

# IBM i Migrate While Active

Henry Vo

Abhishek Anupam

Tanobir Ahmed

Ben Huntsman

Marcin Jaworski

Tim Simon



**IBM Power**

**Cloud**





IBM Redbooks

February 2026

**Note:** Before using this information and the product it supports, read the information in “Notices” on page v.

**First Edition (February 2026)**

This edition applies to IBM i Version 7.4, 7.5 and 7.6  
Power10, IBM Power11  
Power Virtual Server

© Copyright International Business Machines Corporation 2026. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

# Contents

<b>Notices</b> .....	v
Trademarks .....	vi
<b>Preface</b> .....	vii
Authors .....	vii
Now you can become a published author, too! .....	viii
Comments welcome .....	ix
Stay connected to IBM Redbooks .....	ix
<b>Chapter 1. Overview of Migrate While Active on IBM i</b> .....	1
1.1 Introduction to IBM i Migrate While Active .....	2
1.2 Migrating an IBM i System Manually versus using Migrate While Active .....	3
1.2.1 Manual migration .....	3
1.2.2 Migrate While Active (MWA) .....	4
1.3 Relationship with IBM Db2 Mirror for i .....	5
1.4 Supported Environments .....	7
1.4.1 Enhancement Timeline and PTF-Level Mapping .....	7
1.4.2 Cloud Support .....	8
1.5 Licensing considerations .....	10
1.6 Data transfer options .....	12
1.7 IBM i Object Operations: Changes in IBM i 7.5 and IBM i 7.6 .....	14
1.7.1 MWA Enhancements: IBM i 7.5 and IBM i 7.6 .....	14
1.7.2 Performance & Wait Analysis (Granular Visibility in 7.6) .....	15
1.7.3 Automated Object Management via SQL Services .....	16
1.7.4 Object Security & Operational Guardrails .....	16
1.7.5 Summary of Enhancements .....	17
<b>Chapter 2. Migration Planning and Prerequisites for Minimizing Downtime</b> .....	19
2.1 Hardware and Software requirement .....	20
2.2 Minimizing Downtime During IBM i LPAR Migrations .....	23
2.2.1 Migration methods and downtime characteristics .....	23
2.2.2 Minimizing Downtime with Partition Mirroring .....	24
2.2.3 Minimizing Downtime with Manual Migration .....	25
2.2.4 Managing Downtime with Assisted Save and Restore .....	27
2.2.5 Continuous Synchronization and Final Cutover .....	29
2.3 Determining Migration Path and Prerequisites .....	30
2.3.1 Migration paths you can choose .....	30
2.3.2 How and why to choose a migration method .....	31
2.3.3 Method-specific prerequisites and validation .....	33
2.4 Enhanced Db2 Mirror GUI .....	35
2.4.1 Recent GUI Enhancements (2024–2025) .....	36
2.4.2 Core GUI Functions Supporting Migration .....	36
2.4.3 Migration Methods Supported Through the GUI .....	37
2.5 Installing RSYNC using the Open Source Package Management interface .....	38
2.6 Security highlights of IBM i MWA .....	40
2.6.1 Secure Administration & Control Planes .....	40
2.6.2 Protect the Data Paths .....	41
2.6.3 Control Identities, Authority, and Keys .....	43
2.6.4 Cloud Connectivity & Network Isolation for IBM Power Virtual Server .....	45

<b>Chapter 3. Migration method</b> .....	47
3.1 Partition Mirroring .....	48
3.2 Assisted Save and Restore .....	49
3.2.1 Save and restore use cases .....	50
3.2.2 Anatomy of save and restore migration .....	51
3.2.3 Automated Transfer Operations in Migrate While Active .....	53
3.3 Manual Migration .....	55
3.3.1 Clone Types for Storage Replication .....	58
3.3.2 Method Selection Guide .....	58
3.3.3 Preparation Summary .....	58
3.3.4 Configuration via Db2 Mirror GUI .....	59
<b>Chapter 4. SYSBAS Replication and IASP Migration Concepts and Procedures</b> . . . .	61
4.1 SYSBAS Replication and Object Migration .....	62
4.1.1 Object Classification in IBM i Migrate While Active .....	62
4.1.2 Synchronization Jobs .....	63
4.1.3 Tracking Lists .....	64
4.2 Migrating an IASP During Migrate While Active .....	64
4.2.1 Migration Concept .....	65
4.2.2 PowerHA Requirements, Geographic Mirroring Setup, Network & Bandwidth . . .	65
4.2.3 IASP Migration Scenarios .....	67
4.2.4 Migration Timeline and Coordination with Migrate While Active .....	68
4.2.5 Step-by-Step Procedure .....	69
4.2.6 Test, Validation, Management, and Troubleshooting .....	71
<b>Related publications</b> .....	75
IBM Redbooks .....	75
Other Publications .....	75
Help from IBM .....	76

# Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

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

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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

## COPYRIGHT LICENSE:

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

# Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at <https://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

AIX®	IBM Cloud®	Redbooks®
Db2®	IBM Partner Plus®	Redbooks (logo)  ®
DS8000®	Passport Advantage®	System z®
FlashCopy®	PowerHA®	SystemMirror®
IBM®	PowerVM®	z/OS®

The following terms are trademarks of other companies:

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Red Hat, RHCE, are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

# Preface

This book provides practical guidance for planning and executing IBM i migrations with minimal business disruption. It brings together the core concepts, technologies, and processes required to move IBM i workloads across new systems, data centers, and cloud environments. Topics include assessing current and target configurations, preparing for migration, leveraging IBM i Migrate While Active (MWA), and applying best practices to ensure a smooth and reliable transition.

The content is intended for IBM i administrators, architects, and technical professionals who are responsible for maintaining system availability while enabling infrastructure modernization. By using the methods and recommendations described in this book, organizations can confidently migrate IBM i environments while preserving service continuity and supporting evolving business needs.

## Authors

This book was produced by a team of specialists from around the world working at IBM Redbooks, Remote Center.

**Henry Vo** is an accomplished IT professional at IBM® with experience of cross-functional teams members and delivering enterprise-scale innovation. He is widely recognized for his deep expertise as a test specialist, with strengths in technical analysis, test architecture, and strategic planning. Starting his career as a backend programmer on z/OS® systems, Henry has contributed to mission-critical solutions across diverse industries. Over the years, he has successfully guided teams through complex technical challenges, fostering collaboration and driving impactful results. In his current role as IBM Redbooks® Project Leader, Henry spearheads the development of advanced technical content that bridges emerging technologies such as hybrid cloud, automation, and AI-ready infrastructure on IBM Power11 with practical implementation for clients worldwide. His leadership is defined by a commitment to excellence, mentorship, and a passion for translating innovation into real-world value.

**Abhishek Anupam** is a Quality Analyst and a Test Developer with over a decade of experience in qualifying enterprise server and data storage products. He earned his Bachelor of Engineering degree in Bangalore, India in 2015 and has since worked with leading companies in the industry. Abhishek has expertise in IBM i, having previously worked on qualifying IBM i over PowerVS and currently contributing to System Verification Test for IBM i OS. Passionate about technology, he enjoys reading and writing technical blogs. Outside of work, Abhishek loves spending time with his kids and exploring new places.

**Tanobir Ahmed** is an IBM i and Power Systems specialist based in Adelaide, South Australia. He began his IBM i career in 2007 with an IBM Business Partner in Bangladesh and later served as a Senior Technical Lead at Saud Bahwan Group in Oman and at Dynatrade in the UAE, supporting SAP workloads and in-house applications running on IBM i. His responsibilities included IBM i & AIX administration, VIOS and PowerVM®, DS8000 and V9000 storage, FlashCopy®, Copy Services Toolkit (now iASP Manager), and PowerHA®. He subsequently worked with Saudi British Bank in Riyadh, Saudi Arabia (HSBC Group) on core banking platforms running on IBM Power and IBM i, focusing on IBM i system administration, VIOS, IBM Power Systems evergreen lifecycle management, high availability, multi-site storage replication with iASP Manager and CSM, and Safeguarded Copy for cyber resiliency. Tanobir is the Founder of Code24 Pty Ltd, an IBM Partner Plus® (Silver) company in

Australia, where he works to expand IBM Power, AIX, Linux, and IBM i solutions in the Australian market. With nearly 18 years of hands-on experience across IBM i systems administration, migrations, upgrade projects, HA/DR architecture, enterprise storage, and automation—and as a Red Hat Certified Engineer (RHCE)—this is his first IBM Redbooks residency. He is excited to collaborate with global experts, contribute to the Migrate While Active residency program, and, most importantly, learn from the highly experienced team involved in this project.

**Ben Huntsman** is a Senior Consultant and Solutions Architect with Peller Technologies in the United States. He has over 15 years of experience working extensively with IBM solutions including AIX®, IBM i, and Linux on Power, as well as IBM storage systems including FlashSystem, Spectrum Virtualize, and DS8000®. In his current role he has been focusing on helping enterprise customers move large IBM i workloads to the cloud, including the IBM Power for Google Cloud (IP4G), a Google Cloud Marketplace offering operated by Peller. He also has worked extensively with Intel-based workloads and operating systems, and works with customers to strategize and gain a holistic view of their IT footprint and the best ways to leverage existing investments in a modern landscape.

**Marcin Jaworski** is an IBM i Shared Infrastructure Specialist at the Kyndryl Collaborative Poland in Wroclaw. He started his IBM i career in 1997 in a Polish supermarket as an AS/400 administrator, working with OS/400 version V3R7. He then spent 10 years in the banking sector as an OS/400 System Administrator. In 2012, he joined IBM as an IBM i Specialist. Currently he manages over 20 different Power servers and hundreds of IBM i partitions for multiple customers, acting both as an infrastructure specialist and as third level support for IBM i. Throughout his nearly 30 year career, he has been involved in numerous small and large multi-platform projects related to IBM i infrastructure, including storage and software-based migrations, system and firmware upgrades, high availability, consulting, and Power server configurations. His main areas of interest within the IBM i domain are migrations, high availability, and virtualization. He was a speaker at the Common Poland conference, and this is his first Redbooks residency.

**Tim Simon** is an IBM Redbooks Project Leader in Tulsa, Oklahoma, USA. He has over 40 years of experience with IBM, primarily in a technical sales role working with customers to help them create IBM solutions to solve their business problems. He holds a BS degree in Math from Towson University in Maryland. He has worked with many IBM products and has extensive experience creating customer solutions using IBM Power, IBM Storage, and IBM System z® throughout his career.

## Now you can become a published author, too!

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

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

[ibm.com/redbooks/residencies.html](http://ibm.com/redbooks/residencies.html)

## Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks form found at:

[ibm.com/redbooks](http://ibm.com/redbooks)

- ▶ Send your comments in an email to:

[redbooks@us.ibm.com](mailto:redbooks@us.ibm.com)

- ▶ Mail your comments to:

IBM Corporation, IBM Redbooks  
Dept. HYTD Mail Station P099  
2455 South Road  
Poughkeepsie, NY 12601-5400

## Stay connected to IBM Redbooks

- ▶ Find us on LinkedIn:

<https://www.linkedin.com/groups/2130806>

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

<https://www.redbooks.ibm.com/subscribe>

- ▶ Stay current on recent Redbooks publications with RSS Feeds:

<https://www.redbooks.ibm.com/rss.html>





# Overview of Migrate While Active on IBM i

This chapter introduces Migrate While Active, a strategic IBM i function that enables system migrations with minimal application downtime. Migrate While Active allows organizations to move workloads to new environments while maintaining the availability of critical business applications. This capability supports a variety of modernization scenarios, including migrating IBM i partitions to new hardware, integrating cloud-based infrastructure with reduced service disruption, and enhancing disaster recovery implementations.

In addition to providing an overview of how the technology functions, the chapter outlines the common foundational elements relevant to every Migrate While Active implementation. The chapter concludes with a summary of the hardware and software prerequisites required to prepare an IBM i environment for a successful Migrate While Active deployment.

This chapter will cover the following topics:

- ▶ “Introduction to IBM i Migrate While Active”
- ▶ “Migrating an IBM i System Manually versus using Migrate While Active”
- ▶ “Relationship with IBM Db2 Mirror for i”
- ▶ “Supported Environments”
- ▶ “Licensing considerations”
- ▶ “Data transfer options”
- ▶ “IBM i Object Operations: Changes in IBM i 7.5 and IBM i 7.6”

## 1.1 Introduction to IBM i Migrate While Active

Modern enterprises increasingly expect their core business systems to remain accessible with minimal planned downtime. For clients running IBM i, this expectation often becomes a key consideration when migrating workloads to new hardware, upgrading partitions, or moving to modern storage infrastructures. Traditional save-and-restore migration techniques, while reliable, require extended outage windows that may be unacceptable for production environments operating under strict service-level objectives.

IBM i Migrate While Active (MWA) addresses this challenge by providing a controlled and repeatable method for migrating an IBM i environment while ensuring that the source system remains operational for the majority of the migration process. Built on proven IBM i technologies such as journaling, remote journaling, and object replication, MWA enables organizations to move to a new system with near-continuous availability, reducing the final cutover outage to minutes instead of hours.

MWA is designed to support migrations across a range of scenarios, including movement to new Power servers, transitions between LPARs, and consolidation or reorganization of workloads using Independent Auxiliary Storage Pools (IASPs). Because MWA uses a synchronization model, the target environment is continuously updated as changes occur on the source system. This allows administrators to validate and test the target partition long before business operations are redirected—significantly reducing risk and enabling a more predictable migration window.

From an architectural perspective, MWA establishes a replication relationship between the source and target systems. After the initial copy phase, the technology captures ongoing changes at the object and journal entry level and transmits these updates to the target system. The result is a target environment that closely mirrors the active production system, synchronized in near real-time. When migration is ready to proceed, administrators perform a planned, short outage during which final updates are applied and the target partition is brought online as the new production system.

This approach provides several benefits. First, it minimizes the operational impact on the business. Production workloads continue to run during the bulk of the migration process, allowing the organization to maintain normal operations. Second, the ability to test the target system before the final cutover significantly reduces cutover-day risk. Third, MWA can serve as a transitional high-availability mechanism during the migration period, providing a degree of resilience that is not available with traditional save-and-restore approaches.

MWA is not intended to replace full high-availability solutions or long-term replication strategies. Instead, it focuses specifically on migration events—situations in which a planned and predictable transition is required. For customers who need a low-risk, low-downtime method of moving IBM i workloads to new infrastructure, MWA provides an effective, efficient, and repeatable solution.

As organizations modernize their IBM i environments—adopting newer generations of IBM Power systems, integrating flash-based storage, or consolidating workloads—MWA helps them achieve these goals without compromising availability. By combining automation with the inherent reliability of IBM i, Migrate While Active enables smoother transitions, better planning, and higher confidence throughout the migration lifecycle.

## 1.2 Migrating an IBM i System Manually versus using Migrate While Active

This section compares using Migrate While Active to the more traditional migration strategy using a save-and-restore methodology. It provides a high level description of both options and discusses some of the reasons that you might choose to do a manual migration versus using Migrate While Active.

### 1.2.1 Manual migration

A manual migration of an IBM i environment follows a traditional save-and-restore methodology in which administrators create a complete backup of the source system and then restore that backup onto new hardware or a newly provisioned partition. After the restore operation is completed, the administrator reconfigures system attributes such as network settings, device descriptions, and system values to align the target environment with the original system's operational requirements. Once the foundational configuration is in place, the environment undergoes functional and application testing to verify that workloads perform as expected before the new system is placed into production. This approach, often referred to as the classic "save-restore cutover," provides a predictable and well-understood migration path but typically requires a dedicated outage window during which production processing is halted.

The typical steps involved in a manual migration are:

1. Do a Full System Save (SAVSYS, SAVLIB \*NONSYS, SAVDLO, SAV).
2. Record and export configuration (network attributes, TCP/IP, device descriptions, subsystem configs, SBSDB, job descriptions, etc.).
3. Prepare the target system
  - Install same or newer IBM i OS version
  - Apply PTFs
  - Configure storage and LPAR resources
4. Restore data and objects (RSTLIB, RSTOBJ, RSTUSRPRF, RSTAUT, etc.)
5. Validate and test
6. Perform the cutover (production downtime required)
  - Take the production system down.
  - Do a final save and restore data to new system.
7. Go live

Manual migration offers several advantages that make it a practical and flexible option for many IBM i environments. Because the process is entirely administrator-driven, it provides complete control over the sequence of tasks, the timing of each step, and the specific components included in the migration. This level of control can be especially beneficial when the migration coincides with broader system restructuring, such as redesigning library organization, modifying IASP layouts, updating storage architectures, or changing network configurations. Manual migration also avoids additional licensing costs, making it suitable for organizations that must manage expenses carefully or that operate in environments where features such as Live Partition Mobility or PowerHA are not available. However, manual migration also carries several disadvantages that organizations must consider. The approach requires a production outage that can extend for several hours, particularly on large systems, and the overall process can be operationally complex. The extensive manual steps increase the potential for human error, and some system objects require precise restore sequencing to ensure proper application behavior. Administrators must also plan carefully for data and configuration synchronization, especially

when coordinating changes that occur shortly before the cutover. As system size and complexity grow, these factors can lengthen the required outage window and increase the effort needed to ensure a successful migration.

Manual migrations are best when:

- ▶ Downtime is acceptable
- ▶ You're rearchitecting the system
- ▶ Hardware gen changes prevent Live Partition Mobility (LPM)
- ▶ You do not have IASPs or HA replication tooling available

## 1.2.2 Migrate While Active (MWA)

In contrast, Migrate While Active (MWA) performs the migration using a replication-based model that keeps the source and target systems synchronized while production workloads continue to run. Instead of restoring large volumes of data during a downtime window, MWA performs an initial copy of libraries, objects, and configuration elements to the target environment and then begins incremental replication to maintain near-real-time alignment. Because the target system is continuously updated with changes occurring on the source, administrators can validate application behavior, confirm configuration settings, and ensure operational readiness before any cutover takes place. This significantly reduces migration-day risk, as most validation and troubleshooting occur prior to the final outage.

MWA requires only a brief interruption during the final stage of the migration. When administrators are ready to move production workloads, they schedule a short outage to quiesce the source system and allow MWA to apply the last set of changes to the target. Once synchronization completes, the target system is brought online as the new production environment. This approach minimizes business disruption and shortens the overall cutover window from potentially several hours to only a few minutes in most scenarios.

While manual migration is fully administrator-driven, MWA automates much of the data movement and eliminates many of the manual restore steps associated with the traditional method. There is no need to sequence restores, rebuild complex configurations from exported settings, or coordinate multiple save operations near the cutover. MWA also reduces the risk of human error because it uses IBM i's integrated replication capabilities rather than relying on manual orchestration of each migration task.

MWA is particularly valuable when managing environments where downtime must be kept to an absolute minimum. By enabling organizations to migrate systems with ongoing replication, MWA supports a smoother transition to new hardware or new partitions without extended outages. It is well suited for production workloads with tight service-level objectives or for customers who require multiple test cycles before final cutover. However, unlike manual migration, MWA is not designed for extensive system redesign or large-scale restructuring, as it focuses on replicating the existing environment rather than rebuilding it.

In summary, manual migration remains a practical solution when downtime is acceptable, rearchitecture is planned, or replication technology is unavailable. Migrate While Active, however, provides a low-downtime, lower-risk alternative by maintaining continuous synchronization between the source and target systems and limiting the outage window to only the final switch. Together, these methods offer flexibility so organizations can choose the approach that best aligns with their technical, operational, and availability requirements.

## 1.3 Relationship with IBM Db2 Mirror for i

IBM Migrate While Active (MWA) builds on the foundational technologies originally delivered with Db2 Mirror for i. Although the two functions address different business requirements—Db2 Mirror providing continuous availability and MWA enabling low-downtime system migrations—they share common architectural underpinnings. MWA extends the internal replication framework introduced with Db2 Mirror, adapting it to support controlled, migration-oriented workload movement rather than active-active database processing. Understanding how MWA exploits and augments this replication infrastructure is critical when planning modern IBM i migration strategies, including migrations to IBM Power servers and transitions to IBM Power Virtual Server environments.

IBM i Migrate While Active is implemented as part of the IBM Db2 Mirror for i licensed program (5770DBM) and builds directly on its underlying tracking and synchronization technologies. Although it shares infrastructure with Db2 Mirror, Migrate While Active is an independent capability delivered as 5770DBM Option 2 (Migrate While Active).

**Important:** If a system is already configured with Db2 Mirror Option 1 (Db2 Mirror Enablement), that system cannot also use Migrate While Active for migration purposes.

### Shared Architectural Foundations

MWA directly inherits the core replication engine originally designed for Db2 Mirror for i. This inheritance gives MWA access to the same internal OS mechanisms for tracking data and object changes in real time. Db2 Mirror originally implemented deep hooks into the IBM i operating system to detect changes across data, eligible objects, and security attributes, creating the infrastructure necessary for synchronous data replication between active nodes. MWA reuses these capabilities but re-purposes them for asynchronous, migration oriented replication rather than HA-driven synchronization.

### Transformation of the Db2 Mirror Node Model

Db2 Mirror for i introduced a dual-node architecture, where two IBM i partitions remain synchronized through extremely low-latency RDMA (RoCE) connections. This architectural model is retained but reinterpreted in MWA. Instead of mirroring two active production systems, MWA employs the same node pairing concept to establish a source system and a migration copy node. The copy node is provisioned with a matching IBM i release, identical PTF levels, and synchronized system attributes requirements that mirror Db2 Mirror's standards for node equivalence.

While Db2 Mirror maintains synchronous, sub-millisecond coherence, MWA applies these mechanisms to maintain ongoing asynchronous synchronization between the source system and the migration target until the final cutover event. This shift from synchronous to asynchronous behavior is one of the fundamental differences between the two technologies.

### Expansion of Object Coverage Beyond Db2 Mirror

Db2 Mirror supports replication for a defined subset of IBM i object types, optimized specifically for continuous availability operations. Migrate While Active (MWA) expands this scope substantially. MWA is designed to migrate all save-eligible objects within SYSBAS, including objects stored in libraries as well as the Integrated File System (IFS). This capability far exceeds the replication-eligible subset supported by Db2 Mirror and is fundamental to enabling full partition migration. The broadened object coverage reflects MWA's shift from availability-focused replication toward comprehensive system relocation.

Independent ASP (IASP) data continues to be handled separately through IBM PowerHA SystemMirror® for i technologies, which remain complementary to MWA rather than being replaced by it. Together, these technologies support flexible migration and resiliency strategies across a wide range of IBM i deployment models.

