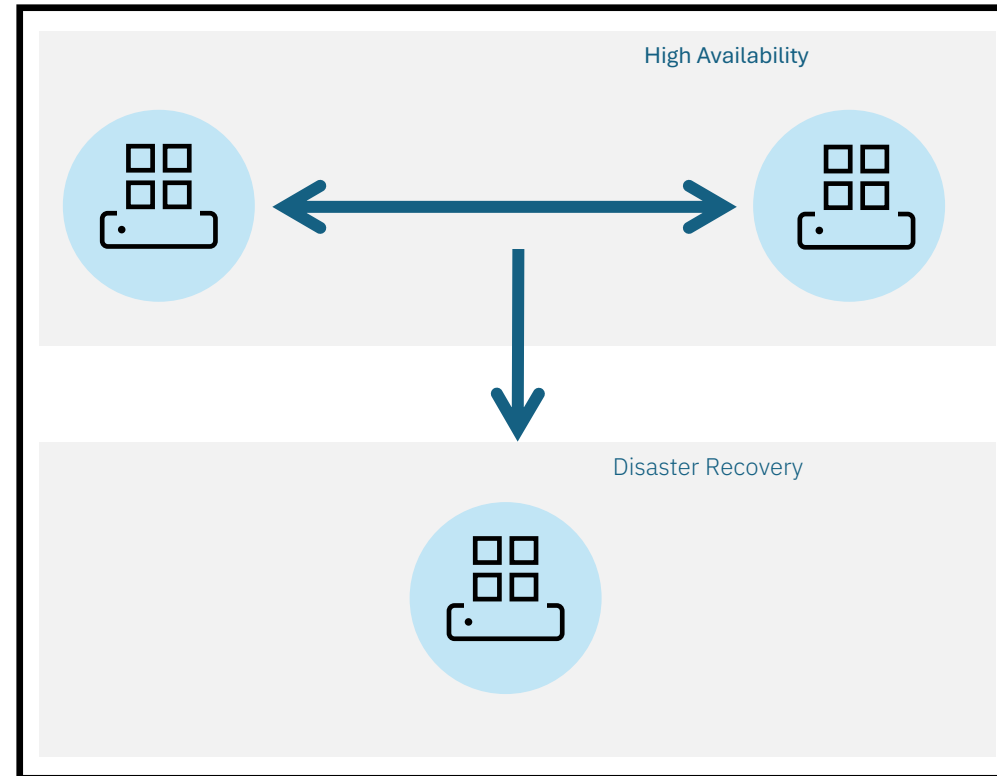
Built-in management, no external software

+

100% Availability Guarantee

+

Meet regulatory and business requirements



Replication
has never
been easier

- Add replication to a partition

The screenshot displays the IBM FlashSystem 5200 management interface for a storage partition named 'myPartition'. The interface is divided into several sections:

- Navigation Sidebar (Left):** Lists various management options: myPartition, Volume groups, Volumes, Volume mappings, Hosts, Replication policies, and Snapshot policies. An 'Exit myPartition' button is at the bottom.
- Header:** Shows 'IBM FlashSystem 5200', the host name 'pine-c', and the current storage partition 'Storage partition: myPartition'. Utility icons for notifications, lists, help, and user profile are on the right.
- Main Content Area:**
 - myPartition:** A 'Manage partition' button is located in the top right.
 - Replication overview:** Features a storage icon and the text 'Replication is not configured for this partition.' with a blue 'Configure replication +' button.
 - Storage components:** A list of components with expandable arrows:
 - Hosts (1): Online (with a green checkmark)
 - Volumes (10)
 - Volume groups (4)
 - Connectivity:** Features a storage icon and the text 'There is no partnership assigned to this partition.'
- Bottom Status Bar:** Displays performance metrics: Latency 0 ms, Read 0 ms, Write 0 ms; Bandwidth 0 MBps, Read 0 MBps, Write 0 MBps; and IOPS 0, Read 0, Write 0.

Replication
has never
been easier

- Configure HA, DR or HA+DR right out of the box
- Or start with 2-site HA or DR today and add a 3rd site in future

IBM FlashSystem 5200 pine-c Storage partition: myPartition

Configure policy-based replication for myPartition

Complete the steps to configure replication for the storage partition

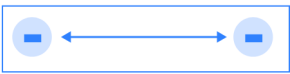
- Select replication type
- Configure replication
- Summary Review

Select replication type

Select the replication types to be used by this storage partition.

