



IBM i Migrate While Active

Methods, architecture, cutover strategies, and enhancements

Nicolae Chirea

System Architect

Seidor Iberia

nchirea@seidor.com



#whoami

Nicolae Chirea



- Infrastructure solutions architect – design and implement solutions based on IBM Power and IBM Storage
- Covering (What I mostly do)
 - IBM Power system - PowerVM
 - IBM i, Linux, AIX
 - Block storage: Flash Systems, SVC and DS8K
 - AI on Power
 - HA, Performance
 - Open-Source on IBM i
 - Automation: Ansible, PowerVC
- 25+ years of working experience with IBM i
- Joined Seidor in 2007 as IBM Power specialist
 - IBM i – HA/DR solutions, Performance analysis, DB2 for i
 - Linux – HANA, DB2 Warehouse
 - IBM Power automation projects – infrastructure provisioning, HA/DR management, security
- CEAC member since 2024

Agenda

- What Migrate While Active solves
- Architecture, roles, software prerequisites
- Partition mirroring
- Assisted save and restore
- Manual migration
- Project plan, validation, key risks, and takeaways for real migrations

Guiding idea

Keep the source productive as long as possible, validate the copy before the cutover, and reserve downtime for the shortest and safest final step.

Audience lens

Business continuity
Infrastructure team
Db2 Mirror / IBM i admins
Cloud migration
stakeholders

01

OVERVIEW

Why IBM i Migrate While Active matters

A migration framework built to reduce planned downtime while preserving control, validation, and predictable cutover.



Executive summary

What to remember before diving into method details

MWA is a migration product, not a HA/DR platform.

The target copy is only coherent after final synchronization and cutover.

Three methods are now available.

Partition mirroring, assisted save and restore, and manual migration.

2025 added major flexibility.

Partition mirroring and test mode arrived in 2025, plus the manual migration method and better estimation/performance updates.

The design center is short outage.

The longer phases happen while the source stays active; downtime is concentrated at the end.

Migration methods at a glance

Three ways to create and keep the copy node current

Partition mirroring

- No initial outage for media creation
- Host-based replication of SYSBAS
- Test mode before cutover

Assisted save and restore

- Initial outage required
- Install media and user data saves via NFS
- Db2 Mirror + tracking lists maintain changes

Manual migration

- Storage replication, save/restore, or both
- More control over system-copy strategy
- Reduced downtime vs assisted save/restore in many cases

Common endpoint

- Continuous synchronization while source stays active
- Final synchronization and cutover
- Validation is a core part of the approach

02

ARCHITECTURE

Roles, nodes, software, ports, and readiness

Every method uses the same architecture language: source node, copy node, optional GUI node, and method-specific staging components.



Core architecture roles

Source node

Current production IBM i partition and origin of migrated data

Copy node

Target IBM i partition that becomes the new production node after cutover

GUI node

System running the Db2 Mirror GUI; can be the source or a separate node

NFS / staging node

Needed for assisted save and restore to hold install media and migration saves

Software and package prerequisites

Component	Where required	Minimum / key note	Why it matters
5770DBM *BASE	Source + GUI	Required	Provides Db2 Mirror framework
5770DBM Option 2	Source only	Required for MWA	GUI and copy do not need the Option 2 license
5770SS1 Opt. 1/3/12/30/33	Source + GUI	Required	Base support, host servers, Qshell, PASE
IBM HTTP Server + Java	GUI node (or source if GUI runs there)	Required / optional depending on where GUI runs	Needed for the Db2 Mirror graphical interface
cloud-init / python39 / python39-ibm_db	Source node	1.5.1 / 3.9.10 / 2.0.5.12	Open-source dependencies used by MWA flows

Readiness checklist before the first migration run

Technical readiness

- Supported IBM i release and hardware path
- Disk sizing at least as large as the source SYSBAS / user ASP scope
- LAN and service tools connectivity confirmed

Operational readiness

- Change freeze window defined
- Validation owners named
- Cutover decision criteria documented

Data readiness

- Object classes understood
- Replication criteria reviewed
- Large IFS areas inventoried

Risk controls

- Firewalls opened
- Rollback expectations set correctly
- Stakeholders aligned that MWA is not HA/DR

03

PARTITION MIRRORING

The lowest-disruption method when SYSBAS can be mirrored

Host-based storage replication inside IBM i avoids the initial media-creation outage and adds a unique test-mode story.

