IBM Z Cyber Vault and IMS

—

Tracy Dean, Product Manager,
IMS Tools and z/VM Tools
tld1@us.ibm.com

August 2023

Logical data corruption

- Hardware components are working as expected
- Data becomes destroyed or corrupted on a content level, including
  - Deletion
  - Encryption
  - Selective manipulation
- Cannot be prevented with traditional HA/DR solutions
  - HA/DR is not content-aware
  - Continuous replication solutions quickly propagate any content level corruption to all copies
- Undetected data corruption, also known as **silent data corruption**, results in the most dangerous errors as there is no indication that the data is incorrect.

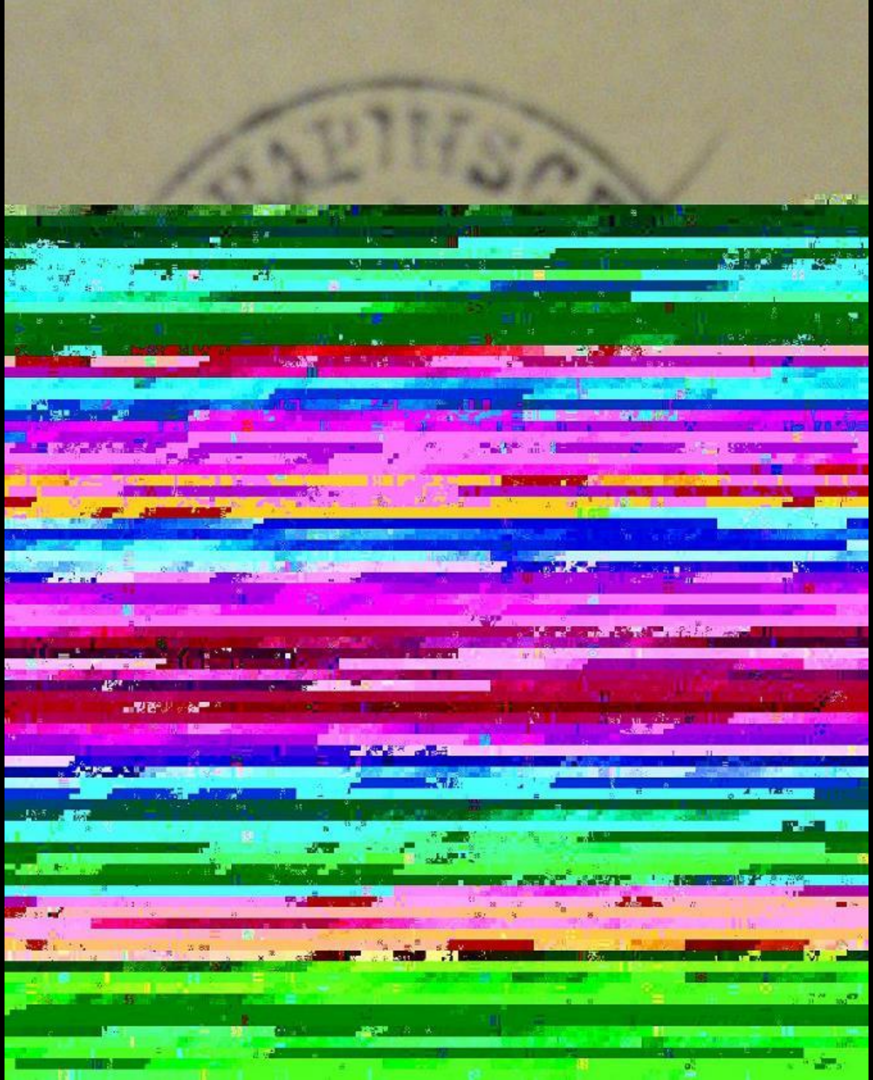

Why traditional resiliency solutions will not protect you from logical data corruption

	What you have	What is required
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 validation 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 can not propagate
Recovery scope	Continuous availability and disaster recovery	Forensic, surgical or catastrophic recovery capabilities

