Thinking about your availability and recovery strategy shouldn't give you sleepless nights. Db2 understands the importance of keeping your database online 24x7 and with Db2 11.1 / 11.5 and its continuous delivery of new features and enhancements, you are on the right path of achieving such goals. This presentation will give you a detailed look into some of the new and upcoming features that have been rolled out in the area of availability and recovery. Learn about some of the new and upcoming features that we are sure will make managing your databases easier and in turn make your databases more available to service your every day demands.

Objectives

Describe and discuss the following:

- Overview of continuous delivery
- Discuss new features that have been delivered in Db2 11.1 and 11.5
- Discuss the current top customer feature requests
- Discuss what is next in the area of availability and recovery

Safe Harbor Statement

Copyright © IBM Corporation 2019. 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 I BM. 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

3

© IBM Corporation 2019. All Rights Reserved.

The information contained in this publication is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this publication, it is provided AS IS without warranty of any kind, express or implied. In addition, this information is based on IBM's current product plans and strategy, which are subject to change by IBM without notice. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this publication or any other materials. Nothing contained in this publication 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 the applicable license agreement governing the use of IBM software.

References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates. Product release dates and/or capabilities referenced in this presentation may change at any time at IBM's sole discretion based on market opportunities or other factors, and are not intended to be a commitment to future product or feature availability in any way. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

IBM, and the IBM logo, are trademarks of International Business Machines Corporation in the United States, other countries, or both.

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

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel Centrino, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

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

Agenda

- **Db2 Versioning and Common Code Engine**
- Refresher – What's New in Db2 11.1
- What's New in Db2 11.5 GA
  - Support for Storage Devices that use 4KB Sector Sizes
  - pureScale Currently Committed Semantics Across Members
  - Reduced Logging
  - Advanced Log Space Management
- Coming Soon – Post 11.5 GA
- The Future

Db2 Versioning

- The official Db2 product signature consists of 4 parts and has the format **VV.RR.MM.FF** where:
  - **VV** = Version number
  - **RR** = Release number
  - **MM** = Modification number
  - **FF** = Fix pack number

9.7.0.11	10.1.0.6
10.5.0.10	11.1.4.4
11.5.0.0	

- Until now, the modification value (**MM**) for Db2 LUW has always been **0** (zero)
  - Traditionally, interfaces that return the product signature have supplied only 3 elements – **VV**, **RR**, and **FF**
  - It has not always been obvious when a Fix Pack contains new functionality