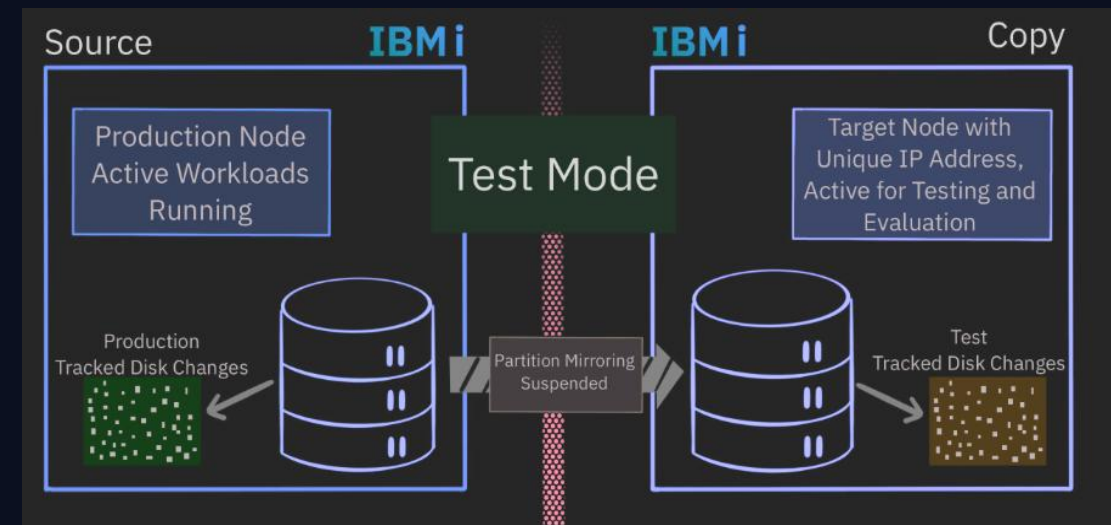



Partition mirroring in one slide

Advantages, limitations, and first decisions

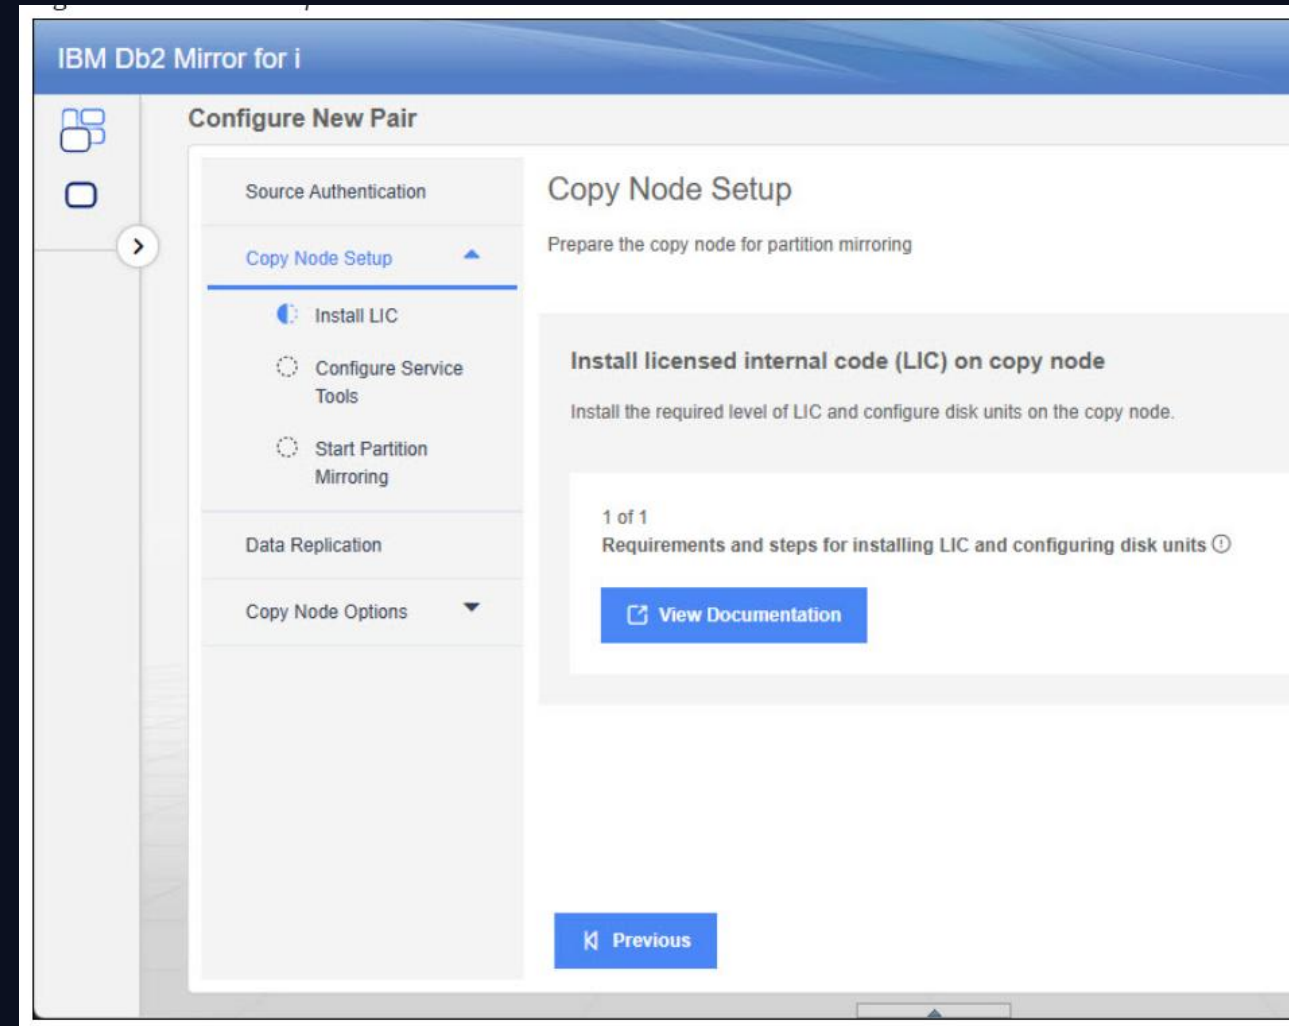
Key messages

- Copies SYSBAS from source node to copy node using host-based storage replication.
- No initial outage is needed because install media and user-data save media are not created first.
- Works with any storage technology because replication occurs inside IBM i.
- Not supported for encrypted user ASPs; IASPs are not replicated by this method.
- Test mode allows the checking of copy node readiness.



Preparing the copy node

- Install Licensed Internal Code at a compatible level with the source.
- Configure storage with SYSBAS and user ASP capacity equal to or greater than the source scope.
- Configure the service tools LAN adapter so the nodes can communicate.
- Keep the copy node in manual / DST-oriented mode until the process explicitly advances.



Replication flow and visibility

Operational behavior

- Initial asynchronous storage replication runs while production continues on the source.
- After full synchronization, replication continues as changes occur on the source node.
- The Db2 Mirror GUI is the primary lens for progress, status, and estimates.
- Replication priority can be adjusted, paused, and resumed as the project requires.

Method detail

IBM Db2 Mirror for i

Configure New Pair

Source Authentication

Copy Node Setup

Data Replication

Copy Node Options

Data Replication

View data replication status before choosing test mode or final cutover

Copy Node Preparation

State: **Completed**

Complete: 100%

Continuous Data Replication

Replication rate is set to: **High** Medium Low Average transfer rate is 212.18 Mbps

State: **Synchronizing** Data Remaining: 11.3 GB Estimated remaining time to sync: 6 minutes

Complete: 50%

Copy Node Options

Copy node options are available when data replication state is synchronized

Test Mode

Use test mode to validate the copy node can run workloads successfully. Data replication is temporarily paused and the copy node is started. Production workloads can optionally continue running on the source node.

Final Cutover

Confirm and finalize the network attributes for the copy node before completing the final synchronization and migration. When complete, the copy node is ready to run the production workloads.

[Previous](#) [Next](#)

Test mode is the signature feature

Active source scenario

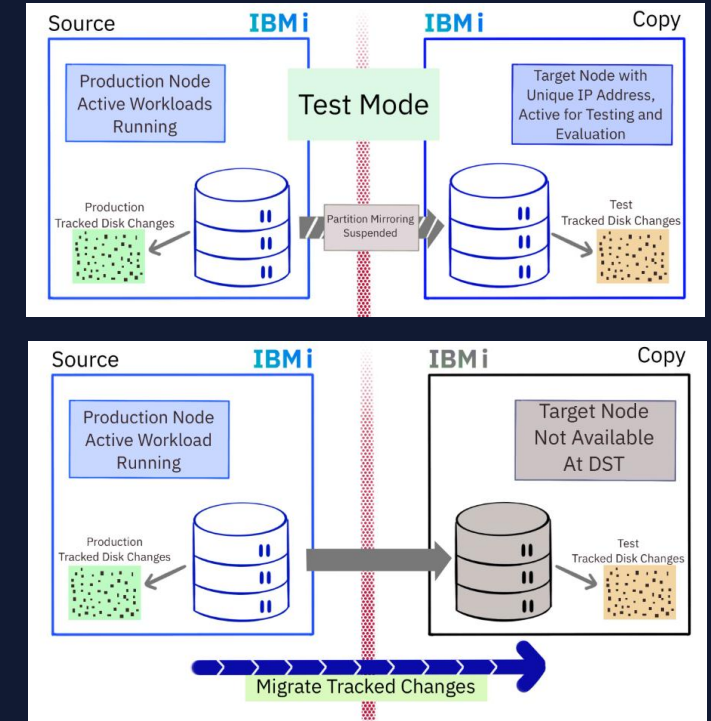
Production keeps running on the source while the copy node is activated with its own unique IP configuration for testing and evaluation.

