




CCDUG 2026



**What's new and coming soon in
History File, Backup and Recovery, and
HADR**

Ryan Xi

ryan@ca.ibm.com

Safe Harbor Statement

Copyright © IBM Corporation 2026. 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

- History file
 - Autoprune configuration
 - Asynchronous autoprune (NEW)
 - Non-integrated backup (future)
- Backup and recovery
 - Online backup caches log files (NEW)
 - Online backup includes log earlier (future)
- Topology/HADR
 - Restore of database with different topology across version (NEW)
 - HADR mixed topology (NEW)
 - 3+ Standby (future)
 - Backup on standby / pS multiple standby (future)

Question

How do you managing old backup images and archived log files?

- Manually delete based on timestamp?
- Script/cron job to delete?
- Pruning?
- Script/cron job to prune?
- Automatic pruning?


↳ We don't recommend manually deleting log files

Is the log really safe to delete from archive?

How do you managing old backup images and archived log files?

- Manually delete based on timestamp?
- Script/cron job to delete?

When deleting Backup 1, are log 0-3 safe to delete?



↳ Let Db2 can manage for you!

History file

- Record of recovery and administrative events on database
 - Create, alter, drop, rename tablespace
 - Reorganize table
 - Backup
 - Restore, rollforward
 - Archive log
 - Load

```
(ryan@db2i-ryan-fh2du-x86) /home/ryan/ryan/NODE0000/SQL00001
```

```
$ ls
```

db2rhist.asc	LOGSTREAM0000	SQLLOGCTL.GLFH.2	SQLSPCS.2
db2rhist.bak	MEMBER0000	SQLSGF.1	
db2rhist.lock	SQLDBCNF	SQLSGF.2	
HADR	SQLLOGCTL.GLFH.1	SQLSPCS.1	

History file

- LIST HISTORY command

```
list history backup all for db SAMPLE
```

```
List History File for SAMPLE
```

```
Number of matching file entries = 4
```

Op	Obj	Timestamp+Sequence	Type	Dev	Earliest Log	Current Log	Backup ID
B	D	20260508142952001	N	D	S0000028.LOG	S0000028.LOG	

```
Contains 2 tablespace(s):
```

```
00001 SYSCATSPACE
```

```
00002 USERSPACE1
```

```
Comment: DB2 BACKUP SAMPLE ONLINE
```

```
Start Time: 20260508142952
```

```
End Time: 20260508142953
```

```
Status: A
```


```
EID: 20 Location: /home/ryan/rtest/run_log/tcase51.work/TC51_TEMP
```

History file

- Without maintenance, more disk space will be taken up
 - History file itself
 - Backup image size
 - Database objects (Backup images, log files)
- Cause of poor performance in...
 - Backup
 - Rollforward
 - Archive log

History file: Pruning

PRUNE HISTORY/LOGFILE command:



- Prune (remove history file entry) based on timestamp
- Entries/objects needed for recovery are protected
- AND DELETE deletes log files in archive
 - Configure **auto_del_rec_obj** to delete backup/load copy images

Note: PRUNE LOGFILE is for log files in active path, and is deprecated

History file: Automatic pruning

- Automatically trigger pruning based on number of backup taken and time.

num_db_backup [default 12]

- Number of backup to be kept

rec_his_retentn [default 90]

- Number of days the history are retained for

auto_del_rec_obj [default OFF]

- Whether recovery objects are deleted when history entry is pruned
 - Backup images, load copy images, archived logs

History file: Automatic pruning behaviour

For any value of num_db_backups:

rec_his_retentn	auto_del_rec_obj	History file entry pruned?	Database objects deleted?
-1	OFF	Yes. Entries older than configured value in num_db_backups are pruned.	No
-1	ON	No	No
0	OFF	Yes. All entries older than the latest backup entry are deleted, regardless of num_db_backups value.	No
0	ON	Yes. Entries older than configured value in num_db_backups are pruned.	Yes
> 0	OFF	Yes. Entries older than configured value in rec_his_retentn and num_db_backups are pruned.	No
> 0	ON	Yes. Entries older than configured value in rec_his_retentn and num_db_backups are pruned.	Yes

History file: Automatic pruning - when does it happen?

