Trying to figure out the ins and outs of Db2 Log Space Management? Spent too much time trying to figure out what settings are best for your environment? With the introduction of Advanced Log Space Management the time has come to unleash it's potential. Fully supported, the journey has begun in making log space management more autonomic. This presentation will introduce you to some of the concepts and principles behind it all. We will do a deep dive into the technology and transformation behind making log space management more hands off, which in the end will make life simple for you and your organization.

This presentation is valid as of 11.5.9.0.

Safe Harbor Statement

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.

Agenda

- **Advanced Log Space Management**
  - **Overview**
  - Monitoring
  - Future

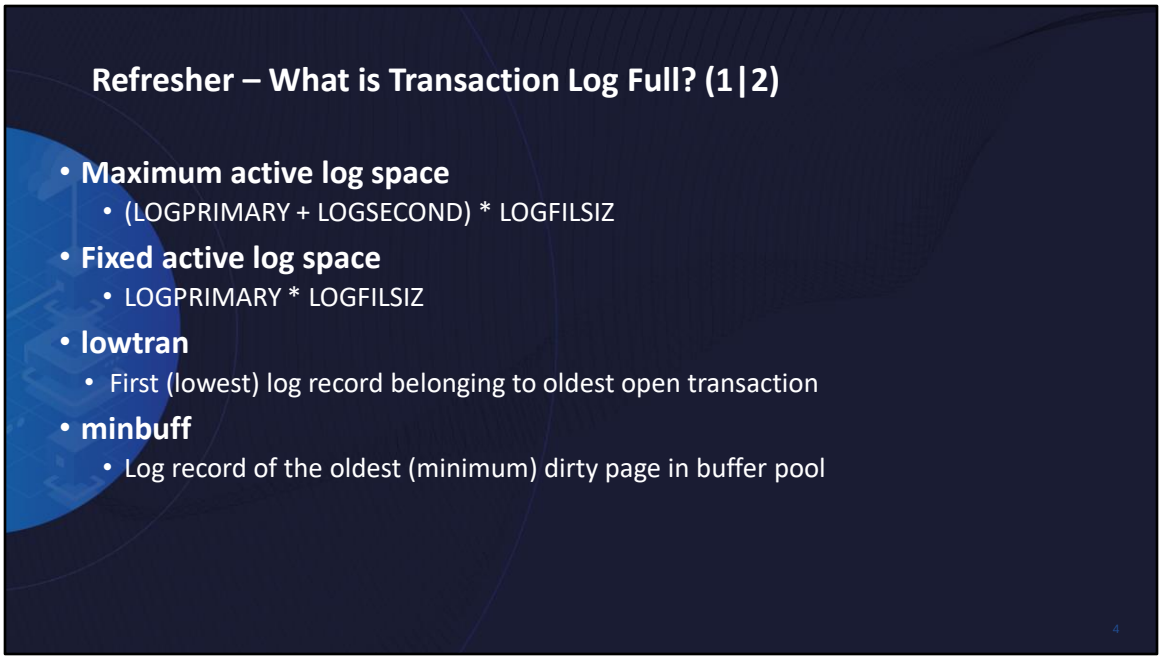

How to solve transaction log for the databases is full (SQL0964C)

- https://www.ibm.com/support/pages/db2-how-solve-transaction-log-databases-full-sql0964c

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: OFF, ON under DB2_WORKLOAD=SAP, 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.

Refresher – What is Transaction Log Full? (2|2)

- Db2 saves log files from min(lowtran, minbuff) called head extent for rollback/crash recovery
- Transaction log full is when Db2 needs to create a new log file above LOGPRIMARY+LOGSECOND but cannot because lowtran and/or minbuff do not move up
  - lowtran => open transaction; minbuff => bufferpool flushing slow

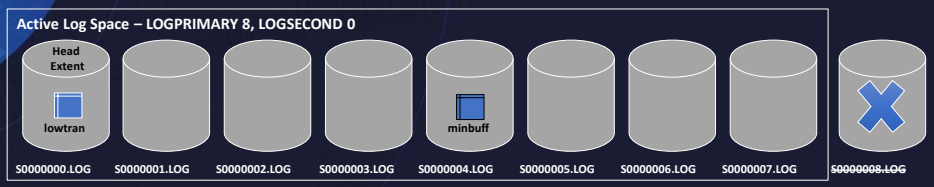

5

How to solve transaction log for the databases is full (SQL0964C)

- https://www.ibm.com/support/pages/db2-how-solve-transaction-log-databases-full-sql0964c

Refresher – Infinite Logging (1|2)

- **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
  - Avoid rogue transactions by using configuration parameters
    - NUM_LOG_SPAN and/or MAX_LOG
  - Rollback and crash recovery may have to retrieve log files from archives
    - Major performance pain point

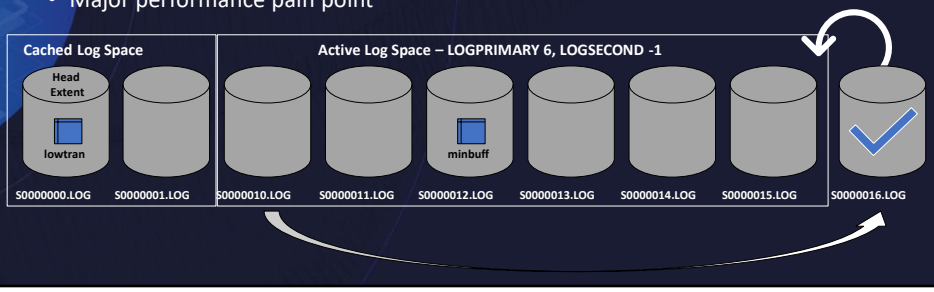

logsecond - Number of secondary log files configuration parameter

- https://www.ibm.com/docs/en/db2/11.5?topic=parameters-logsecond-number-secondary-log-files

Refresher – Infinite Logging (2|2)

- Online backup has to include many more log files
  - Increased image size
  - Longer running backups
- Db2 caches some files (up to 8) above active log space to mitigate need to retrieve log files from archives

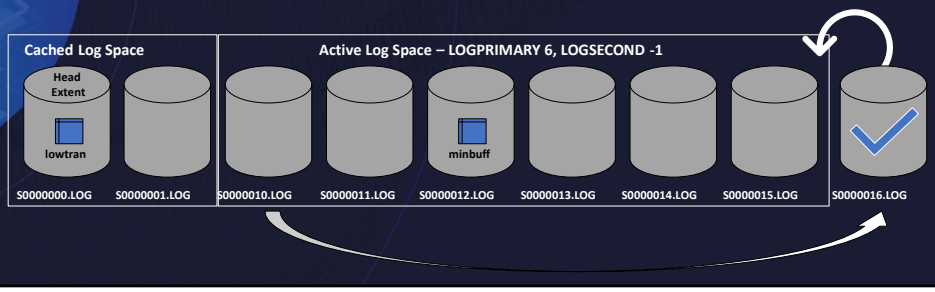

logsecond - Number of secondary log files configuration parameter

- https://www.ibm.com/docs/en/db2/11.5?topic=parameters-logsecond-number-secondary-log-files