Inactive source scenario

A colder validation path is also possible when the source can be quiesced or stopped during the test event.

Exit and reconcile

When test mode ends, tracked changes on the copy are undone and source-side tracked changes are migrated back into sync.



Partition mirroring checklist

Do

- Match LIC and supported software levels
- Confirm unsupported encrypted-ASP / IASP cases early
- Use test mode to rehearse operations

Watch

- Network identity planning for the copy node
- Long-running sync on large storage sets
- Change accumulation during validation windows

Avoid

- Assuming test mode equals rollback
- Treating GUI progress alone as the only validation metric
- Under-sizing the target storage footprint

Outcome

- Very low initial disruption
- Strong rehearsal capability
- Shorter, cleaner cutover story

04

ASSISTED SAVE AND RESTORE

The classic MWA pattern when install media and user-data saves are acceptable

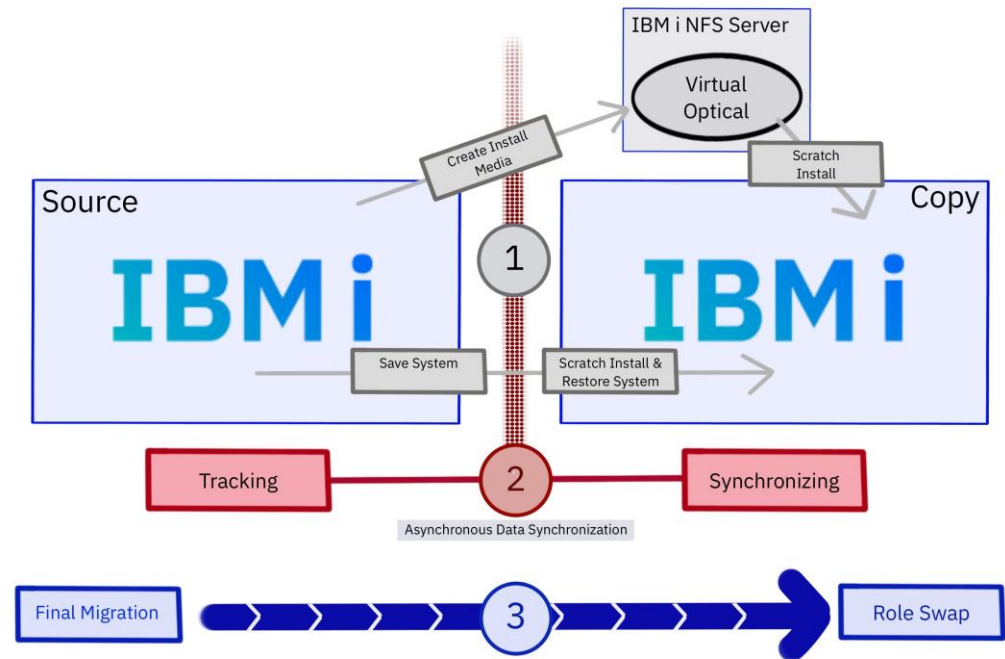
This method begins with a controlled outage to create the copy, then tracks and synchronizes what changed afterwards.



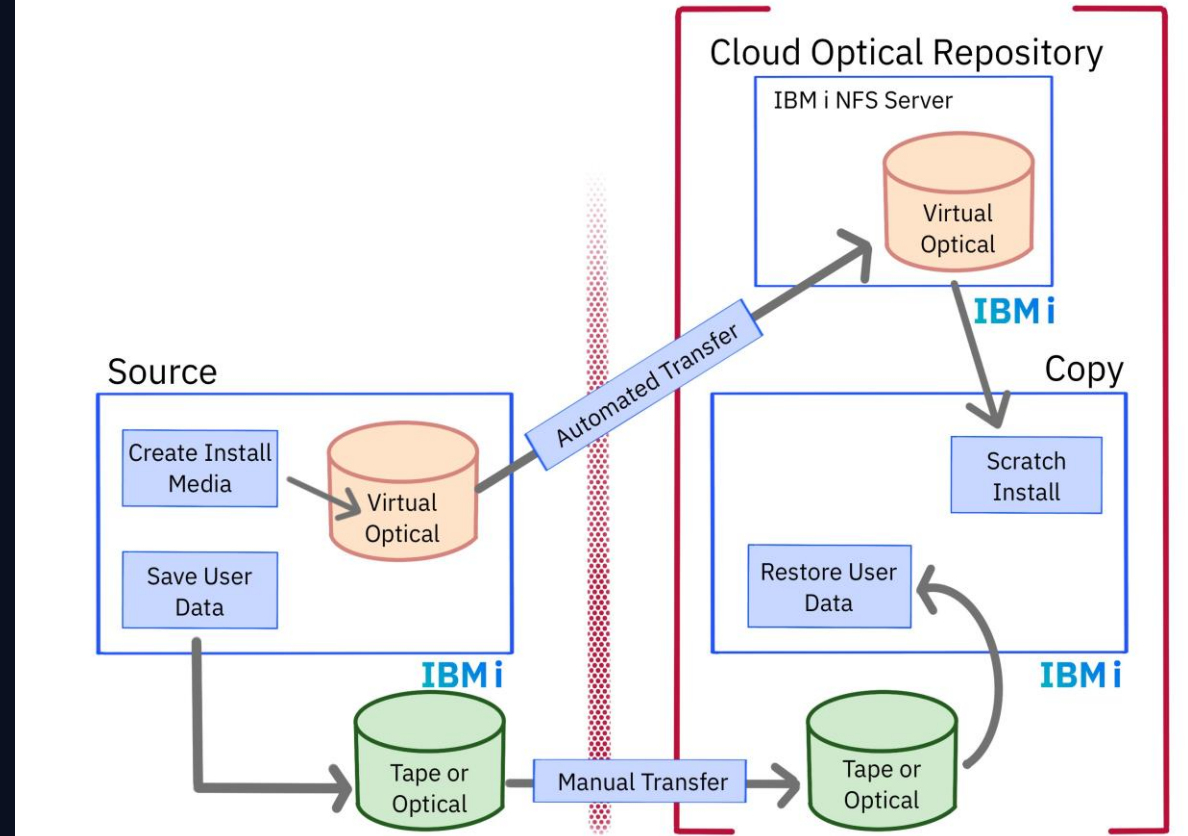
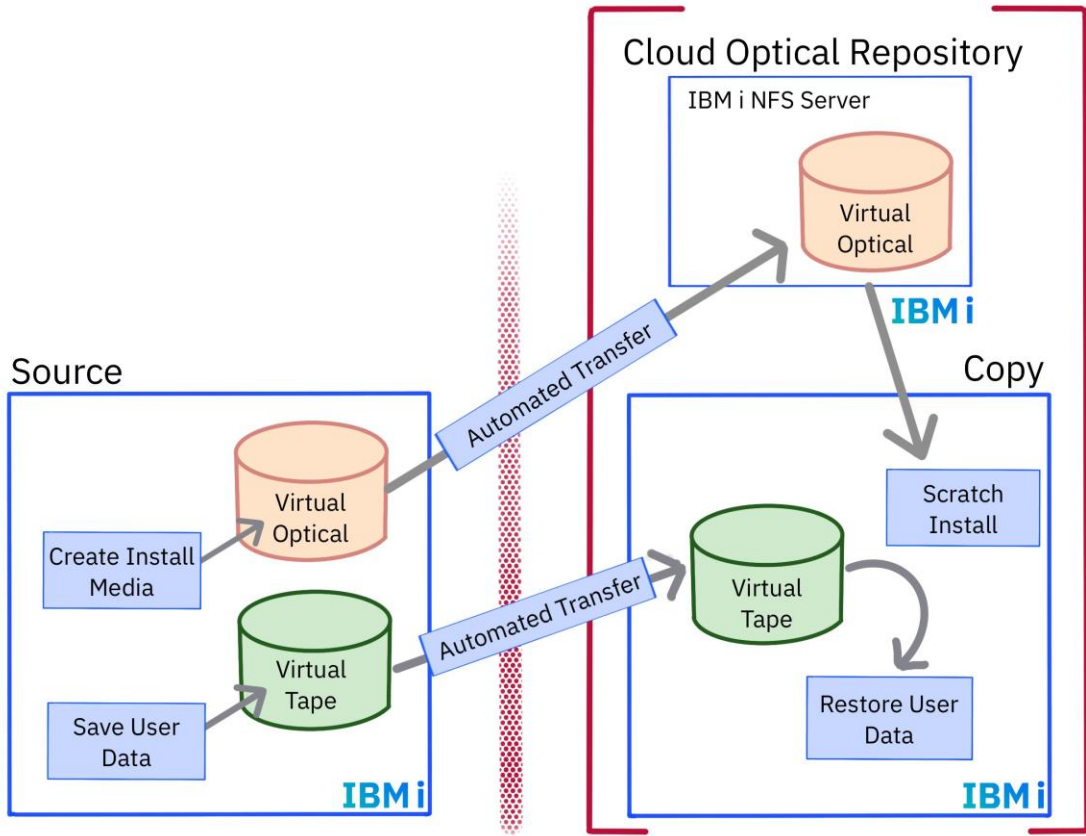
Assisted save and restore flow