IBM Z Cyber Vault

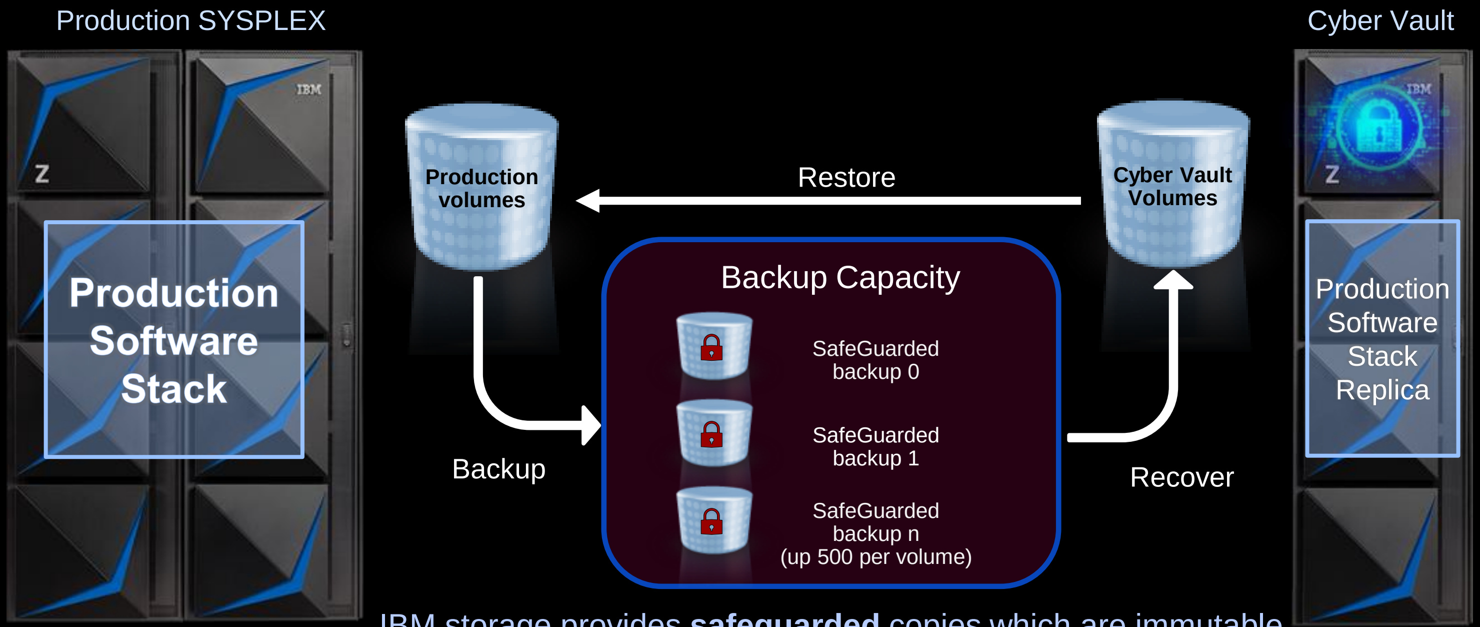

1. Establish analysis
2. Determine data validity
3. If valid, stop and document results
4. Run diagnostic tools
5. Perform forensic analysis
6. Perform recovery

IBM storage provides **safeguarded** copies which are immutable

IBM Z Cyber Vault
Focus areas for IMS

Data and data structure validation

Forensic analysis

Surgical recovery

Catastrophic recovery

Offline backup

Validation - Forensic analysis – Surgical recovery in the Z Cyber Vault

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

- What, when and how data was corrupted
- Can not be automated
- Tools may help, application knowledge required