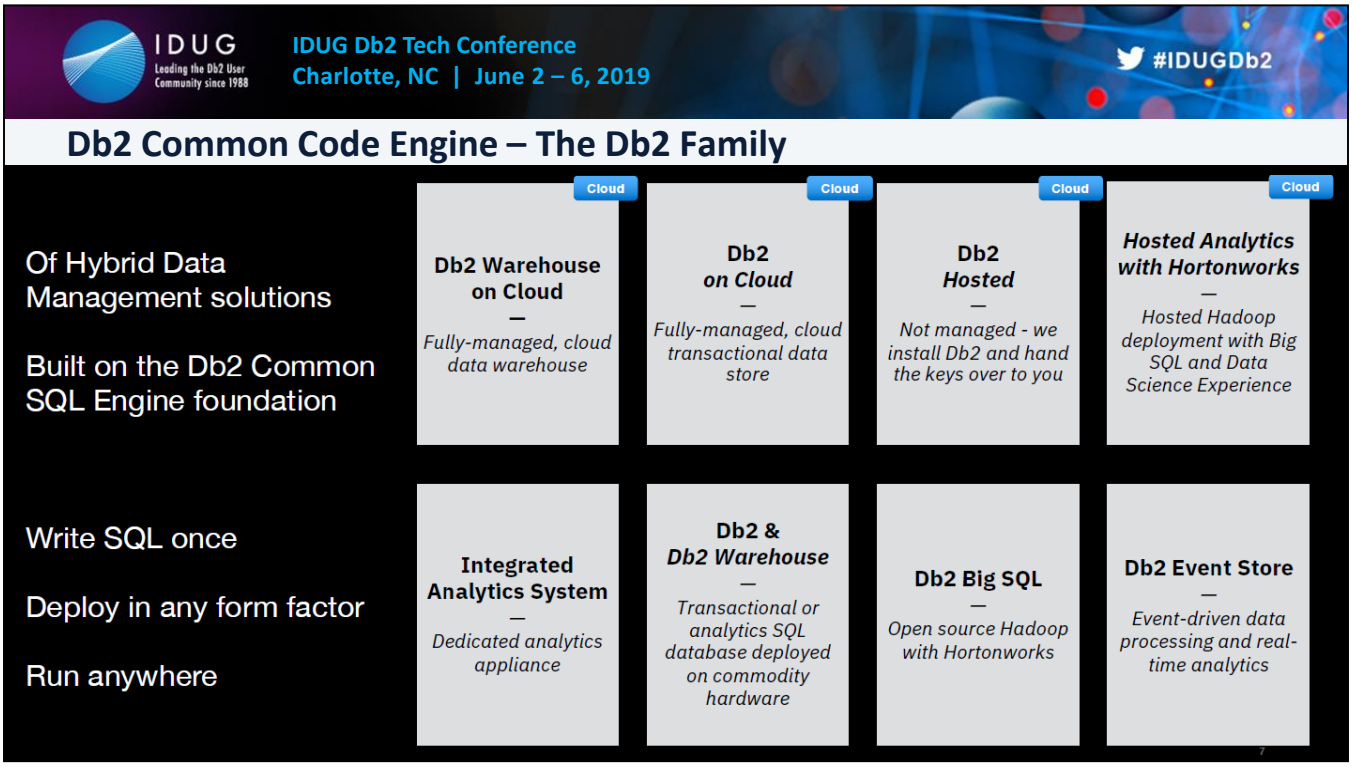

There are a number of products offered under the Db2 umbrella. The primary reason for the differences is how they are packaged and offered.

Db2 offers you the ability to choose from a wide variety of deployment options, from the traditional DIY software, to modern docker images, prepackaged appliances, and an assortment of private and public cloud options. Which option is best for you depends on where you want to deploy it and what level of control you want over the deployment. The traditional Db2 software is available as bother a self-install or as a docker image that you can run alone or in the IBM Cloud private environment. It is also offered in the public cloud in two ways: Db2 hosted, where IBM sets it up and you manage it or as Db2 on Cloud, where IBM installs and manages it all for you. The Db2 Warehouse offering is offered as a docker image or as a managed public cloud offering. And the new IIAS appliance is offered for those who want the simplicity of a prepackaged hardware/software option for on-premise installation.

Db2 Common Code Engine

- **Operational consistency**
  - Security and encryption
  - Integrated workload management
  - Common monitoring interfaces
- **Ecosystem compatibility**
  - One ISV product certification for all platforms
    - Test once, certify all
- **Write once, run anywhere**
  - Reduced feature latency
  - Application portability
  - Common programming model & interfaces

All of these Db2 family of products share the same underlying common SQL engine within the different Db2 form factors. The same Db2 code base provides the underlying database technology used by the Db2 software product and its additional form factors as well as the Db2 Warehouse offering and its additional ones. The Db2 Big SQL product also is embraced by this strategy.

It is the existence of this common layer that enables Db2 to support the “Write once, run anywhere” objective.

Agenda

- Db2 Versioning and Common Code Engine
- **Refresher – What's New in Db2 11.1**
- What's New in Db2 11.5 GA
  - Support for Storage Devices that use 4KB Sector Sizes
  - pureScale Currently Committed Semantics Across Members
  - Reduced Logging
  - Advanced Log Space Management
- Coming Soon – Post 11.5 GA
- The Future

Refresher – What's New in Db2 11.1 – Page 1 of 8

11.1.0.0

- **Fast pre-allocation for log file creation and resize**
  - DB2_USE_FAST_LOG_PREALLOCATION
- **Backup database has a new option**
  - NO TABLESPACE
- **pureScale HADR sync/nearsync support**
- **(non-pureScale) HADR upgrade with no standby re-initialization from 10.5fp7**
- **Recovery through migration from 10.5fp7**
- **Db2 backup and log compression on POWER 7+/ 8 using NX842**
  - AIX only
  - compress comprlib libdb2nx842.a
  - DB2_BCKP_COMPRESSION=NX842

Refresher – What's New in Db2 11.1 – Page 2 of 8

11.1.1.1

- **Increase of limit on LOGFILSIZ (to 64GB)**
  - Maximum theoretical size is 256 log files * 64GB = 16TB
- **(pureScale) HADR upgrade with no standby re-initialization from 10.5fp9**

