

In the ever-evolving landscape of database management, ensuring the availability and swift recovery of critical data has become paramount. This presentation dives into the latest enhancements and changes upcoming in availability and recovery with Db2's Next Release to help you get the most out of your investment and defend against potential disruptions.

1. Discuss and explore in detail the new features and changes in availability and recovery that are coming soon with Db2's Next Release.
2. Discuss what is next in availability and recovery.

3. Discuss the current top customer feature requests in availability and recovery.

Safe Harbor Statement

Copyright © IBM Corporation 2025. All rights reserved.

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

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON CURRENT THINKING REGARDING TRENDS AND DIRECTIONS, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. FUNCTION DESCRIBED HEREIN MAY NEVER BE DELIVERED BY IBM. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS AND/OR SOFTWARE.

IBM, the IBM logo, ibm.com and Db2 are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml

Agenda [1|2]

- What's New

- Backup Intra-Table Space Parallelism (ITP) Support
- History File Improvements
 - Provide Backup Size Info In History File
 - Store Encrypted Backup Information in History File
 - Increase Size of COMMENT Field
- Update db2adutil to display object sizes (12.1.2)
- Log File Validation During Archive
- Support streaming to object/remote storage for log archives (12.1.2)
- HADR Major Version Upgrade with Reads on Standby (RoS)
- Remove the Restriction of Restore and Rollforward Through a Topology Change in pureScale
- UPGRADE DATABASE now works with topology change (12.1.2)
- HADR SSL support in pureScale
- Enable Hostname Validation Feature in pureScale
- Add Support for S3-compatible (Remote) Storage to Automatic Db2 Backup
- Automatic History File Pruning for Automatic Db2 Backup

- Logical BAR Enhancements

- Support Storage Access Alias in Place of Access Keys (Use DB2REMOTE)
- List Backup Image Contents (New LOGICAL_BACKUP_DETAILS_TAB Table Function)
- Config Setting to Control Default Schema Type (Enable Row Modification Tracking by Default)

Agenda [2|2]

- **What's Changed**

- New Columns Added / New Formatting for BAR Stats Diag Entry
- Make DB2_USE_FAST_LOG_PREALLOCATION=YES the Default
- Allow index creation to exceed MAX_LOG / NUM_LOG_SPAN (12.1.2)
- Make DB2_BCKP_PAGE_VERIFICATION=ON the Default
- Make DB2_REDUCE_FLUSHING_DURING_BACKUP ON by Default (12.1.2)
- DB2_FAIL_RECOVERY_ON_TABLESPACE_ERROR Crash Recovery Support
- Externalize DB2_HADR_BLOCK_ON_DISKFULL (formerly DB2_HADR_DISKFULL_NONBLOCKING)
- Externalize DB2_DPS_RLOG_SHR_MEMBYTES
- Advanced Log Space Management (ALSM) on by Default Under DB2_WORKLOAD=SAP

- **Deprecated**

- DB2_HADR_NO_IP_CHECK (the Old HADR NAT support)

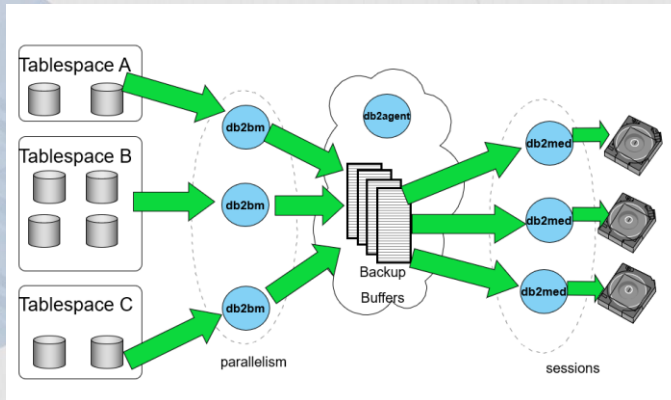
- **Discontinued**

- db2tapemgr
- XBSA Support for Backup/Restore

What's New

- Backup Intra-Table Space Parallelism (ITP) Support
- History File Improvements
 - Provide Backup Size Info In History File
 - Store Encrypted Backup Information in History File
 - Increase Size of COMMENT Field
- Update db2adult to display object sizes (12.1.2)
- Log File Validation During Archive
- Support streaming to object/remote storage for log archives (12.1.2)
- HADR Major Version Upgrade with Reads on Standby (RoS)
- Remove the Restriction of Restore and Rollforward Through a Topology Change in pureScale
- UPGRADE DATABASE now works with topology change (12.1.2)
- HADR SSL support in pureScale
- Enable Hostname Validation Feature in pureScale
- Add Support for S3-compatible (Remote) Storage to Automatic Db2 Backup
- Automatic History File Pruning for Automatic Db2 Backup

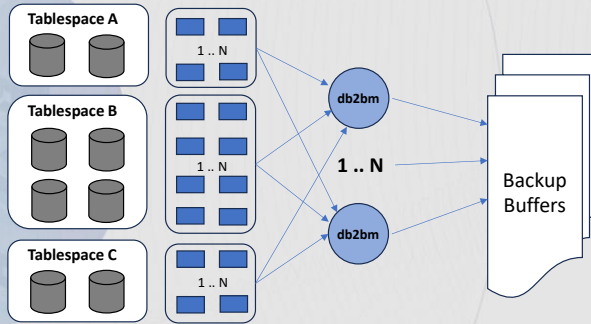
Backup Intra-Table Space Parallelism (ITP) Support [1 | 4]



- Today, backup assigns one db2bm EDU to a table space (tuned based on PARALLELISM)

- If a few large table spaces and many small table spaces, the small table spaces will be backed up quickly but then the large table spaces are read using a single EDU → overall backup time is proportional to the size of the largest table space

Backup Intra-Table Space Parallelism (ITP) Support [2 | 4]



- Solution is to break up table spaces into smaller segments (1 .. N) that can be assigned to BMs to be processed in parallel
- Split table space into an arbitrary number of chunks that 1 .. N db2bms will read
 - N determined by resource factors, like CPU, memory and parallel IO
- External registry variable:
 - DB2_BACKUP_ITP=ON/OFF
 - Default is ON