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

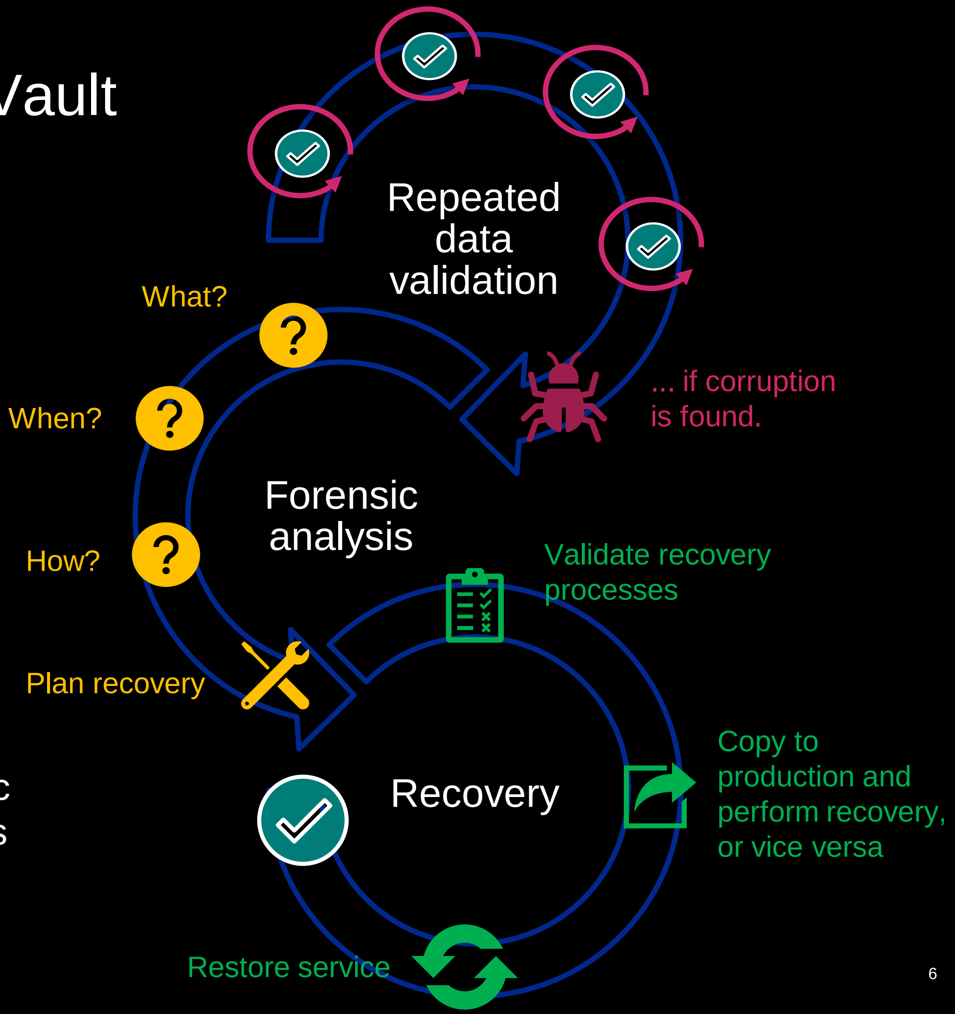
Repeated data validation
What?
When?
How?
... if corruption is found.
Forensic analysis
Validate recovery processes
Plan recovery
Recovery
Copy to production and perform recovery, or vice versa
Restore service