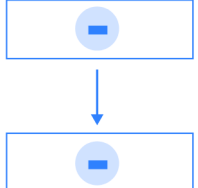
High availability

Active/active high availability between two independent storage systems.



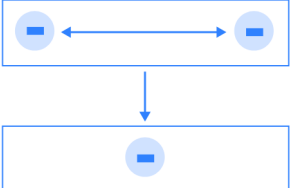
Disaster recovery

Asynchronous replication to another storage system for disaster recovery.



High availability and disaster recovery

Active/active high availability between two independent storage systems, with asynchronous replication to a third system for disaster recovery.



Cancel Skip Continue

Simple graphical setup

- Step-by-step guided configuration

IBM FlashSystem 5200 | pine-c | Storage partition: myPartition

Configure policy-based replication for myPartition

Complete the steps to configure replication for the storage partition

- Select replication type
- Configure replication**
- Summary Review

Configure replication

Choose one or both the replication types that are available to configure on your storage partition.

myPartition

High availability

pine-c
Logged in here

pecan-c

- Partnership configured
- Pool links configured
- IP quorum application downloaded


Cancel | Skip | Continue

HA storage partition

The screenshot displays the IBM FlashSystem 5200 management console interface for a storage partition named 'myPartition'. The interface is divided into several sections:


- Header:** Shows 'IBM FlashSystem 5200', 'pine-c', and 'Storage partition: myPartition'. A 'Manage partition' button is located in the top right.
- Left Navigation Panel:** Lists various management options: myPartition, Volume groups, Volumes, Volume mappings, Hosts, Replication policies, Snapshot policies, IP quorum, Partnerships, and an 'Exit myPartition' button.
- Replication overview:** A diagram showing two hosts, 'pine-c' and 'pecan-c', connected by a green line labeled 'HA established'. 'pine-c' is marked as 'Logged in here'.
- Storage components:** A list of components with expandable arrows:
 - Hosts (1): Online
 - Volumes (10)
 - Volume groups (4)
- Connectivity:** A list of connectivity settings:
 - HA partnership (pecan-c): Configured
 - IP quorum: Connected
- Footer:** Performance metrics including Latency (0 ms), Bandwidth (0 MBps), and IOPS (0).

HA storage partition with async DR

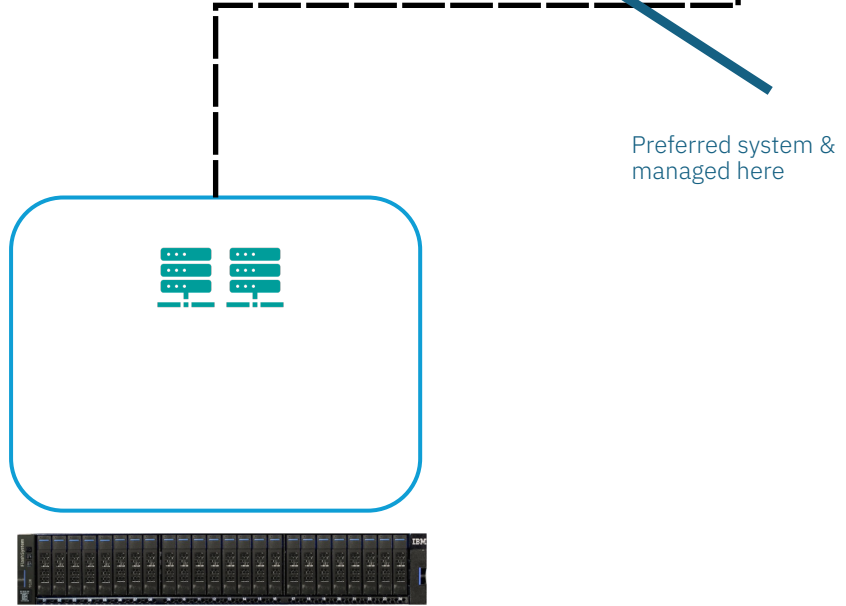


Preferred system & managed here

HA storage
partition with
async DR



Linked partitions



Preferred system &
managed here

Starting with high availability

The screenshot displays the IBM FlashSystem 5200 management interface for a storage partition named 'myPartition'. The interface is divided into several sections:

- Navigation Panel (Left):** Lists various management options including Volume groups, Volumes, Volume mappings, Hosts, Replication policies, Snapshot policies, IP quorum, and Partnerships. An 'Exit myPartition' button is at the bottom.
- Header:** Shows 'IBM FlashSystem 5200', the current host 'pine-c', and the selected partition 'Storage partition: myPartition'. Utility icons for notifications, clipboard, help, and user profile are on the right.
- myPartition Overview (Center):** Features a 'Replication overview' diagram showing two hosts, 'pine-c' (logged in) and 'pecan-c', connected by a green line labeled 'HA established'. A 'Manage partition' button is in the top right.
- Storage components (Right):** A list of system components with their status and expandable arrows:
 - Hosts (1): Online
 - Volumes (10)
 - Volume groups (4)
- Connectivity (Right):** Shows the status of high availability and quorum settings:
 - HA partnership (pecan-c): Configured
 - IP quorum: Connected
- Footer:** Performance metrics for Latency, Bandwidth, and IOPS, all currently at 0.

Click to add disaster recovery

The screenshot displays the IBM FlashSystem 5200 management interface for a storage partition named 'myPartition'. The interface is divided into several sections:

- Header:** Shows 'IBM FlashSystem 5200', 'pine-c', and 'Storage partition: myPartition'. There are also notification, list, help, and user icons.
- Left Navigation Panel:** Lists various management options: myPartition, Volume groups, Volumes, Volume mappings, Hosts, Replication policies, Snapshot policies, IP quorum, Partnerships, and an 'Exit myPartition' button.
- main Content Area:**
 - myPartition:** A title bar with a 'Manage partition' dropdown menu.
 - Replication overview:** A diagram showing two hosts, 'pine-c' (logged in here) and 'pecan-c', connected by a line labeled 'HA established'. A context menu is open over this diagram, with the 'Add disaster recovery' option highlighted. Other menu items include 'Manage preferred location' and 'Remove high availability'.
- Storage components:** A list of components with expandable arrows:
 - Hosts (1) Online
 - Volumes (10)
 - Volume groups (4)
- Connectivity:** A list of connectivity components:
 - HA partnership (pecan-c) Configured
 - IP quorum Connected

At the bottom of the interface, there is a performance summary bar with the following data:

- Latency: 0 ms
- Read: 0 ms
- Write: 0 ms
- Bandwidth: 0 MBps
- Read: 0 MBps
- Write: 0 MBps
- IOPS: 0
- Read: 0
- Write: 0

Click to add
disaster recovery

Configure Partnership

IBM FlashSystem 5200 | pine-c | Storage partition: myPartition

Configure policy-based replication for myPartition

Complete the steps to configure replication for the storage partition

- Configure replication
- Summary Review

Configure replication

Choose one or both the replication types that are available to configure on your storage partition.