Backup Intra-Table Space Parallelism (ITP) Support [3 | 4]

- In cases where the sizes of table spaces are highly skewed (e.g. one tablespace is >70% of the total database size) we are seeing performance improvements of 2x or more
- In cases where the table spaces are balanced, we see much smaller improvements
- The amount of improvement we can get is capped by the I/O capacity of the channel between the disk and the CPU
- ITP works by making unused I/O capacity in the system available. Once that capacity is filled, you hit the maximum improvement that ITP can provide
- We were able to greatly improve ITP performance by adding I/O capacity to our test system
 - For example: install another fiber channel to your disk

Backup Intra-Table Space Parallelism (ITP) Support [4|4]

Backup DB SRC 468GB – Backup Parallelism 4 - Target 2 - Compress

db2 "backup database SRC to /backup/db2backup/SRC,/backup/db2backup/SRC with 8 buffers parallelism 4 compress"

Largest TS: 263GB

All TS: 468GB => **Largest TBSP 56% of all tablespaces**

ITP OFF – 5879s – 100%

2024-07-08-16:23:02.164386+120 E9451776A2152 LEVEL: Info

```
Parallelism = 4
Number of buffers = 8
Buffer size = 16781312 (4097 4kB pages)


BH# Total I/O Comp# MrgQ WaitQ Throttle LoadP# Buffers MBytes Comp#
-----
000 1:27:59 0:07:00 1:30:12 0:00:00 0:00:00 0:00:00 0:00:00 16778 203900 228975 0
001 1:27:59 0:01:01 0:20:13 0:00:00 1:16:52 0:00:00 0:00:00 2363 65676 65276 0
002 1:27:59 0:02:00 0:20:04 0:00:00 1:16:52 0:00:00 0:00:00 1940 71132 70434 0
003 1:27:59 0:02:04 0:20:01 0:00:00 1:16:52 0:00:00 0:00:00 2132 71909 71620 0
TOT - - - - - 23118 478409 436206 0

MCF Total I/O MrgQ WaitQ Buffers MBytes
-----
000 1:27:59 0:06:15 1:31:43 0:00:00 11307 181740
001 1:27:59 0:06:31 1:31:28 0:00:00 11763 188221
TOT - - - - - 23120 369962
```

Disk Usage in %



CPU Usage in %



ITP ON, PARA 4 – 2654s – 45%


2024-07-08-16:16:10.539351+120 E971386A2152 LEVEL: Info

```
Parallelism = 4
Number of buffers = 8
Buffer size = 16781312 (4097 4kB pages)


BH# Total I/O Comp# MrgQ WaitQ Throttle LoadP# Buffers MBytes Comp#
-----
000 0:44:14 0:04:22 0:39:37 0:00:00 0:00:00 0:00:00 0:00:00 5800 117028 109349 0
001 0:44:14 0:04:32 0:39:29 0:00:00 0:00:00 0:00:00 0:00:00 5792 121601 107956 0
002 0:44:14 0:04:28 0:39:33 0:00:00 0:00:00 0:00:00 0:00:00 5828 118105 104372 0
003 0:44:14 0:04:27 0:39:32 0:00:00 0:00:00 0:00:00 0:00:00 5588 120973 107964 0
TOT - - - - - 23118 478409 429644 0

MCF Total I/O MrgQ WaitQ Buffers MBytes
-----
000 0:44:14 0:08:09 0:36:05 0:00:00 11598 185587
001 0:44:14 0:08:07 0:36:06 0:00:00 11522 184365
TOT - - - - - 23120 369962
```

Disk Usage in %



CPU Usage in %



9

- POWER 8
- FC: 2x 8GBit ports – SDD and HDD on different ports
- AIX 7200-05-07-2346
- Backup target: HDD
- Backup source: SDD
- Db2 software compression used to slowdown I/O rate

History File Improvements [1 | 5]

Improvements:

1. Provide Backup Size Info In History File
2. Store encrypted backup information in history file
3. Increase size of COMMENT field
 - Increased from 30 to 254 characters

Motivation:

- More customers want better reporting or to script around inventory of objects, like images and log files:
 - Object sizes
 - Compressed? Encrypted? Backup includes logs?
 - Build restore estimations based on backup image size

History File Improvements [2|5]

- History file entry for backup:
 - Total Size: total size of all backup image sequences in bytes
 - Sequence Size: size of an individual backup image sequence in bytes
 - Compression Lib: ' ' (Not Compressed), 'C' (Default Compr Lib), 'U' (User Defined Custom Lib), 'N' (NX842), 'Z' (ZLIB)
 - Encrypted: Yes/No for 'LIST HISTORY', 1/0 for 'SYSIBMADM.DB_HISTORY'
 - Include Logs: Yes/No for 'LIST HISTORY', 1/0 for 'SYSIBMADM.DB_HISTORY'
- History file entry for load copy:
 - Total Size: total size of all load copy image sequences in bytes
 - Sequence Size: size of an individual load copy image sequence in bytes
 - Encrypted: Yes/No for 'LIST HISTORY', 1/0 for 'SYSIBMADM.DB_HISTORY'
- History file entry for log:
 - Total Size: actual size of the archived log in bytes
 - Compression Lib: ' ' (Not Compressed), 'C' (Default Compr Lib), 'N' (NX842), 'Z' (ZLIB)
 - Encrypted: Yes/No for 'LIST HISTORY', 1/0 for 'SYSIBMADM.DB_HISTORY'

History File Improvements [3|5]

- **SYSIBMADM.DB_HISTORY** displays **Total_Size**, **Seq_Size**, **Compression_Library**, **Encrypted**, and **IncludeLogs** for an online backup image made up of three sequences with compression(ZLIB) and encryption enabled:

```
$ db2 "select EID,TOTAL_SIZE, SEQ_SIZE, COMPRESSION_LIBRARY, ENCRYPTED, INCLUDE_LOGS from SYSIBMADM.DB_HISTORY where OPERATION='B'"
```

EID	TOTAL_SIZE	SEQ_SIZE	COMPRESSION_LIBRARY	ENCRYPTED	INCLUDE_LOGS
4	48308224	24166400	Z	1	1
5	48308224	12070912	Z	1	1
6	48308224	12070912	Z	1	1

3 record(s) selected.

