



ADVANCED TECHNOLOGY GROUP (ATG)



Accelerate with ATG Webinar: Storage Virtualize 8.7.1 Technical Update

Byron Grossnickle

IBM Storage Virtualize/FlashSystem SME

byrongro@us.ibm.com

Chris Bulmer

Storage Virtualize High Availability and Replication Architecture



Accelerate with ATG Technical Webinar Series

Advanced Technology Group experts cover a variety of technical topics.

Audience: Clients who have or are considering acquiring IBM Storage solutions. Business Partners and IBMers are also welcome.

To automatically receive announcements of upcoming Accelerate with IBM Storage webinars, Clients, Business Partners and IBMers are welcome to send an email request to accelerate-join@hursley.ibm.com.



2024 Upcoming Webinars – Register Here!

[Announcing the 10th Generation IBM DS8000 Enterprise Storage System](#) – October 17th, 2024

Important Links to bookmark:



ATG Accelerate Site: <https://ibm.biz/BdSUFN>

ATG MediaCenter Channel: <https://ibm.biz/BdfEgQ>

Offerings

Client Technical Workshops

- IBM DS8900F Advanced Functions – November 6-7 (Virtual)
- IBM Cyber Resiliency with IBM Storage Defender – November 13 (Virtual)
- IBM Fusion & Ceph: A Deep Dive into Next Gen Storage – November 19-20 (Virtual)
- IBM FlashSystem Deep Dive & Advanced Functions – December 3-4 (Herndon, VA)
- IBM Storage Scale System & Storage Scale Workshop

TechZone Test Drive / Demo's

- IBM Storage Scale and Storage Scale System GUI
- IBM Storage Virtualize Test Drive
- IBM DS8900F Storage Management Test Drive
- Managing Copy Services on the DS8000 Using IBM Copy Services Manager Test Drive
- IBM DS8900F Safeguarded Copy (SGC) Test Drive
- IBM Cloud Object Storage Test Drive - (Appliance based)
- IBM Cloud Object Storage Test Drive - (VMware based)
- IBM Storage Protect Live Test Drive
- IBM Storage Ceph Test Drive - (VMware based)

Please reach out to your IBM Representative or Business Partner for more information.

***IMPORTANT* The ATG team serves clients and Business Partners in the Americas, concentrating on North America.**

Registration Open!

Storage @ IBM TechXchange Conference 2024

October 21-24, 2024
Mandalay Bay | Las Vegas
#IBMTechXchange

Key Learnings

- Practical how-to advice
- Patterns and best practices
- Success stories, IBM PoV, proven techniques

Featured Products

- IBM Storage Defender
- IBM Storage Fusion
- IBM Storage Scale + IBM Storage Ceph
- IBM Tape + IBM SAN
- IBM Storage FlashSystem + IBM Storage DS8000

Collaborate. Learn. Play.



Community

- IBM Champions
- User Groups
- Tech Peers
- Business Partners



Sandbox

- Network
- Learn
- Collaborate
- Play





Accelerate your Career

- Labs (Instructor-Led, Self-paced)
- IBM Certification Testing
- Earn up to 25 hours in CPE credits

Breakout Sessions

- Trends and Directions
- User Groups
- Product Deep Dives
- Meet the Expert
- Professional Development
- Show the Code
- Birds of a Feather
- Academic/Research



Roadmaps

Go deep with people in the know and set the stage for where IBM is going in the future



<https://www.ibm.com/community/ibm-techxchange-conference/>

Game On!



Accelerate with ATG Survey

Please take a moment to share your feedback with our team!

You can access this 6-question survey via [Menti.com](https://www.menti.com) with code 1708 6924 or


Direct link <https://www.menti.com/alwhyze7z1gz>

Or


OR Code



Join at [menti.com](https://www.menti.com) | use code 1708 6924



Fun Question: What do you most enjoy about Autumn / Fall?





ADVANCED TECHNOLOGY GROUP (ATG)



Accelerate with ATG:





ADVANCED TECHNOLOGY GROUP (ATG)



Accelerate with ATG Webinar: Storage Virtualize 8.7.1 Technical Update

Byron Grossnickle

IBM Storage Virtualize/FlashSystem SME

byrongro@us.ibm.com

Chris Bulmer

Storage Virtualize High Availability and Replication Architecture



Meet the Speakers



Chris Bulmer is the software architect for IBM Storage Virtualize with responsibility for the high availability and replication features and he is based at IBM Hursley in the UK. Chris has worked on developing IBM Storage products since 2013 and has led the development of high-availability and replication features for IBM Storage FlashSystem, IBM SAN Volume Controller, and IBM Storage Virtualize since 2017. His recent projects include policy-based replication, policy-based high availability and Flash Grid.



Byron Grossnickle is an IBM Storage Technical Specialist concentrating on Storage Virtualize software. This includes FlashSystem, SVC, and Storage Virtualize for Public Cloud. Byron has been with IBM 19 years exclusively in storage. Prior to working for IBM, Byron spent 6 years engineering storage in the Telcom Industry. Prior to that he worked 8 years in healthcare IT. Byron lives in the Kansas City area and is available to travel to customer engagements.

IBM Storage Virtualize 8.7.1 Agenda



- Release Schedule
- PBR and PBHA Review
- PBHA + DR (PBR)
- Support Statements

Release Schedule

- RFA Announce – 8.7.1 – August 20, 2024
- eGA – 8.7.1 – September 20, 2024
- pGA – There is no hardware associated with this release

8.7.1 is a Non-LTS Release. Non-LTS releases are tested identically to LTS releases. Non-LTS releases will not get any patches and will not be maintained long term. Those patches will be available in the next Non-LTS or LTS release.

– [Release FAQ](#)



Replays

- [Accelerate with ATG Storage webinar Series](#)

VGS/Async DR 8.5.1/8.5.2

Policy-based Replication provides a simple to use, high performance DR solution. Volume Group Snapshots provides simple/secure point in time copies

–vVol 2.0 replication enabled (8.6.2)

PBHA 8.6.1

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

PBHA+DR 8.7.1 - New

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

HA replication is so light and 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 can testify the improvement in performance and stability provided.

- Large Bank in France

>4X

Higher
performance

32K


Replicated
volumes

4PiB

Replicated
capacity

IBM Storage Virtualize

PBR and PBHA Review



Why Use Policies?

Consistency

- Policies define how replication must be configured, which the system implements, to ensure that every volume is configured correctly.
-

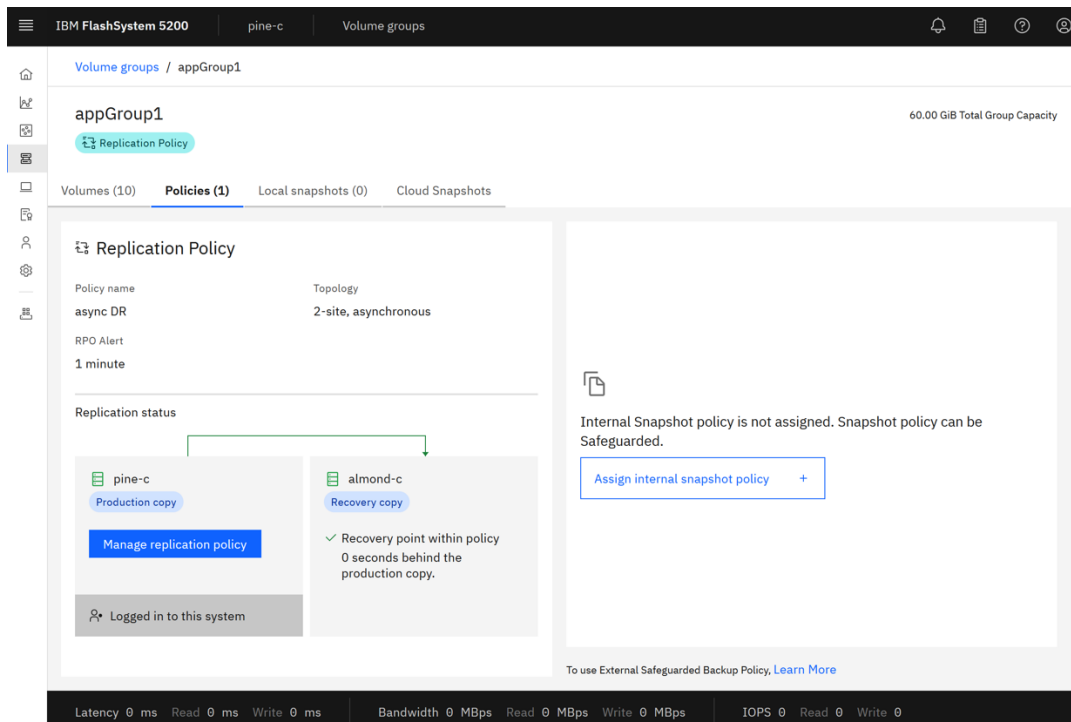
Simplicity

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

Automated

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

Policy-based replication simplifies the configuration, management and monitoring of volume groups

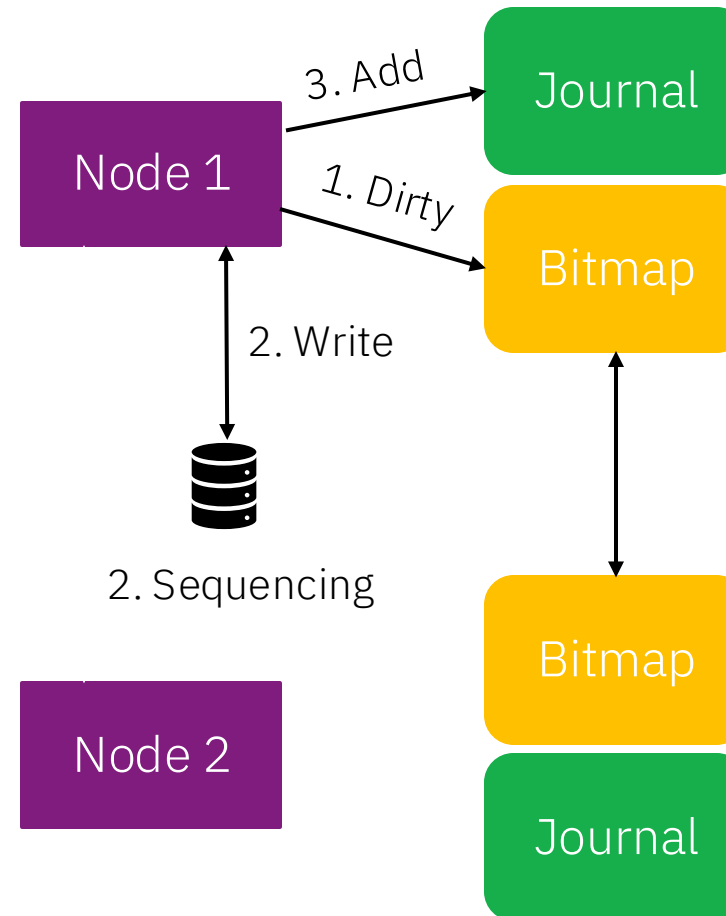


As seen in 8.7.0

- Create and populate
- Create a volume group and add volumes.
- Policy-based
- Add a DR replication policy to the volume group and all the volumes will automatically be configured for replication, with all remote provisioning handled automatically.
- Disaster recovery
- In a disaster, access can be enabled to the recovery copy in a few clicks.
- Flexible
- Create different replication policies for different RPOs
- Replication can be restarted just as easily, in either direction.

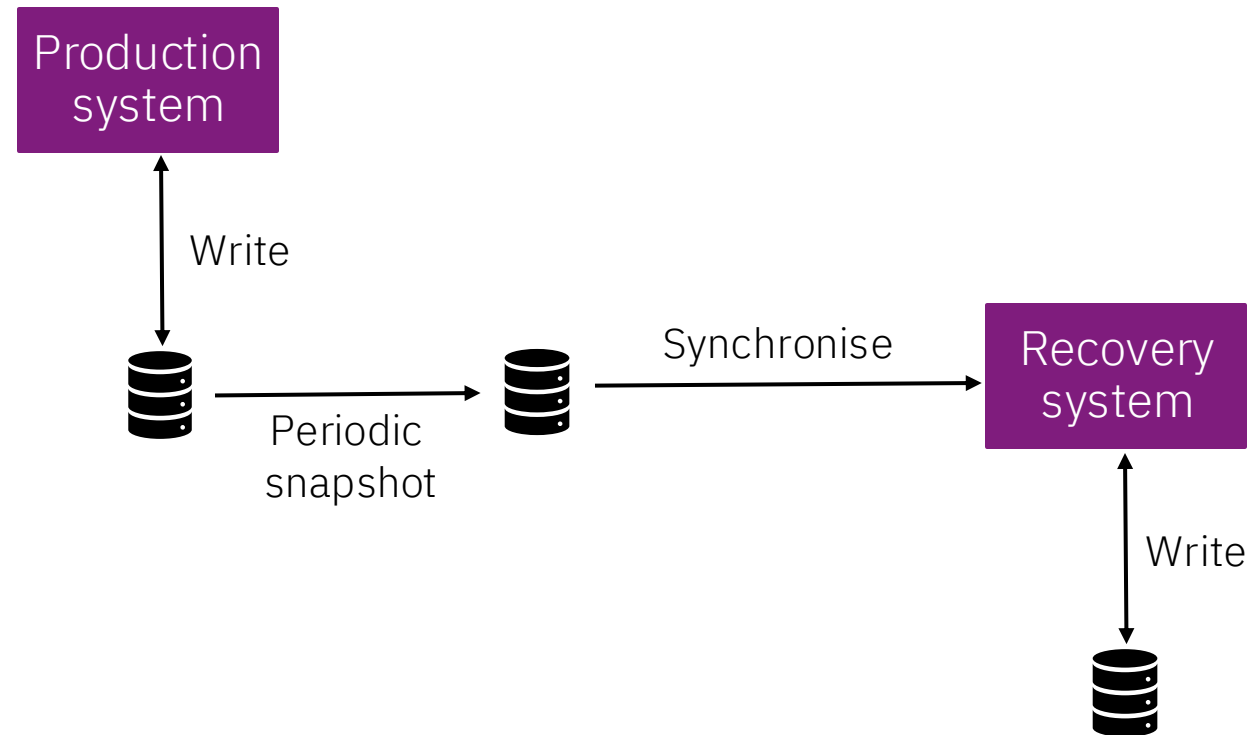
Journaling mode

- Writes are buffered in an in-memory journal that allows for many seconds of writes to absorb replication performance problems without host impact
- Journal sizes vary from 1 GiB to 32 GiB per node
- Using a large journal helps us ride through temporary problems with the link
- Volume groups will always try to replicate in this mode as it provides a better recovery point



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
- Reduces the bandwidth required for replication
- Volume groups may automatically be switched to this mode if journaling mode cannot be sustained



Performance and scalability significantly reduce TCO

>4x

- More than four times the host throughput compared to Global Mirror, significantly reducing the number of I/O groups required.

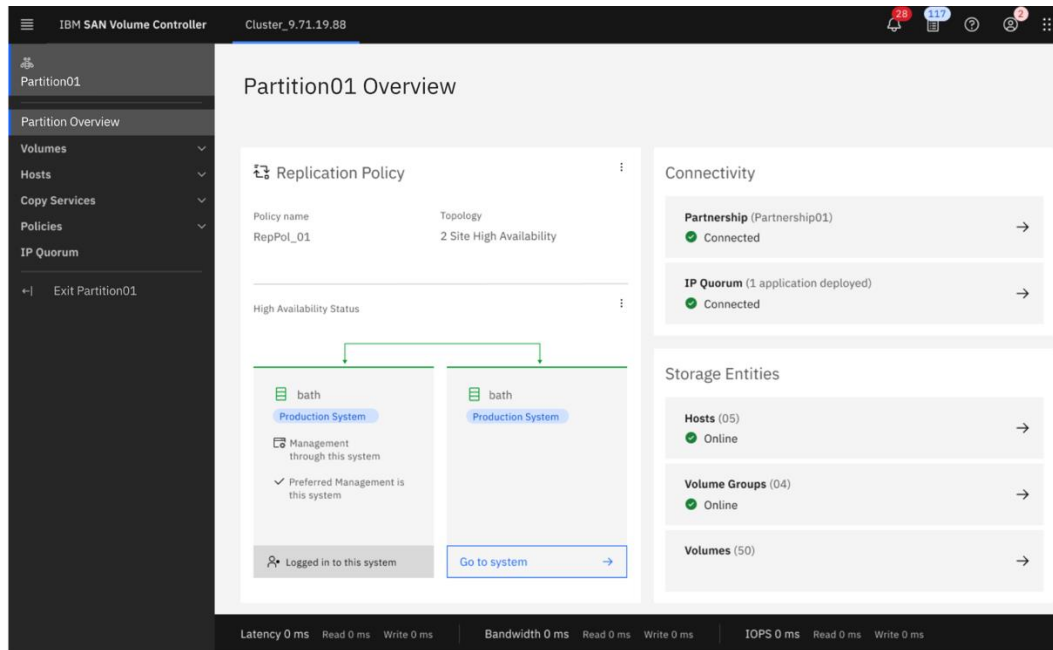
32k

- Replicate up to 32,500 host-mappable volumes from a single system, versus 2,500 for Global Mirror with Change Volumes (and 2000 for HyperSwap) – reducing the number of systems needed

4PiB

- Replicate up to 4PiB per I/O group, reducing the I/O groups needed

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



As seen in 8.7.0

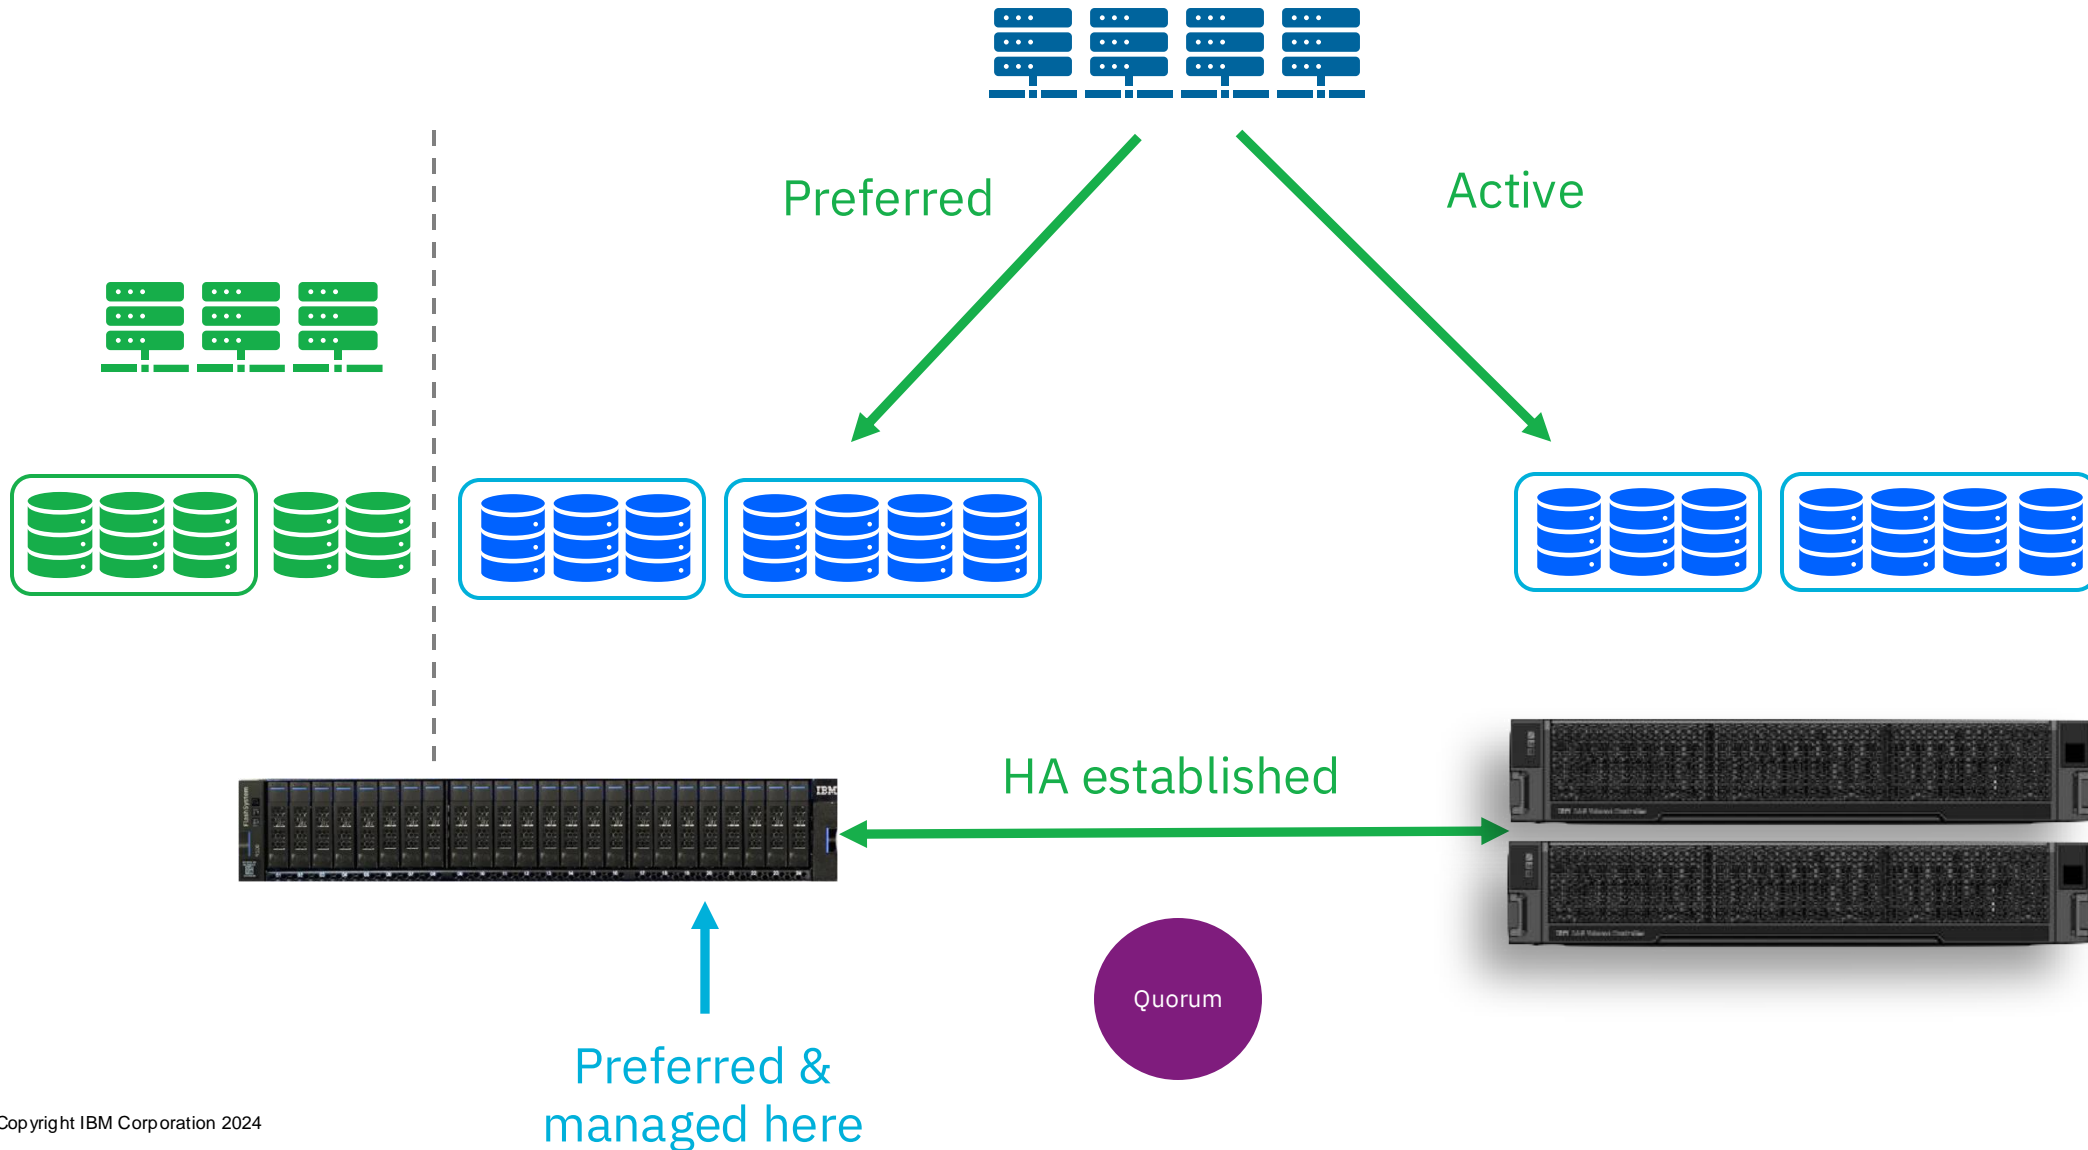
- Create and populate
- Create a **storage partition** in a few clicks, and create or add any number of hosts, volume groups and volumes.