myPartition

pine-c (Logged in here) — HA established — pecan-c

Disaster recovery

DR Partner

- Configure partnership >
- Create partition >
- Link Pools >

Cancel | Skip | Continue

click to add
disaster recovery

Step-by-step
configuration

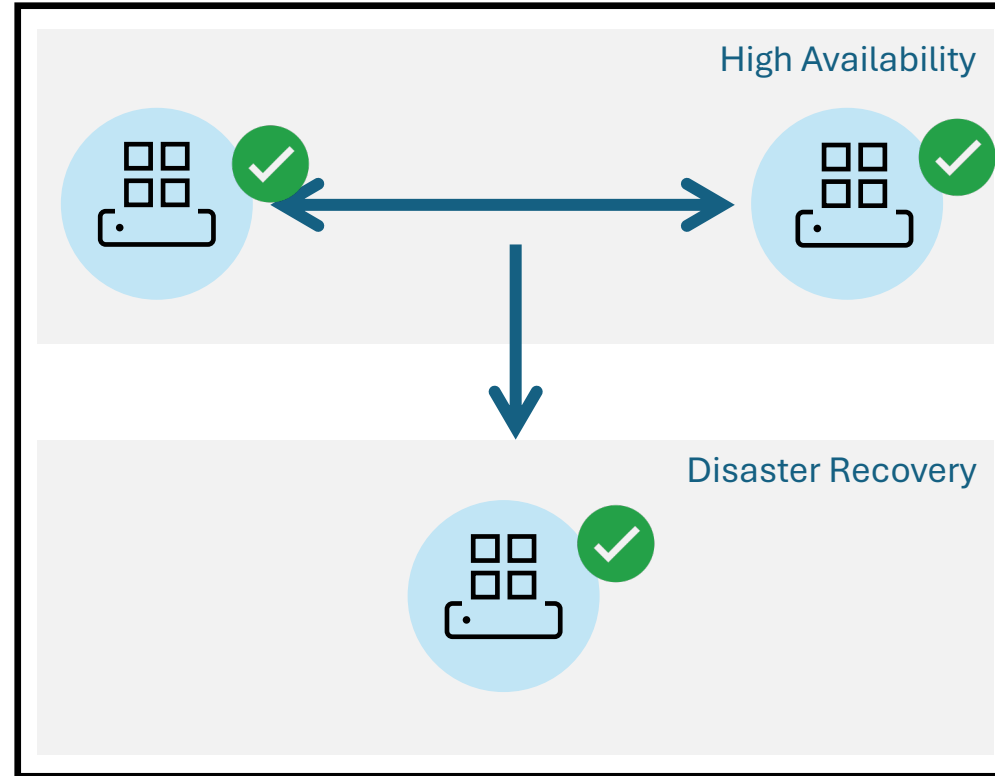
Ready for use

The screenshot displays the IBM FlashSystem 5200 management interface for a storage partition named 'myPartition'. The interface is divided into several sections:

- Navigation Panel (Left):** Lists various configuration options including Volume groups, Volumes, Volume mappings, Hosts, Replication policies, Snapshot policies, IP quorum, and Partnerships. An 'Exit myPartition' button is at the bottom.
- Header:** Shows 'IBM FlashSystem 5200', the current host 'pine-c', and the selected 'Storage partition: myPartition'. It includes notification, list, help, and user icons.
- myPartition Overview (Top Right):** A 'Manage partition' button with a dropdown arrow.
- Replication Overview (Center):** A diagram showing two hosts, 'pine-c' and 'pecan-c', connected by a green line labeled 'HA established'. Below them is a 'DR unused' status box. A third host, 'almond-c', is shown in a separate section labeled 'myPartition_DR'.
- Storage Components (Right):** A list of system components:
 - Hosts (1):** Online.
 - Volumes (10):** (No details shown)
 - Volume groups (4):** No disaster recovery policies configured.
- Connectivity (Right):** A list of configured partnerships:
 - HA partnership (pecan-c):** Configured.
 - IP quorum:** Connected.
 - DR partnership (almond-c):** Configured.
- Performance Metrics (Bottom):** A bar at the bottom shows metrics for Latency, Bandwidth, and IOPS, all currently at 0.

Enterprise-grade business continuity

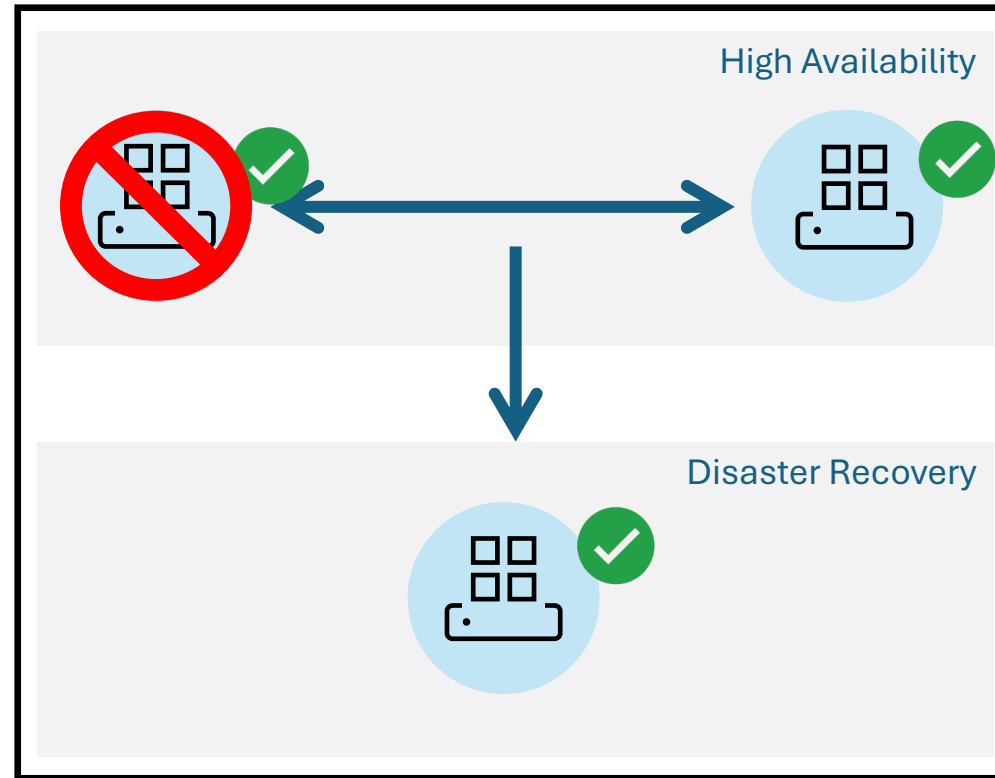
3-Site configurations provide organizations with real peace-of-mind which enable them to focus on other important tasks.



