# Protecting your Db2 for z/OS environment against Cyber Attacks

Jørn Thyssen
jthyssen@rocketsoftware.com
Senior Technical Staff Member, Rocket Software

SHARE
EDUCATE • NETWORK • INFLUENCE

# Agenda

01 **Challenges to maintain resiliency in Db2 for z/OS**
Why recovery is an exception, never a routine — and what the delay costs

02 **Addressing recovery challenges with recovery tooling**
Health checks, recovery plans, cost-based estimates and optimized JCL

03 **Cyber Vault from a Db2 for z/OS perspective**
Immutable copies, data validation, forensic analysis and recovery scope

04 **Strategies to optimize the Recovery Point Objective for Db2 for z/OS in Cyber Vault**
Roll forward with archive logs, BSDS and RESTORE SYSTEM LOGONLY

SHARE
EDUCATE • NETWORK • INFLUENCE

# Agenda

01 **Challenges to maintain resiliency in Db2 for z/OS**
Why recovery is an exception, never a routine — and what the delay costs

02 **Addressing recovery challenges with recovery tooling**
Health checks, recovery plans, cost-based estimates and optimized JCL

03 **Cyber Vault from a Db2 for z/OS perspective**
Immutable copies, data validation, forensic analysis and recovery scope

04 **Strategies to optimize the Recovery Point Objective for Db2 for z/OS in Cyber Vault**
Roll forward with archive logs, BSDS and RESTORE SYSTEM LOGONLY

# Maintaining Resiliency in Db2 for z/OS

**Resiliency is built before the storm — through redundancy, rehearsal and routine maintenance.**

## 01 Avoid single points of failure

- Data Sharing / Parallel Sysplex
- Coupling Facility redundancy
- GDPS
- Dual logging
- z/OS WLM workload distribution

## 02 Disaster recovery planning

- Tooling support to eliminate think time
  - Recovery tooling
- Practice recoveries and site failovers

## 03 Maintenance

- Online schema changes
- Rolling maintenance strategy

And many more...

SHARE
EDUCATE • NETWORK • INFLUENCE

# Agenda

01 **Challenges to maintain resiliency in Db2 for z/OS**
Why recovery is an exception, never a routine — and what the delay costs

02 **Addressing recovery challenges with recovery tooling**
Health checks, recovery plans, cost-based estimates and optimized JCL

03 **Cyber Vault from a Db2 for z/OS perspective**
Immutable copies, data validation, forensic analysis and recovery scope

04 **Strategies to optimize the Recovery Point Objective for Db2 for z/OS in Cyber Vault**
Roll forward with archive logs, BSDS and RESTORE SYSTEM LOGONLY

# Backup, Recovery and the cost of Availability

**What is the cost of availability to your business?**

Safeguarding business applications and data is critical

Many backup strategies incur significant ongoing cost to guarantee a certain level of recoverability

Backup processes are seldom adjusted with new technologies and strategies

Actual recovery is rare

**$5,600** **per minute**
average cost of IT downtime to the business

# Recovery Complexity in Db2 for z/OS

**01**

### You don’t do it every day

Recovery is an exception, never a routine.

**02**

### It’s always at an inconvenient time

Pressure is highest exactly when clarity matters most.

**03**

### Requires extensive research

Dropped object, application error, hardware failure, human error, bad schema change, etc.

**04**

### Research to determine scope

Working out what actually needs recovering takes time.

**05**

### Good tools are available with Db2

Recover, Check, DSN1COPY, DSN1LOGP, etc.

**06**

### Still hard to do — and seldom practiced

Capability that is never rehearsed cannot be relied on.

# What is your Recovery Time Objective (RTO)?

RECOVERY TIMELINE

- Validate jobs — some fail on missing recovery assets
- Job execution
- Can jobs run in parallel?
- Generate recovery jobs
- Determine best strategy
- RTO
- Determine scope of recovery
- Error detected

PRESSURE

- Management makes an appearance
- More calls from management
- Phone rings

**RTO NOT MET**
loop until recovery succeeds

**Can you improve recovery time and reduce ongoing backup cost at the same time?**

- What is your Recovery Time Objective?
- How often do you practice it?
- What does the delay actually cost?

**THINK TIME + EXECUTION TIME**
**= $$$**

# Intelligent recovery tooling

## Reduces the complexity

- Analyzes recovery resources
- Easy to choose recovery points
- Creates detailed recovery plans
- Provides cost-based recovery estimates
- Detects and recovers related objects
- Detects objects not updated since the last recovery point and will not be recovered
- **Utilizes new technologies and recovery strategies**
- Generates the optimized JCL
- Provides reporting for review and control

## Minimizes exposure and risk

- **Pro-actively creates recovery plans to validate the ability to recover**
- Health check functionality validates the ability to recover to selected points in time

**Validate the ability to recover**
before the incident, not during it.

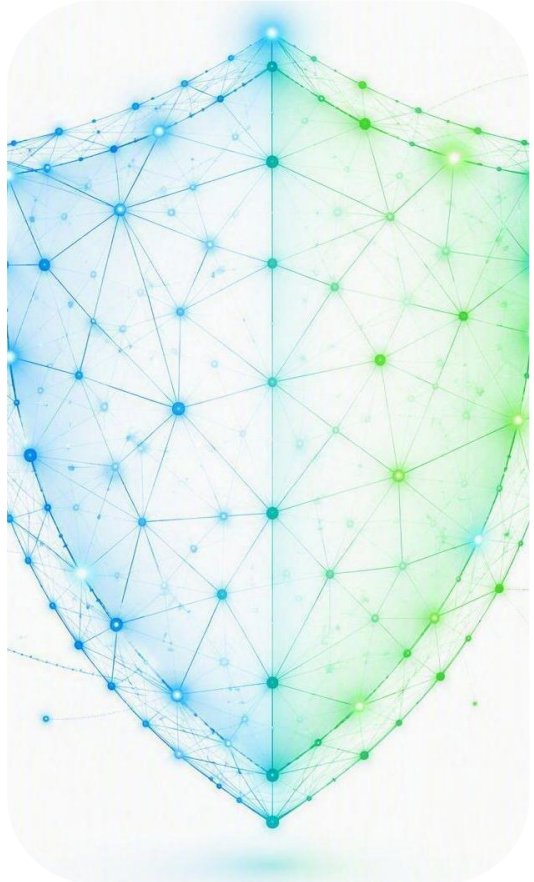

# Several types of recoveries

### Point in Time

Recover an object or a set of objects (application) to a point in time before data was corrupted.

### Dropped Object

Recover an object, all its dependent objects, DCL and rebinds if an object is accidentally dropped.

### Transaction

Back out a specific transaction that caused data corruption without taking a system outage.

### Redirected

Recover an application and its data from production to dev for diagnostic purposes.

### Application Group

Recover several applications in a specified order in the event of a system-wide outage caused by a hardware failure.

### System Level

Recover your entire Db2 system data and possibly the logs at the volume level using fast replication storage, allowing for quick recovery.

### Disaster

Recover your entire Db2 system to an offsite facility in the event of a disaster at the local site, using system level backups or image copies.

### One tool set

From a single object to an entire system — the same recovery discipline applies at every scope.

# Recovery-assisting information

01

## Recovery Health Check

Analyses existing recovery assets and provides information about recoverable and non-recoverable objects and timeframes.

02

## Recovery Suggestions

Suggests required metadata cleanup and informs about deprecated objects.

03

## Recovery Monitoring

Enables you to monitor recovery progress to provide informed projections about your recovery time objective.

