Data and AI

The steps and techniques needed for mass recovery including Db2 recovery from cyber-attacks

Anthony Ciabattoni
Executive IT Specialist
Db2 for z/OS Development SWAT Team

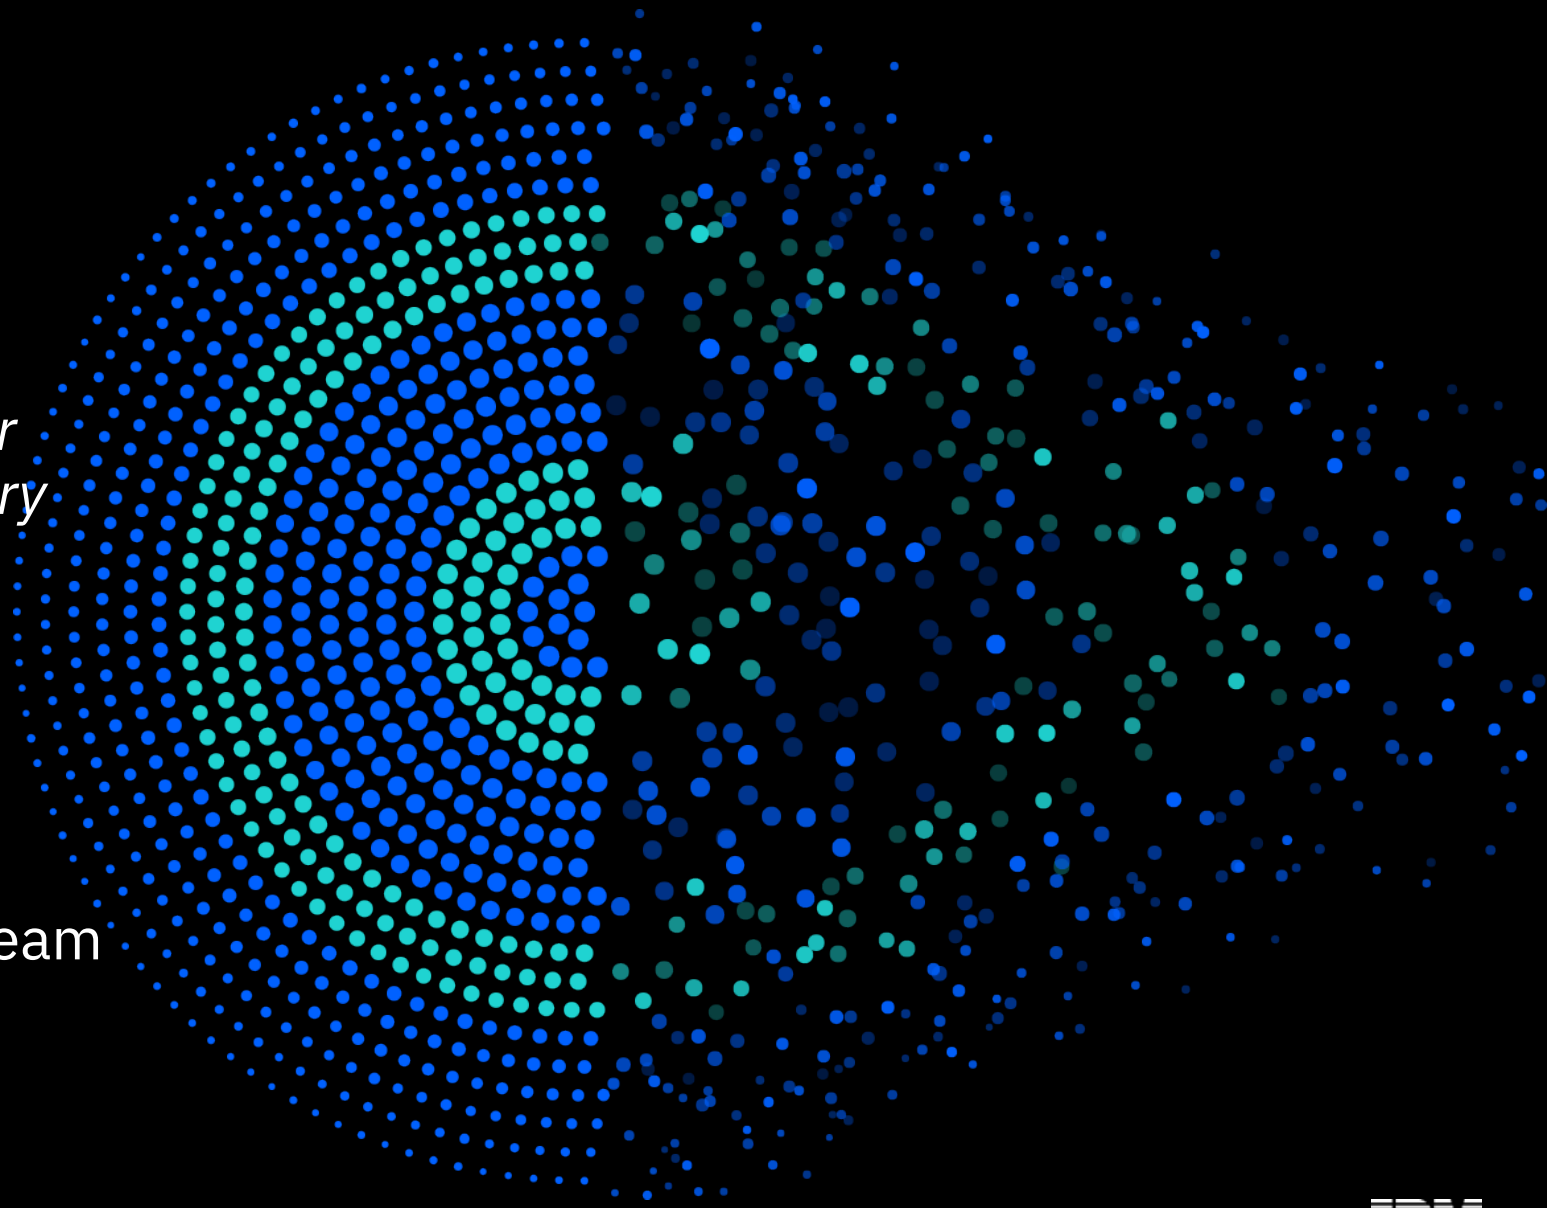

IBM Data and AI

IBM

Agenda

- Background
- Common Issues
- Db2 Logging Considerations
- Db2 Catalog/Directory Best Practices
- Backups
- Recovery
- Data Integrity Testing
- Recovery Testing
- Application Design
- Cyber Resiliency
- Summary

Db2 recovery background

- Db2 log-based recovery of *multiple objects* may be required when…
  - Catastrophic DASD subsystem failure and no second copy
  - Plan B for disaster recovery
    - Mirror is damaged/inconsistent
    - Bad Disaster Restart e.g., using stale CF structures in data sharing
  - Data corruption at the local site caused by…
    - ‘Bad’ application program
    - Operational error
    - Db2, IRLM, z/OS, third-party product code failure
    - CF microcode failure, DASD subsystem microcode failure

- Scope of the recovery may be more or less extensive
  - One application and all associated objects
  - Part of the system (including a random list of objects across multiple applications)
  - Or, in the worst case, the ‘whole world’

Db2 recovery background …

- **Db2 log-based recovery of multiple objects is a very rare event …**

  … but statistically, it is more frequent than a true DR event (flood, fire, etc.)

- Taking regular backups is necessary but far from sufficient for anything beyond minor application recovery
- If not prepared, practiced and optimized, will lead to extended application service downtimes – possibly many hours to several days
  - Things to consider
    - Are my procedures up to date?
    - Configuration changes? Db2 release?
    - Are image copies and recovery jobs created based on object priority?
      - ✓ How long will the “recover” take?
    - Are the image copies on DASD, VTS or physical tape?
    - Are all my objects backed up?
    - If not practiced “what do you not know?”