History File Improvements [4|5]

- **SYSIBMADM.DB_HISTORY** displays Total Size, Compression Library and Encrypted for a log archive with both compression and encryption enabled:

```
$ db2 "select EID, TOTAL_SIZE, COMPRESSION_LIBRARY, ENCRYPTED from SYSIBMADM.DB_HISTORY where OPERATION='X'"
```

EID	TOTAL_SIZE	COMPRESSION_LIBRARY	ENCRYPTED
2	356352	C	1

1 record(s) selected.

History File Improvements [5 | 5]

- **LIST HISTORY** displays **Total Size**, **Compression Library** and **Encrypted** for a log archive with both compression and encryption enabled:

```
$ db2 list history archive log all for db sample

Op Obj Timestamp+Sequence Type Dev Earliest Log Current Log Backup ID
-----
X D 20240711195435      1 D S00000000.LOG C00000000
-----

Comment:
Start Time: 20240711195435
End Time: 20240711195436
Status: A
Total Size: 356352 (bytes)
Compr Lib: C
Encrypted: Yes
-----

EID: 2 Location: /home/wlarry/logs/wlarry/SAMPLE/NODE0000/LOGSTREAM0000/C0000000/S00000000.LOG
```

Update db2adult to display object sizes

- Tool to help manage database recovery objects stored in IBM Storage Protect (formerly IBM Spectrum Protect / Tivoli Storage Manager (TSM))
- Backup images, load copy images and log archives now have size, in bytes, available
 - Assist with recovery estimation
 - Inventory tracking
 - The new default for 12.1.2



Query for database TESTDB

Retrieving FULL DATABASE BACKUP information.

1 Time: 20250119181811, Size: 302096384 (bytes), Oldest log: S0000003.LOG, DB Partition Number: 0, Sessions: 1
2 Time: 20250119174807, Size: 268533760 (bytes), Oldest log: S0000000.LOG, DB Partition Number: 0, Sessions: 1

...

Retrieving LOAD COPY information.

1 Time: 20250119180412, Size: 278528 (bytes), DB Partition Number: 0, Sessions: 1

Retrieving LOG ARCHIVE information.

Log file: S0000000.LOG, Size: 12288 (bytes), Chain Num: 0, Log stream: 0, Taken at: 2025-01-19-21.00.25
Log file: S0000001.LOG, Size: 57344 (bytes), Chain Num: 0, Log stream: 0, Taken at: 2025-01-19-21.00.28
Log file: S0000002.LOG, Size: 573440 (bytes), Chain Num: 0, Log stream: 0, Taken at: 2025-01-19-21.06.08

Log File Validation During Archive [1 | 5]

- **Customer Ask:**
 - **When we archive a log file do verification similar to db2cklog**
 - **To service archive log file corruption reports**
 - Need a way to see if Db2 or file system issue
- **We had basic internal way to validate on archive (11.5), time to externalize**
 - `DB2_VALIDATE_LOG_ON_ARCHIVE=ON/OFF`
 - Default is ON
- **Solution:**
 - Perform validation on archive
 - If log file invalid retry the validation in case the problem is transient
 - The retry logic will stop when the next log file is produced and ready to be archived
 - Conclude this log file is invalid and archive with validation error before moving on to the next log file
 - Notify user that recoverability is compromised, and a new database backup is needed

Log File Validation During Archive [2|5]

- Add new indicator ("VE:") to the Comment field in the history file entry

```
Op Obj Timestamp+Sequence Type Dev Earliest Log Current Log Backup ID
-----
X D 20240213115521 1 D S0000003.LOG C0000000
-----
Comment: VE:-servername=server1
Start Time: 20240213115521
End Time: 20240213115618
Status: A
-----
EID: 4 Location: /db2/user/SAMPLE/NODE0000/LOGSTREAM0000/C0000000/S0000003.LOG

Op Obj Timestamp+Sequence Type Dev Earliest Log Current Log Backup ID
-----
X D 20240213115521 1 D S0000004.LOG C0000000
-----
Comment: VE:
Start Time: 20240213115722
End Time: 20240213115833
Status: A
-----
EID: 5 Location: /db2/user/SAMPLE/NODE0000/LOGSTREAM0000/C0000000/S0000004.LOG
```

17

This allows us to later examine the problematic log file for PD, without having to keep it in the active log path which can cause disk space issue.

Log File Validation During Archive [3|5]

- New admin message written to the notify log when detecting a log file is invalid during archive:

```
ADM1841W  The database manager was either unable to locate log
"<logFile>" or found a validation issue with the log while
attempting to archive it. Your existing recovery strategy may
not work if this log file is required during recovery.
A database backup is required to ensure the recoverability of the
database. However, the backup should be taken after the First
Active Log file (LOGHEAD) DB configuration parameter is beyond the
specified log. You might want to consider deactivating the
database now to have the First Active Log file (LOGHEAD) move up,
and then take a backup.
```

- New element in MON_GET_TRANSACTION_LOG monitor table function to report the last invalid Log file detected

```
>> db2 "Select MEMBER, last_log_validation_error from table(mon_get_transaction_log(-1)) as t"
```

```
MEMBER LAST_LOG_VALIDATION_ERROR
-----
```

```
0 3
```

```
1 record(s) selected.
```

Log File Validation During Archive [4 | 5]