- 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.

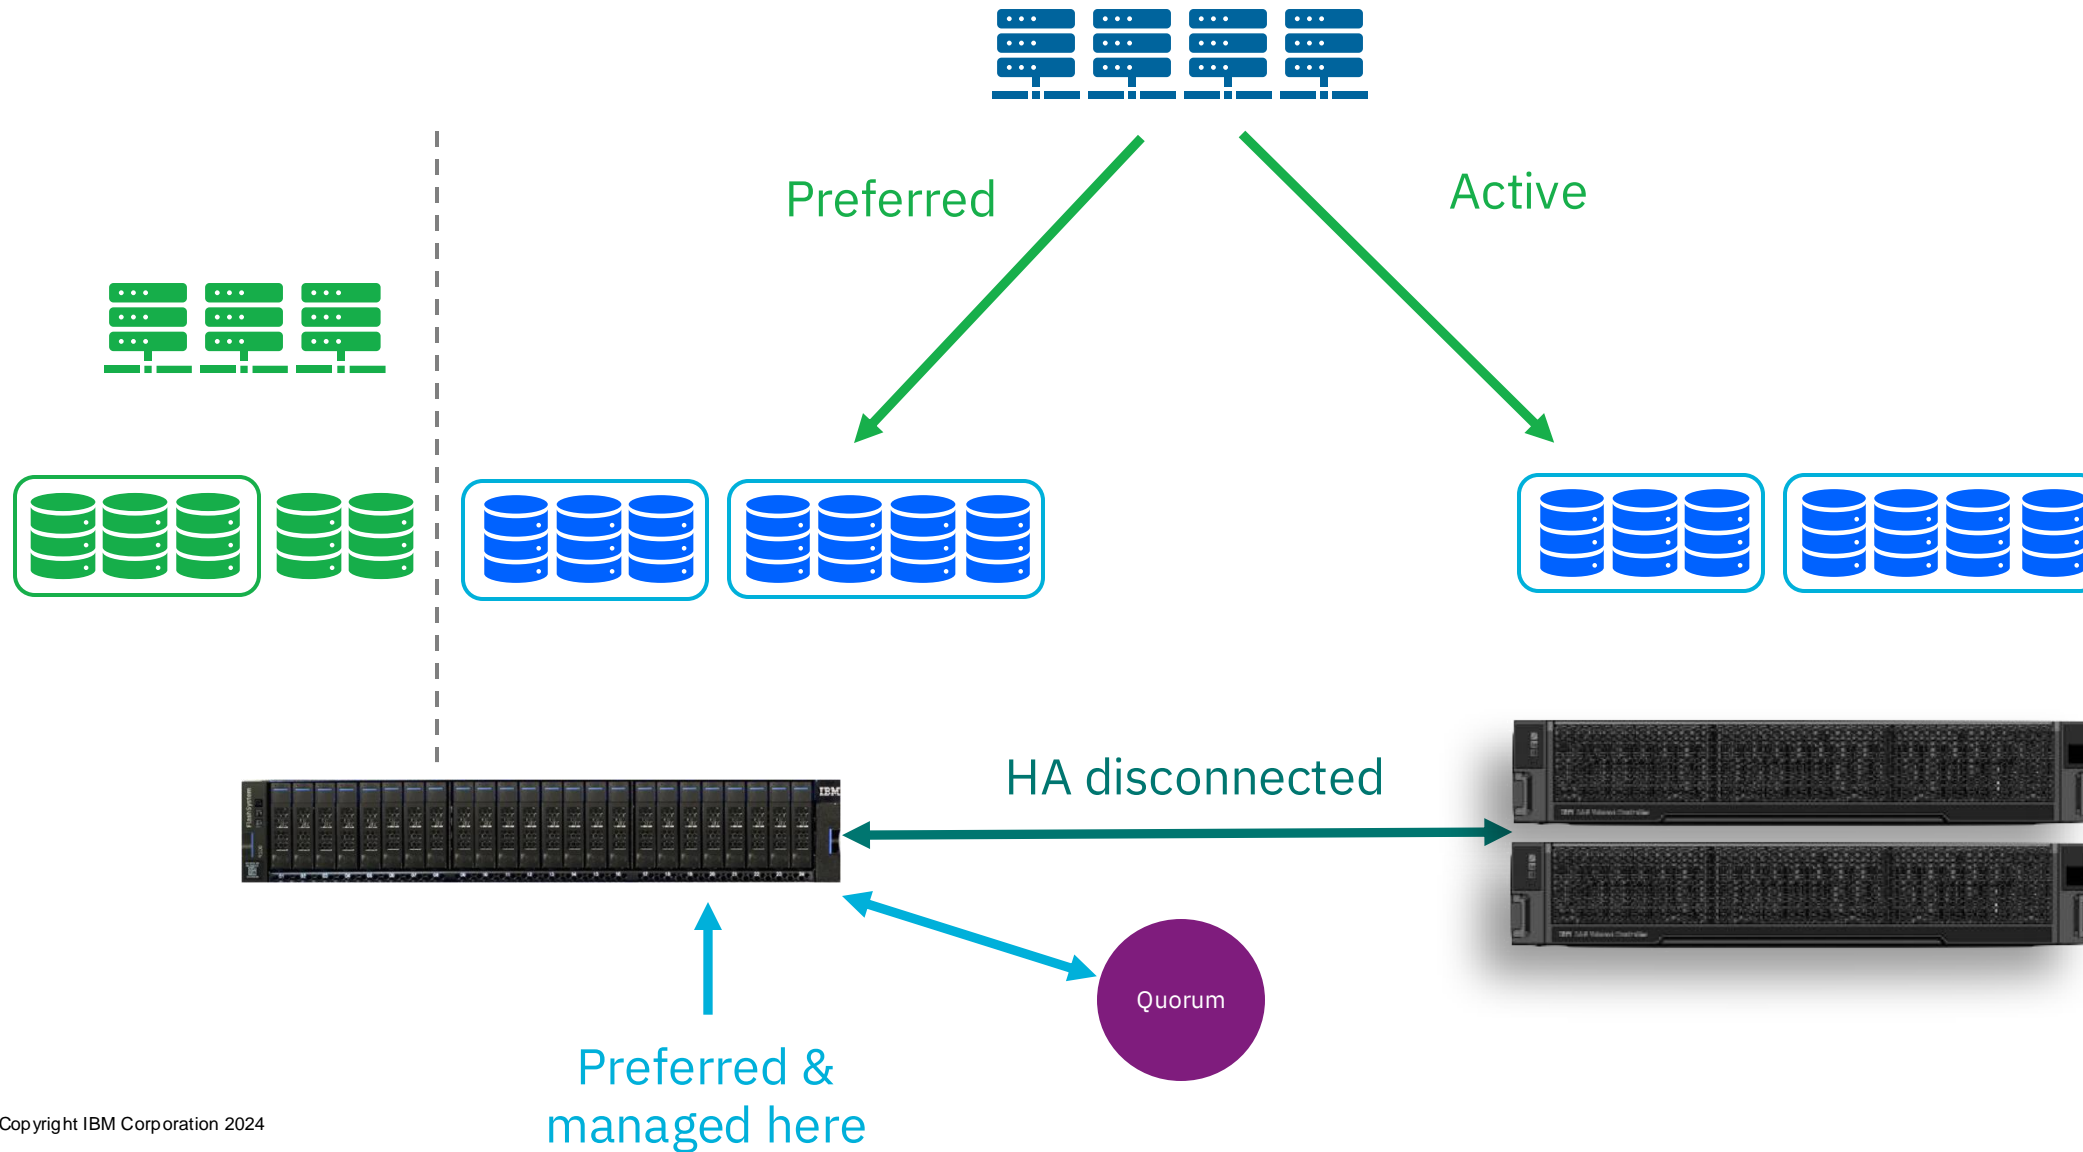
- Highly available
- In a disaster, the storage partition automatically manages its availability ensuring that applications remain accessible, with per-partition control over which system is preferred in the event of a loss-of-connectivity.

- 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 HA Storage Partitions – Normal Running