Know what you can recover, what you cannot, and how long it will take — before you need to.

SHARE
EDUCATE • NETWORK • INFLUENCE

# Agenda

01 **Challenges to maintain resiliency in Db2 for z/OS**
Why recovery is an exception, never a routine — and what the delay costs

02 **Addressing recovery challenges with recovery tooling**
Health checks, recovery plans, cost-based estimates and optimized JCL

03 **Cyber Vault from a Db2 for z/OS perspective**
Immutable copies, data validation, forensic analysis and recovery scope

04 **Strategies to optimize the Recovery Point Objective for Db2 for z/OS in Cyber Vault**
Roll forward with archive logs, BSDS and RESTORE SYSTEM LOGONLY

# A closer look:
# Cyber Vault from a Db2 for z/OS perspective

Immutable Safeguarded Copies, repeated data validation, forensic analysis, and the choice between surgical and catastrophic recovery.

# Worldwide regulation

**United States**
- Interagency paper 'Sound Practices to Strengthen Operational Resilience'
- National Cybersecurity Strategy
- SEC Proposed Ruling for Cybersecurity Risk

**Europe**
- Digital Operational Resiliency Act (DORA)

**United Kin**
- FCA PS21
- statement
- Bank of England Operational

**Global**
- Basel Committee on Banking Supervision issued 'Principles for Operational Resilience' and 'Principles on

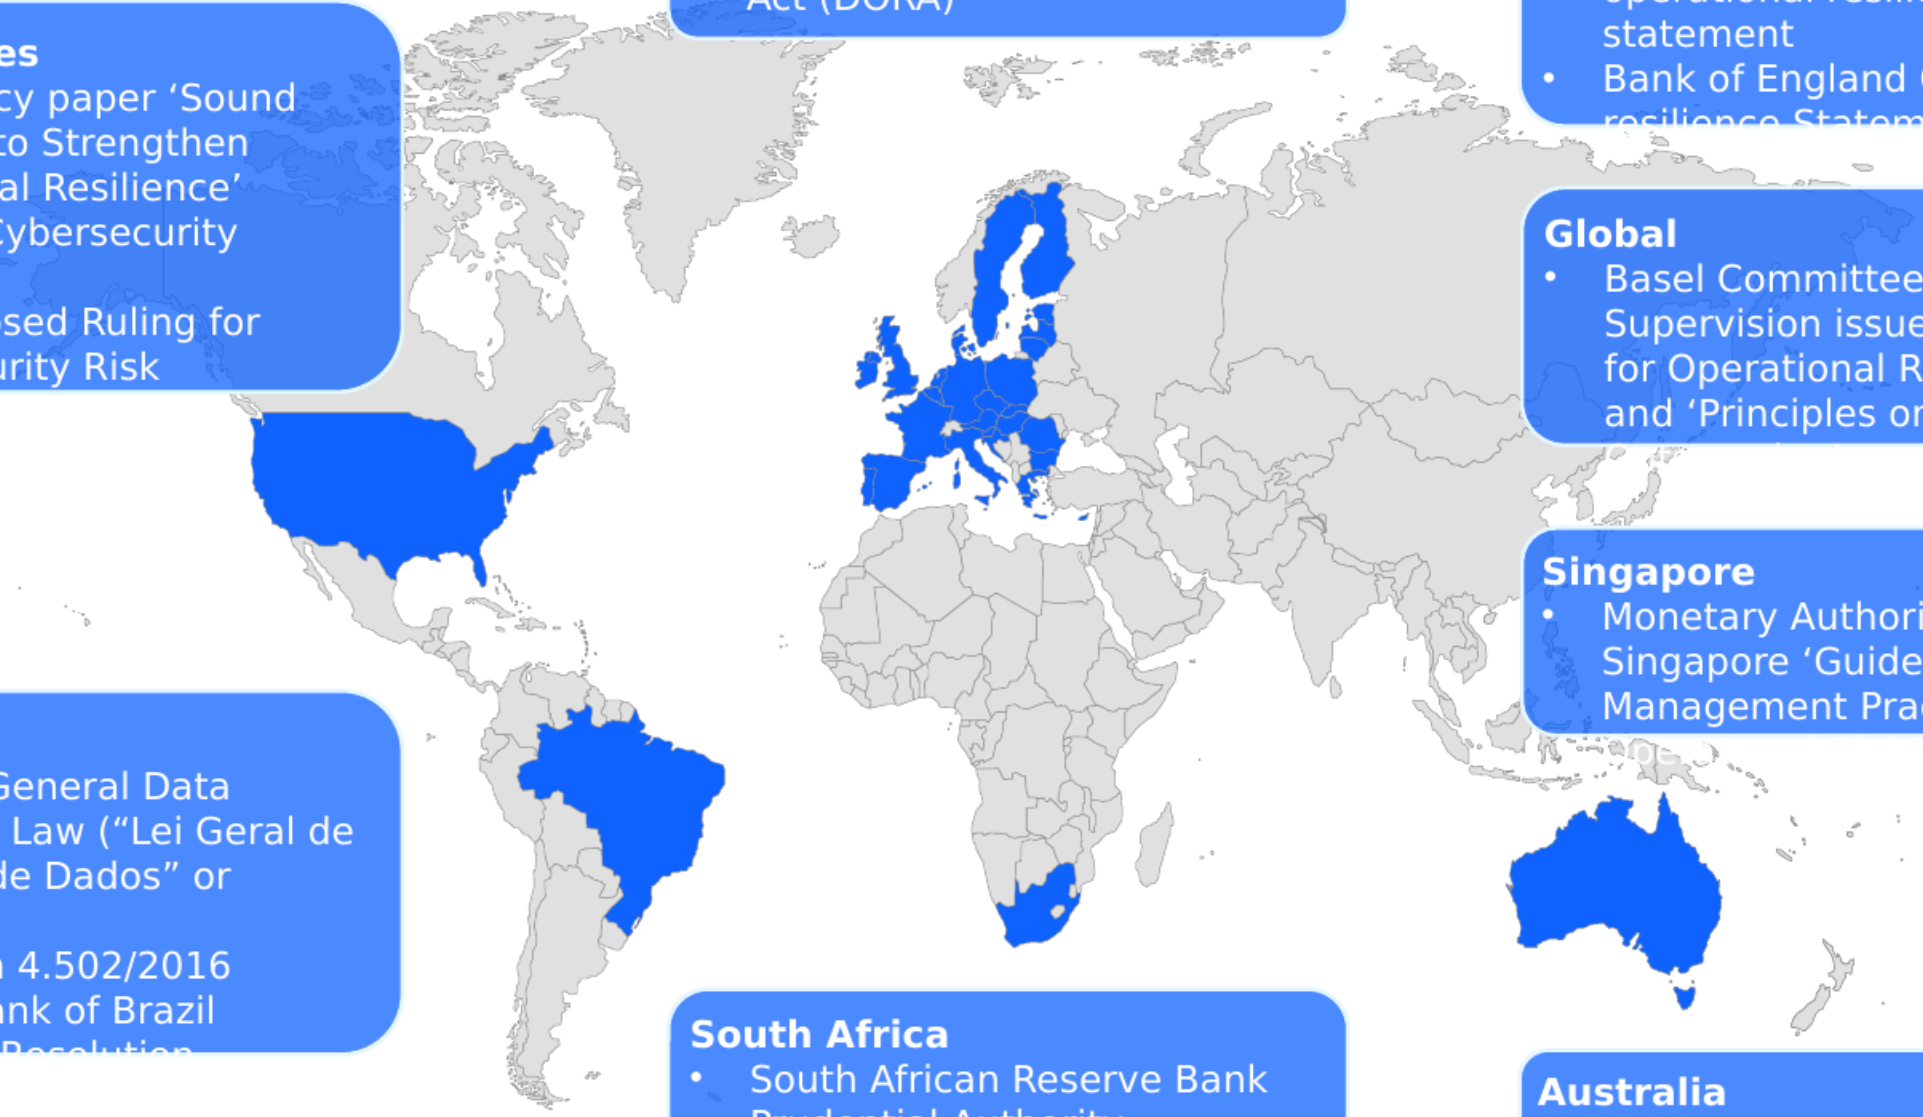

**Singapore**
- Monetary Authority of Singapore 'Guidelines on Risk Management Practices –

Brazil
- Brazilian General Data Protection Law ("Lei Geral de Proteção de Dados" or "LGPD")
- Resolution 4.502/2016
- Central Bank of Brazil