Enterprise-grade business continuity

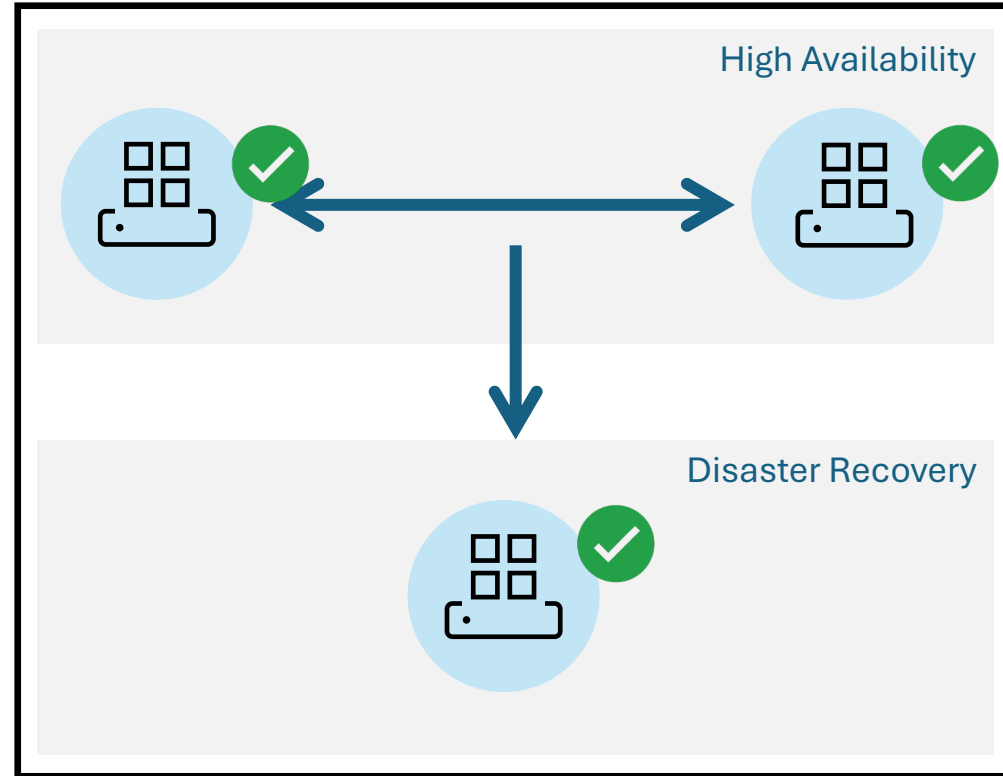
3-Site configurations provide organizations with real peace-of-mind which enable them to focus on other important tasks.

3-site configurations allow organizations to survive the loss of one system and keep the business moving.



