

IBM Data Tech Summit

Oct 3-5, 2023

IBM Silicon Valley Lab

New Dimension of Resiliency and Recovery using Db2, Cyber Vault and Safeguarded copies

Anthony Ciabattoni
IBM Db2 for z/OS SWAT Team


October 5, 8:30 – 9:30



Agenda

- Cyber Resiliency
- Cyber Vault overview
- Safeguarded Copies
- Recovery preparation/process
- Recovery Options
- Data Validation
- Forensic Analysis
- Surgical Recovery
- Questions

Guaranteed absolute resiliency or security is impossible



Systems need to be built for Cyber Resiliency

- ✓ The ability to continuously deliver the intended outcome despite any adverse event or attacks
- ✓ Do everything you can to prevent downtime and attacks, plus minimize the impact and potential loss when an event does happen

Sixty percent of businesses victimized by a cyber attack go out of business within six months.¹



\$4.54M

Average cost of a ransomware attack, not including the cost of the ransom itself.²



29 Days


Savings in response time for those with extended detection and response (XDR) technologies.²

\$1M

Average difference in cost where remote work was a factor in causing the breach versus when it wasn't.²

83%

Of organizations studied have had more than one data breach.²



¹CNBC <https://www.cnn.com/2019/10/13/cyberattacks-cost-small-companies-200k-putting-many-out-of-business.html>

²Cost of a Data Breach Report 2022 <https://www.ibm.com/reports/data-breach>

Cyber resiliency solutions should handle a wide range of possible scenarios to reduce the risk of financial losses

Cyber threats to enterprise data are increasing from a range of different sources

*External Malware Infection
External Hacking
Insider Threats*

Depending on the platform different risks are seen as more or less likely

Privileged Insider

For core systems running on IBM Z, many organizations believe the greatest risk

Similar loss or corruption of data is still also possible from other causes

*Application error
Operational error*

Furthermore, the pandemic and massive increase in remote working changes some of the risk vectors as access to mainframe systems is more likely from outside the organization-controlled networks

New Dimension of Resiliency is Required

Cyber resilience



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
 - Site Swap capabilities
 - “Gold Copies”/Snapshot copies
 - Various frequencies (daily/weekly)



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 and image copy datasets
 - Leading to incomplete recovery assets needed for rec
 - “Gold Copies”/Snapshot copies
 - Point-in-time copies
 - Only as current as the time they are executed
 - Will lead to data loss




Traditional resiliency solutions don't protect you from logical data corruption... more is needed

Protect your data with immutable point in time copies

+

Detect, Respond, and Recover faster with these critical capabilities



Data validation

Regular analytics on the data copy to provide early detection of a problem, or reassurance that everything is



Forensic analysis

Start a copy of the production systems from the copy and use it to investigate the problem and determine the recovery action



Surgical recovery

Extract data from the copy and logically restore back to the production environment



Catastrophic recovery

Recover the entire environment back to the point in time of the copy as this is the only recovery option



Offline backup

Copy the copy of the environment to offline media to provide a second layer of protection

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-medium term retention
- RPO policy governed snapshot frequency for medium to long term retention and fast restore

IBM Z Cyber Vault

Cyber Resiliency:


The ability to *anticipate, withstand, adapt to, and recover* from adverse conditions, events, stresses, attacks, or compromises on systems that use or are enabled by cyber resources

Principle Idea:

Reduce the time to recovery from days to minutes, by implementing a Data Corruption Protection solution as part of your disaster recovery strategy

Cyber Vault Environment:

- IBM DS8K with Safeguarded Copy provides immutable, consistent point-in-time copies of data
- GDPS LCP manages the creation, recovery, and restoration of the copies and provides automation to manage those processes
- IBM zSystems hardware and software provides a secure, isolated environment to perform data validation, forensic analysis, and create offline backups



IBM storage

Data volumes and active copies generated and maintained

DS8000 Safeguarded Copy

Immutable backups

TS7700 Virtual Tape with Encryption and/or WORM

Secure air-gapped data vault

IBM Z and Software

The only System with a 99.99999% availability

EAL 5+ certified IBM Cyber Vault for Z LPAR for validation, testing and forensics

