


# IBM Db2 Mirror for i Getting Started

Jean-Luc Bonhommet

Ingo Dimmer



Power Systems





IBM Redbooks

**IBM Db2 Mirror for i Getting Started**

December 2019

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

**First Edition (December 2019)**

This edition applies to IBM i 7.4 (5770-SS1) and IBM Power Systems servers featuring IBM POWER8 and IBM POWER9 technology.

**© Copyright International Business Machines Corporation 2019. 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 .....	viii
Stay connected to IBM Redbooks .....	viii
<b>Chapter 1. Solution overview</b> .....	1
1.1 Introduction to Db2 Mirror for i .....	2
1.2 Positioning .....	3
1.3 Architecture .....	4
1.3.1 Technology overview .....	4
1.3.2 Remote Direct Memory Access over Converged Ethernet .....	6
1.3.3 Network redundancy group .....	6
1.3.4 Replication states .....	6
1.3.5 Replication criteria list .....	7
1.3.6 Database IASP versus IFS IASP .....	7
1.4 High availability and disaster recovery solutions .....	9
1.4.1 Db2 Mirror and full-system Metro or Global Mirror .....	10
1.4.2 Db2 Mirror and full-system HyperSwap .....	11
1.4.3 Db2 Mirror and PowerHA IASP replication .....	12
1.4.4 Db2 Mirror and ISV logical replication .....	14
1.4.5 Db2 Mirror and Live Partition Mobility .....	14
1.5 Interfaces .....	15
1.5.1 Graphical user interface .....	15
1.5.2 SQL services .....	17
1.5.3 Db2 Mirror setup tool .....	17
<b>Chapter 2. Planning</b> .....	19
2.1 Db2 Mirror setup overview .....	20
2.2 Hardware and software requirements .....	21
2.2.1 IBM Power System hardware requirements .....	21
2.2.2 Hardware Management Console requirements .....	22
2.2.3 Software requirements .....	22
2.3 Objects eligible for replication .....	25
2.3.1 IBM i objects .....	25
2.3.2 Definition only .....	26
2.3.3 Integrated File System objects .....	26
2.4 Replication rules .....	26
2.4.1 Default inclusion state .....	26
2.4.2 System rules .....	27
2.4.3 Replication criteria list .....	28
2.5 Object tracking list suggested priorities .....	29
2.6 Application considerations .....	30
2.7 Quorum .....	31
2.8 Cluster monitors .....	31
2.9 Network Time Protocol server .....	32

2.10 Performance . . . . .	33
2.11 Supported storage. . . . .	33
2.12 Migration considerations . . . . .	33
2.13 Backup and recovery considerations . . . . .	34
2.14 IBM Lab Services offerings . . . . .	35
<b>Chapter 3. Implementation overview . . . . .</b>	<b>37</b>
3.1 Db2 Mirror setup process . . . . .	38
3.2 Initial Db2 Mirror setup by using the GUI setup wizard . . . . .	38
3.2.1 Preparing the setup copy node environment . . . . .	39
3.2.2 Configuring the source and copy node . . . . .	39
3.2.3 Configuring the network redundancy group . . . . .	44
3.2.4 Configuring the storage server . . . . .	44
3.2.5 Configuring the replication criteria list . . . . .	46
3.2.6 Validation and cloning . . . . .	50
<b>Chapter 4. Management overview . . . . .</b>	<b>53</b>
4.1 Db2 Mirror GUI interface . . . . .	54
4.1.1 Connecting to the GUI . . . . .	54
4.1.2 Db2 Mirror GUI overview. . . . .	55
4.2 SQL services. . . . .	59
4.3 Monitoring and alerting . . . . .	61
4.4 Recloning . . . . .	65
<b>Chapter 5. Use cases . . . . .</b>	<b>69</b>
5.1 Planned outages . . . . .	70
5.2 Unplanned outages . . . . .	73
5.2.1 Secondary node outage . . . . .	74
5.2.2 Primary node outage. . . . .	75
5.2.3 Node communication loss. . . . .	78
<b>Related publications . . . . .</b>	<b>81</b>
IBM Redbooks . . . . .	81
Online resources . . . . .	81
Help from IBM . . . . .	82

# 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 <http://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.

Db2®	IBM Spectrum®	pureScale®
DB2®	POWER®	Redbooks®
DS8000®	POWER8®	Redbooks (logo)  ®
FlashCopy®	POWER9™	Storwize®
HyperSwap®	PowerHA®	System Storage™
IBM®	PowerVM®	

The following terms are trademarks of other companies:

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

VMware, and the VMware logo are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

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



# Preface

IBM® Db2® Mirror for i provides a new solution for continuous availability for an IBM i environment based on an active-active clustering design that uses a low-latency communication protocol for synchronous database replication.

With Db2 Mirror, IBM i customers can benefit from continuous application availability for both planned and unplanned outages. Db2 Mirror can help reduce or eliminate application downtime for regular maintenance operations such as program temporary fix (PTF) installations, operating system (OS) upgrades, or for planned server outages.

This IBM Redpaper publication provides a broad overview and understanding of this new solution by covering its architecture, positioning, planning, and implementation aspects.

It provides an introduction reference for a seller or technical specialist audience to become familiar with the new Db2 Mirror solution.

## Authors

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

**Jean-Luc Bonhommet** is an IBM i IT specialist at IBM France. He has more than 20 years of experience as a Technical Specialist in IBM i specializing in external storage on IBM i. He started the IBM i European Benchmark Center in 2000 at the IBM Client Center of Montpellier. He works on the New Technology Introduction on IBM Power Systems for European customers. He manages European customers participating in the Early Ship Program ((ESP), which is also known as beta tests) on Power Systems before product general availability. As part of the IBM Client Center in Montpellier, he delivers customer briefings, and workshops on Power Systems and IBM i subjects. He developed Db2 Mirror for i labs in October 2019 for the IBM Technical University in Prague, Czech Republic.

**Ingo Dimmer** is an IBM Senior Certified IT Specialist and a Project Management Institute (PMI) Certified Project Manager Professional at the EMEA IBM System Storage™ Advanced Technical Support organization that is based in Kelsterbach, Germany. Ingo has 20 years of experience in enterprise storage. He provides pre-sales and post-sales technical support focusing on the Power platform. His areas of expertise include IBM i external disk and tape storage solutions, high availability (HA), I/O performance, and virtualization. Ingo is the author of several IBM technical papers and IBM Redbooks® publications.

The project that produced this publication was managed by:  
**Scott Vetter, PMP**

Thanks to the following people for their contributions to this project:

Aaron Fistler, Scott Forstie, Andy Mallan, Brandon Pederson, Tim Rowe, Kris Whitney  
**IBM**

## 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 papers to be as helpful as possible. Send us your comments about this paper 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 Facebook:

<http://www.facebook.com/IBMRedbooks>

- ▶ Follow us on Twitter:

<http://twitter.com/ibmredbooks>

- ▶ Look for us on LinkedIn:

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

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

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

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

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



# Solution overview

This chapter provides an introduction to the IBM Db2 Mirror for i solution for continuous availability, including its architecture, concepts, and interaction with high availability (HA) and disaster recovery (DR) (collectively known as HA/DR) solutions.

## 1.1 Introduction to Db2 Mirror for i

With the IBM i 7.4 release, IBM offers a new licensed program for the IBM i portfolio: IBM Db2 Mirror for i. This new offering enables near-continuous availability through an IBM i exclusive Db2 active-active 2-system configuration.

Db2 Mirror for i synchronously mirrors database updates between two separate nodes by using a Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) network. Applications can be deployed in *active-active* or *active-passive* mode (updates on one node and read access on the secondary). Db2 Mirror supports applications that use either traditional record-level access or SQL-based database access. It also supports both Java Database Connectivity (JDBC) -attached applications servers and the traditional 5250 approach. The applications and databases may be in either a System Auxiliary Storage Pool (SYSBAS) or as part of an Independent Auxiliary Storage Pool (IASP).

Db2 Mirror for i provides the following benefits:

- ▶ Continuous Application Availability: Provides a recovery time objective (RTO) of zero, enabling continuous application availability.
- ▶ Supports multiple technologies: Flexibility to choose between active-active and active-passive configurations for the application layer.
- ▶ Easy monitoring and management: Intuitive GUI for easy monitoring and management of Db2 Mirror for i.
- ▶ Support rolling upgrades.
- ▶ No journaling is required for replication.

Figure 1-1 shows a general overview of the architecture of a Db2 Mirror environment that uses RoCE as a low-latency communication protocol.

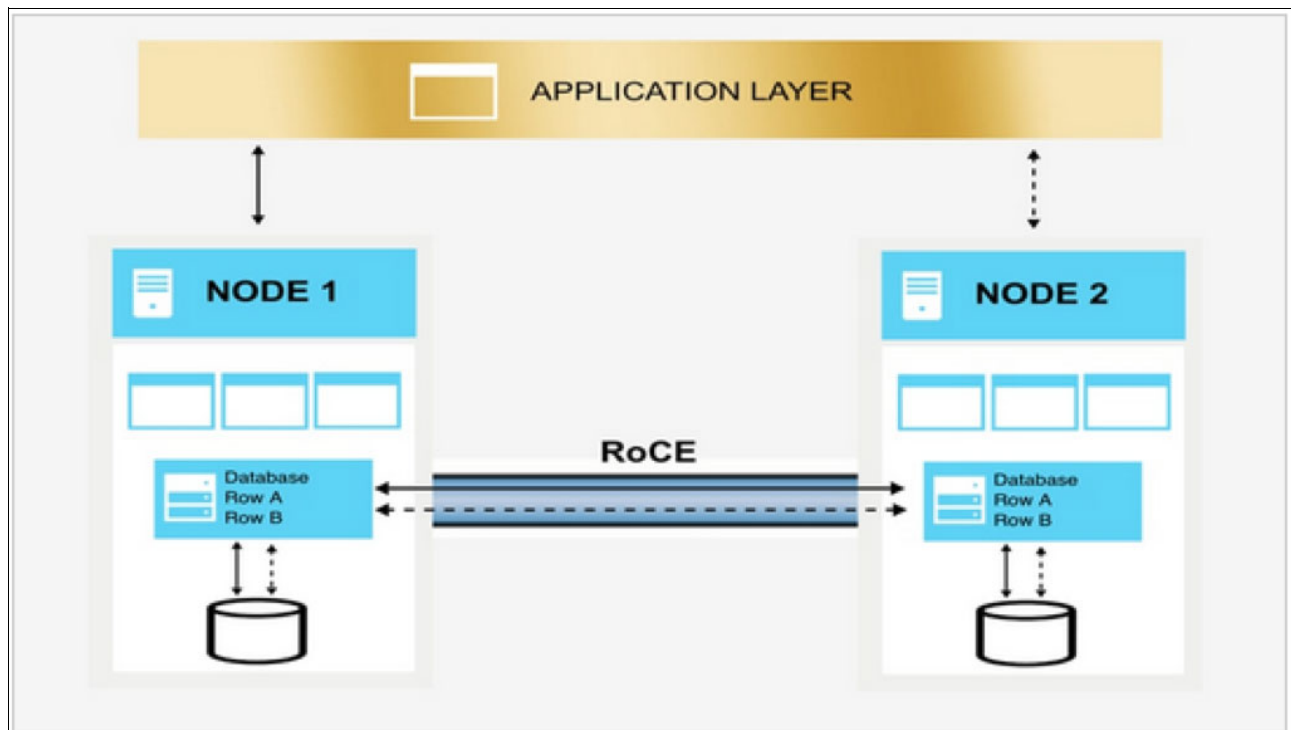


Figure 1-1 Db2 Mirror for i architecture overview

## 1.2 Positioning

Db2 Mirror for i is designed to address business needs that require continuous application availability. It provides an RTO and a recovery point objective (RPO) of zero.

Figure 1-2 shows a spectrum of three different types of HA classes. The illustrations represent 2-node shared-storage configurations for conceptual simplicity. There are many other topologies and data resiliency combinations.

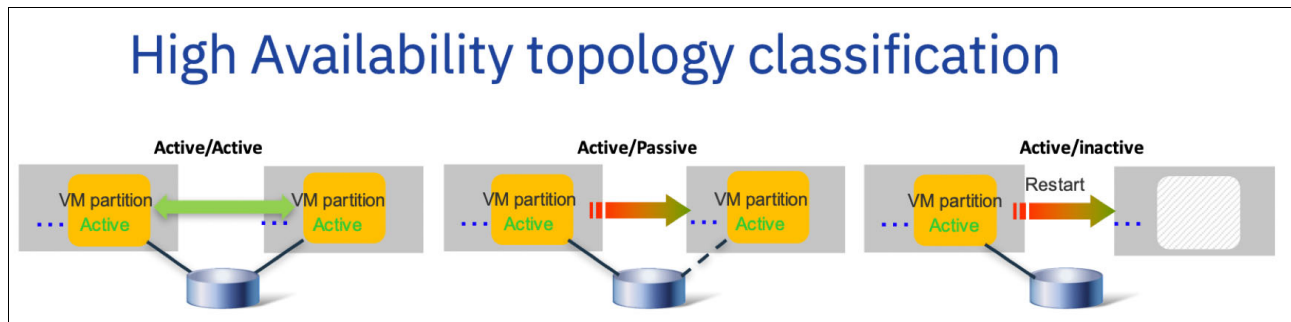


Figure 1-2 High availability classification

The positioning of IBM Db2 Mirror for i within this HA spectrum, including the level of provided outage protection, RPO/RTO, and licensing considerations, is shown in Table 1-1.

Table 1-1 Db2 Mirror for i positioning

Technology	Active-active clustering	Active-passive clustering	Active-inactive
Definition	<i>Application-level clustering:</i> Applications in the cluster have simultaneous access to the production data. Therefore, there is no application restart after an application node outage. Certain types enable read-only access from secondary nodes.	<i>Operating system (OS) -level clustering:</i> One OS in the cluster has access to the production data and multiple OS instances on all nodes in the cluster. The application is restarted on a secondary node after an outage of a production node.	<i>Virtual machine (VM) -level clustering:</i> One VM in a cluster pair has access to the data, one logical OS, and one or two physical copies. The OS and applications must be restarted on a secondary node after a primary node outage event. Live Partition Mobility (LPM) enables the VM to be moved non-disruptively for a planned outage event.
Outage types	Software, hardware, HA, planned, unplanned, RTO 0, and limited distance.	Software, hardware, HA/DR, planned, unplanned, RTO > 0, and multi-site.	Hardware, HA/DR, planned, unplanned, RTO > 0, and multi-site.
OS integration	Inside the OS.	Inside the OS.	OS-neutral.
RPO	Synchronous mode only.	Synchronous and Asynchronous.	Synchronous and asynchronous
RTO	0.	Fast (minutes)	Fast Enough (VM restart).

Technology	Active-active clustering	Active-passive clustering	Active-inactive
Licensing <sup>a</sup>	N+N licensing.	N+1 licensing.	N+0 licensing.
Industry examples	Oracle RAC, Db2 Mirror, and IBM pureScale®.	IBM PowerHA®, Red Hat HA, and Linux HA.	VMware, VM Recovery Manager (VMR) HA, and LPM.

a. N = number of licensed processor cores on each system in the cluster.

## Db2 Mirror advantages

Db2 Mirror is easy to deploy and adopt. Users and applications are largely unaware that synchronous replication is keeping data between the two nodes identical. The data on both nodes is available for production use.

Db2 Mirror can run in two different environment modes:

- ▶ *Active-passive mode*: One node runs your applications while the other node is in standby mode. You can run queries or reports on your database.
- ▶ *Active-active mode*: Both nodes run your applications. Updates to the replicated data are synchronous, and users on both nodes see the same data.

These two modes are not fixed, and you can choose to adapt your environment to meet your needs.

Db2 Mirror can reduce and in some cases eliminate production outage due to planned initial program loads (IPLs) and updates. With the ability to suspend replication, one node can continue processing transactions while the second node is restarted or maintenance is applied. After the IPL is complete and communication is reestablished between nodes, the changes that were tracked by the active node are resynchronized with the secondary node to make the database file identical again. Then, the process can be reversed by suspending replication again and restarting or performing maintenance on the first node. For more information, see 5.1, “Planned outages” on page 70.

This process of using Db2 Mirror to avoid downtime for the business can also be used to upgrade your server or storage hardware.

## 1.3 Architecture

This section provides an overview of the Db2 Mirror architecture and explains some Db2 Mirror concepts.

### 1.3.1 Technology overview

A Db2 Mirror configuration consists of two IBM Power Systems servers in close proximity that are connected through a RoCE network. The application workload updates to the Db2 for i database are replicated in real-time synchronicity between the two systems, either bidirectionally for active-active application deployment, or unidirectional for active-passive application deployments. The active-passive configuration is ideal for query, Business Intelligence (BI), or artificial intelligence / machine learning (AI/ML) workloads that use real-time accurate data on the secondary node.

IBM Db2 Mirror is a licensed program, 5770DBM, that requires IBM i 7.4 on Power Systems servers that are based on IBM POWER8®, IBM POWER9™, and higher technology.

Figure 1-3 shows two IBM i logical partitions (LPARs) which are called *nodes* in Db2 Mirror, each with their own separate database. The Db2 Mirror product provides the ability to have the OS keep native and SQL database objects synchronized between the two nodes. As data change requests occur on either node, Db2 Mirror automatically uses synchronous replication to keep the two instances identical. When one of the nodes is not available, Db2 Mirror tracks every change to replicated objects on the active node. After the mirrored pair is reconnected, all changes are synchronized between the nodes so that once resynchronized both databases are again identical. Db2 Mirror supports libraries and objects that are either in a System Auxiliary Storage Pool (SYSBAS) or an IASP. Simultaneous access for the IBM i Integrated File System (IFS) from both nodes is also supported in an IASP environment.

Figure 1-3 shows a Db2 Mirror environment architecture with applications running on IBM i and a replication layer that uses RoCE technology.

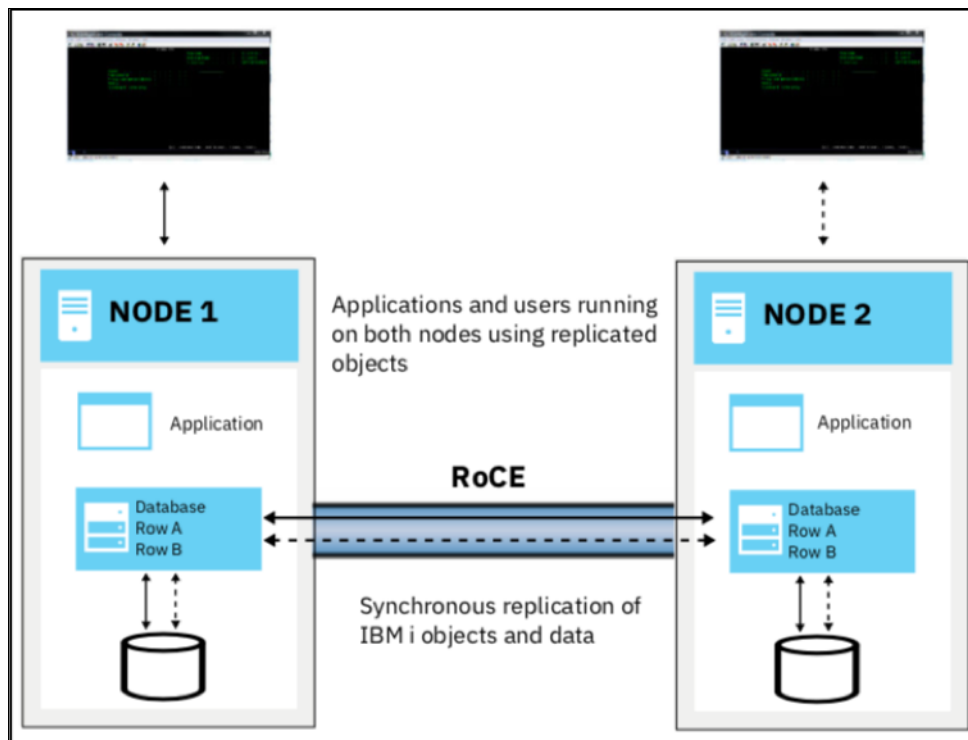


Figure 1-3 Db2 Mirror for i architecture

The configuration of a Db2 Mirror environment begins with one IBM i node that can be an existing production environment. During the setup process, the second node is initialized as an identical copy of that source node.

Db2 Mirror for i can be combined with existing HA/DR solutions like PowerHA or independent software vendor (ISV) logical replication solutions for HA/DR. Some ISVs developed complementary solutions to Db2 Mirror for i. For more information, see 1.4, “High availability and disaster recovery solutions” on page 9.

## 1.3.2 Remote Direct Memory Access over Converged Ethernet

The Db2 Mirror for i solution requires a high-speed network connection with low latency to perform synchronous data operations. RoCE as a RDMA-based protocol meets these requirements. RoCE adapters are required on both nodes for Db2 Mirror for i. The ports of the RoCE adapters can be cabled directly together or connected through a switch. A pair of IP addresses, one from each node, is used to identify each physical RDMA link between the two nodes.

For more information about the details on the hardware requirements, see 2.2.1, “IBM Power System hardware requirements” on page 21.

**Note:** Db2 Mirror does not use remote journals to accomplish synchronous replication. Db2 Mirror synchronous replication is built into the IBM i OS.

## 1.3.3 Network redundancy group

A network redundancy group (NRG) is a logical group of physical ports. It hides the complexity that is associated with automated physical network redundancy, failover, and load balancing of connections between the two nodes.

An NRG consists of one or more (up to 16) physical RDMA-enabled links between nodes. It is a single logical link to the target node. The RDMA network stack transparently handles all processing that is related to which physical link is used at any time based on the NRG configuration.

By default, five NRGs are defined:

- ▶ Db2 Mirror Environment Manager
- ▶ Database replication
- ▶ System object replication
- ▶ IFS replication
- ▶ Resynchronization

Priority and load balancing can be modified for the different NRGs.

## 1.3.4 Replication states

Replication states must be understood and monitored. A replication state indicates whether replication is in one of the following states:

- ▶ Active replication state

Normal state of Db2 Mirror. Replication is active, both nodes are in the active status, and they accept any creates or updates on replicated objects.

- ▶ Tracking / Blocked replication state

When a node is not available or if communication is interrupted between the nodes, active replication is suspended, which causes Db2 Mirror to go into the TRACKING/BLOCKED replication state to prevent conflicting changes to objects. One node may continue operations that affect replicated objects and its replication state is set to TRACKING. The other node is blocked from changing replicated objects and its replication state is set to BLOCKED. The node that is BLOCKED still may change non-replicated objects and query replicated objects.



## Replication role

The determination of which node is set to TRACKING and which node becomes BLOCKED is made by primary and secondary node designation. During the initial setup of Db2 Mirror, the *setup source node* is the *primary node* and the *setup copy node* becomes the *secondary node*. So, the primary node goes to the TRACKING state and the secondary node goes to the BLOCKED state. Primary and secondary node roles can be changed by the administrator.

Database IASPs (DB IASPs) have a separate replication role from SYSBAS. During normal operation, the role of an active DB IASP matches the replication role of SYSBAS. If an IASP is varied off on either node, the replication role of the IASP is adjusted separately from the SYSBAS replication role.

## Object tracking list

The object tracking list (OTL) is the mechanism that is used to track changes that are made on replicated objects when the replication state is TRACKING. Because IASPs might have a separate state from SYSBAS, there is an OTL for SYSBAS and a separate OTL for each IASP. Each OTL on a node contains the changes that must be sent to the other node during the resynchronization process. It is possible to define priorities on objects during resynchronization by using the GUI interface or Db2 Mirror SQL services. For more information, see 2.5, “Object tracking list suggested priorities” on page 29.

### 1.3.5 Replication criteria list

The replication criteria list (RCL) is a rules engine that is used by Db2 Mirror. It consists of a set of *rules* identifying group of objects that should or should not be replicated. The rules combined with the default inclusion state (including the entire SYSBAS or not), selected during initial setup, and system rules, provide a concise process to determine the replication status of any existing or future object.

Because Db2 Mirror supports objects in SYSBAS and in IASPs, an independent RCL exists for each IASP that is registered as a DB IASP with Db2 Mirror.

## Replication-eligible objects

The Db2 Mirror environment supports replication of the object types that are updated in user applications. For more information, see 2.3, “Objects eligible for replication” on page 25.

### 1.3.6 Database IASP versus IFS IASP

IBM Db2 Mirror for i supports existing IASPs to be registered as either DB IASPs or IFS IASPs that are simultaneously accessible from both nodes:

**Database IASPs** Similar to SYSBAS object replication by Db2 Mirror, a DB IASP supports replication of all objects that are eligible for Db2 Mirror replication, as described in 2.3, “Objects eligible for replication” on page 25. Like for non Db2 Mirror environments, DB IASPs can be configured to provide different database name spaces for workload isolation or for extending Db2 Mirror by using an HA/DR solution with remote replication. Although it is simultaneously available on both Db2 Mirror nodes, a DB IASP is not switchable between the two local nodes.