**Restriction:** The migration of an independent auxiliary storage pool (IASP) is not supported by IBM i Migrate While Active at this time. It is suggested to use MWA to migrate SYSBAS to the new copy node and concurrently use tools like PowerHA with geographic mirroring to migrate any IASP attached to the system.

## How Migrate While Active builds on Db2 Mirror

Migrate While Active uses Db2® Mirror’s event-tracking infrastructure to capture changes occurring on the source node and later synchronize those changes to the copy node. Configuration, orchestration, and monitoring are performed through the Db2 Mirror GUI, with operational controls also surfaced through the WRKMIGSTS (Work with Migration Status) command. Key workflow transitions such as initiating or completing final synchronization can additionally be executed anagrammatically through the QSYS2.CHANGE\_MIGRATION\_MANAGER procedure.

Db2 Mirror contributes essential technology components, including:

- ▶ Change tracking services used to identify and record modifications on the source node during migration.
- ▶ Integration with the Db2 Mirror GUI, which becomes the primary interface for managing migration stages.
- ▶ Operational visibility and control through GUI dashboards, WRKMIGSTS panels, and SQL services such as:
  - MIGRATION\_MANAGER\_INFO
  - LIBRARY\_MIGRATION\_LIST
  - ESTIMATE\_FINAL\_SYNCHRONIZATION\_TIME

These services provide consistency, visibility, and an integrated management experience across the migration lifecycle.

Table 1-1 outlines the major distinctions between Db2 Mirror for i (Option 1) and Migrate While Active (Option 2).

*Table 1-1 Key differences between Migrate While Active and traditional Db2 Mirror*

Aspect	Db2 Mirror for i (Option 1)	IBM i Migrate While Active (Option 2)
Replication model and transport	Bidirectional, synchronous continuous availability using RDMA-capable adapters to keep two nodes in sync	Asynchronous migration over TCP/IP between source and copy nodes
Object scope	Replicates a defined set of replication-eligible object types	Migrates all save-eligible object types in SYSBAS, including objects in libraries and the integrated file system (IFS)
Independent ASP (IASP) data	Not applicable. Db2 Mirror description focuses on SYSBAS replication	Does not migrate IASP data; use IBM PowerHA for i IASP replication alongside the migration when needed

Aspect	Db2 Mirror for i (Option 1)	IBM i Migrate While Active (Option 2)
HA/DR suitability and coherence	Designed for continuous availability (HA) architectures (active/active, active/passive, active/read-only)	Not an HA/DR solution. Synchronization is not time-ordered and the copy node becomes coherent only after final cutover

During the system migration stage, the Db2 Mirror GUI guides the configuration of the source and copy nodes, while automated save/restore operations are coordinated using WRKMIGSTS on both consoles. After the initial save completes, object tracking begins on the source node to record ongoing changes. These changes are synchronized later during the continuous data synchronization stage.

For final synchronization and cutover, administrators can initiate the last replication pass using the Db2 Mirror GUI or the WRKMIGSTS command interface. Advanced users can also manage transitions using QSYS2.CHANGE\_MIGRATION\_MANAGER, enabling automation or scripted orchestration.

### Management and observability through Db2 Mirror tools

All migration status information is surfaced directly within the IBM i Migrate While Active Guide section of the Db2 Mirror GUI. This includes dashboards for migration progress, synchronization status, and actionable workflow steps. WRKMIGSTS provides a command line interface for system states where the GUI is unavailable.

A comprehensive suite of SQL services complements both GUI and command line tools, offering programmatic access to migration metadata, library level granularity, and cutover estimation functions.

In summary, although Migrate While Active (MWA) is architected on top of the Db2 Mirror for i technology stack, it is not a Db2 Mirror high-availability solution. Instead, MWA provides a controlled, asynchronous framework for migrating IBM i partitions from a source node to a copy node, leveraging Db2 Mirror services for change tracking, synchronization, and operational visibility. The copy node becomes a coherent and usable system only at the point of cutover, enabling a predictable, staged migration process that minimizes application downtime.

## 1.4 Supported Environments

IBM i Migrate While Active was announced in October 2024. The product is supported by operating system version 7.4 and above. It means that customers running IBM i 7.4, 7.5, or later releases can benefit from assisted migration workflows, improved orchestration, and enhanced migration capabilities of the product.

Support for Migrate While Active continues to evolve with each Db2 Mirror PTF group. However, at the time of writing this book, the partition mirroring migration option is not yet available for IBM i version 7.6 of operating system.

### 1.4.1 Enhancement Timeline and PTF-Level Mapping

IBM continuously enhances Migrate While Active through PTF updates, ensuring compatibility with current POWER hardware and providing ongoing improvements in performance and usability. To maintain full support, customers should keep their systems updated with the latest Db2 Mirror PTF groups.

The enhancement table below provides a chronological overview of key updates delivered for IBM i Migrate While Active across successive PTF levels for IBM i 7.4, 7.5, and 7.6. It highlights when major functions became available.

Table 1-2 MWA keys updates

Date	Migrate While Active Enhancement	Group PTF
July 2026	Planned	SF99961 - Level 3 (7.6) SF99951 - Level 11 (7.5) SF99668 - Level 29 (7.4)
November 2025	Manual migration method available Updated tracking and resynchronization process Performance improvements for IFS migrations Updated SQL services New options for IBM i 7.5: Partition mirroring method available Test mode for partition mirroring method	SF99961 - Level 2 (7.6) SF99951 - Level 10 (7.5) SF99668 - Level 28 (7.4)
July 2025	New options only for IBM i 7.4: Partition mirroring method available Test mode introduced	SF99951 - Level 9 (7.5) SF99668 - Level 27 (7.4)
April 2025	Updated DB2 Mirror GUI interface Replacing SCP with RSYNC Partially system save option available	SF99961 - Level 1 (7.6) SF99951 - Level 8 (7.5) SF99668 - Level 26 (7.4)
December 2024	Migrate while Active product is introduced	SF99951 - Level 7 (7.5) SF99668 - Level 25 (7.4)

To check current information about latest group PTF for Db2 Mirror use following link:  
<https://www.ibm.com/support/pages/db2-mirror-i-gui-ptf-information>

For the latest IBM i Migrate While Active PTF Information please visit following link:  
<https://www.ibm.com/support/pages/node/7178565>

## 1.4.2 Cloud Support

As organizations increasingly adopt hybrid and public cloud strategies, IBM i Migrate While Active (MWA) has become an important capability for enabling low-downtime migration into cloud-hosted Power platforms such as IBM Power Virtual Server and Google Cloud environments powered by IBM Power. These offerings provide a managed Power Systems infrastructure that is compatible with on-premises IBM i workloads, allowing customers to extend or relocate their environments without requiring major architectural changes. With MWA, administrators can replicate data and system objects from an on-premises IBM i system to a cloud-based target while production workloads continue to operate, enabling a controlled and predictable cutover with only a brief outage. This approach allows enterprises to modernize their infrastructure footprint, take advantage of cloud scalability and operational flexibility, and maintain the high availability expectations typically associated with IBM i workloads.

### Migrate While Active support for PowerVS

IBM i Migrate While Active (MWA) extends its migration framework to support cloud-based IBM Power Virtual Server (PowerVS) environments. This capability enables organizations to transition IBM i workloads from on-premises IBM Power Systems to PowerVS with minimal

service interruption. MWA uses asynchronous data replication and a unified operational model to provide a controlled, low-downtime migration path into cloud infrastructure.

IBM Power Virtual Server (PowerVS) is an IBM Power–based cloud offering that enables rapid provisioning of virtual servers with flexible, secure, and scalable compute capacity. PowerVS supports enterprise Power workloads across a variety of deployment models, whether hosted in IBM-managed data centers or deployed within client-operated facilities.

PowerVS supports IBM i 7.2 and later, including IBM i 7.6, the most recent release. For Migrate While Active (MWA) operations, this means that migrations targeting PowerVS must originate from partitions running a supported release level at the destination—typically IBM i 7.4 or later. Older versions, such as IBM i 6.1 or 7.1, must be upgraded prior to initiating an MWA migration to PowerVS. This requirement ensures compatibility with the cloud environment and alignment with current IBM i support policies.

### **PowerVS compatibility with MWA**

When building a new system on IBM Power Virtual Server, customers can take advantage of Cloud-Optimized Root (COR) images. These images provide pre-configured IBM i operating system environments that are ready to deploy, significantly simplifying and accelerating the initial system provisioning process.

IBM i 7.4, 7.5, 7.6, and the IBM i Cloud Optical Repository (COR) operating system images are available in the Image Catalog and they support Power10 and later systems. If required, you can also create a custom image or update your IBM i Virtual Server Instance (VSI) by downloading the installation media from the Entitled Systems Support (ESS) portal available under following link: <https://www.ibm.com/servers/eserver/ess/ProtectedServlet.wss>

For more details about available COR images visit IBM Cloud® website: <https://cloud.ibm.com/docs/power-iaas?topic=power-iaas-powervs-faqs#os-versions>

MWA is supported at the operating system level, and the underlying hardware also must be supported for the specific IBM i release. Table 1-3 shows the different PowerVS server options and the IBM i release that are supported for use with MWA.

*Table 1-3 Support matrix of MWA on PowerVS*

<b>PowerVS servers</b>	<b>IBM i support</b>	<b>MWA support</b>
POWER9 S922 (9009-22A), E980 (9080-M9S)	7.4 TR11, 7.5 TR5	Yes
POWER10 E1080 (9080-HEX ,S1022 (9105-22A)	7.4 TR6, 7.5, 7.6	Yes
POWER11 S1122 (9824-22A)	7.4 TR12, 7.5 TR6, 7.6	Yes

### **Prerequisites for MWA to PowerVS**

To successfully migrate an IBM i partition to PowerVS using MWA, you must fulfill several conditions:

► **OS Version Compatibility**

The target VSI must run a PowerVS supported IBM i level for MWA (7.4 and above), and the source system must be at a version compatible with MWA requirements.

- ▶ Network and Connectivity

Reliable IP connectivity between the on-prem source and PowerVS region is necessary. MWA uses TCP/IP for asynchronous replication, making it suitable even for long-distance migrations.

- ▶ Supported Hardware at the Destination

PowerVS infrastructure is based on POWER9, POWER10, and POWER11 systems depending on region, ensuring capability to run IBM i workloads. Not all servers are available in all regions, so you need to decide carefully.

- ▶ Availability of OS Images

As mentioned earlier in this chapter PowerVS offers Cloud-Optimized Root (COR) images that significantly simplify provisioning a target IBM i VSI. Customers may also upload or build custom images if required.

## **Migrate While Active Support for IBM Power for Google Cloud (IP4G)**

IBM i Migrate While Active also supports cloud-based the IBM Power for Google Cloud (IP4G) Google Cloud Marketplace solution environments. This facilitates customers moving their IBM i workloads from on-premises Power Systems to IP4G with minimal service interruption.

IP4G is an IBM Power-based cloud server offering available in the Google Cloud Marketplace, and operated by IBM Business Partner Peller Technologies. IP4G datacenters are located adjacent to Google Cloud regional datacenters and connect to Google infrastructure over secure, dedicated, low-latency networking. IP4G provides an ideal solution for customers already leveraging the Google Cloud for other workloads, who also have IBM Power workloads and can benefit from the synergy of Google and IBM technologies.

IP4G supports IBM i 7.3 and later releases, including the latest IBM i 7.6. Migrate While Active requires a minimum of IBM i 7.4 and therefore customers must upgrade their on-premises systems to IBM i 7.4 or above in order to leverage MWA for migration into IP4G.

### ***IP4G compatibility with MWA***

When building a new system in the IBM Power for Google Cloud environment, customers can leverage pre-built stock OS images. These images provide pre-configured IBM i operating system environments that help accelerate the initial system provisioning process.

IBM i 7.4, 7.5, 7.6 stock operating system images are available in the IP4G Stock Image Catalog and can be imported into a customer's IP4G instance using the pcloud command. Customers can also create custom images, and can obtain updates by downloading the installation media from the Entitled Systems Support (ESS) portal available under following link: <https://www.ibm.com/servers/eserver/ess/ProtectedServlet.wss>

For more details about importing stock images in IP4G, see the IP4G docs site: <https://docs.converge.cloud/docs/cli/cli-add-images-to-catalog/>

## **1.5 Licensing considerations**

IBM i Migrate While Active (MWA) is delivered as a licensed feature of Db2 Mirror for i (5770 DBM) and follows IBM's subscription based licensing model for IBM i availability and migration offerings. This model replaces perpetual licensing and includes software maintenance for the duration of the subscription. Subscription licensing is ideal for migrations as you only need to procure a license for the duration of the migration.

MWA does not require licensing the full Db2 Mirror product. Instead, customers install the Db2 Mirror base product (5770 DBM option 5050) at no charge and then license the feature option that matches their intended use.

Two Db2 Mirror feature options are available:

- ▶ Continuous Availability (5770 DBM option 5101), licensed per core on both participating nodes, is designed for ongoing availability scenarios.
- ▶ Migrate While Active (5770 DBM option 5102), licensed per core on the source node only, is designed specifically for planned, one time migration activities.

At a glance:

- ▶ Db2 Mirror base product (5050): No charge
- ▶ Continuous Availability (5101): Per core licensing on both nodes
- ▶ Migrate While Active (5102): Per core licensing on the source node only

The source system only licensing of Migrate While Active reflects its migration focused design. Unlike Continuous Availability, MWA is not intended for ongoing bidirectional synchronization or post cutover role reversal, and its licensing scope aligns with this temporary use model.

MWA is offered exclusively as a subscription with the following available terms:

- ▶ 28-day try and buy evaluation
- ▶ 90 day subscription
- ▶ 1 year, 2 year, 3 year, 4 year, and 5 year subscriptions

All subscription terms include IBM Software Maintenance (SWMA), providing entitlement to product updates, fixes, and technical support for the duration of the subscription.

After ordering, Migrate While Active can be downloaded from [IBM Entitled Systems Support \(ESS\)](#) under Software Downloads. Subscription license keys are provided through the IBM i Software Keys and Subscription Codes menu and are applied using standard IBM i license management mechanisms. Software entitlement and access are managed through IBM's standard software entitlement processes and accessed via ESS. Figure 1-1 shows the software download option from 'My Certified Software'.

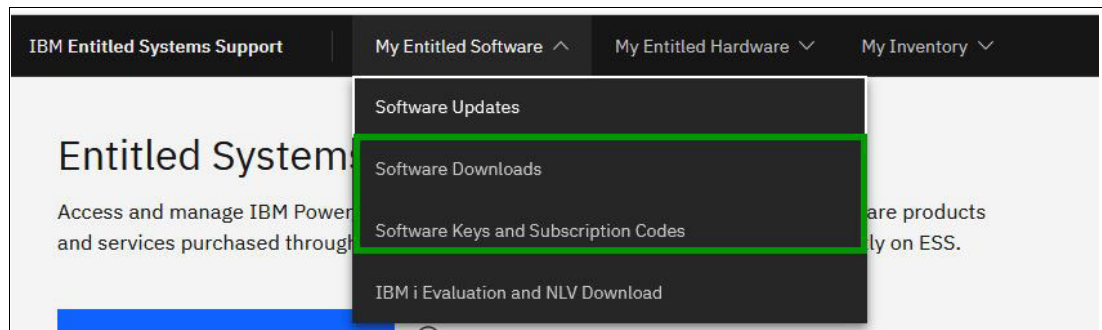


Figure 1-1 Software download options in ESS

For evaluation purposes, MWA can also be downloaded from ESS using the IBM i Evaluation and NLV Download option. The product is included on B\_GROUPx\_05 distribution media for specific IBM i releases; however, ESS is the preferred and recommended distribution method.

Figure 1-2 shows the evaluation download option.

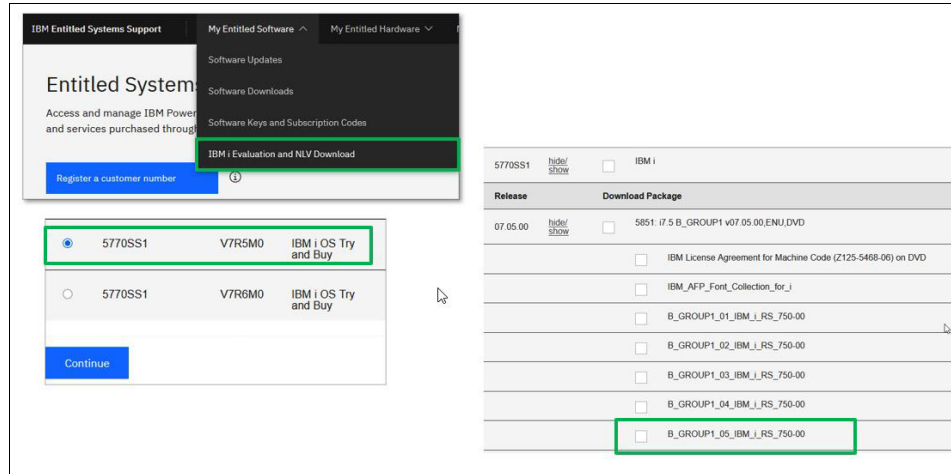


Figure 1-2 IBM i Evaluation software options for MWA in ESS

**Important:** Subscription pricing varies by geography, system configuration, and agreement terms. Refer to your IBM representative or Passport Advantage® agreement for current pricing information.

Before you begin:

- ▶ Verify support for IBM i 7.4, 7.5, or 7.6 and confirm feature availability for the selected release.
- ▶ Ensure required Db2 Mirror GUI and Migrate While Active PTFs are applied.
- ▶ Confirm POWER system compatibility for the planned migration path.
- ▶ Obtain ESS entitlement (subscription or evaluation) and required license keys.

## MWA licensing in PowerVS

One of the primary use cases for Migrate While Active (MWA) is support for migration of IBM i environments into IBM Power Virtual Server and other cloud environments with minimal application downtime. All licensing considerations described in the previous section apply equally to IBM i when deployed in PowerVS. When PowerVS is used as the migration destination, MWA licensing is required only on the source IBM i partition initiating the migration, because all replication, synchronization, and tracking operations are performed from that system. The target PowerVS virtual server does not require an MWA license, simplifying capacity planning and reducing the overall licensing footprint for cloud migrations.

Short-term subscription options—such as the 90-day offering—are particularly well suited for one-time migrations into PowerVS, providing a cost-effective licensing model for organizations transitioning to cloud-based IBM i environments.

## 1.6 Data transfer options

During an MWA operation, the source IBM i partition must transmit specific files—such as IBM i installation media and save files created as part of the migration workflow—to the destination system. Earlier MWA releases performed this transfer using SCP, a basic secure-copy mechanism.

Beginning with the April 2025 update, IBM i Migrate While Active (MWA) uses RSYNC instead of SCP to automatically transfer installation media and saved user data from the source system to the target system during a migration. IBM replaced SCP with RSYNC, providing a more efficient, resilient, and bandwidth-aware transfer method. RSYNC improves performance and reduces the amount of data that must be retransmitted, enhancing the overall reliability of the migration process.

This enhancement significantly increases the reliability of migrations, especially when handling large save images or operating in environments with variable network stability.

Conceptually, the difference can be summarized as:

- ▶ SCP — copies the entire file set every time
- ▶ RSYNC — copies only the data that has changed

RSYNC is optimized for transferring large datasets and provides several operational advantages over SCP in migration workflows. Rather than blindly copying files, RSYNC performs incremental transfers, sending only the portions of data that differ from the target. In addition, RSYNC can resume incomplete transfers without restarting the entire operation, reducing the impact of temporary network disruptions and preventing delays caused by restarting long-running copy processes.

Key RSYNC capabilities that benefit migration workflows include:

- ▶ Incremental transfer: Only changed blocks of data are sent, improving performance for repeated or resumed migrations.
- ▶ Automatic resume: RSYNC can keep partial data and continue an interrupted transfer without starting over, important during long copy operations or variable network conditions.
- ▶ Compression and optimization: optional in-flight compression (-z) can reduce bandwidth consumption on slower or congested links.
- ▶ Preservation of metadata: archive mode (-a) maintains permissions, ownership, timestamps, and symbolic links, critical for restoring IBM i environments accurately.

During a Migrate While Active operation, RSYNC is used to efficiently move the following artifacts from the source to the target system:

- ▶ Installation media generated on the source system
- ▶ User data saved through the assisted save-and-restore workflow

These characteristics make RSYNC a more robust and efficient mechanism for MWA migrations, particularly when transferring large save images or when performing migrations across networks with fluctuating bandwidth or intermittent connectivity.

### **RSYNC prerequisites**

RSYNC is not a native IBM i command (like SAVSYS or RSTOBJ).

Instead, it is delivered through the IBM Open Source environment (5733-OPS or the newer RPM-based Open Source stack under `/QOpenSys/pkg`). Which means:

- ▶ RSYNC must be installed on the source partition.
- ▶ It typically runs in a PASE environment (AIX-like runtime built into IBM i).

MWA uses it automatically, you do not run RSYNC manually.

For details on installing RSYNC see section 2.5, “Installing RSYNC using the Open Source Package Management interface” on page 38

### Role of RSYNC in migration workflows

RSYNC is required when performing migrations using the assisted save and restore method. In this migration path, RSYNC is used to efficiently transfer saved data from the source system to the target system as part of the automated migration workflow.

Other migration methods provided by Migrate While Active do not require RSYNC unless the migration workflow explicitly invokes file-based data transfers. For migrations that do not use assisted save and restore, RSYNC may be present on the system but is not actively used by the Migrate While Active tooling.

By integrating RSYNC into the Migrate While Active workflow, IBM improves the reliability and efficiency of data movement during migration while maintaining alignment with standardized open-source tools available on IBM i.

## 1.7 IBM i Object Operations: Changes in IBM i 7.5 and IBM i 7.6

This section consolidates release specific improvements that influence how IBM i manages objects during migrations and day-to-day operations. For IBM i 7.5 MWA added new migration patterns, better throughput for changed IFS objects, and improved synchronization estimation. IBM i 7.6 extends those gains with deeper performance visibility, file-transfer resiliency, a partial system-save workflow, specific accelerations for delete/move/rename operations, and additional security controls for IFS sharing. Together, these enhancements reduce migration elapsed time, improve predictability at cutover, and make ongoing administration more automated and secure.