Refresher – What's New in Db2 11.1 – Page 3 of 8

11.1.2.2

- **Databases can now be configured to allow connectivity during crash recovery**
  - DB2_ONLINERECOVERY
- **Database restore rebuild supported in pureScale**
- **An HADR Standby tablespace can now be recovered without a full database re-initialization**
- **New STANDBY_TABLESPACE_ERROR flag for HADR_FLAGS monitor element**
- **Databases can now be configured to avoid lock-escalation**
  - DB2_AVOID_LOCK_ESCALATION
- **Faster pureScale member crash recovery**
- **Crash recovery and rollforward replay performance improvements**

Refresher – What's New in Db2 11.1 – Page 4 of 8

11.1.3.3

- **HADR integrity checking of transaction log data during network transmission between the primary and standby servers is improved**
  - When an integrity check failure is detected, seamless retransmission of the log data is performed to auto-correct the condition, with no user intervention required
- **HADR Reads On Standby (ROS) diagnostic improvements allow for easier identification of operations which cause replay-only windows**
  - DB2_HADR_REPLAY_ONLY_WINDOW_DIAGLEVEL
- **The archival of log files using VENDOR or TSM methods can now be configured with a timeout on Unix environments**
  - LOGARCHOPT1/2: --VENDOR_ARCHIVE_TIMEOUT

Refresher – What's New in Db2 11.1 – Page 5 of 8

11.1.3.3 (cont'd)

- **SSL support for the transmission of transaction log data between the HADR primary and standby database servers on all platforms, excluding pureScale**
- **CREATE INDEX operations in a Db2 pureScale environment can now allow concurrent write access to the table**
  - DB2_INDEX_CREATE_ALLOW_WRITE

Refresher – What's New in Db2 11.1 – Page 6 of 8

11.1.4.4

- **HADR Reads on Standby – the availability of the standby database to perform SQL queries is significantly enhanced through the avoidance of the replay-only window**
  - DB2_HADR_ROS_AVOID_REPLAY_ONLY_WINDOW
- **For UNIX databases configured with a mirrored log path a potential performance improvement can be achieved by writing log data to both files asynchronously in parallel**
  - DB2_USE_ASYNC_FOR_MIRRORLOG
- **Monitoring the progress and performance of backup and restore operations can be achieved using the new db2pd -barstats option**

Refresher – What's New in Db2 11.1 – Page 7 of 8

11.1.4.4 (cont'd)

- **The performance of online backup operations can be improved through a reduction in bufferpool page flushing**
  - DB2_REDUCE_FLUSHING_DURING_BACKUP
- **Rollback performance improvements using buffered I/O when reading transaction log file data**
  - Internal tests show 3x improvement
  - DB2_USE_BUFFERED_READ_FOR_ACTIVE_LOG
- **A new db2fmtlog tool can be used to extract and display information from transaction log files for informational and PD analysis**

Refresher – What's New in Db2 11.1 – Page 8 of 8

11.1.4.4 (cont'd)

- **Reducing the size of table spaces through extent movement (reclaim) in pureScale**
  - ALTER TABLESPACE with REDUCE MAX or LOWER HIGH WATER MARK
  - DB2_ENABLE_PS_EXTENT_RECLAIM
- **Support for storage devices that use 4KB sector sizes**
  - Tech Preview (not for production use)
  - DB2_4K_DEVICE_SUPPORT

Agenda

- Db2 Versioning and Common Code Engine
- Refresher – What's New in Db2 11.1
- **What's New in Db2 11.5 GA**
  - Support for Storage Devices that use 4KB Sector Sizes
  - pureScale Currently Committed Semantics Across Members
  - Reduced Logging
  - Advanced Log Space Management
- Coming Soon – Post 11.5 GA
- The Future

What's New in Db2 11.5 GA – Page 1 of 8

Changed default behavior (ON)

- For databases configured with a mirrored log path a potential performance improvement can be achieved by writing log data to both files asynchronously in parallel
- Rollback performance improvements using buffered I/O when reading transaction log file data
- HADR Reads on Standby – the availability of the standby database to perform SQL queries is significantly enhanced through the avoidance of the replay-only window

What's New in Db2 11.5 GA – Page 2 of 8

Changed default behavior (ON) (cont'd)

- **CREATE INDEX operations in a Db2 pureScale environment can now allow concurrent write access to the table**
- **Reducing the size of table spaces through extent movement (reclaim) in pureScale**
  - ALTER TABLESPACE with REDUCE MAX or LOWER HIGH WATER MARK