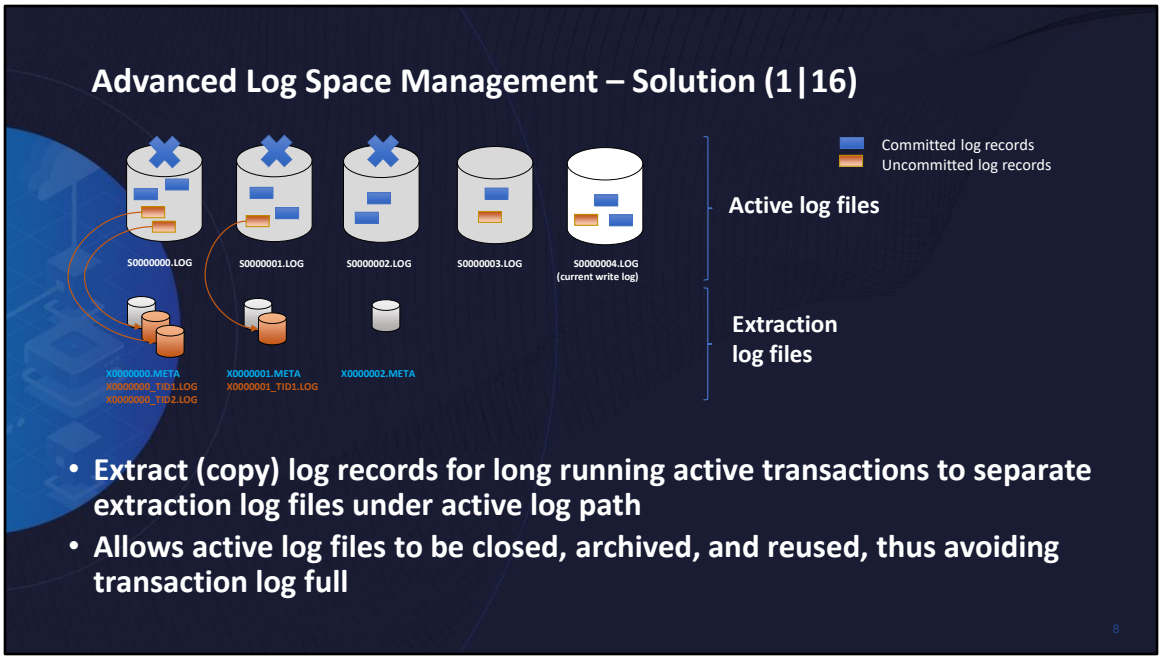

Extraction will kick in when a certain percentage of log space has been consumed and will extract log records starting from head extent ID into separate extraction log files stored in the active log path.

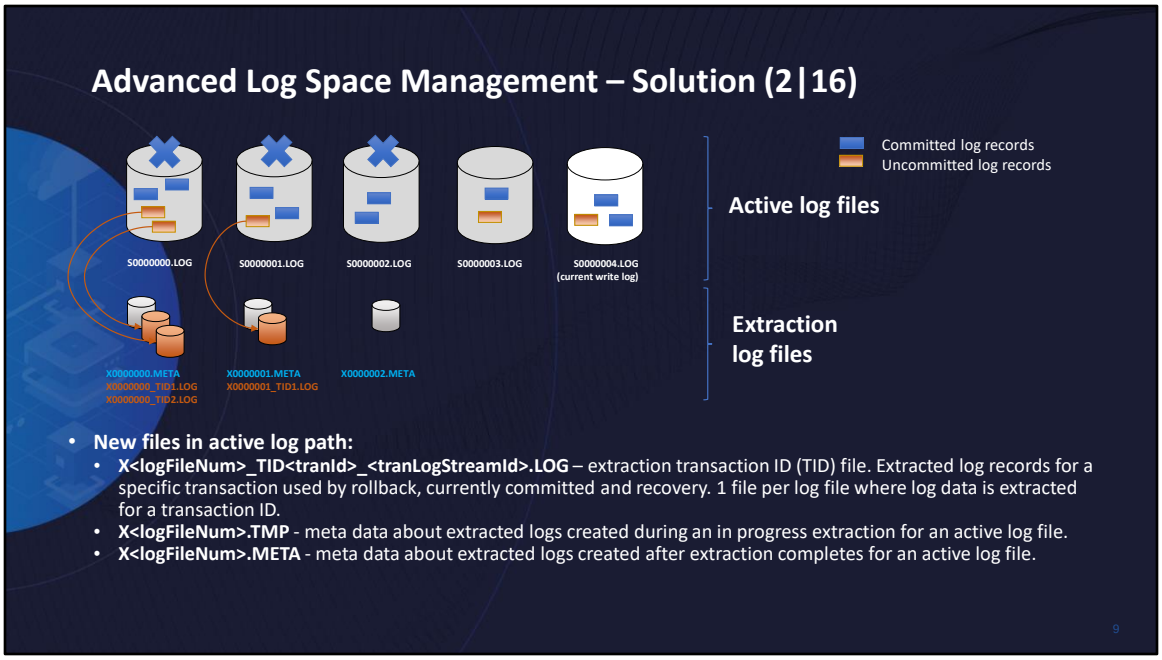

Extraction log files are located in the active log path denoted by the **logpath** database configuration parameter. An extraction file can be thought of as a contained log file, including data for only a single transaction, one that has a potential to cause a transaction log full error. The parent active log file typically contains many more log records, including log records for other transactions that have already ended.

Extraction log files use disk space available outside of the configured upper limit of the number of log files, **logprimary** and **logsecond**. ALSM is designed not to interfere with the log space configured for the database. Extraction log files do not get managed by log archiving. They will only be removed from the active log path once the transaction contained in the particular extraction log file finishes.

There are three types of extraction log files:

•**X<*log file number*>_TID<*tid*>_<*tidLogStreamId*>.LOG**An extraction transaction ID (TID) file. This file contains extracted log records from the log file <*log file number*> for a single transaction identified by <*tid*>_<*tidLogStreamId*>. If available, the TID file is used for rollback, currently committed and all recovery purposes, including crash recovery and database rollforward. There is one TID file per active log file per transaction ID.

•**X<*log file number*>.TMP**A temporary metadata file describing transactions and log records extracted from the log file <*log file number*>. This file gets created while active extraction is in progress and has not yet completed.

•**X<*log file number*>.META**A permanent metadata file describing transactions and log records extracted from the log file <*log file number*> . This file gets created by renaming the aforementioned TMP file after extraction has finished processing the current log file.

Advanced Log Space Management – Solution (3|16)

Timeline:

- 11.5.0.0 – Technical Preview (not for production use)
- 11.5.4.0 – Production support
  - no mirror log support
  - no HADR support
  - no online backup support
  - no pureScale support
- 11.5.5.0 – Basic mirror log support
- 11.5.6.0 – HADR and online backup support
- 

12.1.0.0 – On by default under DB2_WORKLOAD=SAP

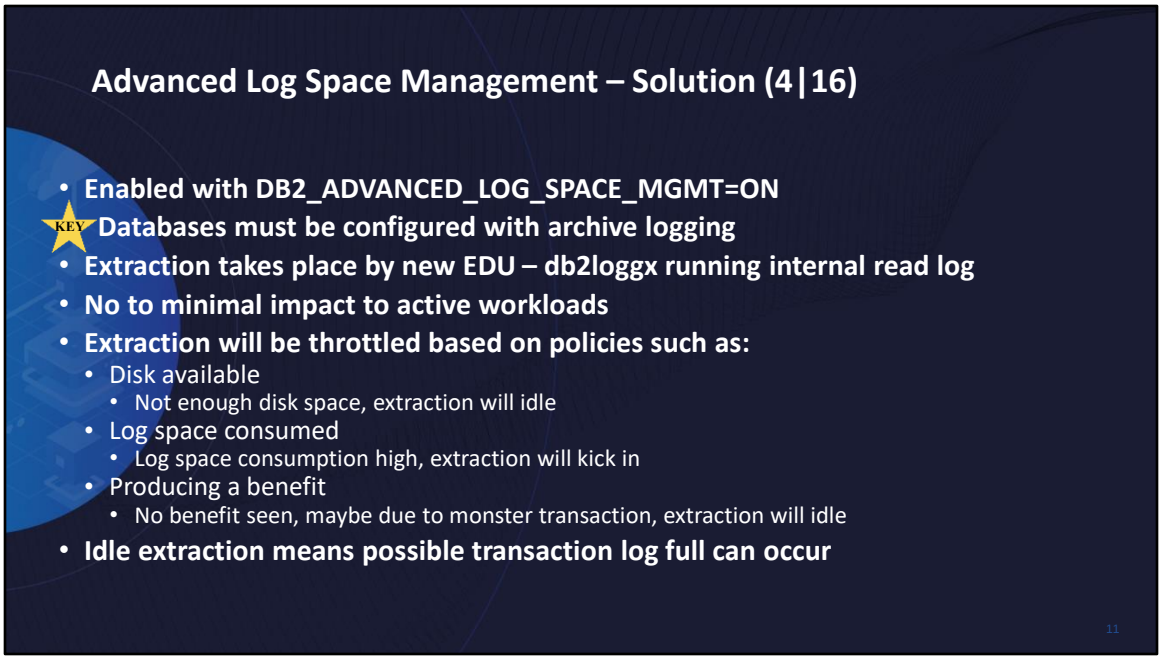

Advanced Log Space Management

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

All ALSM extraction logic is contained within a single-threaded engine dispatchable unit (EDU), **db2loggx**. If configured to run, the log extraction EDU is started as a background agent during database activation or during a user or utility connection such as crash recovery and database rollforward.