**South Africa**
- South African Reserve Bank Prudential Authority

**Australia**
- Prudential Standard CPS 230 -

# Traditional resiliency will not protect you from cyber attack

| | Traditional resiliency for HA/DR | What's required for Cyber Resiliency |
|---|---|---|
| **Replication** | Data is being replicated continuously but logical errors are also replicated instantaneously | Scheduled point in time copies stored in an isolated, secure location |
| **Error detection** | Immediate detection of system and application outages | Regular data analytics on point in time copies to validate data consistency |
| **Recovery points** | Single recovery point that likely will be compromised | Multiple recovery points |
| **Isolation** | All systems, storage and tape pools participate in the same logical system structure | Air gapped systems and storage so that logical errors and malicious intruders cannot propagate |
| **Recovery scope** | Continuous availability and disaster recovery | Forensic, surgical or catastrophic recovery capabilities |

# Recovery Point Objective (RPO)

## Roll Forward

- Additional Recovery Assets are required
  - Db2 archive logs
  - Db2 LOG NO Image Copies
  - ICF Catalog or Metadata
- Advantages
  - Minimize Db2 data loss
  - Use traditionally established reusable disaster recovery/group restart procedures (initial step)
- Operational opportunities
  - 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
  - Db2 Group Restart
  - Roll forward
    - RESTORE LOGONLY
  - Data Validation

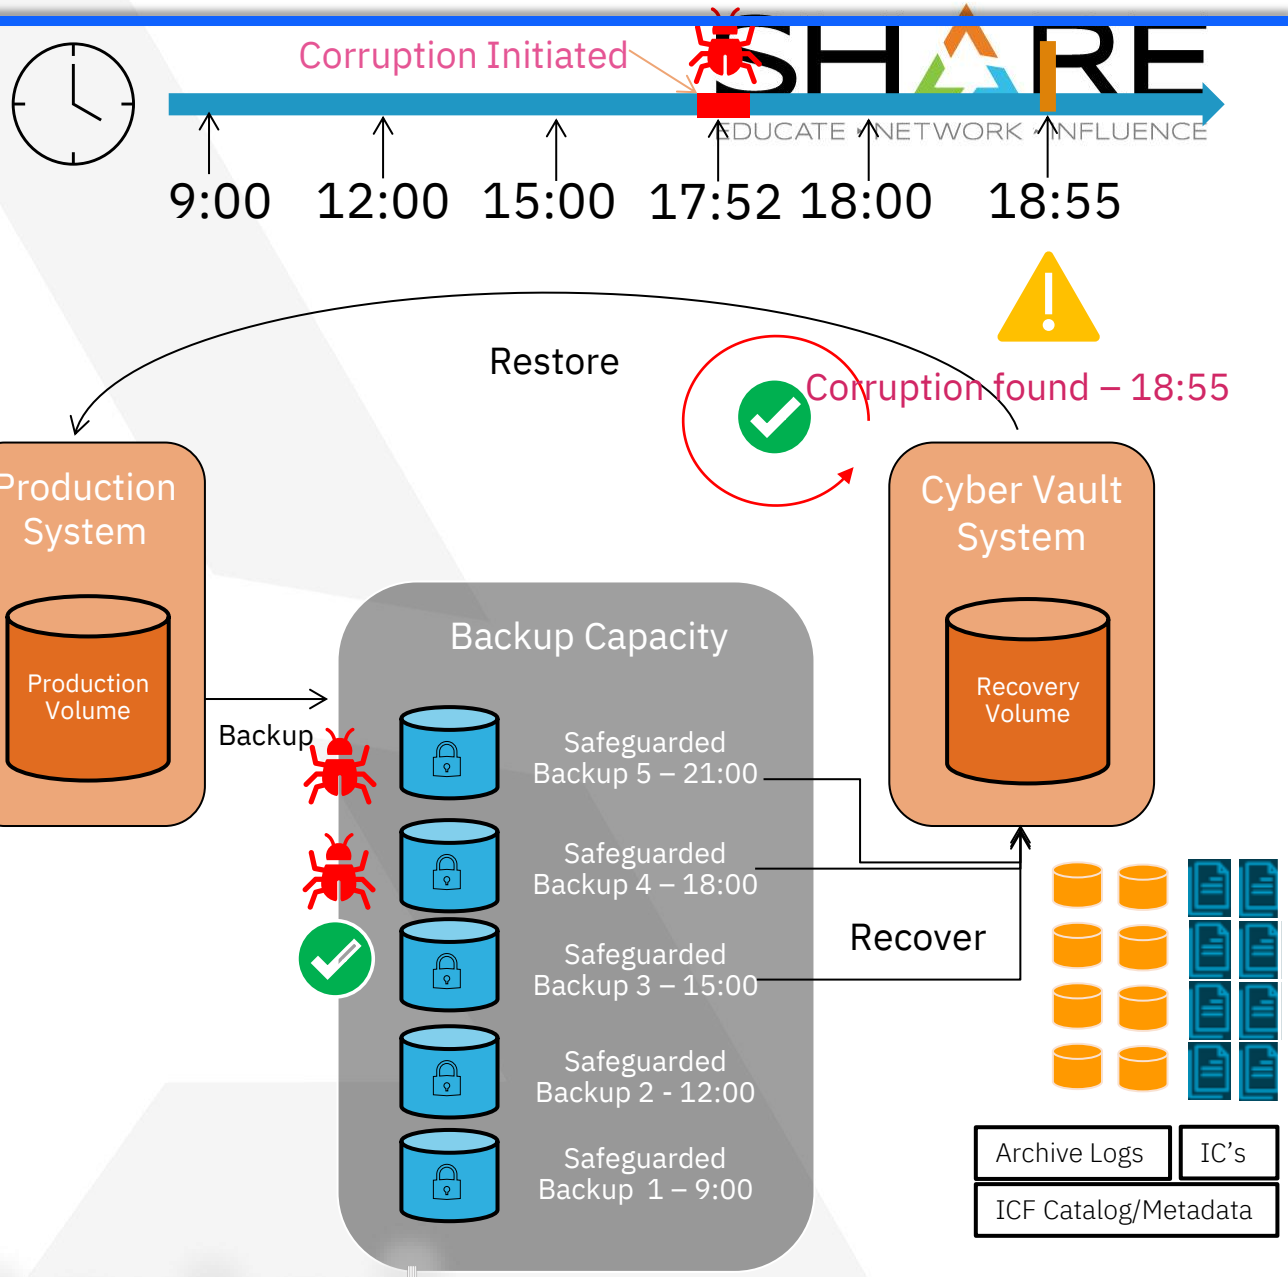

# Data validation

01

**Detect data corruption early, or validate that the copy is clear**

Data validation is the process of executing regular analytics to identify a data corruption situation and determine the most convenient 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.