**IFS IASPs**

An IFS IASP must make IFS objects simultaneously available on both Db2 Mirror nodes. It is used to reference IFS objects only, and at any time is attached to only one of the two Db2 Mirror nodes. It is also must be switchable between both nodes by using PowerHA, but is not replicated by Db2 Mirror. Any database or other non-IFS objects that are included in IFS IASPs are accessible from the local node only where the IASP is varied on, but not from both nodes.

**Note:** An existing IASP with both database files and IFS must be split into two IASPs (one for database and one for IFS) to make both IASP objects and IFS files accessible from both nodes.

The different concepts for simultaneous data access among Db2 Mirror nodes for database and IFS IASPs are shown in Figure 1-4.

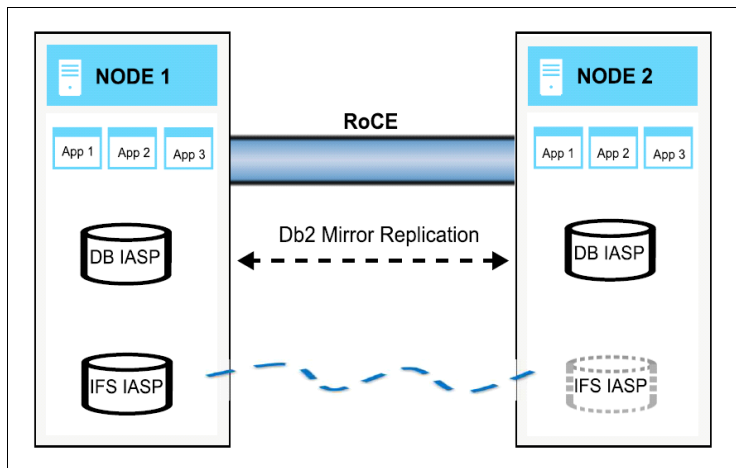


Figure 1-4 Db2 Mirror IASP support

Simultaneous IFS access from both nodes is implemented by using an IFS client/server technology that uses a *mutable file system* architecture like the one that is shown in Figure 1-5.

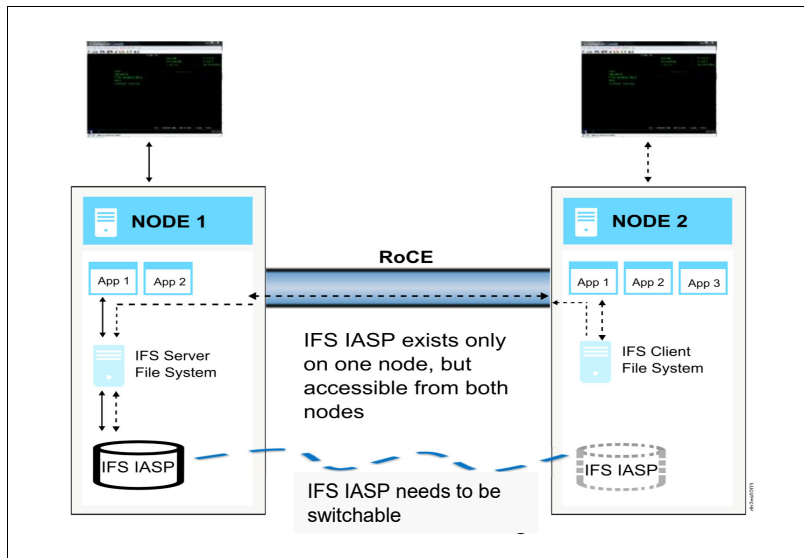


Figure 1-5 Db2 Mirror IFS support

The IFS IASP is accessible for reads and writes from both nodes. Requests to the IFS client node are forwarded to the IFS server node where the IFS IASP is varied on. The IFS IASP must be switchable either by using PowerHA logical unit number (LUN) -level switching or IASP replication. The client/server technology automatically changes when the storage is switched.

## 1.4 High availability and disaster recovery solutions

Db2 Mirror by design provides a local *continuous availability* solution within a data center with a maximum supported distance between two Db2 Mirror nodes of 100 meters, or 200 meters when using an Ethernet switch.

Many customers also require a HA or DR solution to protect them against local or regional outages. For this purpose, a Db2 Mirror local continuous availability solution can be combined with existing HA/DR replication technologies.

For a combination of Db2 Mirror with HA/DR replication technologies, the following considerations apply:

- ▶ Remote replication for DR can be implemented either by storage-based replication or ISV logical replication solutions.
- ▶ Any IFS IASP must remain switchable between both local Db2 Mirror nodes by choosing a DR topology that is supported by PowerHA.
- ▶ Any DB IASP is available on both local nodes (no switch between local nodes).  
A DB IASP is not required for local Db2 Mirror database replication, but might be preferred for implementing a remote replication solution with shorter recovery times compared to SYSBAS replication.
- ▶ For a complete business continuity solution at the DR site, a remote Db2 Mirror node pair can be configured for a 4-node Db2 Mirror PowerHA cluster configuration. Both IFS IASPs and DB IASPs must be registered with the remote Db2 Mirror pair (by using the SHADOW option for the DB IASP to maintain its Db2 Mirror configuration data like default inclusion state and RCL).

Within the following sections, we provide some examples of how a local Db2 Mirror environment might be extended by available HA/DR technologies.

## 1.4.1 Db2 Mirror and full-system Metro or Global Mirror

Db2 Mirror can be combined with storage-based full-system replication, as shown in Figure 1-6.

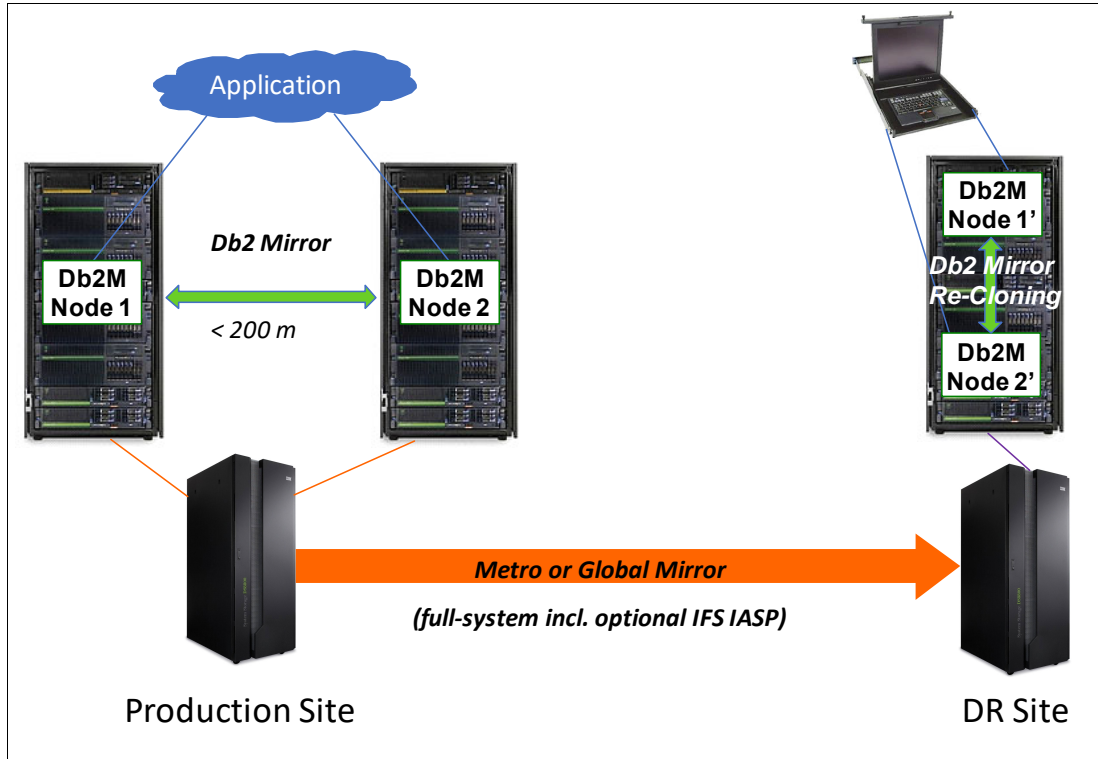


Figure 1-6 Db2 Mirror combined with full-system replication

The following considerations apply for combining Db2 Mirror with full-system replication as a DR solution:

- ▶ Full-system storage-based replication of Db2 Mirror is considered a DR only solution, and it should not be used for regular site swaps.
- ▶ Replication can be either synchronous like with IBM Metro Mirror or asynchronous with consistency like with IBM Global Mirror.
- ▶ The Db2 Mirror node state (active, tracking, or blocked) is replicated.
- ▶ Users typically must start both DR nodes if they are in doubt about the prior node state.
- ▶ If Db2 Mirror on the DR site is wanted, Db2 Mirror reconfiguration and recloning are required.

## 1.4.2 Db2 Mirror and full-system HyperSwap

Db2 Mirror combined with full-system replication (described in 1.4.1, “Db2 Mirror and full-system Metro or Global Mirror” on page 10) also can be combined with IBM Storage Metro Mirror HyperSwap® to provide extra protection against local storage outages without switching to the DR site, as shown in Figure 1-7.

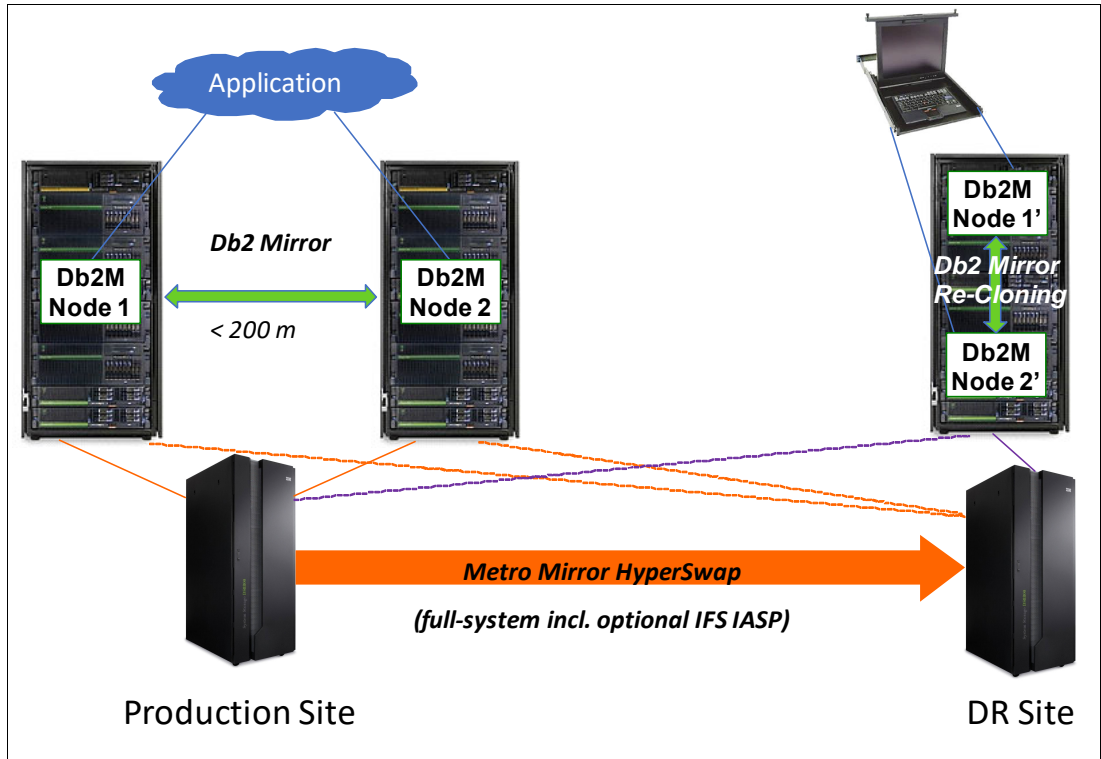


Figure 1-7 Db2 Mirror combined with full-system HyperSwap

### 1.4.3 Db2 Mirror and PowerHA IASP replication

Combining Db2 Mirror with PowerHA IASP replication, as shown in Figure 1-8, instead of full-system replication provides a higher level of availability for DR with a shorter RTO due to the enabled PowerHA and managed IASP failover to the DR site.

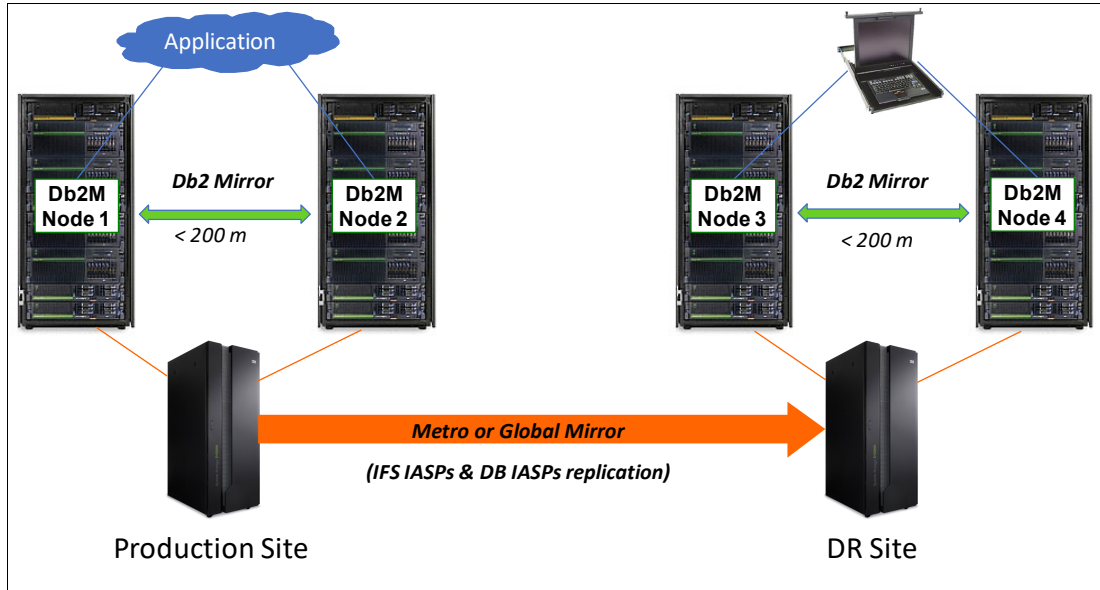


Figure 1-8 Db2 Mirror combined with PowerHA IASP replication

The following considerations apply for this solution:

- ▶ The PowerHA cluster includes two Db2 Mirror node pairs.
- ▶ All data to be replicated between the Db2 Mirror production pair and the Db2 Mirror DR pair should be included in DB IASPs (production) and the IFS IASPs (DR) with their remote replication managed by PowerHA.

PowerHA support for a container cluster resource group (CRG) can be used for simultaneously switching DB IASPs and IFS IASPs as through they were one CRG,

- ▶ The system environment, like user profiles, are kept in sync between both Db2 Mirror pairs by the PowerHA cluster admin domain.
- ▶ A planned site-switch of DB IASPs preserves their replication state.
- ▶ Recloning of the DB IASPs might be required after an unplanned PowerHA failover.

To provide even higher redundancy and availability from a storage perspective, this solution also can be implemented with a separate IBM System Storage DS8000® storage system for each local Db2 Mirror node, as shown in Figure 1-9.

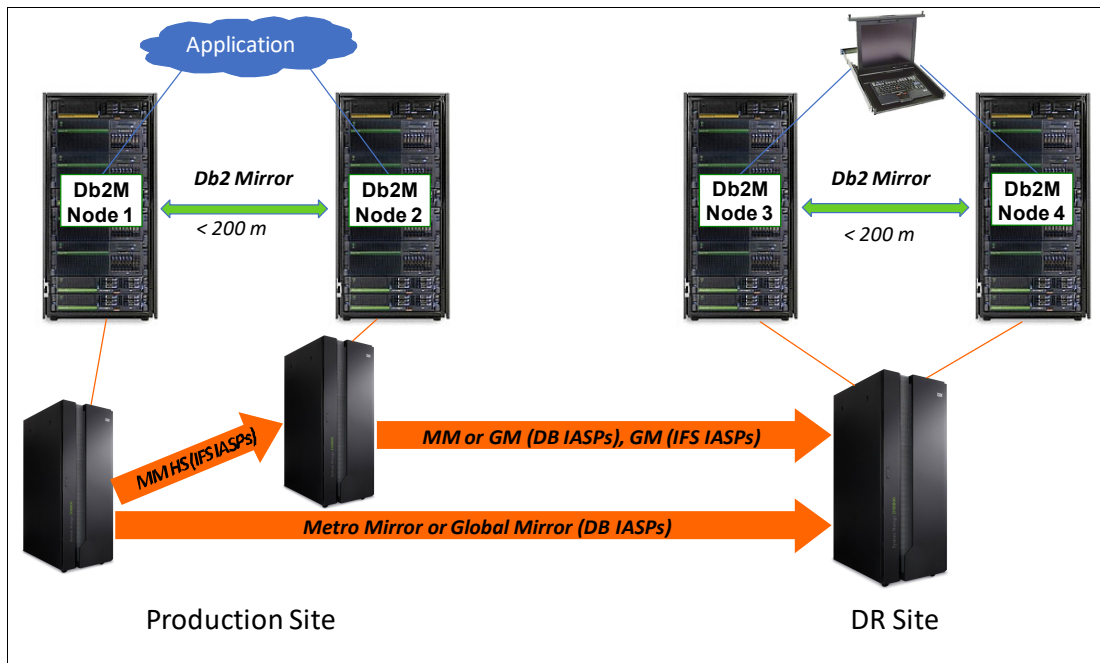


Figure 1-9 Db2 Mirror combined with PowerHA IASP replication with three DS8000 systems

In this example, the DB IASPs are replicated by using Metro Mirror or Global Mirror from the production to the DR site, and the IFS IASPs, which must remain switchable by PowerHA, is replicated by the PowerHA -supported Metro Mirror HyperSwap + Global Mirror 3-site configuration.

## 1.4.4 Db2 Mirror and ISV logical replication

Db2 Mirror also can be combined with an ISV logical replication solution. With this configuration, you can move the source for the logical replication between the production Db2 Mirror pair and replicate to either a single DR node or two DR node configuration, as shown in Figure 1-10.

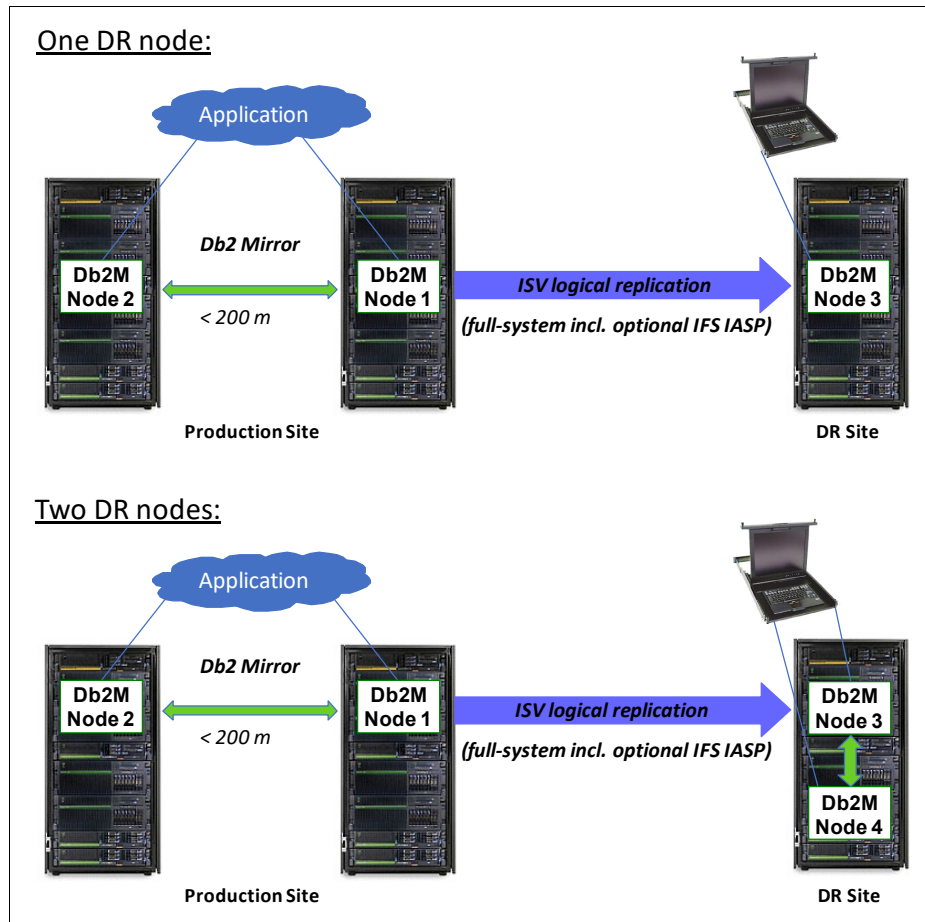


Figure 1-10 Db2 Mirror combined with ISV logical replication

When using a 2-node DR configuration, the DR nodes also should be added to the Db2 Mirror cluster so that the quorum data, including the node status, is replicated outside the Db2 Mirror production node pair.

## 1.4.5 Db2 Mirror and Live Partition Mobility

For planned maintenance or migration without downtime, you can migrate a Db2 Mirror node to another Power Systems server by using LPM.

The following requirements apply for using LPM with Db2 Mirror:

- ▶ IBM i 7.4 TR1 or later.
- ▶ IBM POWER9 processor-based servers with firmware level 940 or later.



- ▶ Single root I/O virtualization (SR-IOV) logical ports are used for migrating the partition instead of a physical RoCE adapter.
- ▶ The *restricted I/O* attribute is set for LPM in the partitions' profile on the Hardware Management Console (HMC).

Using LPM in a Db2 Mirror environment requires special considerations for removing and adding the SR-IOV logical ports, which are described in 5.1, "Planned outages" on page 70.

**Note:** IBM PowerVM® Simplified Remote Restart technology, which is used by some automation tools for remote restart recovery in case of unplanned outages, requires a virtual I/O only configuration for the partition, so it cannot be used for a Db2 Mirror node.

## 1.5 Interfaces

Db2 Mirror can be configured and managed by the *Db2 Mirror GUI* and by various *SQL services* that consist of SQL procedures, functions, and views. The initial setup of Db2 Mirror can be performed by either using the GUI or a combination of SQL services and the Db2 Mirror setup tool **db2mtool**, which is available as a Qshell command.

This section provides a brief overview of the available interfaces to set up and manage a Db2 Mirror for i environment.

### 1.5.1 Graphical user interface

The Db2 Mirror GUI is a web-based GUI that is part of the 5770-DBM licensed program \*BASE option, which does not require a Db2 Mirror license and is running on the IBM i HTTP Apache web server. It can be run from either the Db2 Mirror primary or secondary node for managing and monitoring. However, it is a best practice to run it from a third management node, which is required for the initial setup of a Db2 Mirror pair configuration and also can be used for a Db2 Mirror quorum, as described in 2.7, "Quorum" on page 31.

Figure 1-11 shows an example of one of the Db2 Mirror GUI's initial setup windows.

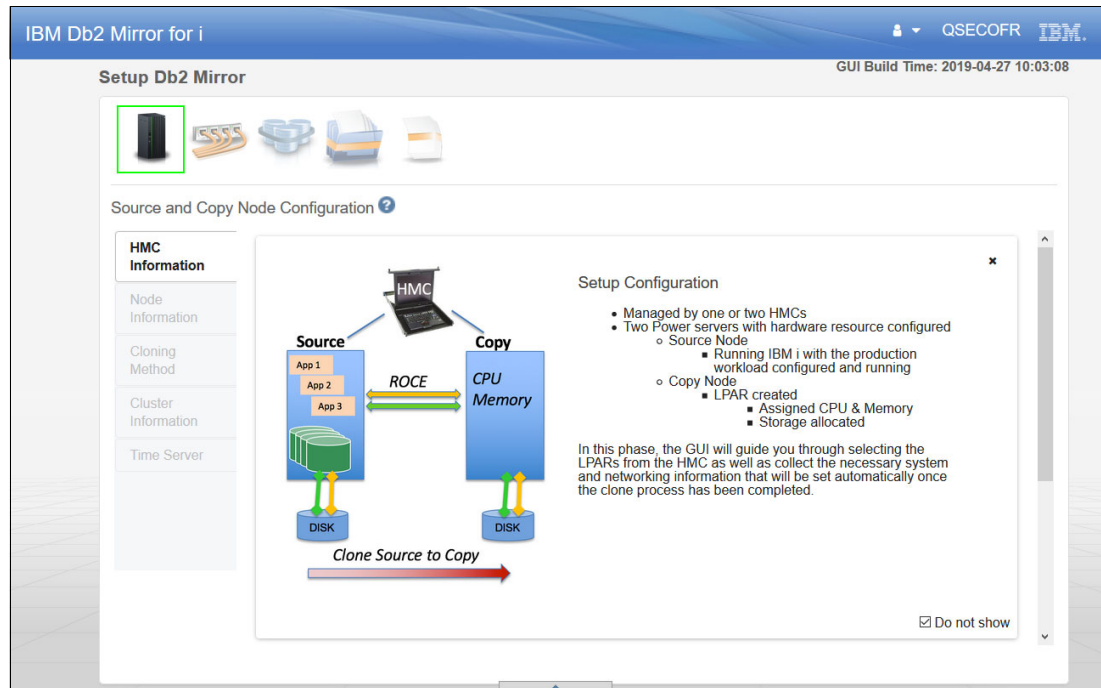


Figure 1-11 iDb2 Mirror GUI initial setup window

Figure 1-12 shows an example of the GUI for managing an existing Db2 Mirror environment with the menu open for OTL tasks.