IBM Z Cyber Vault – data validation

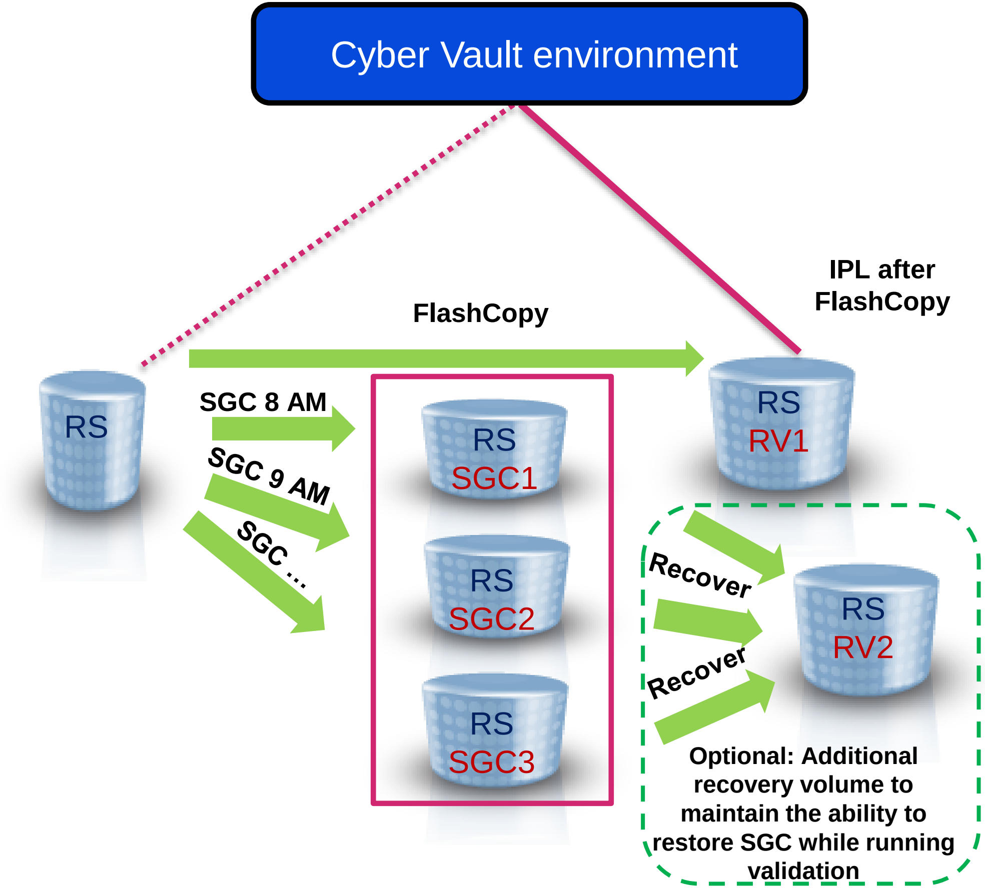

As often as possible ……

Type 1: IPL the FlashCopy of production image on recovery volume (RV) at recovery site (RS)

At least one logical partition (LPAR) per sysplex is required

- System Recovery Boost can be used up to 12 times in any 24 hour period
- Check sysplex infrastructure

Type 2: Data Structure Validation

- Restart IMS
  - If data sharing, must bring up all IMS systems in data sharing group
    - Can be done one at a time in single LPAR in Cyber Vault environment if needed
  - Need Db2 up if same unit of work includes both IMS and Db2 updates
- Run pointer checking
- Validate resources (image copies, logs, etc.) are available for recovery if needed
- Other z/OS validation