What's New in Db2 11.5 GA – Page 3 of 8

New registry variable

- **DB2_HADR_STANDBY_KEEP_UNARCHIVED_LOGS**
  - Controls whether standby will keep log files when the corresponding log files on primary are not archived
  - Set to FALSE on standby, standby will delete log files when the corresponding log files on primary are not archived
  - DEFAULT: TRUE

What's New in Db2 11.5 GA – Page 4 of 8

Changes to database configuration parameters

- **rec_hist_rentn**
  - Current setting 366 days
  - Changed to 90 days for new databases
- **logprimary / logsecond**
  - Increased each from 256 to 4096
    - logprimary + logsecond <= 8192
  - Recoverable databases only
  - Be aware: Db2 instance can only have open 65536 concurrent files
  - Maximum theoretical size is:
    - Circular: 256 log files * 64GB = 16TB
    - Recoverable: 8192 log files * 64GB = 512TB

What's New in Db2 11.5 GA – Page 5 of 8

Changes to database manager configuration parameters

- **HADR now respects SSL_CIPHERSPECS**
  - Specifies the cipher suites that the server allows for incoming connection requests when using SSL protocol
  - Additionally, will now also specify the cipher suites used to communicate between primary and standby

What's New in Db2 11.5 GA – Page 6 of 8

Miscellaneous

- **Deprecated**
  - db2ReadLogNoConn
  - db2ReadLogNoConnInit
  - db2ReadLogNoConnTerm
- **Upgrade**
  - Support from 10.5 and 11.1
    - FYI: 10.5 end of support date April 30, 2020

What's New in Db2 11.5 GA – Page 7 of 8

New features

- **Support for storage devices that use 4KB sector sizes**
  - Fully support in 11.5 GA; off by default
- **pureScale currently committed semantics across members**
- **Reduced database activation times**
  - Tech Preview (not for production use)
  - DB2_DEFER_MEMORY_COMMIT
- **Reduced logging**
  - Reduced undo logging on by default in 11.5 GA
    - Required log space cut in half
  - Reduced redo logging available only in Warehouse installations
    - Up to 95% less logging

What's New in Db2 11.5 GA – Page 8 of 8

New features

- **Advanced log space management**
  - Tech Preview (not for production use)
  - Reduce transaction log full
- **Replication for Db2 Continuous Availability**
  - Tech Preview (not for production use)
  - For details see J06 - Replication for Db2 Continuous Availability
- **Many other pureScale improvements, for details see:**
  - See J04 – Db2 pureScale Enhancements
  - See J13 – Architecture Deep Dive III - pureScale

Agenda

- Db2 Versioning and Common Code Engine
- Refresher – What's New in Db2 11.1
- What's New in Db2 11.5 GA
  - **Support for Storage Devices that use 4KB Sector Sizes**
  - **pureScale Currently Committed Semantics Across Members**
  - **Reduced Logging**
  - **Advanced Log Space Management**
- Coming Soon – Post 11.5 GA
- The Future

Support for Storage Devices that use 4KB Sector Sizes

- **To enable:**
  - DB2_4K_DEVICE_SUPPORT = ON
- **Restrictions and limitations:**
  - DMS Raw containers are not supported
  - Backup images and load copy files will be slightly larger
- **Possible performance penalty when accessing:**
  - LOB data stored on 512-byte sector storage
  - Backup or load copy files created prior to the enablement of 4K device support

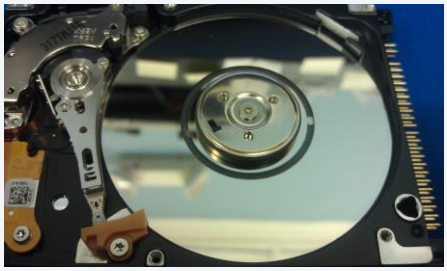

pureScale Currently Committed Semantics Across Members

- **In pureScale environments, the Currently Committed isolation method is supported when:**
  - Application performing row read and application performing row UPDATE or DELETE reside on the same member
  - Application performing row read and application performing row INSERTs reside on any member
- **In current pureScale environments, the Currently Committed isolation method has a significant restriction:**
  - A row-reader application cannot benefit from currently committed semantics when the row is being updated or deleted (locked) by an application on a different member
- **This feature brings completeness to Currently Committed behavior in pureScale environments, by relieving this restriction and allowing a row-reader application in CS isolation to retrieve the currently committed version of a row, even when that row is locked by an application on a different member**