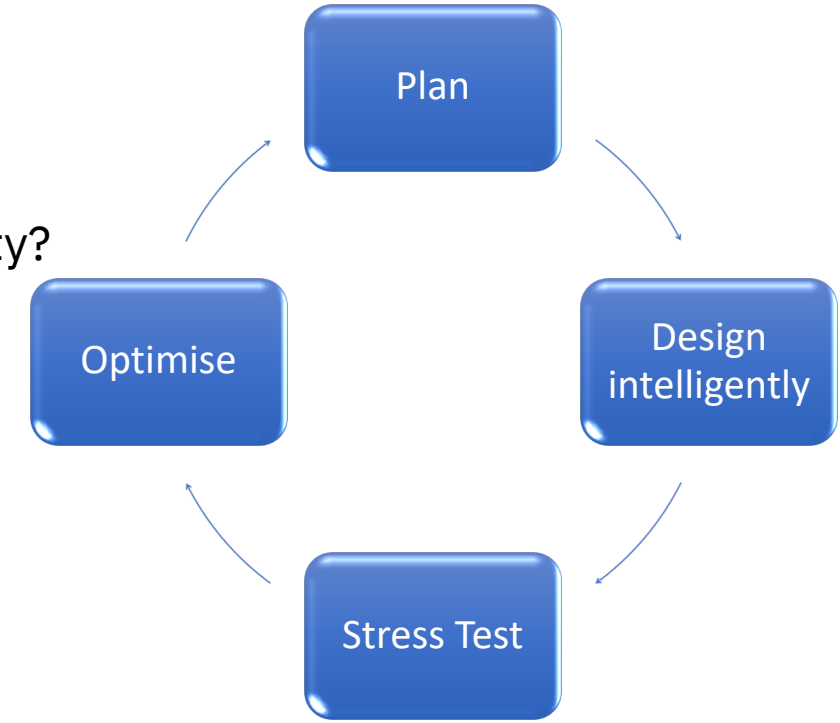

High performance multiple object recovery

Common Issues

- Common issues
  - Lack of planning, intelligent design, optimization, practice & maintenance
  - No prioritized list of application objects and inter-dependencies
    - Limited use of Db2 referential integrity
    - Data dependencies and integrity management are buried in the applications
    - Heavily dependant on application knowledge and support
  - Procedures for taking backups and executing recovery compromised by lack of investment in technical configuration
  - Backup and recovery procedures have not been addressed for years
  - Use of tape including VTS *(“Identity Crisis”)*
    - Cannot share tape volumes across multiple jobs
    - Relatively small number of read devices
    - Concurrent recall can be a serious bottleneck
    - Even though VTS has a disk cache, it is known to z/OS as tape device
      - Same serialization characteristics as all tape devices
      - A single virtual volume cannot be shared by different jobs or systems at the same time

High performance multiple object recovery

Common Issues …

- Results: any or all of the following
  - No estimate of elapsed time to complete
  - Elongated elapsed time to complete recovery
  - Performance bottlenecks so that recovery performance does not scale
  - Breakage in procedures
    - Revert to trial and error approach
  - Surprises caused by changing technical configuration
  - Unrecoverable objects

Factors that affect recovery elapsed time

- ‘Think time’ and preparation of the recovery plan
- Restore
  - Number of pages, number of objects?
  - Backups on tape or DASD? Standard ICs, FCICs, SLBs?
  - Degree of parallelism?
- Log scan
  - Backup frequency
  - Archive logs needed to recover?
  - Archive logs on tape or DASD?
  - Degree of parallelism?
- Log apply
  - Update frequency and update patterns
  - Maximal fast log apply?
- Recover/Rebuild indexes

Db2 logging environment

- Design for availability
  - Keep a minimum of 6 hours of recovery log data in the active log pairs at any time
    - Objective: provide some reaction time in case of archiving problem
    - If there is a stalled active log offload, very little time available to investigate the problem, fix the problem and allow the active log offload processing to catch up
      - ✓ Effectively the Db2 member will stop processing work
      - ✓ Likely to lead to sympathy sickness across the entire Db2 data sharing group
    - Adjust number/size of the Db2 active log pairs
      - ✓ Limit is 93 log pairs
      - ✓ 4GB – 1 Byte (Db2 11), 768GB (Db2 12)
    - Active logs can be added dynamically
      - New -SET LOG NEWLOG option
      - New active log must be IDCAMS defined & preformatted by DSNJLOGF
      - Only a single log dataset at a time
        - ✓ Issue command twice for dual logging
      - No dynamic delete of active log datasets
        - Functionality added to Db2 VNext to dynamically delete active log datasets
- Design for recovery performance
  - Always write archive log COPY1 and COPY2 to DASD, and let DFSMShsm (or equivalent) migrate them away to tape or VTS
    - Eliminate contention on reading the archive logs during recovery
    - Especially important in a data sharing environment
    - Storage needs to be available and processes to recall from tape/VTS prior to recovery

Db2 logging environment …

- Design for recovery performance …
  - Keep at least 48h of recovery log data on DASD

Option #1: Over-configure the active log pairs (number/size)
Write archive log COPY1 and COPY2 to DASD but they can be migrated to tape/VTS at any time

Pros: Optimal log read performance with automatic load balancing for reads between active log COPY1 and COPY2,
Db2 12 increases capacity to 93x768GB
Cons: Maximum capacity in V11 = 93x4GB

Option #2: Keep archive log COPY1 on DASD for 48-72h before migrating it to tape/VTS – archive log COPY2 can be migrated to tape/VTS at any time

Pros: Good log read performance from archive on DASD, potential for less DASD requirements than Option 1

- Be ready to extend the amount of recovery log beyond what is available on DASD
  - Set BLKSIZE=24576 to optimize reads on DASD
  - Prepare a procedure to copy archive logs from tape or VTS to DASD

Db2 logging environment ...

- Design for resiliency
  - Separate COPY1 and COPY2 of the active log pairs and BSDS across different DASD controllers if possible – across different extent pools (RAID arrays) at the minimum
  - Isolate objects into separate ICF user catalogs
    - Separate out the datasets for each Db2 member into separate ICF catalogs
      - Active logs, archive logs, BSDS for member Db2A away from those for member Db2B
      - Result: an ICF Catalog failure would only affect one Db2 member
    - Should also consider further isolation
      - COPY1 of active log, archive log and BSDS into one ICF catalog
      - COPY2 of active log, archive log and BSDS into an alternate ICF catalog
      - Result: an ICF Catalog failure would not affect Db2
    - Additional ICF Catalog considerations for better performance and resilience
      - Isolate Db2 Catalog/Directory objects into a separate ICF catalog
      - Use multiple ICF Catalogs for the Db2 user objects
      - Separate ICF Catalogs for Db2 objects and Db2 image copy backups

Db2 logging environment ...

- Design for serviceability
  - Retain archive log data for 30 days
  - Keep the maximum number of archive logs in the BSDS
    - Set ZPARM MAXARCH=10000
  - At first sign of logical data corruption, stop the deletion of Db2 recovery assets
    - Image copies
    - Archive log datasets
  - Problem diagnosis and resolution

Db2 Catalog/Directory

- Db2 Catalog and Directory objects
  - Db2 Catalog/Directory is “heart” of a Db2 system
  - Take frequent FICs of the Db2 Catalog/Directory
    - At the very minimum daily – best is several times a day using SHRLEVEL(CHANGE)
    - Keep a copy on DASD to speed up recovery
  - Design, build, test and maintain a preplanned job to recover the Db2 Catalog and Directory objects for Db2 12 at the current Catalog Level (CL) and follow-on releases
    - Correctness
    - Recovery timings
    - Influence the number of copies per 24 hours
- Periodically check the integrity of the Db2 Catalog/Directory
  - e.g., using a cloned copy of the Db2 Catalog/Directory into an auxiliary Db2 subsystem
  - See next slide for recommended checks
- Periodically reorganize the Db2 Catalog/Directory
  - Outside of release migration
  - Most importantly, SYSLGRNX should be reorganized at least every quarter
    - Can be run as SHRLEVEL(CHANGE) at a time of low activity
    - Will speed up online REORG, MODIFY RECOVERY, RECOVER, GRECP/LPL recovery