Once started, **db2loggx** takes periodical samples of the active log path and other relevant logging parameters to determine if extraction should begin. Due to the sampling, the log extraction EDU consumes a small amount of system resources even when there is no extraction running. However, the impact of this overhead on performance is negligible.

When the conditions permit for extraction, **db2loggx** initiates a single log stream scan through all relevant log files. Log records from the log stream scan are used to populate the log extraction files as necessary. When the extraction is finished, the log stream scan is closed and the EDU returns back to the default sampling state. The log extraction EDU continues to exist until the database is deactivated. When **db2loggx** is terminated, the deactivation performs a full clean-up of memory resources and extraction files that are not needed.

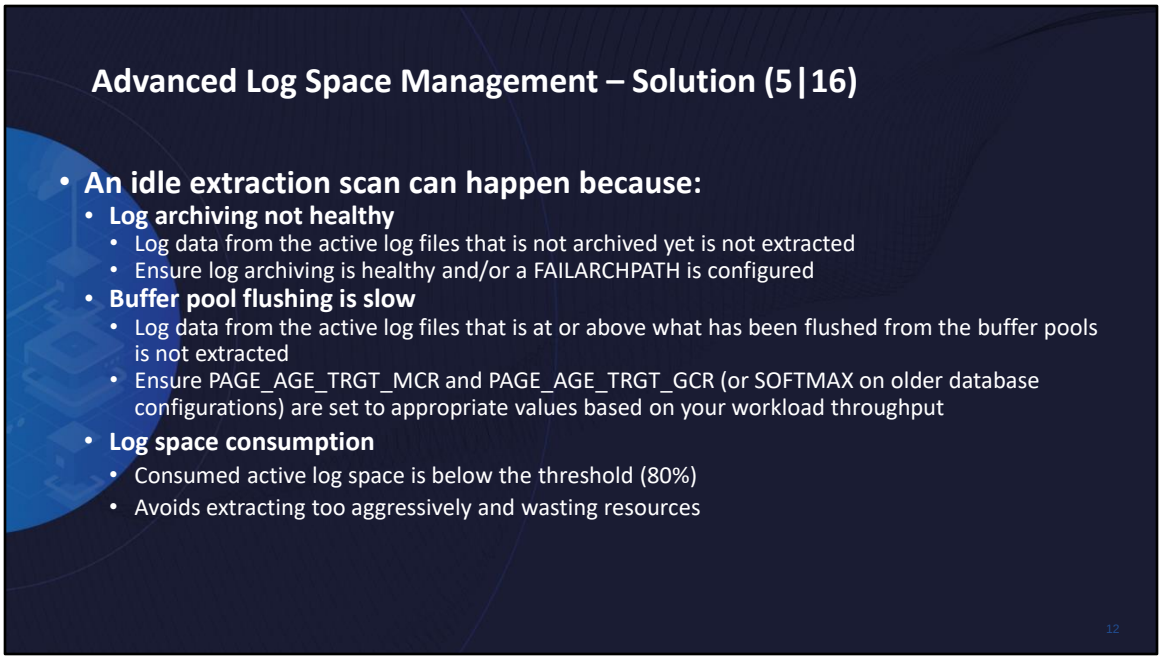

Advanced Log Space Management

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

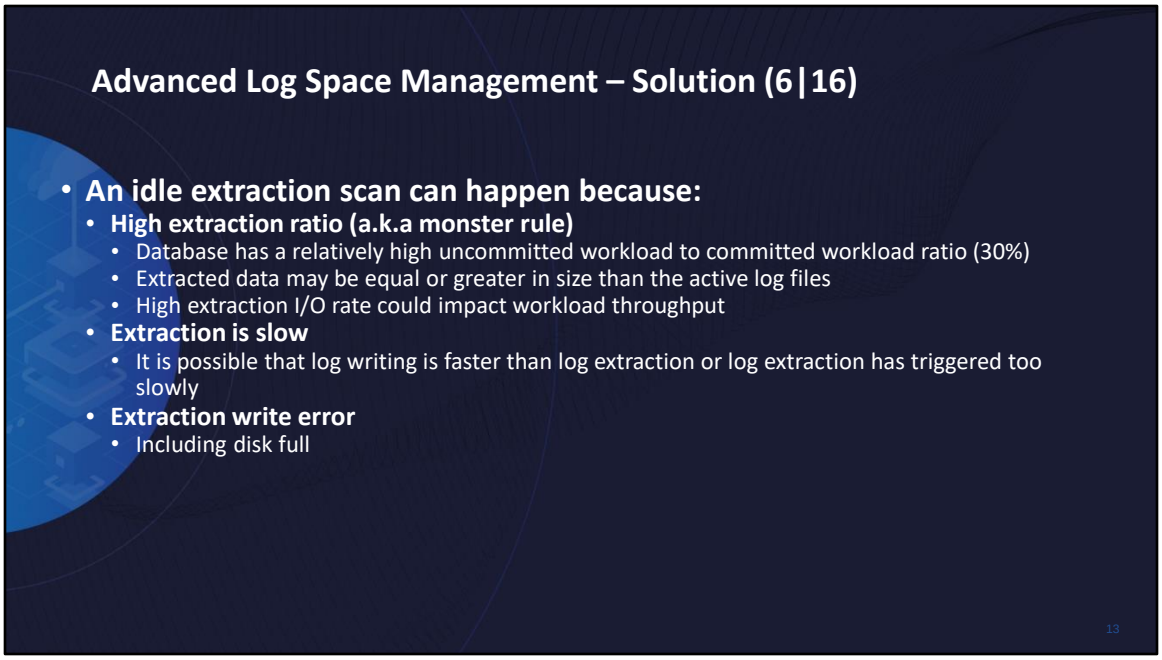

Advanced Log Space Management

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

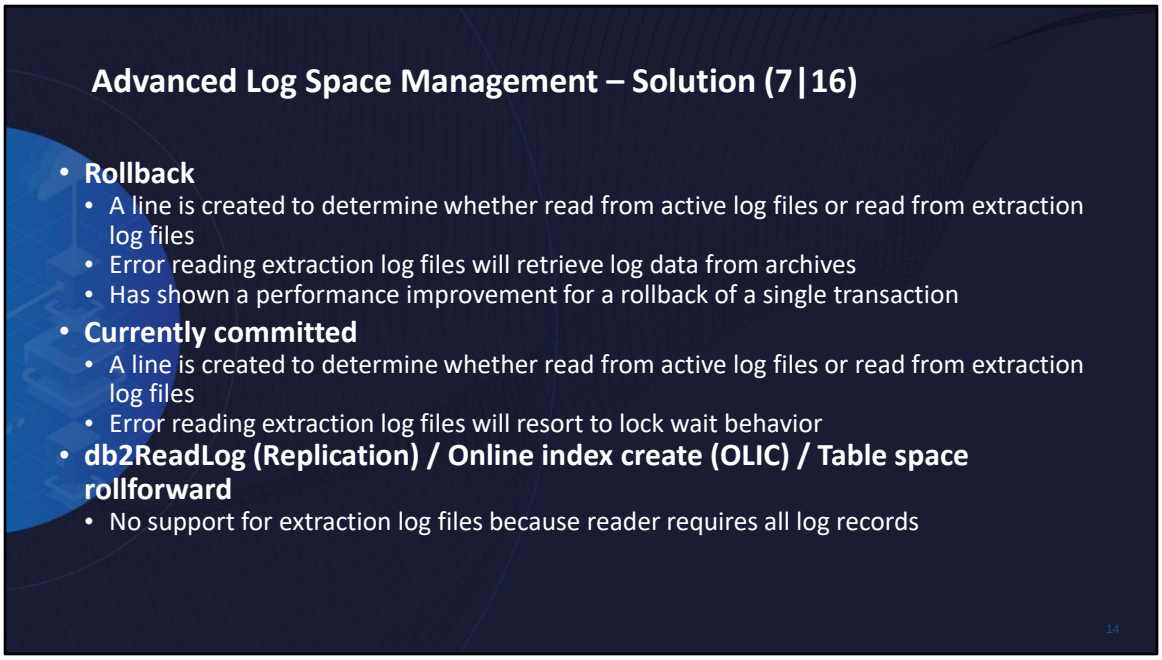

Advanced Log Space Management

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

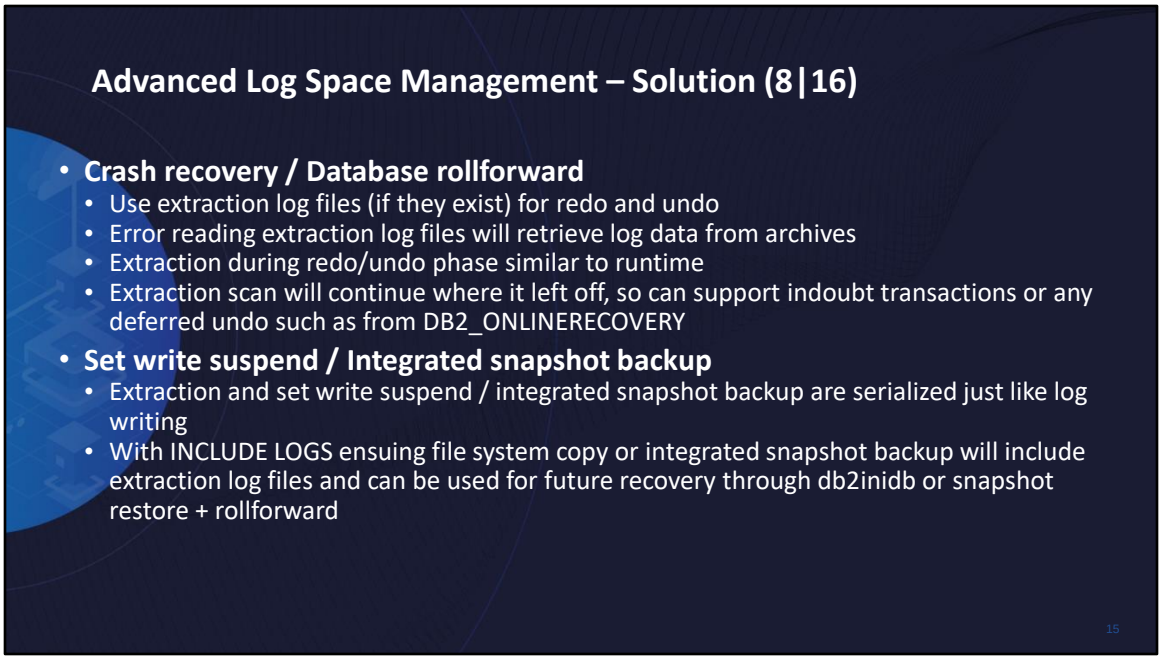