Method logic

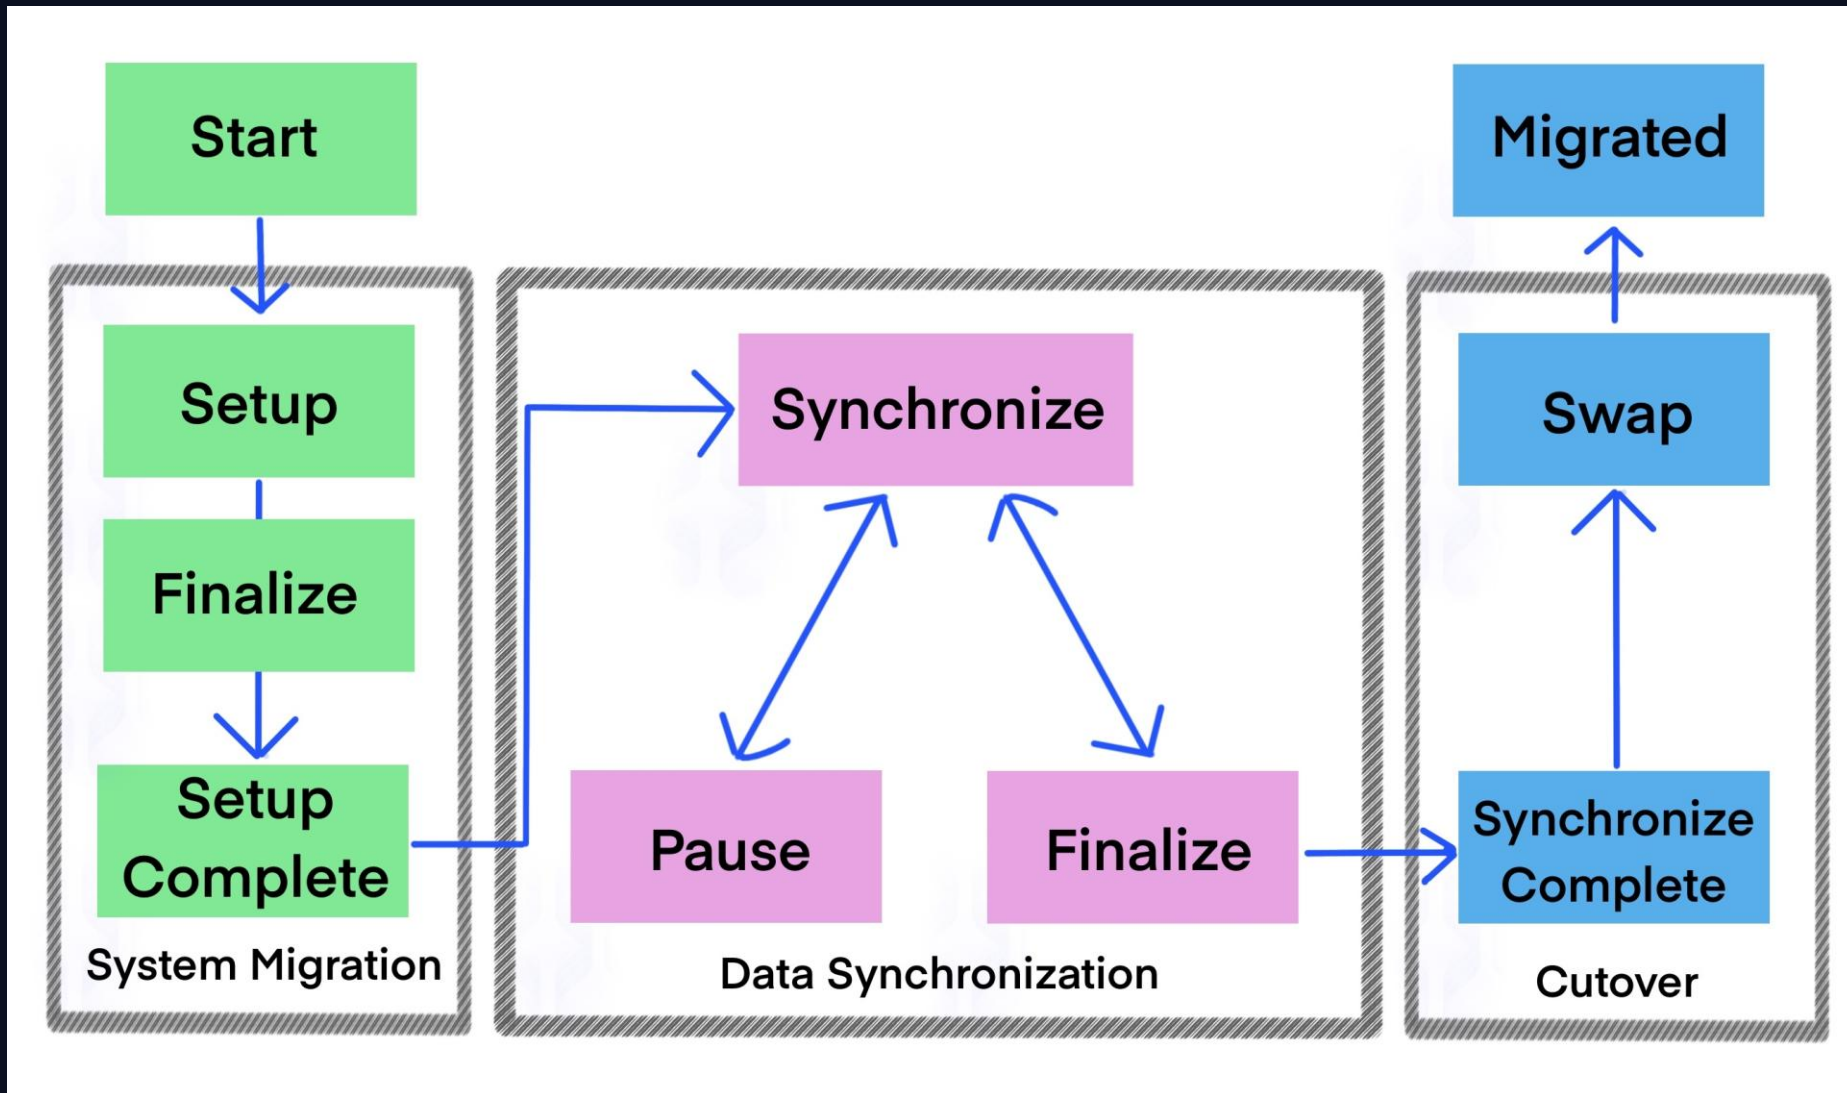
- An initial outage is used to create install media and user data save media.
- Media is staged through an IBM i NFS server or equivalent repository.
- The copy node is scratch-installed, restored, and then kept current through asynchronous data synchronization.
- The migration ends with final synchronization and role swap.