Db2 Catalog/Directory …

- Series of tests that should be run on a regular basis to flush out any latent inconsistency in the Db2 Catalog
  - SQL queries from migration job DSNTESQ
    - Should always return zero rows
  - REPAIR DBD TEST or DIAGNOSE
  - Basic RUNSTATS on all objects
  - CHECK INDEX on all indexes
  - For catalog objects with LOB columns:
    - CHECK LOB
    - CHECK INDEX on AUX index
    - CHECK DATA on base tablespace using SCOPE AUXONLY AUXERROR REPORT

Image copy backups

- Always take dual image copies as part of REORG and LOAD REPLACE (LOG NO events)
- Use as much DASD as possible for optimal recovery
  - If DASD space is an issue
    - Use template switching to write image copies for small objects to DASD and manage by DFSMShsm
      - Objective: Allow fast restore and take pressure off the VTS in case of mass recovery

```
TEMPLATE LRG DSN &DB..&TS..D&DA..T&TI. UNIT=TAPE
TEMPLATE SML DSN &DB..&TS..D&DA..T&TI. UNIT=SYSALLDA LIMIT(20 CYL,LRG)
COPY TABLESPACE SMALL.TS COPYDDN(SML)
COPY TABLESPACE LARGE.TS COPYDDN(LRG)
```

    - Use Db2 data compression for table spaces → COPY does not decompress data
- Consider shortening the full image copy (FIC) cycle time (<= 24 hours) for Db2 Catalog/ Directory and potentially most critical application data
  - Objective: Reduce log apply time
  - Implement a smart image copy process
- Consider use of incremental image copy (IIC)

Image copy backups …

- Exploitation of FlashCopy technology
  - Data set FlashCopy image copies
    - ✓ Potential for significant elapsed time reduction for the RESTORE phase
    - ✓ Can also be used to create a transaction-consistent image copy with COPY SHRLEVEL CHANGE
  - Db2 Backup System Can be restored quickly – if still on DASD
    - Can also be used to create a ‘forensic’ system
      - ✓ Quick cloning if the environment away from main production system
      - ✓ Level restored will be to a point in time where the data is known to be good
      - ✓ Application teams can then analyze and reconcile the data contents of the forensic system vs. current damaged system
- Recommend NOT to use GDGs for image copy datasets
  - Risk of old versions rolling off by accident
    - Especially if using incremental image copies
- Use catalogued datasets instead, with ‘meaningful’ naming convention
  - Adds informational value (e.g., date and time of the backup)

Image copy backups …

- Considerations
  - Schedule a daily production health check process to identify:
    - Unrecoverable objects
    - Existence of two full image copies especially after a LOG NO event e.g., REORG
    - At least 2 image copies present for “read only” tablespaces
    - If IIC are used, ensure a FIC is also available

Identifying the scope of data corruption

- Before developing a recovery plan, it is vital to identify the scope of data the data corruption
- CHECK is a critical tool in case of data corruption
  - Without FlashCopy support, CHECK utilities can be very disruptive
    - Even with SHRLEVEL(CHANGE) – R/O access during creation of the shadow objects
  - CHECK utilities exploiting FlashCopy enables the ability to non-disruptively identify scope of data corruption
- Important FlashCopy parameters:
  - Db2 ZPARM CHECK_FASTREPLICATION
    - PREFERRED (default V9) >> Standard I/O will be used if FC cannot be used
    - REQUIRED (default V10) >> CHECK will fail if FC cannot be used **← strongly recommended whether FlashCopy is available or not**
  - Db2 ZPARM UTIL_TEMP_STORCLAS
    - Optional: can be used to specify a storage class for the shadow data sets
    - If blank, the shadow data sets are defined in the same storage class as the production page set
    - If using DASD-based replication, specify a pool of volumes that are not mirrored
      - ✓ Applies to Metro Mirror (PPRC) without Remote Pair FlashCopy (ZPARM FLASHCOPY_PPRC = REQUIRED), z/OS Global Mirror (XRC) and Global Mirror

Design for intelligent Db2 recover

- Application data recovery
  - It is not just about JCL generation; it is essential to:
    - Plan, design intelligently, stress test and optimise to meet recovery time objective
    - Prioritise most critical applications
    - Understand application and data interdependencies
    - Design for parallel recovery jobs
    - Optimised utilisation of technical configuration
    - Intelligent creation and scheduling of recovery jobs
    - Design for DASD-based recovery for optimal performance
    - Practice regularly

Design for intelligent Db2 recover …

- Build intelligent recovery jobs
  - As a starting point, develop basic automation to generate RECOVER jobs from SYSCOPY for a mass data recovery, in an intelligent fashion, optimizing to reduce overall elapsed time
    - ‘Basic’ starting point, which should be refined after testing to demonstrate achievement
      - ✓ List of 20-30 objects per RECOVER job step, taking into account the stacking of IC datasets on volumes if tape stacking is still used
      - ✓ Maximum of 51 RECOVER jobs per Db2 member for optimal fast log apply (may have to be reduce if too much contention on VTS or DASD)
      - ✓ Concurrent REBUILD indexes (level of parallelism limited by available sort space)
    - Consider the use of INDEX COPY and RECOVER as an alternative to REBUILD INDEX for large NPIs to optimize the overall recovery time
      - ✓ Note: REBUILD INDEX still preferred option after index vs. table mismatches
  - Identify a solution to test and optimize mass application data recovery procedures
    - Take advantage of utilizing the DR environment for testing
      - ✓ “White space” time after a regularly scheduled test
      - ✓ Infrastructure DR testing slots
      - ✓ Need to have access to the VTS grid in read-only mode to access image copy and archive log datasets
    - Local pre-production environment with production like sized objects

Db2 12 Point-In-Time(PIT) Recovery - SCOPE UPDATED

- New in Db2 12, SCOPE UPDATED option for RECOVER utility
  - Goal is to speed up the elapsed time for PIT recovery (TORBA/TOLOGPOINT)
  - SCOPE UPDATED is the default
    - Objects are excluded from recovery that have not changed since the given recovery point
    - Avoids wasting time restoring the image copy for a given object
    - Potentially a great performance optimization
    - Might be known to a few customers by running RECOVER under DIAGNOSE(607) for Db2 11
  - As before, SCOPE ALL forces the recovery of any object even if it has not been changed
  - How does Db2 know if the objects have been changed or not after the recovery point?
    - By reading the entries in SYSLGRNX

Db2 12 Point-In-Time(PIT) Recovery - SCOPE UPDATED ...