Data monitoring, consistency and anomaly detection

Management Software

IBM Security solutions

IBM Services

IBM GDPS provides services, clustering technologies, and server and storage replication and automation

Logical Data Corruption (LCP) and Copy Services Manager (CSM) enhancements manage the entire recovery environment

IBM Lab Services risk assessment and deployment services

Logical corruption protection copies (SafeGuarded Copy)

Safeguarded Copies are secure, point-in-time copies of production data that can later be used for identification, repair, or replacement of production data that has been compromised by either cyber or internal attack or corrupted by system failures or human error



Source devices are where the protection copies are taken from. These could be production devices or taken from a HA/DR copy using data replication




Protection devices provide one or more logical protection copies and are not accessible by any system. Additional security measures aim to protect these from being modified or deleted due to user errors, malicious destruction or ransomware attacks



Recovery devices are used to logically restore back data to the production environment or to investigate a problem and determine what the recovery action should be

IBM Storage provides SafeGuarded Copy

- Prevent sensitive point in time copies of data from being modified or deleted due to errors, malicious destruction or ransomware attacks.
- Create up to **500** SafeGuarded Backups for a production volume stored in SafeGuarded Backup Capacity, which is not accessible to any server.
- The data is accessible only after a SafeGuarded Backup is recovered to a separate recovery volume.
- Recovery volumes are used with a data recovery system for:
 - Data validation
 - Forensic analysis
 - Restore production data




Recovery Point Objective (RPO)

NO Roll Forward




- Advantages
 - Most efficient
 - Operation simplicity
 - Traditionally established reusable disaster recovery/group restart procedures
- Disadvantages
 - Data loss is guaranteed
- Operational steps
 - The data is accessible only after a SafeGuarded Backup is recovered to a separate recovery volume IPL Cyber Vault Recovery LPARs
 - Db2 Group Restart
 - Data Validation



Recovery Point Objective (RPO)

Roll Forward

- Additional Recovery Assets are required
 - Db2 archive logs
 - Db2 LOG NO Image Copies
- Advantages
 - Minimize Db2 data loss
 - Use traditionally established reusable disaster recovery/group restart procedures (initial step)
- Disadvantages
 - Additional procedures and processes
- Operational steps
 - The data is accessible only after a SafeGuarded Backup is recovered to a separate recovery volume IPL Cyber Vault Recovery LPARs
 - Roll forward
 - RESTORE LOG ONLY
 - Db2 Group Restart
 - Data Validation



Preparation

- Recovery Point Objective (RPO) will determine the frequency of the execution of the preparation cycle
 - Installation needs to define RPO (potential data loss) and execute the preparation cycle accordingly
 - Assuming installation can only lose 1 hour of data the following preparation steps would need to be executed hourly
 - 1. Execute queries on SYSIBM.SYSCOPY to identify any Db2 LOG NO events e.g., REORG that would require image copies for system recovery to a point-in-time
 - 2. Issue the ARCHIVE LOG command on each member to create a new archive log pair
 - 3. Using DSNJU004 (print log map) capture Boot Strap Dataset (BSDS) for each member
 - 4. Print out each Boot Strap Dataset (BSDS) needed for diagnosis
 - 5. Execute Db2 stand-alone utility DSN1LOGP on last archive log for each member to find the last log record
 - Lowest Log Record Sequence Number (LRSN) value across all members will be used for the Db2 conditional restart with SYSPITR or SYSPITRT
 - 6. Implement a methodology to save away Db2 recovery assets that are more current than the last successful Safeguarded Copy
 - Needed for a more current recovery point (RPO)
 - Assets need to be transmitted and available in the Cyber Vault environment
 - Db2 Archive Logs
 - Db2 BSDS captured as part of the most recent archive log
 - Image copies needed for a LOG NO event

Recovery Process ...