- New reporting from "db2pd -logs" of the last invalid log file detected

```
>> db2pd -logs -db ese
```

```
Database Member 0 -- Database ESE -- Active -- Up 0 days 00:00:06 -- Date 2023-02-14-10.26.03.119965
Logs:
Current Log Number          0
Pages Written               0
Cur Commit Disk Log Reads  0
Cur Commit Total Log Reads 0
Method 1 Archive Status     Success
Method 1 Next Log to Archive 6
Method 1 First Failure      n/a
Method 2 Archive Status     n/a
Method 2 Next Log to Archive n/a
Method 2 First Failure      n/a
Last Log Validation Error   3      ← NEW FIELD
Extraction Status           n/a (0)
Extraction Throttle Reason  n/a
Current Log to Extract      n/a
Log Chain ID                0
Current LSO                  86182945
Current LSN                  0x00000000000056A62
Address      StartLSN      StartLSO      State      Size      Pages      Filename
0x00007FC0D9A7B126 0000000000000000 86182945 0x00000000 1024    1024      S0000006.LOG
0x00007FC0D9A7B18E 0000000000000000 90356769 0x00000000 1024    1024      S0000007.LOG
0x00007FC0D9A7B1F6 0000000000000000 94530593 0x00000000 1024    1024      S0000008.LOG
```

Log File Validation During Archive [5|5]

Notes: LAST LOG VALIDATION ERROR field

- Kept in sync between HADR primary and secondary
- Will be reset when the log archive entry in the history file is pruned
- Will only report the latest log that failed validation on archive. The latest failure determines if recoverability of the database is compromised

Support streaming to object/remote storage for log archives



- **Support for DB2REMOTE as a log archive method arrived in 11.5.8**
 - Archive and retrieve uses a staging area
- **New for 12.1.2, log archives are now streamed**
 - No longer uses a staging area
 - Log retrieves still use staging area
- **10 - 20% reduction in log archive times**

HADR Major Version Upgrade with Reads on Standby (RoS)

- **Sample Scenario:**

- Single database A in an ESE instance
- Primary on host 1 and standby on host 2
 - Database activated on both



- **High Level Process¹:**

- HADR roles maintained; **standby will NOT need to be re-initialized**
- **Downlevel release MUST be Db2 11.5.8 or later with pre-requisite DT259396**
- Primary and standby must validate log positions in downlevel release
- Primary must be shutdown first, but **standby can stay activated in downlevel release**
- Primary is upgraded first but no logs shipped to standby during upgrade
- Standby can service read only application requests, while in downlevel release
- Primary and standby must be at same code level to communicate (log ship)
- Once primary upgraded, standby can be upgraded and logs will be replayed

- **Summary:**

- Worth noting is log gap that can form as primary is being upgraded and standby in downlevel
- After primary upgraded and opened for business, no HA protection for unplanned outage until standby catches up into peer state
- Trade off we believe customers willing to accept
- With Multiple Standbys, can keep auxiliary downlevel

22

¹ See Performing Db2 HADR Upgrades Made Easy (IDUG 2023) for more in-depth details of base procedure:

<https://roecken.ca/academia/presentations/IDUGEMEA2023-PerformingDb2HADRUpgradesMadeEasy.pdf>

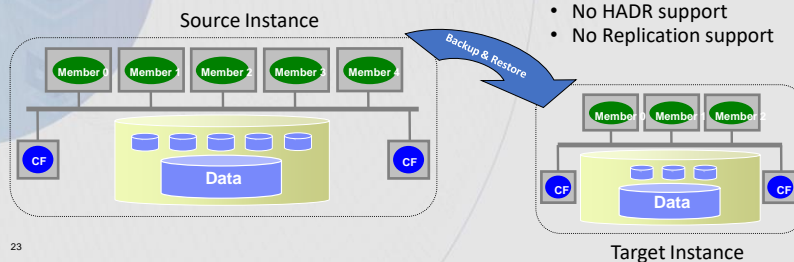
Same procedure can be applied to multiple standby (principal and/or auxiliary standby).

Downlevel release MUST be Db2 11.5.8 or later with pre-requisite DT259396. Ensure you are using a CSB that includes this KI.

Remove the Restriction of Restore and Rollforward Through a Topology Change in pureScale [1 | 2]

Pre-12.1 pureScale to pureScale:

- Online/offline add member
 - Restore + rollforward support (offline/online images)
 - HADR support
- Offline drop member
 - Topology life change event
 - Database must be consistent; no backup / restore / rollforward pending
 - After drop of member must take offline database backup
 - Restore of offline backup images only and without rolling forward
 - No database or table space rollforward support; non-recoverable point
 - No HADR support
 - No Replication support



Db2 10.5 came out with online add member and restore+rollforward support. Drop member is offline with these restrictions:

The last member cannot be dropped.

The catalogued databases must satisfy the following criteria:

- Consistent database.
- Have at least one reconciled member remaining after the drop.
- HADR role standard.
- Not backup pending.
- Not topology life change pending or in progress.
- Not database rollforward pending.
- All table spaces must be in normal state (for example, not rollforward or restore pending).
- Not restore pending.

A drop member is considered a topology life change event: in addition to writing a topology life change log record, this generates a new topology life identifier.

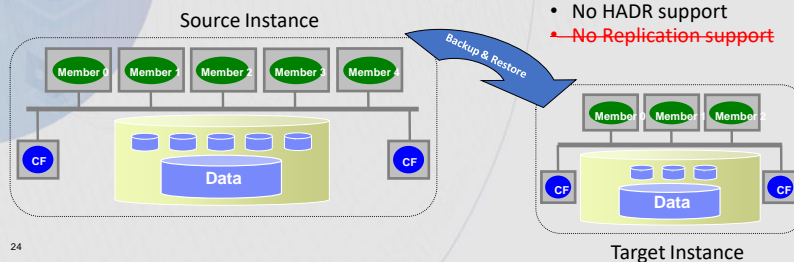
The user is required to take a full offline database backup after one or more drop member operations.

Any restore that results in members dropped from the topology is considered a topology life change event and must have a full offline backup taken after the restore to guarantee recoverability.

Remove the Restriction of Restore and Rollforward Through a Topology Change in pureScale [2 | 2]

12.1 pureScale to pureScale:

- Online/offline add member
 - Restore + rollforward support (offline/online images)
 - HADR support
- Offline drop member
 - ~~Topology life change event~~
 - Database must be consistent; no backup / restore / rollforward pending
 - ~~After drop of member must take offline database backup~~
 - ~~Restore of offline backup images only and without rolling forward~~
 - ~~No database or table space rollforward support; non-recoverable point~~
 - No HADR support
 - ~~No Replication support~~



NOTE: pureScale <-> non-pureScale scenarios/restrictions not changed

Db2 10.5 came out with online add member and restore+rollforward support.
Db2 12.1 came out with offline drop member and restore+rollforward support.