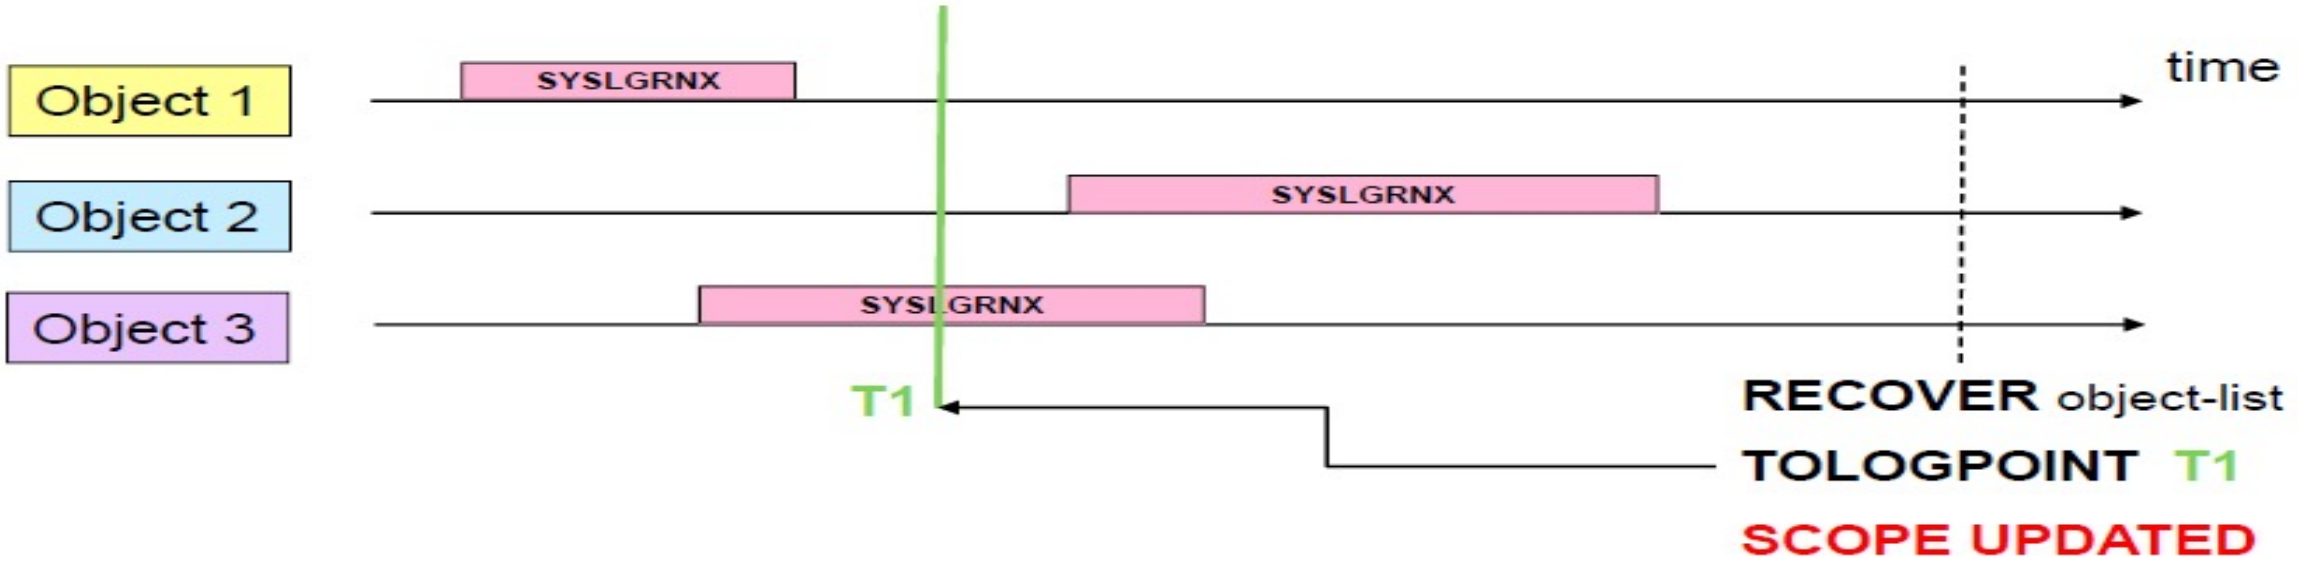


```
DSNU1322I -DB2A 190 04:02:20.42 DSNUCALX - PROCESSING SKIPPED FOR
TABLESPACE ... BECAUSE THE OBJECT DOES NOT NEED TO BE RECOVERED
...
DSNU010I     265 10:10:02.14 DSNUGBAC - ...HIGHEST RETURN CODE=0
```

Object 1

Object 2

Object 3

both objects are consistently recovered to T1

Db2 12 Point-In-Time(PIT) Recovery - SCOPE UPDATED …

- Recovery of unchanged tablespaces
  - Determined at execution time
  - Only applies to point in time recovery
  - Db2 12 RECOVERY SCOPE UPDATED
    - Default
- REBUILD INDEX
  - Default SCOPE ALL
  - Db2 12 SCOPE PENDING (recommended)
- Point in Time Recovery Recommendation
  - RECOVER … SCOPE UPDATED + REBUILD INDEX … SCOPE PENDING

```
DSNU1322I -DB2A 190 04:02:20.42 DSNUCALX - PROCESSING SKIPPED FOR
TABLESPACE ... BECAUSE THE OBJECT DOES NOT NEED TO BE RECOVERED
...
DSNU010I 265 10:10:02.14 DSNUGBAC - ...HIGHEST RETURN CODE=0
```

Db2 12 Point-In-Time(PIT) Recovery - SCOPE UPDATED …

- Problem statement
  - If a PIT recovery of SYSLGRNX is performed, then can run into a big issue when performing a PIT recovery for application objects after the PIT recovery of SYSLGRNX
  - Scenario
    - No data changes have been done to the application object between the time of COPY and the recovery point of SYSLGRNX
    - But data changes are done later after the the recovery point of SYSLGRNX
  - Db2 will not detect that situation because the entries of SYSLGRNX have been eliminated by the PIT recovery of SYSLGRNX
    - THE OBJECT DOES NOT NEED TO BE RECOVERED
    - Db2 will incorrectly exclude the recovery of the application object which will lead to data inconsistencies
    - REPORT RECOVERY utility also reads SYSLGRNX and will not report the changes
    - The RECOVER job ends with RC04 and DSNU1322I message for each excluded object
      - DSNU1322I =D2LC 318 12:10:14.52 DSNUCALX - PROCESSING SKIPPED FOR TABLESPACE DSNDB06.SYSTSLVH BECAUSE THE OBJECT DOES NOT NEED TO BE RECOVERED

Db2 12 Point-In-Time(PIT) Recovery - SCOPE UPDATED …

- Recommendations
  - Apply PTF for APAR PH20056
    - RECOVER will internally change SCOPE UPDATED to ALL for PIT recovery of Catalog/Directory objects
    - RECOVER will internally change SCOPE UPDATED to ALL for any object after SYSLGRNX has been recovered to a prior point in time
      - DSNU124I message will be issued to indicate the override
      - RC will still be 0
    - No plan to change the default from SCOPE UPDATED to SCOPE ALL
  - Until PTF applied for APAR PH20056, explicitly specify SCOPE ALL on PIT recoveries
    - Any pre-existing “canned” job to recover the Catalog/Directory should be modified

“Backout’ Recovery

- Db2 log analysis product to create UNDO SQL to reverse a data changing event
  - Identify and create UNDO SQL to backout a given unit of work, transaction, action
  - Object(s) remain available while SQL is backing out the change
- BACKOUT YES option for point-in-time recovery
  - Backs out both data (except NOT LOGGED) and indexes (if defined as COPY YES)
    - COPY NO indexes must be rebuilt when backout complete
    - You can ALTER indexes to COPY YES and not produce image copies
      - SYSLGRNX entries build up
      - Use MODIFY to delete them (AGE or DATE) even without copies
  - True rollback, not run of generated SQL undo statements
  - Changes are backed out from the current state of the object → not for media recovery
  - Intent: Short backout, not hours/days
    - Fast Log Apply is not used
    - The recovery point must be contained within the Db2 system checkpoints that are recorded in the BSDS for each member
      - Message DSNU1545I-RECOVER does not process any of the objects and ends with RC8

Optimizing Db2 recovery

- Objects with longest end-to-end recovery time need to be recovered first
  - Size of the object
  - Update rate since last image copy
  - Number and size of indexes

- Optimize job scheduling to avoid ‘dead times’

Optimizing Db2 recovery …

- Create automated procedures to create efficient recovery jobs
  - Considerations
    - Table prioritization
    - Virtual Tape/Tape optimization
  - Frequency
    - Execute after nightly backup jobs
      - Recovery jobs updated daily and ready to execute
      - Execute when CPU is available (middle of night)
    - Procedures in place to execute efficiently at time of recover
      - Automated process(REXX) to execute and create recovery jobs at time of recovery
      - Needs to be efficient
      - Procedures need to be tested and proved out periodically

Optimizing Db2 recovery – Stress Test

- Practice regular full-scale ‘fire drills’ for mass recovery of an entire application or even the entire system

- Objectives:
  - Validate that procedures are in working order
    - Both for local and remote DR recovery
  - Maintain readiness on mass recovery execution
  - Find out what the actual service level is
    - Break down the elapsed time of each job: RESTORE/LOG APPLY/REBUILD INDEX
    - If elapsed time needs to be improved further, look for possible optimizations

Optimizing Db2 recovery – Stress Test …

Lot of ‘dead times’ introduced by the job scheduling >> next slide will show how this test would have looked if it took only 1 minute to start REBUILD INDEX after the RECOVERY of all parts + if the job #17 had not been started late