### 1.7.1 MWA Enhancements: IBM i 7.5 and IBM i 7.6

Understanding the evolution of Migrate While Active across IBM i 7.5 and 7.6 is essential for architects and administrators planning migrations, optimizing synchronization workflows, or evaluating upgrade readiness.

The table below provides a side-by-side comparison of the major enhancements in each release, enabling quick assessment of how each version impacts performance, stability, and operational efficiency during MWA-based migrations.

Table 1-4

Category	IBM i 7.5 Enhancements	IBM i 7.6 Enhancements
Migration Methods	Introduces Manual Migration and Partition Mirroring, reducing outages and enabling flexible save/restore or replication-based workflows.	Continues same methods; adds RSYNC-based media and data transfer for improved resiliency and throughput.
Initial Copy Optimization	No partial system-save capability; full system save required for assisted save & restore.	Adds Partial System Save, avoiding saving most Db2 Mirror replication-eligible objects during the initial stage and moving their migration to data-sync phase.

Category	IBM i 7.5 Enhancements	IBM i 7.6 Enhancements
Data Transfer Method	Uses conventional SCP for save media transfer.	Replaces SCP with RSYNC, enabling resumable and incremental transfer of large datasets.
Object Operation Handling	Parallel processing of changed IFS objects improves throughput.	Adds dedicated performance improvements for delete, move, and rename operations of replication-eligible objects.
Rename Handling in Db2 Mirror / MWA	Updated tracking allows rename operations on libraries with tracked objects.	Same capability carries forward.
Sync-Time Prediction	Improved ESTIMATE_FINAL_SYNCHRONIZATION_TIME function provides more accurate predictions.	Same function; benefits enhanced by partial system-save and faster structural ops.
Database / Lock Performance Visibility	No new lock-wait instrumentation.	Adds Object Lock Wait, Record Lock Wait, and Seize Wait metrics with Visual Explain and SQL monitor integration.
IFS / Object Management SQL Services	Supports existing IFS and object authority SQL services.	Adds or expands services such as IFS_OBJECT_PRIVILEGES, IFS_PATH, VERIFY_NAME, and Automated Receiver Cleanup (DELETE_OLD_JOURNAL_RECEIVERS).
Security / Object Governance	Supports RESTRICT ON DROP and DROP IF EXISTS for safer operations.	Adds NetServer authority reinforcement through QIBM_QZLS_NETSVR_SHARE for share-creation control.

## 1.7.2 Performance & Wait Analysis (Granular Visibility in 7.6)

One of the most significant performance enhancements delivered with IBM i 7.6 for database tooling is the introduction of dedicated wait-time metrics that provide clear visibility into where queries stall and why.

► Detailed wait insights

IBM i 7.6 now captures granular wait metrics—including Average Object Lock Wait Time, Average Object Lock Wait Count, Average Record Lock Wait Time, Average Record Lock Wait Count, Average Seize Wait Time, and Average Seize Wait Count. These measurements expose lock- and seize-related contention that previously required indirect analysis or inference from secondary indicators.

► Integrated tooling support

These new wait metrics are fully integrated into Visual Explain (including both Plan Cache snapshots and the Live Plan Cache), Database Monitor, and SQL performance monitors. This native integration enables engineers to rapidly correlate wait events with specific SQL statements, simplifying root-cause analysis and accelerating performance diagnostics.

During migration rehearsal and production cutover windows, the ability to quickly identify and isolate object and record lock waits is especially valuable. Faster detection of application

driven contention helps teams address issues that could otherwise extend final synchronization or impact migration timelines.

### 1.7.3 Automated Object Management via SQL Services

IBM i 7.5 and IBM i 7.6 introduce several enhancements to SQL-based system management that simplify authority analysis, cleanup operations, and IFS-driven scripting—capabilities especially useful during migration planning and validation.

- ▶ IFS authority inspection using IFS nomenclature: The QSYS2.IFS\_OBJECT\_PRIVILEGES table function returns per-object rows with IFS-style authorities (e.g., \*R, \*RW, \*X), owners, primary groups, and authorization lists. Ideal for audits or pre-cutover checks.
- ▶ Automated journal receiver cleanup: The SYSTOOLS.DELETE\_OLD\_JOURNAL\_RECEIVERS procedure deletes detached receivers based on time and filtering criteria, with an optional preview mode. This reduces storage growth and housekeeping complexity.
- ▶ IFS path parsing: SYSTOOLS.IFS\_PATH returns selected parts of a path (file name, extension, prefix, directory path), enabling simpler scripting for IFS file handling or reporting. (Requires ICU option.)
- ▶ Name safety checks: QSYS2.VERIFY\_NAME validates that a string conforms to either SYSTEM (\*NAME) or SQL identifier rules before use, useful for generative admin scripts.

Before and after migration, teams can script authority snapshots, receiver hygiene, and IFS path manipulations, replacing adhoc command sequences with idempotent SQL, which is easier to version, review, and automate.

### 1.7.4 Object Security & Operational Guardrails

IBM i 7.5 and 7.6 introduce several enhancements that strengthen object protection, refine authority governance, and improve replication resiliency—capabilities that directly reduce operational risk during migration preparation and execution.

- ▶ RESTRICT ON DROP / DROP IF EXISTS (continuing in 7.5 and 7.6).
  - RESTRICT ON DROP protects critical tables, functions, and procedures from accidental deletion.
  - DROP ... IF EXISTS allows cleanup scripts to succeed even if a target object is already absent, avoiding spurious failures.
- ▶ IFS file-share controls (7.6).

IBM i NetServer introduces the QIBM\_QZLS\_NETSVR\_SHARE function usage ID. To create, change, or remove shares without \*IOSYSCFG, the profile must own the underlying object and have this function usage authority, tightening share creation privileges.
- ▶ Db2 Mirror rename resiliency (applies to MWA, 7.5/7.6 lines).

The tracking and resynchronization process was updated to allow rename operations on libraries containing tracked objects (other rename restrictions still apply). This removes a former operational friction point when resuming replication.

Stronger drop protections and governed share creation reduce accidental disruptions before/during migration, while improved rename handling decreases manual intervention during resynchronization.

## 1.7.5 Summary of Enhancements

IBM i 7.5 and 7.6 introduce substantial performance and resiliency improvements that strengthen object handling, synchronization throughput, and diagnostic visibility for Migrate While Active (MWA) migrations. IBM i 7.5 delivers key gains in Integrated File System (IFS) throughput through parallel processing of changed objects, accelerating synchronization in environments with high IFS update volume. IBM i 7.6 builds on this foundation with targeted optimizations for structural object operations—such as delete, move, and rename—improving the efficiency and stability of cutover preparation.

IBM i 7.6 further enhances migration predictability by adding granular wait-time metrics that expose lock, record, and seize contention, fully integrated into Visual Explain, Plan Cache, Database Monitor, and SQL performance tools. These metrics provide deeper insight into synchronization delays than what is available in IBM i 7.5.

In addition, IBM i 7.6 modernizes MWA's file-transfer and save-strategy workflow. The shift to RSYNC improves the transfer of installation media and save images through incremental and resumable operations, while the new partial system-save option reduces outage time by excluding most Db2 Mirror-eligible objects from the initial save. These advancements, unique to IBM i 7.6, significantly improve the efficiency, robustness, and end-to-end performance of MWA migrations.





# Migration Planning and Prerequisites for Minimizing Downtime

Chapter 2 provides an overview of the essential foundations required to plan and execute a successful IBM i migration using Migrate While Active. It introduces the core preparatory activities such as assessing both the current and target environments, determining the appropriate migration path, and identifying prerequisites as well as strategies for minimizing downtime during LPAR transitions.

The chapter also highlights key technical enablers, including the use of RSYNC for data transfer and the enhanced Db2 Mirror GUI, and concludes with a brief security overview to ensure a safe and controlled migration process.

This chapter will cover the following session:

- ▶ “Hardware and Software requirement”
- ▶ Table 2-6 on page 22
- ▶ “Determining Migration Path and Prerequisites”
- ▶ “Enhanced Db2 Mirror GUI”
- ▶ “Installing RSYNC using the Open Source Package Management interface”
- ▶ “Security highlights of IBM i MWA”

## 2.1 Hardware and Software requirement

Planning a successful Migrate While Active (MWA) implementation begins with understanding the foundational hardware and software requirements that support the migration process. Because MWA coordinates system save operations, data movement, and synchronization between source and target IBM i environments, both systems must be configured with the appropriate processing capabilities, storage resources, network connectivity, and licensed software components. Ensuring that these prerequisites are in place not only enables MWA to operate reliably but also maximizes performance, reduces migration time, and helps avoid unexpected interruptions during cutover. The Migrate While Active product requires a number of additional IBM i options be installed. Each of these options must be up to date on PTFs as well.

This section describes the essential system-level requirements and environmental conditions needed to prepare an IBM i landscape for MWA deployment. Table 2-1 provides a list of IBM i components on the source node defining whether they are required or optional for MWA. Table 2-2 on page 21 defines the GUI node requirements and Table 2-3 on page 21 defines the requirements for the copy node.

*Table 2-1 Source Node products and options*

Product and Option	Description	Notes
5770SS1 Option 1	Extended Base Support	Required
5770SS1 Option 3	Extended Base Directory Support	Required
5770SS1 Option 12	Host Servers	Required
5770SS1 Option 22	Object Connect	Required
5770SS1 Option 26	Db2 Symmetric Multiprocessing	Required
5770SS1 Option 30	Qshell	Required
5770SS1 Option 33	Portable Application Solutions Environment (PASE)	Required
5770SS1 Option 41	High Availability Switchable Resources	Required
5770SS1 Option 48	Db2 Data Mirroring	Required
5770DBM *BASE	IBM Db2 Mirror for i	Required
5770DBM Option 2	IBM i Migrate While Active	Required. May appear as Reserved – Option 2 or Migration Feature in some interfaces.
5770DG1 *BASE	IBM HTTP Server for i	Optional. Required to run the Db2 Mirror GUI on the Source Node.
5770JV1 *BASE	IBM Developer Kit for Java	Required
5770JV1 Option 16, 17, 19, or 20	Java SE 8, 11, or 17	Optional. Any supported version of Java must be installed to run the Db2 Mirror GUI on the Source Node.
5733SC1 *BASE	IBM Portable Utilities for i	Required

Product and Option	Description	Notes
5733SC1 Option 1	OpenSSH, OpenSSL, zlib	Required
5770TC1 *BASE	IBM TCP/IP Connectivity for i	Required

Table 2-2 GUI Node products and options

Product and Option	Description	Notes
5770SS1 Option 1	Extended Base Support	Required
5770SS1 Option 3	Extended Base Directory Support	Required
5770SS1 Option 12	Host Servers	Required
5770SS1 Option 30	Qshell	Required
5770SS1 Option 33	Portable Application Solutions Environment (PASE)	Required
5770DBM *BASE	IBM Db2 Mirror for i	Required
5770DG1 *BASE	IBM HTTP Server for i	Optional. Required to run the Db2 Mirror GUI on the Source Node.
5770JV1 *BASE	IBM Developer Kit for Java	Required
5770JV1 Option 16, 17, 19, or 20	Java SE 8, 11, or 17	Optional. Any supported version of Java must be installed to run the Db2 Mirror GUI on the Source Node.
5770TC1 *BASE	IBM TCP/IP Connectivity for i	Required

Table 2-3 Copy Node products and options

Product and Option	Description	Notes
5770SS1 Option 1	Extended Base Support	Required
5770SS1 Option 3	Extended Base Directory Support	Required
5770SS1 Option 12	Host Servers	Required
5770TC1 *BASE	IBM TCP/IP Connectivity for i	Required

## Open Source packages

The following Open Source packages must be installed on the Source Node:

- ▶ cloud-init
- ▶ python39
- ▶ python39-ibm\_db
- ▶ rsync

The following Open Source packages must also be installed on the NFS Node if using Assisted Save and Restore:

- ▶ rsync

## PTF Groups

The following PTF Groups must be installed and at the latest level

Table 2-4 Source Node

PTF Group Name	PTF Group Number	Minimum Level	Recommended Level
IBM i Cumulative PTF Package	SF99750	-	Latest
Db2 Mirror PTF Group	SF99951	10	Latest

Table 2-5 GUI Node

PTF Group Name	PTF Group Number	Minimum Level
IBM i Cumulative PTF Package	SF99750	-
Db2 Mirror PTF Group	SF99951	10
IBM HTTP Server for i PTF Group	SF99952	-
Java PTF Group	SF99955	-

## Environment Requirements

If using Partition Mirroring, the following ports need to be enabled between the Source Node and the Copy Node:

- ▶ 936/TCP (Partition Mirroring)
- ▶ Ephemeral/TCP (Partition Mirroring Data Port)

**Note:** Data port can be configured to use a specific port instead of an ephemeral port. This is done using the partition mirroring advanced analysis command and specifying the -PORT parameter.

The following ports are also required between the GUI Node and the Source and Copy Nodes.

Table 2-6 Port option

Port/Protocol	Service Name	Description
449/TCP	as-srvmap	Host servers server mapper
2001/TCP	as-admin-http	HTTP Administration
2006/TCP	as-admin3-http	Db2 Mirror GUI job
8470/TCP	as-central	Host servers central server
8471/TCP	as-database	Host servers database server
8472/TCP	as-dtaq	Host servers data queue server
8473/TCP	as-file	Host servers file server
8475/TCP	as-rmtcmd	Host servers remote command and program call server
8476/TCP	as-signon	Host servers signon server

## 2.2 Minimizing Downtime During IBM i LPAR Migrations

Minimizing downtime is a primary goal of any IBM i LPAR migration. IBM i Migrate While Active (MWA) supports this goal by keeping production workloads running while the target LPAR is synchronized in the background. Once the systems are aligned, migration completes with only a brief cutover window.

MWA provides multiple migration approaches to accommodate different infrastructure models and operational requirements:

- ▶ **Partition mirroring** uses host-based storage replication technology to copy SYSBAS from the source node to the copy node. It is currently available in IBM i 7.4 and IBM i 7.5 only. Partition mirroring offers the lowest downtime option, with no initial outage host-based storage replication. Its test mode capability allows full validation of the target before cutover.
- ▶ **Manual migration** supports warm clone, cold clone, and hybrid storage replication methods, enabling flexibility across diverse storage configurations while minimizing disruption.
- ▶ **Assisted save and restore** requires an initial restricted state save but still achieves a small cutover window through continuous asynchronous synchronization.

Across all methods, MWA uses continuous replication to keep the target closely aligned with the source. The Db2 Mirror GUI provides clear visibility into synchronization progress, allowing administrators to plan migrations confidently and maintain service availability.

By selecting the most appropriate method for the environment and applying recommended best practices, organizations can perform IBM i migrations efficiently and with minimal business interruption.

**Important:** The target LPAR becomes application-consistent only after the final cutover completes. Mid-migration states should not be treated as production-ready, and replication is not reversible after the switch.

### 2.2.1 Migration methods and downtime characteristics

MWA supports three migration methods. Each offers a different balance of initial outage, continuous synchronization behavior, and final cutover time. These are listed in Table 2-7.

Table 2-7 Migration method downtime comparison

Migration Method	Initial Outage	Continuous Sync	Final Cutover
Partition Mirroring	None	Yes - Asynchronous	Minimal
Manual Migration	Quiesce or Cold	Yes - Asynchronous	Minimal
Assisted Save/Restore	Full System Save	Yes - Asynchronous	Minimal

Use this table to select the method that best aligns with your business window, storage architecture, and operational constraints.

## 2.2.2 Minimizing Downtime with Partition Mirroring

Partition mirroring uses host-based storage replication to copy SYSBAS from source to target without requiring a full system save or an initial outage. Replication operates below technology-independent machine interface (TIMI), so active objects do not require explicit locks, and production workloads can continue normally.

Key advantages

- ▶ No initial outage or media creation
- ▶ Below TIMI operation avoids locking conflicts
- ▶ Low overhead via tracking-list change capture (not journal-heavy)
- ▶ Test mode to validate the copy node safely before cutover
- ▶ Observable progress from the Db2 Mirror GUI

### How Partition Mirroring Works

Partition mirroring uses an asynchronous storage replication method based on tracking lists rather than traditional journaling. This approach minimizes the storage overhead required for tracking changes while maintaining efficient data synchronization.

The tracking list approach does not use significant storage relative to journal-based systems where every change must be tracked in sequential order. This efficiency allows partition mirroring to operate with minimal impact on production workloads.

**Note:** Because Migrate While Active uses asynchronous replication and does not track changes in time order, the copy node data is incomplete and not in a coherent state until the final migration is completed. This makes Migrate While Active unsuitable for DR purposes.

### Copy node software requirements

The copy node must meet a minimum LIC/TR level before starting partition mirroring as shown in Table 2-8.

Table 2-8 Minimum TR requirements for partition mirroring

IBM i 7.6	IBM i 7.5	IBM i 7.4
Not yet supported	TR7 or higher	TR12 or higher

Storage Requirements:

- ▶ The storage capacity assigned to SYSBAS on the copy node must be equal to or greater than the amount of allocated (used) storage of SYSBAS on the source node.
- ▶ The individual disk unit sizes and types do not need to match.
- ▶ The number of disk units does not need to match.
- ▶ If the source node has one or more user auxiliary storage pools (ASPs), you must configure the same user ASPs (2-32) on the copy node before starting partition mirroring.
- ▶ Partition mirroring replicates SYSBAS and all basic user ASPs together.

## Test Mode and Migration Validation

One of the most significant downtime reduction features of partition mirroring is the test mode capability. Test mode allows users to thoroughly validate the migration before final cutover, minimizing the risk of issues that would require rollback and extended downtime.

Test mode and validation:

- ▶ Replication is paused and changes on the copy node are tracked during testing.
- ▶ You can IPL and fully validate the copy node while the source remains in production.
- ▶ When testing completes, test mode changes are rolled back, and synchronization resumes.
- ▶ The GUI displays the time to enter/exit test mode, which is a strong predictor of cutover duration.

**Tip:** Use test mode to complete application verification, performance checks, and operator runbooks before scheduling the production cutover window.

Figure 2-1 illustrates the optional transition into Test Mode during a Migrate While Active operation. In this mode, the user can start the copy node, bring it to a functional state, and perform any required testing or validation activities before proceeding with the final migration steps.

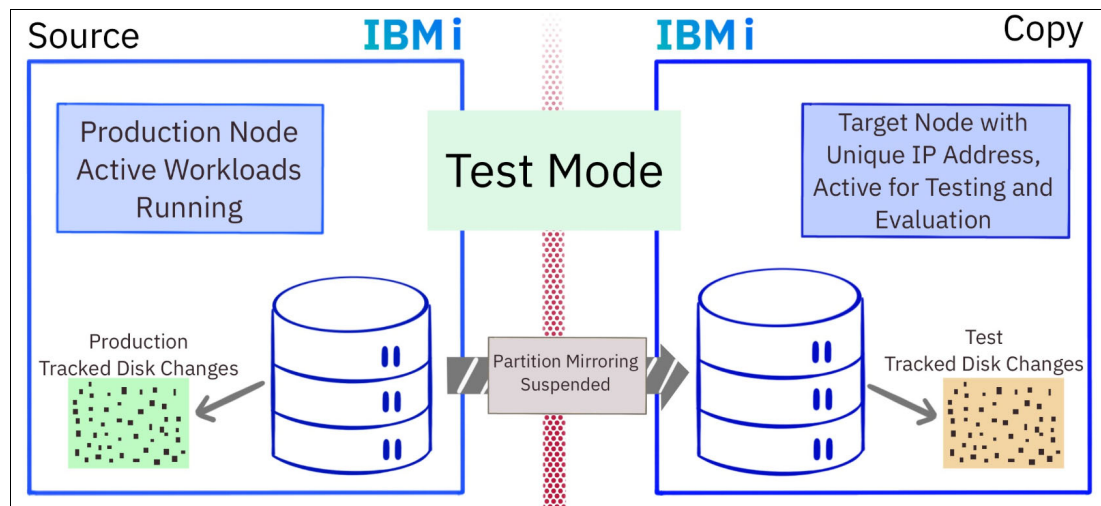


Figure 2-1 Partition mirroring test mode with active source node

Additional details are presented in section 2.1, “Hardware and Software requirement” on page 20.

## 2.2.3 Minimizing Downtime with Manual Migration

IBM i Migrate While Active manual migration provides users several alternative methods for creating a copy of the source node to establish the copy node. After the initial copy is created, IBM Db2 Mirror for i and Migrate While Active technologies are used to continuously synchronize changes in SYSBAS from the source node to the copy node.

### Manual Migration Methods

Manual migration provides the following alternative methods:

- ▶ Storage replication: Uses FlashCopy or remote copy technologies from external storage systems.
- ▶ Save and restore: Uses a full system save of the source node.
- ▶ Storage replication and save/restore: Combines storage replication to a temporary node with save and restore to the copy node.

### **Downtime Reduction with Storage Replication**

Storage replication using FlashCopy or remote copy provides options for minimizing downtime during the initial copy stage:

#### **Warm Clone (Minimal Downtime):**

- ▶ The source node remains operational during most of the cloning process.
- ▶ A brief quiesce period is required to ensure consistency when creating the storage copy.
- ▶ Migrate While Active starts tracking data changes during the quiesce.
- ▶ After the storage copy is complete, the source node resumes normal operations.

#### **Cold Clone (Scheduled Downtime):**

- ▶ The source node is powered off during the storage copy operation.
- ▶ Provides the most consistent copy but requires a full outage window.
- ▶ Tracking begins just before shutdown.
- ▶ The source node can be restarted immediately after the storage copy completes.

**Note:** Before quiescing or powering off the source node, hold scheduled jobs. This allows you to control when scheduled jobs begin to run on the copy node. Release them after the migration stage is complete. Note that changes to scheduled jobs on the source node are not migrated by Migrate While Active during data synchronization.

### **Reduced Downtime with Storage Replication and Save/Restore**

The manual storage replication and save/restore migration method provides users with reduced source node downtime compared to the assisted save and restore method. This method is designed for environments where:

- ▶ The source node uses external storage.
- ▶ The copy node does not use external storage or uses different external storage.
- ▶ Minimizing downtime on the production source node is critical.

#### **Process:**

1. Use storage replication to copy source node volumes to a temporary node.
2. Take a full system save from the temporary node (not the production source).
3. Use the save media to scratch install and restore the copy node.
4. Begin continuous synchronization from source to copy node.

This approach reduces downtime on the source node because the time-consuming full system save is performed on the temporary node rather than the production system.