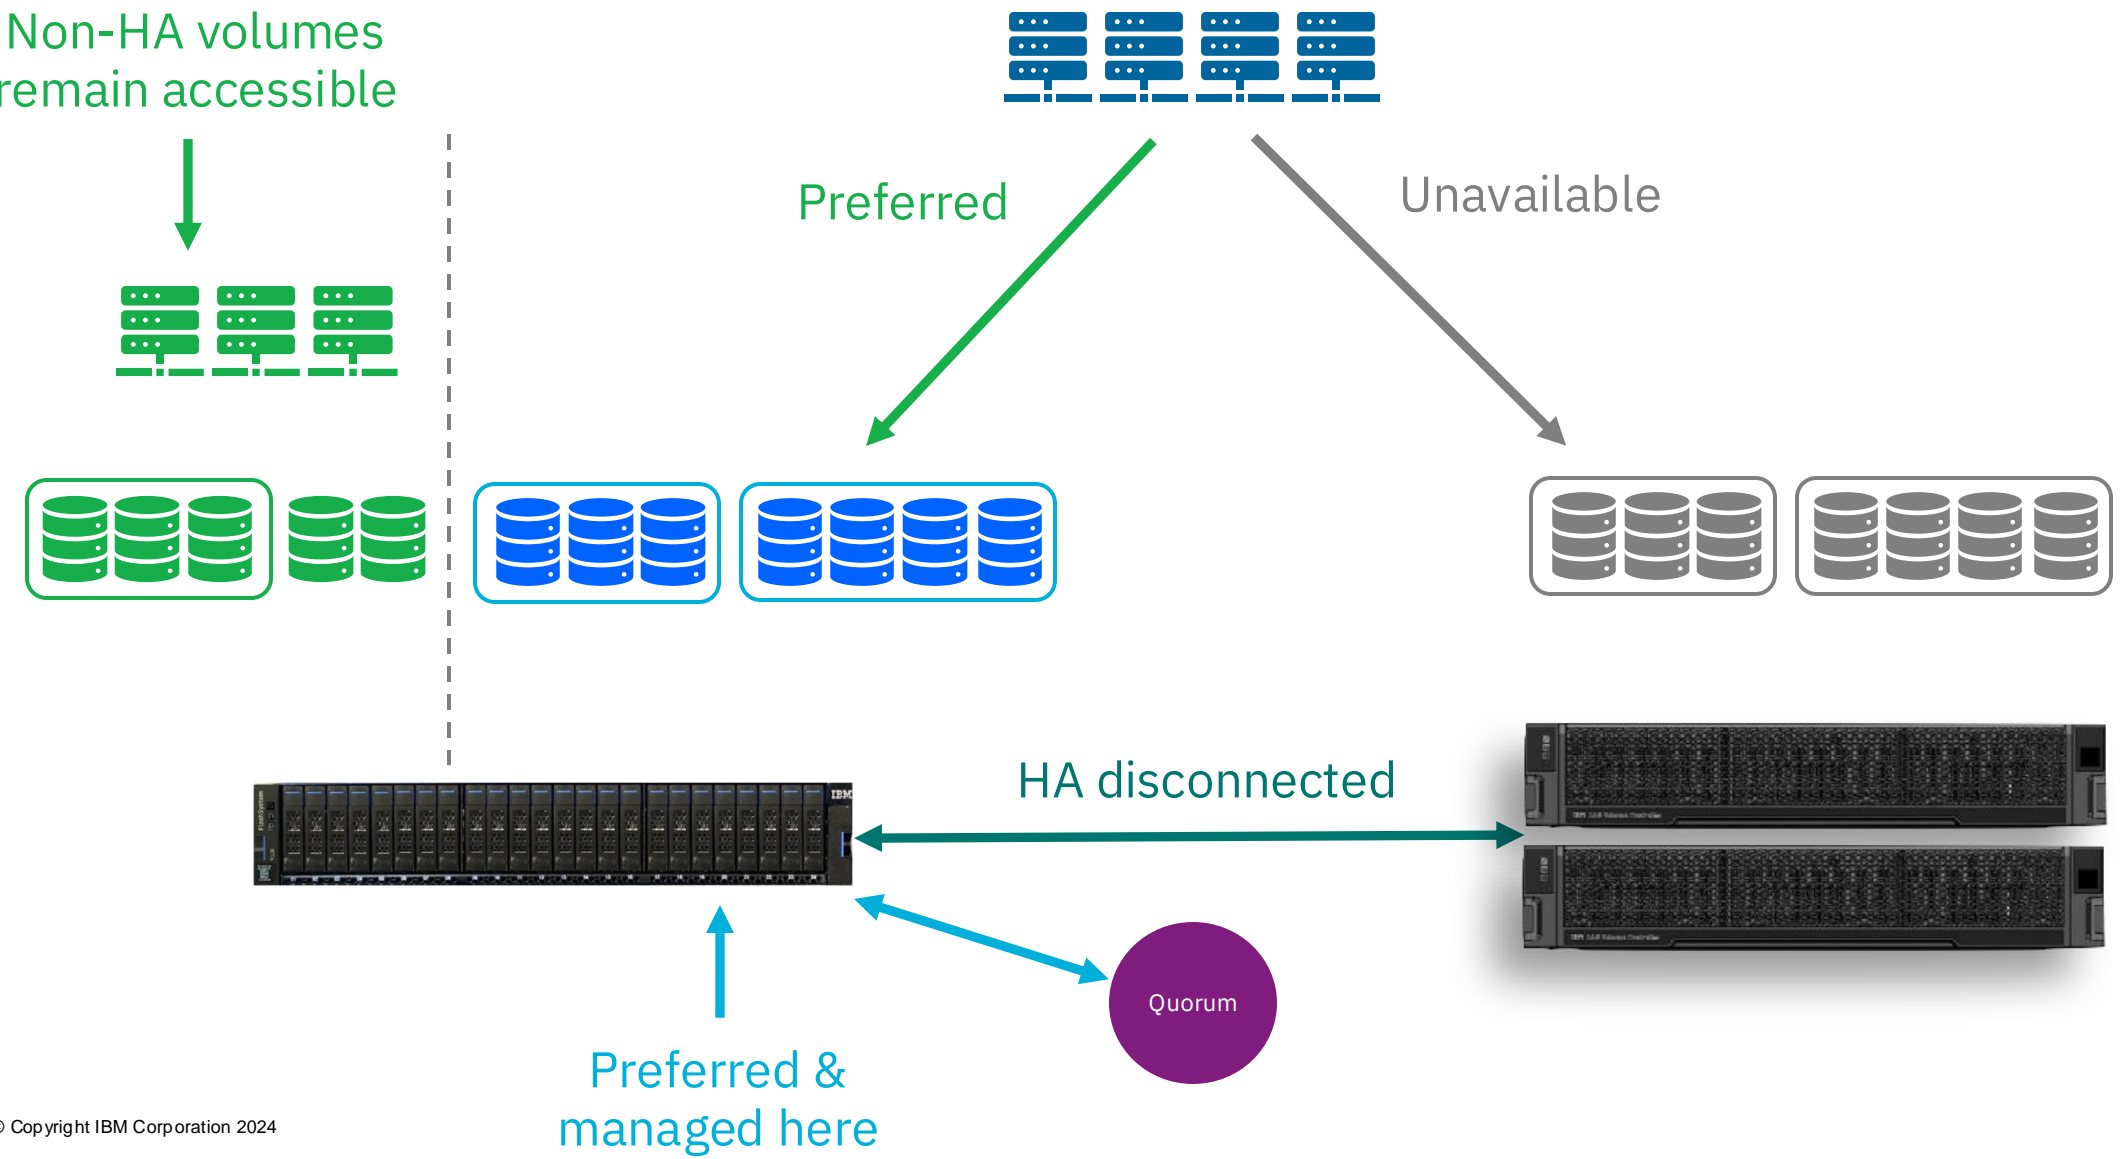


HA Storage Partitions - Disconnecting



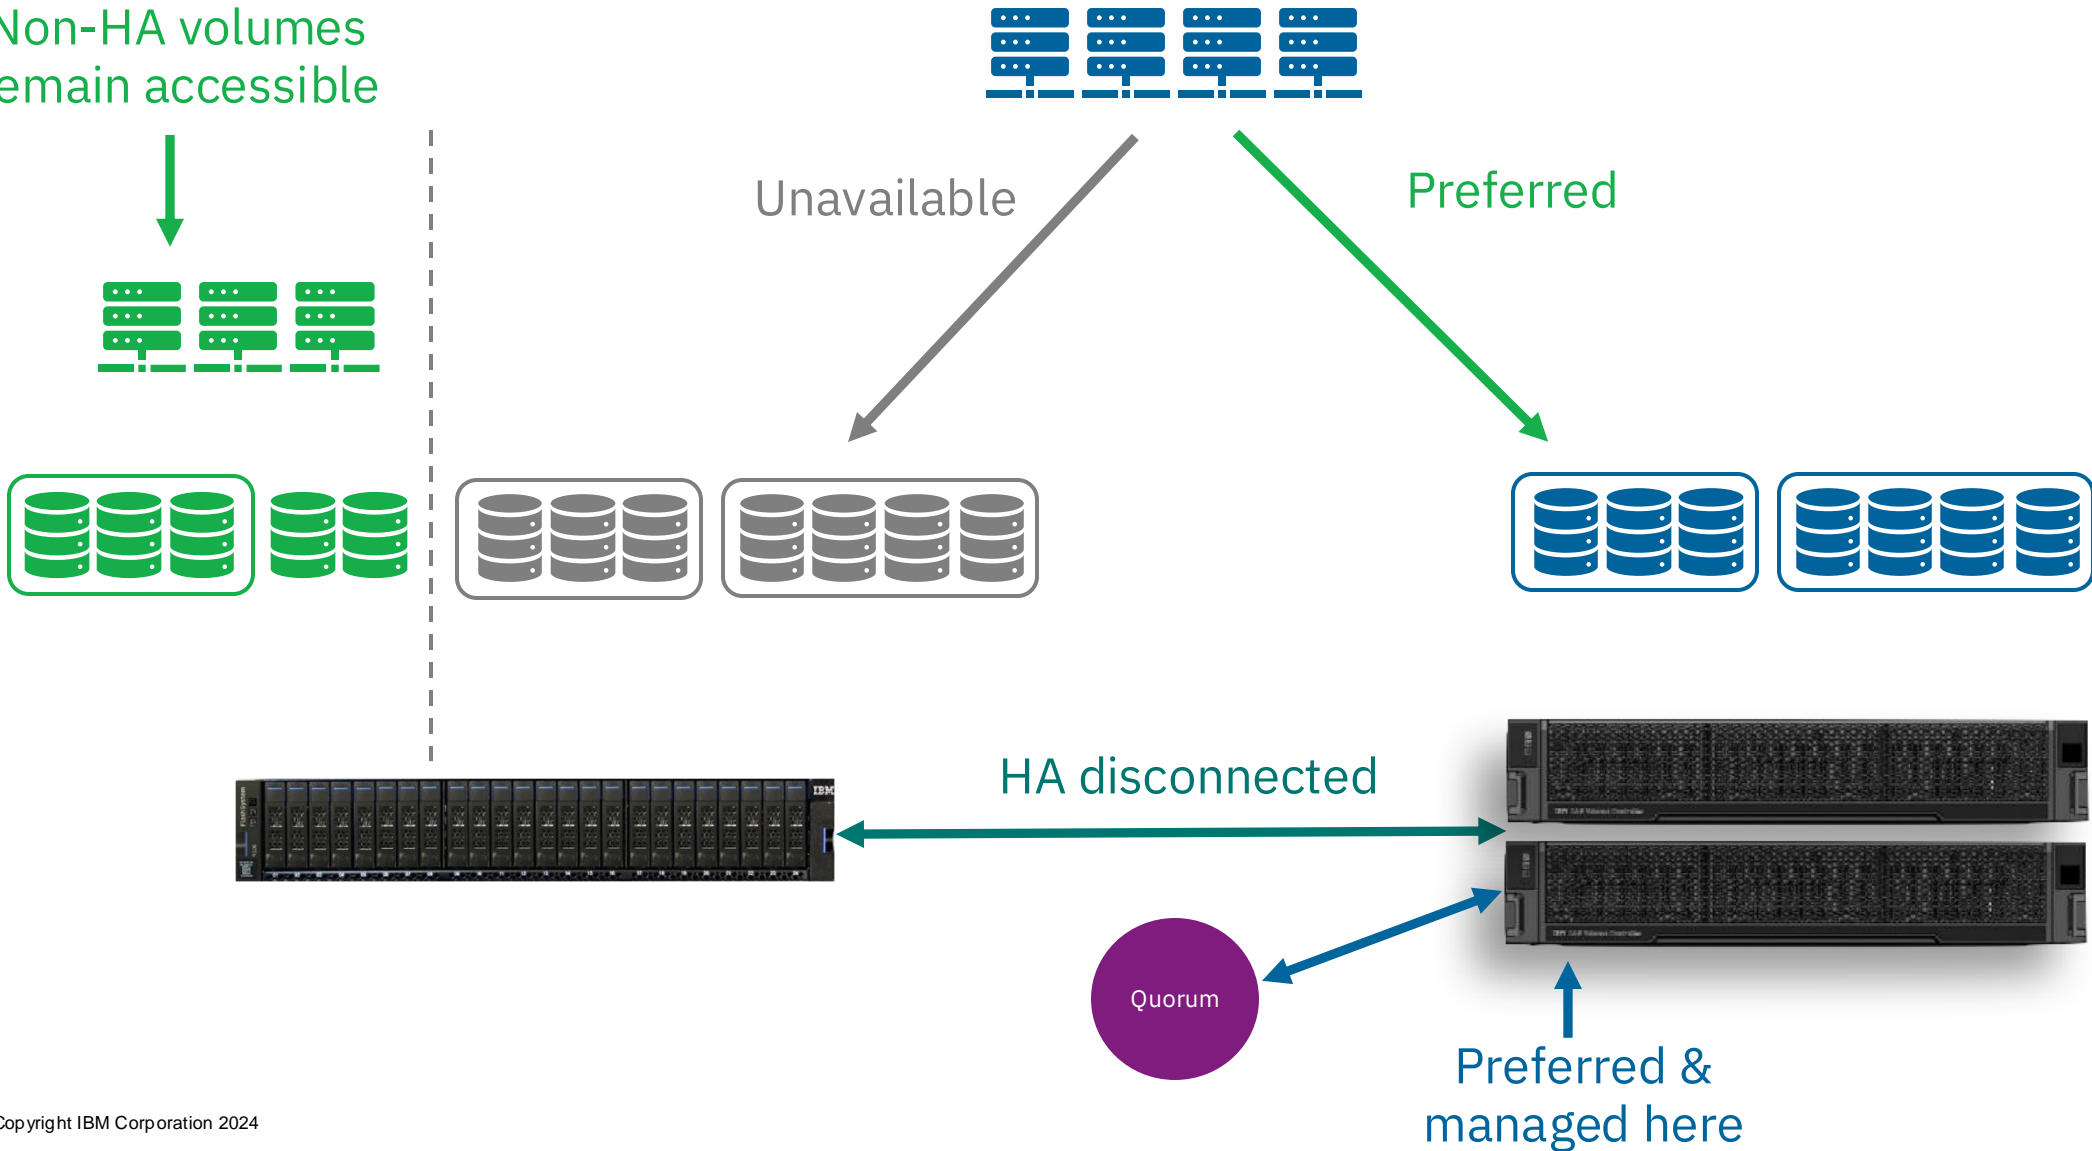
HA Storage Partitions - Disconnected

Non-HA volumes remain accessible

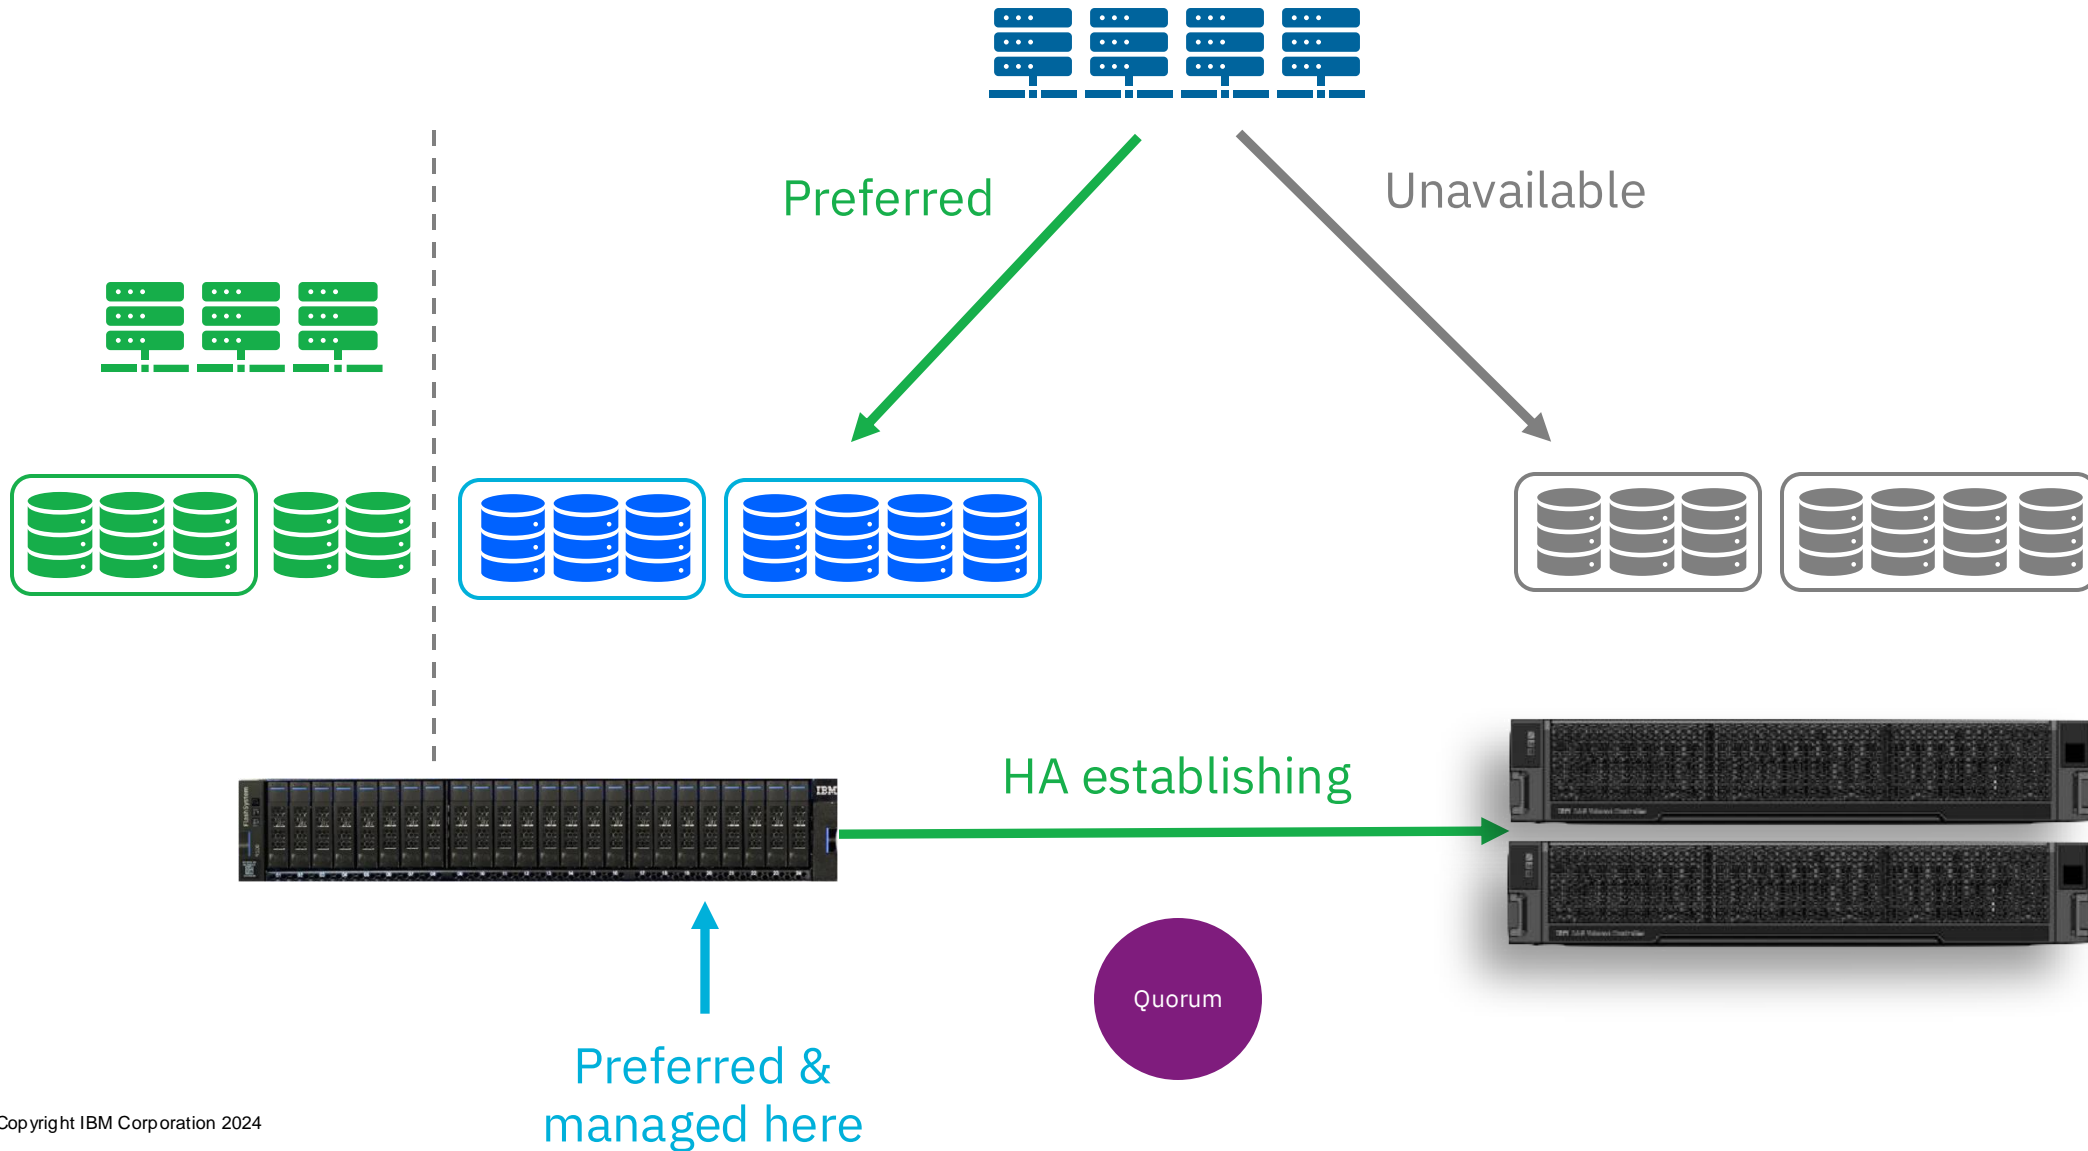


HA Storage Partitions – Independent Races

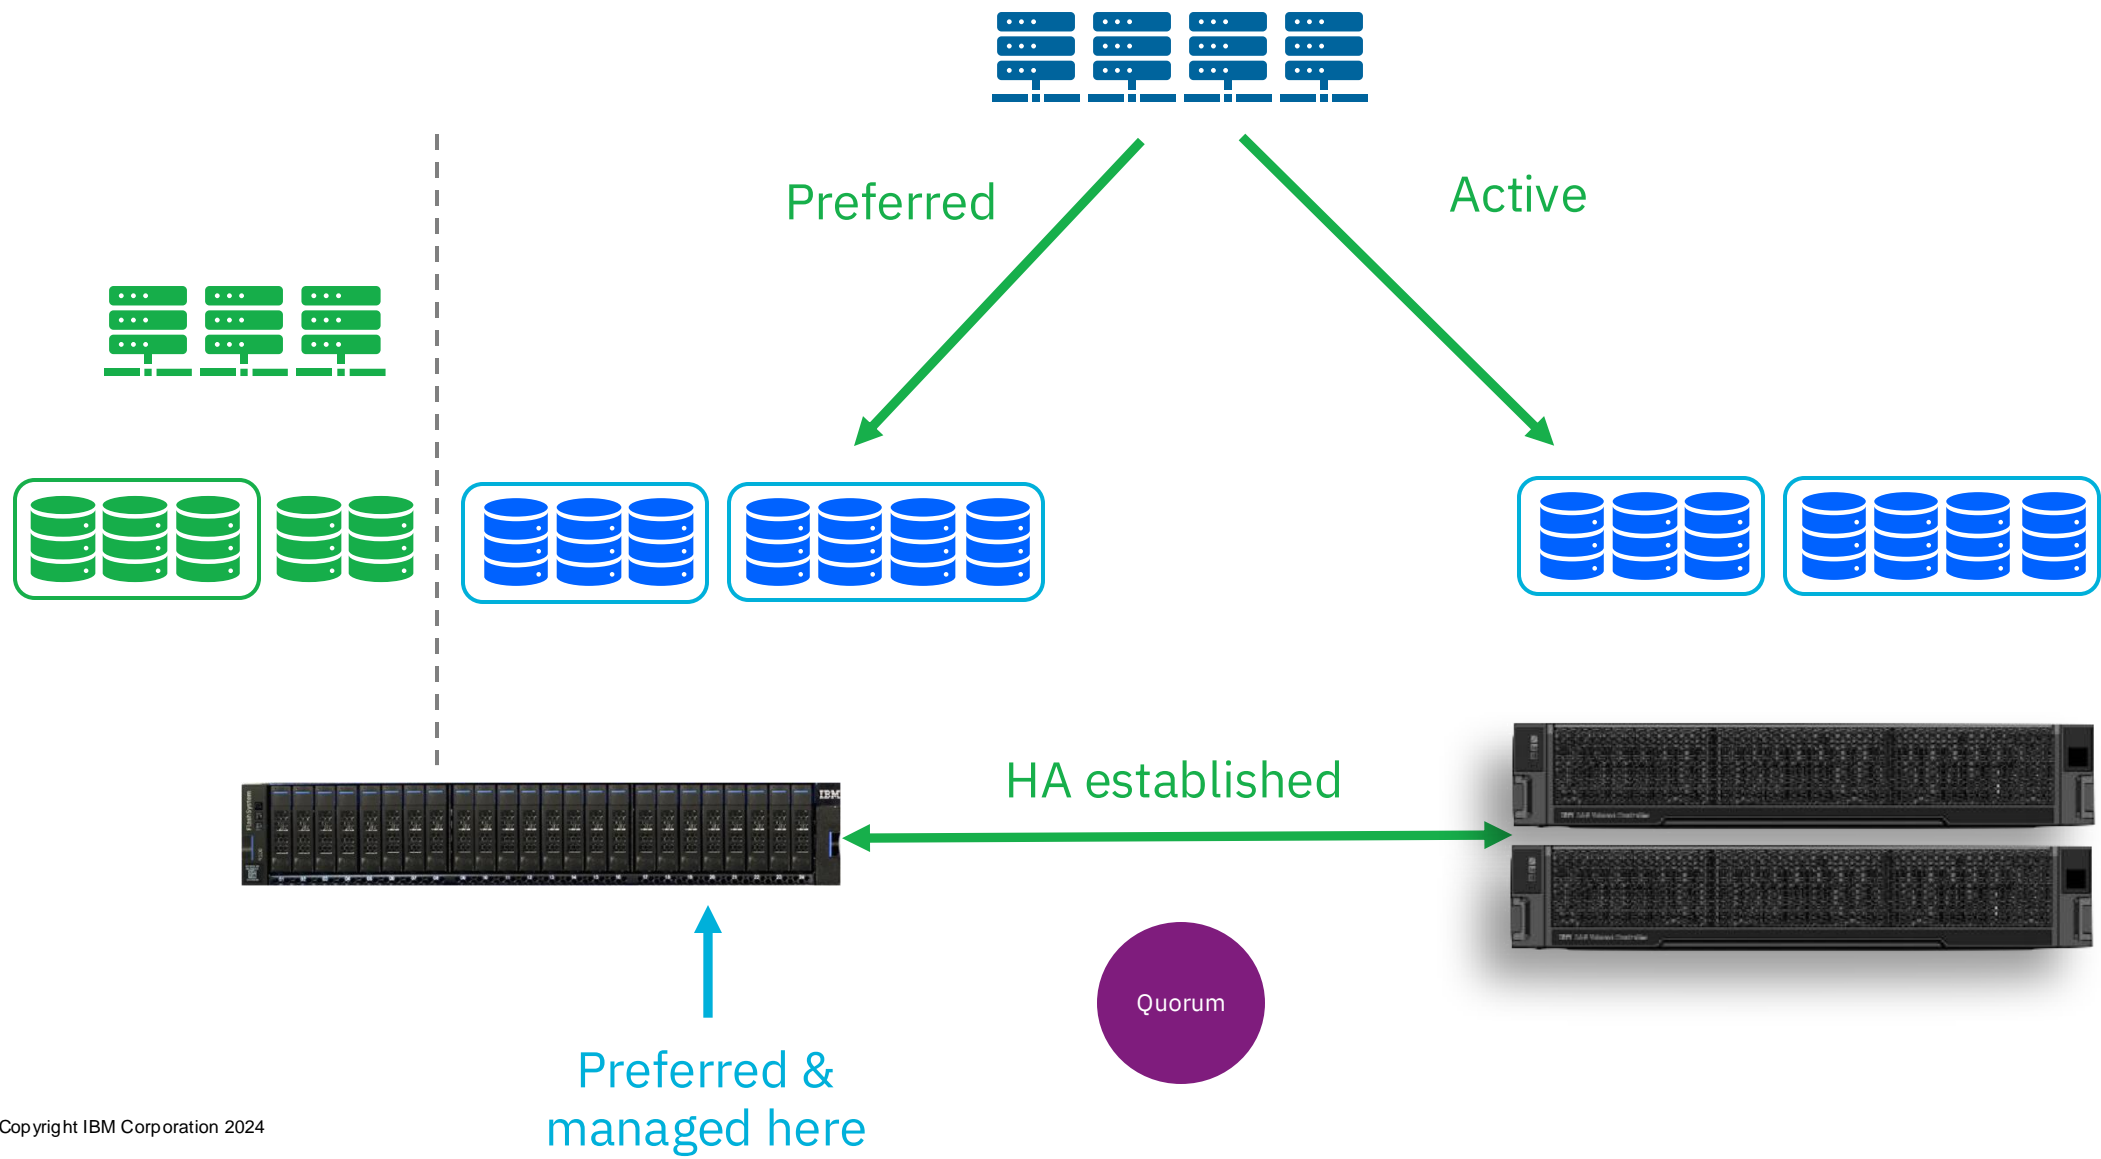
Non-HA volumes remain accessible



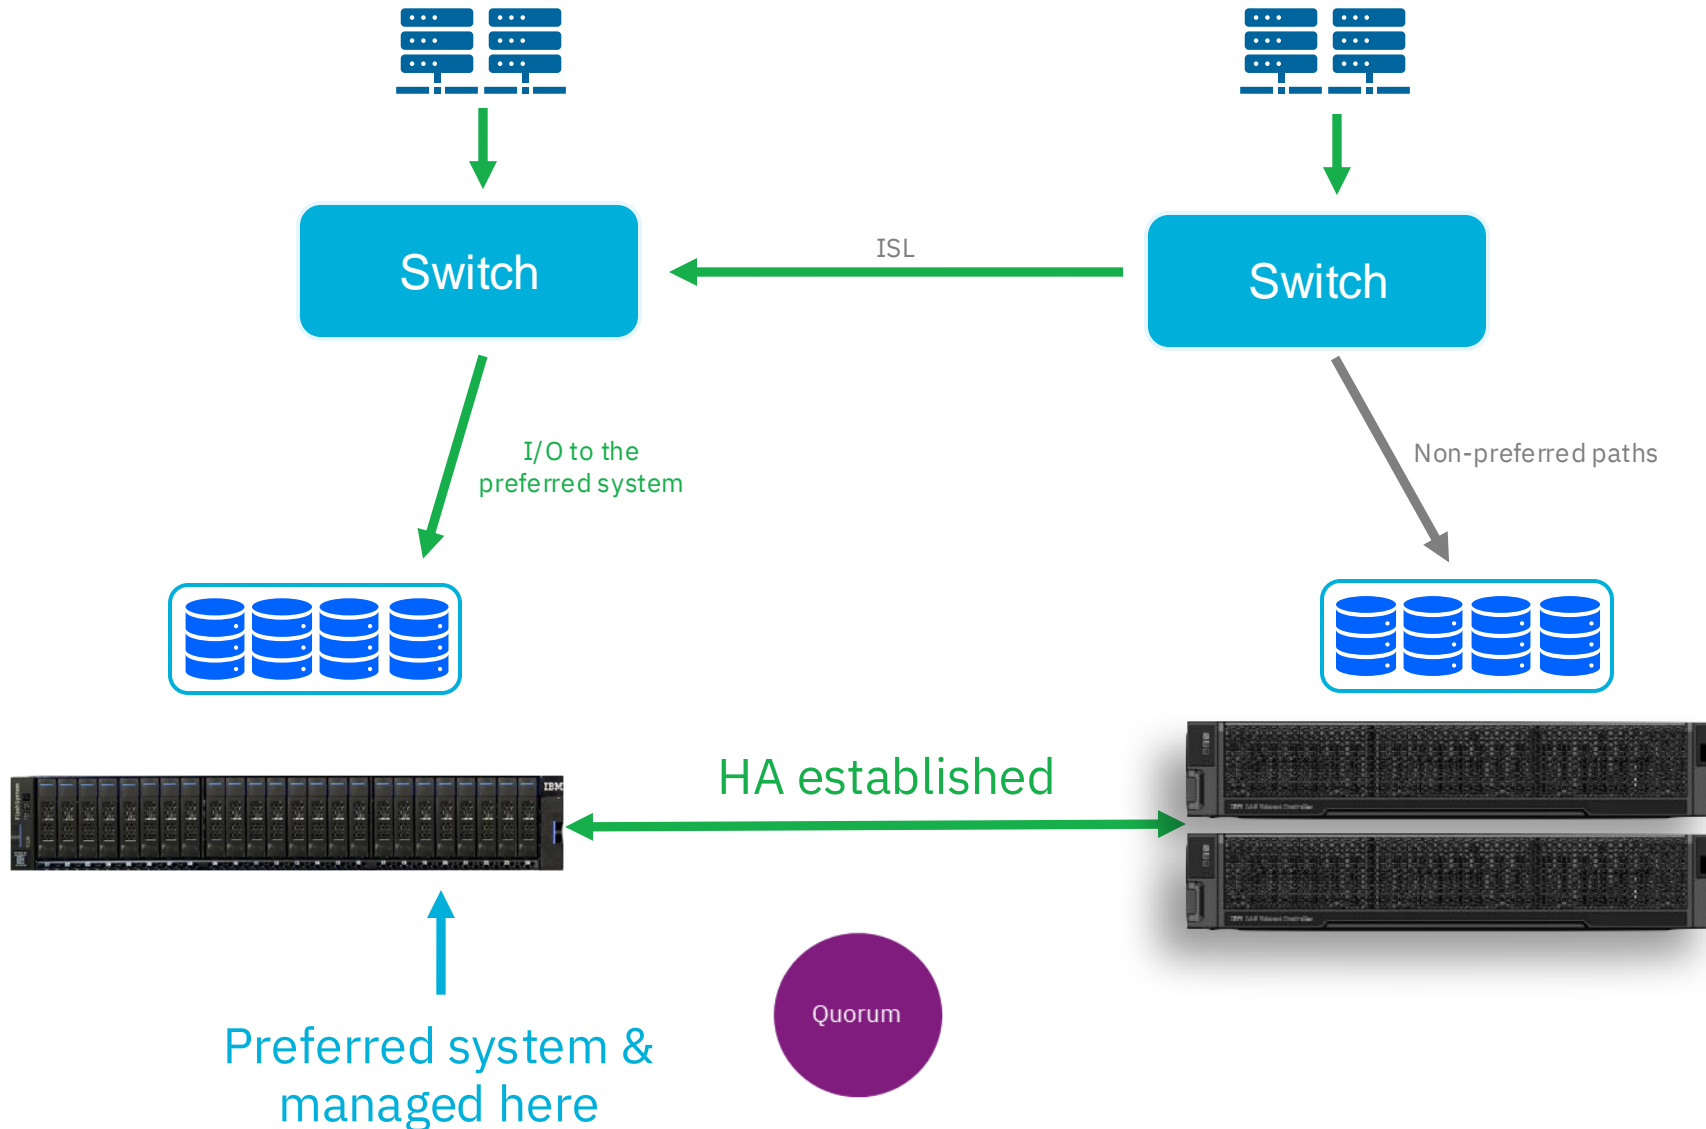
HA Storage Partitions – Re-establishing



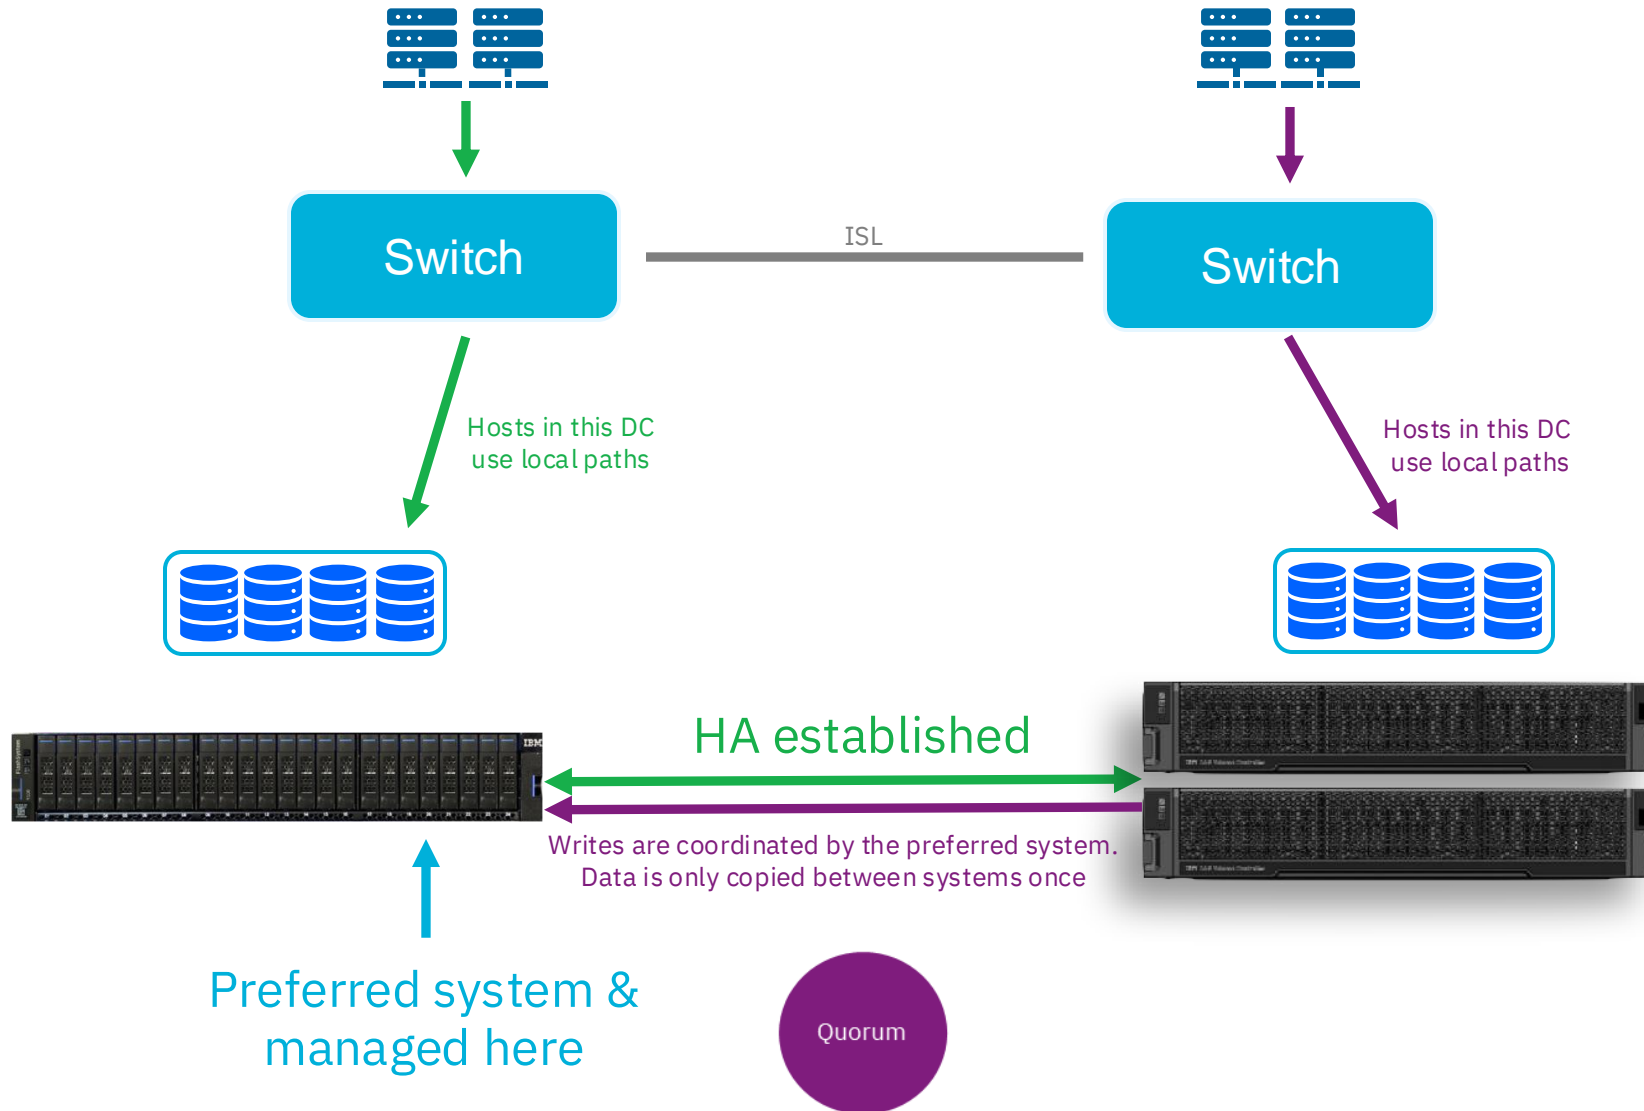
HA Storage Partitions – Normal Running



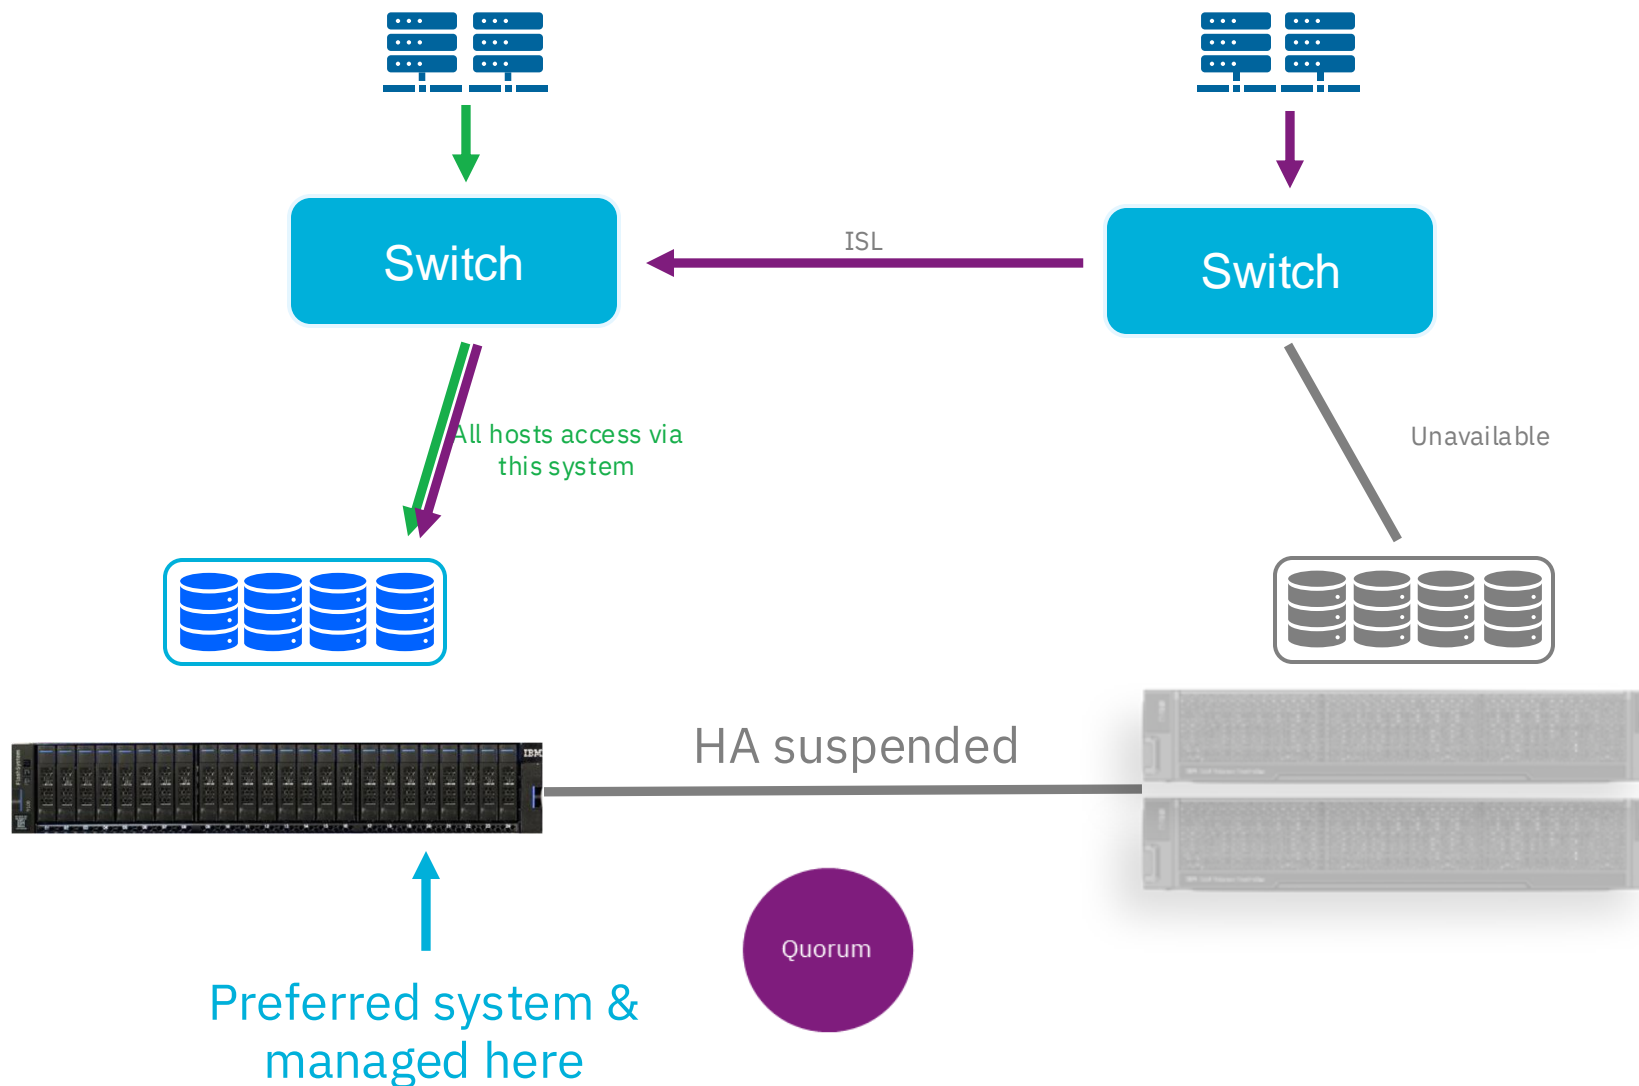
IO Flow - HA Without Host Locations



IO Flow - HA With Host Locations



IO Flow - HA With Host Locations – Storage Failure



Managing The Host's Location

IBM FlashSystem 7200 | pandabear | Storage partition: myPartition - Hosts

myPartition

Storage partition overview

Volumes

Hosts

Copy services

Policies

IP quorum

Exit myPartition

Hosts

+ Add Host | Actions | Download

Default | Contains | Filter

Name	Status	Host Type	# of Ports	Host Mappings	Host Cluster ID	Host Cluster Name	Location	Protocol Type
pandabear-h0	Online	Generic	1	Yes			pandabear	SCSI
pandabear-h1	Online	Generic	1	Yes				SCSI

Edit and rename host

Name

pandabear-h1

Assign location

Host location

sunbear

Advanced ^

Cancel Save

Active/Active HA


- Volumes are active/active when HA is established.
- Writes to the non-preferred system have an additional round trip, but the data only crosses the ISL once.
- If a location is set for a host, the volumes report preferred access to the system in the same location.
- Allows for the hosts located at a specific physical location to read and write to the locally connected system, reducing ISL traffic and latency.
- Allows for both 'uniform' and 'non-uniform' setups.
- Support for the following Fibre Channel SCSI operating systems:
 - Red Hat Enterprise Linux 7 and later
 - VMware ESXi 7 and later
 - AIX 7.2 and later (**new in 8.7.0**)
- Support planned for additional operating systems and functionality using FC SCSI:
 - *Microsoft Windows, IBM i and others*
 - *SCSI persistent reservations*

IBM Storage Virtualize

Policy-based High Availability with Async DR

PBHA + DR

8.7.1



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:** Lists various configuration options such as 'Volume groups', 'Volumes', 'Volume mappings', 'Hosts', 'Replication policies', and 'Snapshot policies'. An 'Exit myPartition' button is also present.
- myPartition Overview:** Shows the partition name and a 'Manage partition' dropdown menu.
- Replication overview:** Contains a message stating 'Replication is not configured for this partition.' and a prominent blue button labeled 'Configure replication +'. Above the message is an icon representing a storage component.
- Storage components:** A list of components associated with the partition:
 - Hosts (1) - Online (with a green checkmark and a right-pointing arrow)
 - Volumes (10) (with a right-pointing arrow)
 - Volume groups (4) (with a right-pointing arrow)
- Connectivity:** Contains a message 'There is no partnership assigned to this partition.' and a storage component icon.
- Performance Metrics:** A footer bar displays various 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

The screenshot shows the IBM FlashSystem 5200 management interface. The top navigation bar includes the system name 'IBM FlashSystem 5200', the user 'pine-c', and the current storage partition 'Storage partition: myPartition'. A sidebar on the left lists navigation options: myPart, Volume, Hosts, Replica, Snapsh, and Exit.