Type 3: Data Content Validation

- Other IBM products
- Customer application program

If no issue found (optional): Create tape copy

IBM Z Cyber Vault – data validation

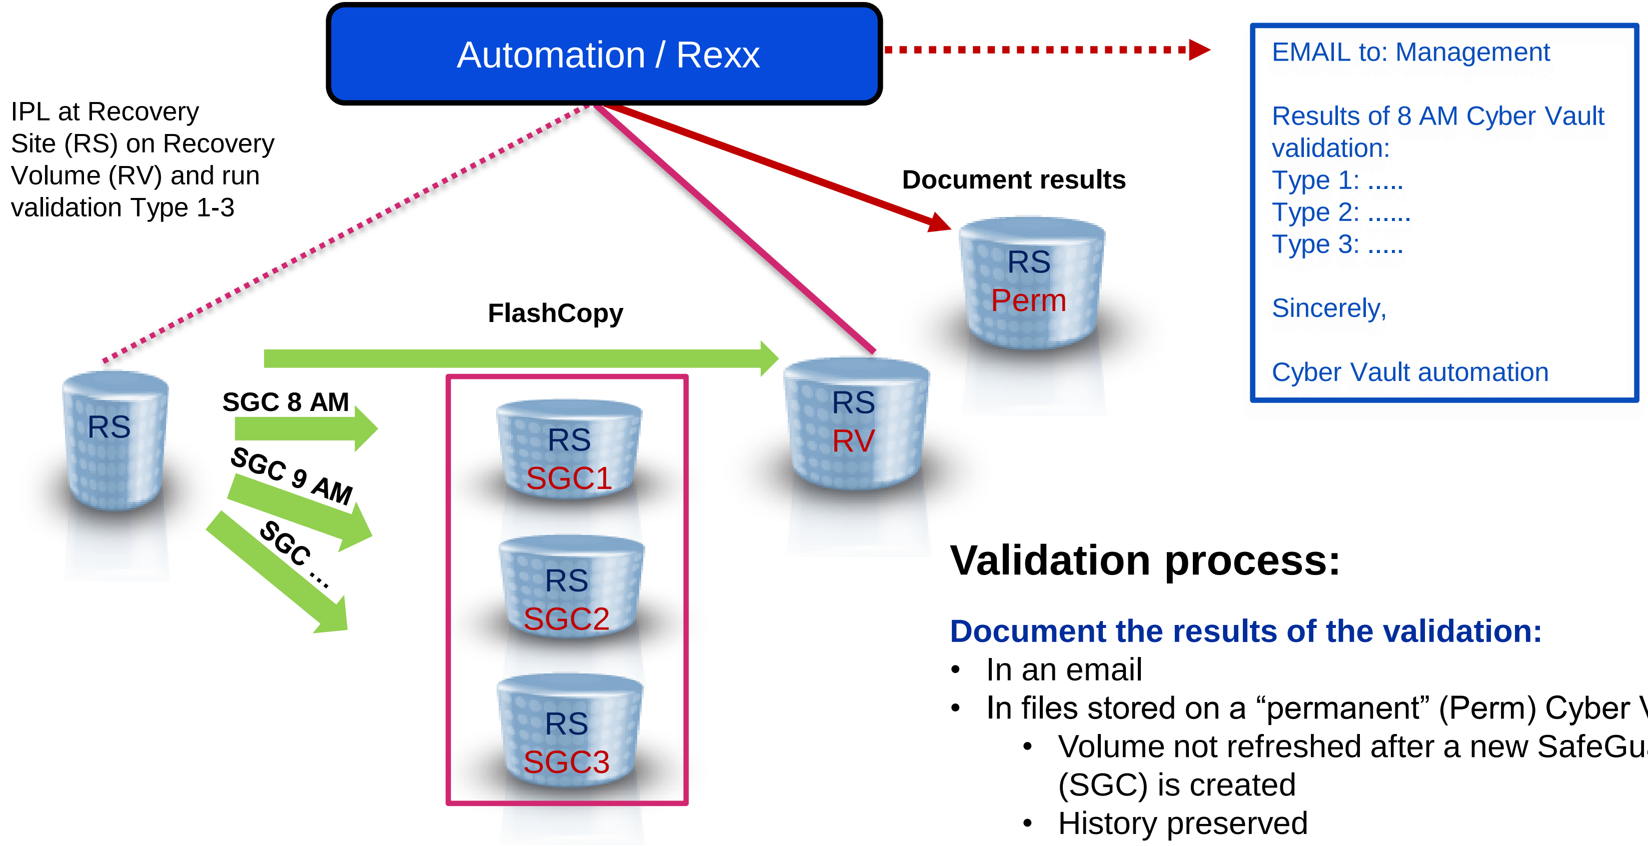

EMAIL to: Management

Results of 8 AM Cyber Vault validation:
Type 1: .....
Type 2: ......
Type 3: .....

Sincerely,

Cyber Vault automation

Validation process:

Document the results of the validation:

- In an email
- In files stored on a “permanent” (Perm) Cyber Vault volume
  - Volume not refreshed after a new SafeGuarded Copy (SGC) is created
  - History preserved