# Data validation

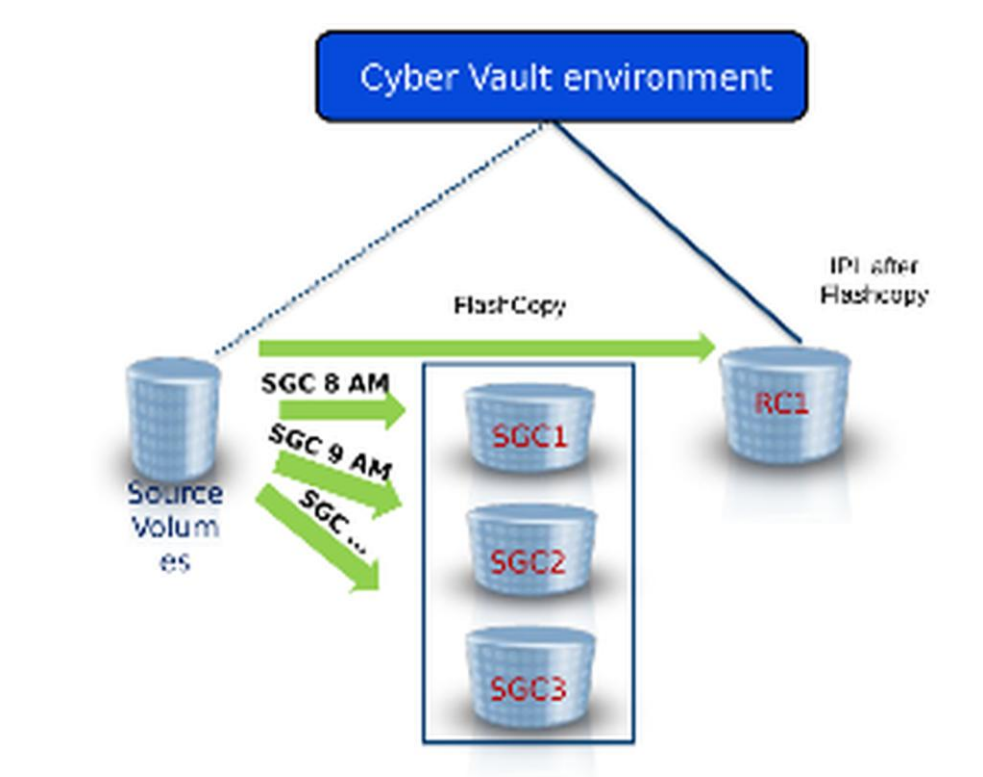


## Early identification of potential issues

### Type 1: Infrastructure Validation

- IPL of FlashCopy of production sysplex to Recovery Copy set (RC1)
- Check sysplex infrastructure and subsystem restart

### Type 2: Data Structure Validation

- **Db2 utilities** — CHECK DATA/INDEX, log analysis
- **IMS utilities** — pointer checker
- **Catalog tools** — VVDS, ICF, IDCAMS, ISV products
- **Storage and security** — VSAM Indexcheck and Datacheck, DFSMShsm, DFSMSrmm, RACF (IRRUT200), zSecure Audit
- **ISV software** — IWS, CA1, CA7 and others

### Type 3: Data Content Validation

- Customer application program

# Forensic analysis

02

**Investigate the problem and determine the best recovery action**

The forensic analysis determines what data is corrupted, when the corruption occurred, and which of the available protection copies is the last good one.

**Based on this analysis, it can be determined how to proceed:**

- Fix the corruption from within 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)

# Forensic analysis

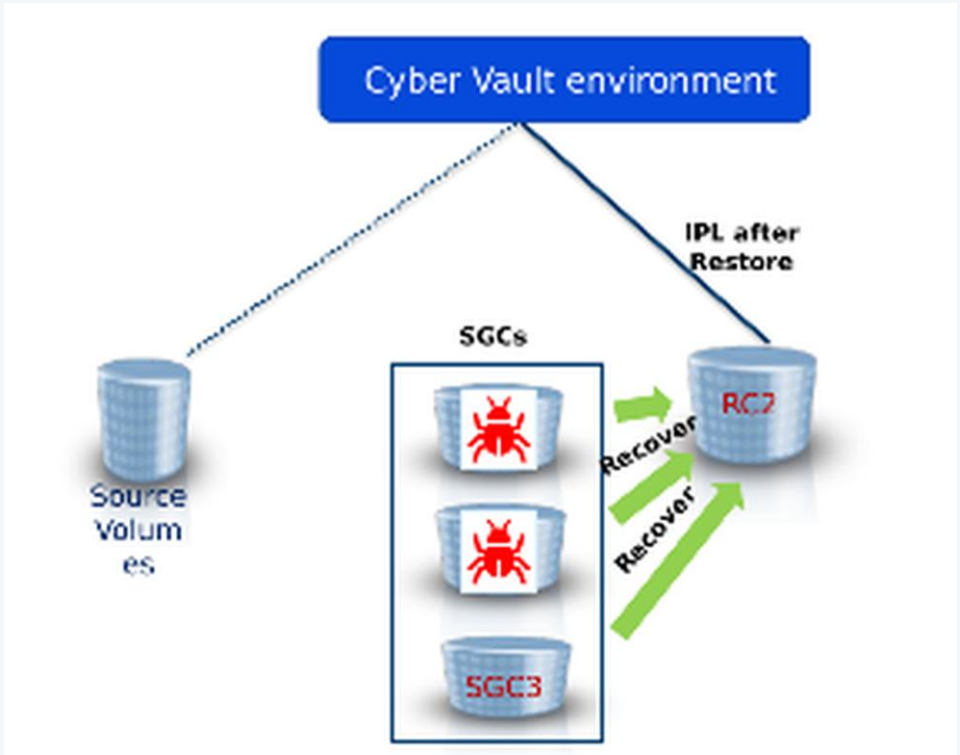


## Determine the start of data corruption

- **IPL** one Safeguarded Copy after the other to the Cyber Vault Recovery Copy set (RC2) 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 analyse databases and logs to fully embrace the scope of the problem
- **Identify** steps forward
  - Create a strategy for recovery dependent on the availability of database image copy files

# Surgical recovery

03

**Extract data from the copy and logically restore it back to production**

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 for this kind of recovery 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 their present state, and just copy replacement data to correct corrupted data.