The main content area displays a configuration wizard titled 'Configure policy-based replication for myPartition'. The wizard's progress bar shows three steps: 'Select replication type' (completed), 'Configure replication' (current step), and 'Summary Review' (pending). The 'Configure replication' step includes the instruction: 'Choose one or both the replication types that are available to configure on your storage partition.'

The configuration diagram shows two nodes: 'pine-c' (labeled 'Logged in here') and 'pecan-c'. A line connects them with a 'High availability' label. To the right of the 'pecan-c' node, a list of completion status items is shown:

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


At the bottom of the wizard, there are three buttons: 'Cancel', 'Skip', and 'Continue'.

HA Storage Partition


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 myPartition, 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', 'pine-c', and 'Storage partition: myPartition'. A 'Manage partition' dropdown menu is located in the top right.
- Replication overview:** A diagram showing two hosts, 'pine-c' (logged in here) and 'pecan-c', connected by a line labeled 'HA established'.
- Storage components:** A list of components with their status and a 'Manage' arrow:
 - 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 for Latency, Bandwidth, and IOPS, all showing 0 values.

HA Storage Partition With Async DR



HA Storage Partition With Async DR



Starting With High Availability

The screenshot displays the IBM FlashSystem 5200 management interface for a 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', 'pine-c', and 'Storage partition: myPartition'. It includes notification, list, help, and user icons.
- myPartition Overview:** Features a 'Replication overview' section with a diagram showing two hosts, 'pine-c' (logged in) and 'pecan-c', connected by a green line labeled 'HA established'. A 'Manage partition' dropdown menu is in the top right.
- Storage components:** A list of components with expandable arrows:
 - Hosts (1): Online
 - Volumes (10)
 - Volume groups (4)
- Connectivity:** A list of connectivity status:
 - HA partnership (pecan-c): Configured
 - IP quorum: Connected
- Performance Metrics (Bottom):** A row of metrics: Latency 0 ms, Read 0 ms, Write 0 ms, Bandwidth 0 MBps, Read 0 MBps, Write 0 MBps, IOPS 0, Read 0, Write 0.

Starting With High Availability

Click to Add
Disaster
Recovery

The screenshot displays the IBM FlashSystem 5200 management interface. The top navigation bar shows 'IBM FlashSystem 5200', 'pine-c', and 'Storage partition: myPartition'. The left sidebar lists navigation options: myPartition, Volume groups, Volumes, Volume mappings, Hosts, Replication policies, Snapshot policies, IP quorum, and Partnerships. The main content area is titled 'myPartition' and features a 'Management partition' dropdown menu. The 'Replication overview' section shows a diagram with two hosts, 'pine-c' (logged in) and 'pecan-c', connected by a line labeled 'HA established'. A context menu is open over the diagram, with 'Add disaster recovery' highlighted. The 'Storage components' section lists: Hosts (1) Online, Volumes (10), and Volume groups (4). The 'Connectivity' section shows: HA partnership (pecan-c) Configured and IP quorum Connected. The 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.

Starting With High Availability

Click to add disaster recovery

Step-by-step 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

- 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

Starting With High Availability

Click to add disaster recovery


Step-by-step configuration

Ready to add DR partition policies

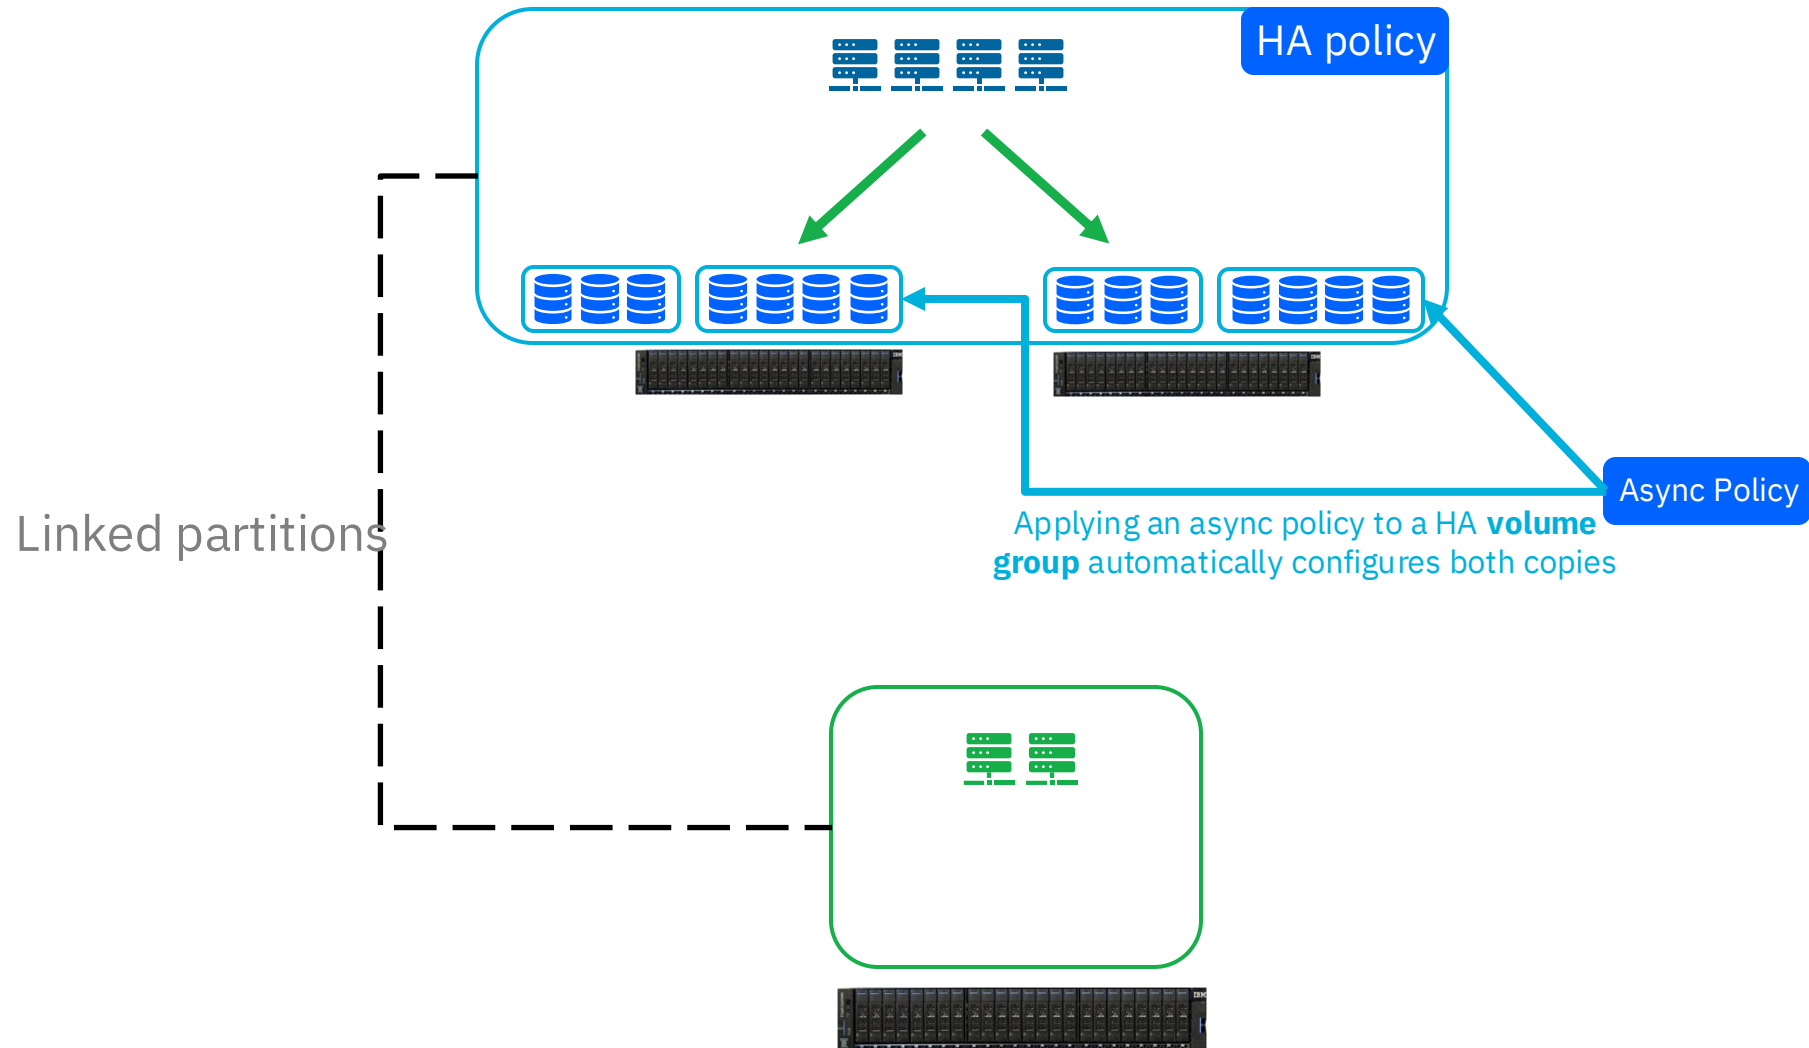
The screenshot displays the IBM FlashSystem 5200 management console for a partition named 'myPartition'. The interface is divided into several sections:

- Navigation Panel (Left):** Lists various configuration options such as Volume groups, Volumes, Volume mappings, Hosts, Replication policies, Snapshot policies, IP quorum, and Partnerships.
- Replication overview (Center):** Shows a diagram with two nodes, 'pine-c' and 'pecan-c', connected by a line labeled 'HA established'. 'pine-c' is marked as 'Logged in here'.
- Storage components (Right):** Lists 'Hosts (1)' as 'Online', 'Volumes (10)', and 'Volume groups (4)'. A warning icon indicates 'No disaster recovery policies configured'.
- Connectivity (Right):** Shows 'HA partnership (pecan-c)' as 'Configured', 'IP quorum' as 'Connected', and 'DR partnership (almond-c)' as 'Configured'.
- DR Configuration (Bottom Center):** A diagram shows 'myPartition_DR' connected to 'almond-c' via a link labeled 'DR unused'.
- Performance Metrics (Bottom):** Displays real-time statistics for Latency, Bandwidth, and IOPS, all currently at 0.

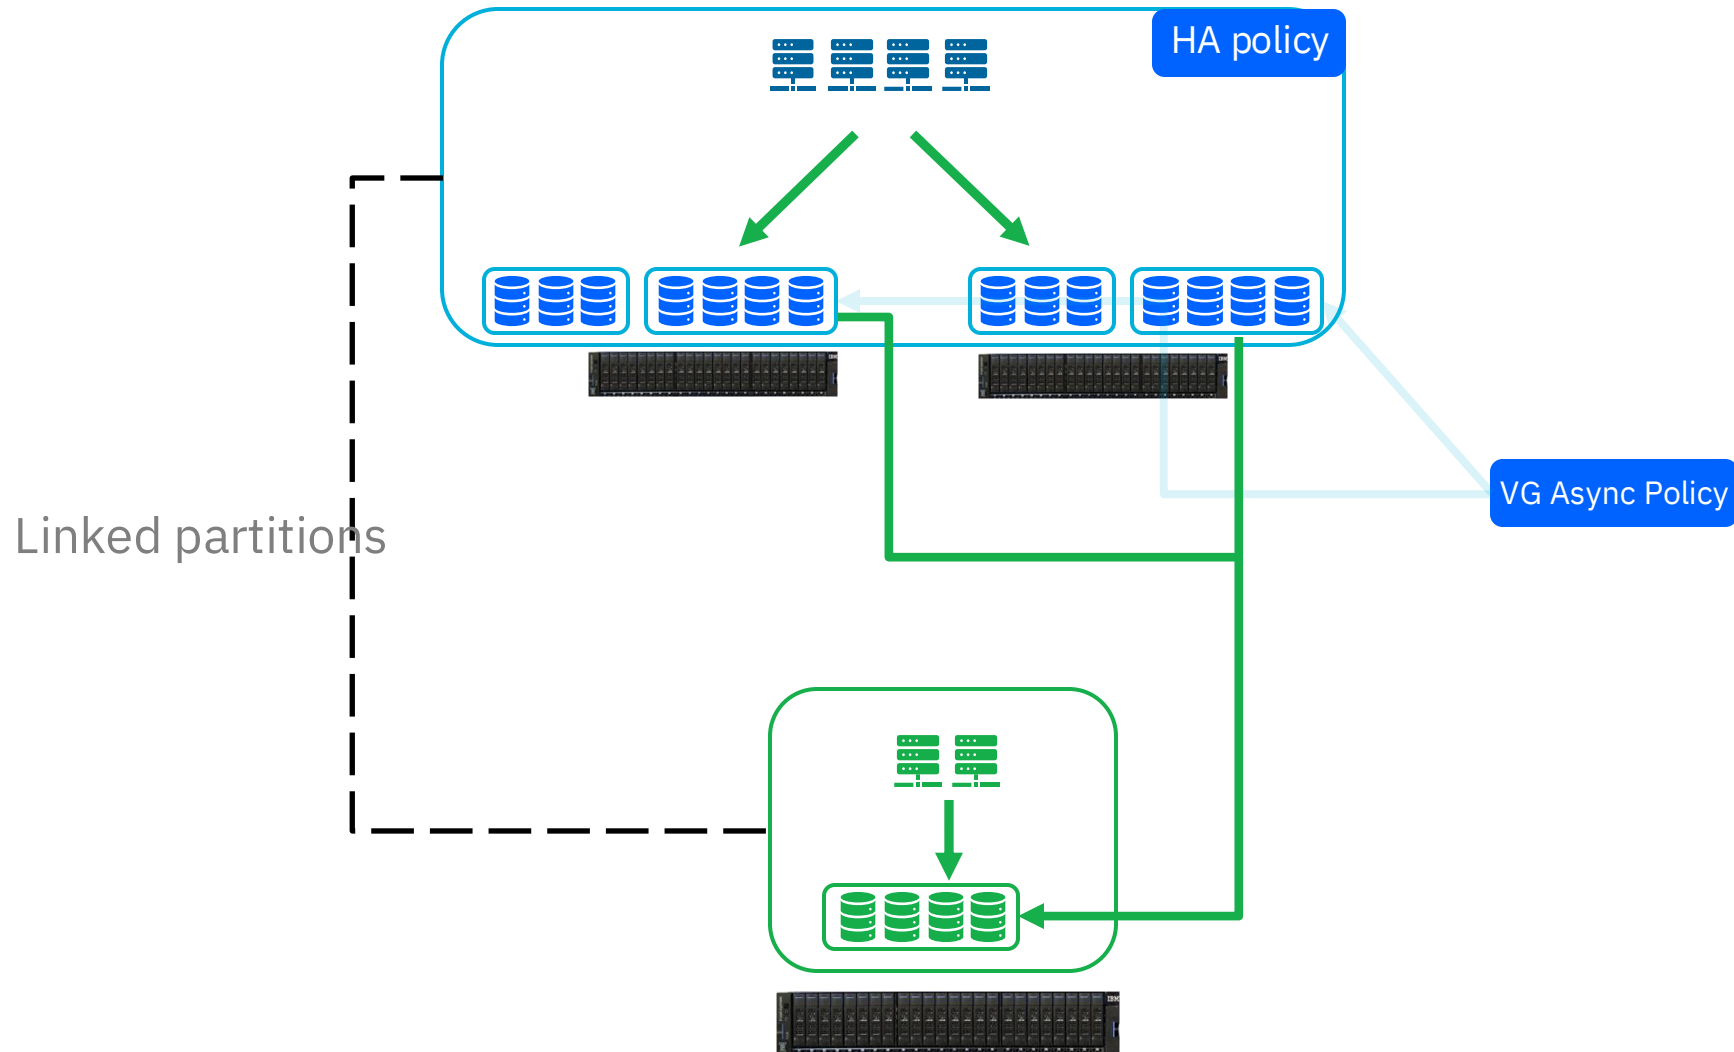
HA Storage Partition With Async DR



HA Storage Partition With Async DR



HA Storage Partition With Async DR



Add Disaster Recovery To A Volume Group

- Select a volume group

The screenshot shows the IBM FlashSystem 5200 management interface. The top navigation bar includes the system name 'IBM FlashSystem 5200', the user 'pine-c', and the current page 'Storage partition: myPartition - Volume groups'. A left-hand navigation menu lists various options: myPartition, Volume groups (selected), Volumes, Volume mappings, Hosts, Replication policies, Snapshot policies, and an 'Exit myPartition' button. The main content area is titled 'Volume Groups' and features a search bar, a 'Create Volume Group' button, and a table with the following data:

Name	Replication State	Volume Count	Replication Policy
appVG1	-	10	-
appVG2	-	5	-
appVG3	-	12	-
appVG4	-	5	-
demoVG	-	0	-
dev and test	-	0	-

At the bottom of the interface, there are 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. The footer shows 'Items per page: 10' and '1-6 of 6 items'.

Add Disaster Recovery To A Volume Group

- Click to add disaster recovery

The screenshot displays the IBM FlashSystem 5200 management interface. The top navigation bar shows the system name 'IBM FlashSystem 5200', the user 'pine-c', and the current page 'Storage partition: myPartition - Volume groups'. The left sidebar contains a navigation menu with options: myPartition, Volume groups (selected), Volumes, Volume mappings, Hosts, Replication policies, Snapshot policies, IP quorum, Partnerships, and Exit myPartition. The main content area shows the configuration for volume group 'appVG1' with a total capacity of 50.00 GiB. A 'Replication Policy' button is visible. Below this, tabs for 'Volumes (10)', 'Policies (0)', 'Local snapshots (0)', and 'Cloud Snapshots' are shown. The 'Policies (0)' tab is active, displaying a message: 'Disaster recovery is not configured for this volume group.' with a blue button labeled 'Add disaster recovery +'. To the right, another message states: 'Internal Snapshot policy is not assigned. Snapshot policy can be Safeguarded.' with a button labeled 'Assign internal snapshot policy +'. At the bottom of the main area, there is a link: 'To use External Safeguarded Backup Policy, [Learn More](#)'. 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, Read 0, Write 0.

Add Disaster Recovery To A Volume Group

- Select a replication policy

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

Volume groups / appVG1

Assign DR Policy

Select a replication policy that you want to assign to the selected volume group.

The volume group and associated volumes will be replicated from the local partition to the remote partition.

Name	RPO Alert
<input checked="" type="radio"/> Gold Tier DR	60
<input type="radio"/> Silver Tier DR	300
<input type="radio"/> Bronze Tier DR	3600

1-3 of 3 items

myPartition (Production) — HA established — pecan-c

Disaster recovery

myPartition_DR (almond-c)

Cancel | Assign

To use External Safeguarded Backup Policy, [Learn More](#)

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

Add Disaster Recovery To A Volume Group

Done!

(we did say it was easy)

The screenshot shows the IBM FlashSystem 5200 management interface. The top navigation bar includes the system name 'IBM FlashSystem 5200', the host 'pine-c', and the storage partition 'Storage partition: myPartition - Volume groups'. The left sidebar contains navigation options: myPartition, Volume groups (selected), Volumes, Volume mappings, Hosts, Replication policies, Snapshot policies, IP quorum, and Partnerships. The main content area displays 'Volume groups / appVG1' with a '50.00 GiB Total Group Capacity' indicator. Under 'appVG1', there is a 'Replication Policy' button and tabs for 'Volumes (10)', 'Policies (1)', 'Local snapshots (0)', and 'Cloud Snapshots'. The 'Policies (1)' tab is active, showing a 'Replication overview' diagram. The diagram illustrates a 'High availability' setup between two hosts: 'pine-c' (Production) and 'pecan-c' (Production). A 'Disaster recovery' arrow points from this setup to a 'Recovery' site 'myPartition_DR' on host 'almond-c'. A status message at the bottom of the diagram reads '✓ Recovery point within policy'. On the right side of the interface, a message states 'Internal Snapshot policy is not assigned. Snapshot policy can be Safeguarded.' with a button to 'Assign internal snapshot policy +'. At the bottom of the interface, performance metrics are shown: 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.

Flexibility to use different DR policies or no policy to match different requirements

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

myPartition

- Volume groups
- Volumes
- Volume mappings
- Hosts
- Replication policies
- Snapshot policies
- IP quorum
- Partnerships
- Exit myPartition

Volume Groups

Search table...

Create Volume Group +

Name	↑	Replication State	Volume Count	Replication Policy	
appVG1		✓ Running	10	Gold Tier DR	⋮
appVG2		✓ Running	5	Gold Tier DR	⋮
appVG3		✓ Running	12	Silver Tier DR	⋮
appVG4		✓ Running	5	Silver Tier DR	⋮
demoVG		✓ Running	0	Silver Tier DR	⋮
dev and test		-	14	-	⋮

Items per page: 10 | 1-6 of 6 items | 1 | 1 of 1 page


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

IBM Storage Virtualize

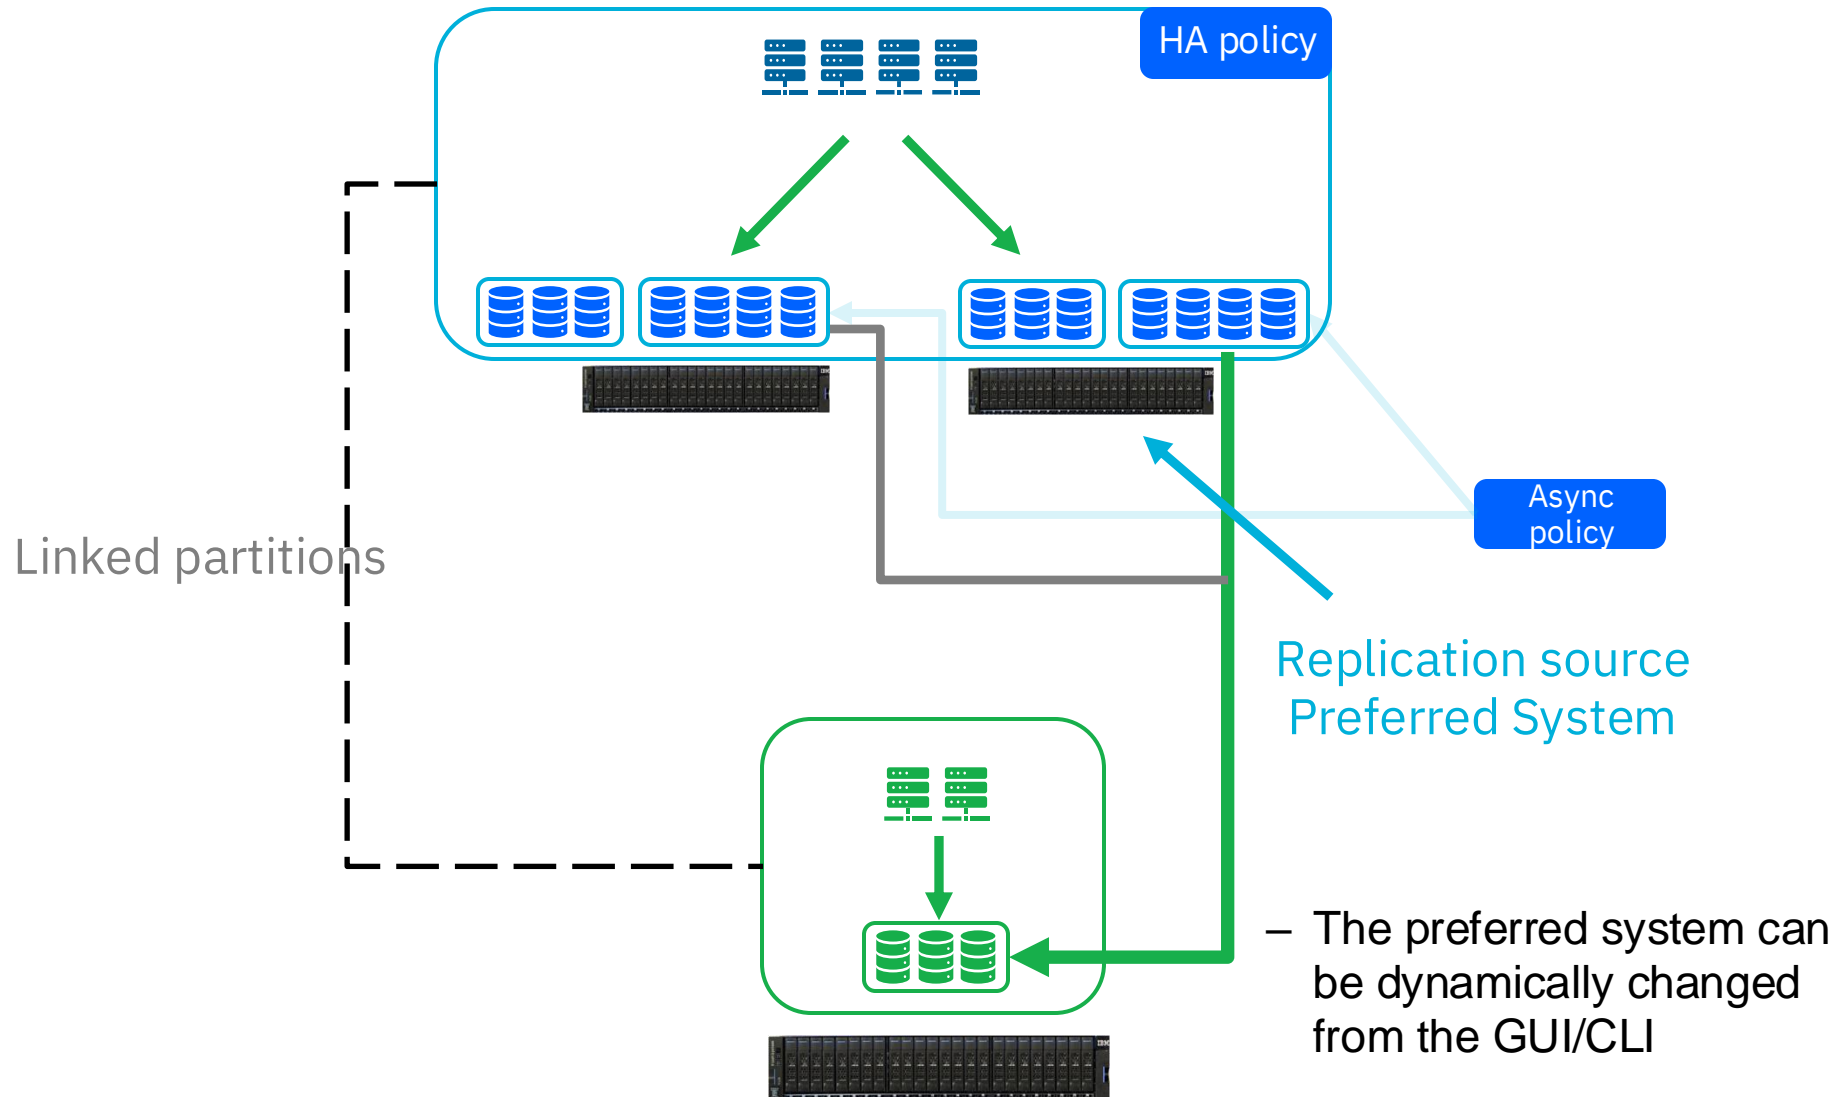
Policy-based High Availability with Async DR

