

SESSION Db2 Mirror for i : Disponibilité continue des applications

Bertrand Guibert Spécialiste Power Systems & IBM i bertrand_guibert@fr.ibm.com

Un évènement organisé par :



Avec:

FRESCHE

tesiliane



IBM Db2 Mirror for i



IBM Db2 Mirror for i: Enables Continuous Availability

- High speed synchronous replication of Db2 for i (Data Center Solution)
- Access Db2 objects from either LPAR

Application Availability Enablement

- Two Nodes read and write to the same DB Files
- Enables quickly moving all work to one node, for planned maintenance or node failure

Enables Business Continuity for Disruptive System Upgrades

- Nodes can be at different OS levels
- Nodes can be on different Power Hardware Generations
- Rolling upgrades for no downtime
- Roll a node back a release with minimal impact if Active/Active applications are deployed

Requires POWER8 or later and IBM i 7.4 New IBM i LPP 5770DBM



Recovery Time Objective

Recovery Point Objective

Near 0

High Availability topology classification & positioning

	Active/Active	Active/Passive	Active/inactive		
	M partition Active	VM partition Active	Restart		
Technology	Active/Active Clustering	Active/Passive Clustering	Active/Inactive		
Definition	Application level clustering; applications in the cluster have simultaneous access to the production data therefore no app restart upon an app node outage. Certain types enable read-only access from secondary nodes	OS level clustering ; one OS in the cluster has access to the production data, multiple active OS instances on all nodes in the cluster. Application is restarted on a secondary node upon outage of a production node.	VM level clustering, One VM in a cluster pair has access to the data, one logical OS, one or two physical copies. OS and applications must be restarted on a secondary node upon a primary node outage event. LPM enables the VM to be moved non-disruptively for a planned outage event.		
Outage Types	SW,HW,HA, planned, unplanned RTO 0, limited distance	SW,HW,HA,DR, planned, unplanned, RTO>0, multi-site	HW,HA,DR, planned, unplanned, RTO>0, multi-site		
OS integration	Inside the OS	Inside the OS	OS agnostic		
RPO	Sync mode only	Sync/Async	Sync/Async		
RTO	0	Fast (minutes)	Fast Enough (VM Reboot)		
Licensing*	N+N licensing	N+1 licensing	N+0 licensing		
Industry Examples	Oracle RAC, Db2 Mirror , pureScale	PowerHA, Redhat HA, Linux HA	VMware, VMR HA, LPM,		

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

• illustrations represent two-node shared-storage configurations for conceptual simplicity. There are many other topologies and data resiliency combinations

Architecture

- Remote Direct Memory Access (RDMA)
- RDMA over Converged Ethernet (RoCE)
- Maximum distance 100m or 200m



Db2 Mirror Statement of Direction

2020 iSight

IBM Db2 Mirror for i planning insights:

IBM plans to enhance IBM Db2 Mirror for i (5770-DBM) to extend the physical distance allowed between two nodes, allowing the two nodes to be located up to 10 kilometers apart.

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.















Db2 Mirror – Active Passive, Green Screen Application



Db2 Mirror – Active Active, Web Clients with Load Balancer



Db2 Mirror – Database Supported Objects

Database replication eligible objects Native:

- Database Physical & Logical File **SQL:**
- Alias
- Function
- Global Variable
- Index
- Procedure
- Schema
- Sequence

- SQL PackageTable
- Table
- Trigger
- User Defined Type
- View
- XML Schema
 - Repository

Included with File support:

- Row Permission
- Column Mask
- Temporal Table
- Constraint
- Etc...

DDS / Record Level Access SQL / Set Based Access



Db2 Mirror – Other Supported Objects

- Other Objects
 - User profiles
 - Authority
 - Ownership
 - Security
 - PGM/SRVPGM
 - Data Areas
 - Data Queues (DDL Only)
 - SYSVALs
 - ENVARs
 - LIB
 - JOBD
 - Journals
 - Files (also has DDL Only option)
- Special Handling
 - OUTQ / Spool
 - Job Queue

Objects can be in either **SYSBAS** or **IASP**s



Db2 Mirror – Security Implementation

Security changes are always replicated:

- •User Profiles
- Public & Private Authorities
- •Function Usage
- •Ownership
- Security Attributes
- Authorization Lists
- Security related System Values

Synchronous replication of Security Implementation is **not** optional

Determining Rules for Replication

Replication Criteria List

- Rules based engine
- Hierarchical

Default Inclusion State

- Include replicate every eligible object that exists plus any new objects as they are created
- Exclude replicate nothing

Then add hierarchical rules at lower levels

- Library
- Object type within library
- Specific object

Example 1:

Default Inclusion State: Exclude

Include everything in library Sales except output queues,

but include output queue Receipts

Library	Object Type	Object	Incl/Excl		
Sales	*all	*all	Include		
Sales	*OUTQ	*all	Exclude		
Sales	*OUTQ	Receipts	Include		

Example 2:

Default Inclusion State: Include

Include everything in library Sales except output queues,

but include output queue Receipts

Library	Object Type	Object	Incl/Excl		
Sales	*OUTQ	*all	Exclude		
Sales	*OUTQ	Receipts	Include		