# Surgical recovery

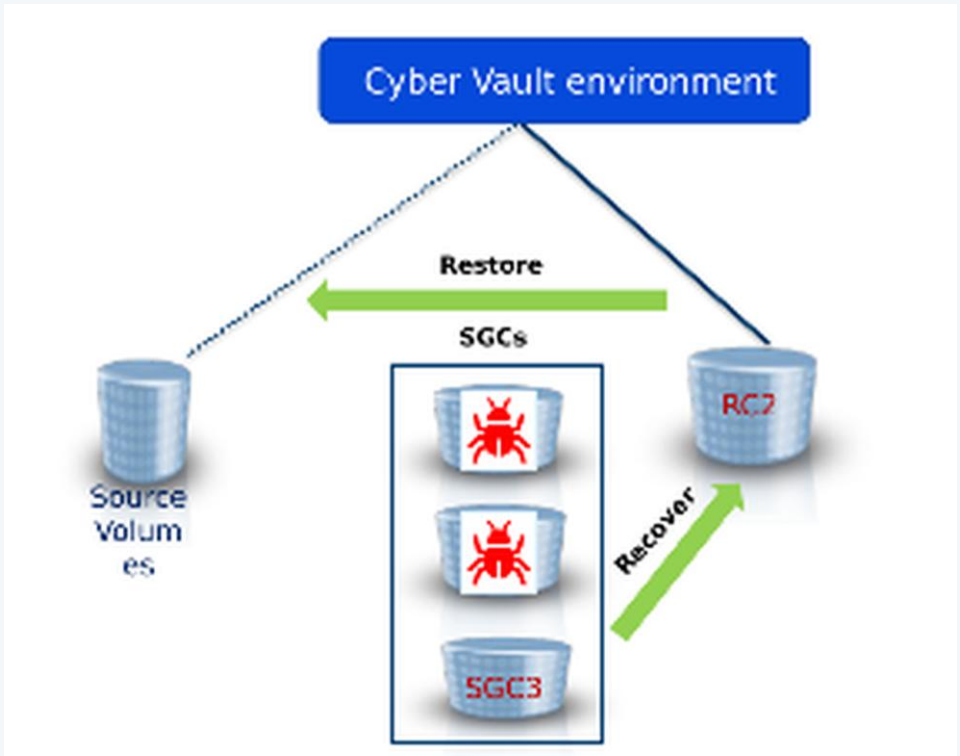


## Restore confirmed “good” objects

- **Identify** the specific point-in-time backup to be used as the restore point
- **Recover** the backup in the Cyber Vault environment (RC2)
- **Analyse** the backup to determine what data is required
- **Restore** only the required Db2 for z/OS objects back
- **Resolve** any inconsistency between backup data and production data

**Warning:** be careful when copying Db2 VSAM datasets — versioning, restrictive states and similar

# Catastrophic recovery

04

**Recover the entire environment back to a point in time copy**

In the case of massive corruption to all or most of the data in the environment, a catastrophic recovery needs to take place.

This means a full restore of a “clean” copy from Safeguarded Copy into the production environment needs to be done.

# Catastrophic recovery

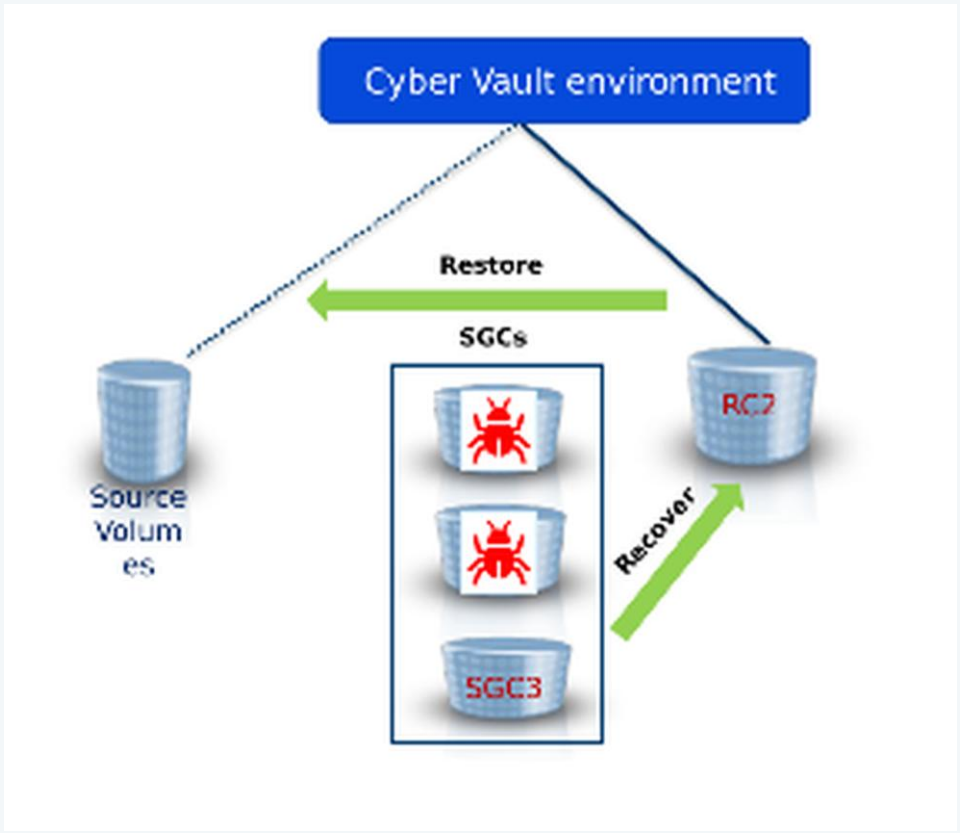


## Restore identified “good” data

- **Identify** the specific point-in-time backup to be used as the restore point
- **Recover** the backup in the Cyber Vault environment (RC2)
- **Incrementally restore** the backup into the production environment
- **Resolve** any inconsistency between the new production environment at T-X hours and any external dependencies

# Process summary


Plan recovery › Validate process › Copy to production › Back in production

## Backup and Validation

- Repeatable and automated
- Time consistent copy is clean
- System is operational

## Forensic Analysis

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

## Recovery

- Execute recovery actions — surgical or catastrophic
- Use existing templates and predefined procedures

SHARE
EDUCATE • NETWORK • INFLUENCE

# Agenda

01 **Challenges to maintain resiliency in Db2 for z/OS**
Why recovery is an exception, never a routine — and what the delay costs

02 **Addressing recovery challenges with recovery tooling**
Health checks, recovery plans, cost-based estimates and optimized JCL

03 **Cyber Vault from a Db2 for z/OS perspective**
Immutable copies, data validation, forensic analysis and recovery scope

04 **Strategies to optimize the Recovery Point Objective for Db2 for z/OS with Cyber Vault**
Roll forward with archive logs, BSDS and RESTORE SYSTEM LOGONLY

# Recovery tooling

**All existing features** of your recovery tooling can be used in a restored Db2 for z/OS subsystem in a **Cyber Vault recovery environment**.