Migration patterns



Assisted save and restore stages



Infrastructure setup for assisted save and restore

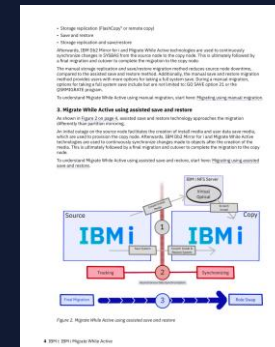
What must exist before media is created

Source and copy preparation

- Install the copy node at the same operating-system level as the source.
- Configure the DST IP interface on the copy node.
- Create a migration profile on both nodes, including SST-level setup.
- Establish SSH connectivity from source to the NFS server.

Staging and restore infrastructure

- Provide an NFS server on another IBM i system or an equivalent staging node.
- Verify enough space for install media and user data saves.
- Plan network boot and scratch-install procedures in advance.
- Use WRKMIGSTS and MWA GUI as the operator cockpit during restore and synchronization.



IBM Db2 Mirror for i

Primary: apop89
Secondary: apop90
migrate

- Dashboard
- Single Node Dashboard
- Home
- Security
- Setup
- Network Redundancy Groups
- Independent Auxiliary Storage Pools
- Replication List
- Object Tracking List
- Application Evaluation
- Serviceability
- Health Center

Setup Migrate While Active

Source and Copy Node Configuration

Specify the system information for each node in this Db2 Mirror pair.

Source APOP89

Host Name:

IP Address:

Communication Port:

Line Description:

Subnet Mask:

Gateway:

System Name:

Domain Name:

DNS Servers:

Copy APOP90

ⓘ The copy node will be powered off during the cloning process.

Host Name:

IP Address:

Communication Port:

Line Description:

Subnet Mask:

Gateway:

System Name:

Domain Name:

DNS Servers:

Service Tool Server Configuration

System Name:

IP Address:

Gateway:

Subnet Mask:

Resource Name:

>>

⏪ Previous
Next ⏩

Stage 2 - Data synchronization

IBM i Migrate While Active Guide

- Step 1**
Configure Migration Settings
- Step 2**
System Migration
- Step 3**
Continuous Data Synchronization
- Step 4**
Final Data Synchronization
- Step 5**
Complete Migration
- Step 6**
Exit Migration

Continuous Data Synchronization

Inactive
Synchronization paused Start

[View Data](#)

Synchronization Jobs

Mirrored Objects	Non-mirrored Objects	IFS Objects
Inactive Synchronization paused	Inactive Synchronization paused	Inactive Synchronization paused

OK

Tracking lists: OTL, LTL, and ITL

Why the synchronization phase is broader than simple journaling

OTL - object tracking list

- Tracks Db2 Mirror replication-eligible objects
- Managed by Db2 Mirror resynchronization
- Typical app-data change path

LTL - library tracking list

- Tracks eligible library objects that are not Db2 Mirror replication-eligible
- Typically synchronized with save/restore logic
- Examples include many library-based object types

ITL - integrated file system tracking list

- Tracks IFS objects in SYSBAS excluding /QSYS.LIB and /QDLS
- Important for large file-system migrations
- IFS behavior can dominate elapsed time

Operational implication

- Until final synchronization completes, the target is not a coherent production replacement
- This is why MWA is not HA/DR
- List sizing and estimation matter

Stage 3 - Cutover

IBM i Migrate While Active Guide

Step 1 ✓
Configure Migration Settings

Step 2 ✓
System Migration

Step 3 ✓
Continuous Data Synchronization

Step 4 ✓
Final Data Synchronization

Step 5
Complete Migration

Step 6
Exit Migration

Complete Migration

Review cutover time estimation, ineligible objects, and object replication errors. [View Data](#)

[Launch](#)

data synchronization, execute option 22

Confirm cutover

This action causes the copy node APOP75C to become the primary node and allows changes to its objects. Once this has completed, the Migrate While Active processing terminates and cannot be canceled.

Are you sure you want to cutover to the copy node APOP75C?

[No](#) [Yes](#)

[OK](#)

Assisted save and restore - strengths and constraints

Good fit when teams are comfortable with save/restore discipline

Strengths

- Well understood by IBM i administrators
- Uses explicit media and restore steps
- Good fit for organizations that want guided orchestration

Constraints

- Requires an initial outage
- More moving parts than partition mirroring
- NFS/media logistics must be managed carefully

Where it shines

- Predictable migration projects
- Teams already accustomed to save/restore operations
- Environments without a partition-mirroring fit

Message to sponsors

- Downtime is reduced, not eliminated
- Validation still matters after restore
- The final cutover remains the key event

05

MANUAL MIGRATION

The new flexible path introduced in 2025

Manual migration lets clients bring their own initial-copy strategy while MWA tracks changes and manages the final handover.