Reduced Logging

- **Applies to:**
  - Column organized tables only
  - Any bulk operation (e.g. upgrade or ingest) which drives insert internally
- **Reduced Undo logging improvements:**
  - Available in 11.5 GA by default
  - Avoid need to reserve log space for undo log records
    - Log space required cut in half
- **Reduced Redo logging improvements:**
  - Available only in Warehouse installations
  - Log meta data changes but skip logging of page contents
  - Similar to "Not Logged Initially" tables but with improved recoverability and concurrency
- **Table contents will be preserved during:**
  - Rollback
  - Crash recovery
  - Database rollforward recovery to end of backup ONLY

Total impact: 95% reduction in required log space

29

Db2 Warehouse columnar deployments typically focused on workload that make heavy use of Extract Load and Transform within the database (ELT) operations.

Need to support massive data volumes in single INSERT statements simply can not spool the necessary log space for such large transactions.

Historically NLI would be leveraged but has a number of properties not friendly to append or update operations – failures take table off line thus backs before/after required or at least recommended

Reduce Logging is a friendly version of NLI. In Db2 Warehouse is enabled by default Implicitly kicks in when a bulk INSERT or UPDATE operation is detected.

Significantly reduces the amount of logging by not logging data pages but still logs all necessary meta data changes to support rollback and crash recovery

While reduce logging can benefit performance due to the reduce logging it is not the main reason for introducing it to Db2 Warehouse.
And it does add additional impact as it requires flush on commit semantics to be enforced.

However it does mean that log based operations like Point-In-Time recovery are no longer available.

Advanced Log Space Management – Problem and Use Case

- **Reduce transaction log full**
  - Often complaints about transactions hitting transaction log full
    - #1 request from many big customers
- **Quick short running transactions running in parallel with:**
  - Long running transactions
    - Low logging rate
    - Log one or two log records and then sit idle for some time
  - LOAD
  - CREATE INDEX
- **Databases must be configured with archive logging**
- **Long running (monster) transactions won't see much benefit**
  - High logging rate
  - Log volume issue
  - Same behavior as before ➔ transaction log full

Advanced Log Space Management – Objective

- **First objective**
  - Log full avoidance with monitoring tools to help manually tune active log space
- **Future objective**
  - More autonomic log space management
    - Today log space is fixed size
    - Move towards Db2 managing log space based on log path file system provided (e.g. LOGPRIMARY, etc. can become automatic)

Advanced Log Space Management – Logging Today

- **LOGPRIMARY * LOGFILSZ**
  - Fixed active log space
- **(LOGPRIMARY + LOGSECOND) * LOGFILSIZ**
  - Maximum active log space available
- **lowtran**
  - First (lowest) log record belonging to oldest open transaction
- **minbuff**
  - Log record of the oldest (minimum) dirty page in buffer pool

---

- **Db2 saves log files from min(lowtran, minbuff) called head extent for rollback/crash recovery**
- **Transaction log full is when lowtran and/or minbuff do not move up**
  - lowtran => open transaction
  - minbuff => flushing slow
- **One way to avoid transaction log full is use infinite logging (LOGSECOND = -1)**
  - Files from head extent and onwards not guaranteed to be in active log path
  - Rollback and crash recovery may have to retrieve log files from archives
    - Major performance pain point

Advanced Log Space Management – Solution – Page 1 of 6

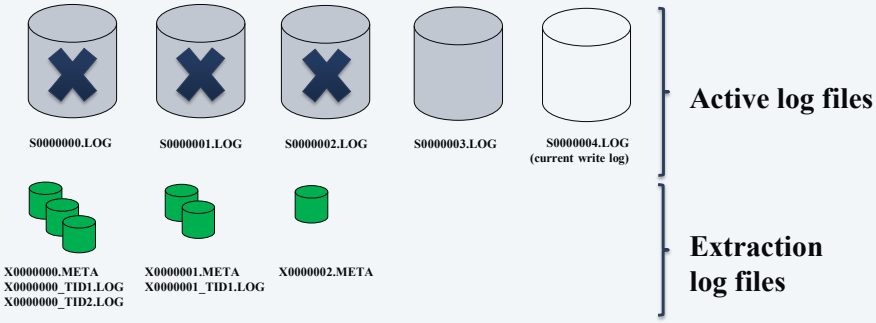

- Extract log records for long running active transactions to separate extraction log files under active log path
- Allows active log files to be closed, archived, and reused, thus avoiding transaction log full

Extraction will kick in when a certain percentage of log space has been consumed and will extract log records starting from head extent ID into separate extraction log files stored in the active log path.