Enterprise-grade business continuity

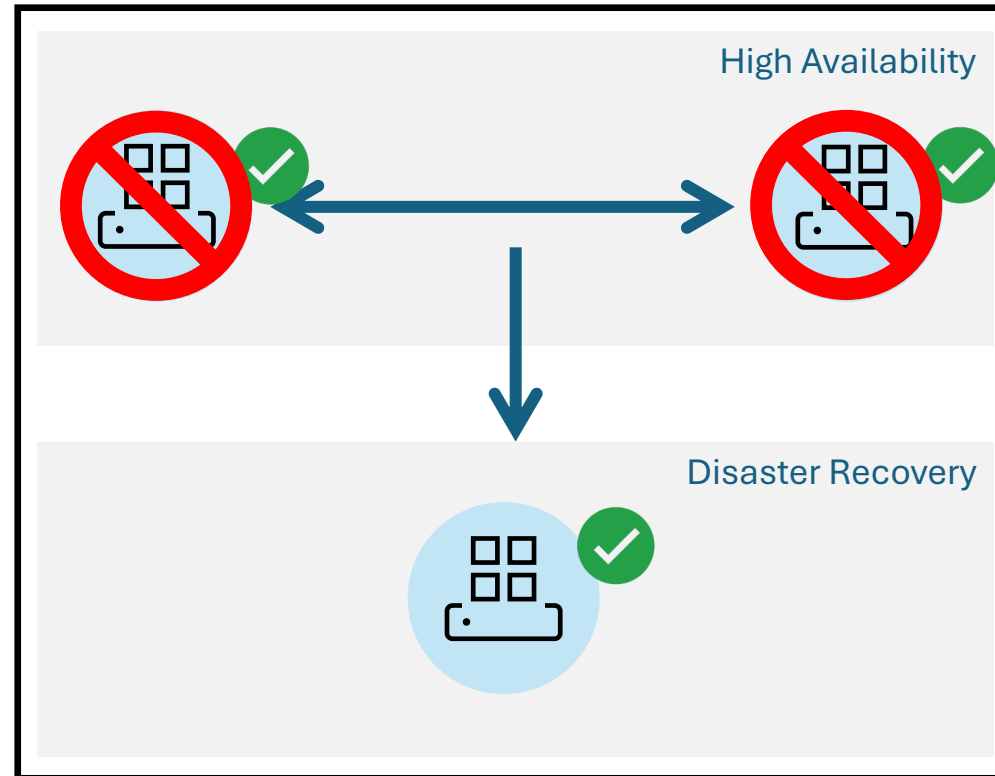
3-Site configurations provide organizations with real peace-of-mind which enable them to focus on other important tasks.



Enterprise-grade business continuity

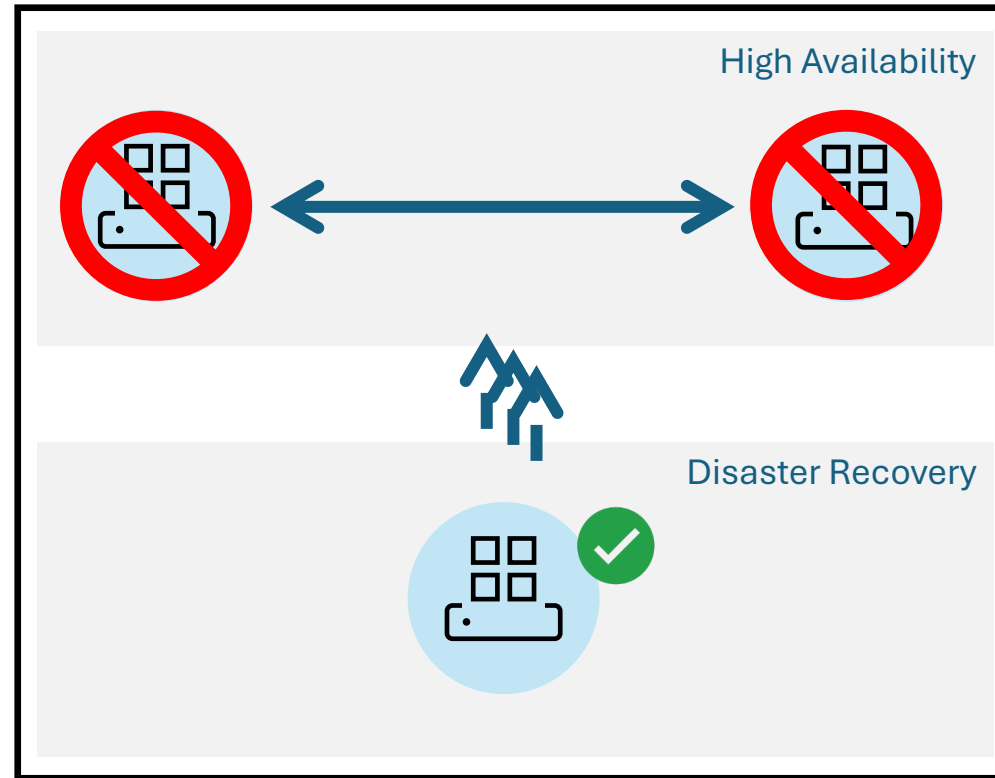
3-Site configurations provide organizations with real peace-of-mind which enable them to focus on other important tasks.