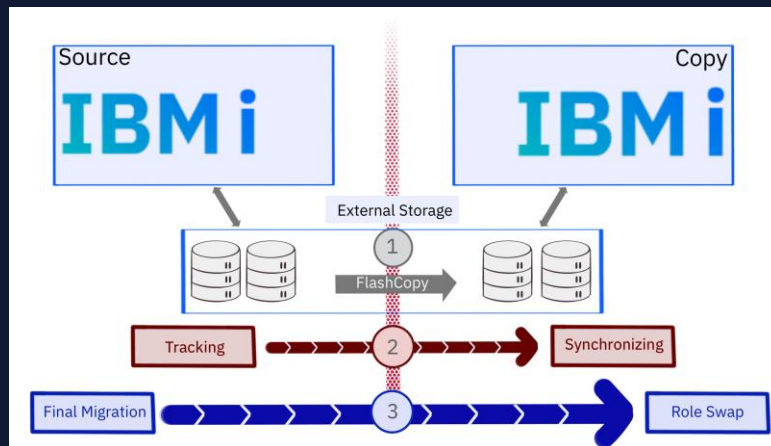


Manual migration variants

The copy-node foundation can be built in more than one way

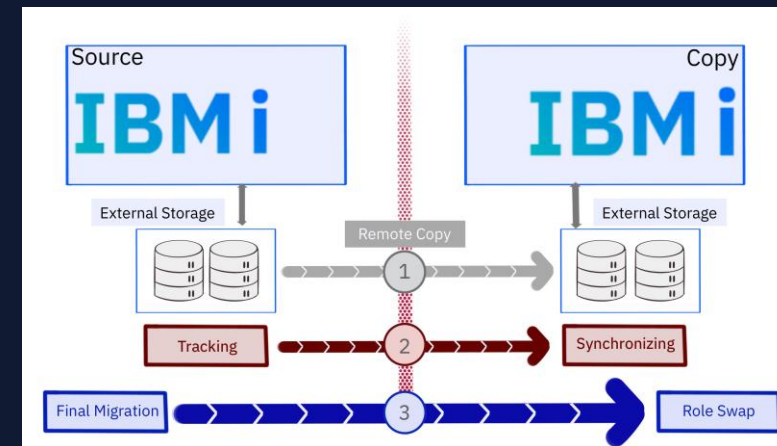
How the copy is created

- Storage replication such as FlashCopy or remote copy.
- Traditional save and restore.
- A hybrid of storage replication plus save/restore.



Why this matters

- You can reduce downtime compared with assisted save and restore.
- You gain more freedom in how to take a full system save.
- GO SAVE option 21 and QSRMIGRATE are possible paths.



The screenshot shows the 'Setup Wizard' for 'IBM Db2 Mirror for i'. The main heading is 'Configure New Pair' with the instruction 'Select the type of pair to configure'. There are two radio button options: 'Db2 Mirror' (unselected) and 'IBM i Migrate While Active' (selected). Below these are three 'Partition migration methods' options, each with a radio button: 'Partition Mirroring' (unselected), 'Assisted Save and Restore' (unselected), and 'Manual Migration' (selected). A 'Next' button is located at the bottom right of the wizard area.

IBM Db2 Mirror for i

Setup Wizard

Configure New Pair

Select the type of pair to configure

Db2 Mirror
IBM i continuous availability solution using real time synchronization between a 2 node pair

IBM i Migrate While Active
Migration of an IBM i node from one location to another while continuing to use the source node

Partition migration methods

Partition Mirroring
Host-based storage replication over a network connection

Assisted Save and Restore
Assisted migration using save media (e.g., Optical, Tape)

Manual Migration
User completed system migration while Migrate While Active tracks data changes. System migration options include storage replication or save and restore

[Next](#)

Object Synchronization Setup

Configure the system and network details to use on the copy node during object level synchronization

Import node details

Import system name and IP addresses from the copy node

Copy node configuration details for object synchronization [View Source \(MICK72\)](#)

System Name: ×

Host Name: ×

Domain Name: ×

DNS Servers: ×
 ×
 ×

Copy node network details for object synchronization [View Source \(MICK72\)](#)

Copy node IP address

IP Address	Line Description	Communication Port/MAC Address	Subnet Mask	Gateway
- Select -				

Replication network details for object synchronization [Refresh](#)

Select the replication IP communication type

RDMA Over TCP/IP
 Use RDMA protocols over TCP/IP. Encryption of RDMA over TCP/IP is not available. To secure the network traffic, it is highly recommended to use another form of encryption such as a VPN.

RDMA Over Converged Ethernet (RoCE) [?](#)
 Use RoCE adapters. Encryption is available on RoCE adapters that support cryptography

Encrypted RDMA required [?](#)

Search: Delete Add IP Address Pair

Source IP Address	Copy IP Address	RDMA Type	Status
No additional IP addresses added.			

<< < > >>

Object Synchronization Setup

Configure the system and network details to use on the copy node during object level synchronization

Select manual migration type
 Select the manual migration method being used to establish the copy node

Storage Replication

Select this option to use storage server technology to copy the volumes of the source node to the copy node

Save/Restore

Select this option to take a full system save of the source node while in restricted state and then use the full system save to scratch install the copy node

Storage Replication and Save/Restore

Select this option to use storage server technology to copy the volumes of the source node to a temporary node and then take a full system save of the temporary node. The full system save is then used to scratch install the copy node.

Migration options

Source node state
 Select the state you want the source node to be in when you use storage server technology to copy the source node

Quiesced (Warm Clone)
 Select this option to quiesce the source node and start tracking data changes before you initiate the FlashCopy or end storage replication with consistency

Set quiesce timeout to: minutes

Powered Off (Cold Clone)
 Select this option to power off the source node before you initiate the FlashCopy or end storage replication with consistency. Tracking data changes starts after it powers on.



Stage 1 - System migration – warm copy

Manual steps required for warm cloning

Select the storage copy type and then perform the manual steps for warm clone in sequential order

i The manual steps can start after the object tracking preparation and readiness checks are complete. Click View Steps to preview the manual steps. ✕

Select the storage copy type

[Preview Steps](#)

FlashCopy Remote Copy

Save

Complete these steps in sequential order

1 of 5

Power off copy node

Power off the copy node to continue

[View Documentation](#) 🔗

I have powered off the copy node

2 of 5

Quiesce source node and initiate FlashCopy

Initiate the FlashCopy while the source node is in a quiesced state

[View Documentation](#) 🔗

i The source node has up to 5 minutes to reach a quiesced state. After reaching the quiesced state, you have up to 60 minutes to initiate the FlashCopy. ✕

Quiesce and start tracking

Start Quiesce

Cancel

Initiate FlashCopy

Initiate the FlashCopy to continue

I have initiated the FlashCopy

3 of 5

End quiesce on source node

End the quiesce on the source node after initiating the FlashCopy

[View Documentation](#) 🔗

End Quiesce

Manual steps required for warm cloning

Select the storage copy type and then perform the manual steps for warm clone in sequential order

i The manual steps can start after the object tracking preparation and readiness checks are complete. Click View Steps to preview the manual steps. ✕