Advanced Log Space Management – Solution – Page 2 of 6

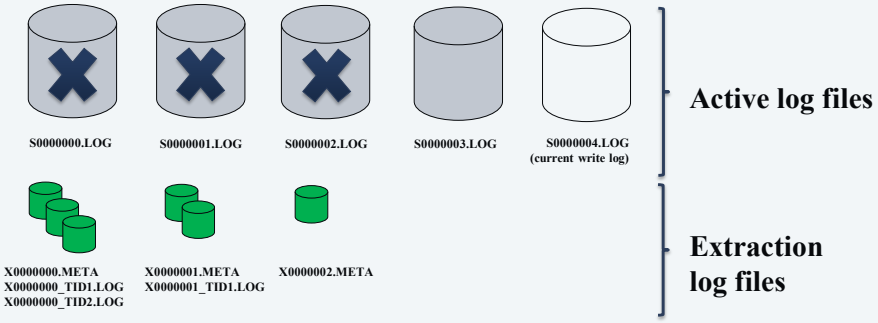

- **New files in active log path:**
  - **X<logFileNum>_TID<tranId>.LOG** - extraction transaction ID (TID) file. Extracted log records for a specific transaction used by rollback and crash recovery. 1 file per log file where log data is extracted for a transaction ID.
  - **X<logFileNum>.TMP** - meta data about extracted logs created during an in progress extraction for an active log file.
  - **X<logFileNum>.META** - meta data about extracted logs created after extraction completes for an active log file.

Advanced Log Space Management – Solution – Page 3 of 6

- **Extraction takes place by new EDU – db2loggx**
- **No to minimal impact to active workloads**
- **Extraction will be throttled based on policies such as:**
  - Disk available
    - Not enough disk space, extraction will idle
  - Log space consumed
    - Log space consumption high, extraction will kick in
  - Producing a benefit
    - No benefit seen, maybe due to monster transaction, extraction will idle
- **Idle extraction means possible transaction log full can occur**

Advanced Log Space Management – Solution – Page 4 of 6

- **An idle extraction scan can happen because:**
  - **Log archiving not healthy**
    - Log data from the active log files that is not archived yet is not extracted
    - Ensure log archiving is healthy and/or a FAILARCHPATH is configured
  - **Buffer pool flushing is slow**
    - Log data from the active log files that is at or above what has been flushed from the buffer pools is not extracted
    - Ensure PAGE_AGE_TRGT_MCR and PAGE_AGE_TRGT_GCR (or SOFTMAX on older database configurations) are set to appropriate values based on your workload throughput
  - **Extraction is slow**
    - It is possible that log writing is faster than log extraction or log extraction has triggered too slowly
  - **Extraction write error**
    - Including disk full

Advanced Log Space Management – Solution – Page 5 of 6

- **Rollback**
  - A line is created to determine whether read from active log files or read from extraction log files
  - Error reading extraction log files will retrieve log data from archives
  - Has shown a performance improvement for a rollback of a single transaction
- **Currently committed**
  - A line is created to determine whether read from active log files or read from extraction log files
  - Error reading extraction log files will resort to lock wait behavior
- **Crash recovery**
  - Use extraction log files for redo and undo
  - Post crash recovery, extraction scan will continue where it left off, so can support indoubt transactions or any deferred undo such as from DB2_ONLINERECOVERY
  - Error reading extraction log files will retrieve log data from archives

Advanced Log Space Management – Solution – Page 6 of 6

- **Set write suspend**
  - Extraction and set write suspend are serialized just like log writing
- **Encryption aware**
  - If database encrypted, extraction log files will be encrypted
- **Monitoring:**
  - MON_GET_TRANSACTION_LOG
  - MON_GET_UNIT_OF_WORK
  - MON_GET_UNIT_OF_WORK_DETAILS
  - db2pd -logs

Advanced Log Space Management – 11.5 GA Restrictions

- **Tech Preview**
  - **Do not use in production**
  - Enable with DB2_ADVANCED_LOG_SPACE_MGMT=ON
- **Databases configured with circular or log retain logging (LOGARCHMETH1/2)**
  - No plan to support
- **Databases configured with mirrored logging (MIRRORLOGPATH)**
  - Future support
- **Databases configured with HADR**
  - Future support
- **Databases in pureScale environments**
  - Future support