Async Replication Handling

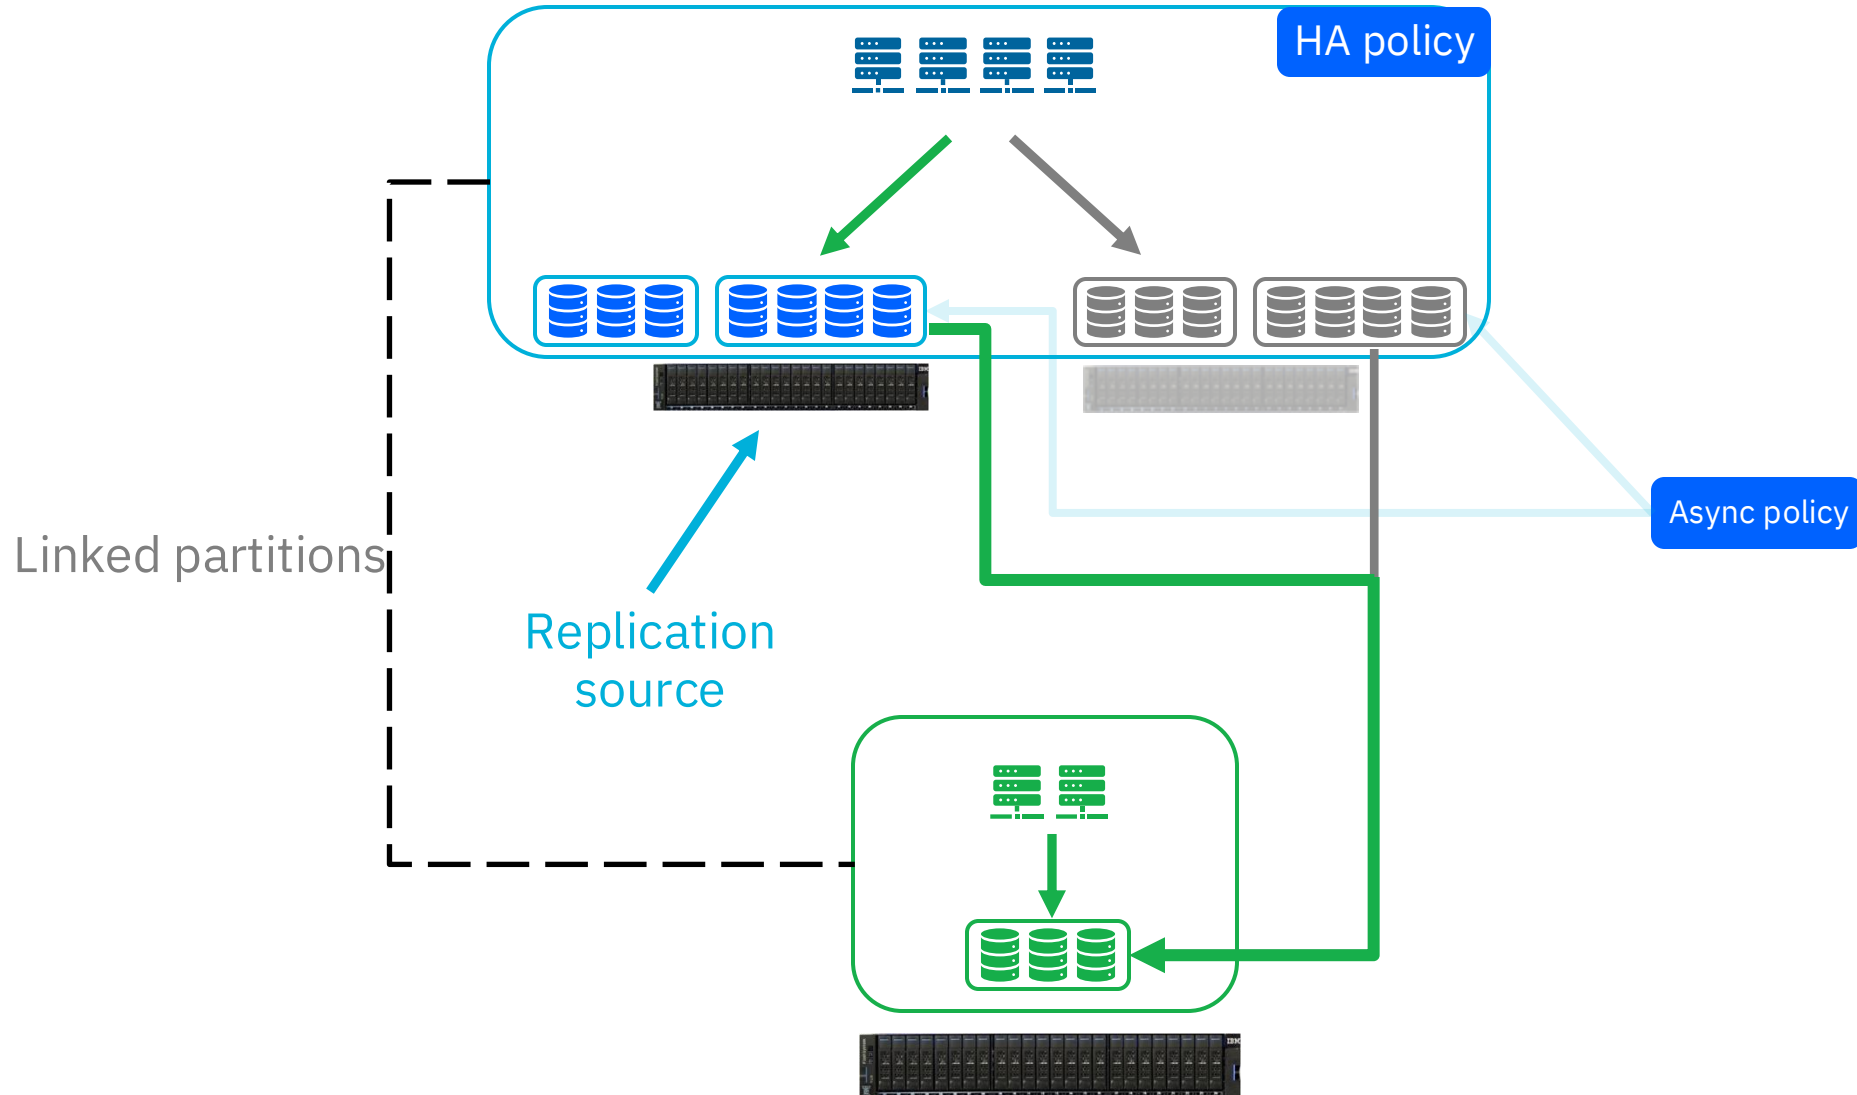
8.7.1



Normal Running Replicates From The Preferred System




Replication Continues Automatically Even If HA Has A Problem



IBM Storage Virtualize

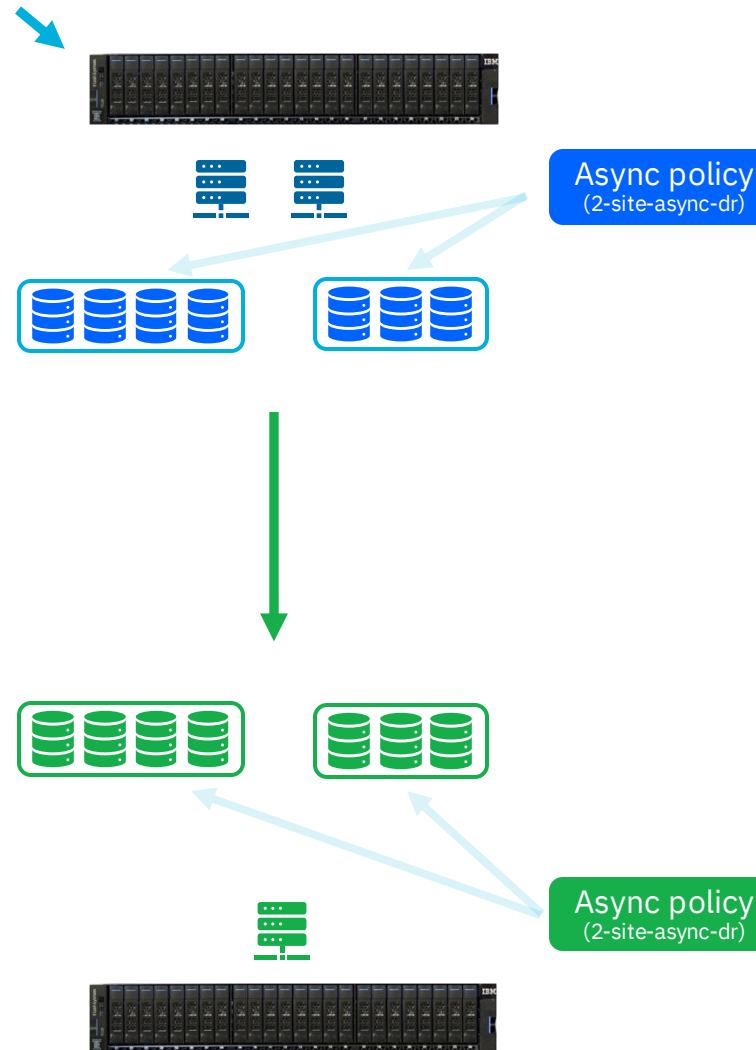
Adding High Availability to an Existing PBR Setup

8.7.1



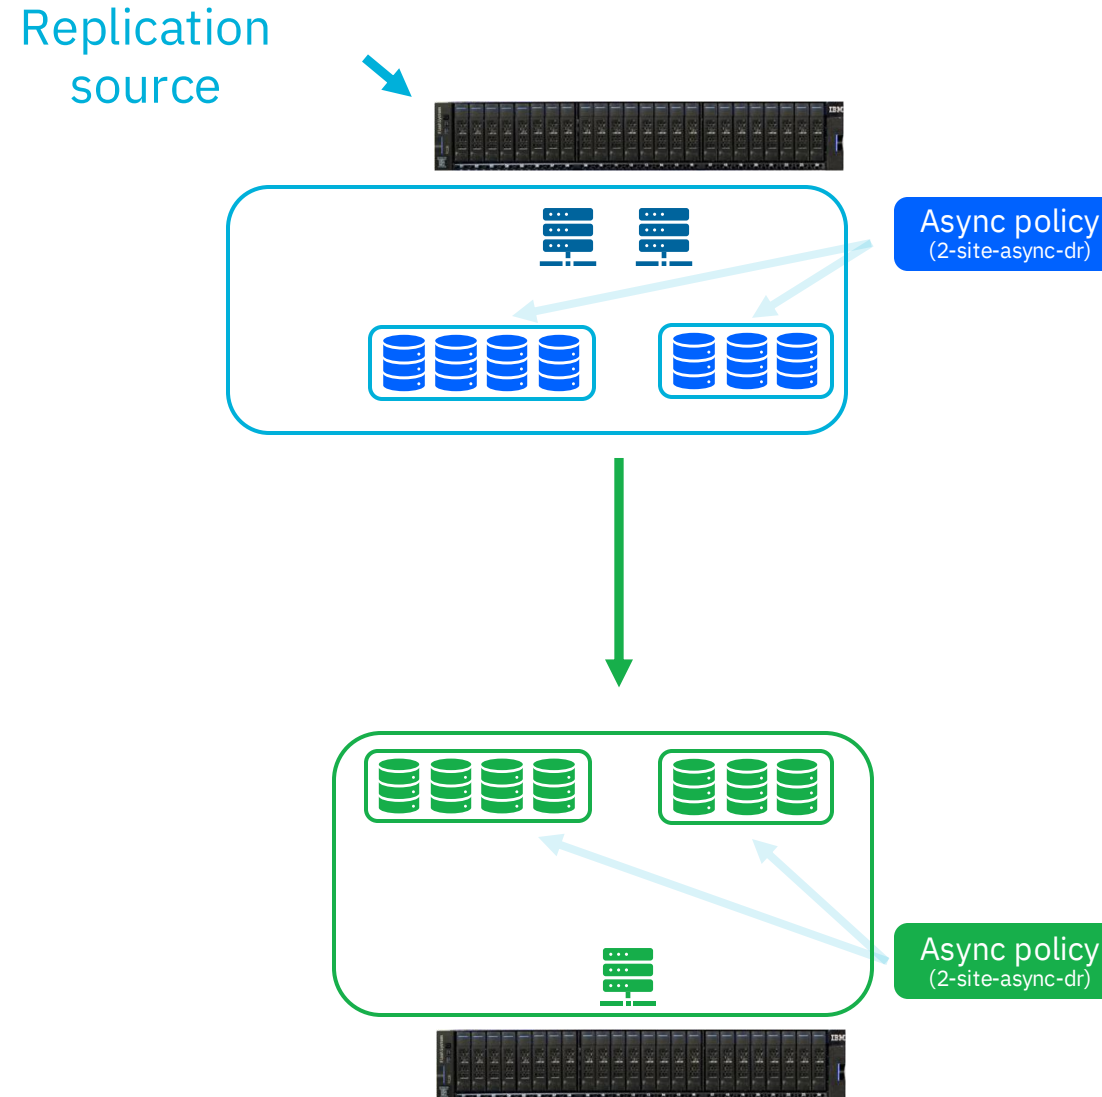
Migrating Existing Async DR Volume Groups

Replication source

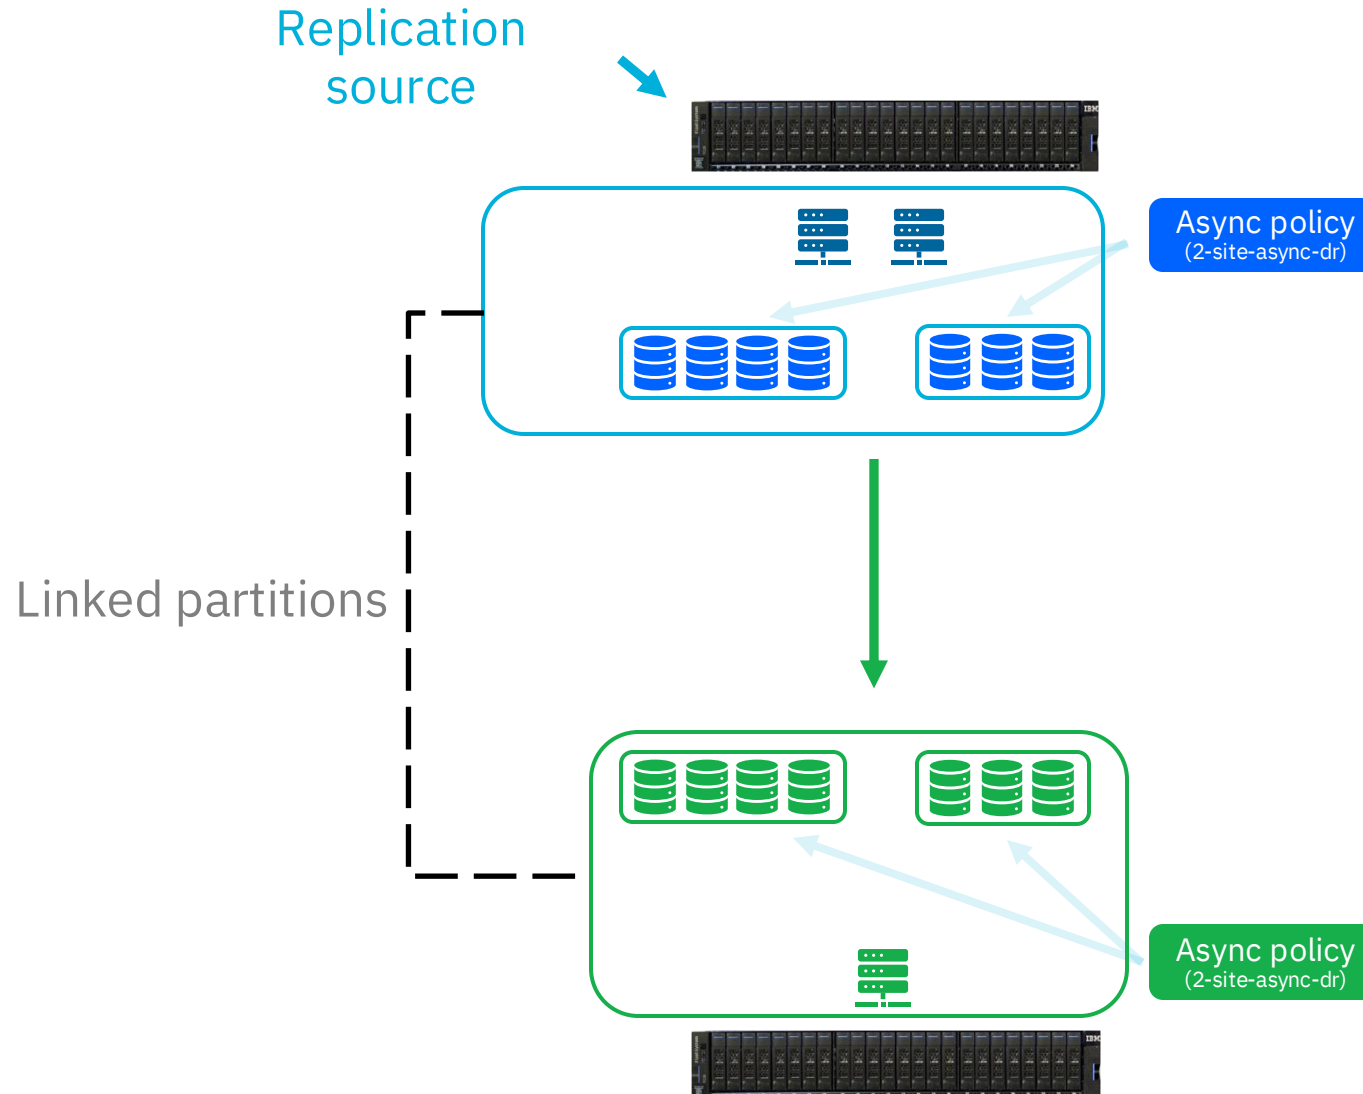


– During this process replication continues and is not stopped

New Partitions Are Seeded On Each System





Partitions Are Linked Together




Replace Existing Replication Policies With Ones Created In The Partition

Replication source

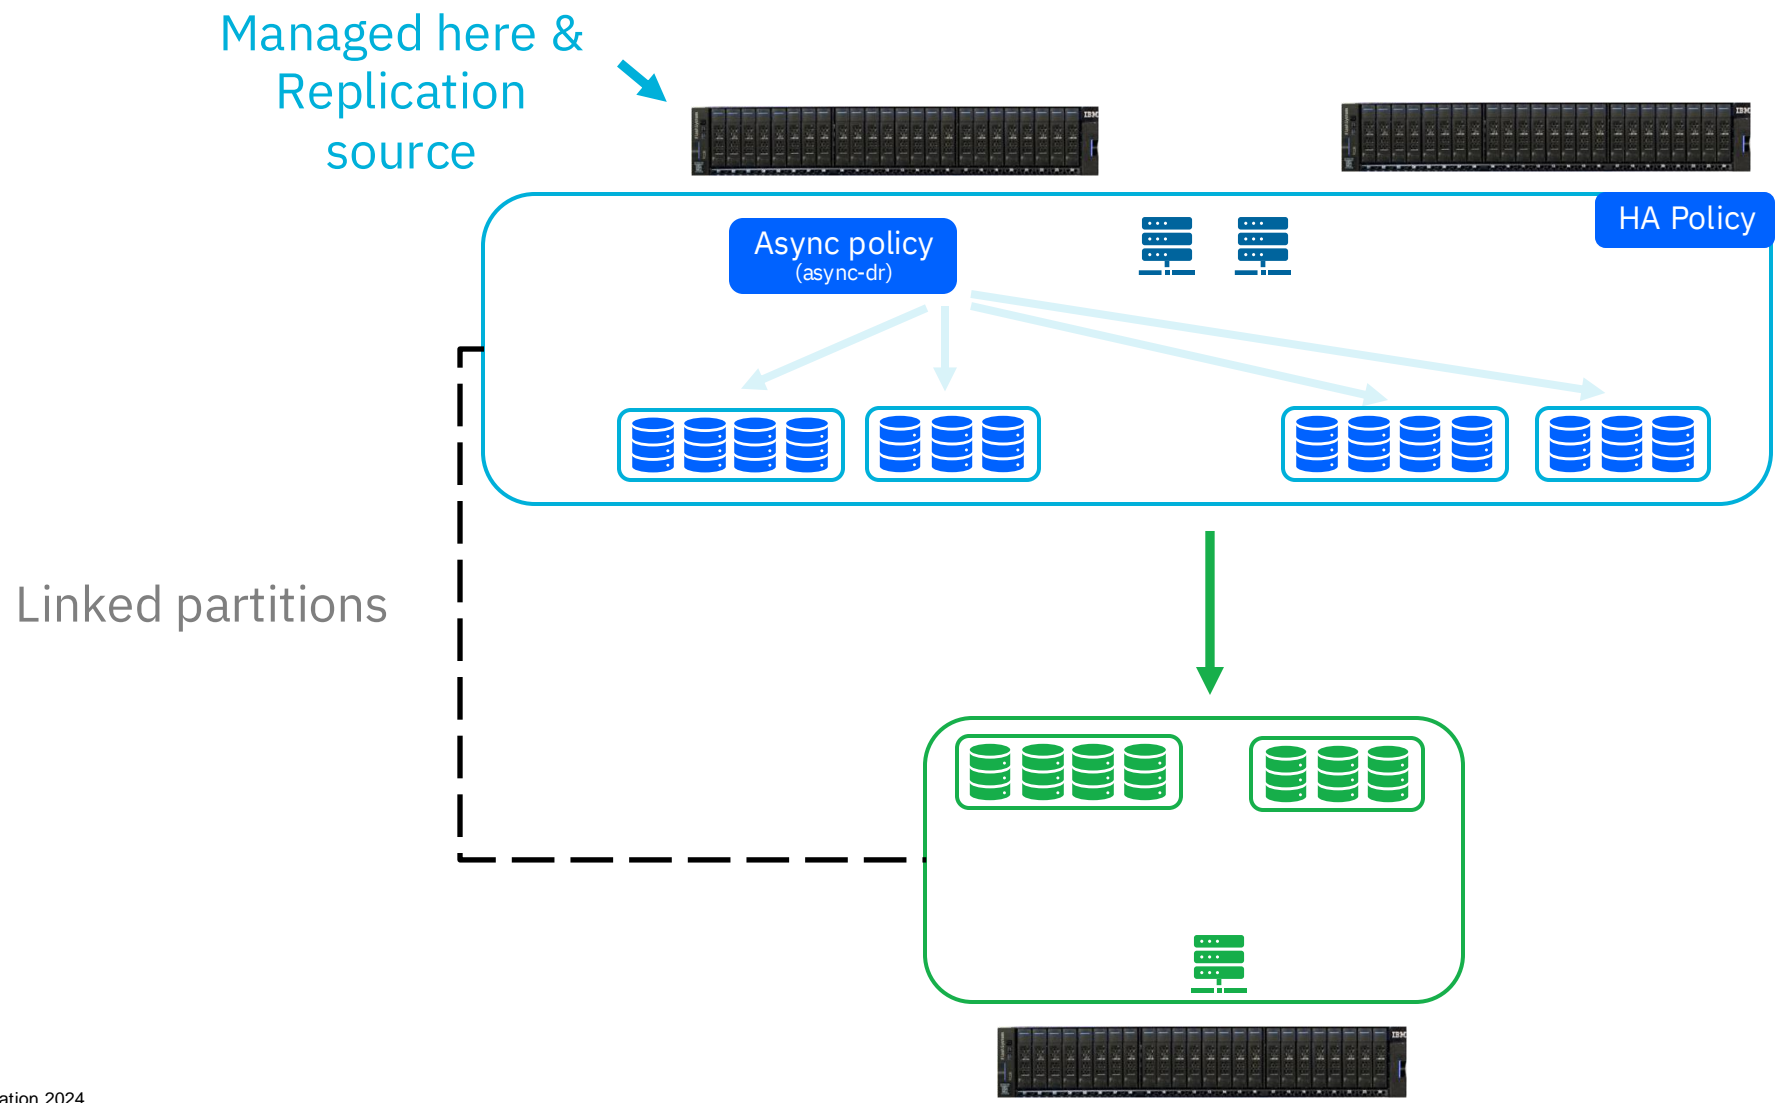


Linked partitions



- This step is not forced, but trying to add HA will cause errors

Add HA Policy To Partition (If HA+DR Is Desired)



Disaster Recovery Using Storage Partitions

- Configure disaster recovery using partitions to position for adding HA in the future and compatibility with Flash Grid

IBM FlashSystem 5200 | pine-c | Storage partition: Prod Partition

Configure policy-based replication for Prod Partition

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

Disaster Recovery Using Storage Partitions

- Consistent look and feel for easy management

The screenshot displays the IBM FlashSystem 5200 management interface. The top navigation bar shows 'IBM FlashSystem 5200', 'pine-c', and 'Storage partition: Prod Partition'. A left sidebar contains a menu with options: Prod Partition, Volume groups, Volumes, Volume mappings, Hosts, Replication policies, Snapshot policies, Partnerships, and 'Exit Prod Partition'. The main content area is titled 'Prod Partition' and features a 'Management partition' button. The 'Replication overview' section shows a diagram with 'Prod Partition' at the top, connected to 'DR Partition' at the bottom. The 'Prod Partition' is associated with host 'pine-c' (logged in here), and the 'DR Partition' is associated with host 'almond-c'. A green arrow labeled 'DR healthy' indicates the replication status. The 'Storage components' section on the right lists: Hosts (1) Online, Volumes (12), and Volume groups (1) with 1 within RPO. The 'Connectivity' section shows a 'DR partnership (almond-c)' as 'Configured'. The 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.