Select the storage copy type

[Preview Steps](#)

FlashCopy Remote Copy

Save

Complete these steps in sequential order

1 of 6

Power off copy node

Power off the copy node to continue

[View Documentation](#) 🔗

I have powered off the copy node

2 of 6

Start storage replication

Start the storage replication to continue

[View Documentation](#) 🔗

I have started the storage replication

3 of 6

Quiesce source node and end storage replication with consistency

After the storage replication is in sync, end the storage replication with consistency while the source node is in a quiesced state

[View Documentation](#) 🔗

i The source node has up to 5 minutes to reach a quiesced state. After reaching the quiesced state, you have up to 60 minutes to end the storage replication with consistency. ✕

Quiesce and start tracking

Start Quiesce

Cancel

End storage replication with consistency

End storage replication with consistency and ensure the target volumes are consistent copies to continue

I have ended the storage replication

Manual steps required for warm cloning

Select the storage copy type and then perform the manual steps for warm clone in sequential order

i The manual steps can start after the object tracking preparation and readiness checks are complete. Click [View Steps](#) to preview the manual steps. ✕

1 of 10

Power off temporary node

[View Documentation](#) 🔗

Power off the temporary node to continue

I have powered off the temporary node

2 of 10

Quiesce source node and initiate FlashCopy

[View Documentation](#) 🔗

Initiate the FlashCopy while the source node is in a quiesced state

i The source node has up to 5 minutes to reach a quiesced state. After reaching the quiesced state, you have up to 60 minutes to initiate the FlashCopy. ✕

Quiesce and start tracking

[Start Quiesce](#)

[Cancel](#)

Initiate FlashCopy

Initiate the FlashCopy to continue

I have initiated the FlashCopy

3 of 10

End quiesce on source node

[View Documentation](#) 🔗

End the quiesce on the source node after initiating the FlashCopy

[End Quiesce](#)

4 of 10

Power on temporary node

[View Documentation](#) 🔗

Power on the temporary node to continue

I have powered on the temporary node

Stage 2 - Data synchronization

IBM i Migrate While Active Guide

Step 1
Configure Migration Settings ✓

Step 2
System Migration ✓

Step 3
Continuous Data Synchronization

Step 4
Final Data Synchronization

Step 5
Complete Migration

Step 6
Exit Migration

Continuous Data Synchronization

Active
Synchronization in progress

[View Data](#)

Pause synchronization

Pause Finalize

Synchronization Jobs

Mirrored Objects	Non-mirrored Objects	IFS Objects
☀ Active Synchronization in progress 060132/QSYS/QMRDBESYNC	☀ Active Synchronization in progress 060134/QSYS/QMRDBNSYNC	☀ Active Synchronization in progress 060135/QSYS/QMRDBISYNC

OK

SEIDOR Objects and migration semantics

Manual migration makes object behavior impossible to ignore

Db2 Mirror replication-eligible

- Typical application data path
- Handled by QMRDBESYNC
- Replicates tracked actions and changed rows

Non-Db2-Mirror library objects

- Tracked on the library tracking list (LTL)
- Handled by QMRDBNSYNC
- Moved by save and restore logic

IFS objects

- Most SYSBAS IFS content is supported
- IFS behavior often dominates elapsed time
- Parallel processing improvements arrived in 2025

Pseudo objects

- System values, security attributes, env vars, function usage IDs
- Need deliberate review because not all behave like normal objects
- Some items are not migrated during the initial system migration

Decision frame

Stage 3 - Cutover

Data Synchronization - Cutover Estimate

*SYSBAS

Synchronizing

Estimated On: 11/20/2024, 4:05:03 PM

Migration Guide

Click to view the migration process

161

Cutover Estimate

Total size of objects remaining:

10,646.1 MB

Summary

Current Calculated Transfer Speed:

878.79 MB per second

All

Estimate of the outage needed to complete the final synchronization for IBM i Migrate While Active:

3.1 minutes

Library Objects

Rebuild

IFS Objects

Synchronized: 809

Estimation Details

Synchronizing: 387

Ineligible: 1

Deferred: 4

Failed: 2

Tracking List ↑↓	Estimated Total Time (minutes) ↑↓	Object Count ↑↓	Data Size (MB) ↑↓
<input type="text" value="Filter"/>	<input type="text" value="Filter"/>	<input type="text" value="Filter"/>	<input type="text" value="Filter"/>
OBJECT TRACKING LIST	3.11	150	9505.40
LIBRARY TRACKING LIST	0.17	10	1140.72
IFS TRACKING LIST	0.02	1	0.00

06

CLOUD TARGETS AND REAL-WORLD
DEPLOYMENT

**MWA is increasingly part of
modernization stories, not just
hardware refreshes**

Congress edition



SEIDOR PowerVS target pattern

From on-prem IBM i to cloud-based Power infrastructure with limited outage

Why it fits

- Source can stay productive during migration
- Controlled cutover into cloud infrastructure
- Good modernization story for aging on-prem systems

Requirements

- Target PowerVS image on a supported IBM i level
- Reliable IP connectivity to the region
- Destination hardware / release compatibility

Provisioning model

- Cloud-Optimized Root images simplify initial setup
- Custom images and ESS media remain options

Architect message

- Cloud does not remove the need for validation
- Version support still governs the path
- Network quality directly affects elapsed time

Decision frame

SEIDOR Remote location and IP4G considerations

MWA is a migration framework that travels well, but connectivity still rules

Remote data center

- Remote-location migration is a valid use case
- Long distance is acceptable because synchronization is asynchronous

IP4G

- IBM Power for Google Cloud destination
- Source must still meet MWA minimum release requirements

What changes

- Latency becomes a planning variable
- Firewall and routing work become first-class tasks
- Image/catalog choices matter more

What does not change

- Need a copy node, validation, and final cutover
- Need to understand object classes and list behavior
- Need short, disciplined production outage at the end

07

PROJECT CONTROL

Risk, validation, and the migration runbook

A technically correct migration can still fail without clear governance, cutover criteria, and ownership.



Most failures are project-design issues, not product defects

Unsupported scope

- Encrypted ASPs or IASP assumptions in partition mirroring
- Old source releases for cloud targets
- Missing software / PTF prerequisites

Hidden elapsed-time drivers

- Large IFS trees
- Non-Db2-Mirror object volume
- Underestimated network transfer windows

Cutover-night chaos

- Unclear roles and command ownership
- No pass/fail validation checklist
- Business sign-off criteria missing

Bad expectation setting

- Treating MWA as HA/DR
- Assuming instant rollback after swap
- Skipping rehearsal because test mode exists



SQL services and compare tools

MWA is easier to operate when status is measured, not guessed

– [CHANGE_MIGRATION_MANAGER procedure](#)

The CHANGE_MIGRATION_MANAGER procedure is used to manage actions performed by the Migration Manager within IBM i Migrate While Active.