IFS Support

- Requires IASP
- IFS accessible on both Nodes (R/W)
- Requires PowerHA
 - Db2 Mirror provides the simultaneous access.
 - PowerHA switches the IASP
- Filesystem automatically 'mutates' when the storage is switched





Db2 Mirror – What makes it different

- New integrated IBM i synchronization technology
- Does not leverage any existing availability technology to provide continuous availability
 - But does work with existing technology

Operating System Replication Technology



Db2 Mirror GUI

GUI runs on IBM i

GUI can run on the Db2 Mirror nodes

GUI can run outside of the Db2 Mirror nodes and manage multiple pairs

http://systemname:2006/Db2Mirror



Db2 Mirror for i now supports Direct Attach Storage

- Can configure Db2 Mirror for i regardless if the storage has copy service functionality.
- Db2 Mirror for i can use save and restore to configure mirroring
- Enables support for storage such as NVMe or SAS SSD





GUI Review of Application Readiness

- From an IBM i LPAR with the base 5770-DBM installed it is possible to point to another IBM i LPAR that is at 7.2, 7.3 or 7.4 to gain insight into some considerations for Db2 Mirror.
- Examine Replication Eligibility
- Examine Object Statistics
- Look for Journal Improvements



Topology Options – Disaster Recovery



Db2 Mirror – Where to get more information

IBM Knowledge Center						
Home > IBM i 7.4 > Availability >						
Db2 Mirror						
■ Table of Contents Change version or product ~						
 Start of Db2 Mirror documentation. PDF file for Db2 Mirror Use this to view and print a PDF of this information. Intro and architecture chapter 1 placeholder Db2 Mirror concepts This section describes the basic concepts you need to understand when working with Db2 Mirror. 						
Application considerations chapter 3 placeholder						
• Planning and setup There are many decisions that need to be made as you prepare and step through the Db2 Mirror setup process. Db2 Mirror configuration should be understood before starting in order to make the right decisions for your business.						

www.ibm.com/support/knowledgecenter/ssw_ibm_i_74/db2mi/db2mintro.htm

Thank you!

© Copyright IBM Corporation 2019

Notices and disclaimers

- © 2019 International Business Machines Corporation. No part of this document may be reproduced or transmitted in any form without written permission from IBM.
- U.S. Government Users Restricted Rights use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.
- Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. This document is distributed "as is" without any warranty, either express or implied. In no event, shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity. IBM products and services are warranted per the terms and

conditions of the agreements under which they are provided.

IBM products are manufactured from new parts or new and used parts.

In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply."

 Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.

- Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those
- customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.
- References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.
- Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.
- It is the customer's responsibility to insure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer's business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer follows any law.

Notices and disclaimers continued

- 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 about this publication 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. IBM does not warrant the quality of any third-party products, or the ability of any such third-party products to interoperate with IBM's products. IBM expressly disclaims all warranties, expressed or implied, including but not limited to, the implied warranties of merchantability and fitness for a purpose.
- The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents, copyrights, trademarks or other intellectual property right.

— IBM, the IBM logo, ibm.com and [names of other referenced IBM products and services used in the presentation] are 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: www.ibm.com/legal/copytrade.shtml

RoCE Adapters for Db2 Mirror

Adapter	Feature Code	Platform
PCIe3 2-port 100GbE NIC & RoCE QSFP28	EC3L / EC3M	Power8 & Power9
PCIe3 2-port 10Gb NIC & RoCE SR/Cu	EC2R / EC2S	Power9
PCIe3 2-port 25/10Gb NIC & RoCE SFP28	EC2T / EC2U	Power9
PCIe4 2-port 100GbE RoCE x16	EC66 / EC67	Power9





~____*___*___

Rolling Upgrade Scenario



Suspend the secondary node



Do your maintenance on the Secondary suspended node



Resume Mirroring to get the Systems back in Sync



Mirror Resume Progress

IBM Db	o2 Mirror f	ior i				Primary: ZZ2P28	Secondary: ZZ2P29	4 - U	Jser: qsecofr	IBM.
0	Summar	y - Current		e	Ð		GU	Build Time	: 2019-04-24 22	2:11:39
		*SYSBAS TRACKING	Group 6 (2019-04-27 10:24:25 -current)		0%	≡ 946 op	erations (946 objects) Bl	SYSBAS		
IBM Db2 Mirror for i 😌 Secondary: ZZ2P28 🔮 Secondary: ZZ2P29 🛔 👻 User: qsecofr IBM.						IBM.				
0	Summar	y - Current		(Ð		GU	l Build Time	: 2019-04-24 2	2:11:39
		*SYSBAS TRACKING	Group 6 (2019-04-27 10:24:25 -current)	24.1 seconds	48%	≡ 1000 c objects)	perations (1000 *	SYSBAS L OCKED		

Resync is Complete, Swap roles and repeat



Performance Expectations

- With synchronous replication the complete path length will increase since the action may drive I/O on both nodes in order to finish. This could increase by up to ~(2-3)X
- The ability to run transactions on both nodes will mitigate per transaction overhead and with a target of achieving equal to or greater transactional throughput
- Read workloads will not be impacted since they do not have to be replicated
- Single threaded or serial I/O workloads will be the most impacted.