## Save Options for Manual Migration

The manual save and restore migration method provides users with more options for taking a full system save. During a manual migration, options for taking a full system save include but are not limited to:

1. **GO SAVE option 21:** Traditional full system save.
2. **QSRMIGRATE program:** Streamlined migration save process.
3. Virtual tape libraries (VTL) for network-based saves.
4. Physical tape media or optical media

Additional information is provided in section 3.3, “Manual Migration” on page 55.

## 2.2.4 Managing Downtime with Assisted Save and Restore

When assisted save and restore is used, IBM i Migrate While Active uses Db2 Mirror for i, along with save and restore to migrate the source node. This method requires an initial outage on the source node to facilitate the creation of install media and user data save media, which are used to provision the copy node.

Assisted save and restore has two distinct downtime periods:

- ▶ **Initial outage for media creation:** The source node must be placed in restricted state to create install media and perform the initial user data save.
- ▶ **Final cutover period:** A minimal downtime window during the final synchronization and role swap.

After the initial media creation, IBM Db2 Mirror for i and Migrate While Active technologies are used to continuously synchronize changes made to objects after the creation of the media. This continuous synchronization minimizes the final cutover window.

Assisted save and restore is appropriate when:

- ▶ External storage is not available for storage replication.
- ▶ Partition mirroring requirements cannot be met.
- ▶ A scheduled maintenance window is available for the initial save.
- ▶ Simplicity of setup is more important than eliminating the initial outage.

**Note:** While assisted save and restore requires an initial outage, the continuous synchronization capability still minimizes the final cutover window, making this method significantly better than traditional migration approaches that require extended downtime during the entire restore and testing process.

### Partial System Save option

The Partial System-Save option in Migrate While Active (MWA) reduces the initial outage window by excluding most Db2 Mirror replication-eligible objects from the initial system save. Instead of saving and restoring the full set of objects up front, MWA captures only the required system components and save-eligible content that are not efficiently handled by replication. The remaining objects are transferred during the synchronization stage, where MWA’s asynchronous change tracking can move data with minimal disruption.

Traditional full-system saves are reliable but can prolong the initial outage—especially in environments with large libraries, extensive IFS content, or frequent object churn. The partial system-save option shifts work from the initial save/restore outage phase into the online

synchronization phase. This results in a shorter cut-out (start of migration) and a smaller cutover (final switch), improving overall migration predictability.

What the partial save excludes and includes

- ▶ Excluded from initial save (migrated during synchronization):
- ▶ Most objects that are replication-eligible in the Db2 Mirror framework (for example, many library-based objects and selected IFS content).
- ▶ High-churn items better handled by incremental synchronization than bulk save/restore.

Included in initial save (handled upfront):

- ▶ System configuration and objects not covered by replication eligibility.
- ▶ Installation media and specific save images that must land on the target early to bootstrap the migration workflow.
- ▶ Any explicitly required artifacts per your change-control plan (for example, environment prerequisites, product options, PTF images).

**Note:** Exact eligibility is governed by the same internal classification that underpins Db2 Mirror tracking. Independent ASP (IASP) data remains outside the MWA scope and should be migrated with PowerHA SystemMirror for i as a complementary workflow.

The Partial System Save option is presented on the User Data Save options screen as shown in Figure 2-2.

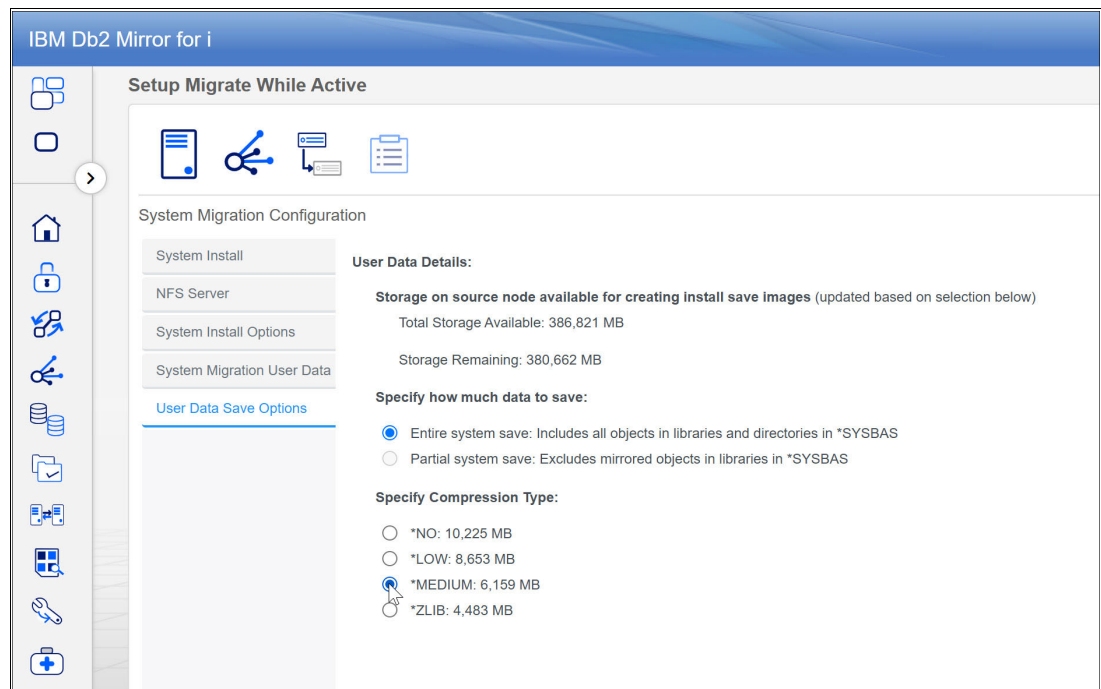


Figure 2-2 User data save screen

This option avoids saving most Db2 Mirror replication-eligible objects during the system save. Instead, these objects are automatically migrated later during the data synchronization stage. Using the Partial System Save option can reduce downtime and the amount of data that must be saved, and subsequently transferred to the remote site to set up the Copy Node.

**Note:** To use the partial system save option requires coordination with IBM. Send an email to [iessspt@us.ibm.com](mailto:iessspt@us.ibm.com), requesting the use of this option. The email will be answered by a member of IBM's Technology Expert Labs (formerly Lab Services).

Additional information is provided in section 3.2, “Assisted Save and Restore” on page 49

## 2.2.5 Continuous Synchronization and Final Cutover

All Migrate While Active migration methods use asynchronous replication during the continuous synchronization phase. This is a key difference from traditional Db2 Mirror, which uses synchronous replication.

Key characteristics of asynchronous replication in Migrate While Active:

- ▶ Uses TCP/IP as the communication protocol instead of RDMA.
- ▶ Migrates all save-eligible object types found in SYSBAS, regardless of whether they reside in libraries or the Integrated File System.
- ▶ Changes are not tracked in time order.
- ▶ The copy node is only in a coherent state after the final cutover completes.
- ▶ Users can monitor the remaining amount of migration needed before the copy node will be in sync.

### Monitoring Synchronization Progress

At any point during the migration process, Migrate While Active users can examine the remaining amount of migration needed before the copy node will be in sync with the source node. This visibility allows for:

- ▶ Accurate estimation of when the migration will be ready for cutover.
- ▶ Planning of maintenance windows for final cutover.
- ▶ Identification of potential issues before the final migration.
- ▶ Communication with stakeholders about migration status.

The Db2 Mirror GUI provides an easy-to-use interface for monitoring:

- ▶ Current replication status.
- ▶ Data transfer progress.
- ▶ Estimated time to completion.
- ▶ Any errors or warnings that require attention.

### Minimizing Final Cutover Downtime

The final cutover is the last phase of migration where the copy node is promoted to become the production system. Migrate While Active minimizes this downtime through:

- ▶ **Extended synchronization:** The continuous synchronization phase can run as long as needed to ensure the copy node is fully synchronized.
- ▶ **Minimal delta:** By the time of cutover, only a small amount of recent changes need to be synchronized.
- ▶ **Automated process:** The cutover process is automated and orchestrated by Migrate While Active.
- ▶ **Role swap:** The final role swap is a quick operation once synchronization is complete

## 2.3 Determining Migration Path and Prerequisites

System migration is a critical decision for any organization. A properly selected migration path allows you to perform the entire process in a way that is safe for the continuity of the source system, avoids unexpected downtime, and minimizes the risk of unpredictable events. When planning an IBM i migration using Migrate While Active, the selection of the migration method depends on several criteria:

- ▶ Size of migrated environment
- ▶ Complexity of data
- ▶ Testing possibility
- ▶ Downtime tolerance
- ▶ Storage type
- ▶ Target location capabilities

### 2.3.1 Migration paths you can choose

IBM i Migrate While Active supports multiple migration methods, and you select the method in the Db2 Mirror GUI migration flow. The Db2 Mirror GUI is the primary interface for the migration journey and guides you through the stages. Some steps must be performed from an IBM i console in restricted state. The primary migration paths are: Partition Mirroring, Assisted Save and Restore, and Manual Migration.

#### **Partition Mirroring**

Partition Mirroring uses host-based storage replication technology to copy SYSBAS from the source node to the copy node. This method avoids an initial outage because you do not need to create install media or perform a full system save to establish the copy node. It also provides test mode, which lets you evaluate the copy node before you complete the migration. When you are ready to complete the migration, you perform final cutover using a cool or warm cutover method (cool cutover is required when migrating IP addresses from the source node).

#### ***Limitations for Partition Mirroring:***

- ▶ If the source node contains an encrypted user auxiliary storage pool (ASP), it is not eligible for Partition Mirroring.
- ▶ If the source node contains independent ASPs (IASPs), the IASPs are not replicated using Partition Mirroring.

#### **Assisted Save and Restore**

Assisted Save and Restore establishes the copy node by taking an initial outage on the source node to create install media and user data save media, then continuously synchronizes changes made after media creation, followed by final synchronization and cutover. IBM i Migrate While Active provides guided, step-by-step operational procedures for the required console actions in this method: saving the source node, installing and restoring the copy node, and completing final synchronization. This method aligns well when you want a documented save/restore baseline and a structured console-driven process (**WRKMIGSTS**) to control the save, restore, and final synchronization stages.

#### **Manual Migration**

Manual Migration lets you create the copy node using one of these approaches: Storage Replication, Save/Restore, or Storage Replication and Save/Restore (using a temporary node and then a full system save to build the copy node). After you establish the copy node, IBM i

Migrate While Active tracks changes and synchronizes them until final synchronization and cutover.

## 2.3.2 How and why to choose a migration method

Choosing the right migration method is a balance of risk, downtime, and operational effort aligned to your business objectives.

This section explains how to translate your environment's characteristics into a clear decision among Partition Mirroring, Assisted Save and Restore, and Manual Migration. It frames the key decision points: data volume, database/IFS complexity, ability to test before cutover, downtime tolerance, storage architecture (internal vs. external and availability of copy services), network/SAN throughput, and target location (on-premises, cross-datacenter, or PowerVS), and calls out constraints that can immediately rule out an option (for example, IASP considerations for Partition Mirroring or capacity demands of save/restore).

Use this guidance to narrow viable paths, verify technical eligibility, and plan the operational steps and cutover style that best fit your environment.

### Size of the migrated environment

The size of the environment is one of the key factors impacting both the migration approach and, most importantly, the total migration duration. Collecting information about system occupancy, available storage space, and disk technology allows you to decide about choosing the right migration method. The total volume of data directly affects synchronization time, bandwidth requirements, and finally the expected migration window and final cutover planning.

- Using save and restore migration, additional disk space near to doubled original capacity is required for storing virtual optical installation media and virtual tapes with user data. In case of large systems and internal storage, it could be hard to achieve.
- Choose Partition Mirroring or Manual Migration using storage copy services (FlashCopy or Remote Copy) for migration of large environments.

### Data analysis before migration

Each migration should begin with a deep analysis of the data to be transferred. Excessive or unnecessary data stored on the system disks extends both the duration of the initial system migration and the later user data synchronization phases.

#### ► Database files

The analysis must include the relationships between database files and libraries. This is especially important for SQL-based applications, where most objects are linked together and cannot be restored in isolation. These dependencies become critical during the data-synchronization phase. A large number of indexes, triggers, constraints, and SQL service programs can significantly affect the restore and synchronization processes.

#### ► System files

Cleaning up the operating system before migration allows you to isolate only the data required for replication and reduce the overall volume of data that must be migrated. These actions accelerate the initial save operation, improve data-synchronization performance, and shorten the cutover process to the migrated system. If not needed, remove the following objects:

- Performance data
- Journal receivers
- PTF save files and cover letters

- User save files
- Files located on IFS (ISO images, ZIP files)
- Spool files

Key points for selecting the migration method based on data analysis:

- ▶ In case of complex Db2 structures, SQL-based applications, and intensive use of IFS files, choose Partition Mirroring or Manual Migration using storage replication technologies.
- ▶ Save and restore migration operations become more complex and require additional attention. This method is more reliable for environments with simple data structures and smaller data volumes.
- ▶ Clean up your system from unused, redundant objects before starting migration.

### **Test migrated data**

Testing of the migrated environment is an essential part of every migration project. It provides a reliable way to verify whether the migrated system or applications operate correctly and whether the data are complete and consistent.

- ▶ Only the Partition Mirroring method offers the ability to start the copy node in test mode, to perform testing within the same migration process.
- ▶ Consider test migration in case of other methods (save and restore, different types of manual migration). This requires additional effort and repeating the same procedure, but it ensures that the migrated environment is fully validated and correctly replicated.

### **Validate source and target environments**

Both the source and target environments must be consistent in terms of disk technology, available server resources such as CPU and memory, and supported storage solutions. Identifying the common technologies and available resources makes it possible to determine the appropriate migration method.

### ***Storage type and location***

Determine whether the target system uses external storage, internal storage, or runs in IBM Cloud on PowerVS.

- ▶ Internal storage: choose between Partition Mirroring and save and restore methods, excluding manual migration using storage copy services. Both methods support migration of workloads from on-premises to IBM Cloud, between different data centers, or within the same data center.
- ▶ External storage: all migration methods can be used, including FlashCopy (within the same storage) and Remote Copy (between two storages) manual migrations.

**Note:** Workload migration from on-premises to IBM Cloud cannot be performed using storage replication services.

### **Migrations based on remote copy services**

IBM i Migrate While Active offers migration methods that use storage-based replication mechanisms such as FlashCopy and Remote Copy. These technologies minimize outage and enable fast system migration. Although FlashCopy and Remote Copy are typically used in PowerHA solutions, they are also successfully implemented for system migration. Main prerequisites for external storage copy services:

- ▶ Remote copy (Global Mirror, Metro Mirror):
  - Storage area network (SAN) infrastructure required.

- At least two IBM i partitions that are separated geographically with at least one external storage unit that is attached to each system.
  - Determine multipath I/O to improve overall resiliency and performance.
  - The same storage type for source and target volumes.
  - The same size of source and target logical volumes (or larger for target).
- FlashCopy:
- At least one IBM i system attached to external storage using Fibre Channel (FC).
  - Both the FlashCopy source and target systems must be attached to the same external storage device.
  - One set of disks for the source and an equal set of disk units for the target.

**Note:** In case of Migrate While Active, FlashCopy is used to create a temporary partition to perform a full system save on it. It is a beneficial method when the source node is using external storage, but the copy node is not. Warm or cold options for FlashCopy reduce downtime on the source system.

► Infrastructure and performance readiness

To consider migration successful, system performance, database response times, and communication between systems must remain at the same level or be better than on the source environment. The target server must provide enough CPU, memory resources, and overall performance capacity to run the migrated workload and maintain continuous synchronization. Before migration, validate key points on the target site:

- Server configuration resources (available CPU, memory, current firmware versions for server and I/O adapters).
- Storage capacity (the same or larger).
- Network throughput (LAN and WAN capabilities). Ensure network throughput is sufficient for the selected migration technique, especially when using save and restore.
- Storage Area Network (SAN) infrastructure and connectivity (switches, directors). Storage-based replication is much faster than typical Ethernet network connections because of Fibre Channel technology that facilitates the transfer of large volumes of raw block data.

### 2.3.3 Method-specific prerequisites and validation

This section translates the general readiness guidance into actionable, method-specific checklists and validation steps you must complete before synchronization and cutover.

It details what to prepare and how to prove readiness for each path: for Partition Mirroring, establishing the copy node and Service Tools LAN in DST, confirming GUI connectivity and required firewall ports, and verifying eligibility constraints (for example, no encrypted user ASP and handling of IASPs outside this method); for Assisted Save and Restore, planning the restricted-state save, orchestrating the automated network install/restore, and driving final synchronization via WRKMIGSTS options; and for Manual Migration, supplying storage replication or save/restore parameters, cluster and replication networks, and monitoring progress in the GUI. Use these steps to surface blockers early, document controls, and ensure a predictable, auditable path to cutover.

#### Partition Mirroring prerequisites and validation

Set up the copy node and configure the Service Tools server LAN adapter by using Dedicated Service Tools (DST) on the copy node as part of the Partition Mirroring flow. The Db2 Mirror GUI migration guide validates connectivity and verifies software requirements on the source

node before you proceed. IBM provides firewall port requirements between the source and copy node and between the GUI node and source/copy nodes for Partition Mirroring.

### Assisted Save and Restore prerequisites and validation

To save the source node, connect to the source console, enter restricted state, and start the automated save by using WRKMIGSTS option 1.

To install and restore the copy node, launch the network install from the Db2 Mirror GUI and then use the copy node console (DST and CL steps) and WRKMIGSTS option 10 to start the automated restore.

To finalize, initiate final synchronization from the Db2 Mirror GUI (Finalize) and then use the source console in restricted state to start and complete final synchronization (WRKMIGSTS options 20 and 21) before proceeding to cutover.

If you use network-attached tape or optical devices, IBM provides additional guidance to start TCP/IP and an interface while in restricted state so the devices are usable during save/restore processing.

### Manual Migration prerequisites and validation

In the Db2 Mirror GUI migration wizard, choose the Manual Migration type (Storage Replication, Save/Restore, or Storage Replication and Save/Restore) and provide node, cluster, and replication network details. Monitor progress in the Db2 Mirror GUI, which surfaces migration progress and provides actions such as Pause and Finalize depending on the stage.

Table 2-9 provides a decision summary to help guide the selection process.

Table 2-9 Decision point base on three migration concept

Decision point	Partition Mirroring	Assisted Save and Restore	Manual Migration
Why you choose it	Avoid an initial outage for creating install media or taking a full system save; optionally use test mode; choose cool or warm cutover.	Use a save/restore-based baseline with guided console steps for automated save, automated restore, and final synchronization.	Choose how you create the copy node (storage replication, save/restore, or both) within the wizard-supported Manual Migration types.
Key constraints that drive the choice	Not eligible with encrypted user ASP; IASPs not replicated using Partition Mirroring.	Requires an initial restricted-state save/restore flow and corresponding operational steps.	Must align to the wizard-supported Manual Migration patterns and required configuration inputs.

Table 2-10 on page 35 is decision aid, it's a compact way to evaluate which IBM i migration method will best work in your environment.

Table 2-10 Migration path comparison by key factors

Criteria	Partition Mirroring	Assisted Save and Restore	Manual Migration
Minimize downtime	No initial outage; creation of install media or a full system save is not necessary.	Outage for full system save.	Quiesced state for warm FlashCopy or Remote Copy; System power off for cold copies.
Type of storage	Independent; storage replication occurs within IBM i.	Internal or external storage.	External storage copy services (FlashCopy, Remote Copy).
Db2 database complexity	All object types supported; host-based storage replication implemented.	Needs special attention, control, and validation of objects and manual intervention.	All object types supported; external storage copy services used.
Test migration	Test mode is available with warm or cold option.	Testing on copy node is not possible. Separate migration needed for tests.	Testing on copy node is not possible. Separate migration needed for tests.
IASP support	Not supported.	Requires additional configuration of PowerHA clusters.	Requires additional configuration of PowerHA clusters.
Target environment	All locations with TCP/IP communication available. PowerVS in IBM Cloud supported.	All locations with TCP/IP communication available. PowerVS in IBM Cloud supported.	Only for external storage: FlashCopy within the same storage; Remote Copy to another storage. Storage copy services not supported between on-prem and IBM Cloud.

## 2.4 Enhanced Db2 Mirror GUI

The Db2 Mirror for i Graphical User Interface (GUI) serves as the primary management and monitoring console for both Db2 Mirror continuous availability environments and IBM i Migrate While Active (MWA) workloads. MWA, delivered in Licensed Program 5770-DBM Option 2, now fully leverages an enhanced Db2 Mirror GUI to orchestrate low downtime migrations of IBM i logical partitions (LPARs). These enhancements introduced progressively through 2024–2025 technology updates and PTF levels significantly modernize usability, automation, and operational visibility across migration workflows.

The GUI is a browser-based interface running on IBM i, deployed on the ADMIN3 Liberty-based application server, and accessed via:

`http://hostname:2001/Db2Mirror`

While the GUI can run on either Db2 Mirror node, IBM recommends hosting it on a third IBM i node to isolate administrative traffic, ensure availability during outages, and streamline multi-pair management. This third node must have adequate processor and memory resources, and Java 1.8 (64-bit) or newer.

## 2.4.1 Recent GUI Enhancements (2024–2025)

A series of improvements to the Db2 Mirror GUI, many delivered alongside MWA enhancements, have elevated the interface from a functional administrative console to a modern, intuitive orchestration platform.

Between April and November 2025, IBM delivered updates providing a clean, modern look, aligning the Db2 Mirror GUI with the broader IBM i management ecosystem. These improvements include updated icons, refined visual layout, and a more responsive user experience for multi-step workflows.

- ▶ The GUI now includes a dedicated Partition Mirroring Wizard, enabling the MWA “partition mirroring” method that uses host based storage replication under the TIMI layer to synchronize SYSBAS. Since this method requires no install media or full system save, the source node does not need to be brought offline during initial copy creation.
- ▶ The GUI supports a Test Mode for partition mirroring migrations, allowing administrators to validate the copy node before performing a final cutover without affecting source production workloads. (User-provided enhancement summary; consistent with MWA’s asynchronous model)
- ▶ Recent PTFs introduce visual progress indicators, improved status tracking, and clearer messaging for synchronization phases, test mode entry, and copy-node readiness. IBM’s GUI PTF updates also improved message formatting, error readability, and cutover related workflows.
- ▶ The GUI now uses RSYNC rather than SCP even for large datasets offering faster incremental transfers, and better efficiency during copy node provisioning. (User-provided enhancement summary; consistent with PTF enhancements focused on migration robustness).