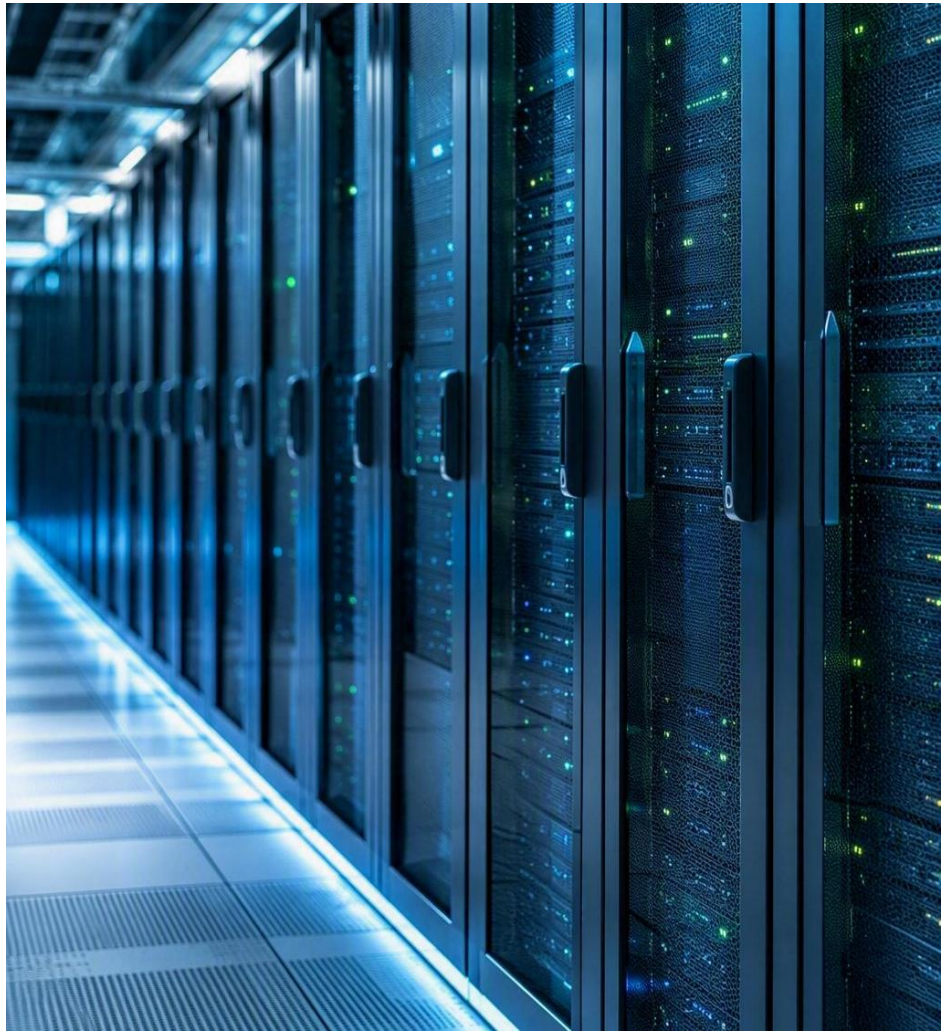

# Problem Statement

### What customers expect

Cyber Vault solutions provide powerful solutions to recover entire z/OS systems to prior points in time in the event of a cyber attack. Many customers desire the ability to restore Db2 subsystems and applications as current as possible.

### What it costs the DBA

A Db2 system administrator must perform many steps and procedures to prepare for and execute Db2 system recovery roll forward to improve the recovery point objective and minimize data loss in a Cyber Vault recovery system.

# Recovery Point Objective (RPO)

## Roll Forward

- Additional Recovery Assets are required
  - Db2 archive logs
  - Db2 LOG NO Image Copies
  - ICF Catalog or Metadata
- Advantages
  - Minimize Db2 data loss
  - Use traditionally established reusable disaster recovery/group restart procedures (initial step)
- Operational opportunities
  - 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
  - Db2 Group Restart
  - Roll forward
    - RESTORE LOGONLY
  - Data Validation

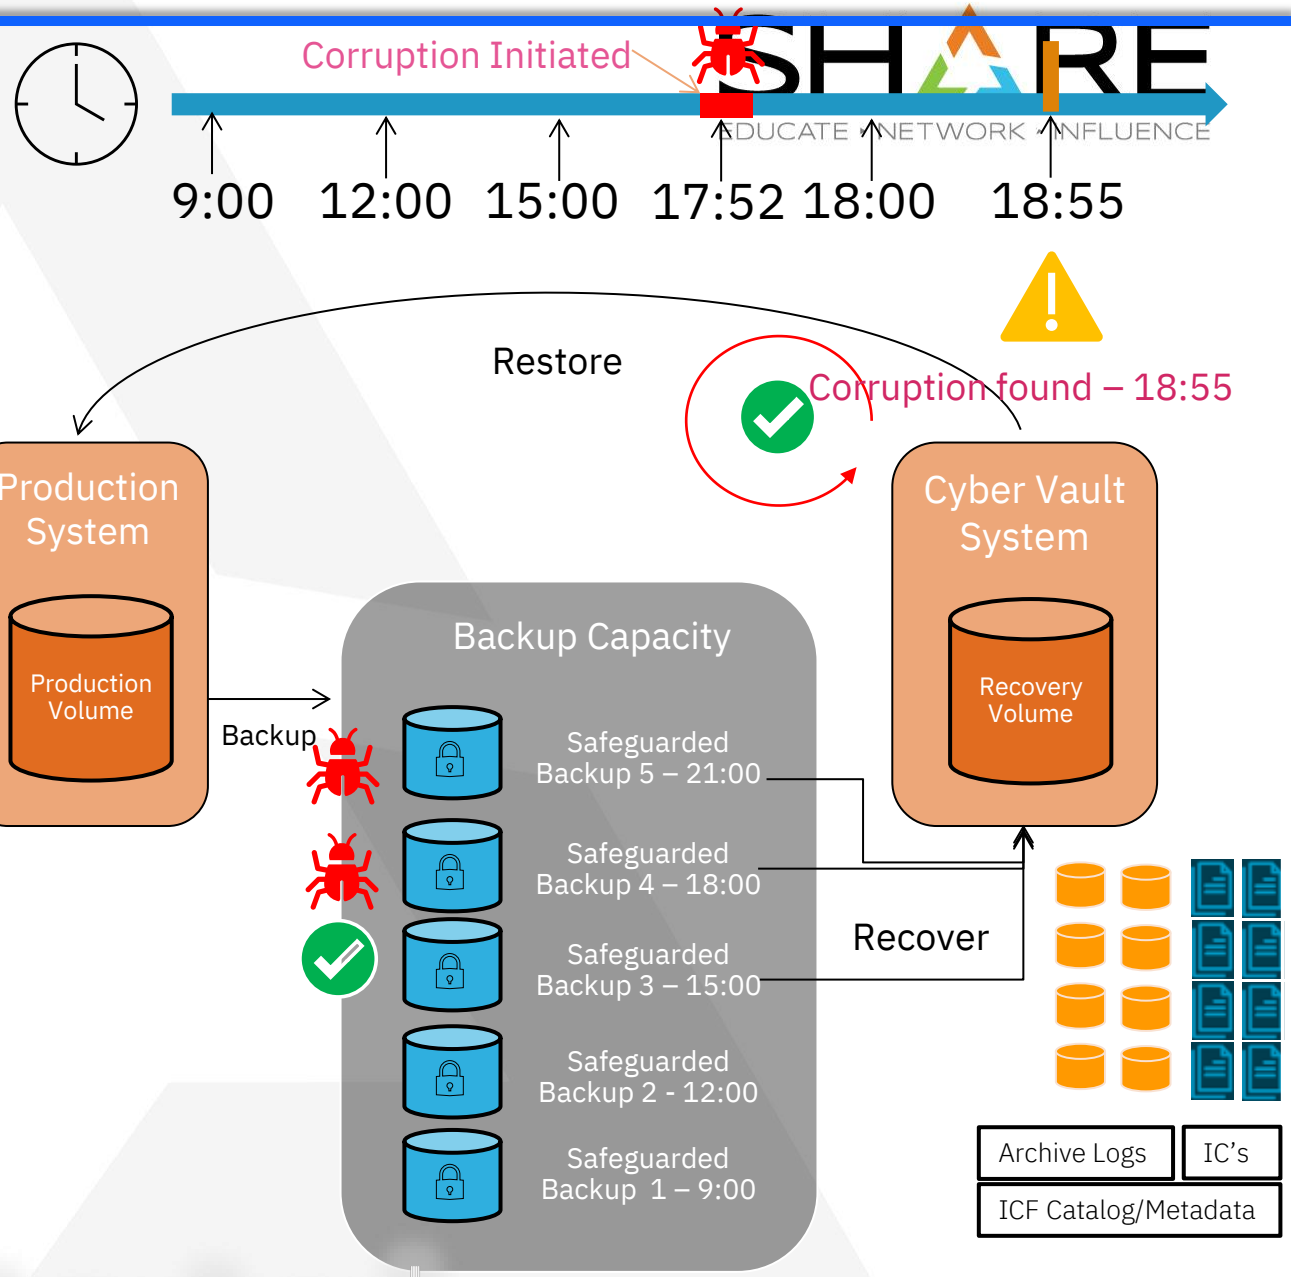

# Roll Forward Recovery

## Establish the recovery window

- **Db2 13 RESTORE SYSTEM LOGONLY**
  A recovery starting point and ending point must be established.
- **Starting point**
  Recovery Base Log Point (RBLP), continuously being incremented.
- **Ending point**
  SYSPITRT or SYSPITR.