Advanced Log Space Management

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

Advanced Log Space Management

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

MON_GET_TRANSACTION_LOG table function - Get log information

- https://www.ibm.com/docs/en/db2/11.5?topic=mpf-mon-get-transaction-log-table-function-get-log-information

num_log_span - Number log span configuration parameter

- https://www.ibm.com/docs/en/db2/11.5?topic=parameters-num-log-span-number-log-span
- Addresses long running or idle transactions with low logging volumes

max_log - Maximum log per transaction configuration parameter

- https://www.ibm.com/docs/en/db2/11.5?topic=parameters-max-log-maximum-log-per-transaction
- Addresses transactions with high logging volumes (consuming high amounts of the log space e.g. monster)

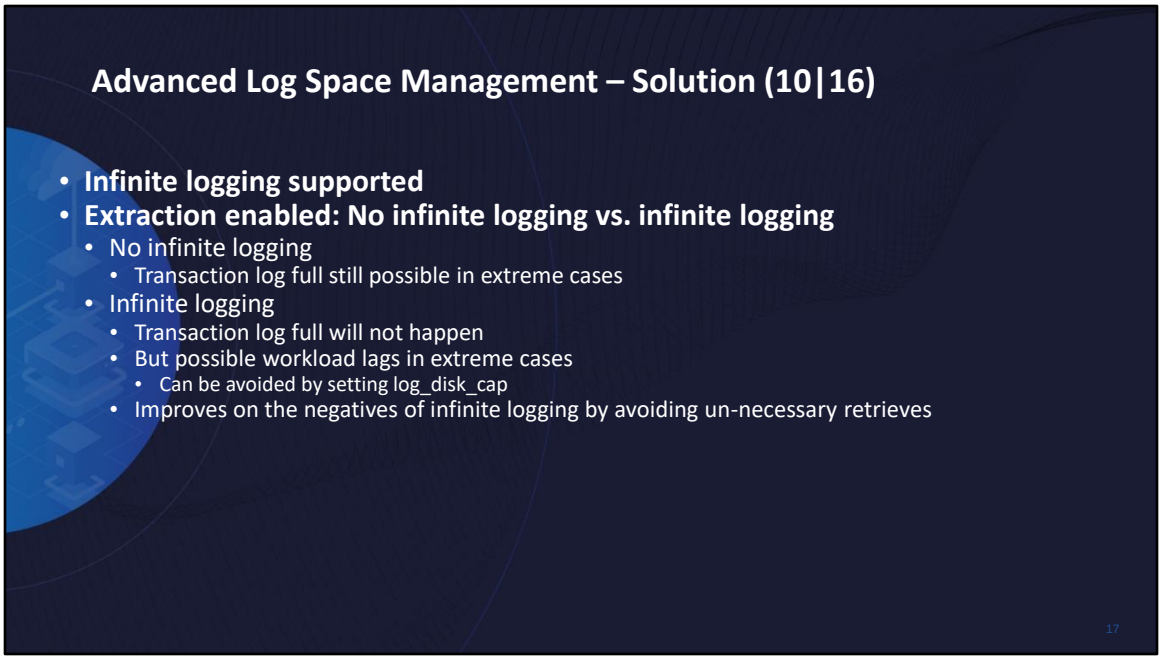

Advanced Log Space Management

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

logsecond - Number of secondary log files configuration parameter

- https://www.ibm.com/docs/en/db2/11.5?topic=parameters-logsecond-number-secondary-log-files

Advanced Log Space Management

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

mirrorlogpath - Mirror log path configuration parameter

- https://www.ibm.com/docs/en/db2/11.5?topic=parameters-mirrorlogpath-mirror-log-path