Figure 1-12 Db2 Mirror GUI managing an existing environment

## 1.5.2 SQL services

As an alternative to setting up or managing Db2 Mirror by using the GUI, there are various SQL services in the form of procedures, functions, and views that are available for the following categories:

- ▶ Communication services
- ▶ Product services
- ▶ Replication services
- ▶ Resynchronization services
- ▶ Serviceability services

A complete list of available Db2 Mirror SQL services can be obtained by querying the SQL SERVICES\_INFO view, as shown in Example 1-1.

*Example 1-1 Available Db2 Mirror SQL services*

---

```
SELECT * FROM QSYS2.SERVICES_INFO WHERE SERVICE_CATEGORY LIKE 'MIRROR%';
```

---

In 4.2, “SQL services” on page 59, we show some examples of using Db2 Mirror SQL services from the 5250 command line and from IBM i Access Client Solutions.

For more information about the available Db2 Mirror SQL services, see [IBM Knowledge Center](#).

## 1.5.3 Db2 Mirror setup tool

The Db2 Mirror setup tool is available as the Qshell command **db2mtool** and is required for some configuration and validation tasks of an initial Db2 Mirror setup if the GUI is not used.

This tool has different action parameters for working with Db2 Mirror JSON configuration files, performing validation and cloning operations, and configuring clustering.

**Note:** The **db2mtool** is required only if the GUI is not used for the initial setup or if PowerHA is not used for managing the Db2 Mirror cluster.

Db2 Mirror uses the following two kinds of JSON configuration files, which describe the setup for a new Db2 Mirror environment, including the setup source node, setup copy node, and storage configurations:

**clone\_info.json file** Contains the storage and cloning attributes that are used by Db2 Mirror to automatically perform the storage cloning actions.

**meta\_data.json file** Contains the setup copy node configuration settings that are used by Db2 Mirror to automatically configure the setup copy node during the first IPL after the storage clone.

If you use the GUI with its setup wizard to set up Db2 Mirror, you do not need to care about these configuration files because they are automatically created, populated, and uploaded to the setup source node with the detailed configuration information specified by you.

Using the Db2 Mirror tool to interactively create and populate a new JSON configuration file for a Db2 Mirror setup with the tool prompting the user to enter all required information, and using it for changing a cluster IP interface, are shown in Example 1-2.

*Example 1-2 The db2mtool usage examples*

---

```
db2mtool action=new jsonType=[clone_info | meta_data] jsonFile=/home/xxx
```

```
db2mtool action=cluster cl=chgcluinterface node=<nodename>  
ip1=<current_IP_address> newip1=<new_IP_address>
```

---

For more information about the Db2 Mirror tool, see [IBM Knowledge Center](#).



# Planning

This chapter introduces into the planning considerations for setting up an IBM Db2 Mirror environment.

It provides information about the requirements; important planning decisions to be made by the user for the configuration of the Db2 Mirror environment, including the scope of replication; and considerations pertaining to application design, performance, backup and recovery, and more.

## 2.1 Db2 Mirror setup overview

Db2 Mirror is initially set up on a single partition that is called the *setup source node*. During the setup process, the setup source node is cloned to create a second node of the Db2 Mirror pair, and it is called the *setup copy node*. The setup copy node is configured and initialized automatically by Db2 Mirror as part of its setup process.

**Note:** To set up Db2 Mirror, a third IBM i node is required, which is used as the *management node* for initiating the initial setup for a Db2 Mirror pair by using the Db2 Mirror GUI or the Db2 Mirror SQL services and setup tool.

Figure 2-1 shows a setup overview for a Db2 Mirror environment, including a third node for the setup. This third node also can be used after the setup as a management node and quorum for the Db2 Mirror cluster.

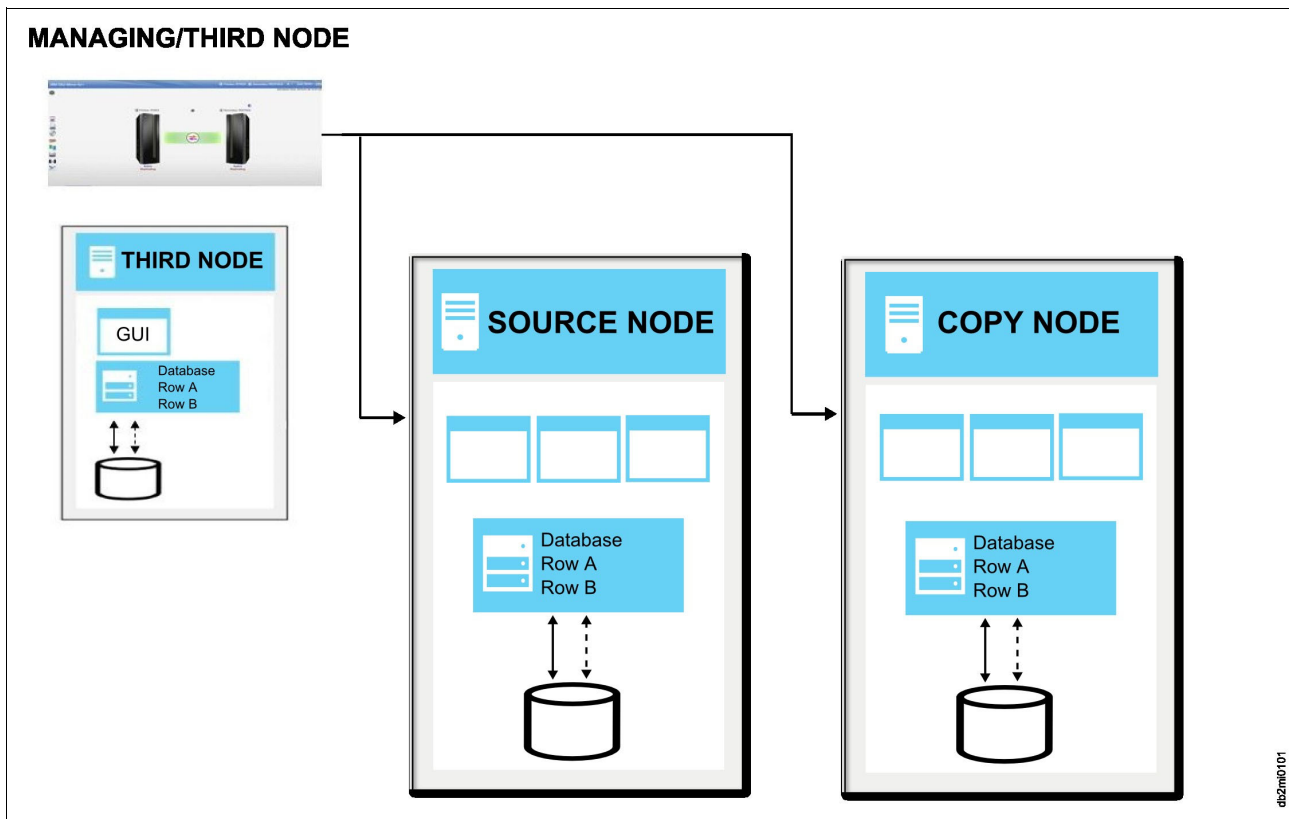


Figure 2-1 Setup process overview

For more information about this setup process, see 3.1, "Db2 Mirror setup process" on page 38.

## 2.2 Hardware and software requirements

This section provides an overview of the Db2 Mirror requirements. For a comprehensive list of the requirements for Db2 Mirror, see [IBM Knowledge Center](#).

### 2.2.1 IBM Power System hardware requirements

You need Power Systems servers and network adapters.

#### Power Systems servers

The two nodes of a Db2 Mirror pair must run on POWER8 processor-based servers or later.

From a resiliency perspective, host the two nodes of a Db2 Mirror pair on two separate Power Systems servers.

There is no requirement that both servers have the same hardware configuration. However, as a best practice, plan for adequate performance capabilities for each of the two servers to support the application workload from both nodes in case either Db2 Mirror node becomes unavailable.

Db2 Mirror can be configured and run on a single Power Systems server, but that diminishes the goal of the continuous availability if the server has an outage. Table 2-1 shows the minimum server hardware and firmware level requirements for Db2 Mirror.

*Table 2-1 Power Systems and firmware levels that are supported with Db2 Mirror*

Power Systems server	Firmware level
POWER8	FW860.60 or later
POWER9	FW930 or later

#### Network adapters

Db2 Mirror requires Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) adapters that provide a fast, high-bandwidth network between the nodes. RoCE adapters, either physical or virtualized, must be installed in both partitions. The adapters can be directly interconnected by using a cable or optionally connected through a RoCE switch. The maximum physical distance between the two nodes is limited by the length of the cables.

Only POWER9 processor-based servers can support RoCE in single root I/O virtualization (SR-IOV) mode, which allows virtualization of the adapter that can then be shared between different IBM i logical partitions (LPARs). Up to eight IBM i LPARs can share one RoCE adapter port.

The network switch must support the RoCE V1 communication protocol.

Table 2-2 lists all the supported network adapters. A 100 Gb adapter is a best practice.

Table 2-2 *RoCE adapters that are supported by Db2 Mirror for i on Power Systems servers*

RoCE adapter	POWER8 processor-based system	POWER9 processor-based system
PCIe3 2-port 100GbE NIC & RoCE QSFP28 Adapter (FC EC3L and EC3M; CCIN 2CEC)	X	X
PCIe3 2-port 10Gb NIC & RoCE SR/Cu adapter (FCEC2R and EC2S; CCIN 58FA)	NA	X
PCIe3 2-port 25/10 Gb NIC & RoCE SFP28 adapter (FC EC2T and FC EC2U; CCIN 58FB)	NA	X
PCIe4 2-port 100GbE RoCE x16 adapter (FC EC66 and EC67; CCIN 2CF3)	NA	X

## 2.2.2 Hardware Management Console requirements

The Db2 Mirror nodes must be managed by a Hardware Management Console (HMC).

The IBM i partition that is used as the setup copy node for the cloning process must be created before beginning the Db2 Mirror configuration process. The partition does not need to be installed with the IBM i operating system (OS), but it does need to be defined on the HMC and have all its hardware assigned to it (such as RoCE adapters and LAN adapters). This partition can be defined on the same HMC or on a different HMC from the setup source node.

Table 2-3 shows the required minimum HMC firmware levels.

Table 2-3 *Minimum HMC firmware levels that are required for Power Systems servers*

IBM POWER® processor-based systems	Minimum firmware level
POWER8	V8R8.6.0
POWER9	V9R1.930.0

## 2.2.3 Software requirements

This section describes that necessary software requirements for Db2 Mirror.

### IBM i operating system

Db2 Mirror requires IBM i 7.4 or higher.

Before beginning the Db2 Mirror setup process, install the latest levels of the following IBM i program temporary fix (PTF) groups on the setup source node:

- ▶ IBM i cumulative PTF package: SF99740 (IBM i 7.4)
- ▶ Db2 Mirror PTF group: SF99668 (IBM i 7.4)
- ▶ Db2 for IBM i PTF group: SF99704 (IBM i 7.4)

On the managing node, install the latest level of the following PTF groups:

- ▶ Db2 Mirror PTF group: SF99668 (IBM i 7.4)
- ▶ IBM HTTP Server for i PTF group: SF99662 (IBM i 7.4)
- ▶ Java PTF group: SF99665 (IBM i 7.4)



## IBM i options and products

Product license keys are required for each of the two nodes of a Db2 Mirror pair.

Add all the license keys that are required for the setup source node and the setup copy node on the setup source node. The keys are automatically applied after the clone completes.

## Required setup source and managing nodes products

The setup source node and the managing node must have the options and products that are shown in Table 2-4 installed.

Table 2-4 Required IBM i products and options for setup source and managing nodes

Product and option	Description	Notes	Setup source node	Managing node
5770SS1 Option 3	Extended Base Directory Support	For jt400.jar.	X	X
5770SS1 Option 12	Host Servers		X	X
5770SS1 Option 26	IBM DB2® Symmetric Multiprocessing	For parallel degree that is used by resynchronization.	X (optional)	
5770SS1 Option 30	Qshell		X	X
5770SS1 Option 33	Portable Application Solutions Environment (PASE)	Required on the setup source node for time synchronization by using Network Time Protocol (NTP). Required on the managing node for cloning IBM Spectrum® Virtualize storage.	X	X (optional)
5770SS1 Option 34	Digital Certificate Manager (DCM)	Required on any node that is running a cluster monitor.	X (optional)	
5770SS1 Option 41	High Availability Switchable Resources	Required for cluster device domain.	X	
5770SS1 Option 48	Db2 Data Mirroring		X	
5770JV1 *BASE	IBM Developer Kit for Java		X	X
5770JV1 Option 16	Java SE 8 32-bit		X	X
5770JV1 Option 17	Java SE 8 64-bit		X	X
5733SC1 *BASE	IBM Portable Utilities for i		X	X
5733SC1 Option 1	OpenSSH, OpenSSL, and zlib		X (optional)	X
5770DG1 *BASE	IBM HTTP Server for i		X	X
5770DBM *BASE	IBM Db2 Mirror for i		X	X

Product and option	Description	Notes	Setup source node	Managing node
5770DBM Option 1	Db2 Mirror Enablement	Db2 Mirror Enablement is not required for the managing node.	X	
DS Command-Line Interface (DS CLI)	DS8000 Command-Line Interface	Required on the managing node for cloning IBM System Storage DS8000.		X (optional)

### Open source packages

These required packages are delivered in Red Hat Package Manager (RPM) packages and installed by using the YUM installer or IBM i Access Client Solutions:

- ▶ python2-six-1.10.0-1.ibm7.1.noarch.rpm
- ▶ python2-ibm\_db-2.0.5.8-1.ibm7.1.ppc64.rpm
- ▶ cloud-init-1.2-0.ibm7.1.ppc64.rpm

Figure 2-2 shows the IBM i Access Client Solutions product managing the open source package installation and update.

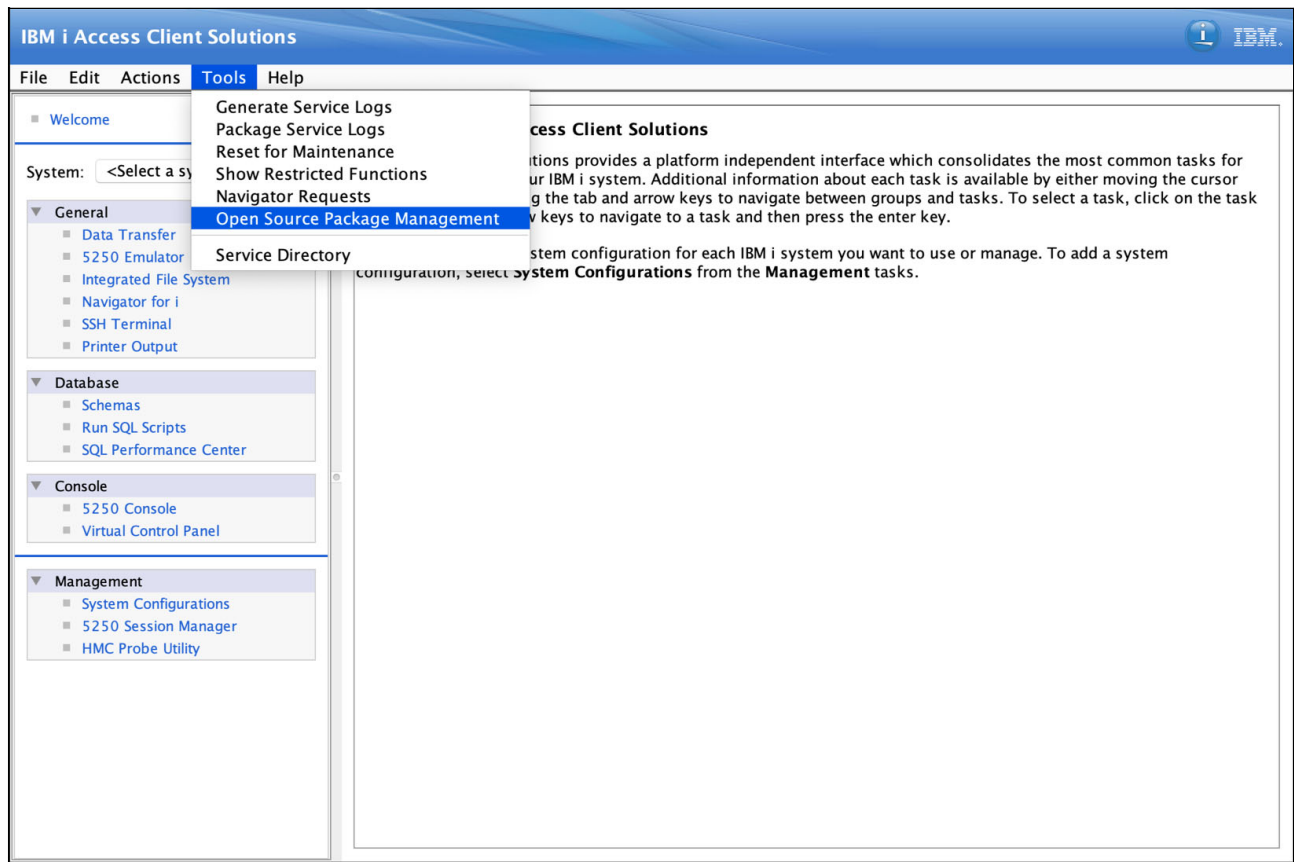


Figure 2-2 Open source package management option from IBM i ACS

## Supported browsers for the Db2 Mirror GUI

- ▶ Chrome
- ▶ Edge
- ▶ Firefox
- ▶ Safari

## 2.3 Objects eligible for replication

The Db2 Mirror environment supports replication of the object types that are typically updated in user applications.

### 2.3.1 IBM i objects

Table 2-5 lists the IBM i object types that are eligible for replication by Db2 Mirror. For more information about the replication for specific object types, see [IBM Knowledge Center](#).

Table 2-5 IBM i object types eligible for Db2 Mirror replication

Object type	Description
*AUTL <sup>a</sup>	Authorization list
*DTAARA	Data area
*DTAQ	Data queue
*ENVVAR <sup>b</sup>	Environment variable
*FCNUSG <sup>ab</sup>	Functional usage
*FILE <sup>c</sup>	File
*JOBQ	Job description
*JOBQ	Job queue
*JRN	Journal
*LIB	Library
*OUTQ <sup>d</sup>	Output queue
*PGM	Program
*SECATR <sup>ab</sup>	Security attribute
*SQLPKG <sup>e</sup>	SQL package
*SQLUDT	SQL user-defined type
*SQLXSR	SQL XML schema repository
*SRVPGM	Service program
*SYSVAL <sup>b</sup>	System value
*USRPRF <sup>a</sup>	User profile

a. Objects of this type are always included in replication through Db2 Mirror.

b. This is a pseudo-object type that is defined for use by Db2 Mirror. The default inclusion state does not apply to this type.

c. Device files, except for SQL aliases, cannot be replicated.

- d. Spooled files are not specified as replicated objects. An output queue can be specified as a replicated object, which causes all spooled files within that output queue to be replicated.
- e. Extended dynamic packages cannot be replicated.

## 2.3.2 Definition only

You might need to replicate only the definition of an object and not replicate the content. This option is supported within Db2 Mirror for the object types of \*DTAQ, \*FILE, and \*LIB. For a data queue, only the definition can be replicated. The content of a data queue is never replicated. A file can be replicated either as DEFINITION or all the content can be replicated too.

## 2.3.3 Integrated File System objects

Integrated File System (IFS) objects are made accessible on both Db2 Mirror nodes by using a different technology than the replicated IBM i object types. To be accessible, the IFS objects must be contained within an Independent Auxiliary Storage Pool (IASP), and the IASP must be part of a PowerHA cluster resource group (CRG).

## 2.4 Replication rules

The definition of the list of objects to be replicated between a Db2 Mirror node pair is determined by three criteria:

- ▶ Default inclusion state, which is defined at initial setup.
- ▶ System rules.
- ▶ The replication criteria list (RCL), which is user-defined.

The combination of these criteria defines replication rules for IBM i objects in the Db2 Mirror environment.

### 2.4.1 Default inclusion state

The *default inclusion state* is defined during the initial Db2 Mirror setup and cannot be changed except by recloning the System Auxiliary Storage Pool (SYSBAS) or a database IASP (DB IASP) and removing all existing rules, as described in 4.4, “Recloning” on page 65.

This is the highest level rule, and it has two possible settings:

- ▶ Exclude: No objects are replicated by default. If you want specific objects or libraries to be included in replication, you must create rules to identify the objects that are to be included.
- ▶ Include: All eligible objects are replicated by default. You can create more rules to exclude objects from replication.

You may select the option *exclude* for the default inclusion state if there is no general need to replicate all eligible objects from SYSBAS or the IASP. Then, you can add rules for objects or group of objects that must be replicated between both nodes.

Figure 2-3 shows the Db2 Mirror for i GUI view of the initial setup where the Default Inclusion State is defined. It lists all configured IASPs on the system, such as PRIMASP. The option to include or exclude IASPs in the Db2 Mirror replication must be selected.

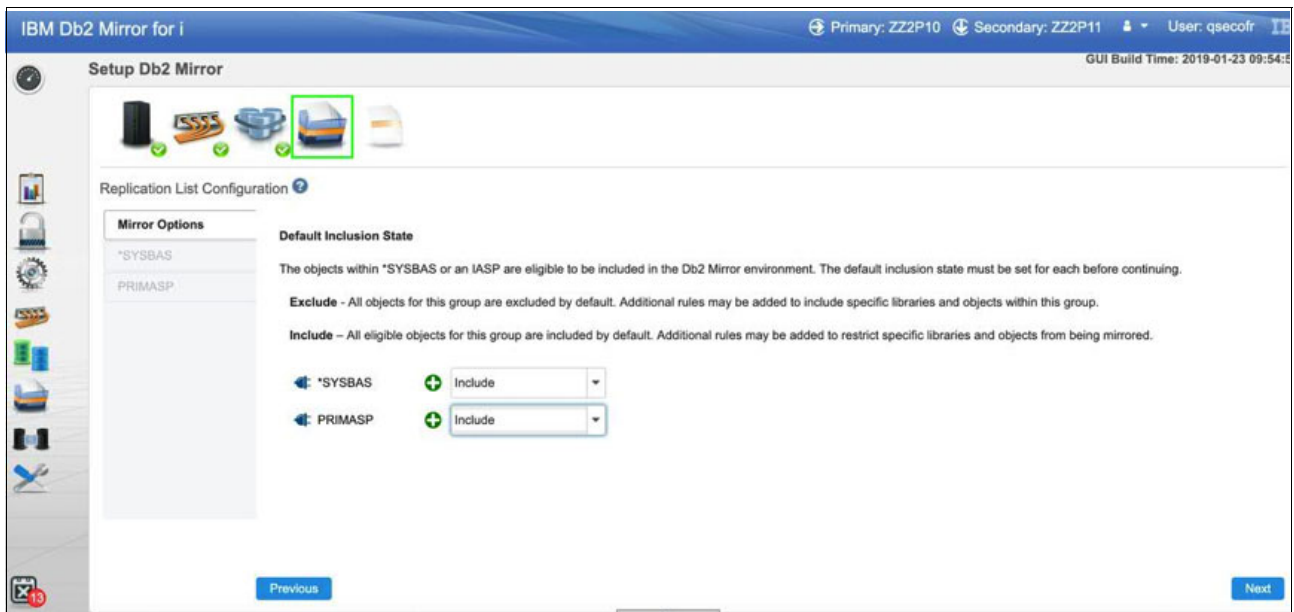


Figure 2-3 Default inclusion state during the initial setup

## 2.4.2 System rules

When the default inclusion state is set for SYSBAS or for a DB IASP, Db2 Mirror adds *system-defined object rules* to the RCL. These rules are for objects that are always replicated by Db2 Mirror or for objects that will never be replicated by Db2 Mirror. These rules cannot be modified or overridden by a user rule. Examples of system rules that are added by Db2 Mirror are related to user profiles, authorization lists, and system values.

Figure 2-4 shows the Manage Replication List view for the system rules, which are filtered by user profile, system value, and include in replication.