IMS data structure validation

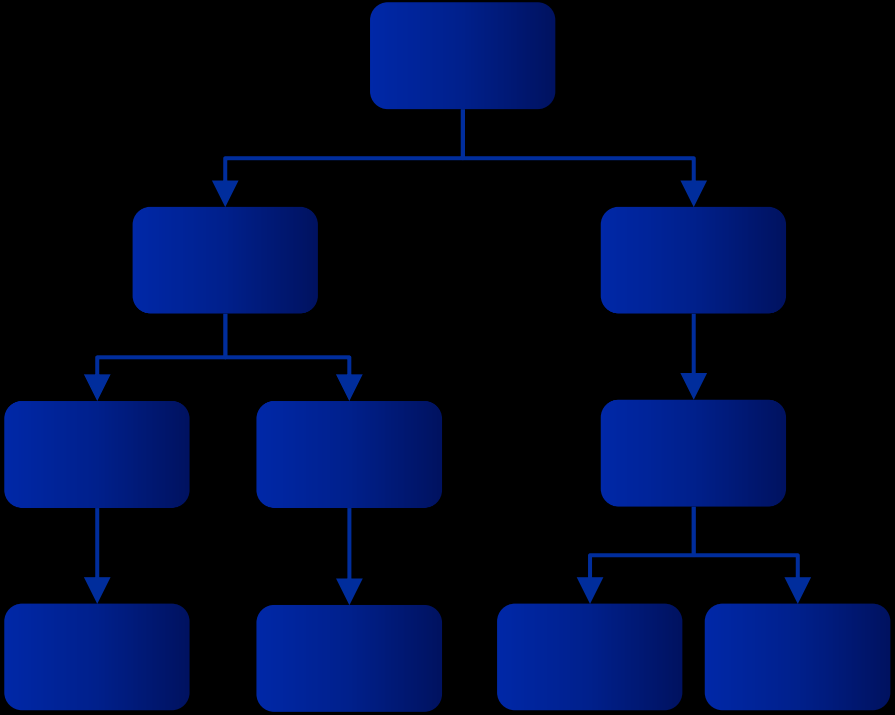

Pointer checking to validate database structure

- Find and report on pointer errors
- Detect changes in database characteristics such as size and number of segments
- IBM solutions
  - **IMS High Performance Pointer Checker** for full function databases
  - **IMS Fast Path Solution Pack** for fast path databases

Recovery readiness

- Verify assets needed for recovery are available
- IBM solution
  - **IMS Recovery Solution Pack** for all IMS databases
  - **IMS Tools Base** for web UI showing recovery readiness exceptions