Offline drop member has these restrictions:

The last member cannot be dropped.

The catalogued databases must satisfy the following criteria:

- Consistent database.
- Have at least one reconciled member remaining after the drop.
- HADR role standard.
- Not backup pending.
- Not topology life change pending or in progress.
- Not database rollforward pending.
- Not restore pending.

UPGRADE DATABASE now works with topology change

A small square icon with a colorful, abstract design and the text "12.1.2" in white.

- **Pre-12.1.2, following commands did not work on a downlevel database that was topology change pending**
 - CATALOG DATABASE
 - UPGRADE DATABASE
 - db2ckupgrade
- **Examples:**
 - pureScale 3 members to pureScale 1 member (drop member pending)
 - pureScale 1 member to pureScale 3 members (add member pending)
 - pureScale any member to non-pureScale single partition (config type mismatch)
- **Expectation:**
 - UPGRADE and use database with no issues
 - If explicit drop member pending or pureScale to non-pureScale involving a pre-12.1 database, post UPGRADE database will be BACKUP PENDING
- **Future**
 - RESTORE+UPGRADE support coming later in 2025

HADR SSL support and Enable Hostname Validation Support in pureScale

- **HADR_SSL_LABEL** – specifies the label of the SSL certificate which encrypts communication between the primary and standby
- **HADR_SSL_HOST_VAL** – specifies whether hostname validation for TLS connections between the primary and standby is enabled
- **Pre-12.1** these two database cfg parms were not supported in pureScale
 - db2checkSD would catch and fail
- **For 12.1** users can choose for **HADR_SSL_HOST_VAL**:
 - To use SSL certificate with hostname validation (BASIC); or
 - To use SSL certificate without hostname validation (OFF)

HADR_SSL_LABEL - Label name in the key file for encrypted communication between HADR primary and standby instances configuration parameter
<https://www.ibm.com/docs/en/db2/11.5?topic=parameters-hadr-ssl-label-label-name-ssl>

HADR_SSL_HOST_VAL - TLS Hostname validation status configuration parameter
<https://www.ibm.com/docs/en/db2/11.5?topic=parameters-hadr-ssl-host-val-label-name-ssl>

Add Support for S3-Compatible (Remote) Storage to Automatic Db2 Backup

- **BACKUP** command currently supports remote/object storage as a media type:

```
DB2REMOTE://<alias>//<storage-path>
```

- There also exists automatic backup through **AUTO_MAINT / AUTO_DB_BACKUP** db cfg parms
- This is a policy-based backup using XML syntax – it does not have syntax to support remote storage as a media type
- In 11.5.7/8, we enhanced our object storage support, this extends that to automatic backup
- Example policy file would contain:

```
<RemoteStorageBackupTarget>  
  <PathName>DB2REMOTE://<alias>//<storage-path></PathName>  
</RemoteStorageBackupTarget>
```

27

Automatic database backup

<https://www.ibm.com/docs/en/db2/11.5?topic=overview-automatic-database-backup>

Automatic History File Pruning for Automatic Db2 Backups

- Applies to automatic backup through AUTO_MAINT / AUTO_DB_BACKUP db cfg parms
- Currently has its own automatic pruning policy of Db2 backup images (keeps most recent)
- This contradicts Db2's automatic pruning feature controlled by AUTO_DEL_REC_OBJ db cfg parm
- Request is to have automatic backup and automatic pruning behave together
- **Solution:**
 - If AUTO_DEL_REC_OBJ is ON then use Db2's automatic pruning feature
 - Otherwise continue to use current automatic backup policy of keeping most recent

28

Automatic database backup

<https://www.ibm.com/docs/en/db2/11.5?topic=overview-automatic-database-backup>

Logical Schema Backup/Restore (LBAR) Enhancements

- Support Storage Access Alias in Place of Access Keys (Use DB2REMOTE)
- List Backup Image Contents (New LOGICAL_BACKUP_DETAILS_TAB Table Function)
- Config Setting to Control Default Schema Type (Enable Row Modification Tracking by Default)

Support Storage Access Alias in Place of Access Keys (Use DB2REMOTE)

- Ability to specify a remote storage alias with LBAR stored procedures and table function instead of specifying bare S3/COS credentials. Example:

```
db2 CATALOG STORAGE ACCESS ALIAS backupAlias VENDOR S3  
SERVER s3.us-east-2.amazonaws.com USER username PASSWORD password CONTAINER mybucket
```

```
# Doing backup with remote storage alias
```

```
db2 call sysproc.logical_backup('-type onl -schema myschema -path  
DB2REMOTE://backupAlias')
```

```
# Doing it with specifying bare credentials
```

```
db2 "call sysproc.logical_backup('-type onl -schema myschema -s3 -access-key  
username -secret-key password -bucket-url mybucket -endpoint s3.us-east-  
2.amazonaws.com') "
```

30

Schema-level and table-level backup and restore

<https://www.ibm.com/docs/en/db2/11.5?topic=recovery-schema-level-table-level-backup-restore>

List Backup Image Contents (New LOGICAL_BACKUP_DETAILS_TAB Table Function)

- Ability to get timestamps for backup images and also list of tenants, schemas, tables combinations for each backup image. Each field of timestamp, tenant, schema, table, etc has their column. Example:

Now run the table function to see list of tables in each schema, and also timestamps, etc.
db2 select * from sysproc.LOGICAL_BACKUP_DETAILS_TAB('-path /home/user/mybackups')