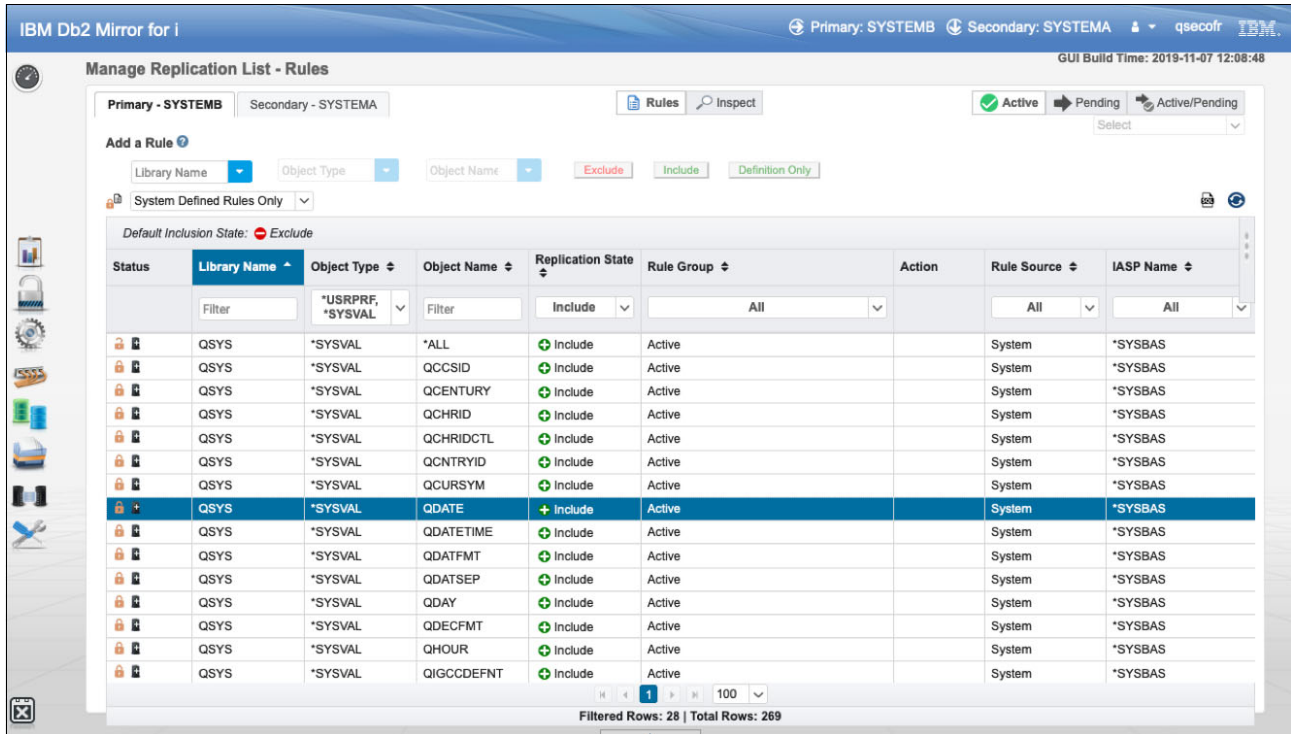


Figure 2-4 Manage Replication List: system rules

### 2.4.3 Replication criteria list

The RCL is a rules engine that is used by Db2 Mirror to determine whether an object should be replicated.

The RCL consists of a set of rules identifying groups of objects that should or should not be replicated. The rules, which are combined with the default inclusion state and system rules, provide a concise process to determine the replication status of any existing or future object.

Because Db2 Mirror supports objects in SYSBAS and objects that are within a DB IASP, an independent RCL exists for each IASP that is registered as a DB IASP with Db2 Mirror.

Figure 2-5 shows the Db2 Mirror GUI dialog box Manage Replication List – Rules for managing the RCL and its rules. As a default inclusion state of *exclude* was configured, the user must add rules to include objects to be replicated between the Db2 Mirror node pair.

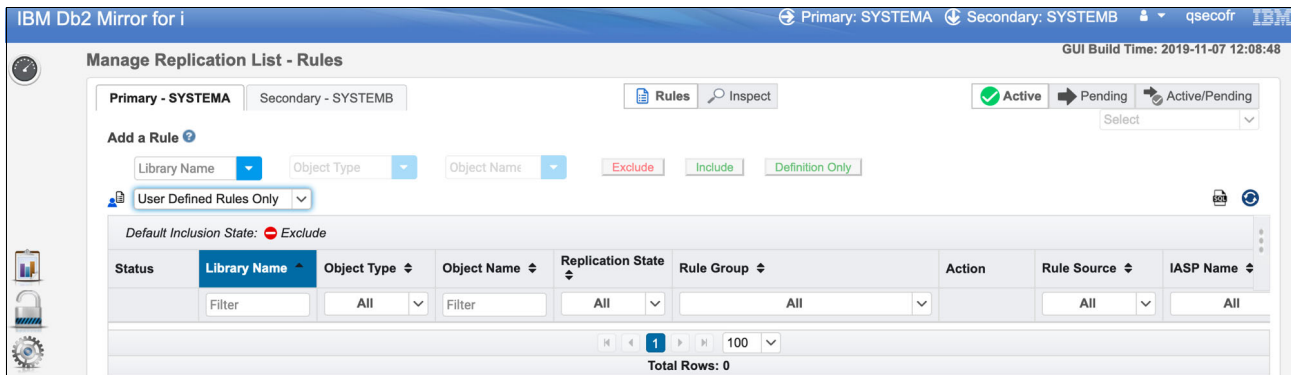


Figure 2-5 Manage Replication List with the default inclusion state of *exclude* for \*SYSBAS

## 2.5 Object tracking list suggested priorities

The user can optionally define *user-suggested priorities* for objects that Db2 Mirror considers when resynchronization is required. By defining suggested priorities, the user can define a data resiliency precedence for selected data based on its importance from a business perspective.

Db2 Mirror determines the order in which objects are resynchronized based on their object type and object dependency on other objects. The suggested priority is defined as a positive integer number with 0 being the highest priority, but it is only a relative and not an absolute priority. The default priority for objects is the null value, which indicates that no user suggested priority is defined.

Figure 2-6 shows an example of a user-suggested priority where the user defined the highest resynchronization priority for all objects from the ITSOLIB library.

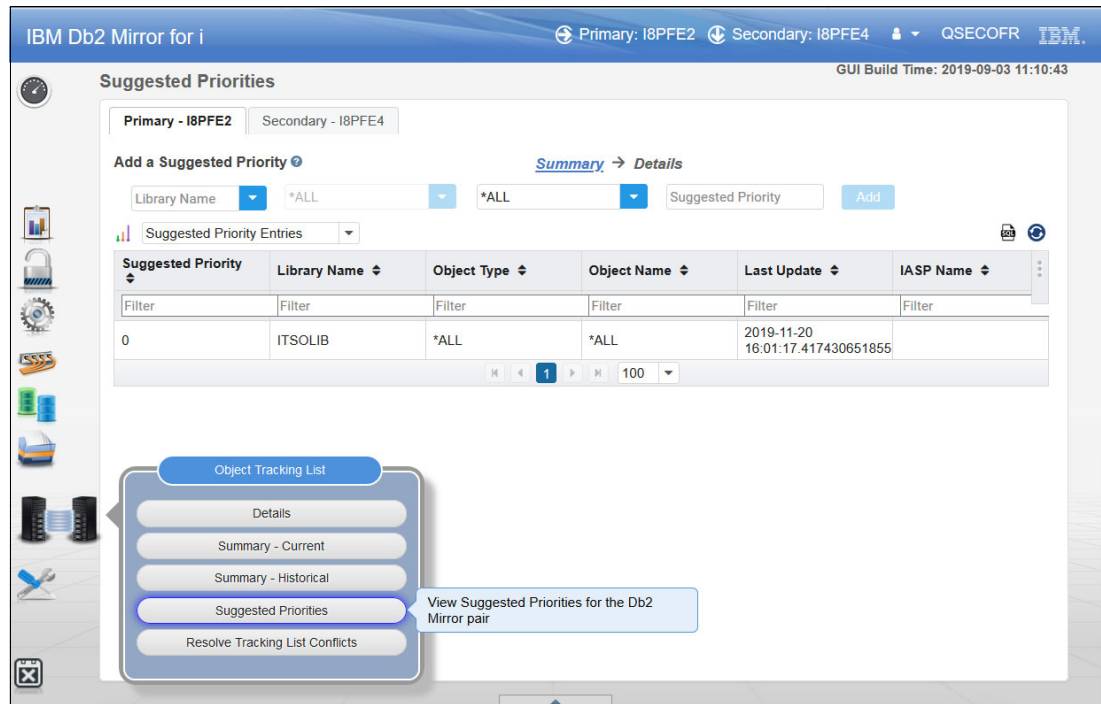


Figure 2-6 Db2 Mirror suggested priorities

## 2.6 Application considerations

Db2 Mirror provides an active-active clustering solution with synchronous data replication between both nodes. It is designed for a recovery point objective (RPO) and recovery time objective (RTO) for outages that can be reduced to near zero.

The design of your production applications in terms of active-active support affects the RTO that is possible with Db2 Mirror for i. Changes to your application architecture can provide added benefits in a Db2 Mirror environment.

Here is a list of items that must be considered:

- ▶ Database replication jobs and activation groups that are used by Db2 Mirror.
- ▶ JTOpen Java Database Connectivity (JDBC) driver support for application availability.
- ▶ Job queues and job scheduled entries considerations with Db2 Mirror.
- ▶ Querying data from a replicated database file and Db2 Mirror status.
- ▶ Database trigger considerations.
- ▶ Db2 Mirror exit points.
- ▶ Application behavior when the replication state is BLOCKED.
- ▶ XA distributed transaction environment.
- ▶ IBM MQ and Db2 Mirror.

You can find more information about these considerations at [IBM Knowledge Center](#).



## 2.7 Quorum

A *quorum node* for a Db2 Mirror cluster serves as an arbitrator for determining a node's correct replication state if both primary and secondary Db2 Mirror nodes were down.

Quorum data, including each node's last replication state, is kept synchronized among all Db2 Mirror cluster nodes so that if the partner node is down when a node performs an initial program load (IPL), the quorum node (a third node of the Db2 Mirror cluster) helps to determine a node's correct replication state of TRACKING or BLOCKED.

**Note:** Configuring a quorum node for a Db2 Mirror cluster is optional but highly recommended because it provides extra resiliency to properly recover from dual-node outages. Without a quorum node, both nodes would be set to a BLOCKED state when they start, and the Db2 Mirror administrator would need to decide to force one of the nodes to a TRACKING state.

The quorum node does not require extra resources for its role as a quorum node or a Db2 Mirror license. A Db2 Mirror management node or a node of a high availability (HA) disaster recovery (DR) (HA/DR) configuration may be designated as a quorum node. However, as the quorum node must be part of a Db2 Mirror cluster, it must be dedicated to a cluster that cannot be shared among multiple Db2 Mirror clusters.

## 2.8 Cluster monitors

A configuration of a cluster monitor on a Db2 Mirror cluster node allows the node to query its HMC through a REST API interface for node failure conditions to help distinguish between a possible communication failure or an actual failure of the other node.

The cluster monitor is also used for automatic file system mutations of registered IFS IASPs, that is, a mutation from a client file system instance to a server file system by an IASP failover or from a server file system instance to a stand-alone file system.

**Note:** Adding a cluster monitor to the quorum node is a best practice because it helps correctly set a node's replication state after recovery from situations where both Db2 Mirror nodes failed concurrently.

A cluster monitor can be configured during the initial Db2 Mirror setup or added later, as shown in Figure 2-7. It requires manually configuring a system certificate store on the node before you use the IBM i DCM GUI to store the HMC's certificate for secure REST API communications.

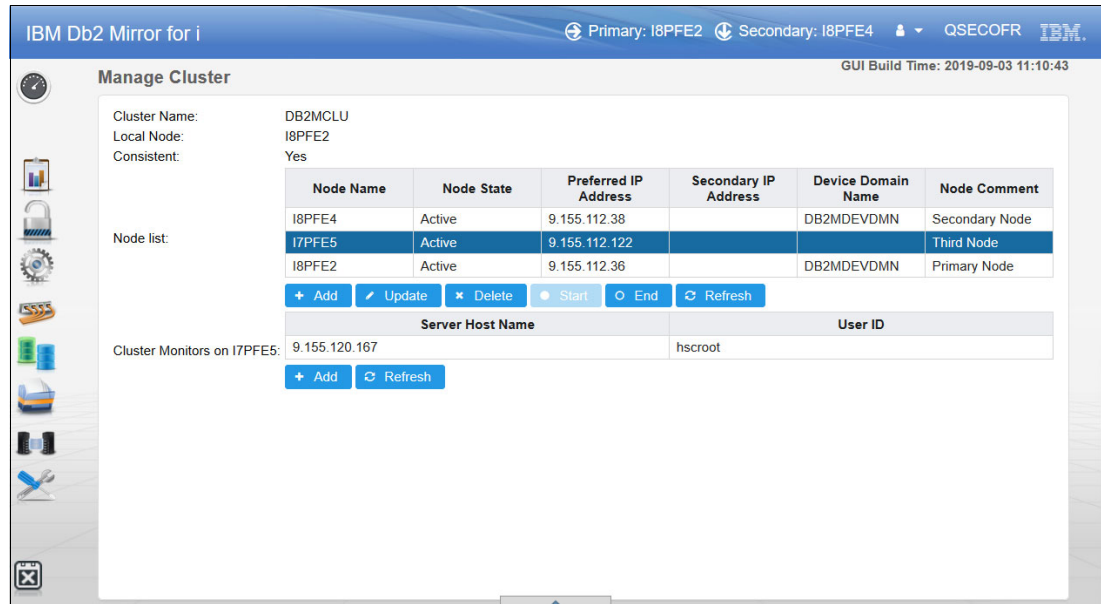


Figure 2-7 Db2 Mirror GUI Manage Cluster

## 2.9 Network Time Protocol server

Db2 Mirror requires NTP-based system time synchronization between the primary and secondary nodes. Both nodes also must be configured with the same time zone setting. The time synchronization between the Db2 Mirror nodes is especially important for logging, temporal tables, and time-based applications.

There are three options for configuring time synchronization at Db2 Mirror setup:

- External time server** Both nodes reference the same external NTP server, which can be either a public one or one within the customer's network.
- Chained time server** The setup source node is configured as a Simple Network Time Protocol (SNTP) server and references an external NTP server. The setup copy node references the setup source node, but also references the external time server as a non-preferred server.
- Peer time server** The setup copy node references the setup source node, which is configured as an SNTP server.

**Note:** If you reference more than one external time server, configure at least four external time servers to help ensure that the times that the IBM i NTP client receives from the different time sources fall into a small enough range so that the client can figure out which one is correct.

## 2.10 Performance

The synchronous replication of Db2 Mirror for database type objects and the IFS mutable file system client/server model demand some considerations about performance:

- ▶ With synchronous replication, the complete I/O path length for database type object change operations increases because the operation drives I/O on both nodes to finish. A typical increase by a factor around 2 - 3x is expected, with single-threaded or serial I/O workloads being most impacted.
- ▶ Database read workloads are not impacted because they are served locally and not replicated.
- ▶ The ability to balance the workload by running transactions on both nodes mitigates the per transaction impact with a target of achieving equal to or greater transactional throughput compared to a single non-mirrored system.
- ▶ For IFS IASPs, the performance might differ depending on which node a file system operation is being initiated from. Users on the server node where the IASP is varied on might experience faster response times than users on the client node.

## 2.11 Supported storage

Db2 Mirror supports external SAN storage only for both of its nodes.

The background for this requirement is that during the initial setup of Db2 Mirror, the so-called *setup source node* is cloned to the *setup copy node*, and this cloning is only supported for external SAN storage.

For planning purposes, the following IBM statement of general direction from the IBM i 7.4 TR1 announcement might be helpful:

“IBM plans to introduce the support for internal disk as a storage option to the Db2 Mirror for i product.”

Db2 Mirror supports automatic cloning for IBM System Storage DS8000 series or IBM Spectrum Virtualize based storage like IBM Storwize® and IBM SAN Volume Controller by either using IBM FlashCopy® (if you use a single storage system for both the setup source and setup copy node) or *remote copy* replication (if you use two storage systems). For other IBM i supported SAN storage systems, Db2 Mirror provides a *manual copy* option for the cloning where the user is required to manually perform the volume copy operations on the storage systems.

## 2.12 Migration considerations

Depending on the configuration of the existing IBM i environment, different considerations apply for the migration to Db2 Mirror.

### IFS migration

As described in 1.3.6, “Database IASP versus IFS IASP” on page 7, if IFS data must be made simultaneously accessible from both Db2 Mirror nodes, it must be included in an IASP that must be registered as an IFS IASP with Db2 Mirror.

## IASP migration

An existing IASP containing both IFS data files and database type objects that must be simultaneously accessible from both Db2 Mirror nodes must be split into two IASPs: one for the IFS data that must be registered as an IFS IASP, and the other one for the database objects that must be registered as a DB IASP with Db2 Mirror.

## Clustering migration

Cluster information on a setup source node from a system that was used in a PowerHA environment before it is preserved by the Db2 Mirror setup because it is deleted on the setup copy node after the cloning to allow it to join the existing cluster.

## 2.13 Backup and recovery considerations

Some special considerations apply for save and restore operations in Db2 Mirror that the user should be aware of to implement a successful backup and recovery strategy.

### Backup considerations

- ▶ When using an IFS IASP, save and restore operations must be performed on the *server* node where the IASP is varied on. Save and restore for the IFS IASP cannot be performed from the *client* node.
- ▶ A *full-system save* (by using **GO SAVE** menu option 21 or BRMS backup control group \*SYSTEM) should be performed for the *primary node*. Unless it is performed from a FlashCopy image, it is considered a planned outage because it involves going into a restricted state, so it must be handled by performing a Db2 Mirror role swap before and after the system save in order for production work to continue while the save is done on the primary.

The reason that the primary is the preferred system for a full system save is that it is the system that is restored and then cloned in case a full recovery for the Db2 Mirror environment is required. The underlying presumption is that the primary's personality and non-mirrored data is the favored system in a typical production environment.

Before running the backup from the primary node, the user should verify that the data area QTEMP/SRMIRCTL does not exist (see the following bullet).

- ▶ For a backup of the secondary node, only non-replicated user data must be backed up (use the **GO SAVE** menu option 23 or BRMS backup control \*BKUGRP).

To omit replicated data from the backup, create a data area that is named QTEMP/SRMIRCTL as follows:

```
CRTDTAARA DTAARA(QTEMP/SRMIRCTL) TYPE(*CHAR) VALUE('1')
```

- ▶ The QUSRBRM library containing the BRMS save history information is excluded from Db2 Mirror replication by a system-defined rule. When using BRMS, it is a best practice to include both Db2 Mirror nodes in a traditional BRMS network for the synchronization of save history information.
- ▶ Independent from Db2 Mirror, for IFS backups it is a best practice to use separate BRMS backup control groups for SYSBAS and each IASP. This approach enforces BRMS to create separately saved items for the IFS data to help prevent duplicate restores causing longer recovery times and undesired restores into SYSBAS if the IASP is not available.

## Restore considerations

To perform a complete system restore for a Db2 Mirror environment from a system backup that was created by using the backup strategy that is described in “Backup considerations” on page 34, complete the following steps:

1. Recover the primary node.
2. Reconfigure Db2 Mirror for the secondary node (which includes recloning).
3. Recover the secondary node’s user data.

For a selected restore of objects that are always replicated by Db2 Mirror like user profiles, authorization lists, and information like function usage information and security attribute information, more considerations apply to keep object ownership and authorities synchronized for replicated objects:

- ▶ Because a restore of all user profiles and authorities requires the IBM i system to be in a restricted state, the **RSTUSRPRF** and **RSTAUT** commands must be run on the primary node, which is in tracking state because the restore of authorities is not allowed for most objects when the replication state is **BLOCKED**.
- ▶ Similarly, a restore of the **QUSRSYS/QUSEXRGOBJ** object, which includes the function usage information, must be restored on the primary node if replication is suspended.

## 2.14 IBM Lab Services offerings

Db2 Mirror is a new technology and a new approach to building resiliency in IBM i environments. To help customers who want to consider and implement IBM Db2 Mirror for i, IBM Systems Lab Services developed an offering that is called *Db2 Mirror for i Readiness Assessment*. It is a customized consulting engagement that provides customers, IBM Business Partners, and ISVs the skills that are required to implement a Db2 Mirror solution, and the ability to test their applications in a Db2 Mirror lab environment. At the end of the workshop, attendees have the skills that are necessary to start planning their Db2 Mirror environment.

For more information, see [IBM Systems Lab Services](#).





## Implementation overview

This chapter provides an overview of the initial setup for IBM Db2 Mirror for i by using the Db2 Mirror GUI setup wizard.

As mentioned in 1.5, “Interfaces” on page 15, the initial setup for Db2 Mirror can also be done by using SQL services and the **db2mtool** utility. For more information about a command-line based setup, see [IBM Knowledge Center](#).

## 3.1 Db2 Mirror setup process

As described in 2.1, “Db2 Mirror setup overview” on page 20, the setup for a new Db2 Mirror configuration requires three nodes with an IBM i management node that is required for the setup. The management node orchestrates the cloning of the IBM i setup source node to the setup copy node for creating the second node of the Db2 Mirror pair.

The setup process that uses the Db2 Mirror GUI setup wizard involves the following steps:

1. Do the pre-configuration (to be done by the user before the Db2 Mirror setup):
  - Check that all requirements for Db2 Mirror are met, which are described in 2.2, “Hardware and software requirements” on page 21.
  - Collect all the required information for the setup by using the [Db2 Mirror checklist](#).
  - Create the setup copy partition on the Hardware Management Console (HMC).
  - Configure the external storage for the setup copy node.
2. Configure the Db2 Mirror settings (by using the GUI setup wizard):
  - Configure the source and copy nodes (including the HMC and node information, cloning method, and cluster and time server information).
  - Configure the network redundancy group (NRG).
  - Configure the storage (including storage server access, System Auxiliary Storage Pool (SYSBAS), and database Independent Auxiliary Storage Pools (DB IASPs) storage configuration information).
  - Configure the replication list (including the default inclusion state selection for SYSBAS and IASPs, and the optional user-defined replication rules).
3. Validate the Db2 Mirror configuration (done automatically at the end of the Db2 Mirror GUI setup wizard after the user enters all the configuration settings).
4. Perform the cloning process (started by the Db2 Mirror GUI setup wizard after successful validation).

The setup source node is cloned either by using the copy services functions of the supported IBM external storage systems or by using a manual copy that is performed by the user for non IBM external storage to create the setup copy node. By using the installed and configured cloud-init software package, the setup copy node is configured and initialized automatically by Db2 Mirror during its first initial program load (IPL).
5. Validate whether the cloning completed successfully and whether active replication began between the two Db2 Mirror nodes (automatically started at the end of the Db2 Mirror GUI setup wizard).

## 3.2 Initial Db2 Mirror setup by using the GUI setup wizard

When setting up the Db2 Mirror environment by using the setup wizard that is part of the Db2 Mirror GUI, you must run the GUI on a third IBM i node. This node can be used after the initial setup as the GUI management node for the newly created Db2 Mirror pair and for many other ones.

The following sections provide a summary of the Db2 setup wizard. For providing a concise overview about the setup process, they do not contain screen captures from all the dialog boxes that you use during the setup.



For a more comprehensive view of the Db2 Mirror setup wizard process, see [IBM Knowledge Center](#).

To connect to the GUI interface on the third node, use the following case-sensitive URL:

`http://hostname:2006/Db2Mirror`

### 3.2.1 Preparing the setup copy node environment

Before setting up Db2 Mirror for i by using the GUI interface, perform the following preliminary configuration:

1. Create the setup copy node logical partition (LPAR) on the HMC.
2. Create host connection and logical volumes on the external storage system.
3. Create appropriate SAN switch zoning.

### 3.2.2 Configuring the source and copy node

This setup section introduces into the following information that is required to be specified during the setup by the user:

- ▶ HMC information
- ▶ Node information
- ▶ Cloning method
- ▶ Cluster information
- ▶ Time server

#### **HMC information**

IBM i LPARs that are part of the Db2 Mirror environment pair can be managed by the same HMC or separate ones.

Figure 3-1 shows the initial setup window for entering the HMC information and which IBM i LPARs are used as a Db2 Mirror nodes pair.