Disaster Recovery Using Storage Partitions

- Click to add high availability

The screenshot displays the IBM FlashSystem 5200 management interface. The top navigation bar shows 'IBM FlashSystem 5200', 'pine-c', and 'Storage partition: Prod Partition'. A left sidebar contains navigation options: 'Prod Partition', 'Volume groups', 'Volumes', 'Volume mappings', 'Hosts', 'Replication policies', 'Snapshot policies', 'Partnerships', and 'Exit Prod Partition'. The main area is titled 'Prod Partition' and features a 'Management partition' dropdown. The 'Replication overview' section shows a diagram with 'Prod Partition' at the top (hosted on 'pine-c') and 'DR Partition' at the bottom (hosted on 'almond-c'). A green arrow labeled 'DR healthy' points from the Prod Partition to the DR Partition. A context menu is open over the Prod Partition, with 'Add high availability' highlighted. The 'Storage components' section on the right lists: 'Hosts (1) Online', 'Volumes (12)', 'Volume groups (1) 1 within RPO', and 'Connectivity' with a 'DR partnership (almond-c) Configured'.

Disaster Recovery Using Storage Partitions

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


- Navigation Panel (Left):** Lists various management options including 'Prod Partition', 'Volume groups', 'Volumes', 'Volume mappings', 'Hosts', 'Replication policies', 'Snapshot policies', 'IP quorum', and 'Partnerships'. An 'Exit Prod Partition' button is at the bottom.
- Header:** Shows 'IBM FlashSystem 5200', 'pine-c', and 'Storage partition: Prod Partition'. A 'Manage partition' dropdown is on the right.
- Replication Overview (Center):** A diagram showing two hosts in the 'Prod Partition': 'pine-c' (logged in) and 'pecan-c'. They are connected by a line labeled 'HA established'. Below this, a 'DR healthy' status is shown with a downward arrow pointing to the 'DR Partition'.
- DR Partition (Center):** Shows a single host 'almond-c'.
- Storage Components (Right):**
 - Hosts (1):** Online.
 - Volumes (12):** (Link to details)
 - Volume groups (1):** 1 within RPO.
- Connectivity (Right):**
 - HA partnership (pecan-c):** Configured.
 - IP quorum:** Connected.
 - DR partnership (almond-c):** Configured.
- Footer:** Performance metrics: Latency 0 ms, Read 0 ms, Write 0 ms; Bandwidth 0 MBps, Read 0 MBps, Write 0 MBps; IOPS 0, Read 0, Write 0.

IBM Storage Virtualize

Policy-based High Availability with Async DR

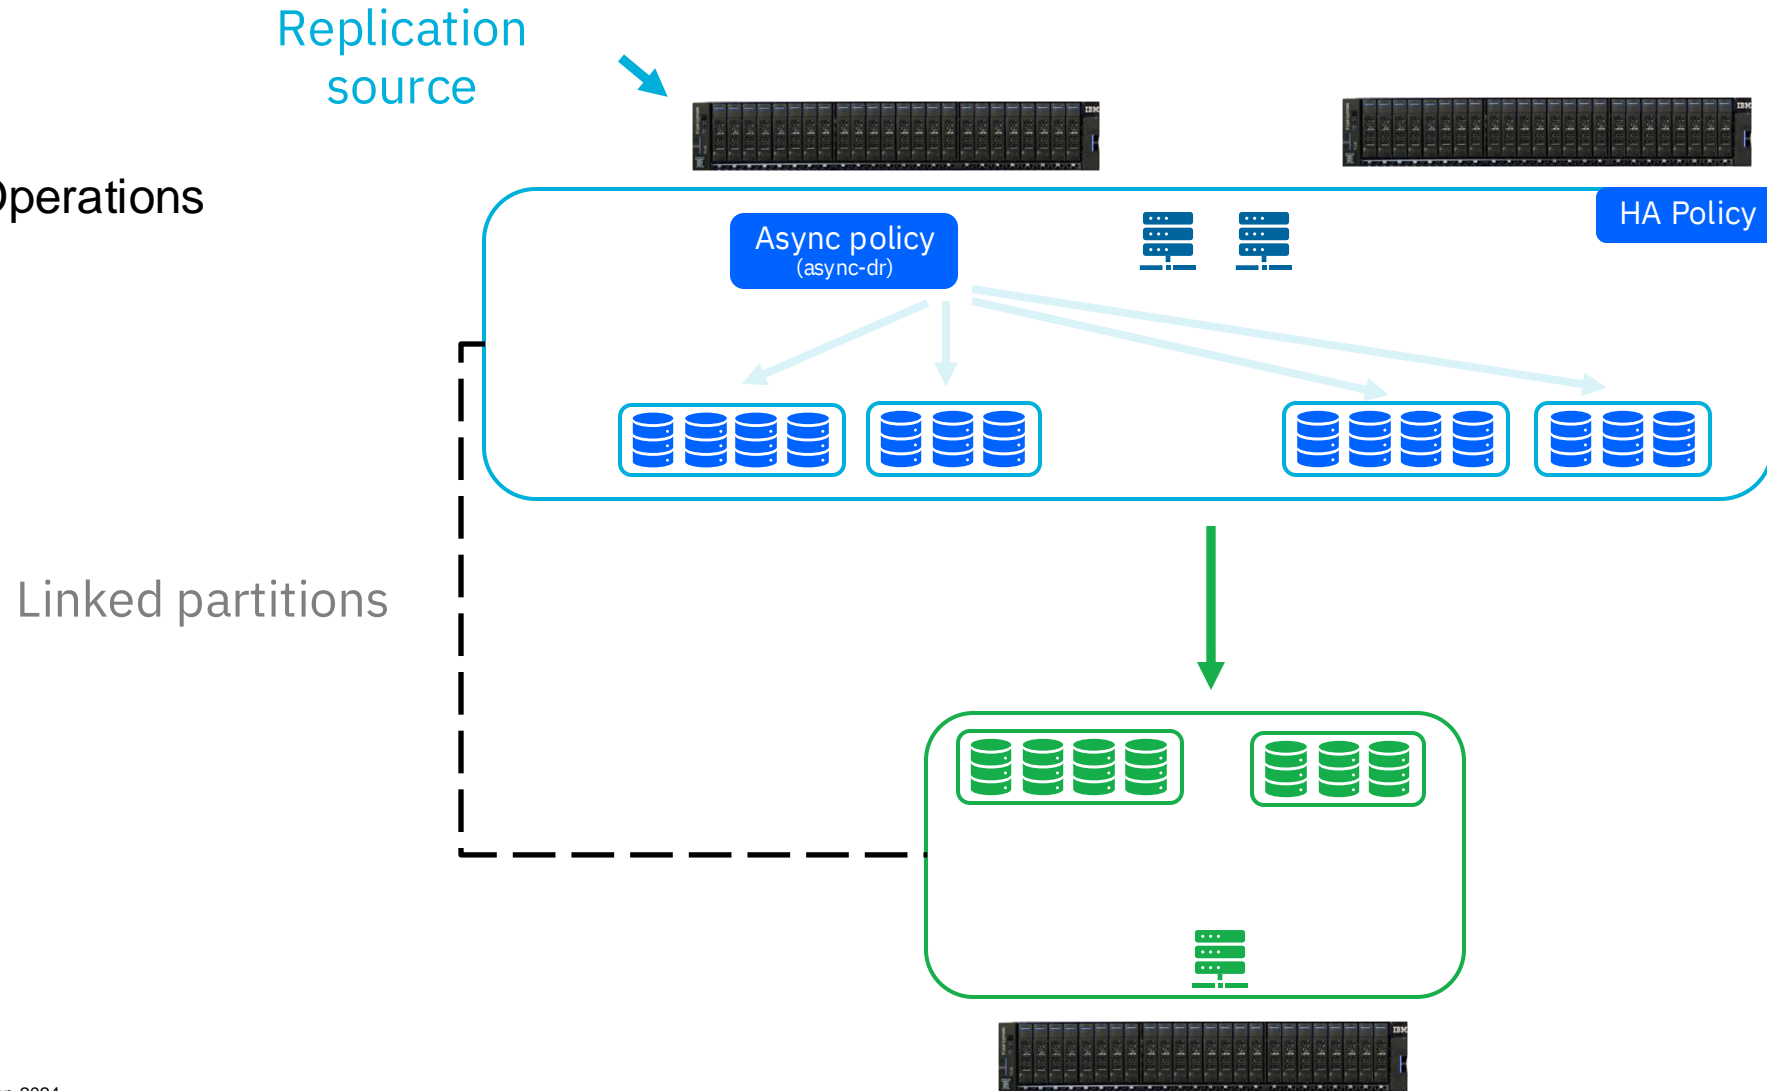
What Happens in a Disaster?

8.7.1



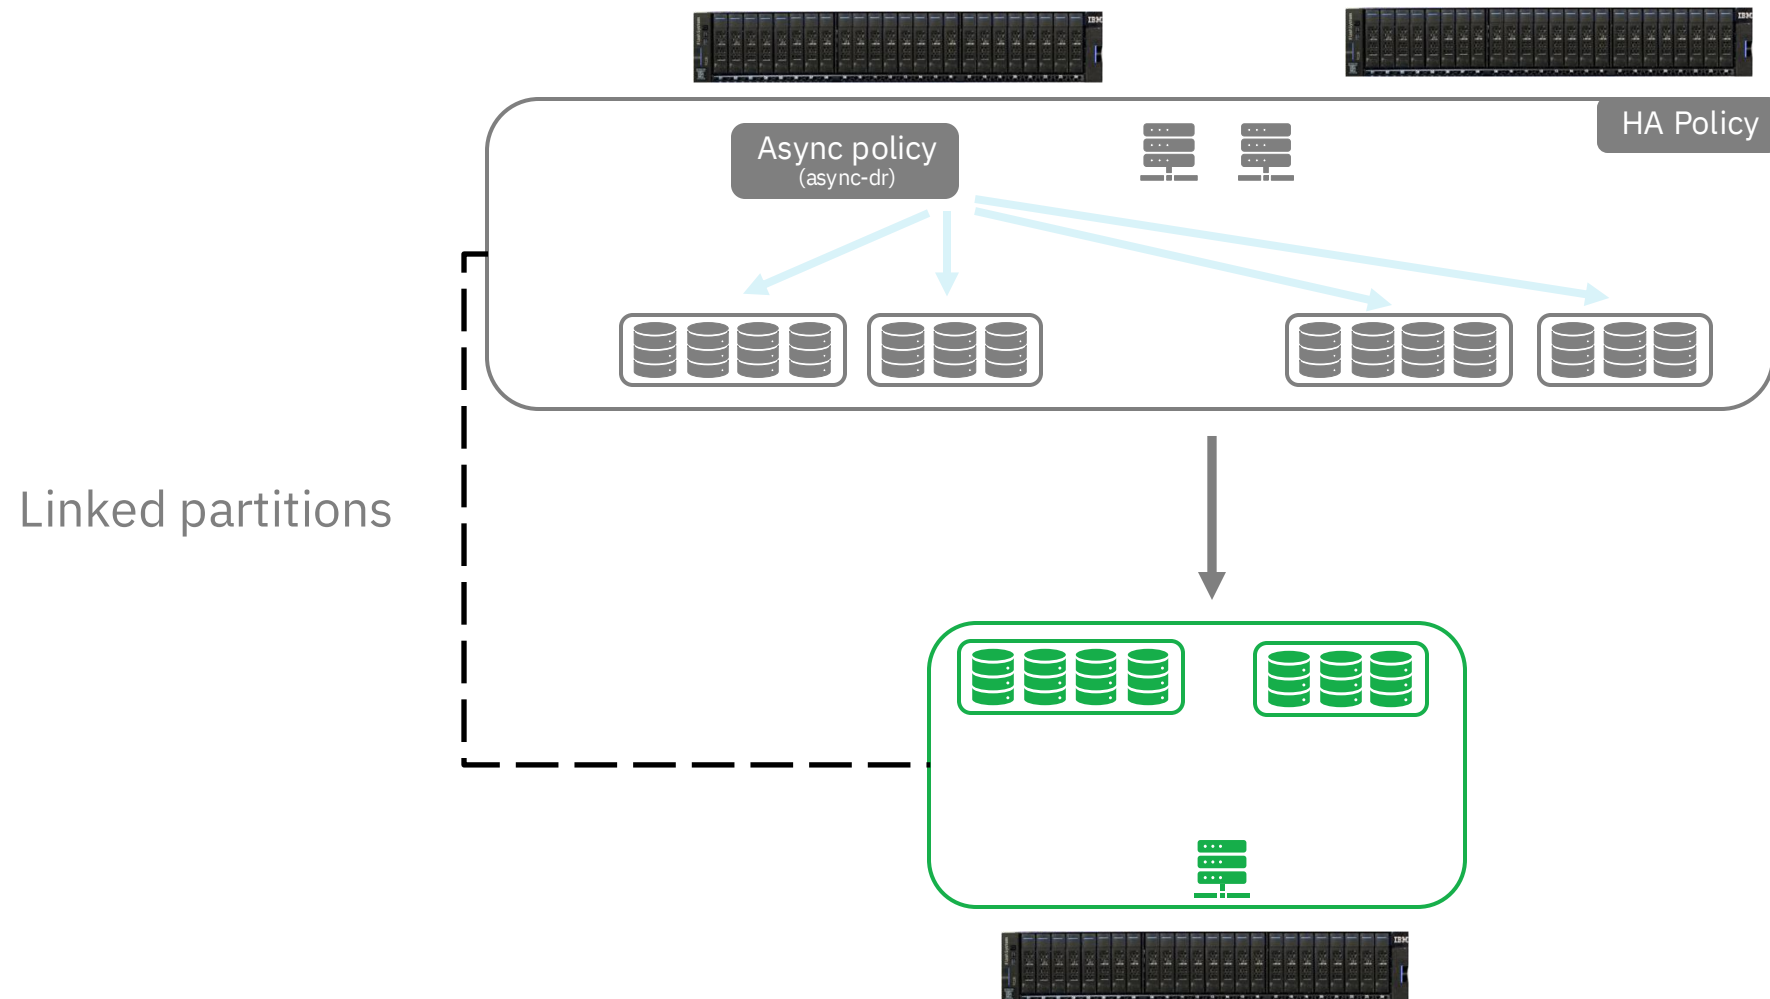
Recovering From Disaster

– Normal Operations



Recovering From Disaster


- Loss of both HA Systems



Recovering From Disaster

- Access Enabled on DR Volume Groups

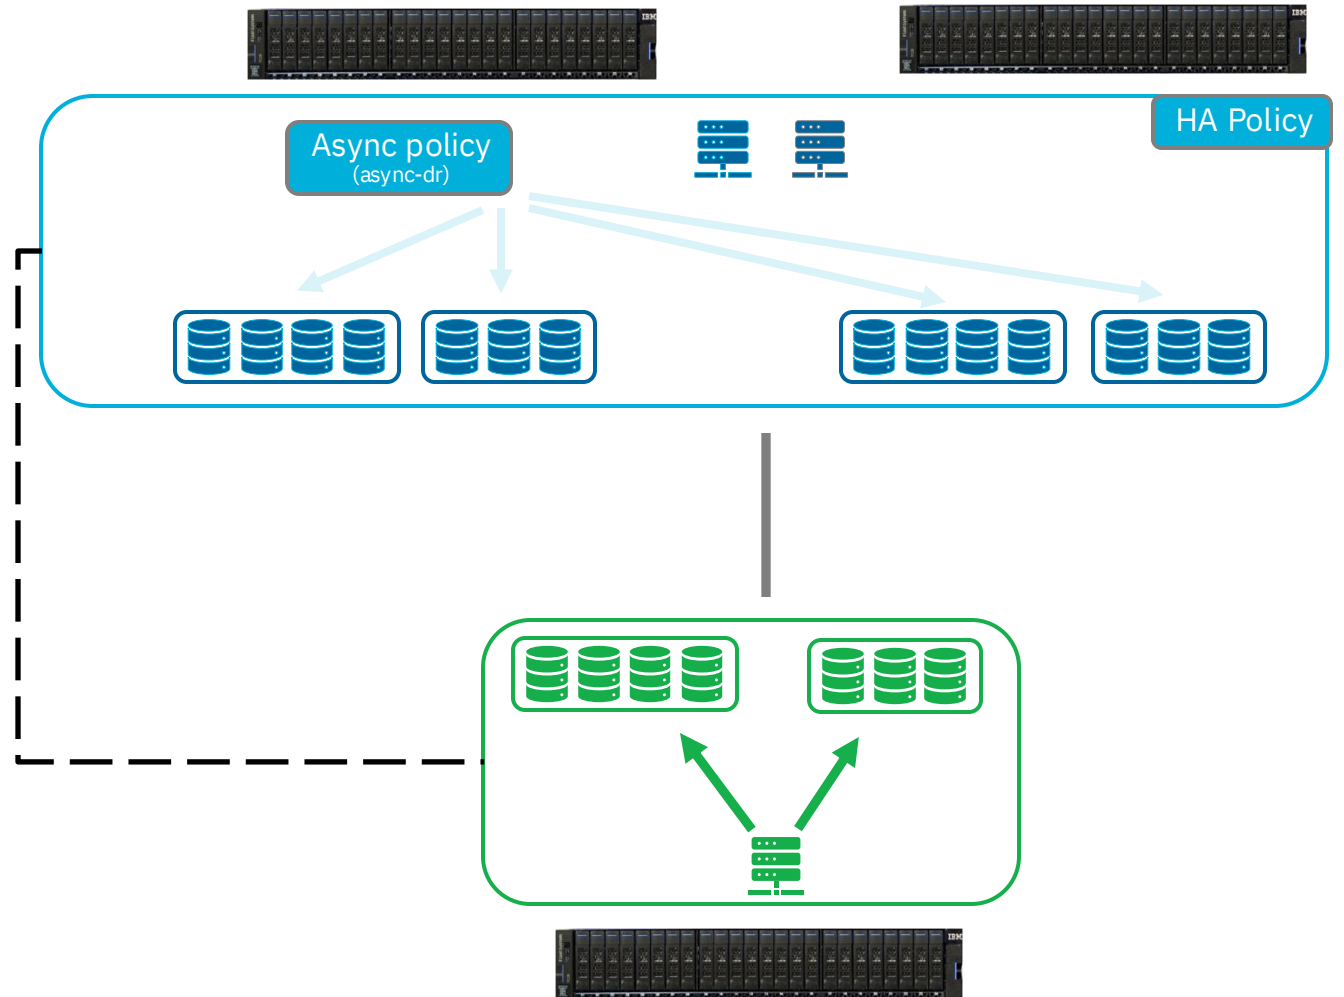
Linked partitions



Recovering From Disaster

- When HA systems return replication is NOT automatically restarted


Linked partitions



Recovering From Disaster

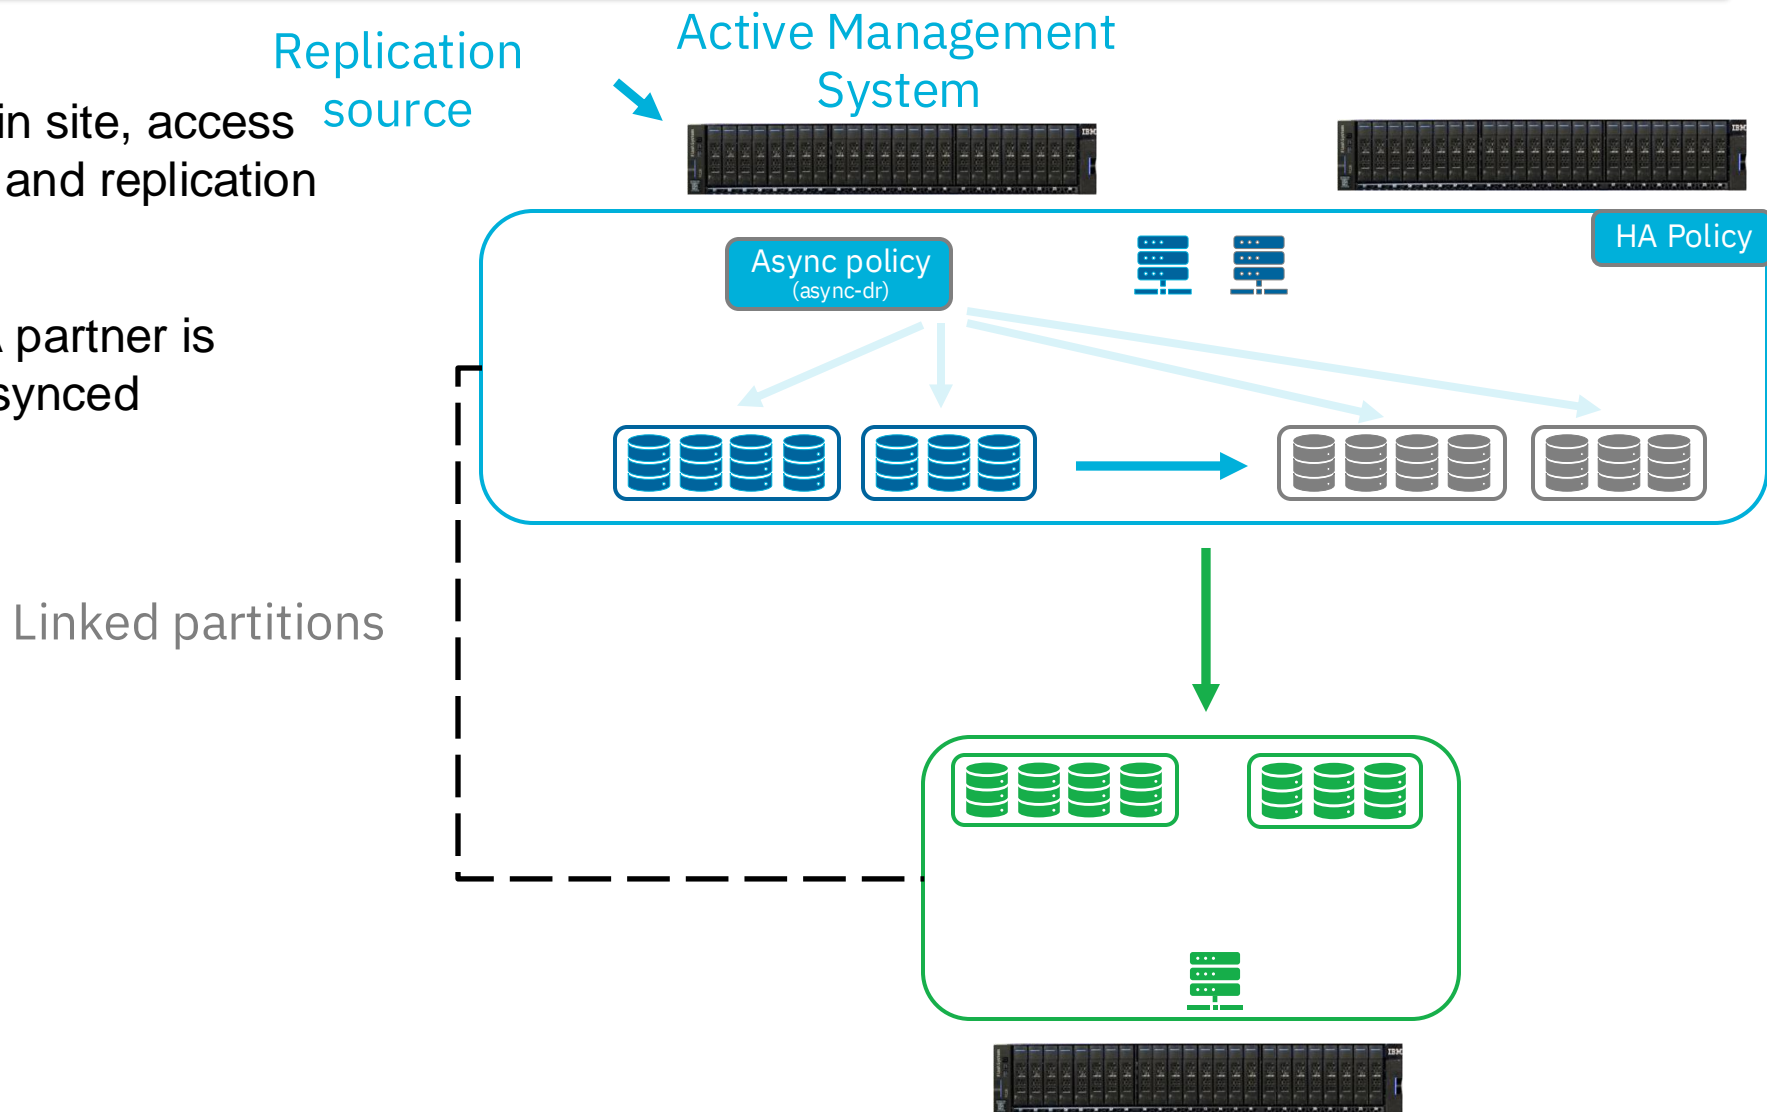
- When ready to move back, replication is restarted from the DR system, replicating back to the active management system.
- HA for the whole partition will show as suspended during recovery
 - The number of volume groups does not matter
 - For a partition to be considered HA, ALL volumes must be synchronized
 - The synchronization of HA does not occur until after all the DR volume groups are moving from primary to DR