## 2.4.2 Core GUI Functions Supporting Migration

The enhanced GUI acts as the operational hub for the entire migration lifecycle under Migrate While Active.

### Setup Wizard

The Setup Db2 Mirror / MWA Wizard guides administrators through:

- ▶ Node pairing
- ▶ Hostname/IP configuration
- ▶ Selecting a migration method
- ▶ Establishing initial replication

The wizard supports pause/resume, multi-pair administration, and is accessible from multiple GUI entry points (initial login, dashboard tray, or left navigation menu).

### Monitoring and Visualization

A dedicated dashboard provides real-time insight into:

- ▶ System Migration progress
- ▶ Continuous Data Synchronization backlogs
- ▶ Node health and communication status

These capabilities build on the GUI's role as a visual management interface for Db2 Mirror environments.

### Operational Controls

Users may take the following actions directly within the GUI:

1. Pause or resume synchronization
2. Enter Test Mode
3. Initiate Finalize / Final Cutover
4. Review error logs and system events

Enhanced PTF functionality improves reliability during quiesce, test mode, and warm cutover conditions.

Management Access Requirements:

- ▶ Accessed via **http://hostname:2001/Db2Mirror** (ADMIN proxy -> Liberty application server).
- ▶ GUI requires Java 1.8 64-bit or newer.
- ▶ ADMIN3 must be active.

## 2.4.3 Migration Methods Supported Through the GUI

The enhanced Db2 Mirror GUI unifies all MWA migration methods into a single interface:

1. Partition Mirroring (Preferred)
  - Uses host based storage replication below TIMI
  - Avoids initial outages, no install media creation required
  - Enables Test Mode prior to cutover
2. Assisted Save and Restore
  - Combines Db2 Mirror replication with save/restore media
  - GUI orchestrates data transfer, restore operations, and ongoing synchronization
3. Manual Migration
  - Supports FlashCopy-based workflows
  - Allows manual save/restore operations
  - MWA continues tracking and synchronizing changed data

**Note:** When the system enters a restricted state where the GUI cannot operate, administrators manage the migration through the WRKMIGSTS (Work with Migration Status) CL command.

The enhanced Db2 Mirror GUI now serves as a unified, modern platform for both continuous availability management and low downtime migration via Migrate While Active. With its updated interface, expanded wizards, RSYNC support, Test Mode integration, and improved progress visualization, the GUI provides an end-to-end management experience that reduces complexity and operator workload during mission-critical IBM i migrations.

## 2.5 Installing RSYNC using the Open Source Package Management interface

RSYNC is delivered as an open-source package on IBM i and is installed using the Open Source Package Management (OSPM) framework available through IBM i Access Client Solutions (ACS). The OSPM interface provides a standardized, graphical way to install, update, and maintain open source utilities such as RSYNC.

The OSPM framework is built on the YUM-based infrastructure introduced for IBM i 7.3 and later. This environment handles dependency resolution, package updates, and repository management.

Before installing RSYNC, the IBM i system must have:

- ▶ A supported IBM i release (7.3 or later supports the OSPM tool) [ibm.com]
- ▶ The /QOpenSys environment configured
- ▶ Access to IBM i Open Source RPM repositories
- ▶ Adequate storage in the Integrated File System (IFS) for package installation

Using ACS, administrators can install RSYNC without manually compiling binaries or managing dependencies. The package management tool automatically installs RSYNC along with any required supporting packages, placing everything into the IBM i Portable Application Solutions Environment (PASE). [ibm.com]

### Detailed Steps to Install RSYNC Using ACS OSPM

#### 1. Launch ACS and Open OSPM

- Start IBM i Access Client Solutions on your workstation.
- Navigate to:  
Tools -> Open Source Package Management (or from the left navigation under “Management,” depending on ACS version).

#### 2. Authenticate Using SSH

OSPM connects to IBM i using SSH.

- When prompted, enter a user profile with \*ALLOBJ special authority.
- If the system’s SSH daemon (SSHD) is not running, IBM provides guidance for starting it.

#### 3. Install the Open Source Environment (If Not Yet Installed)

If this is the first time using OSPM, ACS will prompt to install the underlying YUM infrastructure.

Respond Yes to install it.

After installation, the OSPM window shows the list of installed packages.

#### 4. Install RSYNC

- Select the Available Packages tab.
- Scroll or search for rsync.
- Highlight the package and choose Install.
- OSPM will display a package installation screen showing RSYNC and any dependent packages.

- Type Y to confirm installation.

The installation process automatically handles dependencies and places binaries in the appropriate /QOpenSys locations.

### **Command-line Alternative (YUM in PASE)**

Experienced administrators may choose to use a terminal instead of ACS.

Requirements:

- ▶ Use QP2TERM or SSH to access PASE.
- ▶ Ensure the YUM environment is configured and functional.

Common commands:

```
Shell yum install rsyncS
```

This installs RSYNC and all necessary dependencies automatically.

### **If the System Does NOT Have SSH Access (Offline or Locked-down Environments)**

In restricted environments where SSH cannot be used, IBM community tooling provides an automated bootstrapping method that downloads the required YUM and RPM infrastructure through FTP and initializes OSPM.

A script-based installation method is available and can be run through ACS Run SQL Scripts to bootstrap the open-source environment.

This is typically used for:

- ▶ Air-gapped systems
- ▶ Hardened security environments
- ▶ Systems where administrators cannot enable SSH temporarily

Once the bootstrap is complete, OSPM works normally and RSYNC can be installed through ACS.

### **Repository Configuration and Offline Installation Options**

IBM also provides a method to clone the IBM i Open Source repository to an on-premises IFS location and use it as a local mirror for multiple systems.

This is useful in bandwidth-restricted or enterprise environments requiring controlled changes.

ACS OSPM supports:

- ▶ Cloning the repository
- ▶ Switching package sources
- ▶ Serving local repositories to other partitions

This ensures RSYNC and other packages can be installed even without Internet access.

For detailed instructions on configuring and using the Open Source Package Management environment, including repository setup and basic package operations, refer to the following IBM i documentation:

<https://ibmi-oss-docs.readthedocs.io/en/latest/yum/README.html>

## 2.6 Security highlights of IBM i MWA

This section explains how to keep an IBM i Modern Workload Architecture (MWA) migration secure especially when it spans on-premises systems and IBM Power Virtual Server (PowerVS). Think of the security story in three parts:

- ▶ Protect the control plane:

The web GUI you use to configure and run MWA must be locked down so credentials and settings aren't exposed.

- ▶ Protect the data paths:

Some stages encrypt traffic by design, while others don't. When encryption isn't built in, we add a secure overlay.

- ▶ Control identities, authorities, and keys:

Only the right people and processes should have elevated access, and any stored secrets must be protected.

In practice, this means turning on TLS for the GUI, using SSH where it's available, and applying VPN or strict network segmentation in the stages that don't support encryption. For hybrid connectivity to PowerVS, we rely on private links and encrypted tunnels to keep migration traffic off the public internet and isolated from other workloads.

By the end of this section, you'll know what to secure, when it needs to be secured, and the minimum steps to do it without changing how MWA works.

### 2.6.1 Secure Administration & Control Planes

Secure every admin touchpoint (GUI and admin services) used to operate MWA/Db2 Mirror, and keep those endpoints reachable only over encrypted, private paths

Turn on HTTPS/TLS for the MWA/Db2 Mirror GUI (Admin3)

- Open IBM Web Administration for i at <http://<host>:2001/HTTPAdmin>, select Admin3, and run Configure TLS. The IBM wizard can disable the old 2001 port and bring up the HTTPS port (2010) automatically.
- Db2 Mirror's security guidance explicitly recommends HTTPS for the GUI; without TLS, browser sessions (logins, configs) are sent in clear text.

After successful login to IBM Web Administration for i, select Application Servers and from drop-down list of the servers chose Admin3. From left navigation pane launch TLS configuration wizard as shown in Figure 2-3 on page 41.

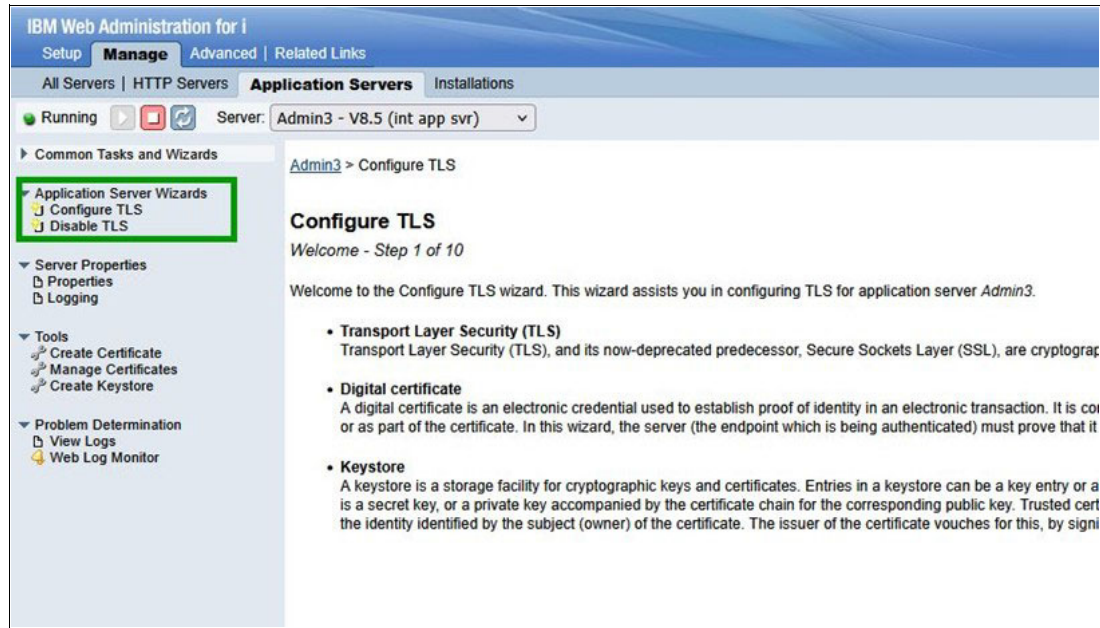


Figure 2-3 IBM Web Administration for i TLS configuration screen

Why this matters: It prevents credential/config interception and sets the baseline for secure administration.

**Note:** If TLS is not configured, non-encrypted connections are used, which means all communication between browser and DB2 Mirror GUI is transmitted in clear text.

For more step on how to config TLS, pls go to this url -

<https://www.ibm.com/support/pages/configure-host-servers-tls-using-navigator-i>

## 2.6.2 Protect the Data Paths

To ensure secure operations, every migration or data-movement hop must be encrypted by design or encapsulated within a private, encrypted overlay when native encryption is not available. IBM i MWA uses the Secure Shell (SSH) protocol to protect all data transfers during the system migration stage.

During automatic installation media transfer, the source node establishes an encrypted SSH session with the NFS server. As part of the migration configuration process, an SSH key must be created on the source node for the migration user profile.

What to do (plain steps):

1. Use SSH during the system-migration stage.

Migrate While Active (MWA) uses SSH for media/NFS transfers in the system-migration stage; configure key-based auth for the migration user so the traffic is encrypted end-to-end.

2. Add a secure overlay for synchronization traffic.

Data synchronization uses RDMA over TCP/IP and cannot be TLS-secured; protect it with a site-to-site VPN (or run only over strictly isolated private networks/VLANs).

3. Switch GUI <-> node connections to TLS as soon as certs exist.

Early in setup the GUI talks to nodes without TLS; once host servers are TLS-enabled, flip those connections to secure (TLS) in the GUI preferences.

4. Keep all admin/data paths on private connectivity to PowerVS.

Use VPC site-to-site VPN for encryption, or IBM Cloud Direct Link for private, off-Internet paths (plan a second, diverse Direct Link for resilience).

5. Enforce modern TLS everywhere TLS is used.

On IBM i, set QSSLPCL to allow only TLS version 1.3/1.2 and disable weak protocols/ciphers per IBM guidance.

Table 2-11 shows quick reference (what's encrypted vs. what to add).

Table 2-11 Quick reference for data protect

Path / Stage	Protocol	Encrypted by design?	Your action
System migration (media/NFS)	SSH	Yes	Use key-based auth; manage keys securely.
Data synchronization	RDMA over TCP/IP	No	Overlay site-to-site VPN / keep on isolated private network.
GUI <-> nodes (early)	Host services (no TLS yet)	No	Keep on private/VPN until certs exist, then enable TLS
Hybrid path to PowerVS	VPC VPN / Direct Link	VPS:Yes; DL:Private path	Prefer policy-based VPC VPN; use Direct Link for private reach. (Order two DLs for diversity.)

Detailed instructions for generating and configuring this SSH key are provided in Step 7 of the [Getting Started section of the MWA documentation](#).

### Secure and Non-Secure GUI communication with MWA nodes.

Db2 Mirror web-based interface with IBM i Migrate While Active option can be run on node that is part of Migrate While Active environment or on separate IBM i system called managing node or GUI node.

During the system migration stage, the GUI node communicates with the source and copy nodes only through non-secure connection to IBM i host servers. During this time the environment is still in initial, pre-synchronization state and several key components required for secure TLS communication don't exist (copy node), are not yet available or fully configured, including certificate store, private keys, trusted CA and assignment of certificates functionality.

Other stages, synchronization and cutover allow you to configure TLS and then secure connection between GUI node and source/copy node.

To properly setup TLS for IBMi host servers see documentation available here: [https://www.ibm.com/docs/en/ssw\\_ibm\\_i\\_76/rzahh/ss1req.htm](https://www.ibm.com/docs/en/ssw_ibm_i_76/rzahh/ss1req.htm)

Figure 2-4 shows sample UI for automatic transfer data.

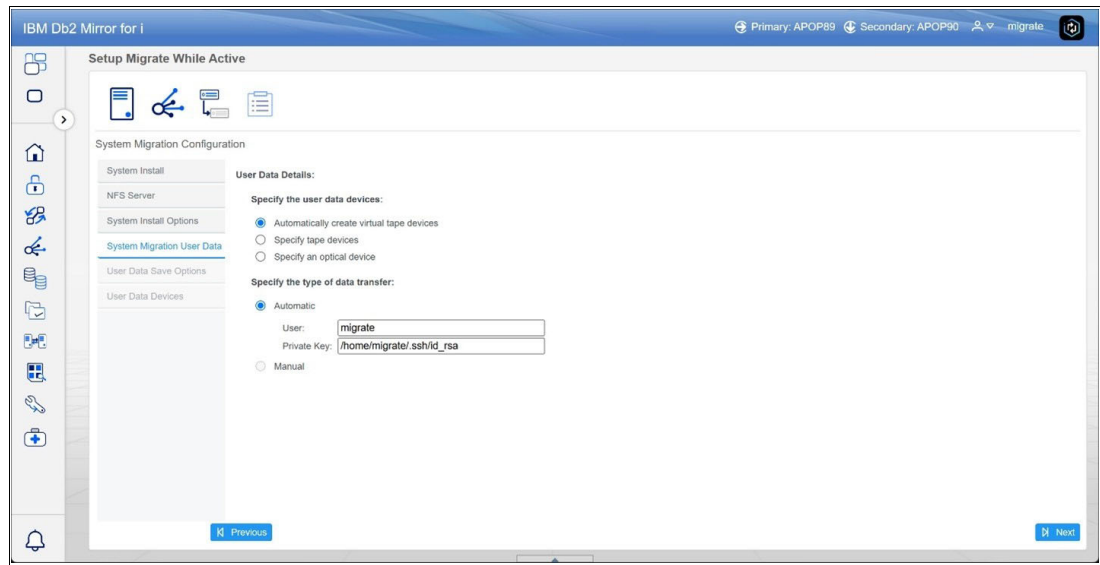


Figure 2-4 Setup MWA transfer method

The partition mirroring migration method uses a protocol that does not encrypt the data transferred from the source node to the copy node and cannot be secured with TLS. You must implement an alternative protection or encryption mechanism, such as virtual private network (VPN) between the nodes or dedicated VLAN's for replication purposes.

During synchronization stage Migrate While Active uses the RDMA over TCP/IP protocol to transfer data from the source node to the copy. This protocol does not provide encryption and cannot be secured with TLS. Similarly to partition mirroring, you must use different protection methods, for instance VPN.

### 2.6.3 Control Identities, Authority, and Keys

To maintain strict operational security, only appropriately authorized users should perform powerful migration actions, and all credentials or stored secrets must be protected. The Migrate While Active function requires specific user authorities and password-security settings to ensure that migration operations are executed safely and reliably. Proper authorization is essential for configuring system components, managing partition mirroring, and accessing the sensitive functions involved in the migration workflow.

MWA leverages IBM Cryptographic Services for i to securely store passwords that are needed during the migration process. In most environments, Cryptographic Services master key 1 is configured by default. However, this setting can be reviewed or updated through IBM Navigator for i, accessible from a web browser at:

[http://server\\_ip:2001/](http://server_ip:2001/)

where **server\_ip** is the host name or IP address of your IBM i system.

What to do (plain steps):

1. Use the right user profiles and least privilege.

The GUI uses the signed-in user's credentials to act on each node; if that profile lacks authority, the request fails so, create dedicated migration profiles with only the access they need (avoid using QSECOFR for routine steps).

2. Enable auditing for privileged activity.

Follow the Db2 Mirror/MWA guidance to review security/auditing and ensure privileged actions during migration are captured for review.

3. Load and set Cryptographic Services Master Key1 before storing any passwords.

Db2 Mirror/MWA store HMC, storage, and other passwords using IBM Cryptographic Services; Master Key1 must be 'loaded and set' on the GUI node and participating nodes first.

4. Manage master keys correctly (custody, backup, change).

IBM i keeps master keys in LIC, uses them (AES-256) to wrap other keys, and records key ops in CY audit records. Apply separation of duties, back up via SAVSYS, and retain passphrases securely. Configure and manage keys in Navigator for i -> Security -> Cryptographic Services Key Management.

5. Keep TLS posture strong for any identity-bearing services.

Where TLS is used (GUI, host servers, APIs), pin to TLS 1.3/1.2 and disable weak ciphers/protocols with IBM's system-wide settings.

The **Security** → **Cryptographic Services** → **Key Management** interface in IBM Navigator for i is used to manage master keys and cryptographic keystore files.

Figure 2-5 shows sample UI for cryptographic Services Key Management.

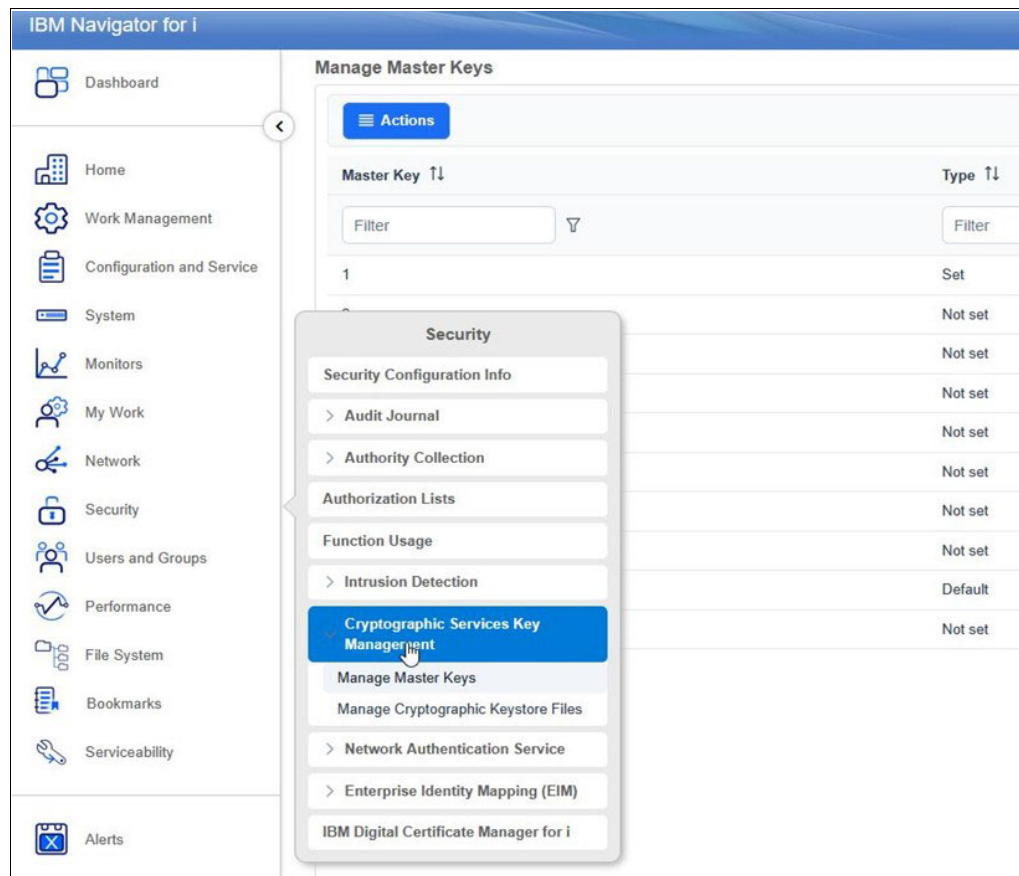


Figure 2-5 Cryptographic Services Key Management options

IBM i Migrate While Active requires proper user authorities depending on migration method. Table 2-12 shows relations between migration methods and user authorities.

Table 2-12 Special user authorities for MWA

Migration method	User Authorities
Partition mirroring	<p>On source node: *SERVICE, *JOBCTL, *ALLOBJ, *SECADM            DST authority on copy node:</p> <ul style="list-style-type: none"> <li>▶ Disk units - operations</li> <li>▶ System partitions - operations</li> <li>▶ System partitions - administration</li> <li>▶ Operating system initial program load (IPL)</li> <li>▶ Install</li> <li>▶ Display/Alter/Dump</li> <li>▶ Dedicated Service Tools (DST) environment</li> <li>▶ Service tool security</li> <li>▶ Start Service Tools</li> </ul>
Assisted save and restore	<ul style="list-style-type: none"> <li>▶ On source node creation of user for migration required instead of using QSECOFR</li> <li>▶ Special authorities:</li> <li>▶ *ALLOBJ, *IOSYSCFG, *JOBCTL, *SAVSYS, *SECADM, and *SERVICE</li> <li>▶ User class:*QSYSOPR</li> </ul>

For more information about managing master keys and using IBM Cryptographic Services, refer to the IBM i documentation:

<https://www.ibm.com/docs/en/i/7.6.0?topic=cryptography-cryptographic-services-key-management>

## 2.6.4 Cloud Connectivity & Network Isolation for IBM Power Virtual Server

Power Virtual Server offers several methods to secure and isolate transfer of data during migration process. Private networking, VLAN segmentation, VPCs, and security groups actively control traffic flows and block unauthorized communication between workloads.