- Triggered at end of successful backup - which could take some time
- Seen backup stuck at 99% progress?

```
db2sd_20220418173716/idsnx/sqlllib_shared $ db2pd -util
```

```
Database Member 4 -- Active -- Up 29 days 13:06:04 -- Date xxxx
```

Utilities:

Address	ID	Type	State	Invoker	Priority	StartTime
DBName NumPhases CurPhase Description						
0x00007FC2A19FDB00 29		BACKUP	0	0	0	xxxxxxxx
DBNAME 1 1		online db				

Progress:

Address	ID	PhaseNum	CompletedWork	TotalWork
0x00007FC2A1B07708 29		1	11036986251136 bytes	11080589182752 bytes

History file: Asynchronous automatic pruning (NEW)

- Automatic pruning is decoupled from backup operation, and done asynchronously
- New EDU: db2autoprun
 - Can be configured through **DB2_AUTOPRUNE_EDU** [Default ON]
- Trigger for automatic pruning:
 - Periodically, once a day by default
 - Once autoprun is done, sets next prune time to +1 day
 - Can be configured through **DB2_AUTOPRUNE_INTERVAL**
 - Successful backup
 - Sets next prune time to current time

History file: Asynchronous automatic pruning monitoring

- New db2pd option: -autohistprune

```
SQLUH_AUTOPRUNE_CB Formatter
Version 12.1.5.0
```

```
apEduID = 777
apEduState = PRUNE_PENDING (3)
apEduInitLatch =
```

```
{
lock      = { 0x00 [ unlocked ] }
fighterCount = 0
identity  = SQLUH_AUTOPRUNE_CB::apEduInitLatch (1180)
}
```

```
apIntervalTime = 86400 seconds
apErrorIntervalTime = 300 seconds
apLastPruneRc = 0
```

```
apTotalLastPruneInPrgsTicks = 1.02 seconds
apTotalPruneInPrgsTicks = 1.02 seconds
```

```
apLastPruneTime = 2026-05-11-12.34.34.000000 (1778528074)
```

```
apNextPruneTime = 2026-05-12-12.34.36.000000 (1778614476)
```

```
apNextPruneTimeByBackup = 2026-05-11-12.34.34.000000 (1778528074)
```

```
apNextPruneMember = 0
```

```
apNextPruneMemberStatus = UNKNOWN (-1)
```

```
apWaitingToBeDesignatedPruneMemberSince = Not set (0)
```

```
apMaxDesignatedPruneMemberWaitTime = 30 seconds
```

```
apEduLockAcquired = false
```

```
apEduLockInfo = Not set
```

```
apPruneHistoryCB = Not set
```

```
apAgentCB = 0x0000000203E101A0
```

```
SQLUH_AUTOPRUNE_CB Formatter
Version 12.1.5.0
```

```
apEduID = 777
apEduState = PRUNE_PENDING (3)
apEduInitLatch =
```

```
{
lock      = { 0x00 [ unlocked ] }
fighterCount = 0
identity  = SQLUH_AUTOPRUNE_CB::apEduInitLatch (1180)
}
```

```
apIntervalTime = 86400 seconds
apErrorIntervalTime = 300 seconds
apLastPruneRc = 0
```

```
apTotalLastPruneInPrgsTicks = 0.02 seconds
apTotalPruneInPrgsTicks = 1.00 seconds
```

```
apLastPruneTime = 2026-05-11-12.36.28.000000 (1778528188)
```

```
apNextPruneTime = 2026-05-12-12.36.29.000000 (1778614589)
```

```
apNextPruneTimeByBackup = 2026-05-11-12.36.28.000000 (1778528188)
```

```
apNextPruneMember = 0
```

```
apNextPruneMemberStatus = UNKNOWN (-1)
```

```
apWaitingToBeDesignatedPruneMemberSince = Not set (0)
```

```
apMaxDesignatedPruneMemberWaitTime = 30 seconds
```

```
apEduLockAcquired = false
```

```
apEduLockInfo = Not set
```

```
apPruneHistoryCB = Not set
```

```
apAgentCB = 0x0000000203E101A0
```

History file: Asynchronous automatic pruning details

- Starting 12.1.4
- If prune operation fails, retry will be attempted every 5 minutes
 - Once successful, it goes back to the default/configured interval
- db2autoprune EDU is not started for super exclusive operations
 - Offline backup, restore, rollforward
- db2autoprune EDU exists for every member in pureScale
 - Only one drives autoprune while one idles