TIMESTAMP	TYPE	PREV_TIMESTAMP	INSTANCE	DATABASE	TENANT	SCHEMA	TABLE	...
202403131114548	F	-	regress1	MYDB	SYSTEM	LBAR1	TABLE1	
202403131114548	F	-	regress1	MYDB	SYSTEM	LBAR2	TABLE1	
202403131114548	F	-	regress1	MYDB	SYSTEM	LBAR2	TABLE2	
202403131114548	F	-	regress1	MYDB	SYSTEM	LBAR3	-	
202403131114548	F	-	regress1	MYDB	SYSTEM	LBAR4	-	
202403131115790	I	202403131114548	regress1	MYDB	SYSTEM	LBAR1	TABLE1	
202403131115790	I	202403131114548	regress1	MYDB	SYSTEM	LBAR1	TABLE2	
202403131115790	I	202403131114548	regress1	MYDB	SYSTEM	LBAR2	TABLE1	
202403131115790	I	202403131114548	regress1	MYDB	SYSTEM	LBAR2	TABLE3	
202403131115790	I	202403131114548	regress1	MYDB	SYSTEM	LBAR3	TABLE1	
202403131115790	I	202403131114548	regress1	MYDB	SYSTEM	LBAR4	-	

31

Schema-level and table-level backup and restore

<https://www.ibm.com/docs/en/db2/11.5?topic=recovery-schema-level-table-level-backup-restore>

Config Setting to Control Default Schema Type (Enable Row Modification Tracking by Default)

- Instead of explicitly specifying 'enable row modification tracking' when schema gets created, we can just set the db cfg parameter so all schemas are created as row modification tracking by default. Example:
- **Old:**
 - `db2 create schema LBAR1 enable row modification tracking`
- **New:**
 - `db2 update db cfg DFT_SCHEMAS_RMT for MYDB using YES`
 - `db2 create schema LBAR1`

32

Schema-level and table-level backup and restore

<https://www.ibm.com/docs/en/db2/11.5?topic=recovery-schema-level-table-level-backup-restore>

What's Changed

- New Columns Added / New Formatting for BAR Stats Diag Entry
- Make DB2_USE_FAST_LOG_PREALLOCATION=YES the Default
- Allow index creation to exceed MAX_LOG / NUM_LOG_SPAN (12.1.2)
- Make DB2_BCKP_PAGE_VERIFICATION=ON the Default
- Make DB2_REDUCE_FLUSHING_DURING_BACKUP ON by Default (12.1.2)
- DB2_FAIL_RECOVERY_ON_TABLESPACE_ERROR Crash Recovery Support
- Externalize DB2_HADR_DISKFULL_NONBLOCKING
- Externalize DB2_DPS_RLOG_SHR_MEMBYTES
- Advanced Log Space Management (ALSM) on by Default Under DB2_WORKLOAD=SAP

New Columns Added for BAR Stats Diag Entry

- Three new fields by default:

- Throttle (added in 11.5.8)
- LogPhase
- LogBytes

FUNCTION: DB2 UDB, database utilities, sqluxLogDataStats, probe:2051
MESSAGE : Performance statistics
DATA #1 : String, 1864 bytes

Parallelism = 10
Number of buffers = 20
Buffer size = 16781312 (4097 4kB pages)

BM#	Total	I/O	MsgQ	WaitQ	Throttle	LogPhase	Buffers	kBytes	LogBytes
---	-----	-----	-----	-----	-----	-----	-----	-----	-----
000	0.59	0.35	0.00	0.23	0.00	0.00	4	16400	0
...									
---	-----	-----	-----	-----	-----	-----	-----	-----	-----
TOT	-	-	-	-	-	-	15	128432	0

- Open: How to account for time using remote storage media?

New Formatting for BAR Stats Diag Entry

- Representing times as seconds makes it difficult to understand the scale of the time

BM#	Total	I/O	Compr	MsgQ	WaitQ	Buffers	GBytes	GBytes
---	-----	-----	-----	-----	-----	-----	-----	-----
000	92873.78	43700.49	9262.30	11644.32	26728.88	406173	6334	11189
001	92862.32	26121.07	6196.32	9588.78	49743.70	274118	4271	6344
...								

- How long is 92000 seconds? These values are for human consumption so they should be human-readable
- We now represent the time in HH:MM:SS format

BM#	Total	I/O	Compr	MsgQ	WaitQ	Buffers	Gbytes	Gbytes
---	-----	-----	-----	-----	-----	-----	-----	-----
000	25:47:53	12:08:20	2:34:22	3:14:04	7:25:28	406173	6334	11189
001	25:47:42	7:15:21	1:43:16	2:39:48	13:49:03	274118	4271	6344
...								

Make DB2_USE_FAST_LOG_PREALLOCATION=YES the Default [1 | 2]

- Cost of creating log files can be expensive
- Using fast pre-allocation improves by reserving logical space up front and defers physical aspects to write; makes create code faster
- Operating systems: AIX and Linux on Veritas VxFS, JFS2, GPFS, ext4 (Linux only) and xFS (Linux only) file systems
- Pre-12.1: Default is OFF
 - SAP runs ON under DB2_WORKLOAD
- 12.1: Default is now ON
 - Removed from DB2_WORKLOAD=SAP

DB2_USE_FAST_PREALLOCATION

- Operating system: AIX and Linux on Veritas VxFS, JFS2, GPFS, ext4 (Linux only) and xFS (Linux only) file systems
- Default: ON for Veritas VxFS, JFS2, GPFS, ext4 and xFS, Values: ON or OFF
- Allows the fast preallocation file system feature to reserve table space, and speed up the process of creating or altering large table spaces and database restore operations. This speed improvement is implemented at a small delta cost of performing actual space allocation during runtime when rows are inserted. To disable fast preallocation, set **DB2_USE_FAST_PREALLOCATION** to OFF. This might improve runtime performance, at the cost of slower table space creation and database restore times, on some operating systems, especially AIX, when there is a large volume of inserts and selects on same table space. Note that once fast preallocation is disabled, the database has to be restarted.

DB2_USE_FAST_LOG_PREALLOCATION

- Operating system: AIX and Linux on Veritas VxFS, JFS2, GPFS, ext4 (Linux only) and xFS (Linux only) file systems
- Default: See **DB2_USE_FAST_PREALLOCATION**, Values: ON or OFF
- Allows the fast preallocation file system feature to reserve space for log files, and

speed up the process of creating or altering large log files, if the underlying file system supports this feature. This speed improvement is implemented at a small delta cost of performing actual space allocation during runtime when log records get written to such preallocated log files. To enable fast preallocation for logs, set **DB2_USE_FAST_LOG_PREALLOCATION** to ON.