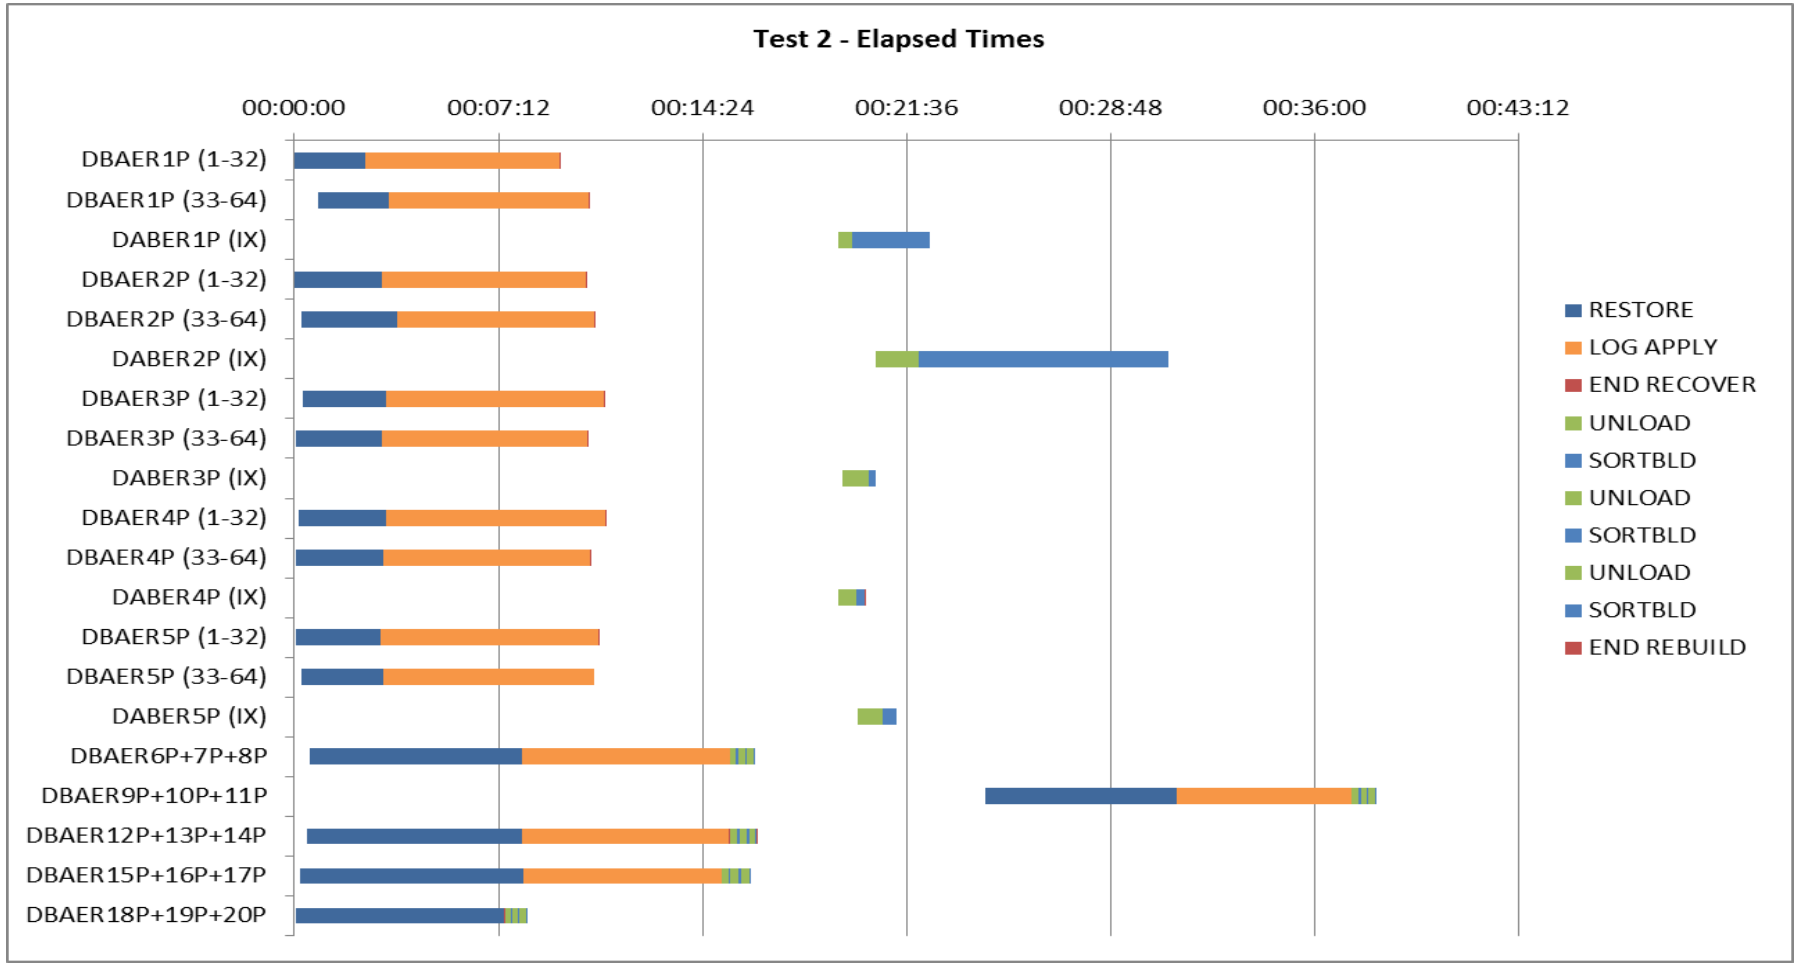

Application design

- Common problems
  - Applications not committing frequently
  - No clear separation between active and inactive data
  - Critical applications tightly coupled to non-critical applications by shared data
  - Data inter-dependencies across multiple data sources (e.g., Db2/VSAM, Db2/IMS)
- Recommendations
  - Frequent commits in long-running batch applications
    - Dynamic, table driven
    - Application must be able to restart from intermediate commit points
  - Separate active from inactive (historical) data
    - Use separate tables
    - Regular, aggressive pruning back of active tables
    - Application toleration of unavailable inactive data
    - Db2 11 – Transparent Archiving
  - Data isolation to de-couple applications
    - Objective – bring back critical applications first to resume availability
      - ✓ Incrementally bring back additional applications

New Dimension of Resiliency is Required

Current infrastructures focus on BC / DR

- HyperSwap
- Snapshots
- Replication
- Backups
- Data “Gold Copies”

Add a focus on Cyber Resilience

- Immutability
- Minimized data loss
- Isolation
- Data Latency

Business Continuity/Disaster Solutions

- Current infrastructures focus on Business Continuity (BC) / Disaster Recovery (DR)
  - Multi-site resilient infrastructure
    - All assets are replicated
      - ✓ z/OS
      - ✓ Db2
      - ✓ DASD storage infrastructure
        - ✓ Metro Mirror, Global Mirror configuration
        - ✓ 4-site DASD Mirror configuration
  - HyperSwap capabilities
    - Planned/unplanned
  - VTS or equivalent is replicated
  - “Gold Copies”/Snapshot copies
    - Various frequencies
      - ✓ Daily/Weekly
  - Site Swap capabilities

Cyber Resiliency

- Cyber exposures in current infrastructures focus on Business Continuity (BC) / Disaster Recovery (DR)
  - Mirroring technology
    - All DASD I/O are replicated
    - All VTS datasets replicated or written in parallel
    - Good, bad and corrupted corrupted activities are replicated (rolling corruption)
      - ✓ Local corruption leads to remote corruption
  - Metro Mirror/Global Mirror/VTS copies commonly do not have access isolation
    - No physical access protection
      - ✓ Mirrored copies can be corrupted independently of primary copy
  - Prior to ransomware event
    - Mischievous attackers perform proactive destructive cyber activities
      - ✓ Delete/corrupt recovery assets at primary site
      - ✓ Destructive activities at mirrored site
      - ✓ Includes:
        - ✓ Archive log datasets
        - ✓ Image copy datasets
      - ✓ Incomplete recovery assets needed for recovery
  - “Gold Copies”/Snapshot copies
    - Point-in-time copies
    - Only as current as the time they are executed
      - ✓ Will lead to data loss