- Recondition BSDS
 - Delete BSDS datasets for each member
 - Re-allocate BSDS datasets for each member
 - REPRO BSDS from most current Db2 archived BSDS dataset into the new re-allocated BSDS dataset
- Defining a point-in-time for Db2 recovery
 - Using DSNJU004 (print log map) print out the Boot Strap Dataset (BSDS) for each member
 - Needed for diagnostics
 - Use DSNJU003 (change log inventory) to set the Db2 conditional restart record (CRESTART) to set the Db2 recovery point (just prior to the corruption)
 - “CRESTART CREATE SYSPITR ...”
 - SYSPITR needs to be set from the lowest Log Record Sequence Number (LRSN) value from all members to be used for the Db2 conditional restart (SYSPITR/SYSPITRT)
 - Step #5 on preparation slide

Recovery Options ...

- Db2 12 Recovery Roll Forward
 - Db2 12 RESTORE SYSTEM LOG ONLY
 - 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
 - A starting point and an ending point needs to be established
 - Recovery starting point
 - Recovery Base Log Point (RBLP)
 - Recovery ending point
 - SYSPITRT
 - SYSPITR
 - Db2 12 RBLP is only updated
 - BACKUP SYSTEM
 - SET LOG SUSPEND/RESUME
 - If RBLP is not updated
 - No starting point for RESTORE SYSTEM LOG ONLY
 - Exhaustive archive search for the starting point will eventually fail
 - Db2 mass application object recovery

Db2 Mass Application Recovery

- Build mass application recovery jobs that exploit the capacity of the entire Db2 data sharing group
 - All recovery assets are available
 - Archive logs
 - Db2 image copies
 - LOG NO events e.g., Online REORG
 - Recovery parameters
 - RECOVER *&dbname.&tsname* LOGONLY
 - Defensive purposes use SCOPE ALL option
 - Indexes need to be rebuilt
 - Maximum parallelism in the RESTORE phase
 - Copy archive logs that are on tape (VTS) to DASD for concurrent access
 - For partitioned tablespaces, use parallelism by part
 - LISTDEF utility statement with the PARTLEVEL option will build a list of partitions for an object and automatically handle partitions that are added or pruned
 - Use PARALLEL for parallel processing from image copies on DASD
 - Use PARALLEL(n) TAPEUNITS(n) for image copies stacked on tape
 - Optimal use of fast log apply (FLA)
 - ZPARM LOGAPSTG has been removed and is set internally to 510MB
 - Schedule up to 51 RECOVER jobs per Db2 subsystem
 - RECOVER a list of objects rather than individual objects
 - No more than 98 objects per RECOVER job for best results (1 partition = 1 object)
 - 20-30 objects per RECOVER job seems to be optimal for FLA use
 - Single pass of the recovery log for all objects in the list
 - Spread the jobs across all Db2 data sharing members


Recovery Options ...

- Db2 13 Recovery Roll Forward
 - Db2 13 RESTORE SYSTEM LOG ONLY
 - 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
 - Db2 13 provides additional RESTORE SYSTEM LOG ONLY capabilities
 - RBLP is updated by Db2 every 5 minutes, starting point always incremented
 - RESTORE SYSTEM LOG ONLY to point in time is now very appealing
 - Recovery starting point
 - Recovery Base Log Point (RBLP) continuously being incremented
 - Recovery ending point
 - SYSPITRT
 - SYSPITR
 - Recovery assets must be available
 - Archive logs
 - Db2 image copies
 - LOG NO events e.g., Online REORG

Recovery Process

- Cyber Vault recovery steps
 1. Start Db2 ACCESS(MAINT)
 2. Respond to WTOR/CRCR record
 3. Execute RESTORE SYSTEM utility with the LOGONLY option on one member
 4. START/STOP (bounce) Db2 in normal mode to remove ACCESS(MAINT)
 5. Issue -DIS UTIL command
 - If any outstanding utilities
 - Issue terminate with -TERM UTIL command to terminate any active utilities
 6. If RESTORE SYSTEM LOGONLY encountered a LOG NO event
 - Tablespace object will be placed in a RECOVER-Pending (RECP) status
 7. Execute the RECOVER UTILITY using the Db2 image copies produced during a LOG NO event
 - Index(s) will be placed in REBUILD-Pending (RBDP) status
 8. In subsequent step in the recovery process execute REBUILD INDEX utility to rebuild the index(s) associated with the tablespace that was just recovered