## What Db2 13 adds

- **Additional RESTORE SYSTEM LOGONLY capabilities**
- RBLP is updated by Db2 every 5 minutes, so the starting point is always incremented.
- RESTORE SYSTEM LOGONLY to a point in time is now very appealing.

## Why object-level recovery falls short

- Db2 mass application object level recovery is technically an option, but for large, active systems the recovery time will likely be significantly longer than restoring an entire system with one log apply event.
- Lack of planning, intelligent design, optimization and large-scale recovery testing.
- Job scheduling not optimized for object size, update rate and the number and size of indexes, elongating recovery time.
- Db2 Fast Log Apply (FLA) not used efficiently.
- No prioritized list of application objects and inter-dependencies.

# Pain Points

## Many steps and procedures for the Db2 system administrator to perform

### STEP 1
### Preparation at the production system

- Set archive log frequency based on asset collection interval
- Periodically save recovery assets to immutable storage accessible at the recovery system
  - Archive logs with BSDS
  - Image copies for table spaces and indexes after LOG NO events

### STEP 2
### Preparation at the Cyber Vault recovery system after Safeguarded Copy is restored

- From immutable storage retrieve archive logs, BSDS and image copies that are more recent than the Safeguarded Copy that was restored
- For each Db2 member: replace the BSDS with the most recent BSDS, then run DSN1LOGP on the most recent archive log data set
- Identify the SYSPITR system recovery point: the lowest LRSN across the members from the DSN1LOGP output. This is the ending point for the roll forward (log apply)
- Adjust the SYSPITR system recovery point

### STEP 3
### Recovery roll forward at the Cyber Vault recovery system

- Conditional restart process with SYSPITR
- Run RESTORE SYSTEM LOGONLY
- Stop and start Db2
- Terminate any outstanding utilities
- Run RECOVER on table spaces and indexes in RECOVER-pending status
- Run REBUILD INDEX on indexes in REBUILD-pending status
- Validation of data