3-site configurations allow organizations to survive the loss of two systems and still keep the business moving.



Enterprise-grade business continuity

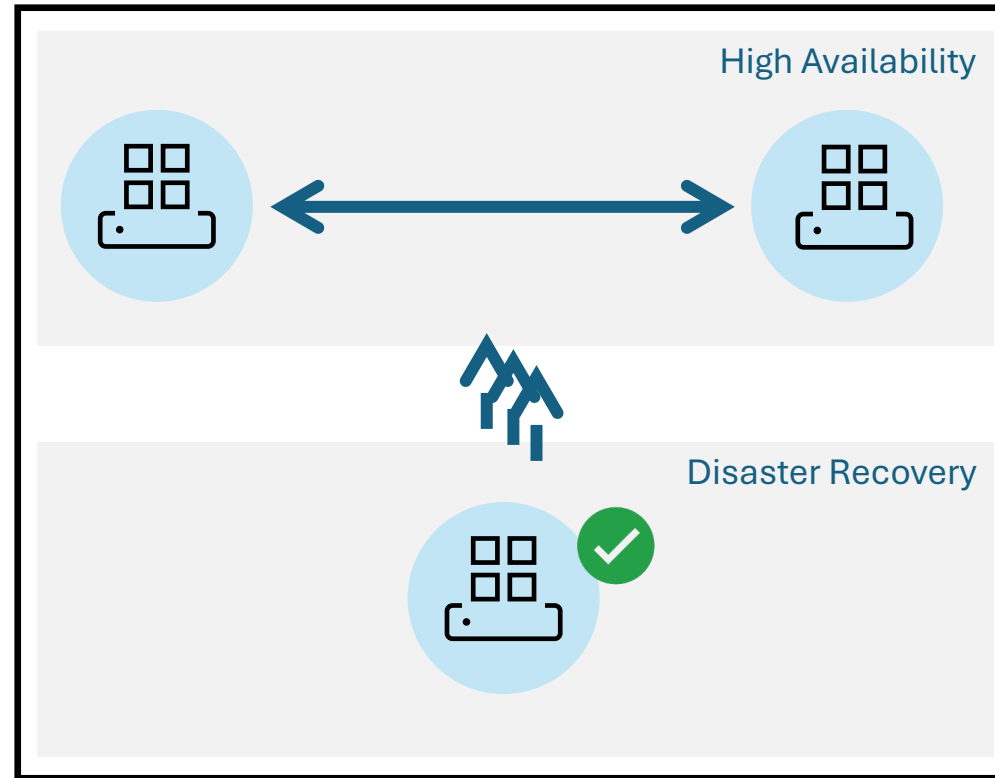
After a failure, while the organization continues to work, the failed systems can be brought back online



Enterprise-grade business continuity

After a failure, while the organization continues to work, the failed systems can be brought back online