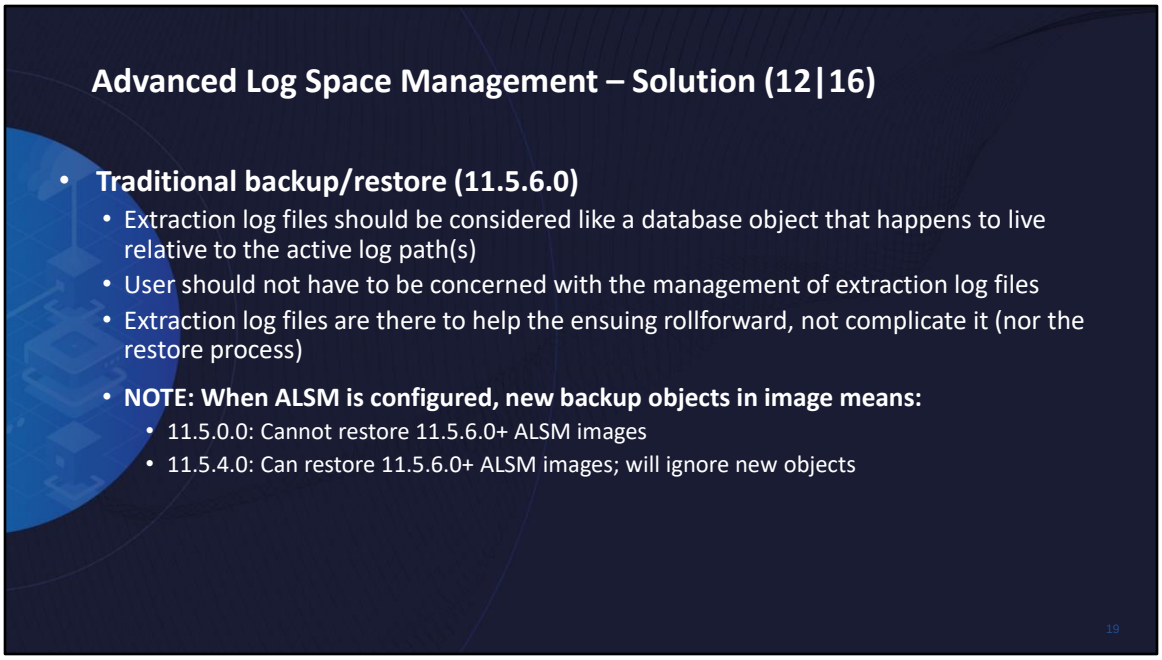

Advanced Log Space Management

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

Advanced Log Space Management – Solution (13|16)

- **Traditional backup/restore (11.5.6.0)**

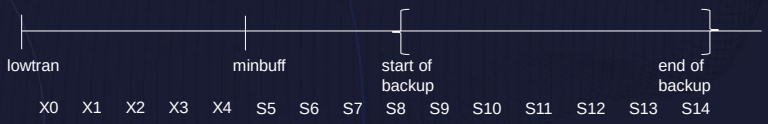

  - **BACKUP .. INCLUDE LOGS**
    - Will include the range of extraction and active log files that will allow you to roll forward to the end of backup:
      - Extraction log files copied at the start of the backup (X0-X4 [META and TID files])
      - Remaining active log files not extracted from yet copied at end of backup (S5-S14)
      - Any error reading (backing up) extraction log files causes us to stop including extraction log files in the image, but backup continues by including the active log files from the error file and onwards; this may involve retrieving active logs from the archives

20

Advanced Log Space Management

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

Advanced Log Space Management

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

Advanced Log Space Management – Solution (15|16)

- **HADR Support (11.5.6.0)**
  - **Recommendation: User sets up a shared archive location between primary and standby**
  - Will still function without a shared archive location, but on extraction log file errors may need manual user intervention if system does not have access to active log file in the archives
  - Primary will have log extraction like a non-HADR database today
  - If during rollback an error is detected with an extraction log file, will attempt to retrieve from the archives; **NOTE: if the active file cannot be accessed, manual user intervention will be required**
  - Standby will make use of extraction log files during redo/undo like database rollforward and crash recovery today
  - Standby will have throttled extraction of inflight data ready for undo keeping a balance between physical disk space available and configured log space available

22

Advanced Log Space Management

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

Advanced Log Space Management

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

Advanced Log Space Management – Current Restrictions

- **Future Support**
  - Databases in pureScale environments
- **No Plan to Support**
  - Databases configured with circular or log retain logging (LOGARCHMETH1/2)

Advanced Log Space Management – Current Limitations

- **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
  - Primary active log path:
    - Dedicated file system (not shared with other data)
    - Space for (LOGPRIMARY+LOGSECOND) * LOGFILSIZ
    - Additional space available for extraction log files
    - Additional space for log retrieval
      - Avoid by using OVERFLOWLOGPATH db config parm
  - Will be addressed in future (uncommitted) by log_disk_cap

Agenda

- **Advanced Log Space Management**
  - Overview
  - **Monitoring**
  - Future

Advanced Log Space Management – Monitoring (1|5)

New columns for MON_GET_TRANSACTION_LOG

Column Name	Data Type	Description
LOG_EXTRACTION_PROCESSED_BYTES	BIGINT	Number of bytes analyzed for extraction
LOG_EXTRACTION_PROCESSING_TIME	BIGINT	Time spent to extract log records
LOG_EXTRACTION_WRITTEN_BYTES	BIGINT	Number of bytes written to extraction log files
LOG_EXTRACTION_WRITE_TIME	BIGINT	Time spent writing to extraction log files
LOG_EXTRACTION_ROLLBACK_READS	BIGINT	Number of lookups in extraction files for rollback
LOG_EXTRACTION_ROLLBACK_TIME	BIGINT	Time spent for rollback lookups in extraction log files
LOG_EXTRACTION_CUR_COMMIT_READS	BIGINT	Number of lookups in extraction files for currently committed
LOG_EXTRACTION_CUR_COMMIT_TIME	BIGINT	Time spent for currently committed lookups in extraction log files
LOG_EXTRACTION_DISK_SPACE_USED_TOTAL	BIGINT	Number of bytes used in extraction log files
LOG_EXTRACTION_DISK_SPACE_USED_TOTAL_TOP	BIGINT	High water mark of LOG_EXTRACTION_DISK_SPACE_TOTAL_USED since database member activation
LOG_EXTRACTION_LAST_EXTRACTED_LOG	BIGINT	Log extent number of the last log file successfully extracted
LOG_EXTRACTION_PROCESSED_LSO	BIGINT	The log sequence offset of last processed log record for extraction
LOG_EXTRACTION_PROCESSED_LSN	BIGINT	The log sequence number of last processed log record for extraction
LOG_EXTRACTION_NUM_DISK_FULL	BIGINT	Number of times log extraction stopped, because there was not enough disk space in active log path
LOG_EXTRACTION_STATUS	SMALLINT	The current status of extraction.
LOG_EXTRACTION_THROTTLE_REASON	VARCHAR(32)	The reason why extraction is throttled, if applicable.

Advanced Log Space Management – Monitoring (2|5)

New column for MON_GET_UNIT_OF_WORK

Column Name	Data Type	Description
LOG_EXTRACTION_DISK_SPACE_USED	BIGINT	Number of bytes used in extraction log files

New XML element for MON_GET_UNIT_OF_WORK_DETAILS

Column Name	Data Type	Description
LOG_EXTRACTION_DISK_SPACE_USED	xs:nonNegativeInteger	Number of bytes used in extraction log files

Advanced Log Space Management – Monitoring (3|5)

- **ALSM integration into SAP DBA Cockpit**

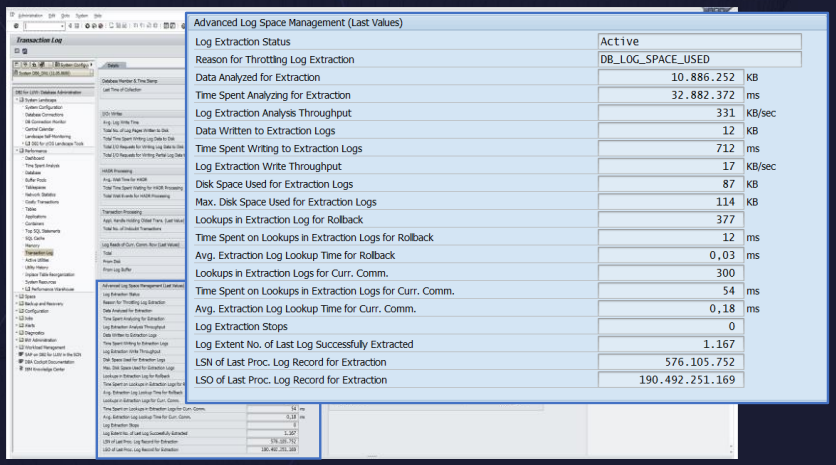

Advanced Log Space Management (Last Values)