# Recovery Point Objective (RPO)

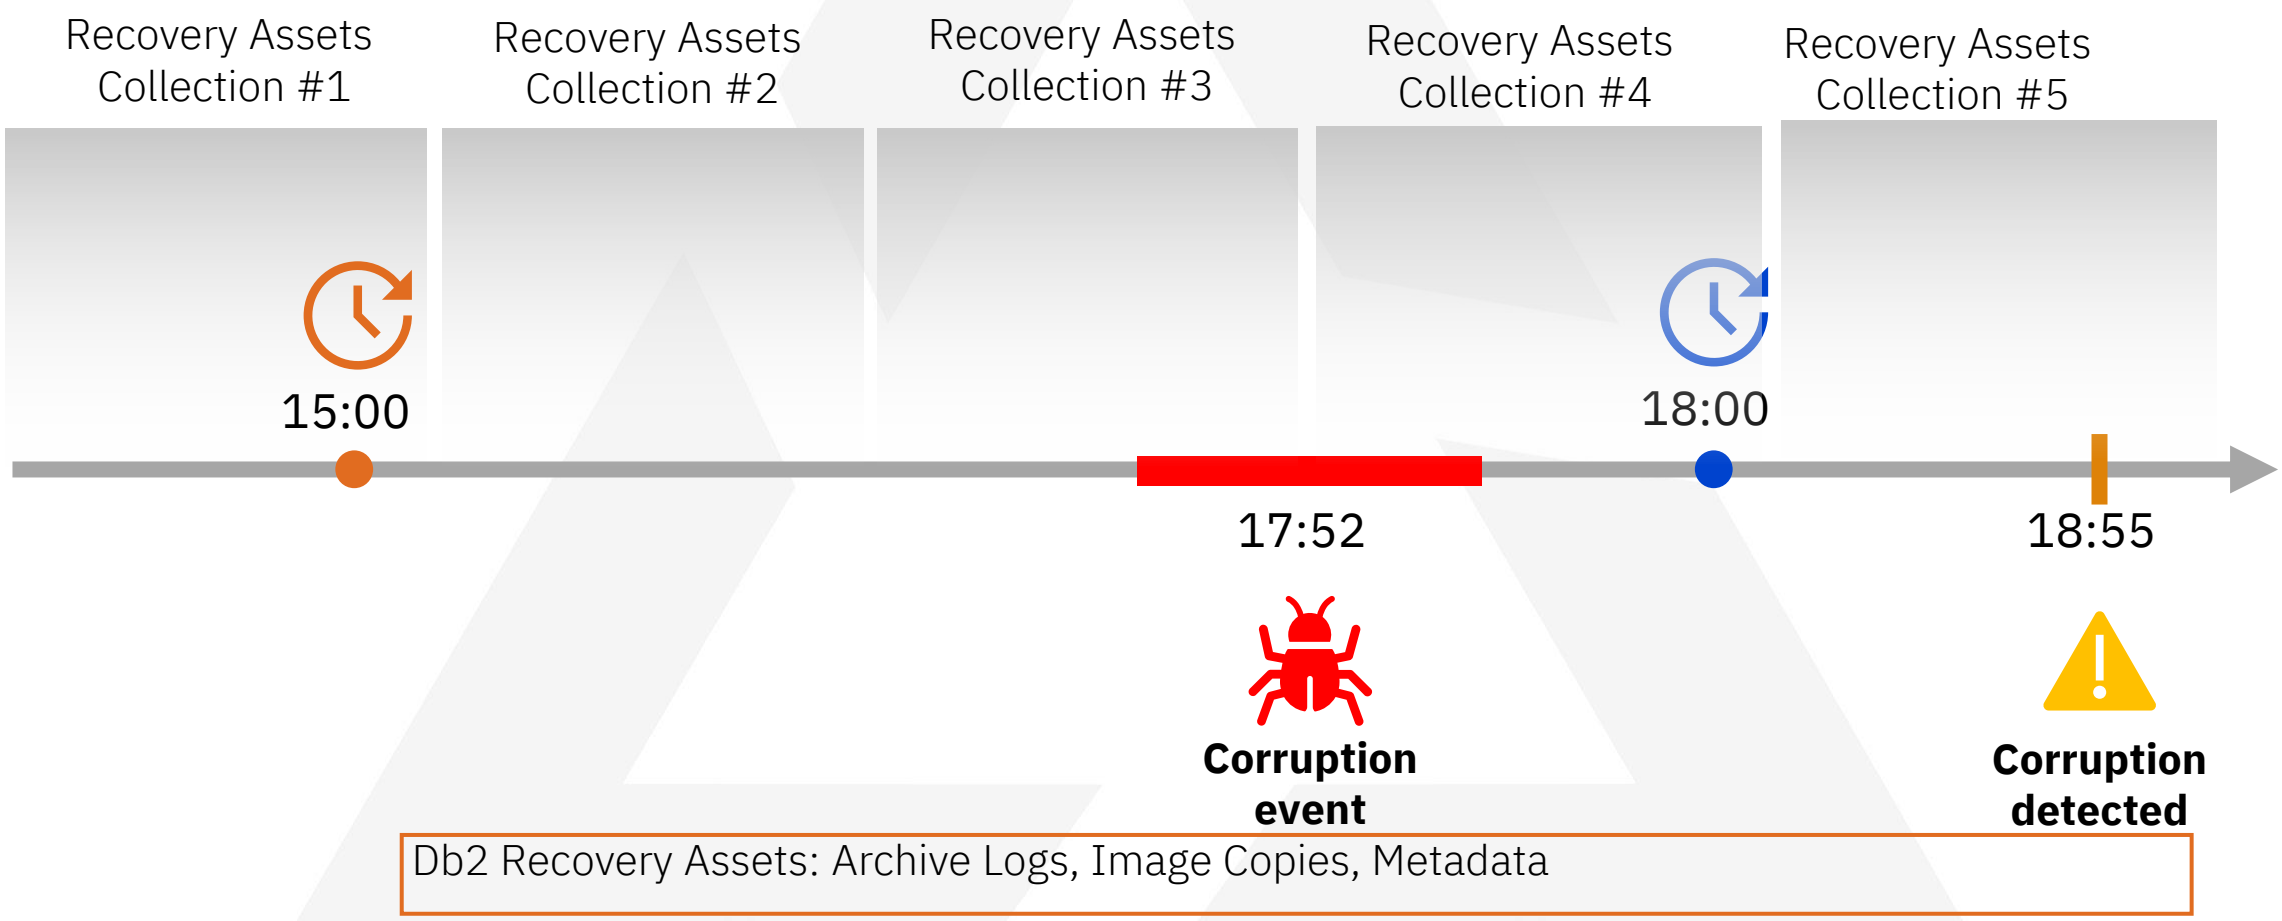

# Resource collection

## Production SYSPLEX

Recover technology started task

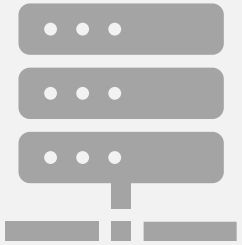

Archive logs

BSDS

Image Copies

Metadata

## Immutable Storage

S3 Cloud Storage

TS7700 LWORM Storage

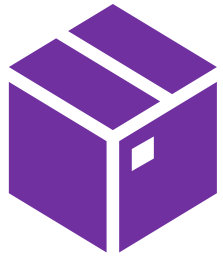

Separate SGC pool

# Resource collection

## Production SYSPLEX

Recover technology started task

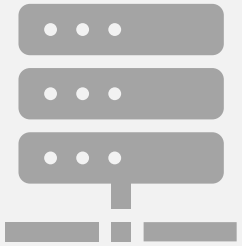

Archive logs

BSDS

Image Copies

Metadata

## Immutable Storage

S3 Cloud Storage

TS7700 LWORM Storage

Separate SGC pool

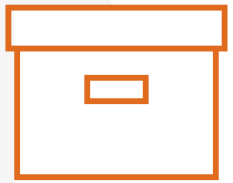

# Resource collection

## Production SYSPLEX

Recover technology started task

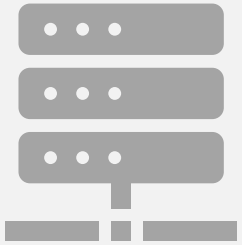

Archive logs

BSDS

Image Copies

Metadata

## Immutable Storage

S3 Cloud Storage

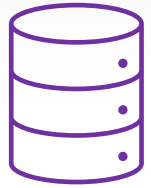

TS7700 LWORM Storage

Separate SGC pool

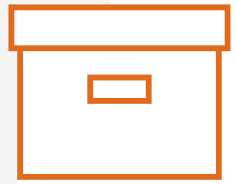

# Resource collection

## Production SYSPLEX

Recover technology started task

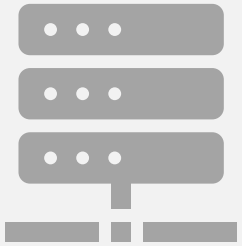

Archive logs

BSDS

Image Copies

Metadata

## Immutable Storage

S3 Cloud Storage

TS7700 LWORM Storage

Separate SGC pool

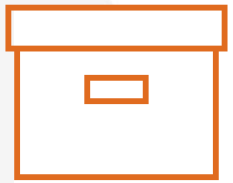

# Resource collection

## Production SYSPLEX

Recover technology started task

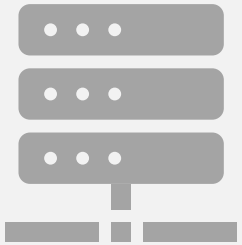

Archive logs

BSDS

Image Copies

Metadata

## Immutable Storage

S3 Cloud Storage

TS7700 LWORM Storage

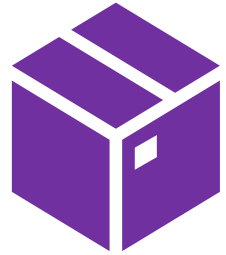

Separate SGC pool

SHARE
EDUCATE • NETWORK • INFLUENCE

# Summary

**Cyber resilience for Db2 for z/OS is a recovery discipline: isolate the copies, prove they are clean, then close the RPO gap with log roll-forward.**

## 01 Think time drives RTO

- Recovery is an exception, never a routine — and research, not execution, consumes the clock
- Downtime averages $5,600 per minute, so think time is a cost line
- Intelligent tooling removes it: health checks, recovery plans, cost-based estimates, dependent-object detection and optimized JCL
- Validate the ability to recover before the incident, not during it

## 02 HA/DR is not cyber resilience

- Continuous replication copies logical corruption instantly — one recovery point, already compromised
- Cyber resilience needs air-gapped, immutable Safeguarded Copies in an isolated location
- Multiple recovery points, not one — backups are inaccessible until recovered to a separate recovery volume
- Detection shifts from outage alerts to regular analytics on point-in-time copies

## 03 The Cyber Vault workflow

- Validation, then forensic analysis, then surgical or catastrophic recovery
- Validation runs at three levels: infrastructure IPL, data structure (CHECK DATA/INDEX, catalog and VSAM tools) and application content
- Forensics IPLs Safeguarded Copies in turn on the recovery copy set to find the last clean one
- Backup and validation automate well; forensics cannot — application knowledge is required

## 04 Closing the RPO gap

- A restored Safeguarded Copy is only as current as the copy — roll forward recovers the hours since
- That needs archive logs, BSDS, LOG NO image copies and ICF catalog metadata collected to immutable storage
- At the vault: replace the BSDS, run DSN1LOGP, take the lowest LRSN as SYSPITR, conditional restart, RESTORE SYSTEM LOGONLY, then RECOVER and REBUILD INDEX
- Db2 13 increments RBLP every five minutes, and a started task automates asset collection — cutting the manual steps

# Your feedback is important!

SHARE
EDUCATE • NETWORK • INFLUENCE

**Submit a session evaluation** for each session you attend:

SHARE mobile app -or- www.share.org/evaluation