Data Resilience inc. Cyber Resilience

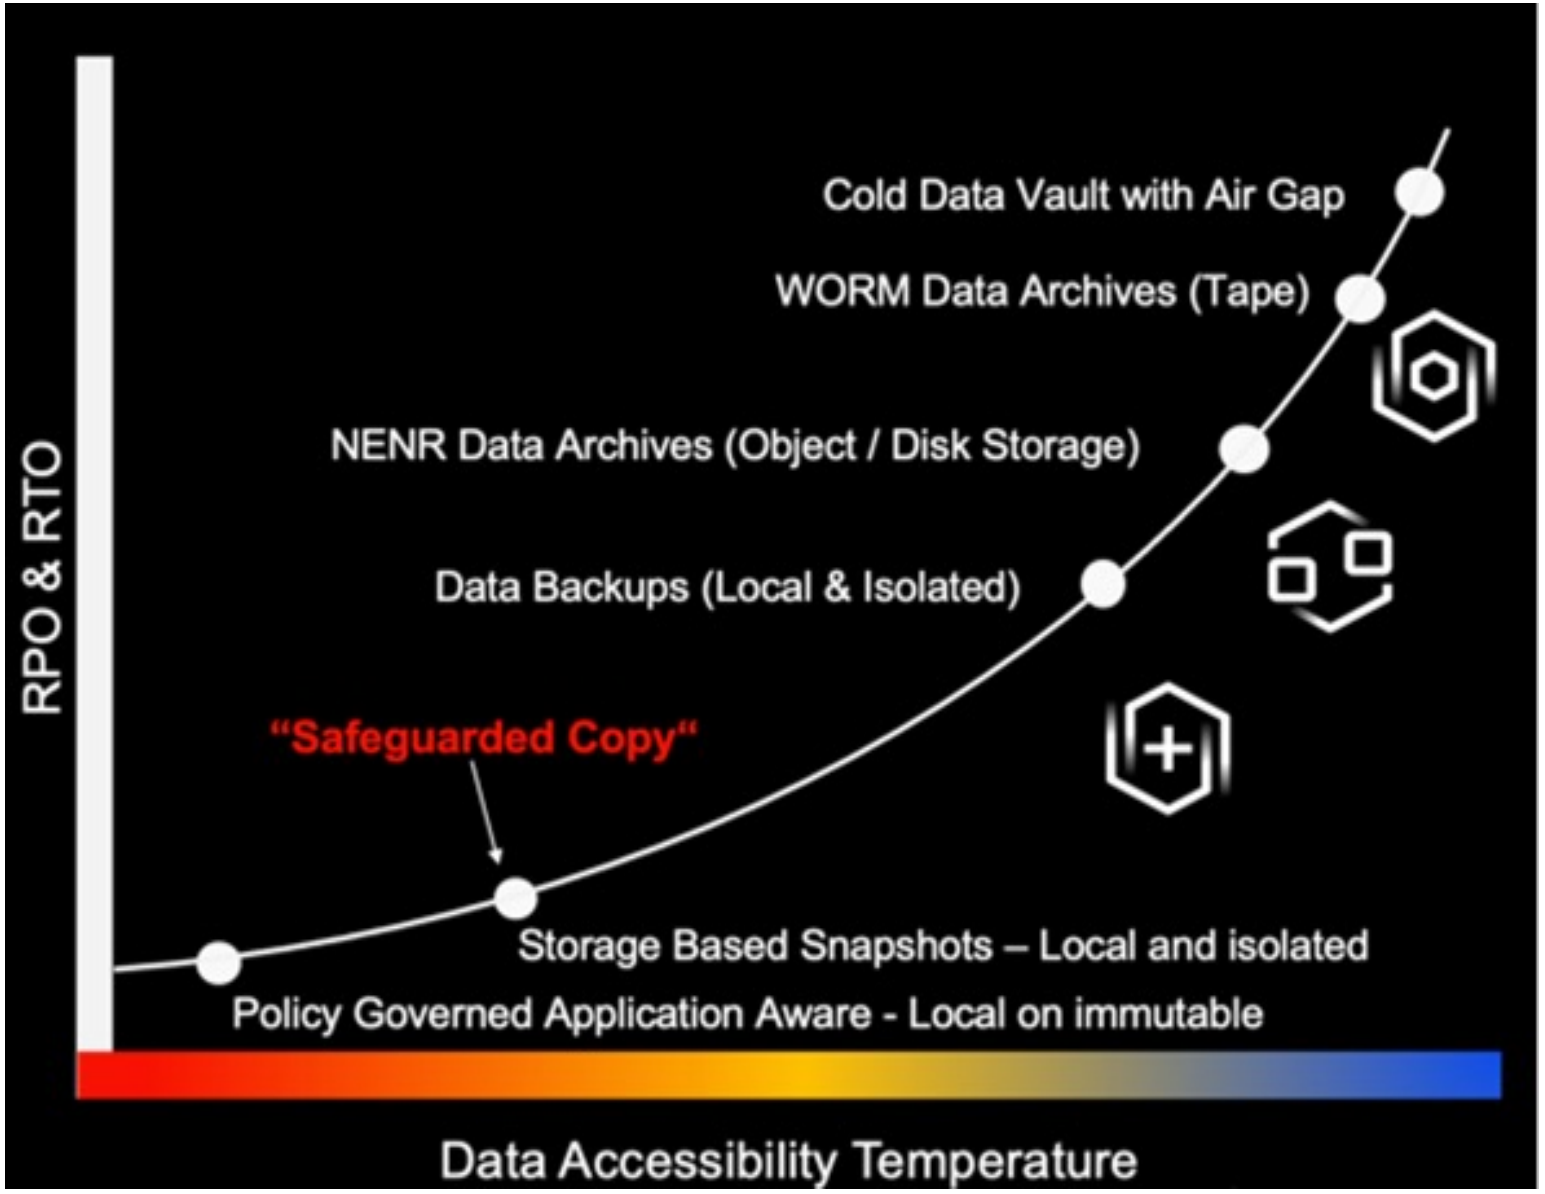

Copy Separation:

Create a structure of data separation across multiple layers and services including;

- Copy Services
- Backup Services
- Separation of security controls

Immutability & Access Isolation

Create a structure of data immutability at multiple layers and services including;

- Logical / physical isolation (Air Gap)
- Non-erasable / Non-rewritable Storage
- Cold Storage / Object Storage
- Data Vaults
- Isolated Infrastructure

Cyber Resilience

Requires short- and long-term retention capability;

- High snapshot frequency & fastest restore for short to medium term retention
- RPO policy governed snapshot frequency for medium to long term retention and fast restore

Cyber Resiliency/ “Safeguarded Copy”