History file: Non-integrated backup

Process for non-integrated snapshot backup

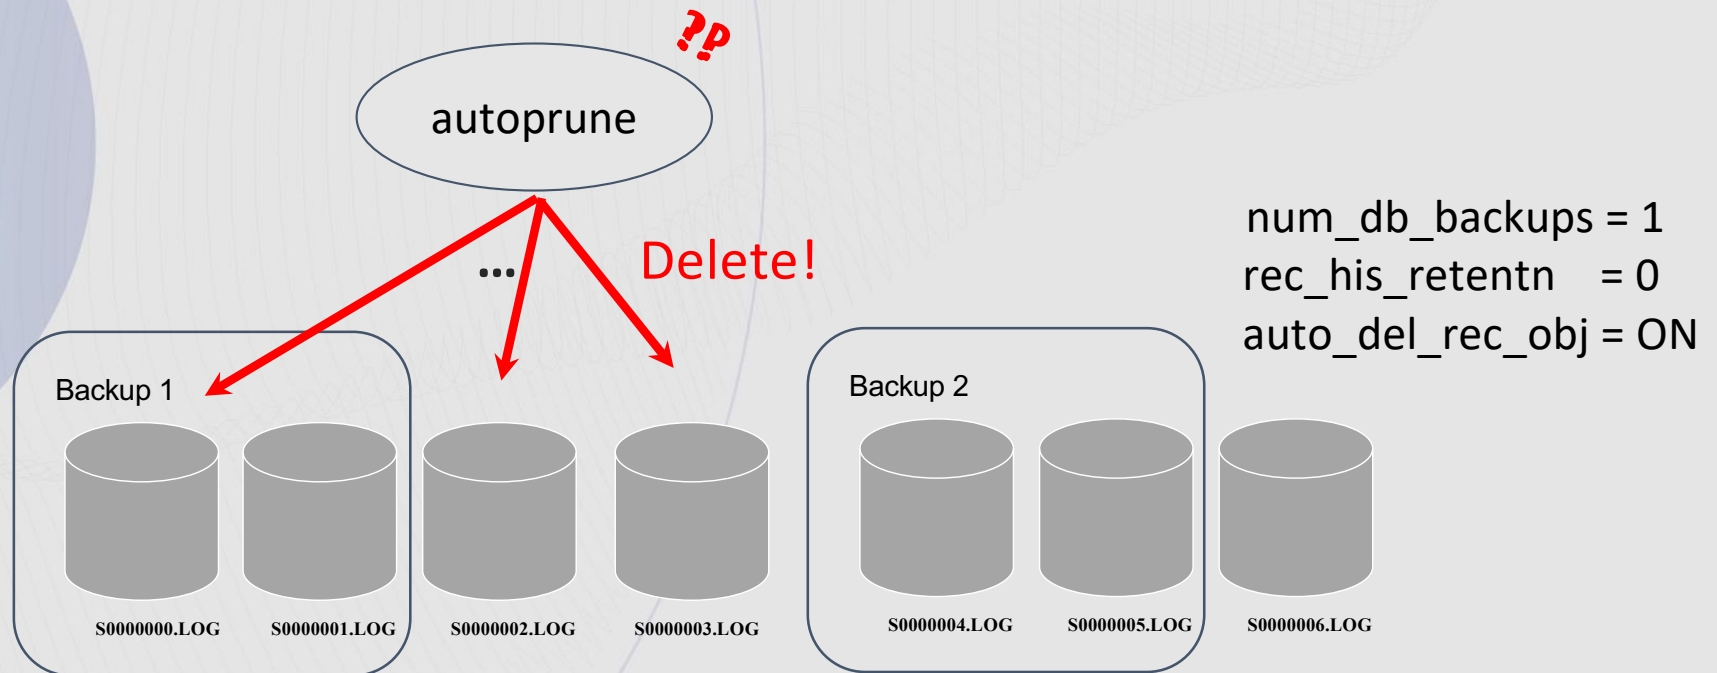
1. db2 SET WRITE SUSPEND FOR DB
2. __copy database and log files using non-Db2 tool__
3. db2 SET WRITE RESUME FOR DB

=> Should it be recorded in history file?

=> How do we know that a backup was taken?

History file: Non-integrated backup

- Non-integrated backup is not recorded in history file
- Backup not in history file => Autoprune is not able to prune anything
 - Including archived logs!