- Changes to this variable take effect immediately on the next log file allocation. Changes are applied without the need to restart the database instance, or the need to deactivate and activate the database.

Make DB2_USE_FAST_LOG_PREALLOCATION=YES the Default [2 | 2]

- **NOTE: Complications with xFS and Red Hat (prior to 8.3) when disk full/fragmentation**
 - TechNote: Db2 can fail with an incorrect file size:
 1. <https://www.ibm.com/support/pages/node/6984753>
 2. <https://www.ibm.com/support/pages/db2usefastlogpreallocation-registry-variable-xfs-can-fail-incorrect-file-size>
 - File systems recommended for Db2:
<https://www.ibm.com/support/pages/file-systems-recommended-db2-linux-unix-and-windows>

Allow index creation to exceed MAX_LOG / NUM_LOG_SPAN

- **MAX_LOG**

- Maximum log per transaction configuration parameter
- Percentage of primary log space one transaction can consume

- **NUM_LOG_SPAN**

- Number log span configuration parameter
- Limit to how many log files one transaction can span

- **Problem**

- Large HADR database, LOGINDEXBUILD configuration is set, this means CREATE INDEX is logged, which can exceed the limit imposed by MAX_LOG and NUM_LOG_SPAN
→ transaction rolled back

- **Solution**

- MAX_LOG / NUM_LOG_SPAN limit will only apply to normal workload
 - Not maintenance operations such as CREATE INDEX, including index recreation (if index was marked bad)



Make DB2_BCKP_PAGE_VERIFICATION=ON the Default

- **Verifies pages during backup**
 - Fails backup if anything does not verify correctly
- **Pre-12.1: Default is OFF**
 - SAP runs ON under DB2_WORKLOAD
 - pureScale runs ON
- **12.1: Default is now ON**
 - Removed from DB2_WORKLOAD=SAP
 - Performance impact: 2% degradation

39

DB2_BCKP_PAGE_VERIFICATION

- Operating system: All
- Default: ON, Values: ON, OFF
- Specifies whether DMS and AS page validation occurs during a backup. This registry variable is dynamic, that is, it can be set or unset without having to stop and start the instance.

Make DB2_REDUCE_FLUSHING_DURING_BACKUP ON by Default

12.1.2

- **Reduces the amount of bufferpool flushing during online backup**
 - Large bufferpools with concurrent workload can spend long time flushing
 - We have seen sometimes 30+ minutes
 - Long running online backups means more logs to retrieve and include into image
 - DB2_REDUCE_FLUSHING_DURING_BACKUP introduced in 11.1.4
- **Pre-12.1.2: Default is OFF**
 - SAP runs ON under DB2_WORKLOAD
- **12.1.2: Default is now ON**
 - Removed from DB2_WORKLOAD=SAP
 - Improved bufferpool flushing can see drastic performance gains for certain scenarios
 - Shorter running backups == less logs to include == quicker rollforward to end of backup

DB2_REDUCE_FLUSHING_DURING_BACKUP

- Operating system: All
- Default: ON, Values: ON | OFF
- At the beginning of an online backup operation, changed pages within the buffer pool must be persisted to table space storage. The persisting of these changed pages can increase the duration of the backup operation in configurations with very large buffer pools. This variable specifies whether an online backup runs reduced flushing of changed pages within the buffer pool.
- Changes to this variable do not require the database instance to be restarted.
- Changes to this variable have no effect on an online backup operation that is running at the time of the variable change.

DB2_FAIL_RECOVERY_ON_TABLESPACE_ERROR Crash Recovery Support

- Completed reg var support to include table space errors during crash recovery

Operating system: All

Default: NO, Values: NO, YES, ROLLFORWARD or **CRASHRECOVERY**

This variable specifies whether a recovery operation (including database/table space rollforward, database/table space restore, HADR standby replay and **crash recovery**) should fail after encountering an error condition on a table space.

If set to NO, Db2 would change the table space state and skip all subsequent recovery actions on the table space. This allows a user to recover the rest of the table spaces for the database first. The affected table space can be dropped or recovered using a new subsequent recovery operation.

If set to YES, the recovery operation would fail without changing the table space state for database/table space rollforward, database/table space restore, HADR standby replay and **crash recovery**.

If set to ROLLFORWARD, the recovery operation would fail without changing the table space state for database/table space rollforward, database/table space restore and HADR standby replay.

If set to CRASHRECOVERY, the recovery operation would fail without changing the table space state for crash recovery.

Changes to this variable do not require the database instance to be restarted and take effect at the start of the recovery operation.

Externalize DB2_HADR_BLOCK_ON_DISKFULL (formerly DB2_HADR_DISKFULL_NONBLOCKING)

- **Default behaviour of HADR standby running into disk full condition in active log path is that the standby database will stay up (similar to the behaviour where blk_on_log_disk_full database configuration parameter is enabled)**
 - This can cause transactions on primary to be blocked
- **With internal registry variable DB2_HADR_DISKFULL_NONBLOCKING set to ON (default OFF), standby database will be brought down in disk full situation, thus taking HADR out of PEER state and unblock the transactions on the primary**
- **Solution:**
 - New registry variable DB2_HADR_BLOCK_ON_DISKFULL
 - Clearer and avoids double negative
 - Default: ON; Values: ON, OFF
 - ON – standby stays up; OFF – standby brought down
 - Discontinue DB2_HADR_DISKFULL_NONBLOCKING
 - Users that have set will see db2diag.log message informing them to use new registry variable
 - Will still honour if DB2_HADR_BLOCK_ON_DISKFULL not set

Externalize DB2_DPS_RLOG_SHR_MEMBYTES

- Internal registry variable to help manage memory allocated to read log data for db2ReadLog (Replication) applications
 - More memory can improve performance
- We externalized in Tech Notes and APAR text and from time to time take questions on it