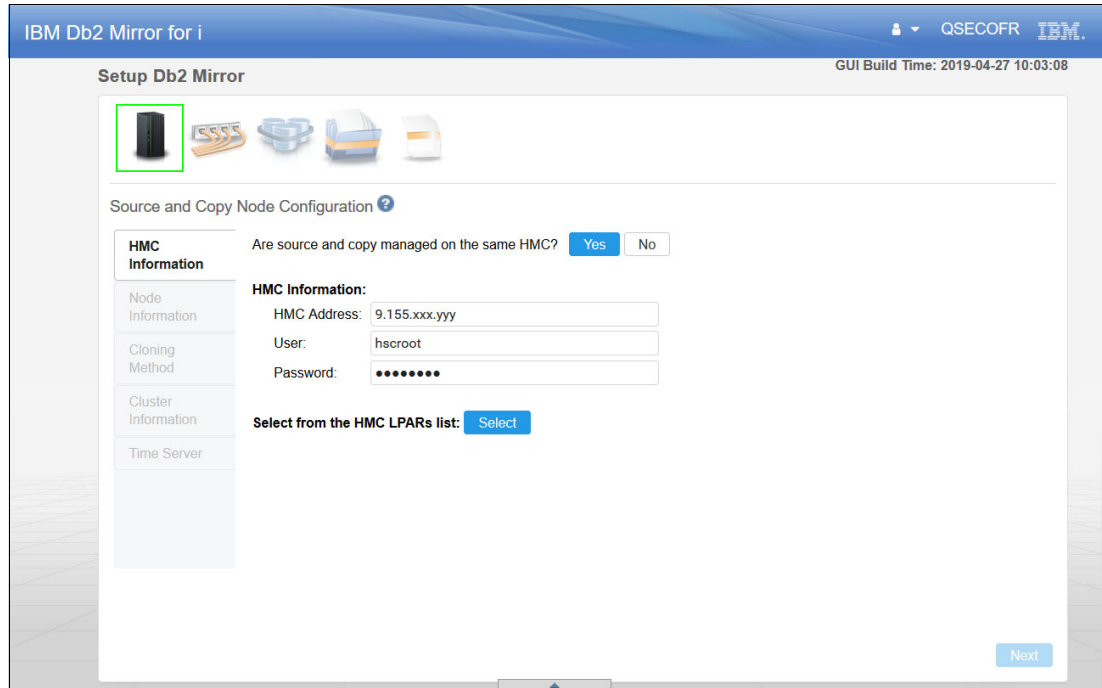


Figure 3-1 Db2 Mirror setup wizard: HMC information

**Note:** HMC REST services are used by Db2 Mirror during the setup and configuration process. The HMC REST services are provided through HTTPS, which might require Db2 Mirror to install a certificate. You see the dialog box that is shown in Figure 3-2 before the certificate is installed. You can check the certificate information before clicking **OK** to add the certificate.

Figure 3-2 shows the dialog box with the Certificate Install Confirmation request.

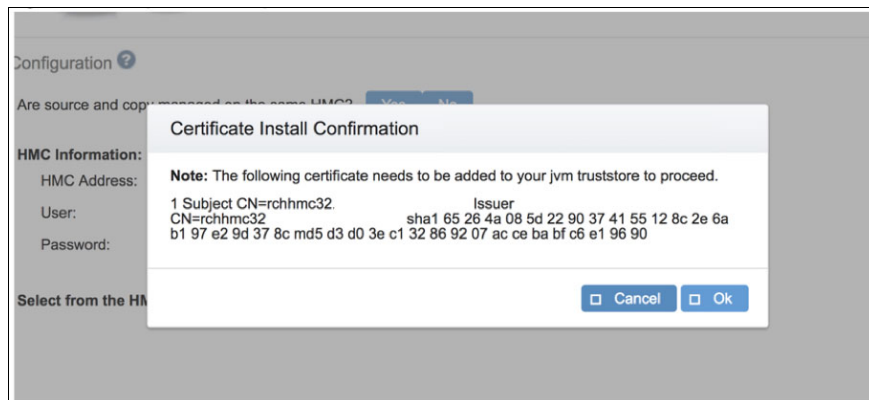


Figure 3-2 Db2 Mirror setup wizard: HMC Certificate installation confirmation

### Node information

This step defines the node information for the setup source node and the setup copy node. The setup source node must be running.

Figure 3-3 shows the initial setup window for the setup source node and setup copy node information.

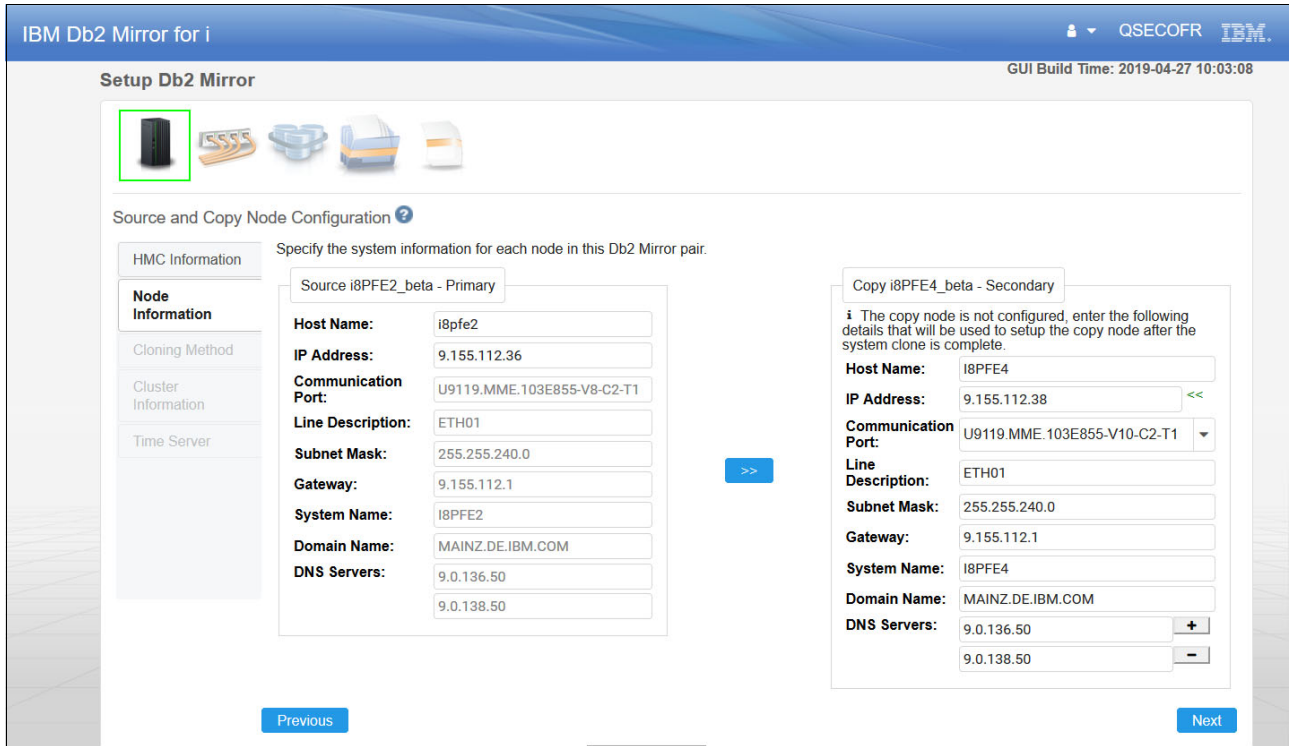


Figure 3-3 Db2 Mirror setup wizard: Node information setup

After clicking **Next**, the Db2 setup wizard verifies the prerequisites and the setup source node Db2 Mirror status.

### Cloning method

The cloning method defines how to copy data from the setup source node to the setup copy node. There are two options that must be specified:

- ▶ Cloning method
- ▶ Cloning process to use

Figure 3-4 shows the cloning method setup window. If using IBM Spectrum Virtualize or IBM System Storage DS8000 series storage, the setup process automates the node cloning by using either of the selected FlashCopy or Remote copy storage copy services functions. After you choose the cloning method, the cloning process method, **Cold Cloning** or **Warm Cloning**, must be selected. A cold cloning process involving an orderly power down of the setup source node is a best practice to ensure a consistent copy.

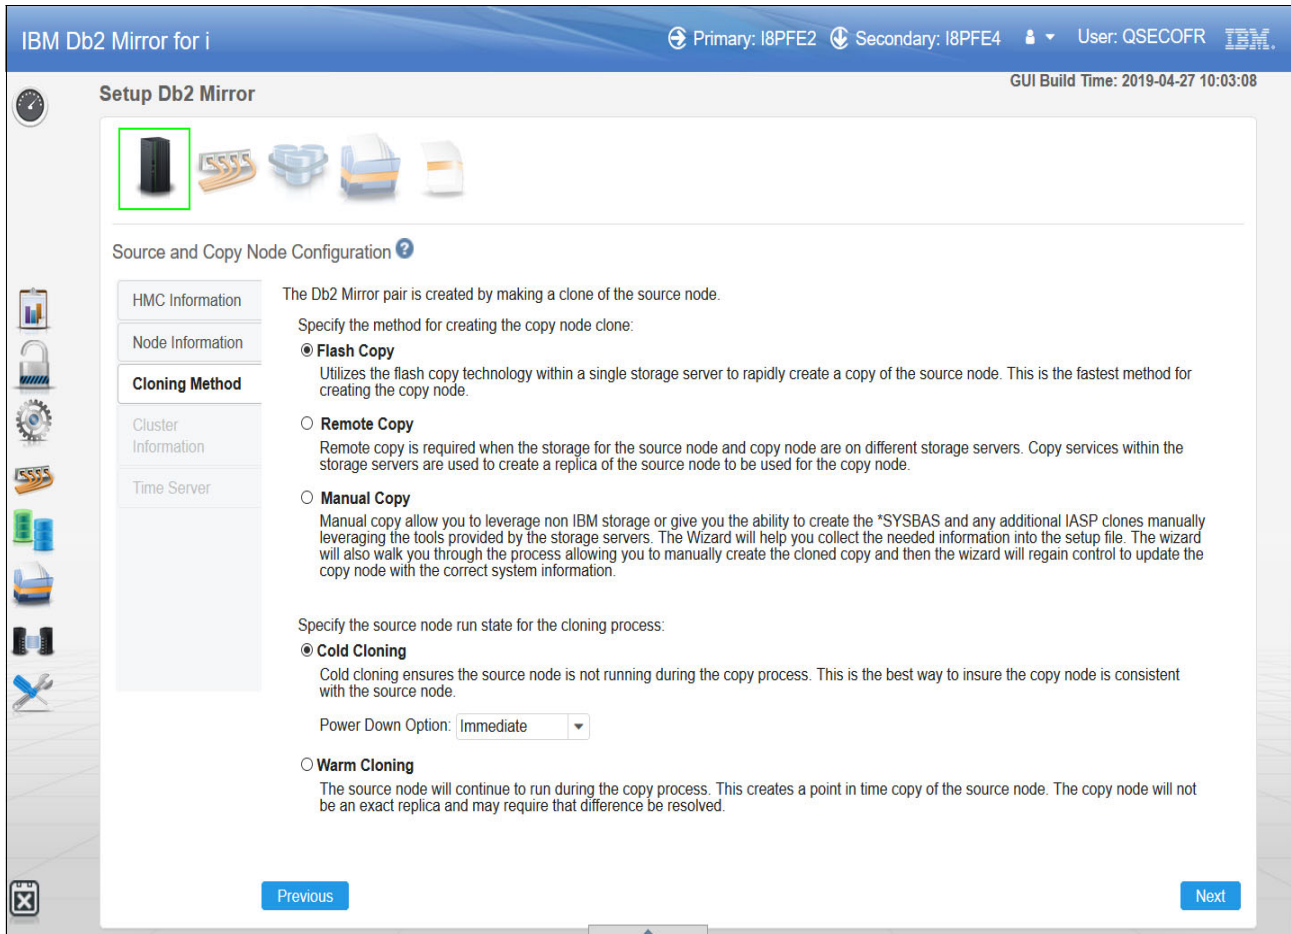


Figure 3-4 Db2 Mirror setup wizard: Setting up the cloning method

## Cluster information

Db2 Mirror uses clustering for maintaining quorum data, advanced node failure detection, and IASP management.

Figure 3-5 shows the clustering setup information. An extra node can be added as a quorum node, which is a best practice because it helps to ensure the reliability of the quorum data, as described in 2.7, “Quorum” on page 31.

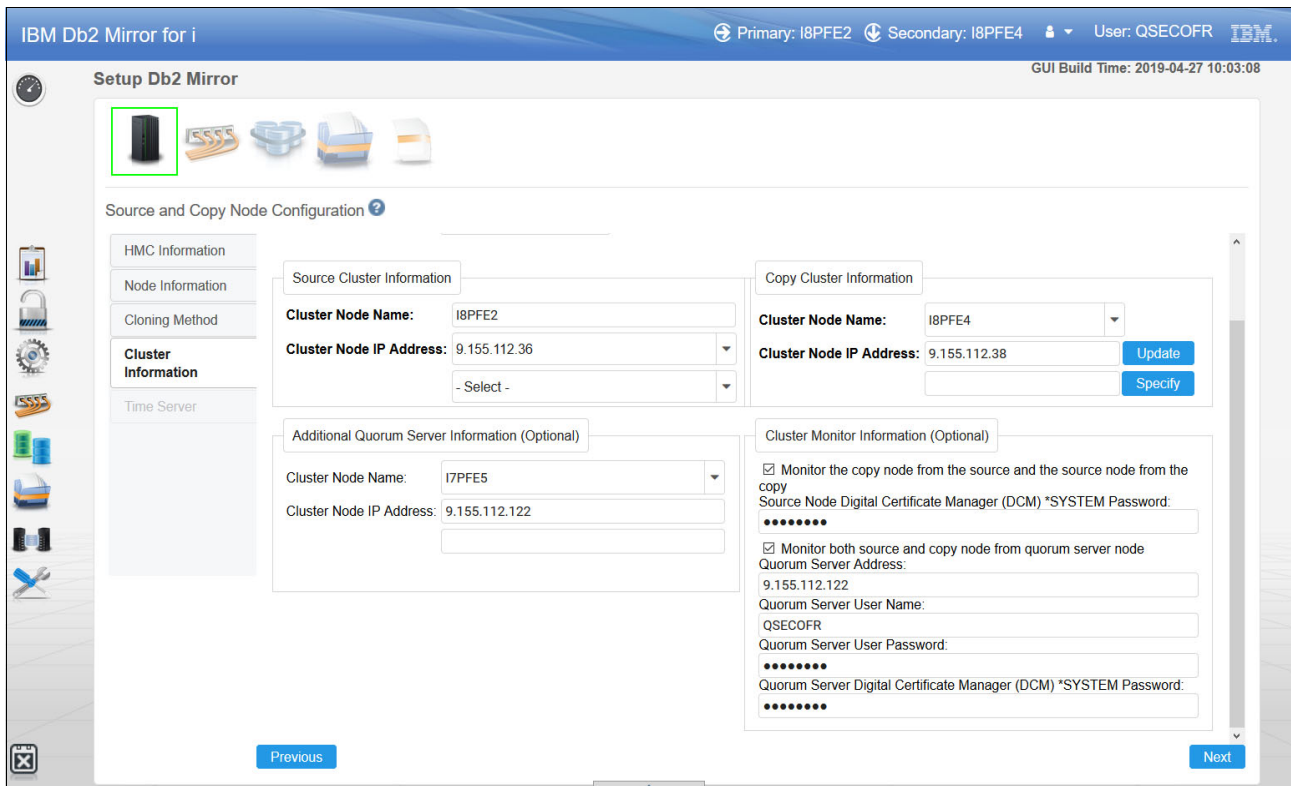


Figure 3-5 Db2 Mirror setup wizard: Cluster information setup

### Time server

A time server must be configured to keep the clocks for the nodes synchronized, as described in 2.9, “Network Time Protocol server” on page 32.

Figure 3-6 shows the setup window for the time server for the setup copy node. An external time server can be used, or you can use the setup source node as a chained time server.



Figure 3-6 Db2 Mirror setup wizard: Time server setup.

### 3.2.3 Configuring the network redundancy group

To configure the NRGs, you must specify the IP addresses for each Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) line.

Figure 3-7 shows the setup window to add an IP address pair between the setup source node and related setup copy node.

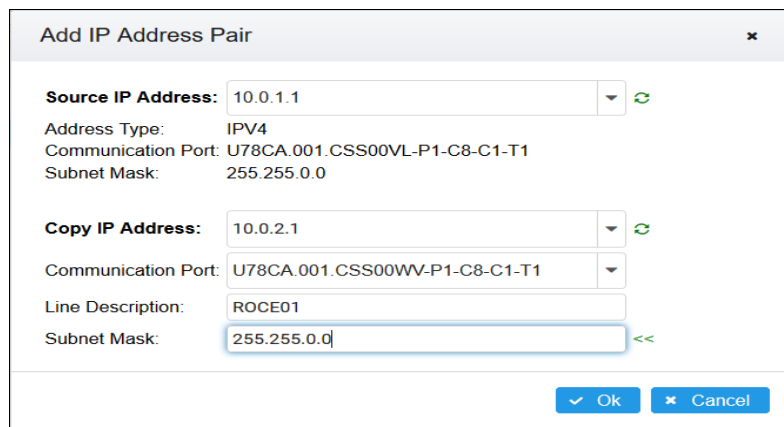


Figure 3-7 Db2 Mirror setup wizard: Add IP address Pair for NRG link

### 3.2.4 Configuring the storage server

Within the Storage Server information window of the setup wizard, the user specifies the storage system that is used by the setup source and setup copy nodes because we selected **FlashCopy** as the automatic cloning method before the setup wizard offers to select **DS8000** or **IBM Spectrum Virtualize** as supported storage systems for automatic cloning.

The following setup windows are related to the IBM Spectrum Virtualize setup. For more information about the DS8000 setup, see [IBM Knowledge Center](#).

Figure 3-8 shows IBM Spectrum Virtualize connection access information where *setup copy node* is defined. The following screen captures are taken from another Db2 Mirror configuration, so the names differ.

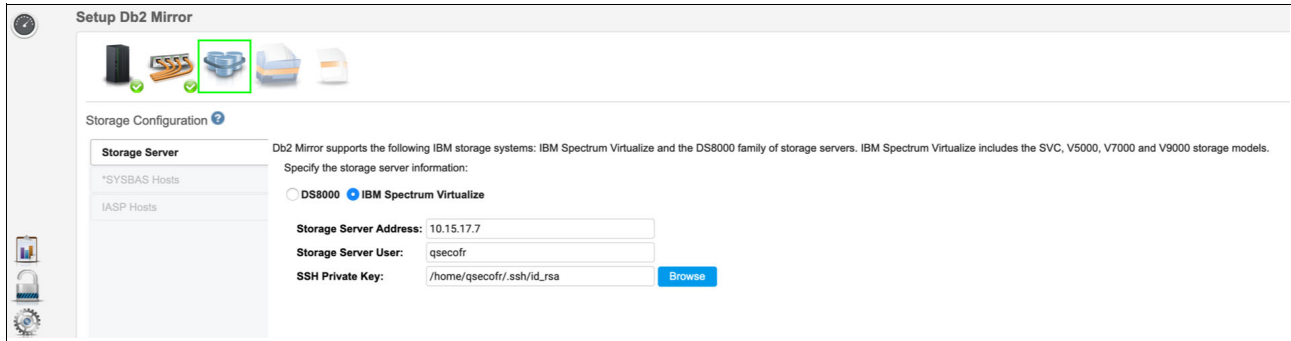


Figure 3-8 Db2 Mirror setup wizard: IBM Spectrum Virtualize access information

The setup and configuration process accesses an IBM Spectrum Virtualize storage server by using an SSH private key and an IBM System Storage DS8000 storage server through the DS Command-Line Interface (DS CLI). To access an IBM Spectrum Virtualize storage server, you must have the SSH private key uploaded to the GUI managing node and specify the path to where the SSH private key file is located or click **Browse** to go to the path where the key file is stored on the management node. For more information about how to set up and configure the key file, see [IBM Knowledge Center](#).

Figure 3-9 shows the result after selecting the Source Host information for the setup copy node. The volume sizes must be the same between the setup source node and the setup copy node.

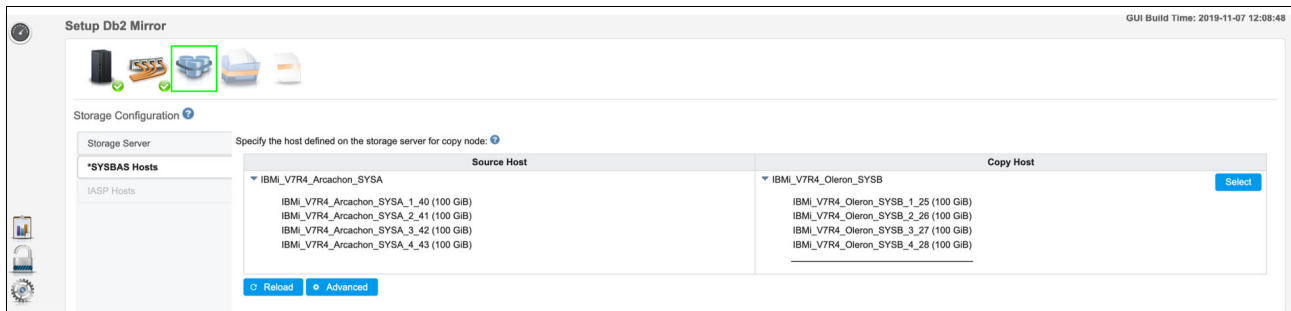


Figure 3-9 Db2 Mirror setup wizard: Storage Host selection for setup copy node.

The next window of the Db2 Mirror setup wizard (Figure 3-10) focuses on the selection of any existing IASPs as either to be cloned and becoming registered as a database IASP (DB IASP) or to be ignored for cloning, which means not becoming simultaneously accessible from both nodes. For illustration purposes, this screen capture was taken from another system configuration that includes IASPs.

**Note:** The initial setup wizard supports cloning and registration for DB IASPs only. Any existing or newly created IASPs to be registered as Integrated File System (IFS) IASPs with Db2 Mirror can still be registered after the initial cloning setup.

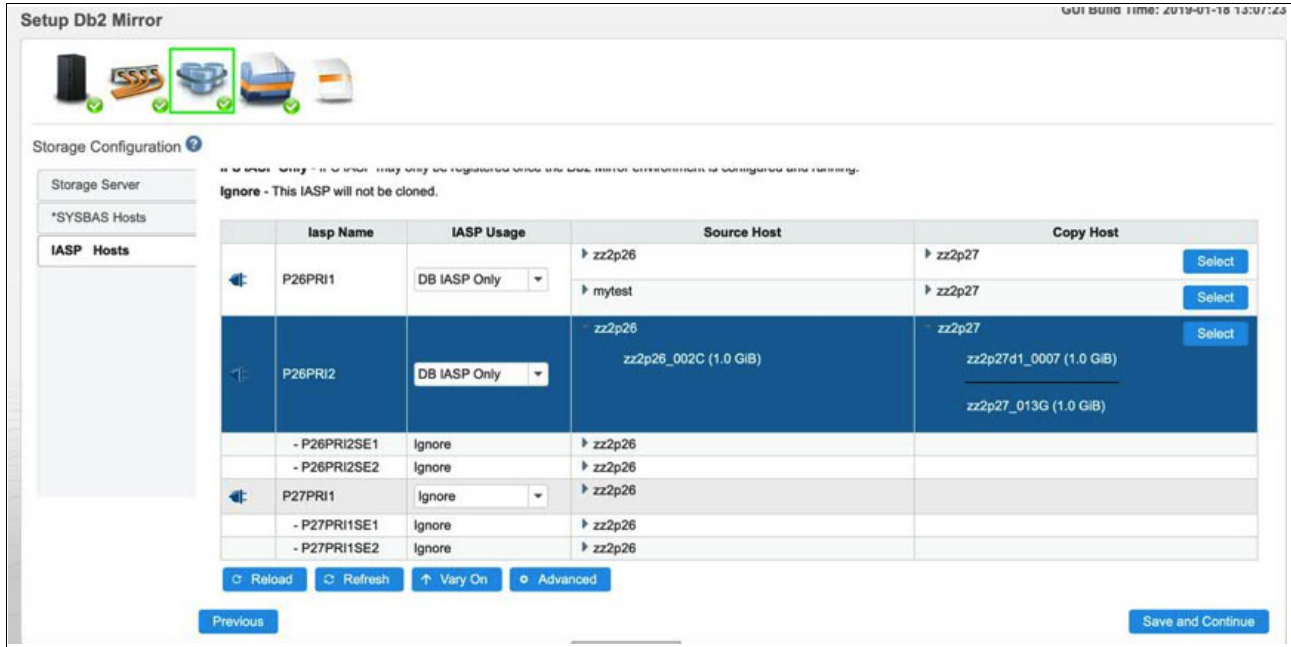


Figure 3-10 Db2 Mirror setup wizard: IASPs selection

### 3.2.5 Configuring the replication criteria list

Before the cloning of the setup source to the setup copy node can be performed, the *default inclusion state* must be selected and the initial replication rules must be defined.

The user must define whether by default if they want to replicate (the Include setting) or not replicate (the Exclude setting) all eligible objects within the SYSBAS and any existing IASP between both nodes. A unique default inclusion state and replication criteria list (RCL) exists for SYSBAS and for each IASP.



Figure 3-11 shows the Mirror Options window for this Db2 Mirror environment with a selected default inclusion state of Include for SYSBAS and one DB IASP.



Figure 3-11 Db2 Mirror setup wizard: Default inclusion state

Afterward, RCL rules can be defined for each of these auxiliary storage pools (ASPs). For this example, because all eligible objects were chosen to be replicated by default, the user can define rules to exclude some objects.