History file: Recording non-integrated backup (NEW)

- User can tell Db2 to record non-integrated backup event
- New RECORD BACKUP option for following commands
 - SET WRITE RESUME
 - RESTART
 - db2inidb
- Also part of APIs under iOptions
 - db2SetWriteForDB
 - db2DatabaseRestart

History file: Recording non-integrated backup (NEW)

- History entry: Op value “S” for non-integrated snapshot backup

```
Op Obj Timestamp+Sequence Type Dev Earliest Log Current Log Backup ID
-- -- -
```

S	D	20260508142954	N		S0000032.LOG	S0000032.LOG	
---	---	----------------	---	--	--------------	--------------	--

```
-----
```

Contains 3 tablespace(s):

```
00001 SYSCATSPACE
00002 TEMPSPACE1
00003 USERSPACE1
-----
```

```
Comment: SET WRITE SUSPEND INCLUDE LOGS
Start Time: 20260508142954
End Time: 20260508142954
Status: A
-----
```

```
EID: 22
```

History file: Recording non-integrated backup - details

- Starting 12.1.5
- Treated as backup entry
- Used to manage archive log files by prune/autoprun
- Limitation: unlike integrated backup, history file entry does not have image location
 - Does not delete backup image on prune
 - Cannot be used as part of RECOVER/REBUILD/RESTORE INCREMENTAL AUTOMATIC command
- Limitation: history entry is tied to SWS command
 - Non-integrated backup taken offline is NOT recorded

History file: Recording non-integrated backup scenarios

- Record successful non-integrated backup
 1. db2 SET WRITE SUSPEND FOR DB
 2. __copy database and log files using non-Db2 tool__
 3. db2 SET WRITE RESUME FOR DB **RECORD BACKUP**
- Non-integrated backup failed and resuming use of database
 1. db2 SET WRITE SUSPEND FOR DB
 2. __copy database and log files using non-Db2 tool__
 3. db2 SET WRITE RESUME FOR DB

History file: Recording non-integrated backup scenarios

- Backup is successful, but Db2 instance crashes before resuming
 1. db2 SET WRITE SUSPEND FOR DB
 2. __copy database and log files using non-Db2 tool__
 3. __Db2 instance crashes__
 4. db2start
 5. db2 RESTART DB <DBNAME> WRITE RESUME **RECORD BACKUP**

History file: Recording non-integrated backup scenarios

- Take backup, and use it to seed a test system on a different host

Host 1:

1. db2 SET WRITE SUSPEND FOR DB
2. __copy database and log files using non-Db2 tool__
3. db2 SET WRITE RESUME FOR DB **RECORD BACKUP**


Host 2:

1. __restore backup taken above__
2. db2start
3. db2inidb <DBNAME> AS SNAPSHOT **RECORD BACKUP**

History file: PRUNE HISTORY without timestamp (NEW)

- PRUNE HISTORY command can now be issued without timestamp
- Prunes based configured **num_db_backups**, **rec_his_retentn**, and **auto_del_rec_obj**
 - Mimics what autopruner does


```
-- PRUNE --+-- HISTORY --+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     +-- timestamp --+-----+-----+-----+-----+-----+-----+ |
|                                     +-- WITH FORCE OPTION -- + +-- AND DELETE --+ |
+----- LOGFILE PRIOR TO -- log-file-name -----+-----+-----+-----+-----+-----+-----+
```



Questions?

Online backup: Log file behaviour

- Transaction logs generated during online backup must be included in the backup image
- Including the log happens at the end of backup



↳ Logs are archived once it's full

↳ Once archived, log in active path can be deleted/reused at any time

↳ This behaviour does not change during online backup

↳ When backup needs to include logs, logs not found locally need to be retrieved

Online backup: Save required logs locally (NEW)

- Logs required by online backup is no longer deleted/reused
- The disk space allowed for the saved log is part of **LOG_DISK_CAP**
 - Logs for online backup is not the only user of LOG_DISK_CAP space
 - Default is 0
- This is 'best effort' for backup performance
 - If there is not enough disk space or space defined by LOG_DISK_CAP, logs will still be removed and retrieved when necessary


NOTE: For pureScale, only the member doing backup will save the logs

Side note: LOG_DISK_CAP

- Intended to be “future” of log management
- Currently, there are many log types and locations they can be in...how much disk space is required?
 - Active/cached/retrieved/extraction logs
 - LOGPRIMARY/LOGSECOND/LOGFILSIZ
 - Active/mirror log path, retrieve subdirectory
 - Overflow log path
 - Failarchpath
- LOG_DISK_CAP: Let us know how much disk space you can give us to work with, we will figure out and manage the rest

Online backup: upcoming

- “Including the log happens at the end of backup”
 - ...does it really need to wait until the end?
- Future release (soon™) will include logs as they are full and closed
 - Logs does not need to be saved for the duration of online backup, less disk space taken up



Questions?