IBM Cloud Direct Link(2.0) provides secure, 10 Gbps connections between on-premises datacenter and IBM Cloud services. It offers predictable network performance and enhanced security, without using the public internet.

There are two types of Direct Link:

- ▶ Direct Link Dedicated provides a physical port connection for private customer access. It serves clients with high volume workloads that require low latency and supports rapidly moving data across environments.
- ▶ Direct Link Connect provides connectivity between clouds datacenters or between on premises and cloud. It supports flexible bandwidth requirements. It also offers IBM Cloud clients a lower cost entry point compared to other Direct Link options. Direct Link uses a connection that is set up and managed by an external service provider.

By default, IBM Cloud Direct Link is not a redundant service. You must order a separate Direct Link Connect instance for redundancy.

Virtual Private Network services provide private and low-cost connectivity to IBM Cloud. You can use SSH connections to access your virtual server instance through private IP address. There are two options for VPN connection:

- ▶ Site-to-site gateway connections (VPN for VPC ) allow to connect on-premises network to your Virtual Private Cloud (VPC) network. This configuration provides secure, highly scalable, encrypted connectivity based on secure VPN tunnels with IPsec configured.
- ▶ Client-to-site server connectivity is dedicated for remote, secure access to VPC network using an OpenVPN software. It is TLS 1.2/1.3-based secure/encrypted connectivity over the internet.

To conclude this section, Table 2-13 shows a summary of IBM i Migrate While Active security aspects.

Table 2-13 Behavior by stage - Security summary

Stage	Primary Protocol(s)	Encryption Native?	Required Compensating Controls	Notes
GUI Setup (Admin3)	HTTPS/TLS (recommended)	Yes (when enabled)		Enable TLS before operations; plaintext if not configured.
Early GUI <-> Nodes	IBM i host servers (plaintext)	No (temporary)	Segmented VLAN, private IP, VPN	TLS becomes available at later stages once certs are ready.
System Migration (media/NFS)	SSH	Yes	Key hygiene, least-privilege user	Use SSH keys for migration user.
Partition Mirroring	Native mirroring protocol	No	VPN, dedicated VLAN, private routing	TLS not supported; isolate strictly.
Synchronization	RDMA over TCP/IP	No	VPN, dedicated VLAN, private routing	TLS not supported; treat as sensitive bulk movement.
PowerVS Hybrid Connectivity	Direct Link, VPN	Yes (VPN), private path (DL)	Redundant DL instances; security groups; private-only IP	DL not redundant by default; order a second instance



## Migration method

As introduced in Chapter 1, IBM i Migrate While Active (MWA) provides three distinct migration modes, each designed to address different operational requirements and workload characteristics. These modes offer flexible approaches for transitioning IBM i environments with minimal disruption, allowing clients to select the method that best aligns with their infrastructure, time lines, and business objectives.

This chapter provides a detailed examination of each method. It describes the underlying technologies, outlines the operational workflow for each approach, and identifies the scenarios in which each method is most effective. By understanding the strengths and intended use cases of each migration mode, clients can plan migrations with greater confidence and predictable outcomes.

The following topic will be covered in this chapter:

- ▶ “Partition Mirroring”
- ▶ “Assisted Save and Restore”
- ▶ “Manual Migration”

## 3.1 Partition Mirroring

Partition Mirroring is the premiere migration method offered by IBM i Migrate While Active. This migration mode leverages the technology underlying PowerHA SystemMirror for i's Geographic Mirroring, but allows the solution to be used for the System ASP (\*SYSBAS). This revolutionary development provides the most simplified, automated solution for migrating IBM i workloads to new systems.

Figure 3-1 is a high level view of the partition mirroring option.

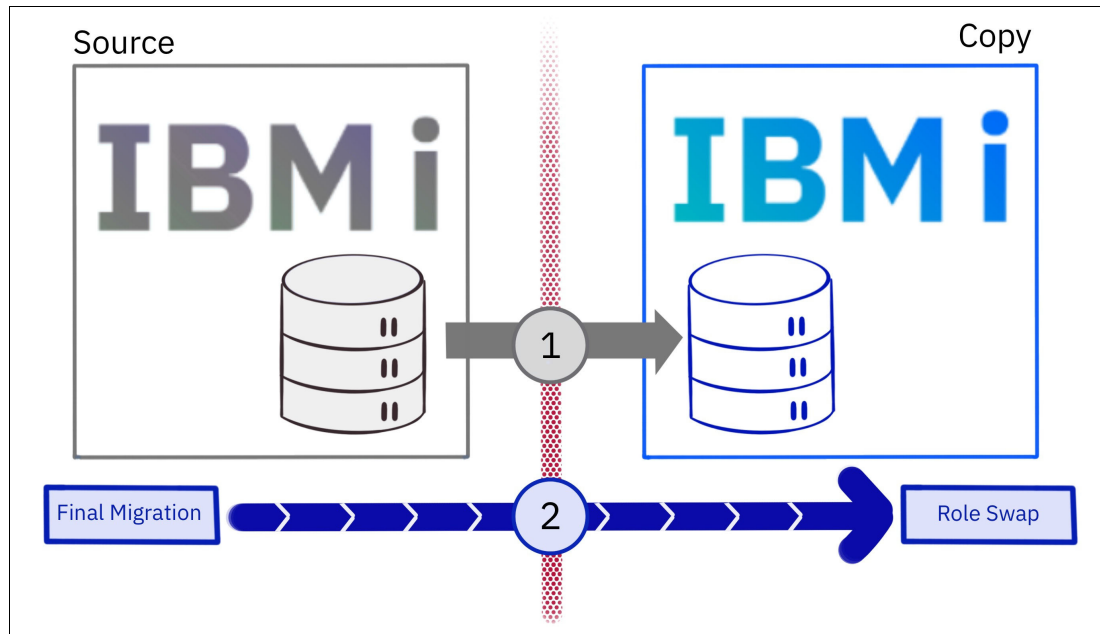


Figure 3-1 Diagram of partition mirroring

Partition Mirroring requires that the Source Node have the latest Migrate While Active and Db2 Mirror for i PTFs installed. As a first step, the IBM i administrator creates SAVSYS media to virtual optical devices. By only needing to create SAVSYS media instead of a Full System Save, the downtime required is minimized. Alternatively, SAVSYS media can be created on a non-production system as long as the required PTFs are installed. The target system, known as the Copy Node, is then installed from this SAVSYS media. Only the Licensed Internal Code (LIC) needs to be installed. Once the DASD and Service Tools LAN Adapter are configured, the Migrate While Active GUI assists in configuring Partition Mirroring between the Source Node and the Copy Node. Data transfer can then begin, and Partition Mirroring will mirror the System ASP to the Copy Node. Once the data is synchronized, Migrate While Active will present the IBM i administrator with the option to perform a test role swap, or full role swap.

Partition Mirroring is the preferred Migrate While Active migration method due to the simplicity of setup needing only LIC to be installed on the target system. IBM recommends that customers leverage Partition Mirroring as the primary migration solution except in specific situations as described in the other migration methods.

## 3.2 Assisted Save and Restore

Assisted Save and Restore was introduced in the first release of IBM i Migrate While Active. This method uses IBM Db2 Mirror for i technology to replicate objects between the source and copy nodes over Ethernet, ensuring the copy remains up to date throughout the migration.

Migrate While Active is particularly well-suited for migrations where traditional D-Mode IPL and physical tape restore operations are not practical. The solution performs the save to optical media, enabling a network-based installation path rather than a tape-based workflow.

The overall process includes three major phases as shown in Figure 3-2:

- ▶ System migration – Establishes the copy node by creating the initial system copy (Step 1).
- ▶ Continuous synchronization – Keeps the copy node updated with changes from the source node (Step 2).
- ▶ Promotion and cutover – When the administrator decides to stop synchronization, the copy node is promoted and made ready for use (Step 3).

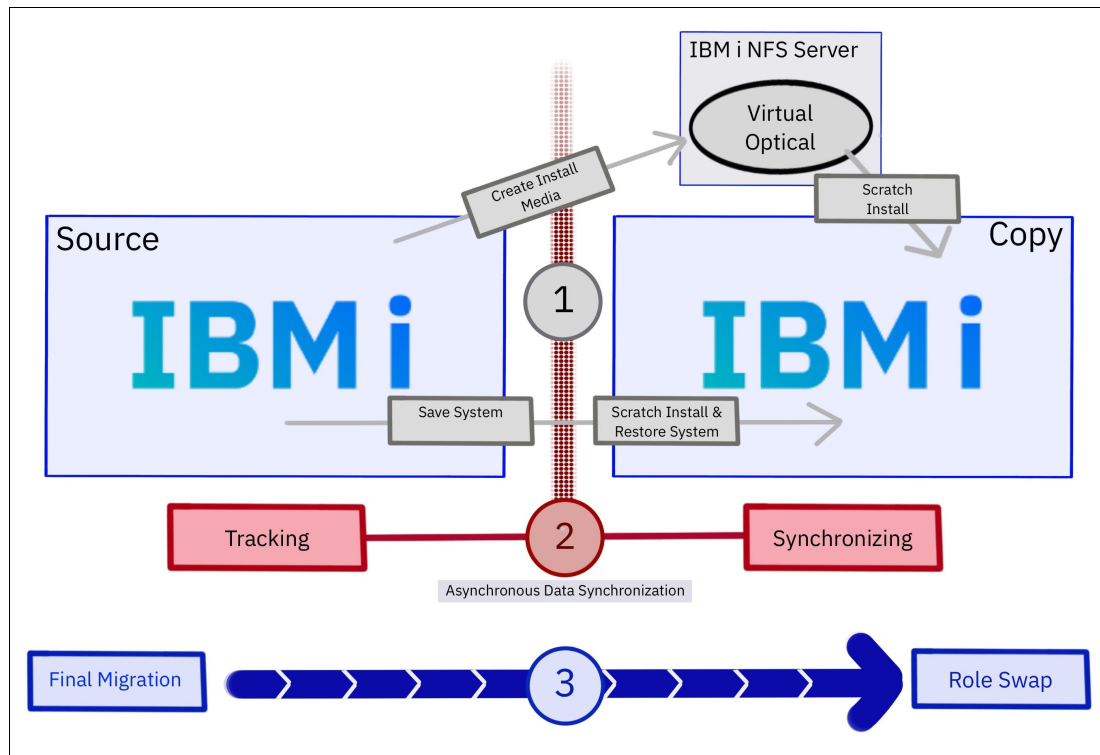


Figure 3-2 Save and restore migration

The Migrate While Active GUI guides the IBM i administrator through configuring the Assisted Save and Restore environment. As part of this setup, the solution deploys the Source Node, the Copy Node, and an IBM i-based NFS Server Node. The NFS Server Node is used to receive a copy of the backup media and later serves as the source for the network installation of the Copy Node.

After the environment is in place, Migrate While Active leads the administrator through creating either a Full System Save or Partial System Save to virtual optical media. When the save operation completes, normal production workloads can resume. Migrate While Active then handles transferring the save media to the NFS Server Node. Once the media is

available on the NFS server, the solution automates the setup of the network-install environment and guides the administrator through performing the installation and restore on the Copy Node. This ensures that both the Source Node and the Copy Node begin from an identical baseline.

When the Copy Node is fully deployed, Migrate While Active automatically configures Db2 Mirror for i functionality to keep the Copy Node synchronized with ongoing changes from the Source Node until the business is ready to perform the role swap.

Assisted Save and Restore is often an attractive option because most of the System ASP is seeded on the Copy Node during initial setup. This can significantly reduce network bandwidth requirements and, for high-throughput systems, can lower the processing load required to maintain synchronization between the two nodes.

### 3.2.1 Save and restore use cases

IBM i Migrate While Active allows organizations to migrate an IBM i partition from one environment to another while keeping production workloads online. The solution supports a wide range of migration scenarios, enabling businesses to modernize infrastructure, adopt cloud strategies, or create reusable clones with minimal disruption. IBM explicitly notes that Migrate While Active can be used to move an IBM i partition within the same server, to another data center, or to remote environments such as IBM Power Virtual Server in IBM Cloud.

Migrate While Active supports several key use cases:

#### 1. Migrate from On-Premises to IBM Power Virtual Server

Organizations can move an on-premises IBM i LPAR to IBM Power Virtual Server, enabling cloud adoption for production, dev/test, or hybrid-cloud deployments. This use case is explicitly documented by IBM, which highlights PowerVS as a common destination for cloud-based IBM i workloads.

Customers increasingly choose this option to eliminate tape-based restores, accelerate cloud adoption, and maintain uptime during the migration.

#### 2. Migrate from On-Premises to Any IBM i Cloud Provider

In addition to IBM's own cloud, Migrate While Active can migrate a production LPAR into any IBM i capable cloud environment. [IBM documentation](#) lists this scenario as a standard supported use case, enabling customers to move to managed IBM i hosting environments or third party cloud providers.

#### 3. Migrate to Newer or More Advanced Power Systems

Enterprises frequently use Migrate While Active to refresh hardware without extended downtime. IBM states that you can migrate an on-premises IBM i to newer hardware within the same data center or server room.

This is particularly valuable when upgrading to modern POWER10 systems or consolidating workloads onto newer infrastructure.

#### 4. Create an IBM i Clone for Testing, Sandboxes, or Projects

Migrate While Active can duplicate a running production system to create a clone. The production node stays fully operational, while the copy node becomes an isolated system for test environments, analytics, development, audits, or short-term projects. IBM lists this scenario as a built-in use case, where the production node remains intact and serves as the authoritative source for the clone.

### 5. Migrate Between Data Centers

According to IBM documentation, Migrate While Active is commonly used to move IBM i workloads between physical data centers, whether to support consolidation, relocation, or disaster-recovery site modernization.

This includes migrations across geographic regions when using TCP/IP connections between nodes.

### 6. Hybrid Cloud Transition and Infrastructure Realignment

Industry reports emphasize that many IBM i customers use Migrate While Active specifically to bridge between on-prem systems and hybrid cloud architectures. The solution helps organizations overcome challenges such as tape logistics, slow network replication methods, or lack of direct FlashSystem-to-PowerVS replication.

### 7. Seamless Platform Modernization with Continuous Availability In Mind

Although not an HA/DR solution, Migrate While Active builds on the same technical foundation as IBM Db2 Mirror for i. This makes it desirable for businesses that need uptime during infrastructure refreshes or cloud transitions while still keeping environments synchronized until cutover.

Migrate While Active enables you to migrate from one IBM i to another. This section explores

As shown in Figure 3-3, Migrate While Active can be used to achieve a wide variety of business needs:

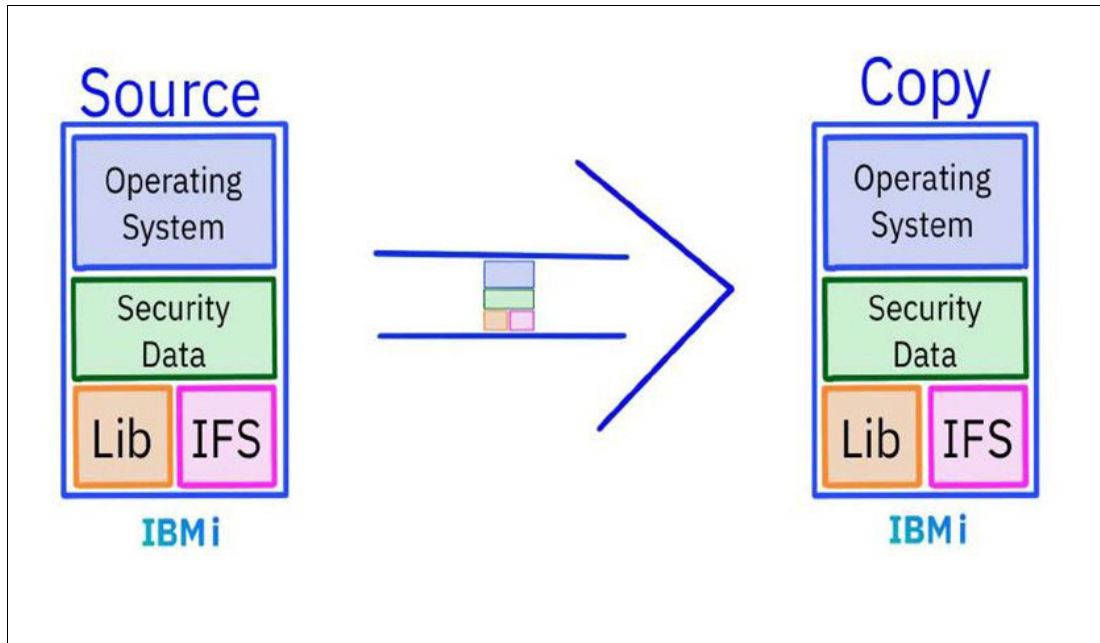


Figure 3-3 MWA use case

## 3.2.2 Anatomy of save and restore migration

Migrate While Active makes it easier than traditional approaches to migrate an IBM i. This section provides the first glimpse into the steps that take place when migrating an IBM i.

As depicted in Figure 3-4 on page 53, the migration of an IBM i is broken down into three stages:

- ▶ Stage 1 - System migration

During the system migration stage, Migrate While Active is configured and a system save is performed on the source node. The save includes the operating system and some or all of the existing user data. The system save is done while in restricted state. The Migrate While Active user indicates how much user data to save. All user data saves are focused on objects within SYSBAS.

A unique version of install media is created during the system save and used to scratch install the copy node. Then any saved user data is restored onto the copy node.

Decisions that need made before system migration include:

- Which migration pattern will be used?
- How much user data will be saved?
- Will the user data be saved on the source node or externally?
- When will the system migration be started?

**Tip:** Use the GUI to understand the storage size estimations for the save options

► Stage 2 - Data synchronization

After the system save of the source node is complete, users and applications can run production workloads on the source node. During this time, objects within SYSBAS libraries and the integrated file system can be created, changed, or deleted. Any change to objects is noticed by Migrate While Active and tracked.

After the system migration is complete, Migrate While Active begins the continuous synchronization of changed objects that is known as data synchronization. During this time, Migrate While Active is continuously keeping the copy node "in sync" with changed objects from the source node.

While in the data synchronization stage, changes to objects on the source node are tracked and asynchronously migrated to the copy node. Data synchronization is not done in time order, but instead occurs in intervals which are controlled by Migrate While Active. When a synchronization interval begins, Migrate While Active starts batch jobs that handle the data synchronization to the copy node.

Insights to be gained during data synchronization include:

- The volume, size, and type of objects being changed and tracked on the source node
- The count, size, and velocity of data synchronization
- The time estimate for the final synchronization outage

► Stage 3 - Cutover

The Migrate While Active user decides when to leave the data synchronization stage and finalize the migration to the copy node.

During the cutover stage, the source node enters restricted state and a final synchronization of any remaining tracked objects occurs. After the final synchronization is complete, the Migrate While Active user decides whether to complete the migration. Upon successful completion of cutover, data synchronization between the source and copy nodes is ended and the copy node is ready to be used.

Considerations within cutover include:

- Examination of any objects that failed to migrate to the copy node
- Examination of any objects that were ineligible to migrate to the copy node

Evaluation of whether to complete the migration to the copy node.

This is shown in Figure 3-4.

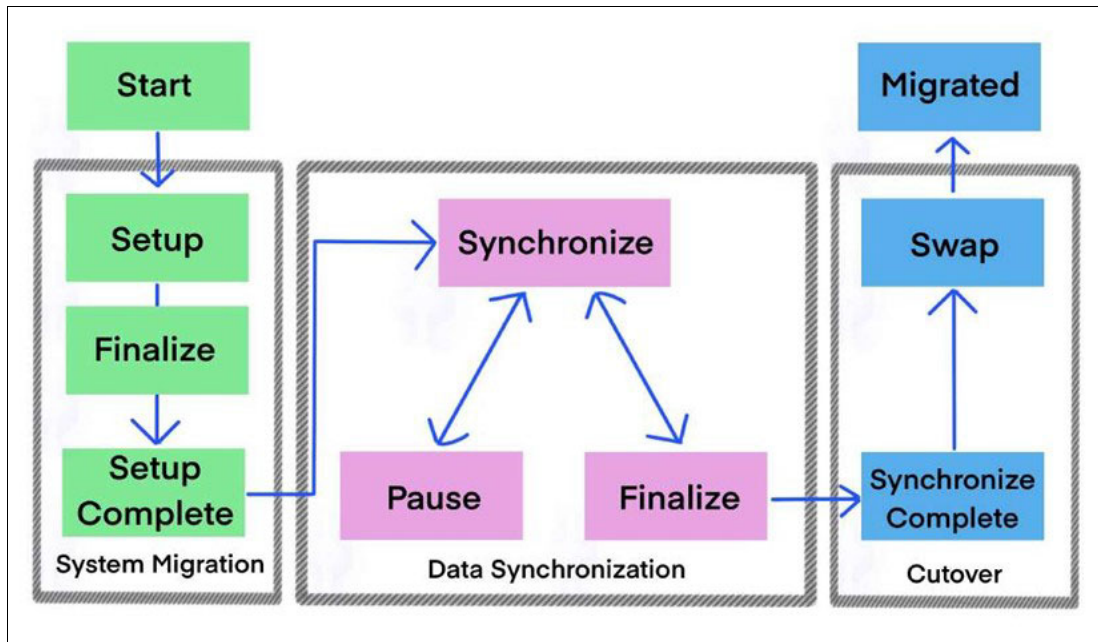


Figure 3-4 Migrate While Active stages and phases

### 3.2.3 Automated Transfer Operations in Migrate While Active

IBM i Migrate While Active (MWA) automates the movement of install media, user-data saves, and other synchronization content between the source node, the Cloud Optical Repository (NFS server), and the copy node. This automation reduces manual effort, accelerates migration activities, and keeps both systems consistently prepared throughout the migration.