The sooner you identify the problem, the smaller the impact



- Repeatability and Automated
- Time Consistent Copy is clean
- System is operational

- What, when and how data was corrupted?
- Can't be automated
- Tools may help, application knowledge is required

- Execute Recovery Actions - Surgical or Catastrophic.
- Use existing templates and predefined procedures




Continuous Data Validation

- Stay one step ahead with continuous automated data validation
 - Data validation is the process of executing regular analytics
 - Identify a data corruption event and scope of corruption
 - Determine the appropriate recovery action
 - Performing corruption detection and validation processes against a copy of data is more practical than doing this in the live production environment
 - Valid data can be sent to offline media to have a reliable and isolated point-in-time copy

Continuous data validation allows the early detection of a problem or reassurance that a given protection copy is uncorrupted.



IBM Z Cyber Vault – Data validation concept



As often as possible

Type 1: IPL with production image

- At least one LPAR per Sysplex is necessary
 - Preferably the same configuration as production infrastructure
- System Recovery Boost Upgrade record used for one IPL per day
- Check Sysplex infrastructure

Type 2: Data Structure Validation

- Db2 restart (all data sharing group members), Utilities, Log analysis
- Db2 Catalog/Directory consistency checks
 - DSN1COPY
 - Check Utility
 - Consistency queries
- Perform CHECK INDEX on application objects
 - In table priority order
- Db2 log-based recoveries utilizing recovery assets that reside on VTS

Type 3: Data Content Validation

- Customer Application Program

If no issue found (optional): Create tape copy


Forensic Analysis

- Back trace a cyber attack by forensic analysis
 - Determine:
 - What data is corrupted (scope)
 - When the data corruption occurred
 - Which available protection copies is the last good one
 - Based on this analysis (corruption scope), it can be determined how to proceed:
 - Fix the corruption from the production environment
 - Extract and recover certain parts of the data from a valid backup copy (surgical recovery)
 - Restore the entire environment to a point-in-time that is known to be unaffected by the corruption (catastrophic recovery)

A forensic analysis identifies the cause and scope of a problem before deciding on a recovery action.



Forensic Analysis



Determine start of data corruption ...

- **IPL** one Safeguarded Copy after the other to find the last clean copy
- **Understand** the problem
 - Run specific data structure and data content analysis on all stored Safeguarded Copies until a “clean” copy is found.
 - Use database tools to analyze databases and logs to fully embrace the scope of the problem
- **Identify** steps forward
 - Create strategy for recovery dependent on availability of database image copy files

Surgical Data Recovery

- Surgical Recovery may be a faster method if only a small portion of the production data is corrupted and if consistency between current production data and the restored parts can be re-established
- Another case may occur if the last known good backup copy is too old to restore the complete environment. It may then be desirable to leave most of the production volumes in its present state, and just copy replacement data to correct actually corrupted data
- 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

Surgical recovery consists of the extraction of specific data from a valid copy and logically restore it back to the production environment.



Surgical Recovery - Scenarios

Surgical Recovery is rather complex and the execution is dependent mainly on which data is available where for restore and recovery. In case Surgical Recovery needs to be done, the first step is to identify the actual scenario.

1. Backups are available in Production

- Image Copy of database exist in the production environment

Test in Cyber Vault, recover in production (business as usual)

2. Backups are available in the Cyber Vault only

- Image Copy of database does not exist in the production environment
- Image Copies exist on DASD in the Cyber Vault environment

Recover and test in Cyber Vault, send to production

3. No Backups are available neither in Production nor in the Cyber Vault environment

- Image Copy of database does not exist in the production environment
- Image Copies do not exist on DASD in the Cyber Vault environment

Recover and test in Cyber Vault, send to production, ensure database integrity

Questions



Thank You

The image features the words "Thank You" in a bold, 3D, metallic grey font. The text is centered and sits on a highly reflective, white surface that creates a clear mirror image of the letters below. The background is a light grey and white geometric pattern of overlapping squares and triangles, some of which are semi-transparent, creating a layered, architectural effect.