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	

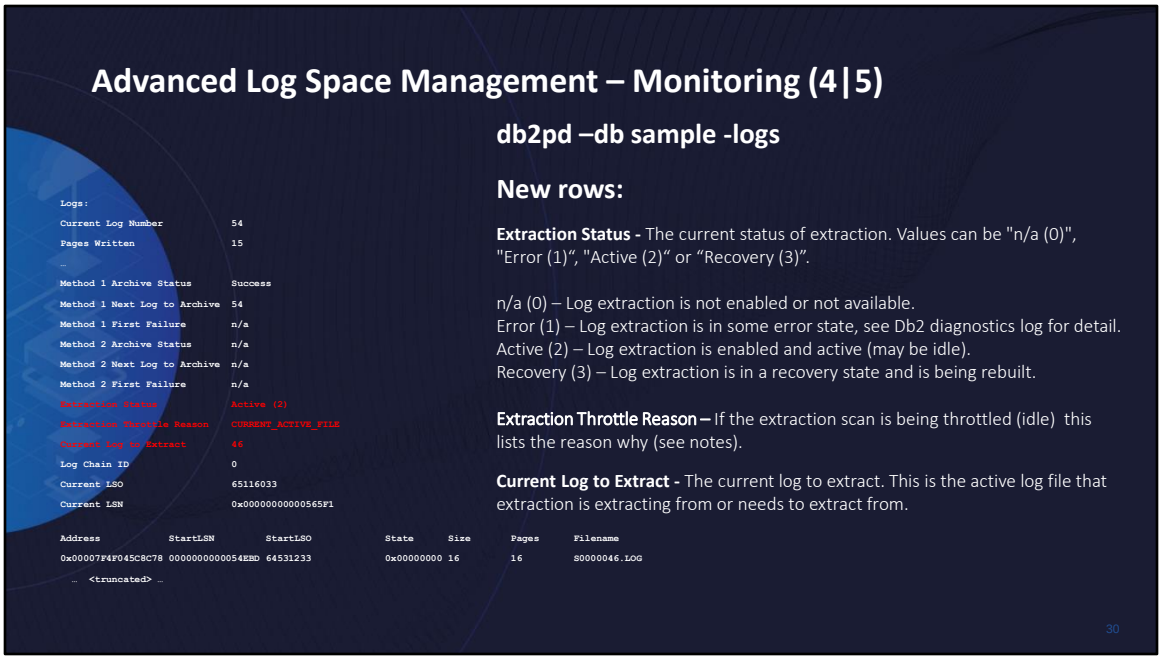

Extraction Throttle Reasons are detailed at https://www.ibm.com/docs/en/db2/11.5?topic=reference-l#r_ext_status.

Throttling can occur for any of the following reasons:

Throttling not enabled

- Throttling not available, for example because ALSM has been disabled.
- Throttle reason: n/a

Disk full

- During runtime and recovery undo, throttle extraction if the extraction log files were to interfere with the upper limit of the configured number of log files for the database. This rule ensures that the disk space that would be needed for extraction files is not taken away from the configured log space for the database. Running out of disk space in the active log path during regular logging activities could result in an unexpected database shutdown, therefore this condition overrules any other throttle reasons.This rule accounts for the **`logprimary`** and **`logsecond`** database configuration parameters, and ensures that extraction log files never breach this space. There are special cases when Db2 may end up with more primary log files than specified by **`logprimary`** (for example, with extraction running), or with less primary log files (for example, when the database is in the process of being activated and log files are being allocated asynchronously). Additionally, the **`logsecond`** parameter can change dynamically. The rule accounts for all such cases.

When log spooling is enabled on an HADR standby database, extraction is throttled to protect disk space for log spooling. If the database **hadr_spool_limit** configuration parameter is set to a fixed value, this rule will ensure that extraction does not take away any disk space needed for log spooling. If **hadr_spool_limit** is set to AUTOMATIC and disk space is limited, this rule forces extraction to share disk space between itself and log spooling. If **hadr_spool_limit** is set to -1, then this rule will not throttle extraction.

- Throttle reason: DISK_FULL

Distance from active log file

- Throttle extraction if the log file being considered for extraction is the current active log file for writing. Extraction only works on closed log files.
- Throttle reason: CURRENT_ACTIVE_FILE

Log archiving

- Throttle extraction if log archiving has not been enabled for the database, or if the log file being considered for extraction has not yet been archived. Active log files waiting to be archived will always stay in the active log path, and extraction would only duplicate disk space without any added benefit.
- Throttle reason: LOG_ARCHIVING

Log space usage

- Throttle extraction when consumed active log space is below a calculated threshold. This rule is designed to save system resources and ensures that extraction will only start when the database is close to running out of the configured log space.
- Throttle reason: DB_LOG_SPACE_USED

Extraction ratio

- Throttle extraction if the sum of the extracted data exceeds the calculated configured log space percentage limit. The purpose of this rule is to prevent extraction for very large transactions, thus potentially duplicating what is found in the active log files.
- Throttle reason: EXTRACTION_RATIO

New extraction zone

- Do not extract data if a new extraction zone has been detected and any extraction log files before this zone are no longer needed. An extraction zone is the range of log records from start to end that the extraction scan will need to process.
- Throttle reason: NEW_EXTRACTION_ZONE

Buffer pool flush needed

- Throttle extraction if the currently processed log record has not yet been flushed to the disk. Log records that have not been flushed from the buffer pool will always be needed for recovery purposes, and thus would always need to be extracted. This can be controlled by the database configuration parameters **PAGE_AGE_TRGT_MCR** and **PAGE_AGE_TRGT_GCR**
- Throttle reason: SLOW_BP_FLUSH

Previous extraction error

- Throttle extraction if the currently processed log file needs to be skipped. Upon encountering certain types of errors, the extraction scanner may decide to skip the currently processed log file and restart the scan in the next log file.
- Throttle reason: SCAN_ERROR

Advanced Log Space Management – Monitoring (5|5)
Is feature enabled?

- **Turn on registry variable:**
  - db2set DB2_ADVANCED_LOG_SPACE_MGMT=ON

- Not enabled:
  - db2diag.log will state reason

```
db2pd -db sample -logs

Extraction Status            n/a
Current Log to Extract       n/a
```

Enabled:

```
db2pd -db sample -logs

Extraction Status            Active
Current Log to Extract       0

db2pd -edus |grep loggx
901       140049278560000      7921                db2loggx (SAMPLE)


EDUID    : 901                  EDUNAME: db2loggx (SAMPLE)
FUNCTION: DB2 UDB, data protection services,
sqlpLogExtractionScanCB::loggxEnableExtractionScan, probe:1410
DATA #1 : <preformatted>
Log extraction under advanced log space management has been enabled for
database.
        Extraction state = IDLE
 Primary extraction path = /db2/NODE0000/SQL00001/LOGSTREAM0000/
  Mirror extraction path = Not set
```

31

Turn on the registry variable DB2_ADVANCED_LOG_SPACE_MGMT and activate database. The db2diag.log will display a message whether log extraction is enabled or not. db2pd –logs will also show state of extraction. All extraction is done by the new db2loggx EDU.