There are two automation examples that will be explained in detail:

- ▶ Fully automated workflow
- ▶ Hybrid automated workflow

Figure 3-5 on page 54 shows a completely automated workflow for both install media and user data transfers.

The source system generates the virtual optical install media, which MWA transfers automatically to the Cloud Optical Repository, enabling the copy node to perform a scratch installation directly from this image. In parallel, the source saves user data to virtual tape, and MWA transfers this tape image to the copy node, where it is immediately ready for restore.

The entire process requires no physical media handling, with all transfers occurring over the network by using NFS and MWA automation. This migration pattern is used when the environment supports:

- ▶ Creation of virtual optical media
- ▶ Creation of virtual tape images
- ▶ Network-based automated transfers using NFS and MWA automation mechanisms

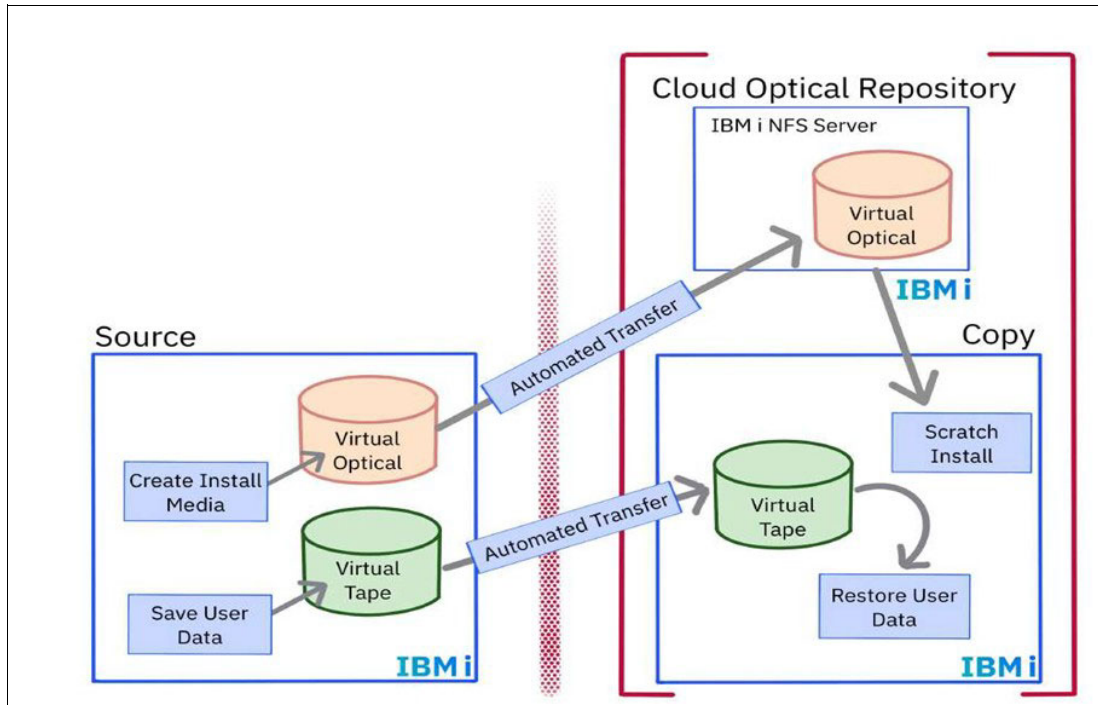


Figure 3-5 . Automated migration of an IBM i to IBM Power Virtual Server using source node local storage

Figure 3-6 on page 55 shows a hybrid Assisted Save and Restore process where the install media is transferred automatically, but the user data requires manual handling. The automated portion works the same way as in the previous example: the source system creates virtual optical media, it is transferred to the Cloud Optical Repository, and the copy node performs the scratch installation from that repository.

The difference is in how user data is managed. Instead of using virtual tape and automated network transfer, the source system saves the user data to physical tape or optical media. The user must then manually move this media to the copy node, where the system restores the data directly from the delivered tape or optical source.

This pattern is used when:

- ▶ On-premise migrations involving physical tape workflows
- ▶ Environments not using virtual tape
- ▶ Limited automation capabilities
- ▶ Migrations where user data is too large for network transfer

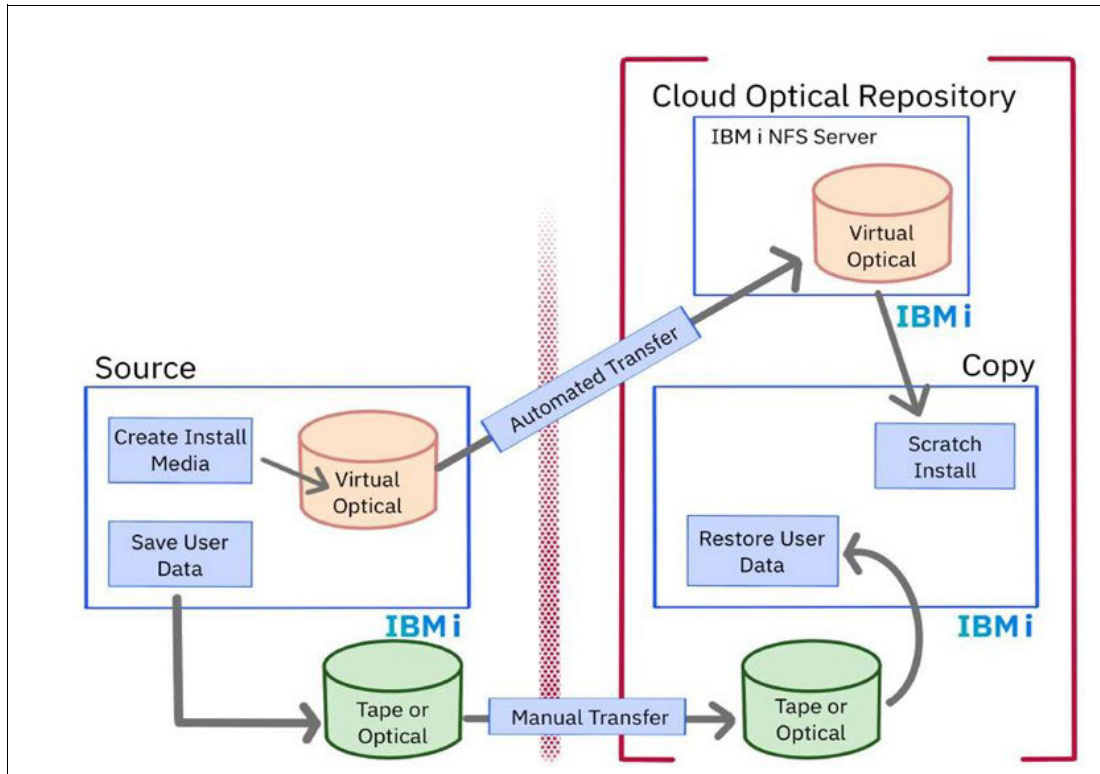


Figure 3-6 Automated migration of an IBM i to Power Virtual Server using user provided devices for user data

Table 3-1 is a side by side comparison between these options.

Table 3-1

Feature	Figure 3-5	Figure 3-6
Install media transfer	Automated	Automated
User data transfer	Automated via Virtual Tape	Manual via Tape/Optical
Physical media	None required	Required for user data
Automation level	Full automation	Partial automation
Best for	Cloud migrations, modern environments	Legacy tape workflows, constrained network

### 3.3 Manual Migration

The Manual Migration method provides the highest level of flexibility, making it ideal for environments with unusual requirements or very large data volumes. This flexibility comes with added complexity, as it relies less on automation and more on administrator-driven processes.

Four manual migration approaches are available, each suited to different storage configurations and downtime considerations. Manual Migration also allows the use of additional techniques during the Copy Node creation stage:

► Storage Replication

When the source and copy nodes share the same external storage, FlashCopy technology is used to create a point-in-time copy of the source node volumes directly to the copy node volumes. This approach is typically used for local migrations (see Figure 3-7).




Figure 3-7 Manual migration using flash copy

When the nodes are connected to different storage systems, Remote-copy technologies, such as Global Mirror or Metro Mirror are used to replicate data between the systems. This approach supports remote migrations (see Figure 3-8).

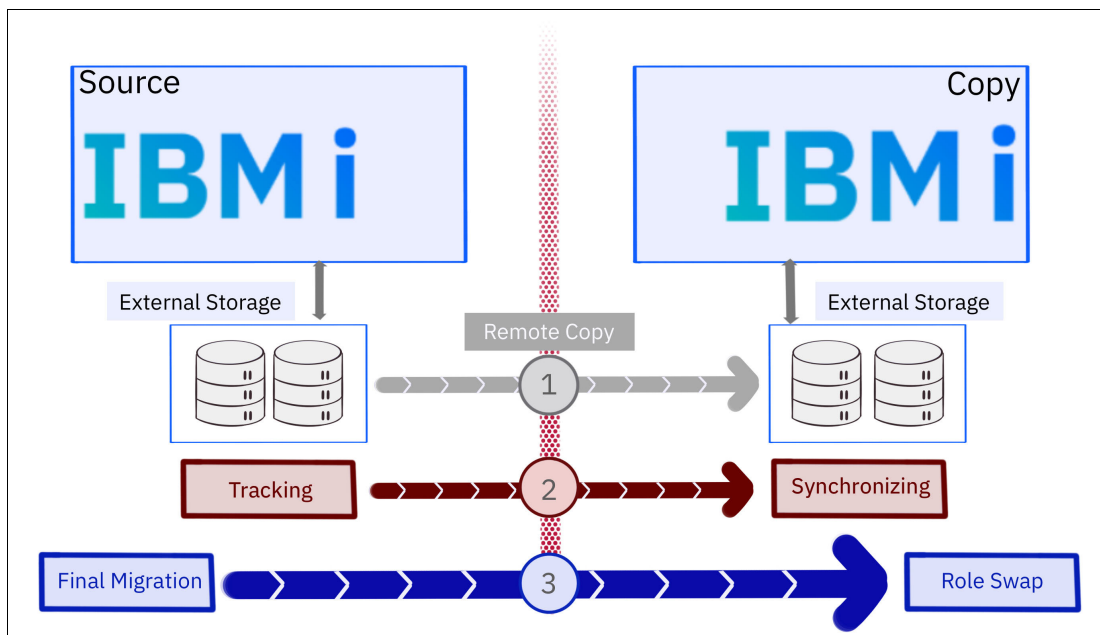


Figure 3-8 Storage Migration using Remote Copy

► Save and Restore

Figure 3-9 shows the approach which uses a full system save of the source node to create a complete copy of the environment. The saved media is then used to perform a scratch installation followed by a full restore on the copy node. Manual Migration leverages the same synchronization technology as Assisted Save and Restore, but gives the IBM i administrator more control over the save media creation, transfer, and environment setup. A Copy Node and NFS Server Node are still required.

This method also requires downtime on the Source Node while the full save is completed. However, this method allows the save media to be created on a tape library or virtual tape library (VTL). The IBM i administrator then can manually manage the data transfer to the NFS Server Node. Especially when used in conjunction with a VTL, this allows the data transfer between sites to be offloaded to the VTL solution. IBM currently supports FalconStor StorSafe as a preferred solution.

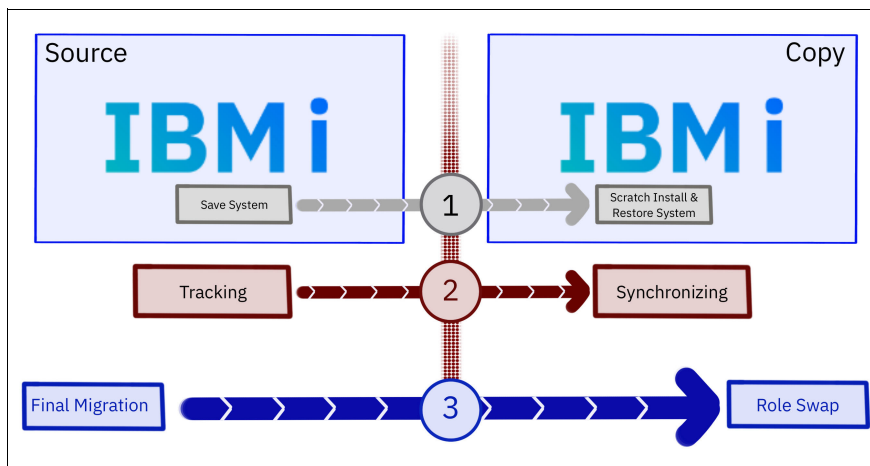


Figure 3-9 Save and Restore diagram

► Storage Replication Combined with Save and Restore

This hybrid method (shown in Figure 3-10) is designed for environments where the source node uses external storage but the copy node does not. It combines storage replication to a temporary node with a save/restore process to build the final copy node. This approach helps reduce downtime on the source system while still supporting mixed-storage architectures.

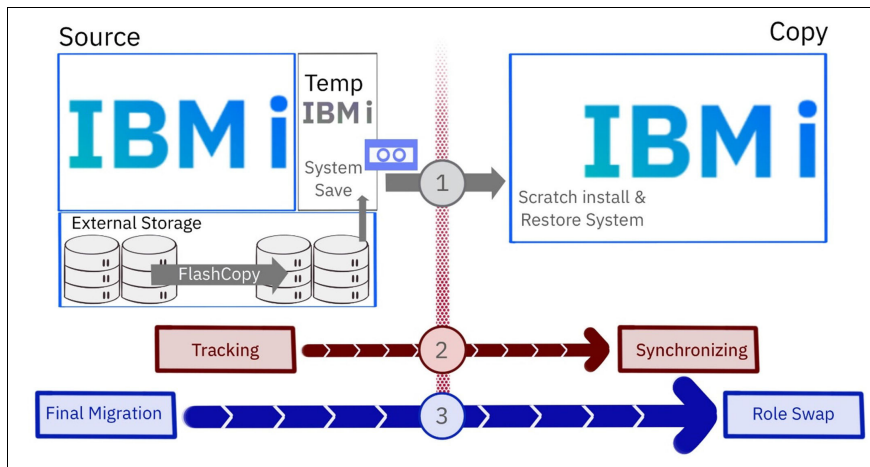


Figure 3-10 Storage Replication and Save/Restore Migration

For detailed step-by-step procedures, see [IBM Documentation](#)

### 3.3.1 Clone Types for Storage Replication

When using storage replication methods (FlashCopy, Remote Copy, or Hybrid), two clone types are available as shown in Table 3-2.

Table 3-2 Clone tupe

Clone Type	Description
Warm Clone (Quiesced)	Source node is quiesced—new database transactions are halted while pending transactions reach a commit boundary. Storage copy occurs while source remains operational. Tracking begins immediately after storage copy completes. Provides shorter outage but requires successful quiesce.
Cold Clone (Powered Off)	Source node is powered off before storage copy operation. Provides maximum data consistency. Tracking begins after source node is powered back on. Longer outage but simpler execution.

**Note:** If the source node fails to reach a quiesced state within the specified timeout during a warm clone, transactions resume on the source node and data changes are not tracked. Consider using cold clone for systems with long-running transactions.

### 3.3.2 Method Selection Guide

Use the following table to select the appropriate manual migration method based on your environment:

Table 3-3

Method	Storage Config	Downtime	Best For
FlashCopy	Same external storage	Minimal (warm) or moderate (cold)	Same data center migrations
Remote Copy	Different external storage	Minimal (warm) or moderate (cold)	Cross-site migrations, DR
Save/Restore	Any (no external required)	Extended (full save duration)	Simple environments, no external storage
Hybrid	Source external, copy internal	Moderate	Mixed storage environments

### 3.3.3 Preparation Summary

Before beginning any manual migration, complete the following preparation tasks:

1. Install required software on source node and GUI node
2. Configure copy node resources (processor, memory, network, storage)
3. Configure secure environment per security guidelines
4. Verify Service Tools user QSECOFR is enabled with non-expiring password
5. Run SYSTOOLS.CHECK\_PRODUCT\_OPTIONS() to identify software product issues

6. Perform method-specific preparations (storage configuration, save media, temporary node)

**Important:** For storage replication methods, ensure target volumes are consistent copies and match source volumes in number, size, and type. For QSRMIGRATE, review additional preparation requirements in IBM Documentation.

### 3.3.4 Configuration via Db2 Mirror GUI

Access the Db2 Mirror GUI at <http://hostname:2001/Db2Mirror> (or port 2007 if TLS enabled). Select IBM i Migrate While Active and choose Manual Migration as the partition migration method. The wizard guides configuration through the following Table 3-4.

Table 3-4 Wizard step for configuratio

Panel	Purpose
Source Authentication	Validate source node connection and software requirements
Node Details	Configure copy node system name, network, and replication IP address pairs
Cluster Details	Review or configure cluster and device domain information
Migration Options	Select migration method (Storage Replication, Save/Restore, or Hybrid) and clone type
Summary	Review all configuration settings before initiating migration





# **SYSBAS Replication and IASP Migration Concepts and Procedures**

Modern IBM i environments rely on flexible storage architectures that support mobility, scalability, and continuous operations. Two core technologies that enable these capabilities are SYSBAS replication and IASP migration.

This chapter outlines the concepts, prerequisites, operational workflows, and detailed procedures for both SYSBAS replication and IASP migration. It explains why IASP migration differs from system-wide replication, describes the strategic decisions involved in selecting an approach, and provides guidance for configuring, executing, validating, and managing these operations. The chapter concludes with testing guidelines, validation strategies, and common troubleshooting techniques that help ensure predictable and repeatable outcomes in production environments.

This chapter will cover the following session:

- ▶ “SYSBAS Replication and Object Migration”
- ▶ “Migrating an IASP During Migrate While Active”

## 4.1 SYSBAS Replication and Object Migration

Understanding which IBM i objects IBM i Migrate While Active can migrate to the copy node is a key consideration when planning and executing a successful migration. This session provides a detailed view of object migration based on IBM i object types and describes the level of support provided during both the system migration phase and the data synchronization phase.

Replication of the SYSBAS environment focuses on duplicating the core system base, which includes essential objects that reside outside of an Independent ASP (IASP). This replication is fundamental in scenarios that require a synchronized system environment for recovery, workload mobility, test environment provisioning, or system cloning. Because SYSBAS contains the IBM i operating system, licensed program objects, middle-ware components, and various system configuration artifacts, the replication method must account for system-wide dependencies and the tight integration characteristic of IBM i architecture.

IBM i Migrate While Active uses a combination of Db2 Mirror for i technology and traditional save/restore mechanisms to synchronize changes in SYSBAS from the source node to the copy node. The migration method used for each object type depends on its eligibility for Db2 Mirror replication. Objects that can be mirrored take advantage of Db2 Mirror's real-time replication capabilities, while non-eligible objects are migrated through complementary save/restore or synchronization techniques.

**Note:** Data within an Independent ASP (IASP) can be migrated outside of Migrate While Active using IASP replication technologies found in the IBM PowerHA SystemMirror for i (5770-HAS) product. For additional resources regarding migrating an IASP in conjunction with Migrate While Active, visit: <https://www.ibm.com/support/pages/node/7176812>

### 4.1.1 Object Classification in IBM i Migrate While Active

IBM i Migrate While Active organizes IBM i objects into four categories based on their location within SYSBAS, their object type, and their eligibility for Db2 Mirror for i replication. Understanding these classifications is essential for determining how each object is migrated and synchronized between the source and copy nodes.

#### Category 1: Db2 Mirror Replication-Eligible Objects in Libraries

These are objects located in SYSBAS libraries that are eligible for replication through the Db2 Mirror for i environment. IBM i Migrate While Active leverages the same Db2 Mirror technology to handle real-time synchronization for object types commonly updated by user applications.

Db2 Mirror-eligible objects include a defined set of library-based types such as:

- ▶ \*Db2 for i physical and logical files (**FILE**)
- ▶ \*Data areas (**DTAARA**)
- ▶ \*Data queues (**DTAQ**)
- ▶ \*Programs (**PGM**)
- ▶ \*Service programs (**SRVPGM**)
- ▶ \*User spaces (**USRSPC**)

These objects are synchronized through the **QMRDBESYNC** process, which detects tracked changes and migrates them to the copy node.

## Category 2: Non-Db2 Mirror Replication-Eligible Objects in Libraries

This category consists of objects that reside in **SYSBAS** libraries but are not eligible for Db2 Mirror replication. For these objects, IBM i Migrate While Active uses *save/restore-based synchronization* techniques rather than real-time mirroring.

Examples include:

- ▶ \*Alert Tables (**ALRTBL**)
- ▶ \*Class objects (**CLS**)

These object types are excluded from Db2 Mirror replication rules, meaning that changes are captured and migrated through save-eligible processing rather than synchronous replication.

## Category 3: Integrated File System (IFS) Objects

IFS objects whose paths reside within **SYSBAS** are largely supported by Migrate While Active. Unlike Db2 Mirror which only covers object types defined in its replication eligibility list Migrate While Active supports all save-eligible objects in **SYSBAS**, including those located in the IFS.

Typical IFS object types include:

- ▶ \*Stream files (**STMF**)
- ▶ \*Directories (**DIR**)

Because IFS objects are outside the traditional library structure, they are synchronized using save/restore operations during the initial migration and during change synchronization.

## Category 4: Pseudo Objects

Pseudo objects are special entries used in the Object Tracking List (OTL) to represent elements of the IBM i environment that do not exist as discrete system objects. These references are required so that Migrate While Active can accurately track and synchronize changes to system configuration attributes that must be kept consistent between source and copy nodes.

Pseudo-object categories include:

- ▶ \*System-level environment variables (**ENVVAR**)
- ▶ \*Security attributes (**SECATR**)
- ▶ \*System values (**SYSVAL**)
- ▶ \*Function usage identifiers (**FCNUSG**)

These pseudo entries ensure that operational characteristics such as security settings, system policies, and runtime environment definitions are fully replicated as part of the migration process.

### 4.1.2 Synchronization Jobs

Migrate While Active uses dedicated synchronization jobs listed in Table 4-1 on page 64 to handle different object categories during the data synchronization stage:

Table 4-1 Job and purpose

Jpb Name	Purpose
QMRDBESYNC	Handles Db2 Mirror replication-eligible objects and pseudo objects from the Object Tracking List (OTL)
QMRDBNSYNC	Handles non-Db2 Mirror replication-eligible objects from the Library Tracking List (LTL)
QMRDBISYNC	Handles integrated file system objects from the IFS Tracking List (ITL)

These jobs start and end on an interval basis. When tracked changes exist in the respective tracking lists, the synchronization job migrates the changes from the source node to the copy node.