### Rules for \*SYSBAS

The \*SYSBAS Rules window of the setup wizard allows the user to define replication rules for objects within SYSBAS, that is, with the system ASP and any user ASPs. As all eligible objects from SYSBAS were to be included in Db2 Mirror replication by the prior selection of Include for the default inclusion state, some rules can be defined to exclude some objects from SYSBAS replication. If \*SYSBAS was excluded from replication by a default inclusion state selection of Exclude, rules could be created to include selected eligible SYSBAS objects in Db2 Mirror replication.

Figure 3-12 shows two rules that are defined for SYSBAS that are excluding two libraries and all their objects from Db2 Mirror replication.

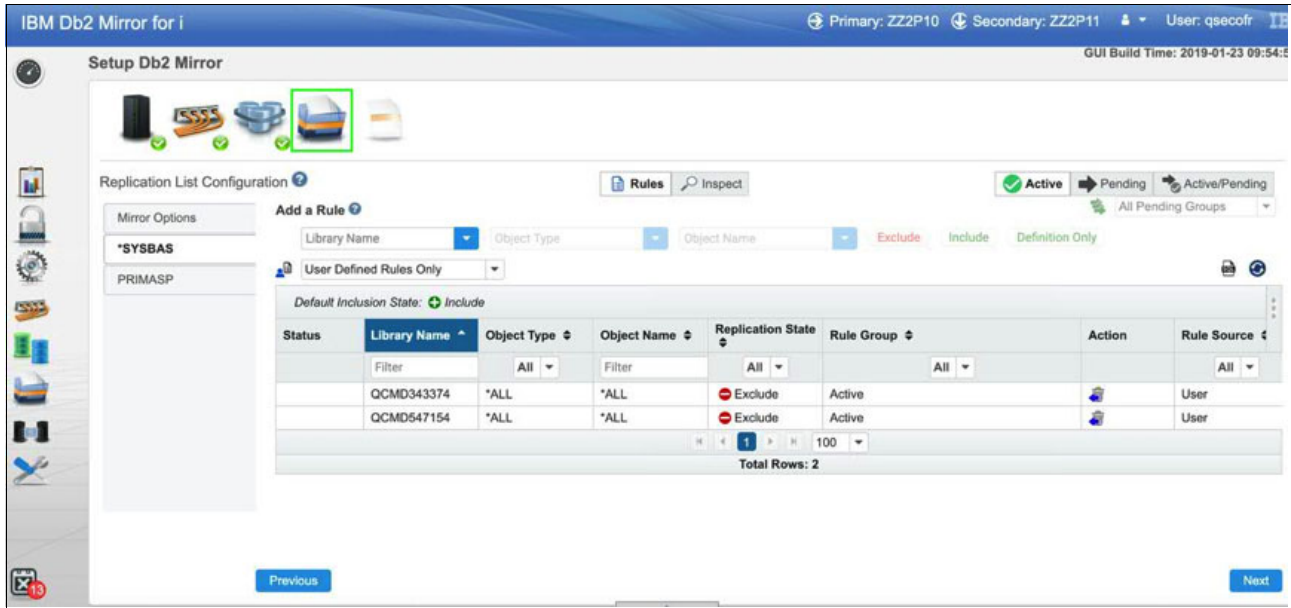


Figure 3-12 Db2 Mirror setup wizard: Rules definition for \*SYSBAS

### Rules: IASPs

For each DB IASP, Db2 Mirror maintains an RCL and the same rules concept applies like for SYSBAS.

Figure 3-13 shows the RCL for an Independent Auxiliary Storage Pool (IASP). Because in this example no user-defined rules are configured, all eligible objects from the IASP PRIMASP are replicated as defined by its default inclusion state of Include.

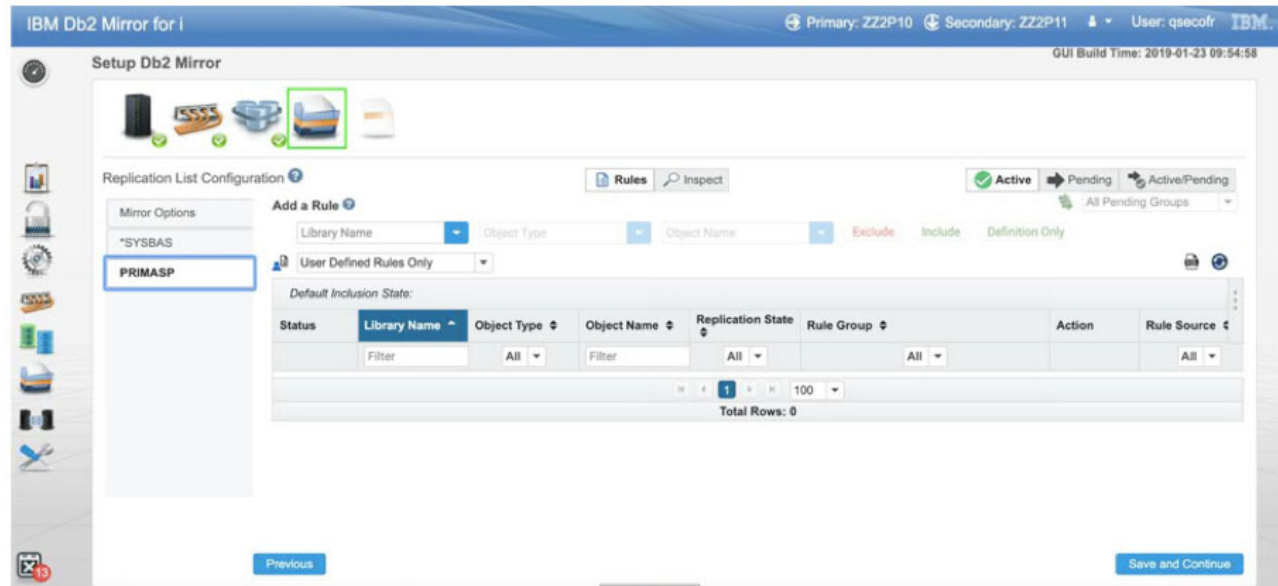


Figure 3-13 Db2 Mirror setup wizard: Rules definition for the IASP

## Inspecting

As you add rules to the RCL, use **Inspect** to understand how the rules affect the objects.

Figure 3-14 shows a view of all replication-eligible objects that are in this IASP (IASP33P) and how the specified rules affect the replication state of each object. This can be a useful tool for determining which objects are replicated and which objects are not replicated.

The screenshot displays the 'Setup Db2 Mirror' wizard interface. At the top, there are several icons representing different components, with a green checkmark next to the 'Db2 Mirror' icon. Below the icons is the 'Replication List Configuration' section, which has two tabs: 'Rules' and 'Inspect'. The 'Inspect' tab is active, showing two tables of data.

Mirror Options	Library Name	Replication State	Object Count	Object Name	Object Type	Object Replication State
IASP33P	Search			Search	Search	
	\$\$D2BSYSCF	Exclude	3	GLOBALV	*FILE	Exclude
	\$\$D2BTMPCF	Exclude	1	OBJLIST	*FILE	Exclude
	\$\$PROD	Exclude	1	QD2BOPTS	*FILE	Exclude
	\$BLDSHIP	Exclude	90	Total Rows: 3		
	\$BLDTEST	Exclude	377			
	\$DFPNAT003	Include	16			
	ZHONGJT	Exclude	16			
	ZJTTSTAA	Exclude	0			
	ZJTTSTC	Exclude	0			
	ZMJTEST	Exclude	36			
	Total Rows: 10					

Figure 3-14 Db2 Mirror setup wizard: Rules inspection

### 3.2.6 Validation and cloning

After defining all the details for your Db2 Mirror environment, the setup wizard presents a configuration summary for the defined Db2 Mirror setup, as shown in Figure 3-15.

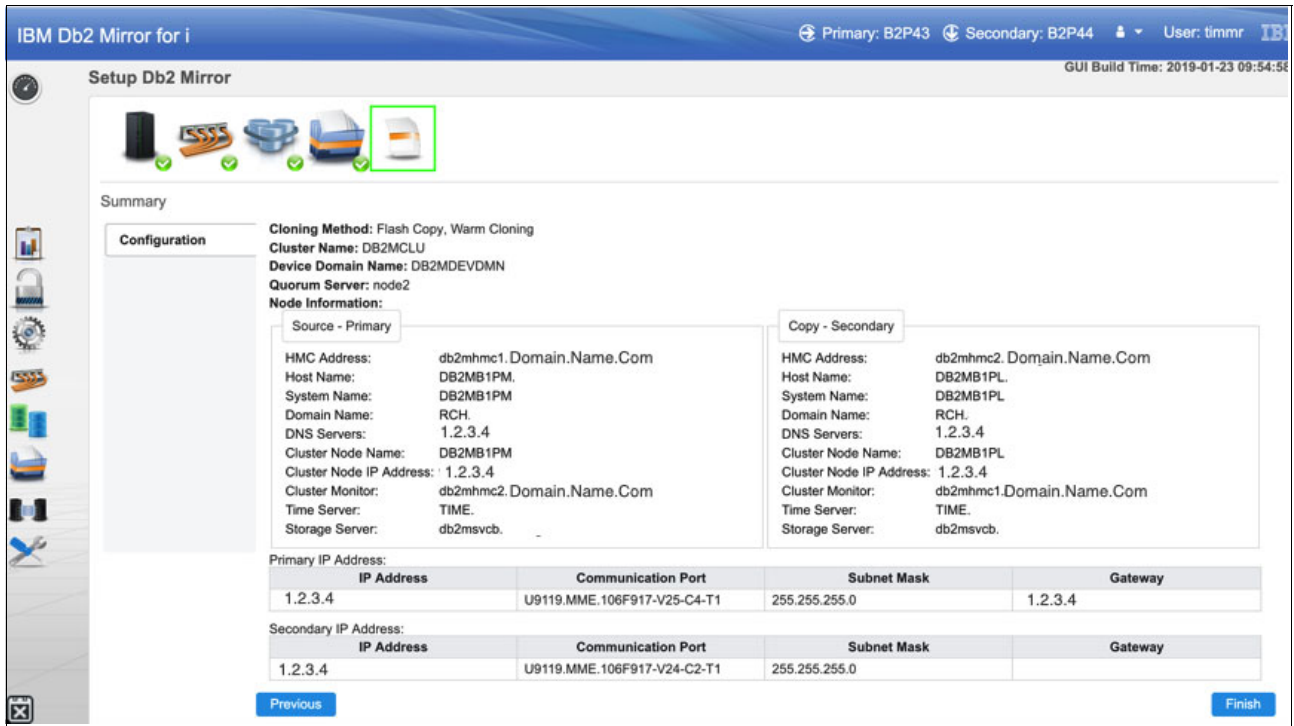


Figure 3-15 Db2 Mirror setup wizard: Setup information summary

Clicking **Finish** starts the validation process before the Db2 Mirror GUI setup wizard proceeds with performing the actual setup, including the cloning.

If the validation does not pass, an error dialog box opens.

Figure 3-16 shows the Setup Progress with some of its history and a successful completion status.

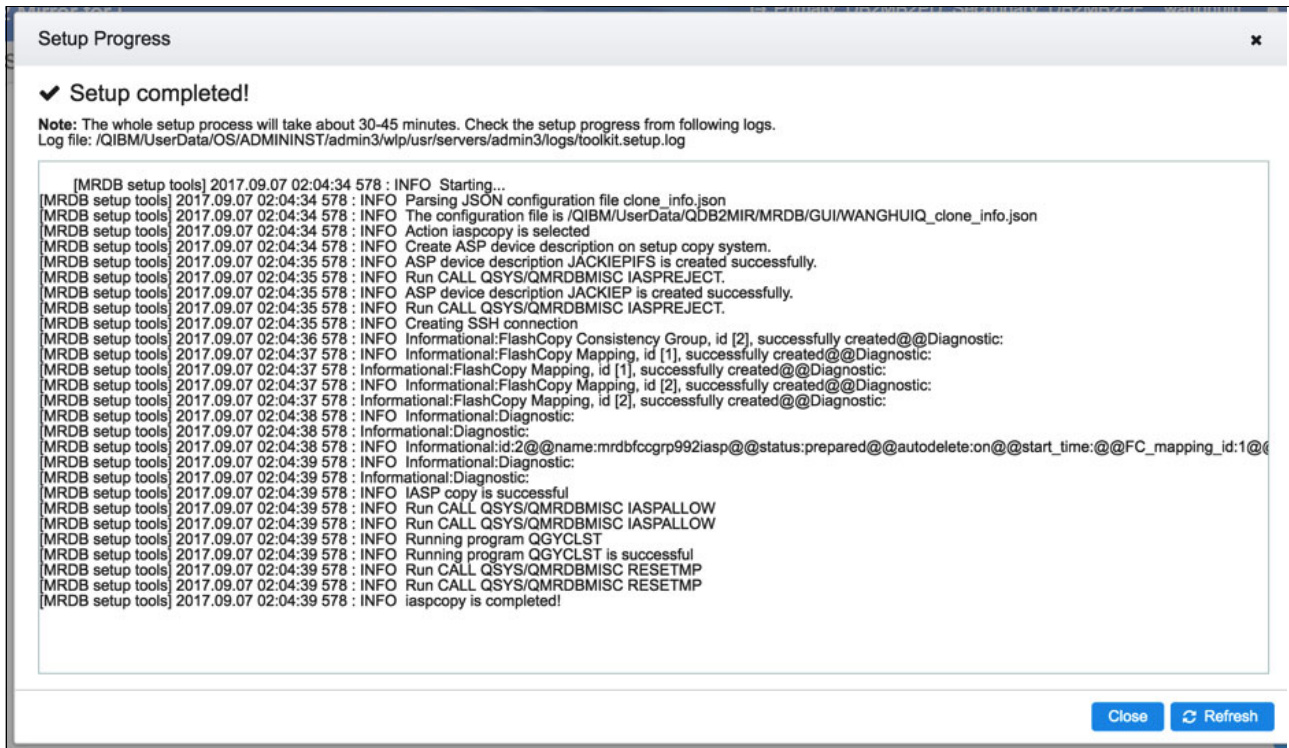


Figure 3-16 Db2 Mirror setup wizard – setup progress status

After the setup process is completed, both Db2 Mirror nodes are actively replicating.





## Management overview

This chapter provides a brief overview about how you can use the Db2 Mirror GUI and Db2 Mirror SQL services to manage a Db2 Mirror environment. It also covers basic monitoring and alerting for a Db2 Mirror environment and illustrates how you can perform a recloning for a Db2 Mirror node.

## 4.1 Db2 Mirror GUI interface

When setting up, managing, and monitoring the Db2 Mirror environment, there are two interfaces that you can use:

- ▶ IBM Db2 Mirror for i GUI
- ▶ SQL services, consisting of SQL procedures, functions, and views

The GUI interface provides the necessary interfaces to set up and manage all aspects of the Db2 Mirror environment.

The Db2 Mirror GUI interface is a web-based interface that runs on IBM i. It can run on either of the two Db2 Mirror nodes or on a separate third IBM i node. The GUI is included with the base Db2 Mirror licensed program (5770DBM \*BASE), which does not require a license key. If the GUI is run on a third IBM i node, no other Db2 Mirror product options are required for the GUI to run and manage Db2 Mirror environments. Although the GUI can run on a Db2 Mirror node for management and monitoring purposes, it must be run from a third node for initial setup or when performing a reconfiguration.

The advantages of running the GUI on a third node are:

- ▶ Required when doing a cold clone.
- ▶ Keeps HTTP traffic on the third node and not on the production nodes.
- ▶ Connects to the active node during planned and unplanned outages.
- ▶ Allows for a single management location if multiple Db2 Mirror pairs are defined.

### 4.1.1 Connecting to the GUI

The default URL to access the Db2 Mirror GUI is:

```
http://hostname:2001/Db2mirror
```

The default URL is routed to the ADMIN HTTP Apache web server running on the GUI node. It is *not* case-sensitive. The default admin Apache server listens for this URL and routes this web request to the ADMIN3 server by using the following URL:

After the ADMIN HTTP Apache web server receives the Db2 Mirror GUI request, it redirects the request to the Db2 Mirror Application server by using the following case-sensitive URL:

```
http://hostname:2006/Db2Mirror
```

The Db2 Mirror GUI by default has TLS/SSL disabled. You can easily configure the Db2 Mirror GUI for HTTPS by using the TLS/SSL wizard from the Web Admin GUI. The Web Admin GUI can be accessed at the following URL:

```
http://hostname:2001/HTTPAdmin
```



From the Web Admin GUI, select the **Manage** tab, select the ADMIN3 server, and click **Configure SSL** from the **Application Server Wizards** menu of the left navigation pane, as shown in Figure 4-1

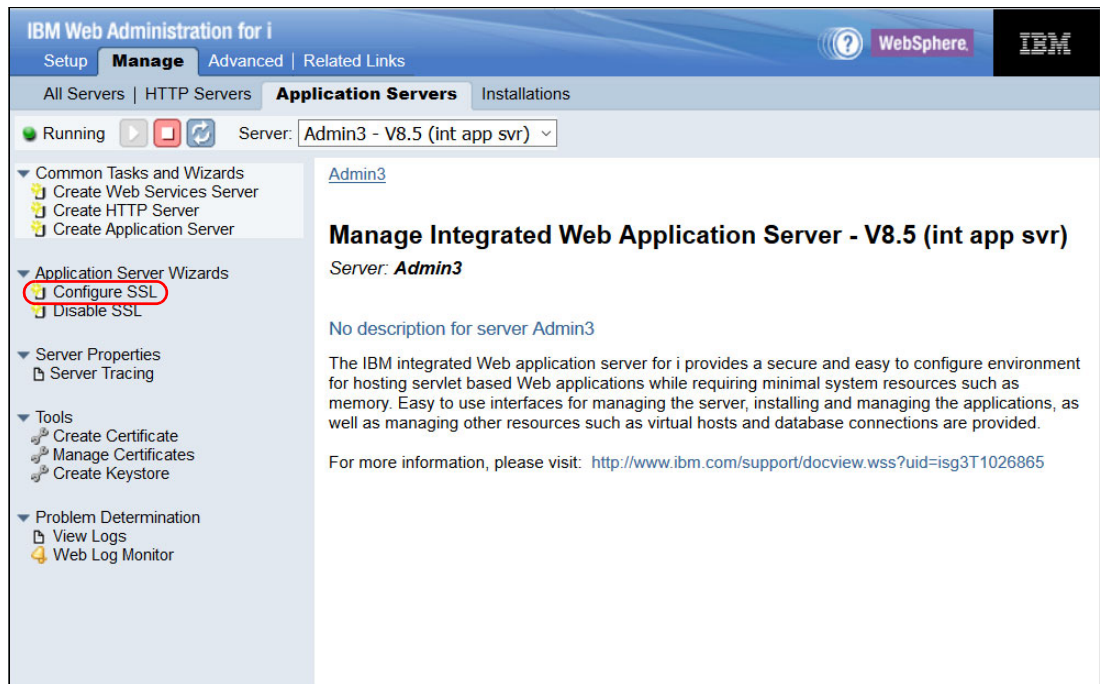


Figure 4-1 IBM i Web Admin GUI: Configure SSL

By using this Web Admin GUI interface, you can reconfigure the ADMIN3 server to use HTTPS by using the certificate of your choice. The interface also can handle all the necessary configuration within the Digital Certificate Manager (DCM).

## 4.1.2 Db2 Mirror GUI overview

The Db2 Mirror GUI interface shows a selected Db2 Mirror pair with its replication status in the middle of the window.

Figure 4-2 shows the Db2 Mirror home page with the Db2 Mirror pair status and icon bar with different options at the left. On the lower left, an icon shows the number of unread Db2 Mirror alert messages that were retrieved from the QSYSOPR message queue from both nodes.



Figure 4-2 Db2 Mirror GUI interface

The same GUI interface can be used to manage another Db2 Mirror pair. Figure 4-3 shows a Db2 Mirror GUI managing many Db2 Mirror pairs.



Figure 4-3 Db2 GUI interface: Managing many Db2 Mirror pairs

You can view information on the Db2 Mirror GUI home page by either by hovering the cursor over a node or by right-clicking a node.

Figure 4-4 shows a screen capture of the Db2 Mirror GUI with all the information shown. This view is a combination of different views because you cannot have all this information shown simultaneously. The two pop-up menus in gray on the left and right of the nodes are each shown by right-clicking each node (primary and secondary). Other pop-up windows, such as about node information or replication state details, open as a result of hovering the cursor over a node object or replication icon.

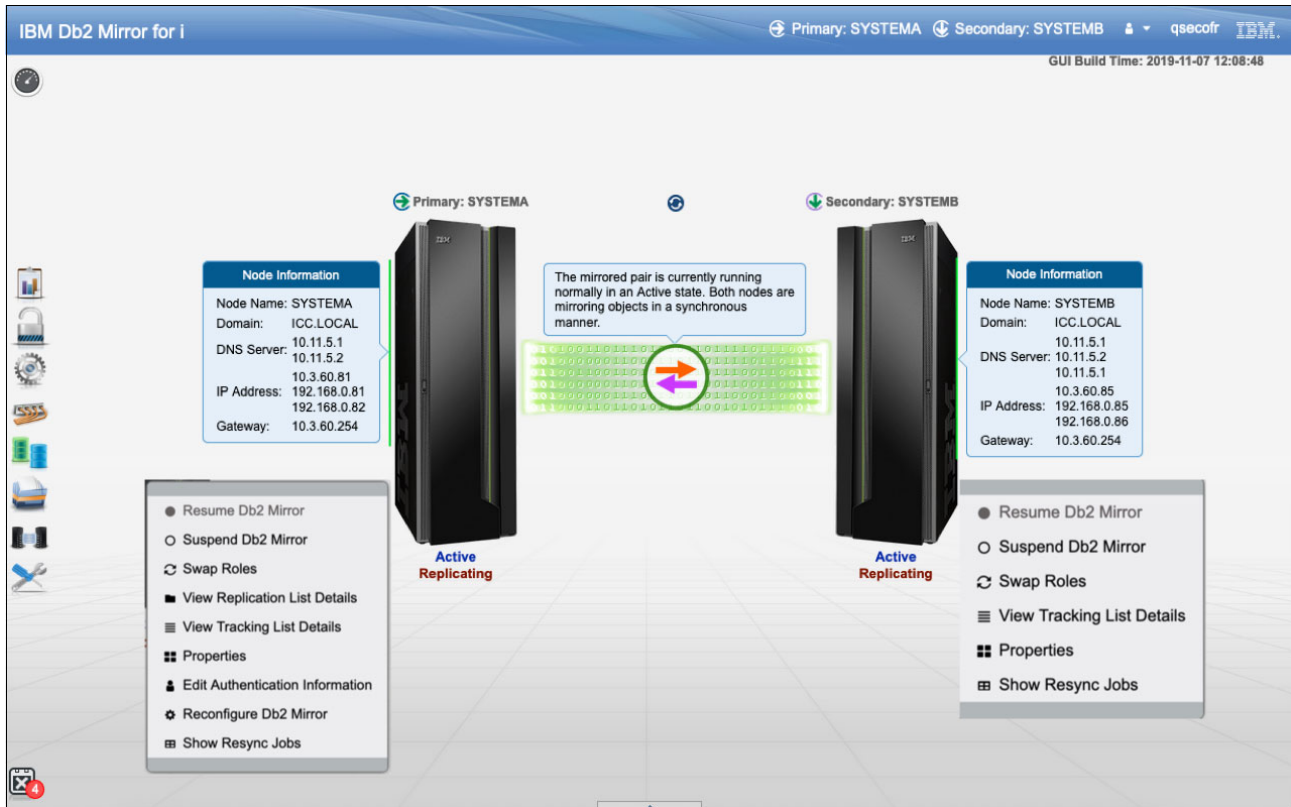


Figure 4-4 Db2 Mirror GUI: Information combined view

The icon bar on the left of the Db2 GUI interface provides a set of options and actions for managing the selected Db2 Mirror pair.

Figure 4-5 shows a view of the left icon bar of the Db2 Mirror GUI interface with an explanation of each icon.