Forensic analysis in general

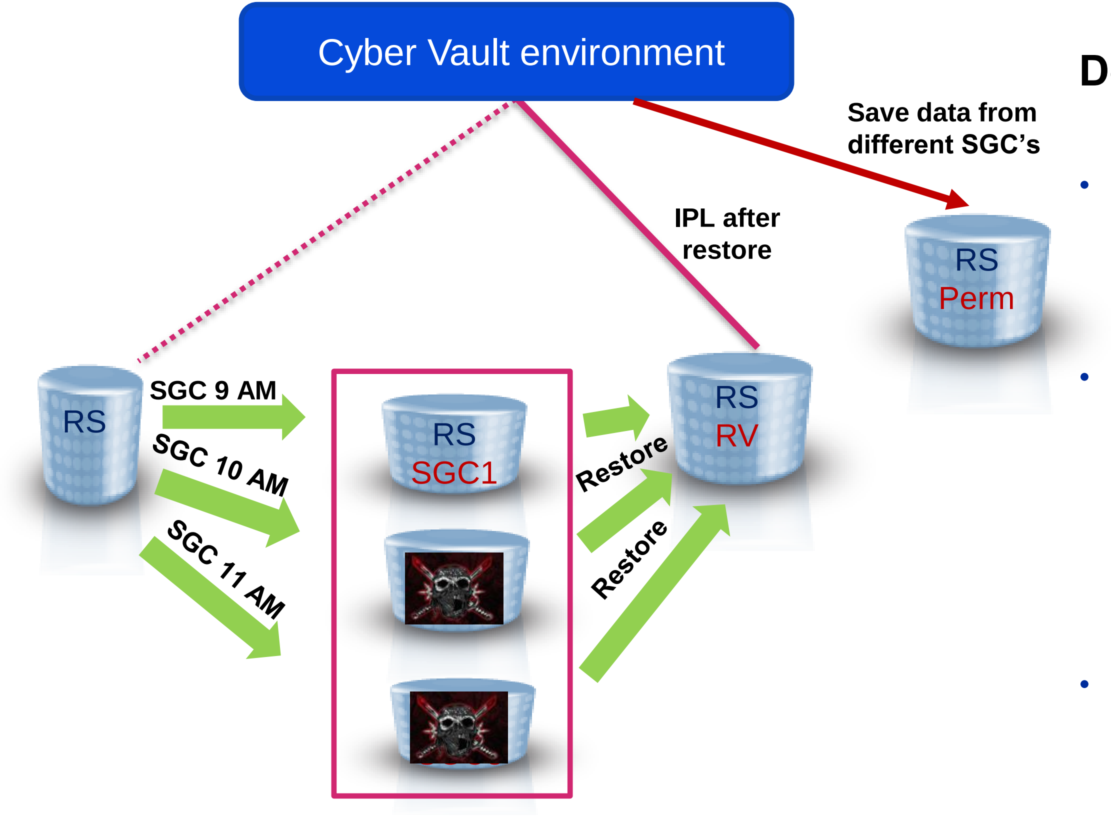

Determine start of data corruption …

- **IPL** each SafeGuarded Copy (SGC) on the Recovery Volume (RV) at the recovery site (RS)
  - Save logs from each SGC on the Cyber Vault permanent volume
- **Understand** the problem
  - Run specific data structure and data content analysis on each stored SafeGuarded Copies until a "clean" copy is found
  - Use tools to analyze databases and logs from corrupted SGC's to fully understand the scope of the problem and when it first occurred
- **Identify** steps forward
  - Create a strategy for recovery dependent on availability of database image copy files and extent of corruption

IMS forensic analysis

Collect data in production environments

- **IMS**: log data created automatically, including all database updates
- **IMS Connect**: requires tools to collect data for applications connecting to IMS via TCP/IP
- IBM solution
  - ➢ **IMS Connect Extensions** (running in production) to collect data for IMS Connect transactions

Create reports

- List transactions processed during a time period, including filters
- IBM solution
  - ➢ **IMS Performance Analyzer** for flexible report creation

Perform analysis

- View details of specific transactions
- IBM solution
  - ➢ **IMS Problem Investigator** for deep dive analysis of IMS transactions