### 4.1.3 Tracking Lists

Migrate While Active uses three tracking lists shown in Table 4-2 to record changes that need to be synchronized to the copy node:

Table 4-2 Tracking list

Tracking list	Description
OTL	Object Tracking List – Tracks Db2 Mirror replication-eligible objects and pseudo objects
LTL	Library Tracking List – Tracks non-Db2 Mirror replication-eligible library-based objects
ITL	IFS Tracking List – Tracks integrated file system objects

## 4.2 Migrating an IASP During Migrate While Active

Migrating an Independent Auxiliary Storage Pool (IASP) requires a fundamentally different approach from SYSBAS migration, and understanding this distinction is essential for planning a complete Migrate While Active strategy. Whereas SYSBAS migration is handled natively by IBM i Migrate While Active, IASP environments fall outside the scope of its built-in replication mechanisms. This architectural separation drives the need for a complementary, storage-based replication technology to complete the overall migration picture.

SYSBAS migration focuses on synchronizing the operating system, licensed program objects, configuration artifacts, and Db2 Mirror–eligible object types. This process is tightly integrated with the IBM i platform and allows the system base environment to be kept current on the copy node throughout the migration.

IASPs, however, represent a different domain: they contain user databases, application libraries, and workload specific objects that are physically isolated from SYSBAS by design. Because of this separation, IASP content is not synchronized through Migrate While Active. Instead, it must be migrated using IBM PowerHA SystemMirror for i with geographic mirroring, which provides asynchronous replication at the storage layer to keep the IASP consistent between source and target systems.

This division of responsibility means that IASP migration introduces its own operational considerations. During a Migrate While Active scenario, the IASP must be quiesced at the appropriate point, synchronized using PowerHA geographic mirroring, and validated for

application readiness on the destination system. Factors such as network throughput, bandwidth stability, storage topology, and the timing of cutover all play critical roles in ensuring a reliable transition.

By understanding how SYSBAS and IASP migration workflows complement each other, administrators can orchestrate a coordinated process that delivers a fully consistent environment system base and application data at the target node.

**Important:** The migration of an Independent Auxiliary Storage Pool (IASP) is not supported by IBM i Migrate While Active at this time. This chapter provides a process to perform IASP migration in conjunction with Migrate While Active using PowerHA geographic mirroring.

Table 4-3 shows key differences between SYSBAS and IASP Migration

Table 4-3 Differences Between SYSBAS and IASP Migration

Aspect	Description
SYSBAS Migration	Handled natively by Migrate While Active using Db2 Mirror technology. Automatic tracking and synchronization of changes.
IASP Migration	Must be migrated separately using PowerHA geographic mirroring. Requires manual configuration and coordination with MWA stages.
Timing	IASP migration occurs AFTER system migration stage but BEFORE final cutover, while SYSBAS data synchronization is in progress.
Technology	PowerHA SystemMirror for i (5770-HAS) with geographic mirroring is required for IASP migration.

## 4.2.1 Migration Concept

The basic process is to:

1. Use Migrate While Active to migrate SYSBAS from source node to copy node
2. During the data synchronization stage (while both nodes are active), configure PowerHA geographic mirroring for the IASPs
3. Synchronize IASP data from source to copy node using geographic mirroring
4. Deconfigure geographic mirroring before MWA final cutover
5. Complete Migrate While Active cutover process

For detailed procedures, see IBM Support:

<https://www.ibm.com/support/pages/node/7177936>

## 4.2.2 PowerHA Requirements, Geographic Mirroring Setup, Network & Bandwidth

Independent ASP (IASP) migration within a Migrate While Active solution relies on IBM PowerHA SystemMirror for i and its geographic mirroring capabilities. Unlike SYSBAS, which Migrate While Active handles natively, IASPs are synchronized using storage-level replication. The following subsections describe the required software components, PTF levels, networking characteristics, and geographic mirroring modes.

## Software Requirement

IASP migration requires that all participating systems support geographic mirroring and cluster resource management follow by Table 4-4. According to IBM PowerHA documentation, geographic mirroring requires PowerHA SystemMirror for i and up-to-date IBM i releases on all nodes.

Table 4-4 Required components

Component	Requirement
PowerHA SystemMirror for i	5770-HAS must be installed on all participating systems (source, copy, and HA nodes). Geographic mirroring is an HA function provided through PowerHA.
IBM i OS Version	IBM recommends systems run the most current release participating in geographic mirroring; systems at different releases must have the backup system on the newer release. Supported HA releases include IBM i 7.4 and 7.5, aligned to current Migrate While Active support.
Clustering Requirements	Systems must be configured in a cluster with device domain support. Option 41 (HA Switchable Resources) is required for managing independent disk pools in clusters.

## PTF Requirements

IBM support documentation notes that all nodes in a PowerHA geographic mirroring environment must maintain current PTF levels related to PowerHA, clustering, and geographic mirroring. Before configuring replication, IBM advises confirming that all required PTFs are installed.

Table 4-5 shows applicable fixes for known issue DT400920 on multiple i version.

Table 4-5 Names for fixed issue DT400920

Known Issue	IBM i 7.5 Fix	IBM i 7.4 Fix	IBM i 7.6
DT400920	MJ02654	MJ02653	Included in base release ("Base")

## Network and Bandwidth Requirements

Geographic mirroring replicates disk writes over TCP/IP, and IBM emphasizes the need for adequate bandwidth and properly segmented network traffic to ensure replication performance. Table 4-6 shows all network requirements.

Table 4-6 Network requirement and description

Requirement	Description
Dataport Connections	Dedicated IP addresses for the geographic mirroring data stream between the source and mirror nodes. These connections carry the replicated disk writes.
Cluster Communication Network	A separate network is recommended for cluster heartbeat, cluster management traffic, and command operations. This prevents replication load from affecting cluster stability.
Bandwidth	Bandwidth must be sufficient to accommodate the IASP write workload. Geographic mirroring transmits every write operation from the production disk to the mirror node.

Requirement	Description
Distance / Latency Considerations	Synchronous mode generally supports shorter distances due to latency sensitivity; industry practice typically keeps synchronous mirroring within tens of kilometers. Asynchronous mode tolerates higher latency and greater distances.

### Geographic Mirroring Delivery Modes

IBM PowerHA geographic mirroring supports synchronous (SYNC) and asynchronous (ASYNC) modes. These modes are described in IBM Redbooks guidance. Table 4-7 shows detail for geographic mirroring support mode descriptions.

Table 4-7 Mode descriptions

Mode	Description
Synchronous (SYNC)	Each write operation is committed on both the production disk and the mirror disk before returning control to the application. This provides maximum data protection but requires low-latency, high-bandwidth connections.
Asynchronous (ASYNC)	Write operations complete locally and are queued for remote transmission. This mode supports longer distances and higher latency links, offering improved performance but with potential exposure to minimal data loss during a failure.

**Note:** For IASP migration during Migrate While Active, ASYNC mode is typically recommended as it provides better performance during the synchronization period and is more tolerant of network variations.

## 4.2.3 IASP Migration Scenarios

The approach used to migrate an Independent Auxiliary Storage Pool (IASP) in conjunction with IBM i Migrate While Active depends directly on the current high availability (HA) configuration of the source environment. Because Migrate While Active does not migrate IASPs natively, all IASP migration operations are performed using IBM PowerHA SystemMirror for i, with geographic mirroring acting as the primary replication method. IBM explicitly states that the IASP migration steps vary depending on the pre-existing HA topology and must be selected accordingly.

The Table 4-8 summarizes the supported approaches, followed by detailed explanations for each scenario.

Table 4-8 Overview of IASP Migration Scenarios

Scenario	Current HA Status	Migration Approach
Standalone IASP	No existing HA	Configure new geographic mirroring for migration
Metro/Global Mirror IASP	DS8000/external storage HA	Add geographic mirroring while preserving storage replication
Multi-target Geo Mirror (7.5+)	Existing geographic mirroring	Add copy node as additional target (requires IBM i 7.5+)

Scenario	Current HA Status	Migration Approach
Geographic Mirror (7.4)	Existing geographic mirroring	Deconfigure existing, reconfigure for migration

#### 4.2.4 Migration Timeline and Coordination with Migrate While Active

IASP migration must be carefully synchronized with the phases of IBM i Migrate While Active (MWA). Because SYSBAS and IASPs are migrated using different technologies, the overall process requires deliberate coordination to ensure data consistency, cluster integrity, and a clean cutover. The following timeline outlines the recommended activities and their alignment with the MWA lifecycle.

##### Phase 1: Preparation Before the MWA System Migration Outage

Prior to initiating the MWA system migration outage, any existing high-availability configuration involving IASPs must be cleanly shut down. This step ensures that clustering metadata and storage ownership are in a known, stable state before SYSBAS migration begins.

Key activities include:

- ▶ Cleanly shut down the existing cluster environment (if present)
- ▶ End all Cluster Resource Groups (CRGs) associated with IASPs
- ▶ Vary off all IASPs on the source node to release ownership and prevent access during the outage

Completing these steps avoids conflicts during the system save and restore operations performed by Migrate While Active and ensures that IASPs are not left in an indeterminate state.

##### Phase 2: After MWA System Migration, Before Cutover

Once the MWA system migration phase has completed, both the source node and the copy node are active and operational. During this stage, SYSBAS synchronization is ongoing, and this window is used to migrate IASPs using PowerHA geographic mirroring.

At this point:

- ▶ Both nodes participate in the same cluster and device domain
- ▶ Geographic mirroring is configured for each IASP that must be migrated
- ▶ Initial synchronization and ongoing replication of IASP data begins
- ▶ Replication is allowed to continue until the IASP is fully synchronized

Optionally, once synchronization is complete, administrators can:

- ▶ Vary on the IASP on the copy node
- ▶ Perform application-level validation or testing to confirm data integrity and application readiness

This phase is critical because it allows IASP migration to occur in parallel with SYSBAS data synchronization, minimizing the total migration window.

### Phase 3: Final Steps Before MWA Cutover

Immediately prior to the MWA cutover, geographic mirroring must be cleanly dismantled so that the copy node can assume full ownership of the IASPs without residual HA dependencies.

Required actions include:

- ▶ Detach geographic mirroring on the copy node
- ▶ De-configure the geographic mirroring session
- ▶ Delete the associated CRG definitions
- ▶ Ensure the IASP is fully consistent and ready to be brought online permanently

Once these steps are complete, the environment is ready for the final MWA cutover.

### Phase 4: MWA Cutover

The MWA cutover completes the migration by transferring production roles from the source node to the copy node. At this point:

- ▶ SYSBAS is fully synchronized
- ▶ IASPs are locally owned and no longer mirrored
- ▶ The copy node becomes the new production system

After cutover, clustering and HA can be re-established as needed, based on the target architecture and operational requirements.

The success of an IASP migration with Migrate While Active depends on using the SYSBAS synchronization window to perform IASP replication. When executed correctly, this approach minimizes downtime, preserves data integrity, and enables a controlled, predictable transition to the target system.

## 4.2.5 Step-by-Step Procedure

**Important:** The following procedures provide a high-level overview. For detailed command syntax and parameters, refer to IBM Support page: <https://www.ibm.com/support/pages/node/7177936>

### Pre-Migration: Before MWA System Migration Outage

If an existing cluster is configured on the source node, perform these steps before the MWA system migration outage: Vary off each IASP on source node as shown in Example 4-1.

*Example 4-1 Vary off IASP*

---

```
VRYCFG CFGOBJ(<iasp-name>) CFGTYPE(*DEV) STATUS(*OFF)
```

---

6. End each CRG as shown in

*Example 4-2 End each CRG*

---

```
ENDCRG CRG(<crg-name>)
```

---

7. End clustering on source node as shown in Example 4-3.

*Example 4-3 End clustering on source node*

---

```
ENDCLUNOD NODE(<source-node-name>)
```

---

### **IASP Migration: After MWA System Migration, Before Cutover**

Perform these steps while the SYSBAS data synchronization stage of Migrate While Active is in progress:

1. Create a CRG for the IASP

The source node should be Primary, copy node should be Backup 1. Specify site names and IP addresses as shown in

*Example 4-4 Create CRG*

---

```
CRTCRG CRG(<crg-name>) CRGTYPE(*DEV) RCYDMN((<source-node-name> *PRIMARY *LAST  
SITE1 ('<source-ip>')) (<copy-node-name> *BACKUP *LAST SITE2 ('<copy-ip>'))  
CFGOBJ((<iasp-name> *DEV *CURRENT))
```

---

2. Vary off the IASP on source node as shown in Example 4-5.

*Example 4-5 Vary off IASP on source*

---

```
VRFCFG CFGOBJ(<iasp-name>) CFGTYPE(*DEV) STATUS(*OFF)
```

---

3. Configure geographic mirroring

The IASP will be automatically varied on and synchronization will start. Use the command shown in Example 4-6.

*Example 4-6 Configure geographic mirroring*

---

```
CFGGEOMIR ASPDEV(<iasp-name>) ACTION(*CREATE) SSN(GEOMIRSSN)
```

---

4. Step 4: Monitor synchronization progress

Wait until synchronization is complete before continuing: The command is shown in Example 4-7.

*Example 4-7 Monitor synchronization*

---

```
DSPASPSSN SSN(GEOMIRSSN)
```

---

5. (Optional): Test applications on copy node

If you want to verify applications work before cutover, detach, test, then reattach and re-sync. This is shown in Example 4-8.

*Example 4-8 Optionally test applications*

---

```
CHGASPSSN SSN(GEOMIRSSN) ACTION(*DETACH) -- Perform testing on copy node --  
CHGASPSSN SSN(GEOMIRSSN) ACTION(*REATTACH) DSPASPSSN SSN(GEOMIRSSN)
```

---

## Cutover Preparation

When ready to complete the Migrate While Active cutover:

1. Detach and deconfigure geographic mirroring ON THE COPY NODE as shown in Example 4-9.

*Example 4-9 Detach and deconfigure mirroring on the copy node*

---

```
CHGASPPSN SSN(GEOMIRSSN) ACTION(*DETACH) CFGGEOMIR ASPDEV(<iasp-name>)  
ACTION(*DELETE)
```

---

2. Step 7: Delete the CRG as shown in Example 4-10.

*Example 4-10 Delete the CRG*

---

```
DLTCRGCLU CRG(<crg-name>)
```

---

3. Complete Migrate While Active cutover process

Proceed with the standard MWA cutover procedure. After cutover, the IASP will exist on both source and copy nodes as independent copies with no further replication.

## 4.2.6 Test, Validation, Management, and Troubleshooting

Thorough validation and monitoring are essential to ensure a successful IASP migration when used in conjunction with IBM i Migrate While Active. Because SYSBAS and IASPs are migrated using different technologies, administrators must validate prerequisites, monitor synchronization health, and confirm application readiness at multiple points in the process.

The following checklist provides recommended validation steps before, during, and after IASP migration.

### Before IASP Migration

Before configuring geographic mirroring or beginning synchronization, verify that the environment is fully prepared:

- ▶ Confirm that PowerHA SystemMirror for i is installed on all participating nodes (source, copy, and any HA or DR nodes)
- ▶ Verify that all required PTFs for PowerHA, clustering, and geographic mirroring are applied and at recommended levels
- ▶ Validate network connectivity between nodes, including:
  - Dataport IP addresses used for geographic mirroring
  - Separate network paths for cluster communication, if configured
- ▶ Ensure that disk capacity on the copy node is sufficient to accommodate the full size of each IASP, including expected growth during migration

Completing these checks reduces the risk of configuration failures and unplanned resynchronization events.

### During IASP Synchronization

While geographic mirroring is active and data is synchronizing, continuous monitoring is required:

- ▶ Use **DSPASPPSN** to monitor synchronization status, progress, and any reported errors

- ▶ Watch for automatic suspend conditions, which can occur due to network instability, disk write delays, or storage issues
- ▶ Monitor network bandwidth and latency to ensure replication can keep pace with the IASP write workload
- ▶ Review job logs and PowerHA messages for warnings that may indicate emerging issues

Addressing issues promptly during synchronization helps prevent extended re-sync times and reduces risk as cutover approaches.

### After Cutover

After Migrate While Active cutover completes and the copy node becomes production, validate operational readiness:

- ▶ Verify that each IASP varies on successfully on the copy node and remains accessible
- ▶ Test application access to all IASP-resident data, including databases, user libraries, and IFS paths
- ▶ Confirm that batch jobs, schedulers, and external interfaces function as expected
- ▶ If ongoing availability is required, configure a new HA relationship, such as geographic mirroring or storage-based replication, from the new production (copy) node to its designated HA or DR node

These post-cutover checks ensure the migrated environment is stable, consistent, and ready for sustained production use.

Effective validation throughout the migration lifecycle helps distinguish transient synchronization behavior from genuine configuration issues. By following a structured checklist at each phase, administrators can reduce migration risk, shorten recovery time for unexpected issues, and confidently transition workloads to the target environment.

Table 4-9 show some common issues and provides some possible resolutions.

Table 4-9 Troubleshooting

Issue	Resolution
Geographic mirroring auto-suspends	Check network connectivity and bandwidth. Review QSYSOPR messages. Use CHGASPSSN ACTION(*RESUME) to restart.
CFGGEOMIR fails	Verify target disk units exist and are non-configured. Check namespace creation for NVMe drives.
Synchronization slow	Verify sufficient bandwidth. Consider ASYNC delivery mode. Check for competing network traffic.
CRG creation fails	Verify both nodes are active in cluster and same device domain. Check site names and IP addresses.

Table 4-10 is common commands for key monitoring

Table 4-10 Monitoring command

Command	Purpose
DSPASPSSN	Display geographic mirroring session status and progress
DSPCLUINF	Display cluster information and node status
WRKCLU	Work with cluster configuration

Command	Purpose
DSPCRGINF	Display CRG information and configuration



# Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

## IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- ▶ *IBM i 7.6 features and function*, SG24-8588
- ▶ *Introducing IBM i Modernization Engine for Lifecycle Integration - Merlin*, SG24-8583
- ▶ *IBM Power Systems Virtual Server Guide for IBM i*, SG24-8513
- ▶ *IBM PowerHA SystemMirror for i: Using Geographic Mirroring (Volume 4 of 4)*, SG24-8401
- ▶ *PowerHA SystemMirror for IBM i Cookbook*, SG24-7994

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

[ibm.com/redbooks](http://ibm.com/redbooks)

## Other Publications

- ▶ IBM i Migrate While Active:  
<https://www.ibm.com/support/pages/ibm-i-migrate-while-active>
- ▶ IBM i Migrate While Active – PTF information:  
<https://www.ibm.com/support/pages/ibm-i-migrate-while-active-ptf-information>
- ▶ Cleaning up disk storage:  
<https://www.ibm.com/docs/da/i/7.5.0?topic=upgrades-cleaning-up-disk-storage-space>
- ▶ Reducing system ASP disk space:  
<https://www.ibm.com/support/pages/reducing-system-asp-disk-space-dasd-storage-used>
- ▶ Manual migration using storage replication and save/restore:  
[https://www.ibm.com/docs/en/ssw\\_ibm\\_i\\_75/mig/migmmgettingstarted.htm#migmmgettingstarted\\_\\_prepstorageandsave](https://www.ibm.com/docs/en/ssw_ibm_i_75/mig/migmmgettingstarted.htm#migmmgettingstarted__prepstorageandsave)
- ▶ PowerHA documentation (supported storage servers):  
<https://www.ibm.com/docs/en/i/7.4.0?topic=resiliency-powerha-supported-storage-servers>
- ▶ Configuring PowerHA 7.5 (external reference):  
<https://fortradocs.atlassian.net/wiki/spaces/IWT/pages/2994765846/Configuring+PowerHA+7.5>

- ▶ IBM i Migrate While Active product documentation:  
<https://www.ibm.com/docs/en/i/7.6.0?topic=availability-i-migrate-while-active>
- ▶ Migrating using manual migration  
<https://www.ibm.com/docs/en/i/7.5.0?topic=active-migrating-using-manual-migration>
- ▶ Manual steps for storage replication  
<https://www.ibm.com/docs/en/i/7.5.0?topic=migration-manual-steps-storage-replication>
- ▶ Manual steps for save/restore  
<https://www.ibm.com/docs/en/i/7.5.0?topic=migration-manual-steps-saverestore>
- ▶ Manual steps for storage replication and save/restore  
<https://www.ibm.com/docs/en/i/7.5.0?topic=migration-manual-steps-storage-replication-saverestore>
- ▶ Software requirements  
<https://www.ibm.com/docs/en/i/7.5.0?topic=concepts-software-requirements>
- ▶ Security guidelines  
<https://www.ibm.com/docs/en/i/7.5.0?topic=concepts-security-guidelines>
- ▶ Migrate an IASP in conjunction with Migrate While Active  
<https://www.ibm.com/support/pages/node/7177936>
- ▶ Additional IBM i Migrate While Active resources  
<https://www.ibm.com/support/pages/node/7176812>
- ▶ PowerHA: Steps to Configure Geographic Mirroring  
<https://www.ibm.com/support/pages/powerha-steps-configure-geographic-mirroring-between-two-nodes>

## Help from IBM

IBM Support and downloads

[ibm.com/support](https://www.ibm.com/support)

Services from IBM Consulting

[ibm.com/services](https://www.ibm.com/services)

IBM Training

[ibm.com/training](https://www.ibm.com/training)



# IBM i Migrate While Active

SG24-8597-00

ISBN



(1.5" spine)  
1.5" <-> 1.998"  
789 <-> 1051 pages



# IBM i Migrate While Active

SG24-8597-00

ISBN



(1.0" spine)  
0.875" <-> 1.498"  
460 <-> 788 pages

**Redbooks**

## IBM i Migrate While Active

SG24-8597-00

ISBN



(0.5" spine)  
0.475" <-> 0.873"  
250 <-> 459 pages

**Redbooks**

## IBM i Migrate While Active

(0.2" spine)

0.17" <-> 0.473"

90 <-> 249 pages

(0.1" spine)

0.1" <-> 0.169"

53 <-> 89 pages



**Redbooks**

# IBM i Migrate While Active

SG24-8597-00

ISBN

(2.5" spine)  
2.5" <-> nnn.n"  
1315 <-> nnnn pages



**Redbooks**

# IBM i Migrate While Active

SG24-8597-00

ISBN

(2.0" spine)  
2.0" <-> 2.498"  
1052 <-> 1314 pages







SG24-8597-00

ISBN

Printed in U.S.A.

Get connected