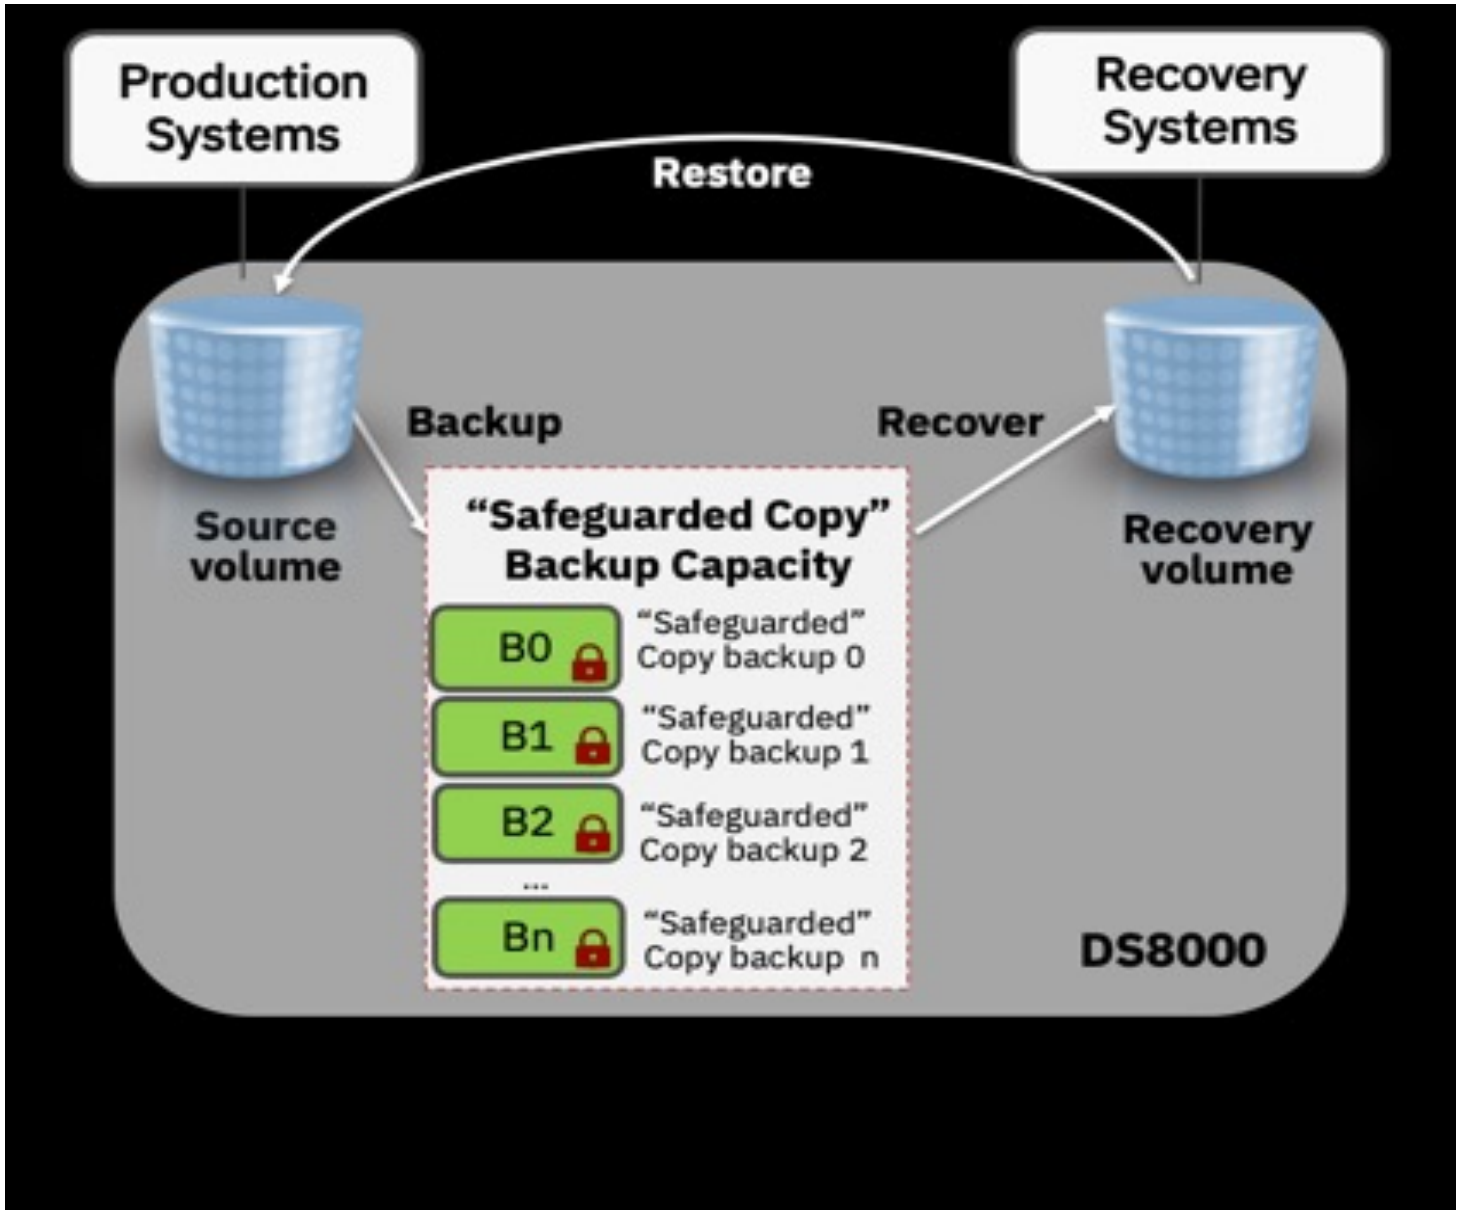

- “Safeguarded Copy” provides numerous protected and immutable copies per source volume, which are hidden and not accessible by any server
- The data can only be accessed after a ‘Safeguarded Backup” is recovered to a separate Recovery volume
- Secure immutable copies are a core of the IBM Z Cyber Vault offering
  - Add “Safeguarded Copy” to existing machines or deploy a physically isolated Cyber Vault

“Safeguarded Copy” solution

- “Safeguarded Copy” solutions provide for “*potentially*” much better point-in-time data recovery
  - Taking more frequent backups with thin provisioning to limit the amount of Db2 log apply which leads to faster data recovery
  - Faster and operationally simpler to restore the complete system in either an isolated “forensic” environment or in the primary environment
    - “Forensic” environment = made available to application teams to investigate data corruption/inconsistency and reconcile against the current production data
- Each “safeguarded” copy in the vault is I/O crash consistent
  - Data is **not *application transaction consistent***
- Danger in recovering individual datasets and subset of datasets unless super confident about understanding application and data dependencies
  - Avoid “jagged” edge in terms of time consistency across related applications and objects

“Safeguarded Copy” solution …

- Db2 crash restart (system) or Db2 RECOVER LOG ONLY (pageset/partition or dataset) is required to make any data extracted from the vault, application transaction consistent
  - Careful considerations need to be taken into account if the objective is to perform a Db2 system restore with using RESTORE SYSTEM LOG ONLY e.g., RBLP
    - Additional capabilities available in Db2 VNext
- Must be super confident about understanding application and data dependencies when recovering individual datasets and subset of datasets
- Solutions designed for recovering individual dataset, subset of datasets, whole system and building a “forensic” environment need to be tested, validated and regularly practiced to make sure they are in correct working order
- Taking Db2 image copy backups is still strongly recommended
  - Drive Db2 data recovery of last resort
  - Need to go back beyond 24–48 hours to support problem determination and advanced data recovery
  - Intra-page integrity checking performed when taking daily image copies
  - Db2 image copy backups are still required as input to some IBM and vendor Db2 tools
  - Db2 Image copy backups need to taken for Db2 LOG NO event

Summary

- Need to design for high performance and reduced elapsed time
  - Plan, design intelligently, stress test and optimise
    - Prioritise most critical applications
    - Understand application and data interdependencies
    - Design for parallel recovery jobs
    - Optimised utilisation of technical configuration
    - Intelligent creation and scheduling of recovery jobs
  - Design for DASD-based recovery for optimal performance
  - Practice regularly
- Applications and data life cycle also have a role to play…
  - Separate active/operational data from inactive/historical data
  - Perform regular aggressive archiving to historical
  - Allow application toleration of unavailable historical data
  - Look at creating 'fire walls' between applications

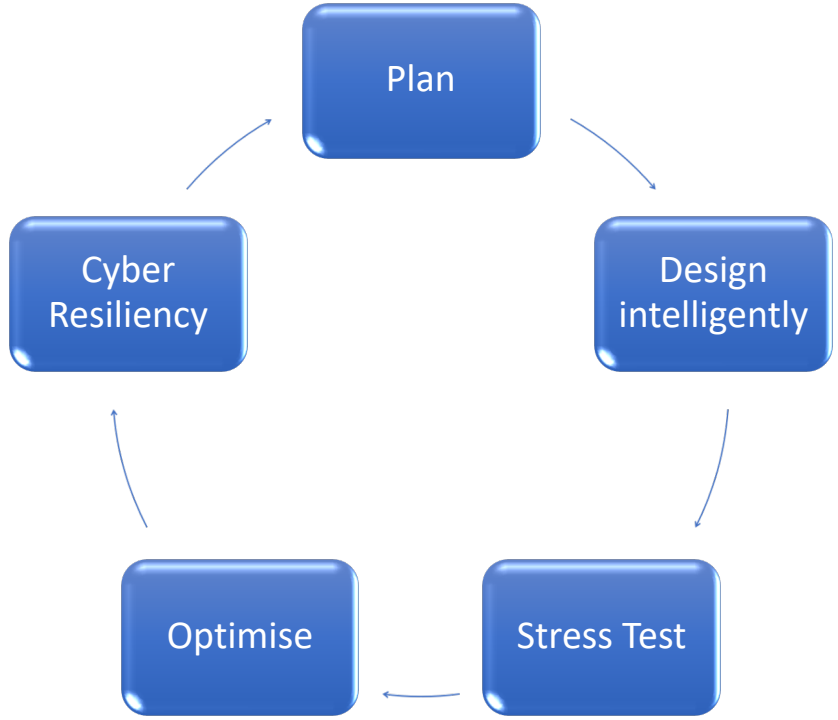

IBM®

Data and AI

The steps and techniques needed for mass recovery including Db2 recovery from cyber-attacks