– [DATA_COMPRESSION_FACTOR table](#)

The DATA_COMPRESSION_FACTOR table contains the values used by IBM i Migrate While Active for estimating compression of objects.

– [ESTIMATE_FINAL_SYNCHRONIZATION_TIME table function](#)

The ESTIMATE_FINAL_SYNCHRONIZATION_TIME table function provides an estimate of the outage needed to complete the final synchronization for IBM i Migrate While Active.

– [IFS_MIGRATION_LIST view](#)

The IFS_MIGRATION_LIST view returns information from the integrated file system tracking list (ITL).

– [LIBRARY_MIGRATION_LIST view](#)

The LIBRARY_MIGRATION_LIST view returns information from the library tracking list (LTL) and object tracking list (OTL).

– [MIGRATED_IFS view](#)

The MIGRATED_IFS view returns a list of the integrated file system objects that have been successfully migrated to the copy node by IBM i Migrate While Active.

– [MIGRATED_LIBRARY view](#)

The MIGRATED_LIBRARY view returns a list of library-based objects that have been successfully migrated to the copy node by IBM i Migrate While Active.

– [MIGRATION_MANAGER_INFO view](#)

The MIGRATION_MANAGER_INFO view returns information about the IBM i Migrate While Active environment.

– [MIGRATION_NETWORK_BANDWIDTH table function](#)

The MIGRATION_NETWORK_BANDWIDTH table function measures bandwidth between the source and copy nodes for IBM i Migrate While Active.

Product services

These services can be used to configure and review details for the Db2 Mirror environment used by IBM i Migrate While Active.

– [CHANGE_MIRROR procedure](#)

– [CHANGE_MIRROR_OBJECTCONNECT procedure](#)

– [MIRROR_CLUSTER_INFO view](#)

– [MIRROR_INFO view](#)

– [SET_MIRROR_CLUSTER procedure](#)

– [SETUP_MIRROR procedure](#)

Serviceability services

These services can be used to configure and work with serviceability details for the Migrate While Active environment.

– [MIRROR_COMPARE_LIBRARY procedure](#)

– [MIRROR_COMPARE_NODE procedure](#)

– [MIRROR_COMPARE_OBJECT table function](#)

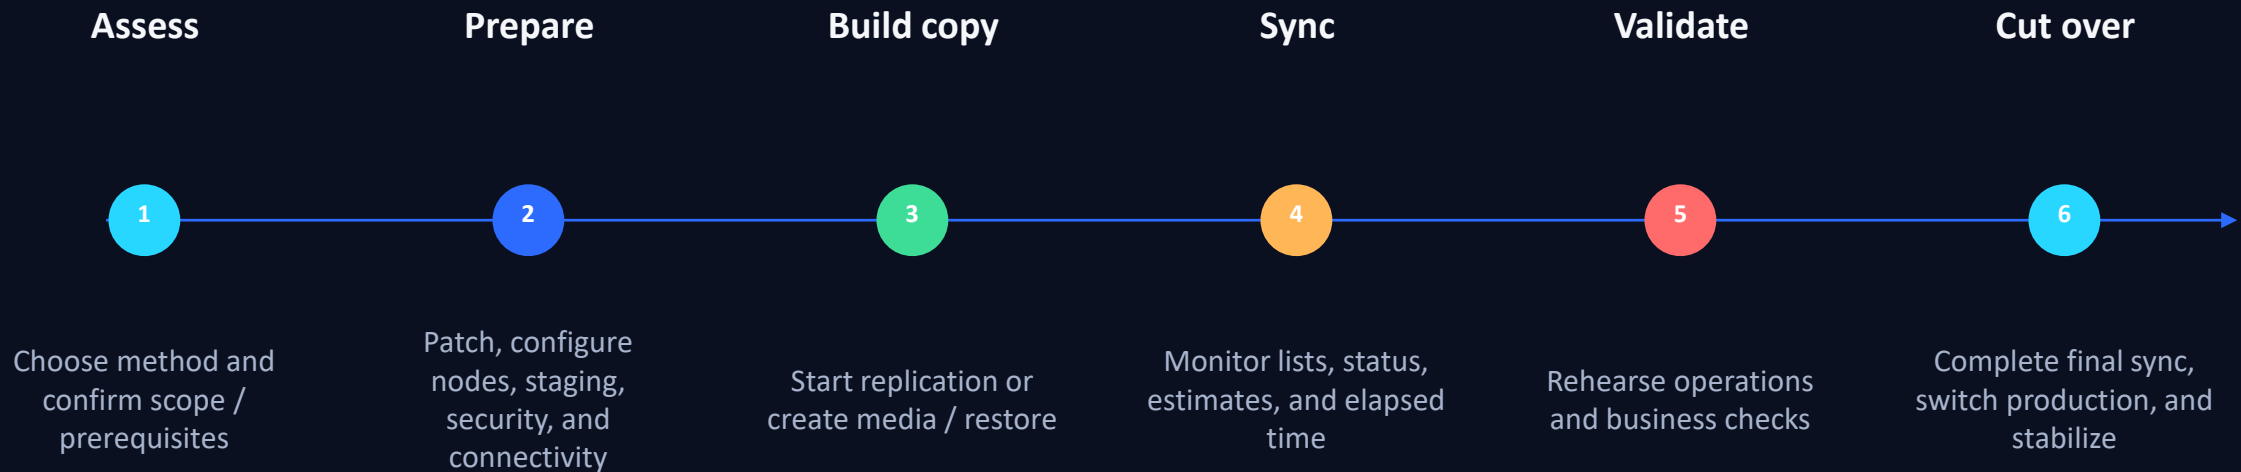
– [MIRROR_DISPLAY_JOURNAL table function](#)

– [MIRROR_FLIGHT_RECORDER_INFO view](#)

– [SET_MIRROR_FLIGHT_RECORDER_INFO procedure](#)

A reference project plan

One way to turn the technology into a migration program



Reference timeline - adapt duration to data volume, link speed, and validation depth.

Lifecycle

When to choose each method

The selection is mainly about downtime tolerance, storage options, and operational control

Choose partition mirroring

- Need zero initial outage
- Can keep scope to SYSBAS
- Want a built-in test mode

Choose assisted save/restore

- Traditional save/restore is acceptable
- Need IBM-guided orchestration
- Can stage media through NFS

Choose manual migration

- Existing FlashCopy/remote-copy process exists
- Need flexible system-save options
- Want to tailor the initial copy method

Avoid false assumptions

- MWA is not long-term HA/DR
- IASP and encrypted-ASP cases need special handling
- Reverse replication after cutover is not the design goal

- MWA is best understood as a migration orchestrator that minimizes downtime, not as a standing HA/DR architecture.
- Partition mirroring is the premium low-disruption path when SYSBAS scope and prerequisites fit.
- Assisted save and restore remains a strong option for disciplined save/restore environments.
- Manual migration broadens the portfolio with storage-replication and hybrid strategies.
- The best migrations are won by preparation, estimation, validation, and stakeholder alignment.

Questions?

IBM i Migrate While Active
methods, cutover discipline, and
modernization paths

Thank you