Agenda

- **Advanced Log Space Management**
  - Overview
  - Monitoring
  - **Future**

Advanced Log Space Management – Future

- **On by default in future release – Uncommitted**
  - Need to behave “well” for all workloads
- **pureScale support – Uncommitted**
  - Integration with active log file management
    - Runtime and merged log recovery operations (e.g. group crash recovery / database rollforward)
- **Full mirror logging support – Uncommitted**
  - Currently only support extraction log files in primary path
  - Full 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

log_disk_cap - Active log space disk capacity configuration parameter (1|2)

- **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**
- **logfilsiz 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**

34

log_disk_cap - Active log space disk capacity configuration parameter

- https://www.ibm.com/docs/en/db2/11.5?topic=dcp-log-disk-cap-active-log-space-disk-capacity

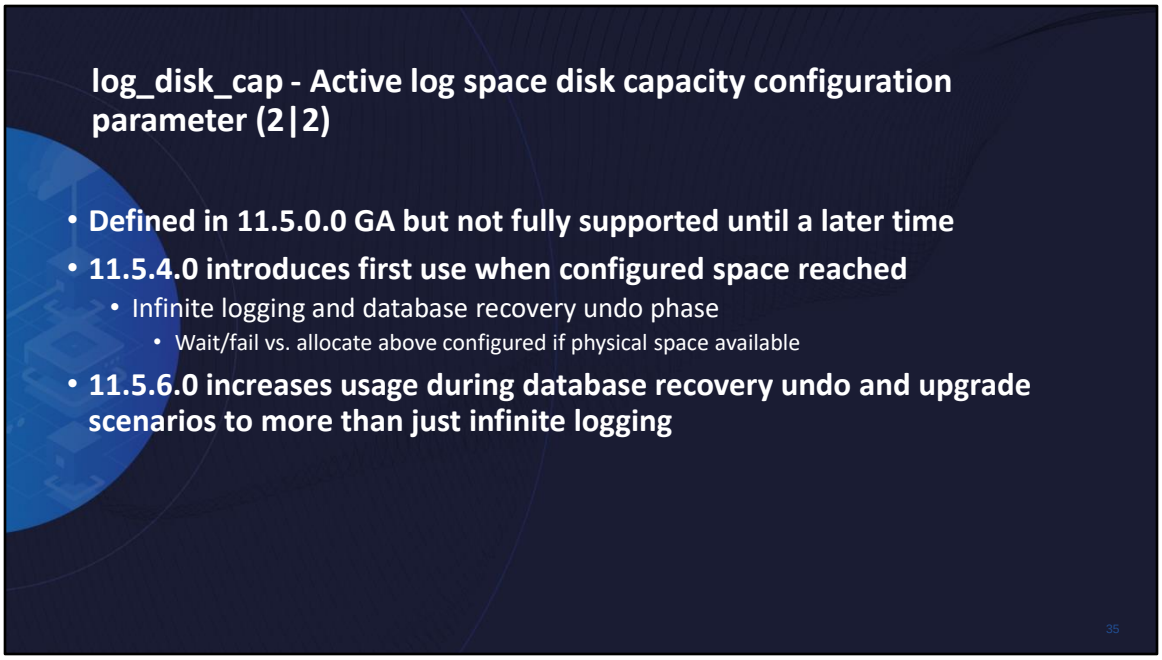

log_disk_cap - Active log space disk capacity configuration parameter

- https://www.ibm.com/docs/en/db2/11.5?topic=dcp-log-disk-cap-active-log-space-disk-capacity

Resources

- **IBM Documentation**
  - https://www.ibm.com/docs/en/db2/11.5?topic=logging-advanced-log-space-management
- **What's New 11.5.6.0**
  - https://www.ibm.com/docs/en/db2/11.5?topic=new-1156
- **SAP Blogs**
  - **https://community.sap.com/t5/technology-blogs-by-sap/advanced-log-space-management-alsm-as-of-db2-for-luw-version-11-5/ba-p/13429732**
  - **https://community.sap.com/t5/technology-blogs-by-sap/using-the-dba-cockpit-for-monitoring-advanced-log-space-management/ba-p/13515740**
  - **https://community.sap.com/t5/technology-blogs-by-sap/ad-hoc-monitoring-of-db2-advanced-log-space-management/ba-p/13545262**
  - **https://community.sap.com/t5/technology-blogs-by-sap/ad-hoc-monitoring-of-db2-advanced-log-space-management-part-2/ba-p/13581445**

Questions ???

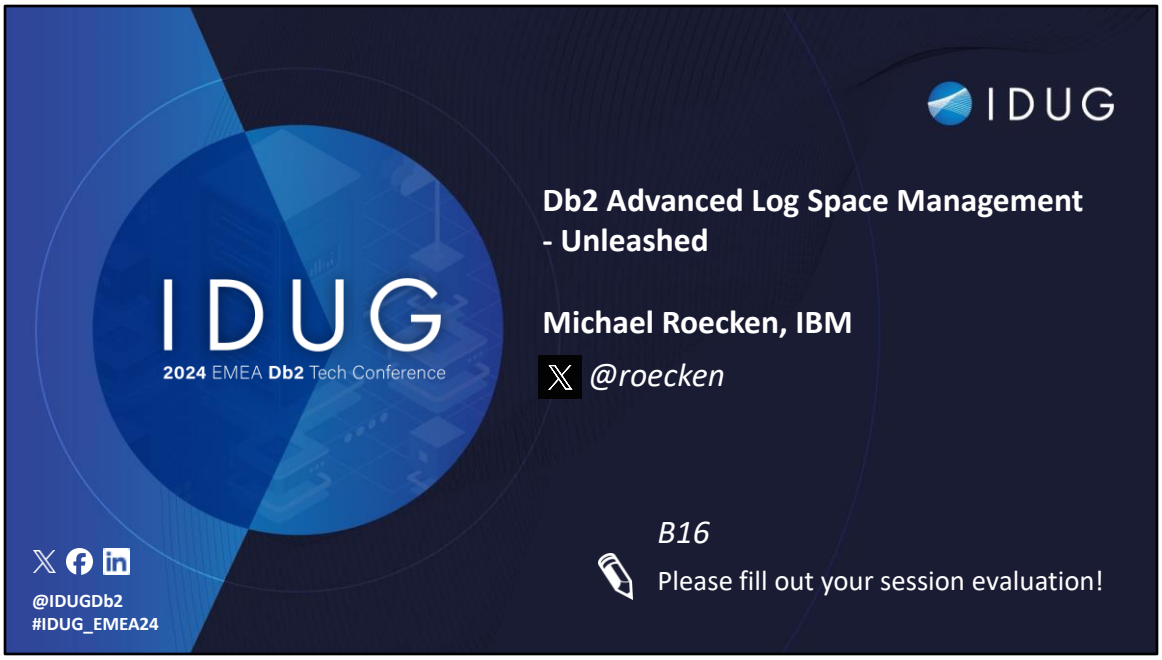

Michael Roecken is a senior software developer with Db2 for Linux, UNIX, and Windows platforms at the IBM Toronto Lab. Michael has worked since 2000 designing, implementing and supporting various features and capabilities in the areas of: backup/restore, crash/rollforward recovery, high availability/disaster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

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

Backup Slides

Advanced Log Space Management – Monitoring
Would feature be beneficial? (1|2)

- **Three considerations of what "beneficial" means:**
  - Avoid transaction log full
  - Disk space consumption
  - CPU and I/O overhead (e.g. impact to system/workload)
- **Ideally, the best of all**
  - Avoid transaction log full by extracting very little
  - Extraction process shown to be little overhead

Advanced Log Space Management – Monitoring Would feature be beneficial? (2|2)

- **Monitor the longest running transactions on the database**
- **(A) Get total amount of log space used by the 5 longest running transactions:**

```
SELECT SUM(uow_log_space_used)
FROM (SELECT uow_log_space_used,
             uow_start_time
      FROM TABLE(MON_GET_UNIT_OF_WORK(NULL,-1)) AS t
      ORDER BY uow_start_time ASC FETCH FIRST 5 ROWS ONLY)
```

- **(B) Get total active log space used by the database:**

```
SELECT total_log_used
FROM TABLE(MON_GET_TRANSACTION_LOG (-1)) AS t
```

- **If ratio of A / B is low then extraction will be beneficial**
  - Indicates long running low logging volume

41

> “If ratio of A / B is low then extraction will be beneficial”

“low” here is relative. The lower the number the better the disk space saving. So need to decide what amount of disk space you are willing to save in order to avoid transaction log full.

Advanced Log Space Management – Monitoring
How much disk space do I need to run optimally with feature?

- **Depends on amount of data that can be extracted**
  - If very little to extract then potentially can reduce active log space
- **Bare minimum is enough disk space to extract from one active log file**
  - Non-infinite: (LOGPRIMARY + LOGSECOND + 1) * LOGFILSIZ
  - Infinite: (LOGPRIMARY + 1) * LOGFILSIZ
- **Recommend ~20% extra disk space**
- **Continue to monitor until you find right fit**

Advanced Log Space Management – Monitoring
How efficient is extraction?