Advanced Log Space Management – 11.5 GA Limitations – Page 1 of 2

- **Disk space**
  - Will consume additional disk space to hold extraction log files
  - Should provide extra disk space otherwise extraction will not take place and log full can occur
- **Online backup**
  - Extraction log files will not be included in backup image
    - May need to retrieve active log files
  - Could increase range of log files that needs to be included
    - Larger image sizes
    - Take longer

Advanced Log Space Management – 11.5 GA Limitations – Page 2 of 2

- **Crash recovery**
  - Redo phase will retrieve log files
  - Undo phase will use extraction log files but may not be optimal performance
- **Restore and rollforward**
  - Deletes all extraction log files
  - Will retrieve log files

Agenda

- Db2 Versioning and Common Code Engine
- Refresher – What's New in Db2 11.1
- What's New in Db2 11.5 GA
  - Support for Storage Devices that use 4KB Sector Sizes
  - pureScale Currently Committed Semantics Across Members
  - Reduced Logging
  - Advanced Log Space Management
- **Coming Soon – Post 11.5 GA**
- The Future

Coming Soon - Post 11.5 GA

Advanced Log Space Management Continuation – Page 1 of 2

- **Full crash recovery support**
  - Redo phase make use of extraction log files
  - Redo / undo performance
- **HADR support**
  - Primary has runtime extraction in progress which requires a log space X
  - Standby does not have these same extraction log files, but would like the same log space requirement
  - Standby has the additional challenge of replay falling behind log shipping, so the disk space challenge on standby is more
    - How to manage when extraction alone is not able to relieve the disk space
      - Drop out of PEER or do we hold up transactions on primary?

Coming Soon - Post 11.5 GA

Advanced Log Space Management Continuation – Page 2 of 2

- **Online backup and database rollforward support**
  - Include extraction log files in backup image
  - Database rollforward makes use of extraction log files initially
  - What to do once extraction log files all read but more to replay?
    - Will retrieve log files but what if not enough disk space to hold all retrieved log files needed for undo?
- **pureScale support**
  - Integration with active log file management
    - Runtime and merged log recovery operations (e.g. group crash recovery / database rollforward)
- **Mirror logging support**
  - Extraction log files in both primary and mirror log path
  - Options?
    - Mirror extraction log files in both log paths
    - One version of extraction log files but exists in whatever path is healthy

Coming Soon - Post 11.5 GA

log_disk_cap - Active log space disk capacity configuration parameter

- **Defined in 11.5 GA but not supported until a later time**
- **Allows you to specify the maximum disk capacity for storing transaction log records in the active log path:**
  - Active and extraction log files needed for inflight transactions
  - Inactive log files that have not been archived yet (and not moved to failarchpath)
  - Retrieved log files (if overflowlogpath parameter is not set)
- **logprimary / logsecond used as guidance**
- **The number of files created on disk for logging of inflight transactions might be adjusted based on other consumption**
- **logfilsz is still used to specify the size of the active log files**
- **Primary and mirror log paths should be able to hold this amount**
- **DPF/MPP and pureScale all partitions/members should be able to hold this amount**

Coming Soon - Post 11.5 GA

DPF/MPP Two-Phase Commit (2PC) Indoubt Transaction Resolution

- **Issue:**
  - 2PC indoubt transaction in DPF/MPP not always resolved automatically by Db2
- **Root Cause:**
  - Indoubt resolution is only triggered by primary connection (application connecting to that partition)
  - Attempts to contact other partitions related to transaction, however if instance not up yet on other partitions indoubt is not resolved and left behind
- **Work around:**
  - Issue "db2_all restart database <dbname>“ to make all partitions available
- **Solution:**
  - Implement a new database agent to execute the indoubt resolution (for both primary and secondary connection), and keep the agent to retry, if some partition is not available yet

Coming Soon - Post 11.5 GA

Improved Table Space Error Handling During Recovery – Page 1 of 2

- **HADR standby recovery of load copy images**
  - **Issue:**
    - Load copy image is not available on standby so table space is placed into restore pending
  - **Scenario:**
    - Customer replicates and/or ships images over to standby but there is a time delay, hence failure
  - **Proposed solution:**
    - Build time based retry attempt to restore load copy image
      - Question:
        - Registry variable or database configuration parameter?
        - Default re-try or only when knob set?