Anthony Ciabattoni
Executive IT Specialist
Db2 for z/OS Development SWAT Team
aciabattoni@ibm.com

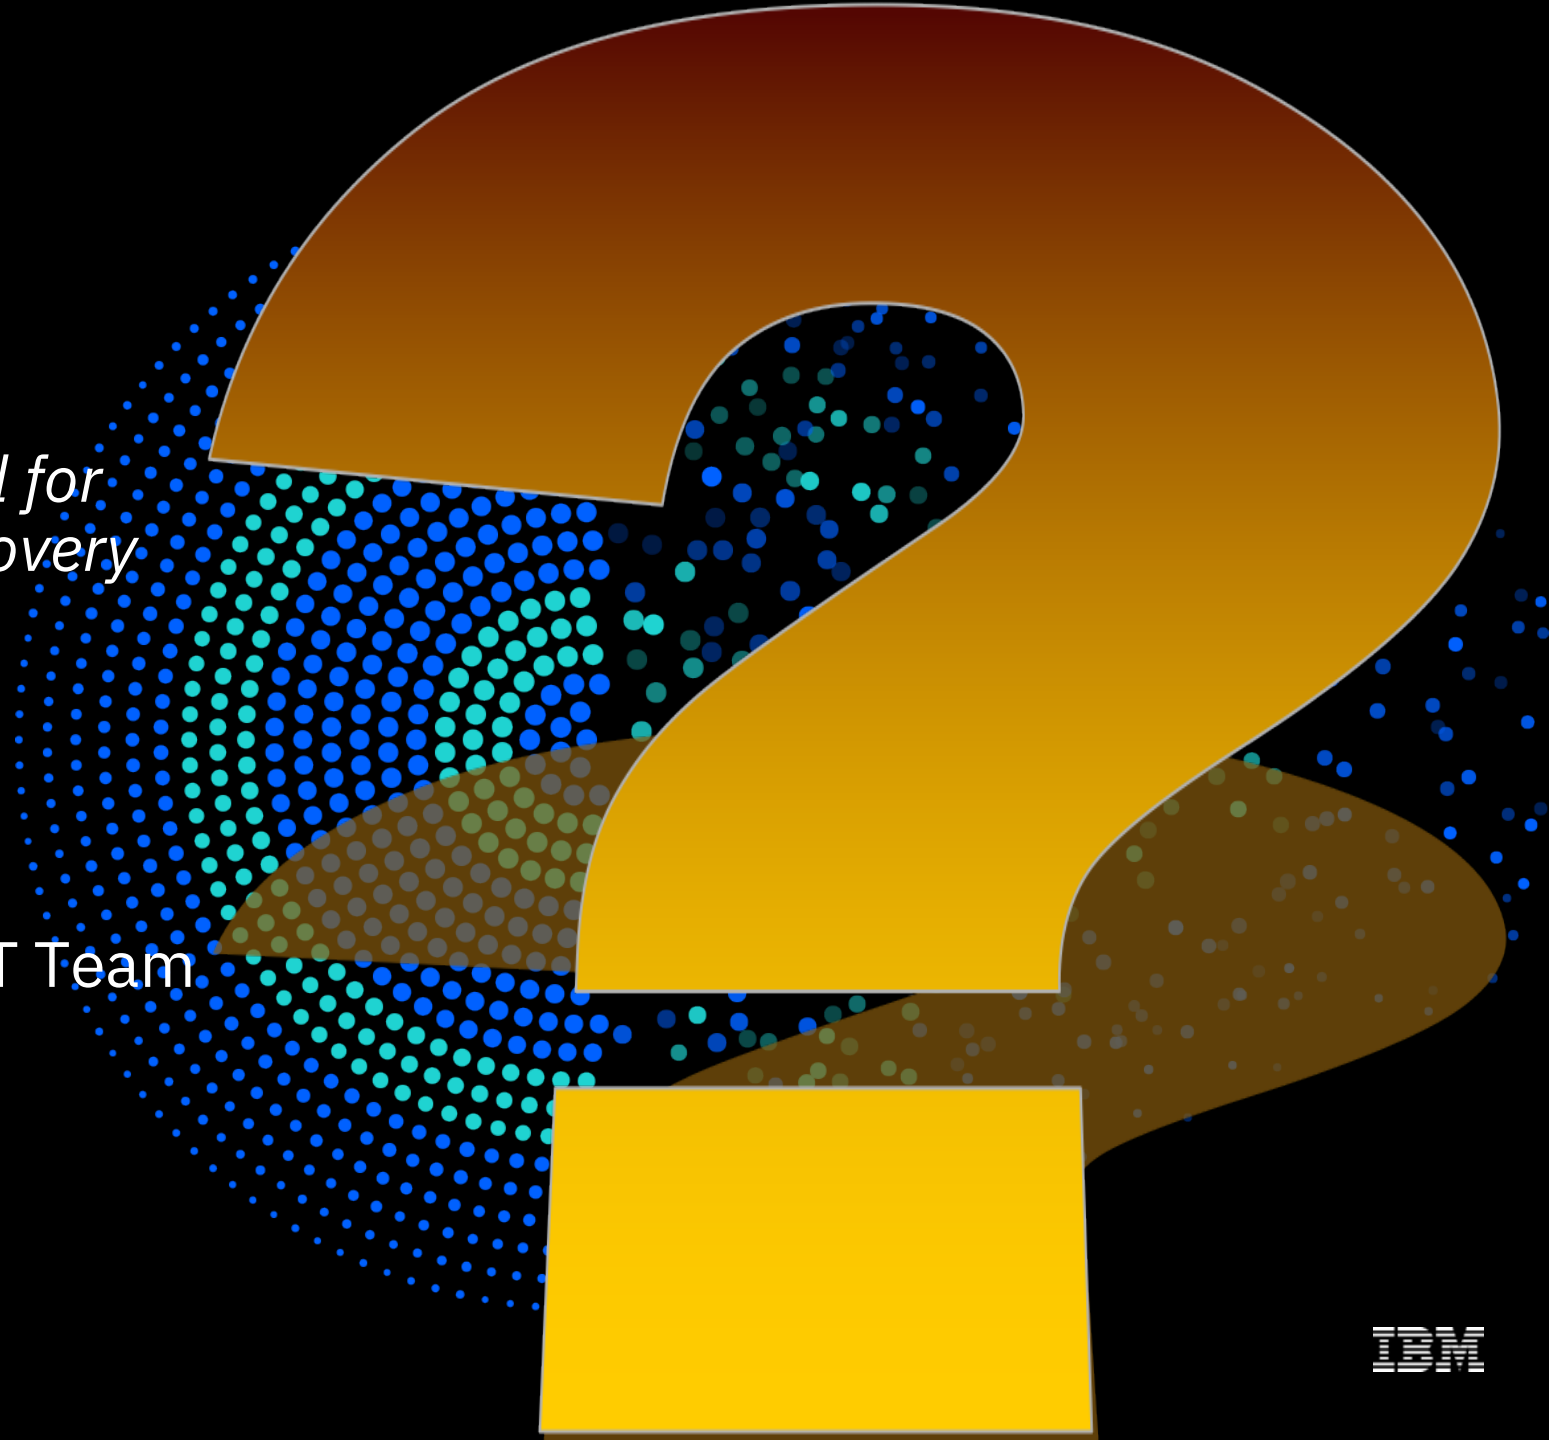

IBM