Topology Change: Timeline from last year

- 12.1.0 - Drop member no longer requires backup
Restore and rollforward from backup images with less members
- 12.1.2 - Upgrade and catalog database support with topology difference
- 12.1.3 - Restore of backup images with diff topology from previous release
- Future - Archive logs for deactivated or dropped member
Support drop member in HADR
Online drop

Topology Change: Timeline from last year

- 12.1.0 - Drop member no longer requires backup
Restore and rollforward from backup images with less members
- 12.1.2 - Upgrade and catalog database support with topology difference
- 12.1.3 - Restore of backup images with diff topology from previous release
- Future - Archive logs for deactivated or dropped member
Support drop member in HADR
Online drop

Upgrade/Restore of image from older release with topology change

- Previously, topology change was not allowed with operations across Db2 versions
- Topology change is supported across versions
 - UPGRADE (12.1.2)
 - CATALOG (12.1.2)
 - db2ckupgrade (12.1.2)
 - RESTORE (12.1.3)

NOTE: If member is dropped involving pre-12.1 database, database will
BACKUP PENDING post UPGRADE

NOTE 2: non-pS to pS change is not supported due to db2checkSD restriction

HADR: Old process in HADR to change topology

Add:

- Add member to Standby
- Add member to Primary
- Activate or connect

↳ Add member log record shipped to standby and replayed

Drop:

- **STOP HADR**
- Stop instance
- Drop member from primary
- Drop database on standby
- Drop member from standby
- **Backup database on primary**
- **Re-initialize standby**

HADR: NEW* process in HADR to change topology

Add:

- Add member to Standby
- Add member to Primary
- Activate or connect

↳ Add member log record shipped to standby and replayed

Drop:

- Stop instance
- Drop member from standby
- Drop member from primary
- Start instance
- START HADR

↳ Drop member log record shipped to standby and replayed

↳ *Actually, there's not even a need for this process now!


HADR: Mixed topology

- Starting 12.1.3, primary and standby can have different number of members
 - Must have at least 1 common member
 - pureScale - non-pureScale scenario/restriction is not changed
- Add or drop member on primary and/or standby at any time
 - If replay member is dropped, standby is bounced first
 - Including preferred replay member

```
ADM12521W The preferred replay member "<member-id>" no longer exists
in the database topology.
```


HADR: Mixed topology, main use case

- X number of members on primary, 1 member on standby



HADR: Mixed topology, what about takeover?


- If unexpected takeover happens:



- Option 1: Fix original primary, and takeover
- Option 2: Add members to new primary. Optionally drop members from standby at any time

HADR: Mixed topology, what about takeover?


- If takeover is planned, topology can be prepped



- Add two members to standby
- Takeover
- Drop two members from (new) standby

HADR: upcoming

- More than 3 standby (v12.1.5)
- Backup on standby (soon™)
 - Option to take backup on standby instead when backup is issued on primary



Questions?



Call for Sessions

Submission deadline: [May 22](#)

IBM TechXchange 2026

October 26–29 | Atlanta, GA

Clients, Partners & IBMers are eligible to submit sessions

Share your expertise with thousands of technical professionals!

What we are looking for

- Real-life use cases and implementation stories
- Technical deep dives (non-promotional)
- Demonstrated expertise with IBM products
- Sessions aligned to 2026 priorities: enterprise AI applications, key products, technical skills

Session types

- A compelling title and concise, clear abstract



Technology Breakout (45 minutes)

In-depth technical session featuring architectures, demos, best practices, client/partner stories



Tech Talk (20 minutes)

Fast-paced spotlight on use cases, technical trends, and special-topic insights.

Client & partner submissions only for Tech Talks.

Technology areas

AI Applications	Data Management	Infrastructure
Cloud & Platform Engineering		
Security, Risk & Compliance		

Special perks

- A complimentary full conference pass (that's a \$1599 USD value!)
- Speaker training and content support
- Social tools to promote your session

Submission deadline: [May 22](#)

Get started