43

DB2 readlog API Slow Performance

<https://www.ibm.com/support/pages/db2-readlog-api-slow-performance>

Advanced Log Space Management (ALSM) on by Default Under DB2_WORKLOAD=SAP [1 | 2]

- **Use ALSM to reduce your likelihood of hitting transaction log full conditions**
 - Enabled by setting `DB2_ADVANCED_LOG_SPACE_MGMT=ON`
- **In 12.1 DB2_WORKLOAD=SAP includes ALSM ON by default**

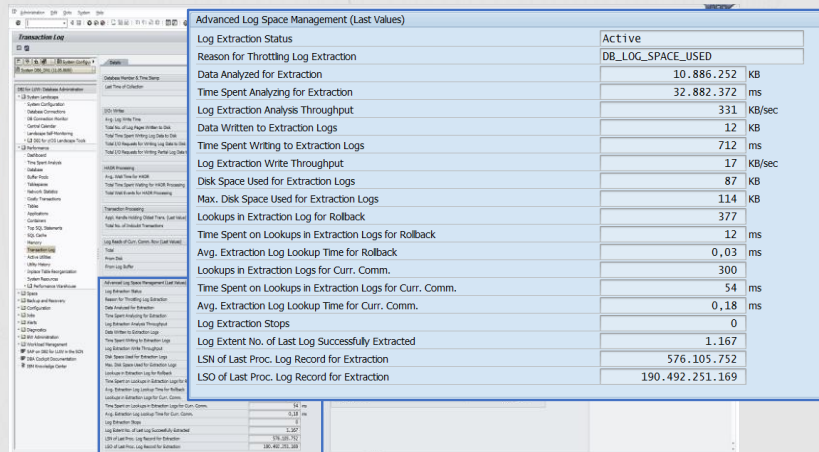
44

Advanced Log Space Management

<https://www.ibm.com/docs/en/db2/11.5?topic=logging-advanced-log-space-management>

Advanced Log Space Management (ALSM) on by Default Under DB2_WORKLOAD=SAP [2 | 2]

- ALSM integration into SAP DBA Cockpit



Advanced Log Space Management (Last Values)	
Log Extraction Status	Active
Reason for Throttling Log Extraction	DB_LOG_SPACE_USED
Data Analyzed for Extraction	10.886.252 KB
Time Spent Analyzing for Extraction	32.882.372 ms
Log Extraction Analysis Throughput	331 KB/sec
Data Written to Extraction Logs	12 KB
Time Spent Writing to Extraction Logs	712 ms
Log Extraction Write Throughput	17 KB/sec
Disk Space Used for Extraction Logs	87 KB
Max. Disk Space Used for Extraction Logs	114 KB
Lookups in Extraction Log for Rollback	377
Time Spent on Lookups in Extraction Logs for Rollback	12 ms
Avg. Extraction Log Lookup Time for Rollback	0,03 ms
Lookups in Extraction Logs for Curr. Comm.	300
Time Spent on Lookups in Extraction Logs for Curr. Comm.	54 ms
Avg. Extraction Log Lookup Time for Curr. Comm.	0,18 ms
Log Extraction Stops	0
Log Extent No. of Last Log Successfully Extracted	1.167
LSN of Last Proc. Log Record for Extraction	576.105.752
LSO of Last Proc. Log Record for Extraction	190.492.251.169

45

Advanced Log Space Management

<https://www.ibm.com/docs/en/db2/11.5?topic=logging-advanced-log-space-management>

Deprecatcd / Discontinued

- Deprecatcd

- DB2_HADR_NO_IP_CHECK (the Old HADR NAT support)

- Discontinued

- db2tapemgr
- XBSA Support for Backup/Restore

Deprecated DB2_HADR_NO_IP_CHECK (The Old HADR NAT support)

- **NAT means Network Address Translation**
- **Old method:**
 - Registry variable DB2_HADR_NO_IP_CHECK
 - Had limitations/issues thus deprecating in 12.1
- **New method:**
 - Configure db cfg values HADR_LOCAL_HOST / HADR_LOCAL_SVC with both the internal and external address as "<X> | <Y>"
 - Since 11.5.6
 - Ensure DB2_HADR_NO_IP_CHECK not set
 - Examples: <https://www.ibm.com/docs/en/db2/11.5?topic=support-hadr-nat>

HADR and network address translation (NAT) support

<https://www.ibm.com/docs/en/db2/11.5?topic=support-hadr-nat>

Discontinued db2tapemgr

- **Was deprecated in 11.5; not aware of any use**
 - SAP environments had most use case scenarios
- **Discontinued in 12.1; executable removed**

48

Log archiving using db2tapemgr

<https://www.ibm.com/docs/en/db2/11.5?topic=management-log-archiving-using-db2tapemgr>

db2tapemgr - Manage log files on tape command



<https://www.ibm.com/docs/en/db2/11.5?topic=commands-db2tapemgr-manage-log-files-tape>

Discontinued XBSA Support

- The XBSA library was removed from the product in 2006, but the docs never updated
- Samples removed


Tentative For the Remainder of 2025


- **Decouple Automatic History File Pruning from BACKUP Utility**
- **Automatic Archive Log Pruning When Non-Db2 Snapshots are Used**
- **RESTORE+UPGRADE Support of a Downlevel Database that is pureScale Add/Drop Member Pending**
- **HADR pureScale Drop Member Support**
- **HADR pureScale Mixed Topology Support**
- **Reduce the need to retrieve active logs files to include into an online backup**



**Db2 Availability & Recovery:
Get Ready For What's Next in 12.1.x**

Michael Roecken

 *@roecken*



Michael Roecken is a senior software developer with Db2 for Linux, UNIX, and Windows platforms at the IBM Toronto Lab. Michael has worked since 2000 designing, implementing and supporting various features and capabilities in the areas of: backup/restore, crash/rollforward recovery, high availability/disaster recovery, logging/transaction management and upgrade/migration.

Please connect with Michael on X (formerly known as Twitter) at [@roecken](#) and LinkedIn.