Linked partitions



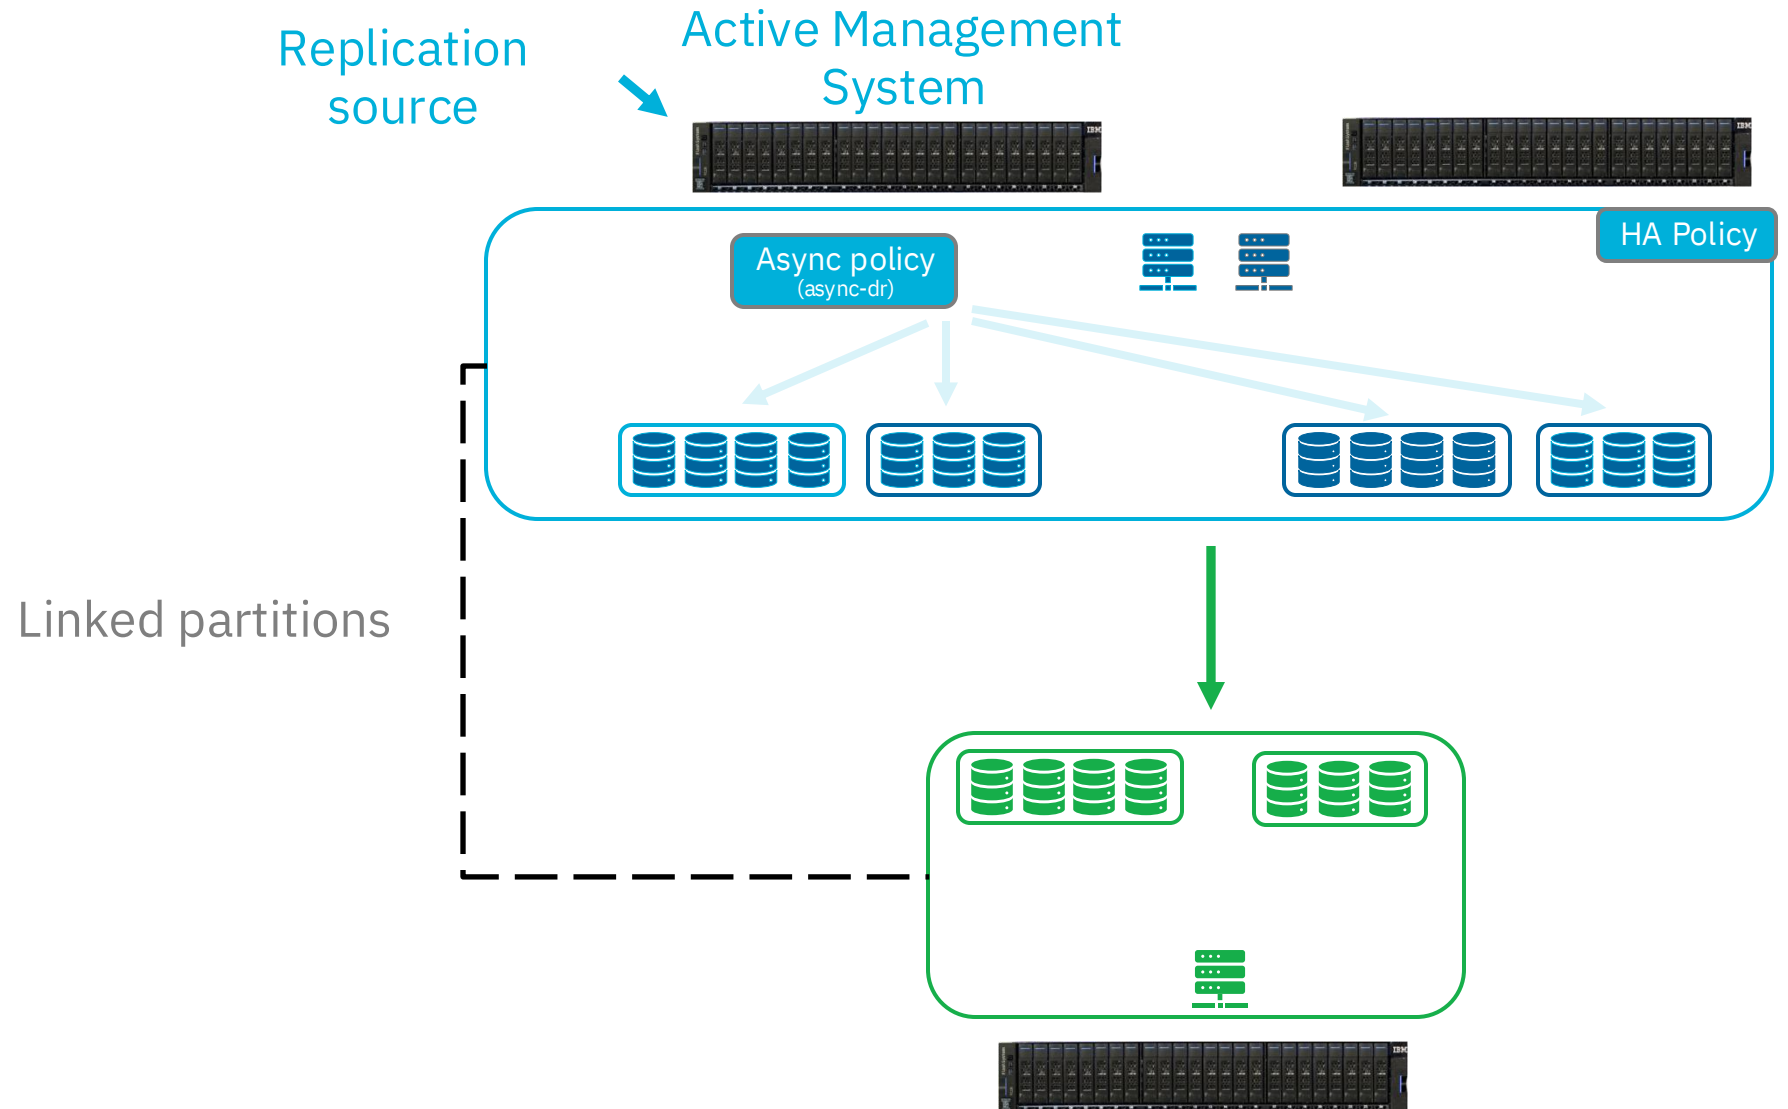
Recovering From Disaster

- Once DR is synced with the main site, access can be enabled at the main site and replication restarted to the DR site
- When the above occurs, the HA partner is automatically, incrementally re-synced



Recovering From Disaster

– Full health restored

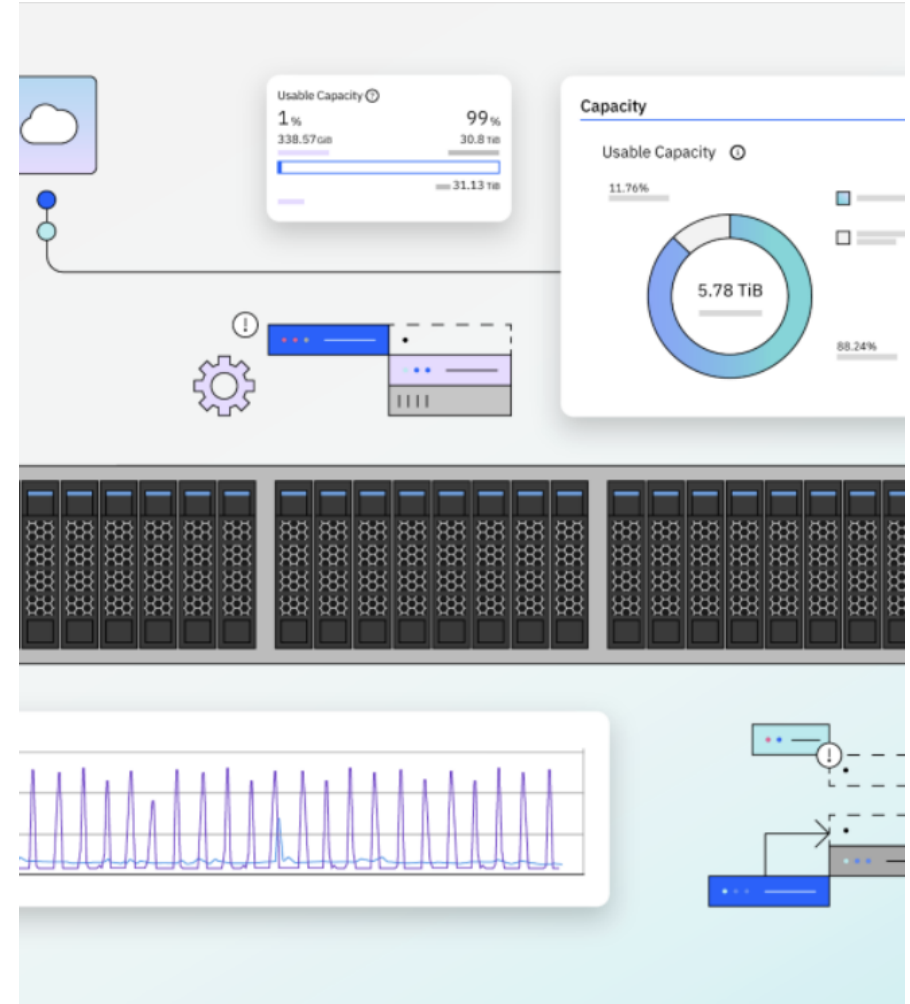


IBM Storage Virtualize

Policy-based High Availability with Async DR

DR Site Testing

8.7.1



Enable Access

- To test a volume group at the DR site, you can enable access to that volume group.
- Pros:
 - Testing on the actual target volumes
- Cons:
 - Replication to DR is stopped during testing and RPO falls behind
 - Replication must be restarted from the production site
 - If replication is started from the wrong site, production data is in danger of being destroyed

Snapshot/Thin Clone

- A snapshot of the target volume group can be taken, and a thin clone provisioned off it to test access to the data
- Pros:
 - Replication is not stopped to test
- Cons:
 - Not testing against the actual target volumes
 - Increased volume count
 - Volumes for snapshots
 - Volumes for thin clones
 - Must manually be cleaned up when finished
 - If not cleaned the copies (if manually taken) will exist forever

Recovery Test


- Storage Virtualize has a recovery test function (released in 8.6.2) that can be used from the CLI/API at the DR site
 - This will allow testing on the PBR recovery volumes while maintaining replication in the background
- **chvoluemgroupreplication -startrecoverytest <VG name>**
- **chvolumeegroupreplication -stoprecoverytest <VG name>**
- Pros:
 - Testing on the actual recovery volume
 - Replication is maintained during testing
 - The system cleans up after – stoprecoverytest is issued
- Cons
 - CLI and API only

IBM Storage Virtualize

Policy-based High Availability with Async DR

Useful Information

8.7.1



Partnerships

Support Notes

- **High availability** can be configured using:
 - Fibre Channel partnerships
 - Short-distance IP partnerships using RDMA
 - Max 1ms RTT (SCORE requests can be submitted for higher)
 - Applies to both FC and RDMA partnerships
- **Disaster recovery** can be configured using:
 - Fibre Channel partnerships
 - Max 250ms RTT
 - Long-distance IP partnerships using TCP
 - Max 80ms RTT
 - Short-distance IP partnerships using RDMA
 - Max 3ms RTT
 - 80ms for all partnerships on the FlashSystem 5045
- Separate ports required for HA, DR and host traffic
- The Background Copy rate setting has been deprecated and hard coded at 100%
- Partnership panels are now under Settings in the 8.7.1 GUI
- Support for 1 HA partnership and up to 2 DR partnerships on a system, or
- Support for 3 DR partnerships on a system
- Up to 2 partnerships can be short-distance IP using RDMA
- **8.7.1 and later will only partner with systems running 8.7.0 or later**

Planning

Support Notes

Model	Replicated volume count	2-site HA or DR capacity	3-site HA+DR capacity
FlashSystem 5015/5035	Not supported	Not supported	Not supported
FlashSystem 5045	4096	200 TiB	Not supported
FlashSystem 5200/5300	7932	1024 TiB	512 TiB
FlashSystem 7200	7932	2048 TiB	1024 TiB
FlashSystem 7300	16050	2048 TiB	1024 TiB
FlashSystem 9200	7932	2048 TiB	1024 TiB
FlashSystem 9500	32500	4096 TiB	2048 TiB
SAN Volume Controller (SA2/SV2)	7932	2048 TiB	1024 TiB
SAN Volume Controller (SV3)	7932	4096 TiB	2048 TiB

- Each volume using HA+DR counts as 2x replicated capacity
- VMware vVol replication is supported for 2-site DR only

8.7.1

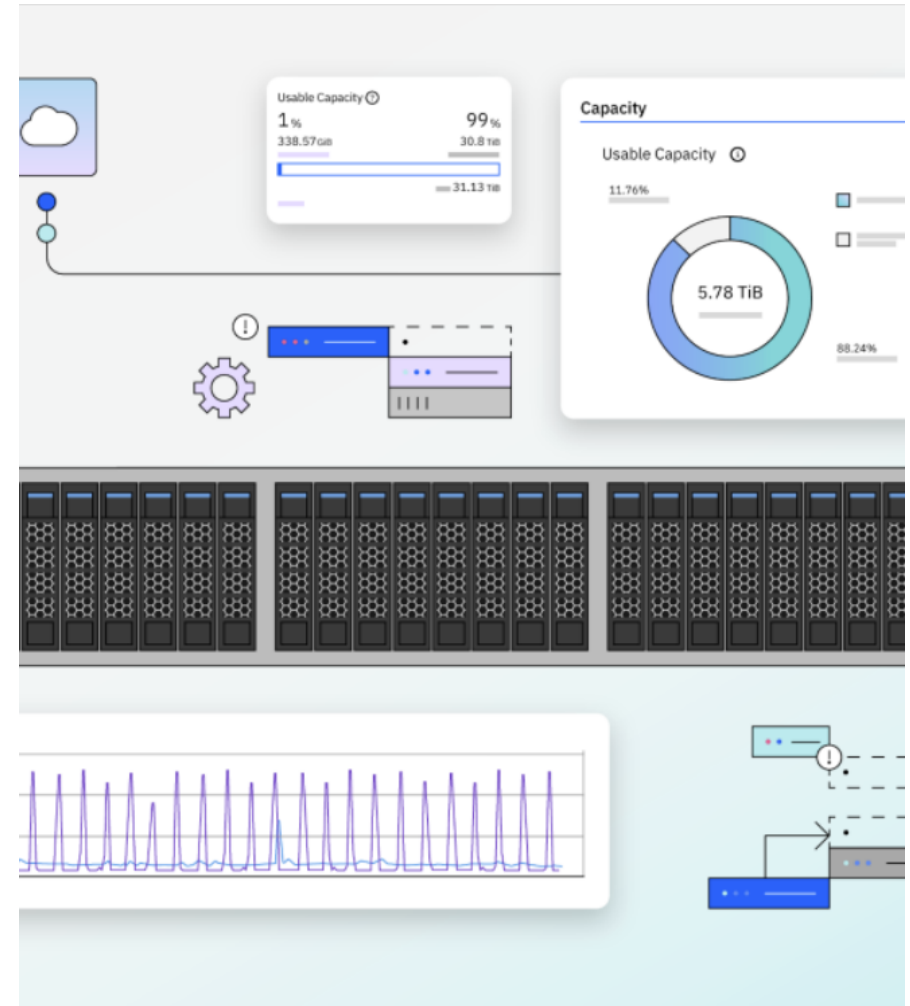
Follow-on releases

Support Notes

- 3-site HA+DR requires high-bandwidth links to DR system
-
- 3-site support for low-bandwidth links to DR system, with automatic mode switching
 - Support for synchronous replication to the DR system (both 2-site and 3-site)
 - Statement of Direction

IBM Storage Virtualize

Support Statements



Storage Virtualize 8.7.1 Update

8.7.0 is the final version that supports these features

FlashSystem 9100

- Released in 2018, it's following the natural product lifecycle.

Remote Copy

- Metro Mirror, Global Mirror, HyperSwap and volume mobility are all superseded by policy-based replication, policy-based high availability and partition migration.

iSER host attach and Ethernet clustering

- Ethernet clustering is only used with HyperSwap, RDMA-based Ethernet available for PBR and PBHA.
- iSER host support was not widely adopted with the industry focused towards NVMe.

TRAIID 10

- DRAID has far improved rebuild times and performance, making it more suitable for modern systems than traditional RAID.

FlashSystem Support For Multiple I/O Groups

- Flash Grid replaces I/O groups as the way to scale-out FlashSystem
- Multiple I/O group FlashSystems can remain on 8.7.0 or consolidate to a single I/O group to use Flash Grid

SAN Volume Controller Support For Multiple I/O Groups

- SVC will continue to support multiple I/O groups on LTS releases (i.e., 8.7.0, 8.8.0).
- Non-LTS releases will only support single I/O group SVC systems
- (i.e., 8.7.1, 8.7.2).

IBM Champions for FlashSystem

Visit <https://community.ibm.com/community/user/champions/ibm-champions-program> or get in touch with Kyle Cline (kcline@us.ibm.com) to begin your nomination to become an IBM Champion.



The Program

We invite you to join the IBM Champions program for FlashSystem, forming a trusted group of users who will play a pivotal role in shaping the future of FlashSystem.

IBM Champions demonstrate practical expertise in IBM technologies, while providing extraordinary support and advocacy in both digital and local communities. They are top experts and advocates around IBM technology.

The IBM Champions program recognizes these innovative thought leaders in the technical community and rewards their contributions by amplifying their voice and increasing their sphere of influence. IBM Champions are enthusiasts and advocates: IT professionals, business leaders, developers, executives, educators, and influencers who support and mentor others to help them get the most out of IBM software, solutions, and services.

Acts of Advocacy

We encourage this enthusiastic group of experts and professional technologists to both work with us and independently to do many of the following acts of advocacy:

- Endorse internally at their own company
- Provide a 1:1 external endorsement
- Create and distribute content on social media
- Create, manage, and actively participate in User Groups or other small, regional events
- Host and execute events, both live and online
- Present at events and conferences
- Provide product feedback, product reviews, and case studies
- proof-Participate as a marketing, sales, or analyst reference
- Actively mentor, teach, and participate in research
- Participate in of-concept opportunities.
- Contribute to a press release
- Actively engage in the IBM TechXchange Community
- Create and distribute content on social media

We are asking our FlashSystem Champions to participate in at least one activity by the end of 2024.

Benefits

To recognise our champions and to develop your technical eminence, FlashSystem skills and knowledge of the wider storage industry, IBM Champions may be eligible to receive:

- Invitations, discounts, and VIP perks at IBM conferences and events.
- Annual package of IBM Champion logo gear to wear, use, and display.
- An IBM Champion-verified digital credential through the IBM Digital Badge Program.
- Visibility, recognition, and networking opportunities.
- A dedicated workshop annually to discuss roadmap and influence upcoming features
- Access to local networking events
- Exclusive access to online IBM communities.
- Exclusive access to IBM product development teams.
- Promotion of content and assets via IBM's channels.
- Special recognition on their IBM TechXchange Community profile.

Thank you!

Accelerate with ATG Survey

Please take a moment to share your feedback with our team!

You can access this 6-question survey via [Menti.com](https://www.menti.com) with code 1708 6924 or

Direct link <https://www.menti.com/alwhyze7z1gz>

Or

OR Code