Coming Soon - Post 11.5 GA

Improved Table Space Error Handling During Recovery – Page 2 of 2

- **General recovery and placing table spaces into restore pending, rollforward pending or offline state**
  - **Issue:**
    - Crash recovery, database/table space rollforward, HADR standby replay if they have an issue places a table space into restore pending, rollforward pending or offline state
    - Original design was to get as much data available as soon as possible
    - Find out about error much later and can complicate recovery of missing pieces
    - Question:
      - What is more important or varies on recovery type?
  - **Proposed solution:**
    - On such errors fail the recovery operation and avoid setting table space state
      - Problem can be solved
      - Avoid data not available on standby
    - Question:
      - Registry variable or database configuration parameter?
      - Setting based on recovery type?

Agenda

- Db2 Versioning and Common Code Engine
- Refresher – What's New in Db2 11.1
- What's New in Db2 11.5 GA
  - Support for Storage Devices that use 4KB Sector Sizes
  - pureScale Currently Committed Semantics Across Members
  - Reduced Logging
  - Advanced Log Space Management
- Coming Soon – Post 11.5 GA
- **The Future**

The Future – Top RFEs / Ideas

Idea #	Abstract	Votes
DB24LUW-I-97	Make logprimary and logfilsiz online configurable	61
DB24LUW-I-202	Decouple archive log purging from BACKUP utility	38
DB24LUW-I-265	Allow offline utilities such as backup, restore, rollforward to automatically force connections	36
DB24LUW-I-232	DB2 LUW: Intra-tablespace parallelism	34
DB24LUW-I-384	Force LOAD COPY location with HADR	31
DB24LUW-I-138	Restore of online backup taken with older Db2 version	29
DB24LUW-I-93	Parallel log file archival and retrieval	27
DB24LUW-I-208	Ability to keep configuration parameters in sync on HADR standby databases	22
DB24LUW-I-23	Include load copy files in online backup images	21
DB24LUW-I-716	Enable backup on a standby database	21
DB24LUW-I-131	LOAD with COPY=YES and TSM needs similar resilience feature like log archiving with VENDOR_ARCHIVE_TIMEOUT	20
DB24LUW-I-525	Ability to archive standby - backup and restore	20

The Future – Open Conversation – Page 1 of 2

- **Dynamic logprimary / logfilsiz**
  - Question:
    - With Advanced Log Space Management, increase in logprimary / logsecond for recoverable databases and the introduction of log_disk_cap still a (high) need?
- **Continued HADR RoS Improvements**
  - Outstanding issues:
    - When replay conflict with query still forces connection
      - Better to not have to kill the connection
      - Question: Just return some error (e.g. a new reason code in SQL0911 to let application handle (retry))
    - Improvement with RUNSTATS
      - Currently still treat this as a conflict, but customers argue it should not be
      - Question: Willing to accept less optimal plan on standby if the plan is generated via older stats
- **HADR Backup on Standby**
  - Question:
    - Snapshot only support vs. non-snapshot support
    - Requirement differences between single standby and multiple standby

The Future – Open Conversation – Page 2 of 2

- **pureScale RoS**
- **Include load copy images in online backup images**
- **Intra Table Space Parallelism**
  - db2bm working across many table spaces at one time
- **Remote storage support for log archive/retrieve**
- **Auto backup feature**
  - AUTO_MAINT / AUTO_DB_BACKUP
  - RFEs / Ideas:
    - Full support for things like compression/encryption
    - Same auto pruning rule as with auto_del_rec_obj
- **What feature from 11.1 (or earlier) would you want the default changed?**
- **Other ???**

Resources

- IBM Data & AI (formerly Analytics) Ideas (RFEs)
  - https://ibmanalytics.ideas.aha.io/?project=DB24LUW
- Db2 Roadmap
  - http://ibm.biz/AnalyticsRoadmaps
- My Blog - Db2 Availability & Recovery Insider
  - https://www.idug.org/p/bl/et/blogid=703

Community Session – IDUG Charlotte

#IDUGDb2

Date: Thursday, June 6, 2019
Time: 13:15 – 15:30
Location: City Lights Roof Top Bar @ Le Meridian

Join members of the Db2 (Common SQL Engine) offering management and development teams to discuss your requirements – ask any questions you may have – share feedback on our launch - and share your future plans and strategy with IBM. This is an informal session and a unique opportunity to talk to and influence those who are making decisions about the direction to take our IBM Hybrid Data Management offerings.

IDUG
Leading the Db2 User Community since 1988

Food and drinks provided !!

We want your feedback! Take 2 minutes - http://sgiz.mobi/s3/IDUG-2019

54

Please complete your evaluations before leaving.

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, and logging/transaction management.

Please connect with Michael on Twitter at @roecken and LinkedIn.