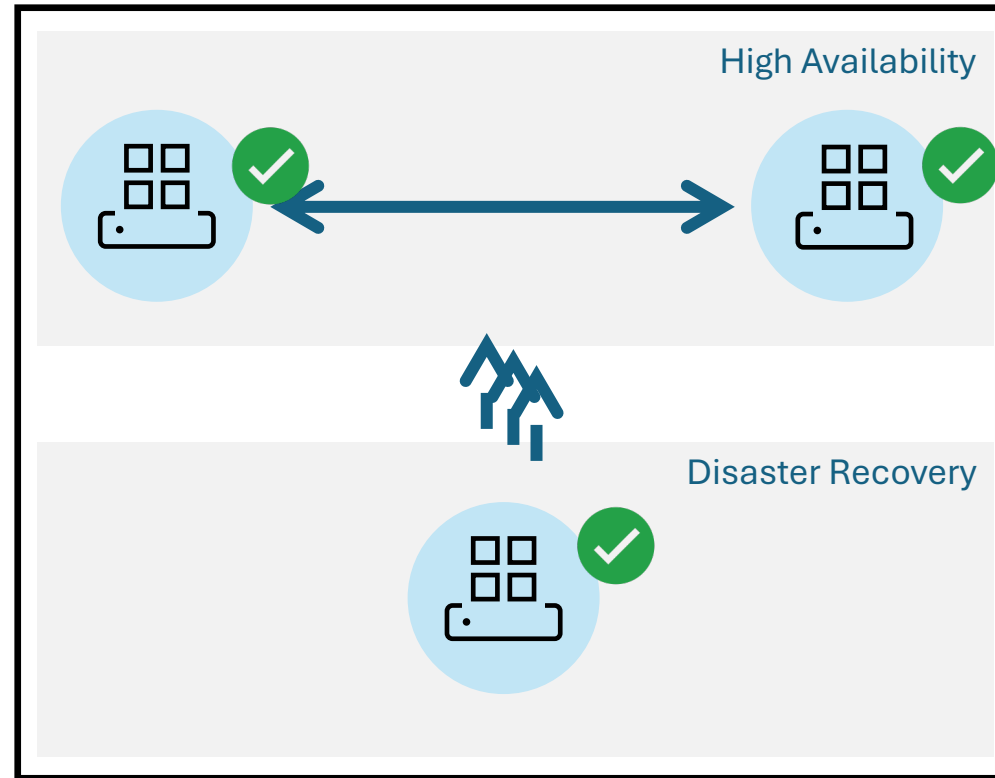
Once access on the production site is restored, the system will incrementally resynchronize changes from the DR site.



Enterprise-grade business continuity

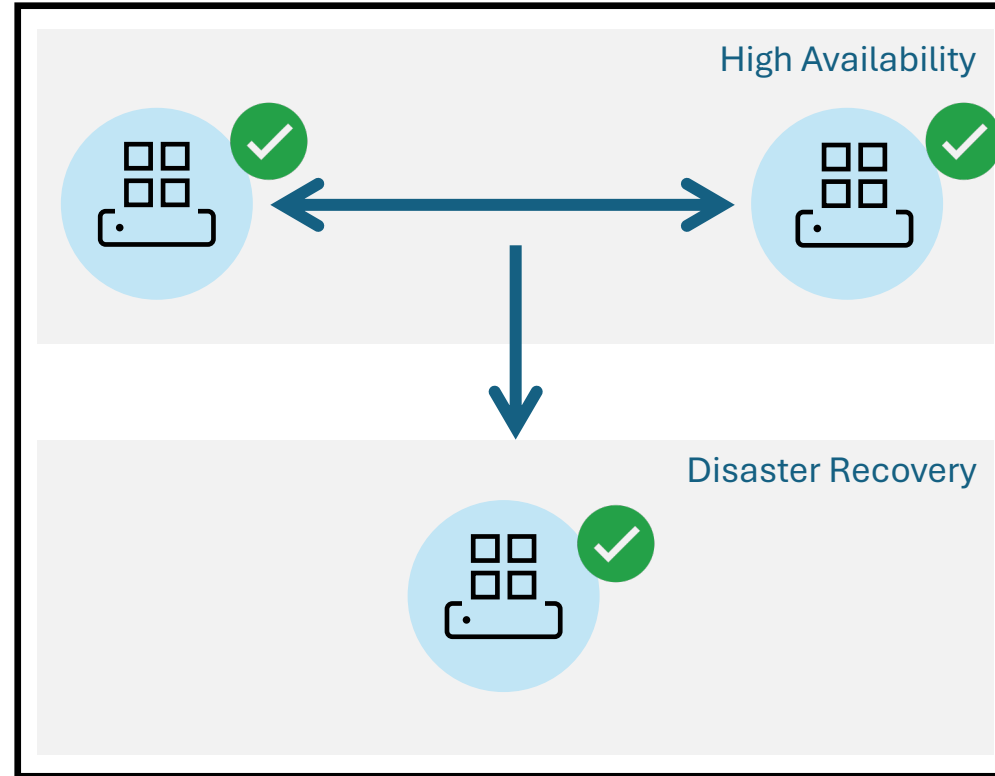
After a failure, while the organization continues to work, the failed systems can be brought back online

Once access on the production site is restored, the system will incrementally resynchronize changes from the DR site.



Enterprise-grade business continuity

3-site configurations can help organizations move forward and succeed.



IBM