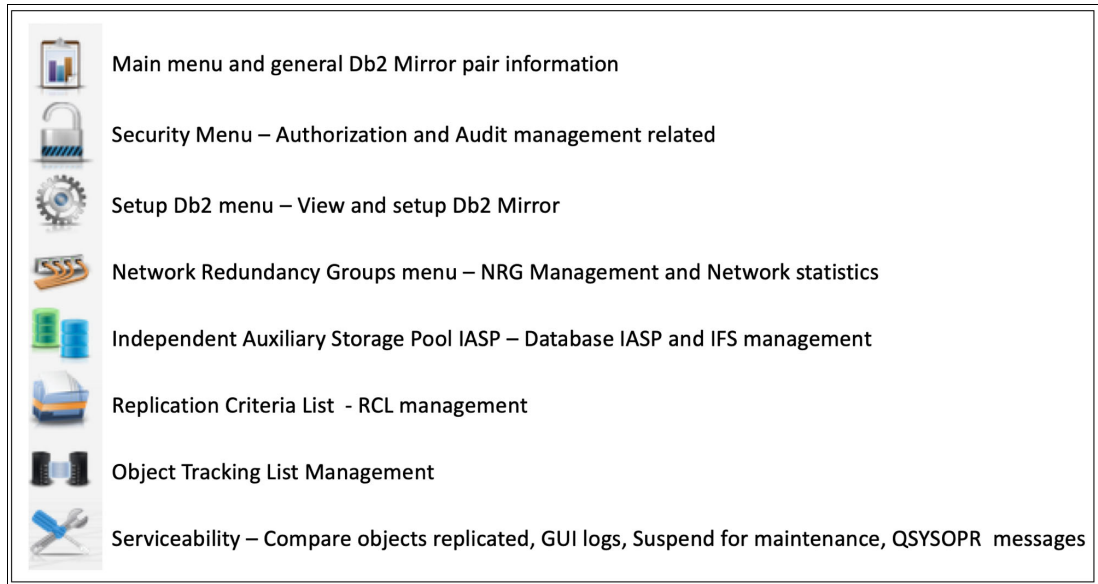


Figure 4-5 Db2 GUI interface: Left bar explanations

## 4.2 SQL services

As mentioned in 1.5.2, “SQL services” on page 17, Db2 Mirror offers various SQL services for the setup, management, and monitoring of a Db2 Mirror environment. This section does not show the full spectrum of available services with all its use cases, but provides brief examples only for demonstrating how to use Db2 Mirror SQL services.

An example of using Db2 Mirror SQL services for managing the replication criteria list (RCL) is to call the `ADD_REPLICATION_CRITERIA` procedure to add a library `ITSOLIB` to the RCL and use the `REPLICATION_CRITERIA_INFO` view to list the active replication criteria rules, as shown in Example 4-1.

*Example 4-1 Using SQL services to manage the RCL*

---

```
CALL QSYS2.ADD_REPLICATION_CRITERIA(INCLUSION_STATE => 'INCLUDE',  
IASP_NAME => '*SYSBAS', LIBRARY_NAME => 'ITSOLIB');
```

```
SELECT * FROM QSYS2.REPLICATION_CRITERIA_INFO WHERE APPLY_STATE = 'ACTIVE'  
ORDER BY IASP_NAME, LIBRARY_NAME, OBJECT_TYPE, OBJECT_NAME;
```

---

Apart from manually running Db2 Mirror SQL services, there also are IBM provided examples that can easily be run and customized by using the IBM i Access Client Solutions. Click **Run SQL Scripts** → **Insert from Examples**, as shown in Figure 4-6.

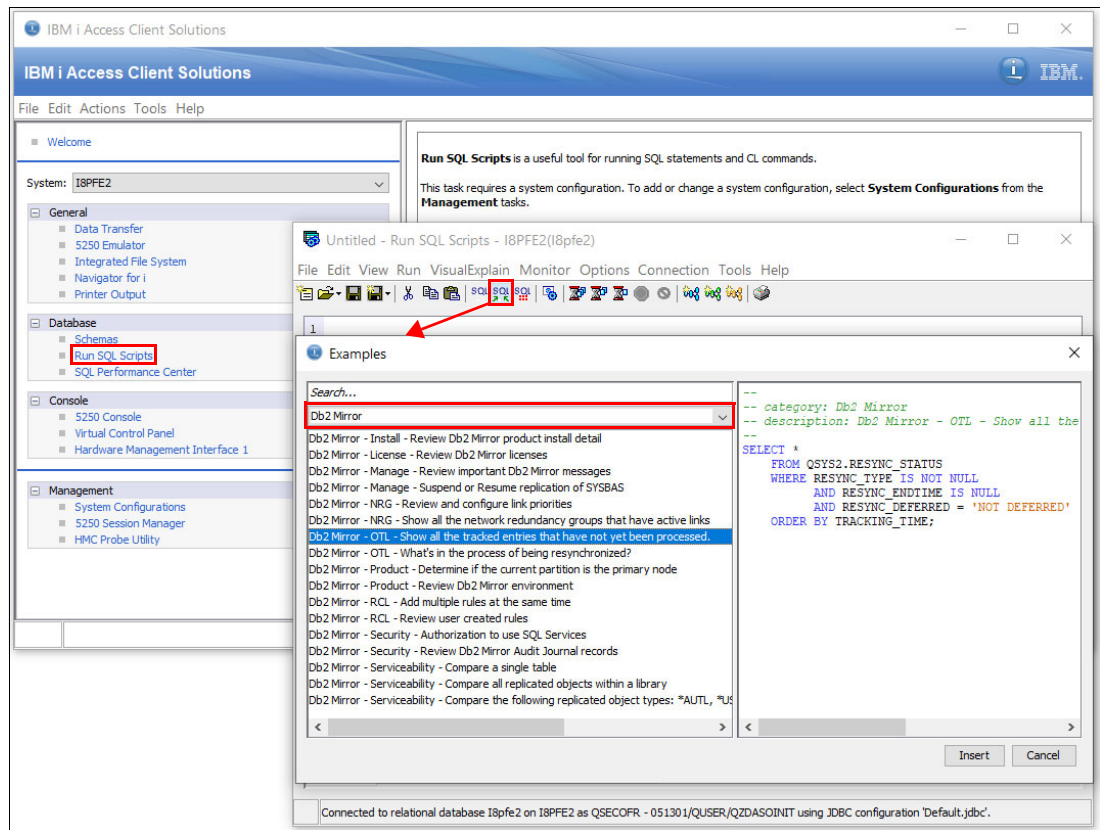


Figure 4-6 Using the IBM i Access Client Solutions Run SQL Scripts Insert from Examples function

## Using Db2 Mirror SQL services with control language program automation

For automation purposes, Db2 Mirror SQL services can be included in control language (CL) programs by using the **RUNSQL** or **RUNSQLSTM** command. Example 4-2 shows the usage of the Db2 Mirror SQL procedure **CHANGE\_MIRROR** for suspending Db2 Mirror replication illustrates how a Db2 Mirror SQL services procedure can be embedded into a CL program or also run interactively from the CL command line.

### Example 4-2 SQL procedure run by a CL command

```
RUNSQL SQL('CALL QSYS2.CHANGE_MIRROR(IASP_NAME => '*SYSBAS', REPLICATION_STATE
=> 'SUSPEND'))')
```

The **RUNSQL** CL command does not support the SQL **SELECT** and **VALUES** statements. For CL program automation that requires Db2 Mirror SQL services views and functions, use the **CREATE TABLE AS** and **INSERT INTO** statements with subselect as a replacement with data from the result table that is retrieved by using the **RCVF** CL command.

## 4.3 Monitoring and alerting

For monitoring and alerting of a Db2 Mirror environment, you can use different sources of information, such as the Db2 Mirror management GUI, SQL services, and QSYSOPR messages.

The GUI provides monitoring and alerting functions for resources, such as the following ones:

- ▶ Node role and replication states
- ▶ Network redundancy group (NRG) link states and usage statistics
- ▶ Remote Direct Memory Access (RDMA) link and connections states and usage statistics
- ▶ Object tracking list (OTL) entries (historic and current entries, including replication progress)
- ▶ OTL conflicts
- ▶ QSYSOPR messages, history log, flight recorder, and GUI log files viewing and export

Figure 4-7 shows an example of using the Db2 Mirror GUI for monitoring and alerting, with its home page showing the selected Db2 Mirror pair's replication state as suspended after a communication loss between the two nodes, and the corresponding alerts.



Figure 4-7 Db2 Mirror GUI home page with replication states

Alternatively, you can view the replication states by using the Db2 Mirror SQL services **MIRROR\_INFO** view, as shown with its invocation and output from an IBM i 5250 session in Example 4-3.

Example 4-3 Db2 Mirror SQL services **MIRROR\_INFO** view

```
SELECT IASP_NAME, REPLICATION_STATE, REPLICATION_DETAIL, REPLICATION_DETAIL_INFO,
REPLICATION_DETAIL_INFO_TEXT, PRIMARY_NODE, SECONDARY_NODE FROM QSYS2.MIRROR_INFO;
```

```

                                Display Data
                                Data width . . . . . :    1102
Position to line . . . . .      Shift to column . . . . .
.....1.....2.....3.....4.....5.....6.....7.....8.....9.....10
IASP      Replication  Replication  Replication  Replication  Replication
Name      State          Detail        Detail        Detail        Detail
                                Info          Info          Info          Info
*SYSBAS   TRACKING      SUSPENDED    220          Suspended due to NRG down on this node
IASP_IFS  -              -            -            -
*****   End of data *****

```

After clicking the highlighted alerts icon at the lower left of the Db2 Mirror GUI home page, as shown in Figure 4-7 on page 61, the user can see the Alerts dialog box with the corresponding QSYSOPR messages that were posted by Db2 Mirror for the communication issue, as shown in Figure 4-8.

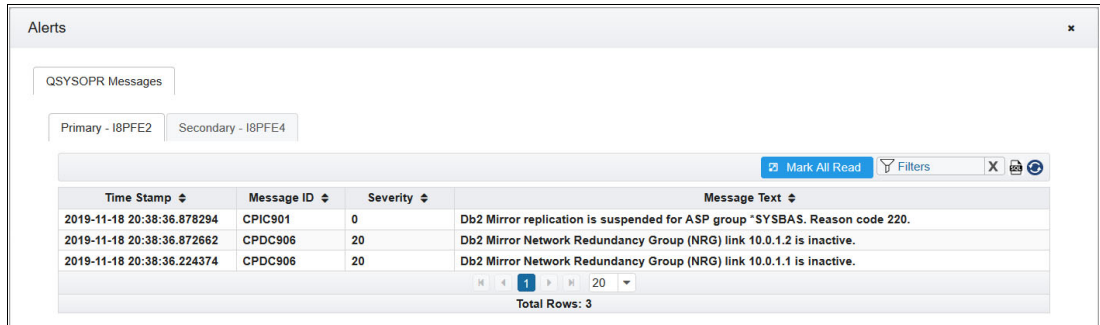


Figure 4-8 Db2 Mirror GUI alerts

The Db2 Mirror product uses the following *state change* QSYSOPR messages:

- CPDC905** Db2 Mirror NRG link <IP address> is active.
- CPDC906** NRG link <IP address> is inactive.
- CPIC901** Db2 Mirror replication is suspended for auxiliary storage pool (ASP) group IASP33P. Reason code <reason code>.
- CPIC902** Db2 Mirror replication is suspended for ASP group <iASP name or \*SYSBAS> due to an error. Reason code <reason code>.
- CPIC903** Db2 Mirror replication is suspended for maintenance operations.
- CPIC904** Db2 Mirror replication is active for ASP group <iASP name or \*SYSBAS>.

Db2 Mirror uses the following *product failure* QSYSOPR messages:

- CPD3E43** DRDA/DDM Db2 Mirror server error occurred with reason code <reason code>.
- CPF32CD** Db2 Mirror resynchronization failed for job <job name or \*ALL>.



The user can also look at further details like object changes that are affected by the suspended replication and tracked on the primary node by displaying the OTL details, as shown in Figure 4-9.

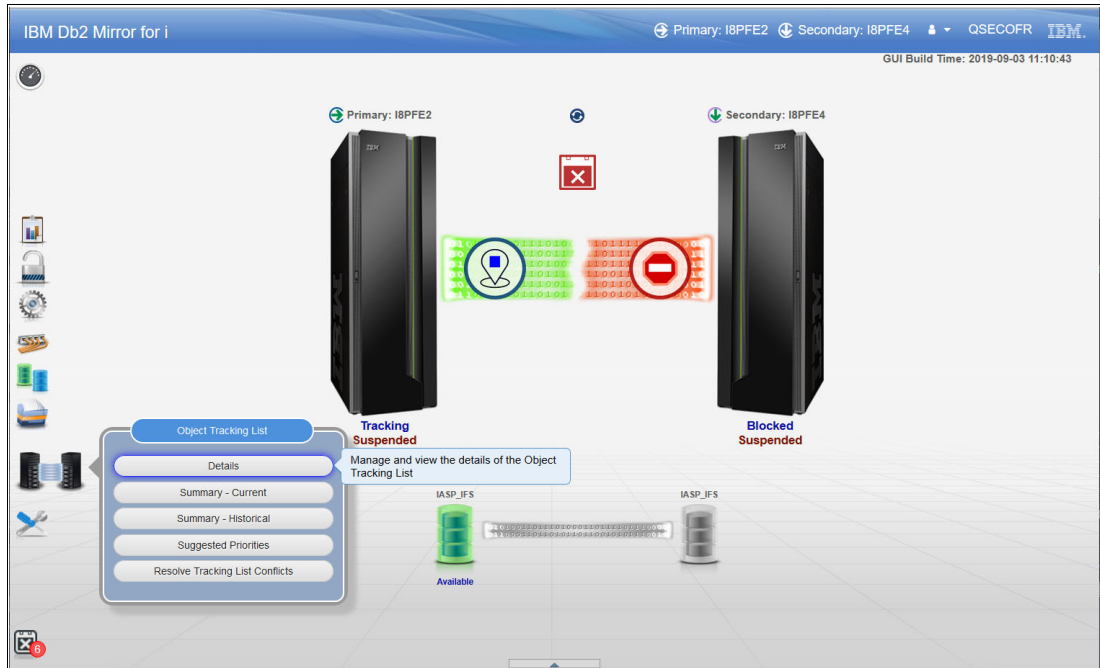


Figure 4-9 Db2 Mirror GUI object tracking list details selection

An example of OTL details with some database and user profile object tracking entries is shown in Figure 4-10.

All Tracking - Details GUI Build Time: 2019-09-03 11:10:43

Primary - I8PFE2    Secondary - I8PFE4

Current Tracking Entries    [Summary](#) → [Details](#)    Filters X

Status	Resync Group	Add Time	Sync Start	Library Name	Object Type	Object Name	Member Name	Resync Type	Tracking Group	Replication Reason
○	6	2019-11-18 23:32:03.265879569824		ITSOLIB	*FILE	ITSODATA3	ITSODATA3	DB I/O	65	UNABLE TO MIRROR
○	5	2019-11-18 23:30:51.810964000000		ITSOLIB	*FILE	ITSODATA3	ITSODATA3	DB DDL	65	UNABLE TO MIRROR
○	5	2019-11-18 23:30:51.591864000000		ITSOLIB	*FILE	ITSODATA3		DB DDL	65	UNABLE TO MIRROR
○	1	2019-11-18 23:28:20.884846018554		QSYS	*USRPRF	QSECOFR		PROGRAM CALL	65	UNABLE TO MIRROR
○	1	2019-11-18 23:28:20.883125342773		QSYS	*USRPRF	QSECOFR		PROGRAM CALL	65	UNABLE TO MIRROR
○	5	2019-11-18 23:28:00.867097000000		ITSOLIB	*FILE	ITSODATA2	ITSODATA2	DB DDL	65	UNABLE TO MIRROR
○	5	2019-11-18 23:28:00.815267000000		ITSOLIB	*FILE	ITSODATA2		DB DDL	65	UNABLE TO MIRROR

Showing 7 of 7

Figure 4-10 Db2 Mirror GUI object tracking list entries

A similar view about the OTL details can be obtained by using IBM i Access Client Solutions menu. Select **Run SQL Scripts...** → **Insert from Examples** → **Db2 Mirror - OTL - Show all the tracked entries that have not yet been processed**, as shown in Figure 4-11.

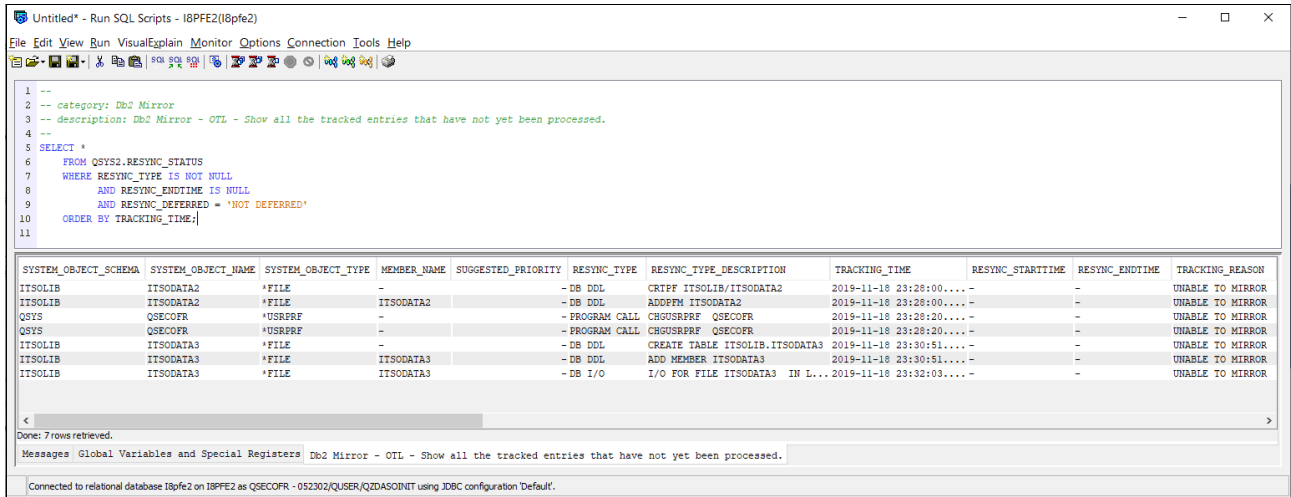


Figure 4-11 IBM i Access Client Solutions Insert from Examples OTL entries not processed

You can monitor the progress of an ongoing Db2 Mirror synchronization from the Db2 Mirror GUI by selecting **Object Tracking List** → **Summary – Current**, as shown in Figure 4-12.

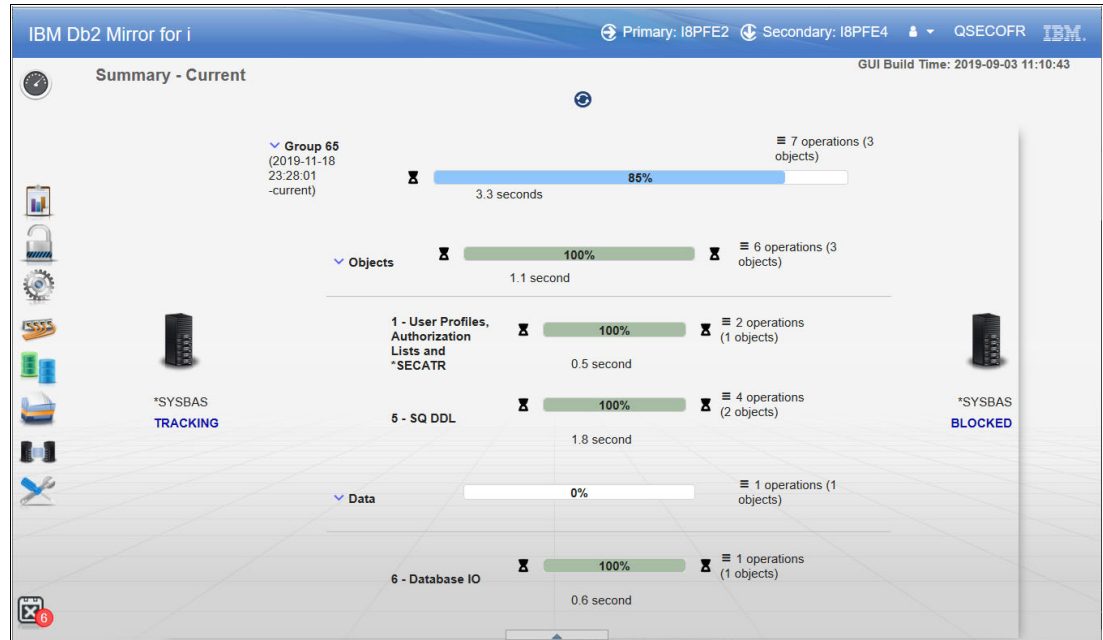


Figure 4-12 Db2 Mirror GUI object tracking list summary

## 4.4 Recloning

This section provides a brief overview about reconfiguring Db2 Mirror, which includes a recloning of the Db2 Mirror pair.

Recloning that involves a full resynchronization of the Db2 Mirror pair is required if the user changes their mind about the default inclusion state setting, as described in 2.4.1, “Default inclusion state” on page 26, or as a recovery after some unplanned outages for which no IBM i main store memory dump was taken and main storage was not preserved.

A Db2 Mirror recloning process can be started from the Db2 Mirror GUI, and it involves the following steps:

1. Ensure that the system that you use as the setup source node for the recloning is the current primary node, and its replication state is set to suspended/tracking.
2. Start the reconfiguration process from the Db2 Mirror GUI by selecting **Setup** → **Reconfigure Db2 Mirror**, as shown in Figure 4-13.

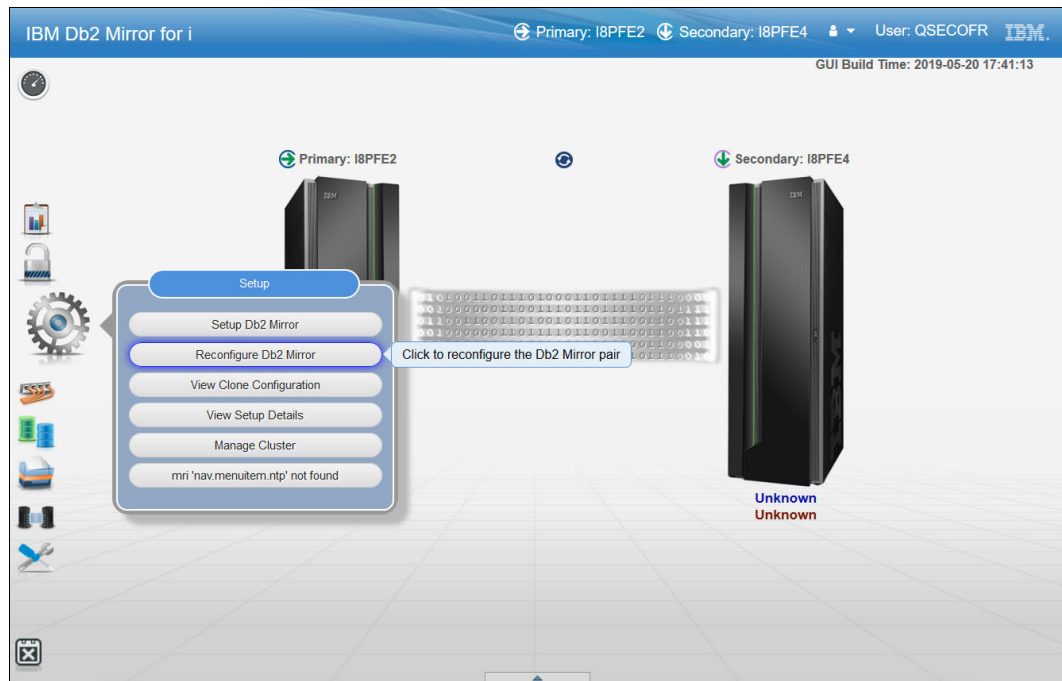


Figure 4-13 Db2 Mirror GUI: Reconfigure Db2 Mirror

3. Select a reconfiguration method, as shown in the Reconfigure Db2 Mirror dialog box in Figure 4-14.

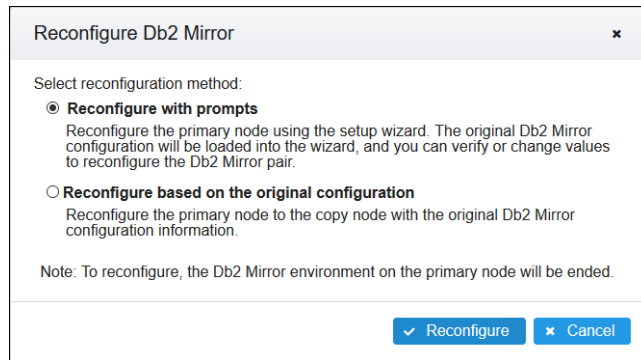


Figure 4-14 Db2 Mirror GUI: Reconfigure Db2 Mirror dialog box

Selecting **Reconfigure with prompts** loads the original Db2 Mirror configuration into the wizard and allows the user to verify or modify the configuration settings, as shown in Figure 4-15. Selecting **Reconfigure based on the original configuration** starts the recloning is immediately after clicking **Reconfigure**.