- **This can be based on the extraction filter rate**
  - Data analyzed vs. data written

```
SELECT log_extraction_written_bytes,
       log_extraction_processed_bytes
FROM TABLE(MON_GET_TRANSACTION_LOG(-1)) as t

    LOG_EXTRACTION_WRITTEN_BYTES LOG_EXTRACTION_PROCESSED_BYTES
    ---------------------------- ------------------------------
                           16589                         647632
```

- **16,589 / 647,632 = ~3% of active log data written has been extracted = efficient**

Advanced Log Space Management – Monitoring
Is feature healthy and working? (1|3)

- **What constitutes extraction health:**
  - No transaction log full ☺
  - Right workload / configuration (extraction filter rate)
    - Using the extraction filter rate can tell if extracting too much
      - Main cause of slow extraction speed
  - Archiving
    - Monitor to ensure not falling behind or sick
    - Use FAILARCHPATH
  - Bufferpool flushing (minbuff)
    - Verify PAGE_AGE_TRGT_MCR / PAGE_AGE_TRGT_GCR (or SOFTMAX)
  - Disk full
    - Verify storage space assigned to active log paths

Advanced Log Space Management – Monitoring
Is feature healthy and working? (2|3)

- **Helpful queries/commands:**

```
SELECT first_active_log,
       log_extraction_last_extracted_log AS last_extracted_log,
       log_extraction_num_disk_full AS num_extract_disk_full,
       log_extraction_status as extraction_status,
       log_extraction_throttle_reason as extraction_throttle_reason,
       case archive_method1_status
          when 1 then 'Good' else 'Failure' end as archive_method1_status,
       method1_next_log_to_archive,
       current_active_log,
       last_active_log
FROM TABLE(MON_GET_TRANSACTION_LOG(-1)) as t
```

```
FIRST_ACTIVE_LOG     LAST_EXTRACTED_LOG   NUM_EXTRACT_DISK_FULL EXTRACTION_STATUS EXTRACTION_THROTTLE_REASON
-------------------- -------------------- --------------------- ----------------- --------------------------
               989                 1032                     0                 2        CURRENT_ACTIVE_FILE

ARCHIVE_METHOD1_STATUS METHOD1_NEXT_LOG_TO_ARCHIVE CURRENT_ACTIVE_LOG   LAST_ACTIVE_LOG
---------------------- --------------------------- -------------------- --------------------
                  Good                        1034                 1034                 1035
```

Advanced Log Space Management – Monitoring
Is feature healthy and working? (3|3)

- db2pd -db sample -logs

```
Logs:
Current Log Number              1034
…
Method 1 Archive Status         Success
Method 1 Next Log to Archive    1034
…
Extraction Status               Active
Extraction Throttle Reason      NONE
Current Log to Extract          1033
…

StartLSN         StartLSO             State      Size       Pages      Filename
00000000000729AB 78259201             0x00000000 4          4          S0001026.LOG
…
0000000000000000 78405937             0x00000000 4          4          S0001035.LOG
```

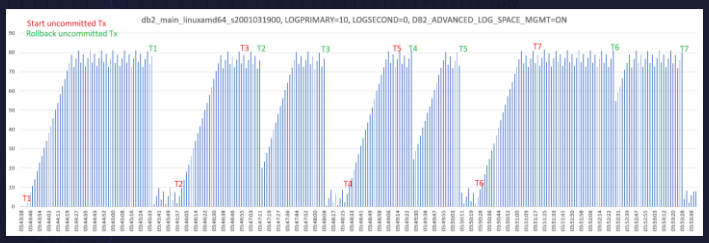

46

No log data will be extracted from an active log file that has not been archived yet. This would duplicate disk space. Ensure methx_status is 1 (healthy), not 0 (error).

No log data will be extracted from an active log file where minbufflsn exists. This is due to recovery algorithm that needs to replay all log records >= minbufflsn. So no benefit of extracting such data as it would duplicate disk space.

Advanced Log Space Management – Monitoring
What is the disk space consumption of extraction?

- Current total extraction disk space consumed
- Maximum total extraction disk space consumed since last activation

```
SELECT log_extraction_processed_bytes AS processed_bytes,
       log_extraction_written_bytes AS written_bytes,
       log_extraction_disk_space_used_total AS disk_space_used_total,
       log_extraction_disk_space_used_total_top AS disk_space_used_total_top
FROM TABLE(MON_GET_TRANSACTION_LOG(-1)) as t
```

```
PROCESSED_BYTES WRITTEN_BYTES DISK_SPACE_USED_TOTAL DISK_SPACE_USED_TOTAL_TOP
--------------- ------------- --------------------- -------------------------
         266882           165                 35165                     54461
```

47

This query tells you that since the last activation extraction processed 266,882 bytes of log data from the active log files. From that amount, 165 bytes of log data was written to extraction TID files. The current total amount of disk space consumed by extraction files, including log data and meta data, is 35,165 bytes. Since the last activation, extraction has taken up 54,461 bytes.

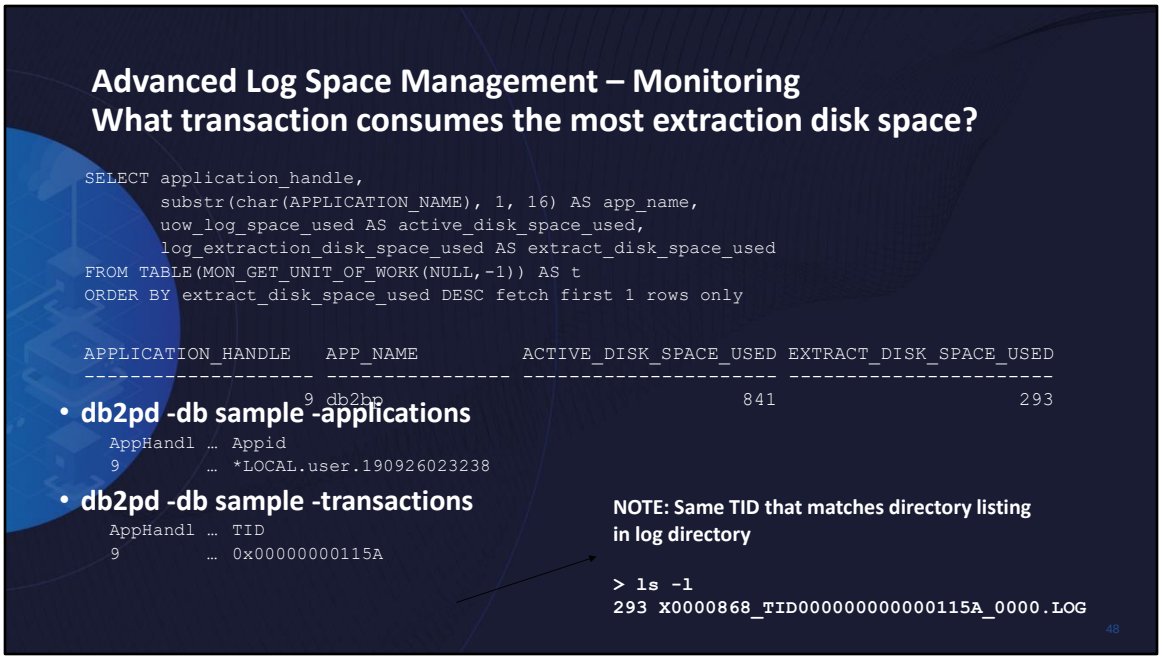

To find the transaction that is consuming the most extraction log space allows one to understand if this is a known expectation or not, maybe possibly a rogue transaction.

You can map a transaction ID (TID) from a directory listing or you can use a combination of commands to track down which application/transaction is consuming the most amount of extraction space.

The above example shows that this particular transaction has written 841 bytes of log data to the active files, but only 293 bytes have been extracted so far.