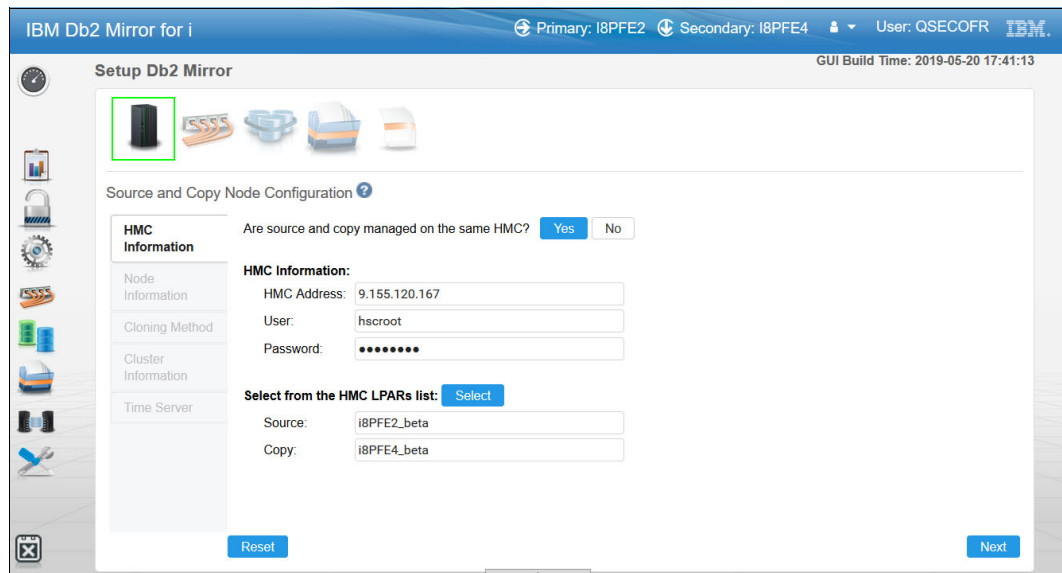


Figure 4-15 Db2 Mirror GUI: Setup wizard

At the end of the setup wizard, the user can still choose to reset the Db2 Mirror reconfiguration data to the defaults, which clears any custom NRGs and replication rules, as shown in Figure 4-16.

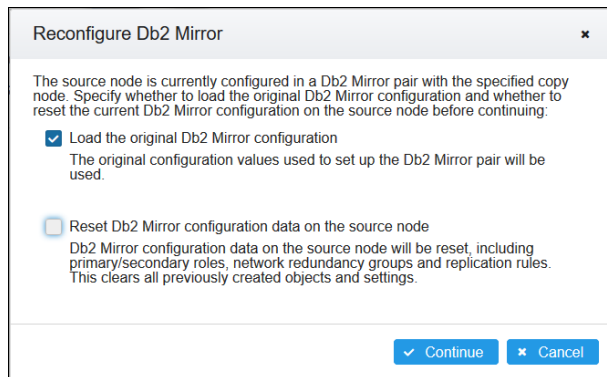


Figure 4-16 Db2 Mirror GUI: Reconfigure config data dialog box

- In the last step of the reconfiguration process, the recloning starts, as shown in Figure 4-17.

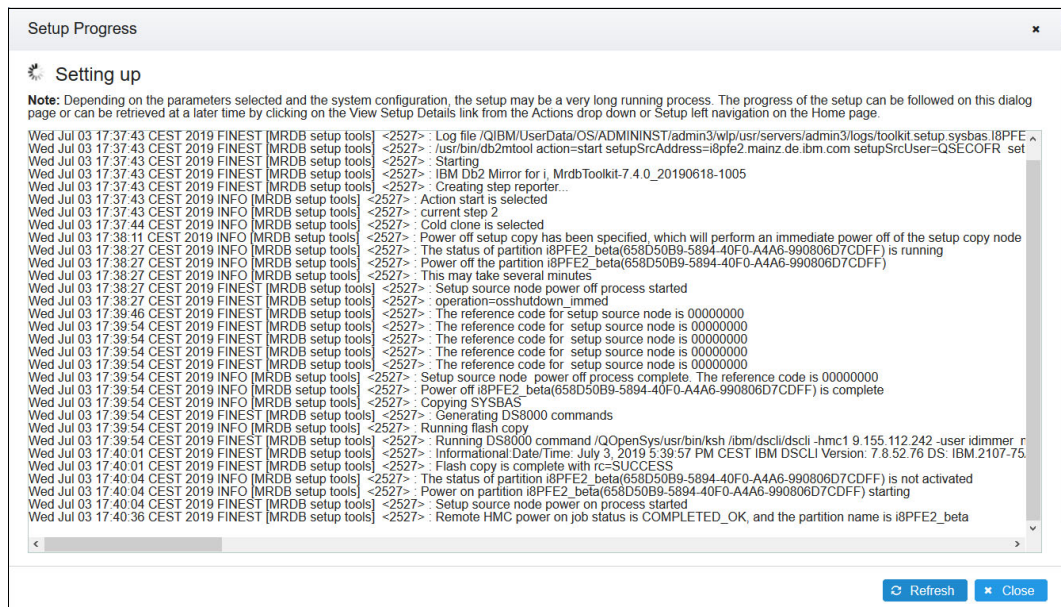


Figure 4-17 Db2 Mirror GUI: Recloning progress





## Use cases

This chapter provides an overview about how IBM Db2 Mirror for i provides continuous availability for different planned and unplanned outages.

## 5.1 Planned outages

For a planned outage scenario such as performing hardware, operating system (OS), or application maintenance on the Db2 Mirror secondary node, Db2 Mirror replication can be suspended, which causes the primary node to track changes and the secondary node to block changes (if the planned maintenance affects the primary node, a Db2 Mirror role swap can be performed). After the planned maintenance on the secondary node, Db2 Mirror can start synchronization again between both nodes.

### Rolling update scenario

This maintenance approach can also be used as follows for a rolling update scenario like planned program temporary fix (PTF) installations on both nodes:

1. Prepare application workloads for a suspension of Db2 Mirror.

If the Db2 Mirror pair is used in an *active-active* operation model, the user might want to ensure that, either by an active-active application design that provides automated workload failover or by manual actions, the workload is repositioned from the secondary node to the primary node before proceeding with suspending Db2 Mirror in step 2.

If the Db2 Mirror pair is used in an *active-passive* operation model, the user might want to end their reporting or analysis-based type of workload on the secondary node before proceeding with suspending Db2 Mirror in step 2. The reasoning is that even though reads continue to succeed after the secondary node reaches the suspended/blocked state, the data might not reflect the state of the business.

2. Suspend Db2 Mirror.
3. Perform maintenance actions PTF installations and initial program load (IPL) on the secondary node.
4. Resume Db2 Mirror.
5. After the resync is complete, swap the Db2 Mirror roles.
6. Repeat steps 1 - 5 to maintenance on the new secondary (original primary) node.

### Live Partition Mobility

Another option to prepare a Db2 Mirror pair for planned maintenance, especially hardware maintenance or server migration, without having to tolerate a longer duration of suspended Db2 Mirror replication is to use PowerVM Live Partition Mobility (LPM) for migrating a Db2 Mirror node to another IBM Power Systems server. For more information about the requirements for using LPM with Db2 Mirror, see 1.4.5, “Db2 Mirror and Live Partition Mobility” on page 14.

The procedure to use for live migration of either the primary or secondary Db2 Mirror node is as follows:

1. Single root I/O virtualization (SR-IOV) logical Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) ports must be varied off and removed by using a dynamic logical partition (DLPAR) before LPM so that the primary Db2 Mirror side tracks and the secondary side blocks changes.
2. Perform the LPM operation for the Db2 Mirror node.
3. Use DLPAR to add the SR-IOV logical port back to the migrated partition.
4. Resume Db2 mirroring.



## Regular maintenance by using suspend and resume

Suspending Db2 Mirror for planned maintenance can be accomplished as shown in Figure 5-1. Right-clicking a node from the Db2 Mirror GUI home page and select **Suspend Db2 Mirror** to stop replication.



Figure 5-1 Db2 Mirror GUI suspend Db2 Mirror

Alternatively, you can suspend Db2 Mirror replication by using the Db2 Mirror SQL services procedure `CHANGE_MIRROR`, as shown in Example 5-1. The `SUSPEND` action requires `*SYSBAS` as the `IASP_NAME` and suspends both `SYSBAS` and any database IASP (DB IASP) replication.

*Example 5-1 Using Db2 Mirror SQL services to suspend replication*

```
CALL QSYS2.CHANGE_MIRROR(IASP_NAME => '*SYSBAS', REPLICATION_STATE => 'SUSPEND');
```

After the planned maintenance is complete, you can resume the suspended Db2 Mirror replication from the Db2 Mirror GUI home page by right-clicking a node and selecting **Resume Db2 Mirror**, as shown in Figure 5-2.



Figure 5-2 Db2 Mirror GUI: Resume Db2 Mirror

Alternatively, you can resume a suspended Db2 Mirror pair by using the SQL services procedure `CHANGE_MIRROR`, as shown in Example 5-2.

*Example 5-2 Using Db2 Mirror SQL services to resume replication*

```
CALL QSYS2.CHANGE_MIRROR(IASP_NAME => '*SYSBAS', REPLICATION_STATE => 'RESUME');
```

After Db2 Mirror resumes, data resynchronization starts between the primary and secondary nodes, as shown in Figure 5-3, until both nodes are in the active-replicating state again.



Figure 5-3 Db2 Mirror GUI synchronizing

## 5.2 Unplanned outages

This section provides a brief overview about how unplanned outages are handled by Db2 Mirror, and which recovery actions are possibly required by the user.

## 5.2.1 Secondary node outage

If the Db2 Mirror secondary node becomes unavailable (for example, an IPL occurs, it goes into a restricted state, or there is a main storage dump as illustrated by Figure 5-4), the Db2 Mirror primary node begins tracking replicated object changes, as shown in Figure 5-5, and user applications (if all their dependent data was included in Db2 Mirror replication before) continue to run.




Figure 5-4 Db2 Mirror secondary node outage



Figure 5-5 Db2 Mirror states after secondary node outage

After the secondary node becomes available again, it is in a blocked state, as shown in Figure 5-6, and it does not allow changes to replicated objects until the two nodes resume mirroring. Db2 Mirror by default automatically resumes mirroring after the secondary node becomes available again.

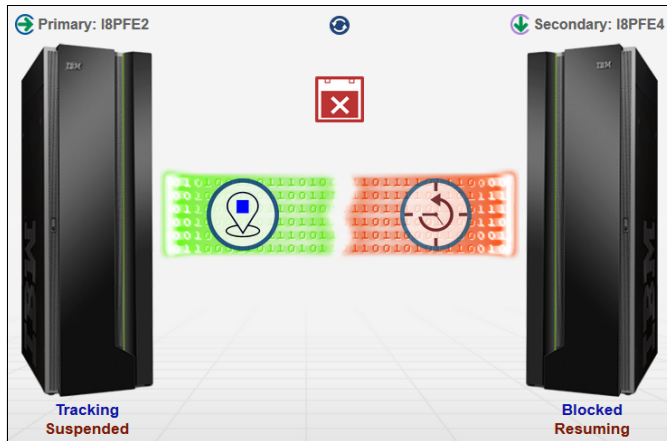


Figure 5-6 Db2 Mirror states after the secondary node becomes available again

## 5.2.2 Primary node outage

This section describes the scenario of a primary node outage in a Db2 Mirror environment, as shown in Figure 5-7.

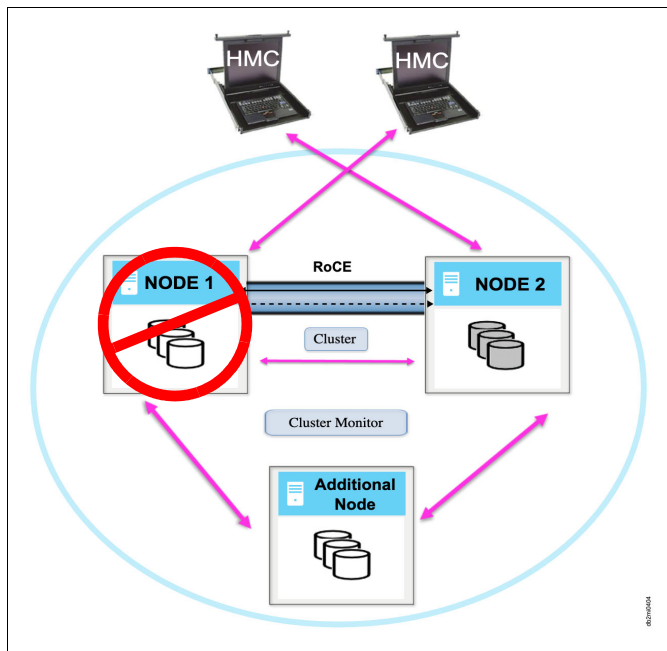


Figure 5-7 Db2 Mirror primary node outage

The handling of a primary node outage by Db2 Mirror depends on the condition that led to the primary node not being available. These conditions are described in the following sections.

## Primary node becomes unavailable due to an IPL or going into a restricted state

The secondary node remains blocked and the primary node tracks while it is in a restricted state or until the IPL completes. An IFS IASP is failed over by PowerHA to the secondary node if it is not prevented by a user-defined PowerHA policy and the cluster resource group (CRG) has not ended before. Figure 5-8 shows the states of the nodes and the IASP for this scenario.

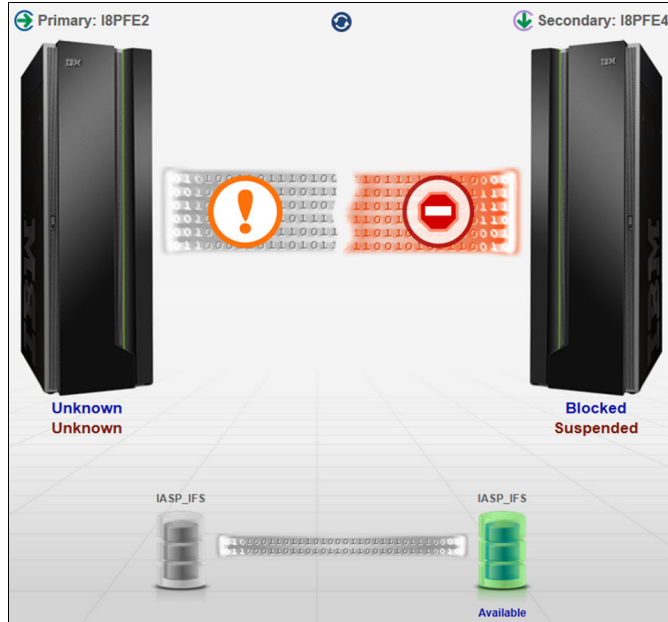


Figure 5-8 Db2 Mirror states after a primary node becomes unavailable

The user can then still decide whether they want to wait for the primary node to become available again. If they prefer to continue working with their database applications, they may perform a Db2 Mirror role swap, as shown in Figure 5-9.

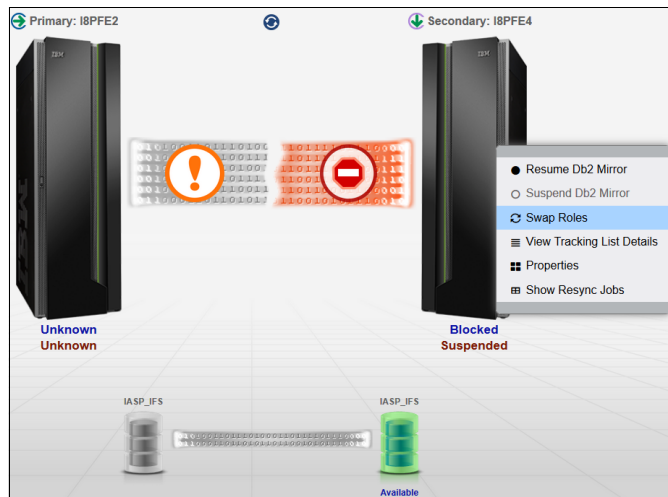


Figure 5-9 Db2 Mirror role swap after the primary node becomes unavailable

The role swap forces the blocked secondary node to become a suspended and tracking primary node. Db2 Mirror requests the user to confirm this action, as shown in Figure 5-10. The blocked secondary node has the potential for object collisions because some changes, such as security-related objects as user profiles and authorization lists or for system values, are allowed in the blocked state.

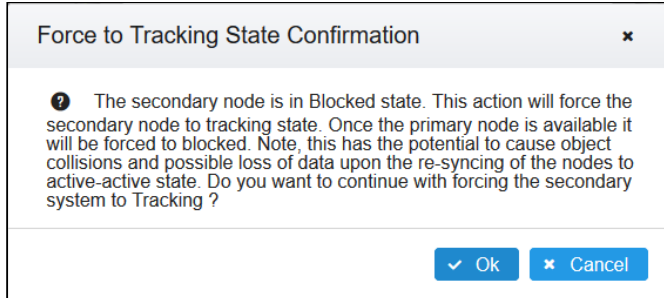


Figure 5-10 Db2 Mirror GUI: Role swap confirmation

After the role swap, the blocked secondary node becomes a tracking primary node, as shown in Figure 5-11. (The Db2 Mirror GUI shows the primary node always on the left of the window.)

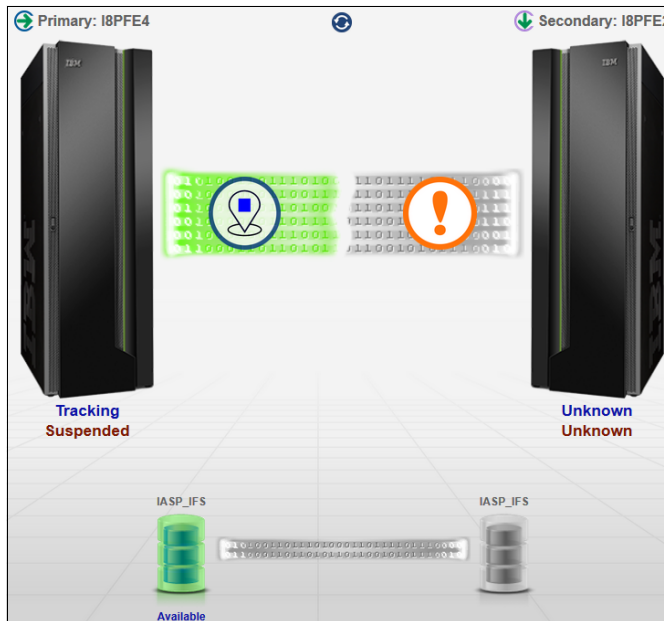


Figure 5-11 Db2 Mirror states after a role swap after the primary becomes unavailable

### Primary node fails with a crash or main store dump

If the secondary Db2 Mirror node can connect to the Hardware Management Console (HMC) and determine that the primary node failed, the secondary node takes over as the primary and begins tracking.

Otherwise, if the secondary cannot detect the failure, it remains blocked, as shown in Figure 5-12.

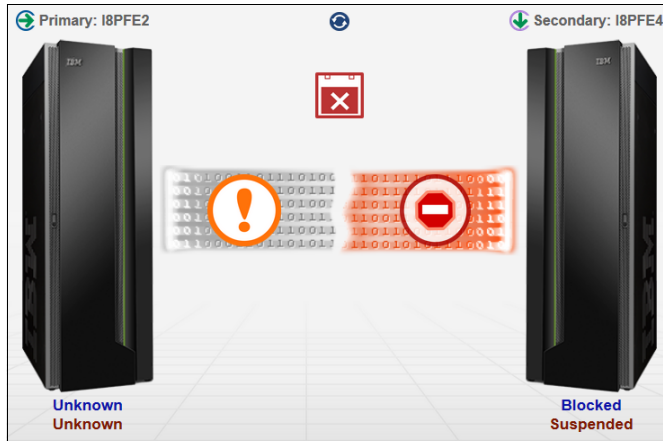


Figure 5-12 Db2 Mirror states after undetected primary node failure

If the user does not want to wait for the primary to become available again to continue to work with their database applications, they may perform a Db2 Mirror node role swap to force the secondary to become the primary, as described in **“Primary node becomes unavailable due to an IPL or going into a restricted state”** on page 76.

### 5.2.3 Node communication loss

If communication is lost between the two Db2 Mirror nodes over the RoCE network, as illustrated by Figure 5-13, the primary tracks replicated objects, and the secondary blocks changes to replicated objects, as shown in Figure 5-14 on page 79, until mirroring can resume. Db2 Mirror by default automatically resumes its mirroring after the communication between the two nodes is reestablished.

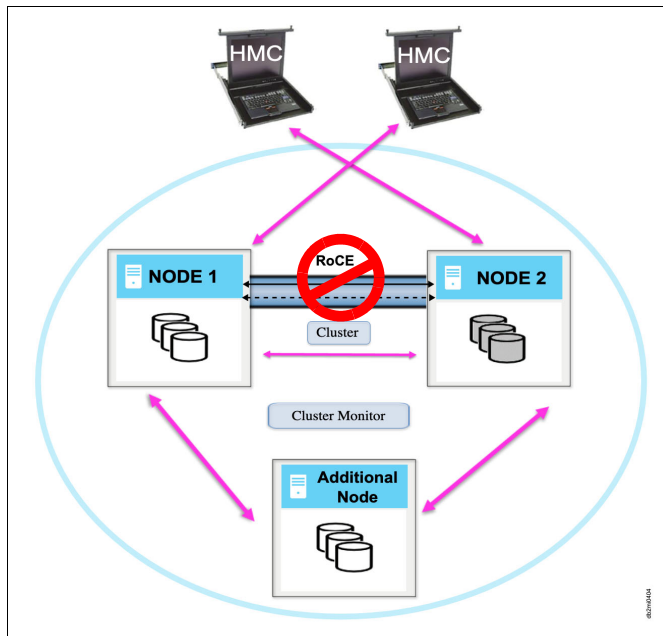


Figure 5-13 Db2 Mirror communication loss





Figure 5-14 Db2 Mirror states after communication loss



# Related publications

The publications that are listed in this section are considered suitable for a more detailed description of the topics that are covered in this paper.

## IBM Redbooks

The following IBM Redbooks publications provide more information about the topics in this document. Some publications that are referenced in this list might be available in softcopy only.

- ▶ *IBM i 7.2 Technical Overview with Technology Refresh Updates*, SG24-8249
- ▶ *Modernizing IBM i Applications from the Database up to the User Interface and Everything in Between*, SG24-8185
- ▶ *PowerHA SystemMirror for IBM i Cookbook*, SG24-7994

You can search for, view, download, or order these documents and other Redbooks, Redpapers, web docs, drafts, and additional materials at the following website:

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

## Online resources

These websites are also relevant as further information sources:

- ▶ Application considerations  
[https://www.ibm.com/support/knowledgecenter/ssw\\_ibm\\_i\\_74/db2mi/db2mappdev.htm](https://www.ibm.com/support/knowledgecenter/ssw_ibm_i_74/db2mi/db2mappdev.htm)
- ▶ Collecting all required information for the setup by using the Db2 Mirror checklist  
[https://www.ibm.com/support/knowledgecenter/ssw\\_ibm\\_i\\_74/db2mi/db2mplancheklist.htm](https://www.ibm.com/support/knowledgecenter/ssw_ibm_i_74/db2mi/db2mplancheklist.htm)
- ▶ Command line-based setup  
[https://www.ibm.com/support/knowledgecenter/ssw\\_ibm\\_i\\_74/db2mi/db2mcmdsetup.htm](https://www.ibm.com/support/knowledgecenter/ssw_ibm_i_74/db2mi/db2mcmdsetup.htm)
- ▶ Comprehensive list of the requirements for Db2 Mirror  
[https://www.ibm.com/support/knowledgecenter/ssw\\_ibm\\_i\\_74/db2mi/db2mplanningandsetup.htm](https://www.ibm.com/support/knowledgecenter/ssw_ibm_i_74/db2mi/db2mplanningandsetup.htm)
- ▶ Db2 Mirror SQL services  
[https://www.ibm.com/support/knowledgecenter/ssw\\_ibm\\_i\\_74/db2mi/db2mservices.htm](https://www.ibm.com/support/knowledgecenter/ssw_ibm_i_74/db2mi/db2mservices.htm)
- ▶ Db2 Mirror tool  
[https://www.ibm.com/support/knowledgecenter/ssw\\_ibm\\_i\\_74/db2mi/db2mdb2mtoolreference.htm](https://www.ibm.com/support/knowledgecenter/ssw_ibm_i_74/db2mi/db2mdb2mtoolreference.htm)
- ▶ Db2 Mirror setup wizard  
[https://www.ibm.com/support/knowledgecenter/ssw\\_ibm\\_i\\_74/db2mi/db2mguisetup.htm](https://www.ibm.com/support/knowledgecenter/ssw_ibm_i_74/db2mi/db2mguisetup.htm)

- ▶ DS8000 setup  
[https://www.ibm.com/support/knowledgecenter/ssw\\_ibm\\_i\\_74/db2mi/db2mguiconfigstorageds8k.htm](https://www.ibm.com/support/knowledgecenter/ssw_ibm_i_74/db2mi/db2mguiconfigstorageds8k.htm)
- ▶ IBM Systems Lab Services  
<https://www.ibm.com/it-infrastructure/services/lab-services>
- ▶ Replication for specific object types  
[https://www.ibm.com/support/knowledgecenter/ssw\\_ibm\\_i\\_74/db2mi/db2mobjreplication.htm](https://www.ibm.com/support/knowledgecenter/ssw_ibm_i_74/db2mi/db2mobjreplication.htm)
- ▶ Setting up and configuring an SSH key file  
[https://www.ibm.com/support/knowledgecenter/ssw\\_ibm\\_i\\_74/db2mi/db2msshkeyfile.htm](https://www.ibm.com/support/knowledgecenter/ssw_ibm_i_74/db2mi/db2msshkeyfile.htm)

## Help from IBM

IBM Support and downloads

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

IBM Global Services

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





REDP-5575-00

ISBN 0738458244

Printed in U.S.A.

Get connected