Forensic analysis for the example scenario

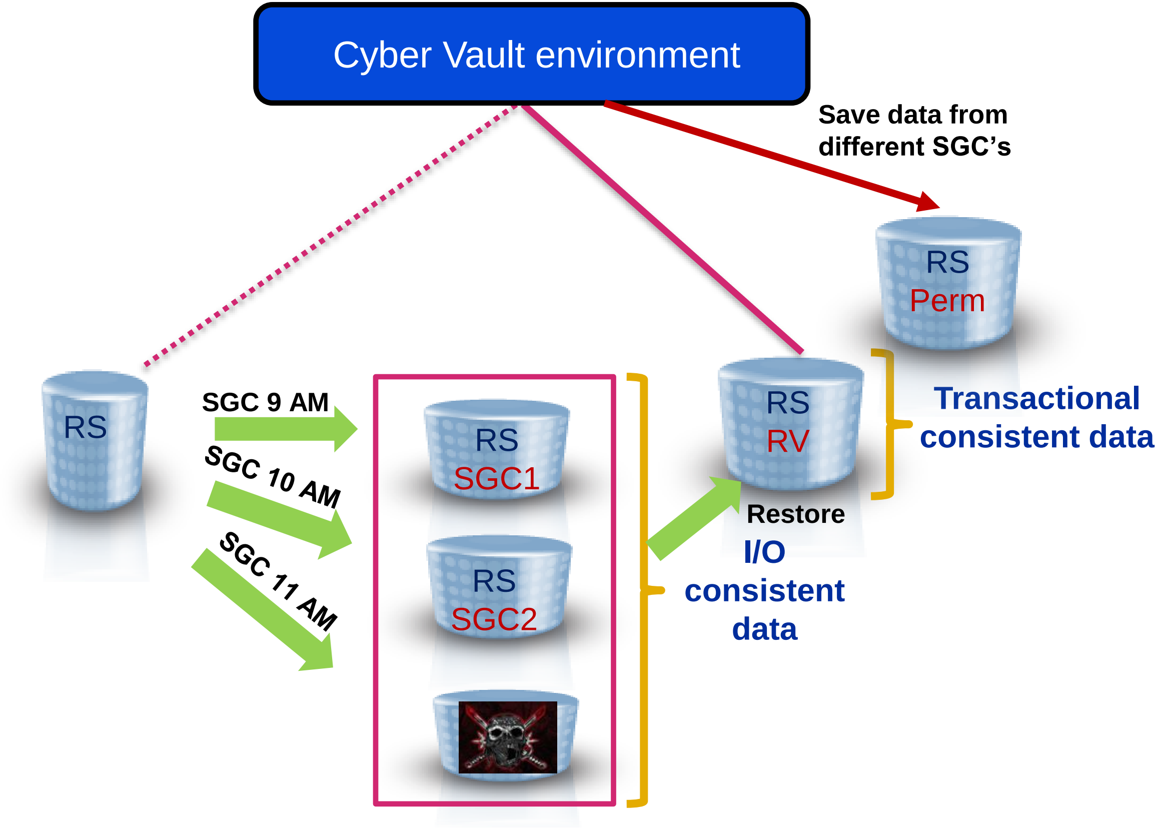

Find clean SafeGuarded Copy and retrieve logs

- IPL with 11:00 AM SafeGuarded Copy (SGC) at the recovery site (RS)
  - At least one Logical Partition (LPAR) per sysplex is required
  - System Recovery Boost Upgrade record used for one IPL per day
- Check sysplex infrastructure
- Save logs to permanent volume
- Restart IMS (all data sharing group members)
- Validate database structure, IMS HP Pointer Checker
- **11:00 AM SGC is corrupted**
- Repeat with previous SGC until clean copy is found

Result of IPL and Subsystem restart on Recovery Volumes (RV)

- All "in flight" transactions are written from log to database or backed out
- RS / RV contains "**Transactional consistent data**" at 10:00 AM
- RS / RV can be used as a base for surgical recovery to production

IMS log analysis

- Use IMS Performance Analyzer and IMS Problem Investigator for logs on permanent volume to identify the exact time when the corruption occurs
- Determine that malicious activity occurred at 10:50 AM

Surgical recovery - scenarios

Surgical Recovery is rather complex and the execution is dependent mainly on which data is available where for restore and recovery. When surgical recovery is required, the first step is to identify the actual scenario

1. Backups are available in production

- Valid image copies of database exist in production environment

2. Backups are available in the Cyber Vault only

- Valid image copies of database do not exist in production environment
- Valid image copies exist on DASD in Cyber Vault environment

3. No backups are available in production nor the Cyber Vault environment

- Valid image copies do not exist in production environment
- Valid image copies do not exist on DASD in Cyber Vault environment

Detailed scenario description (example)

Scenario - physical architecture:

- Assume a 3-site solution: Active sysplex across 2 sites in city A with a Global Mirror to city B far away
- Data center in city B serves as Disaster Recovery (D/R) site
- SafeGuarded Copy (SGC) and Cyber Vault are implemented in city B, copies are taken every hour (8:00 AM, 9:00 AM, ...)
- Validation is done on a consistent frequency in Cyber Vault environment
- The client is running 25 applications called A1 - A25

Problem description:

- 11:25 AM: logical corruption is detected when performing data validation on the 11:00 AM SafeGuarded Copy
- Forensic analysis shows the corruption began at 10:50 AM and two applications (A3 and A4) are impacted, 23 applications are free of errors
- The two impacted applications were stopped in production immediately

Remark:

- Of course the scenarios are a bit different depending on the database (IMS / Db2 / other)
- This presentation focuses on IMS

IBM Z Cyber Vault (3-site solution, virtual isolation)

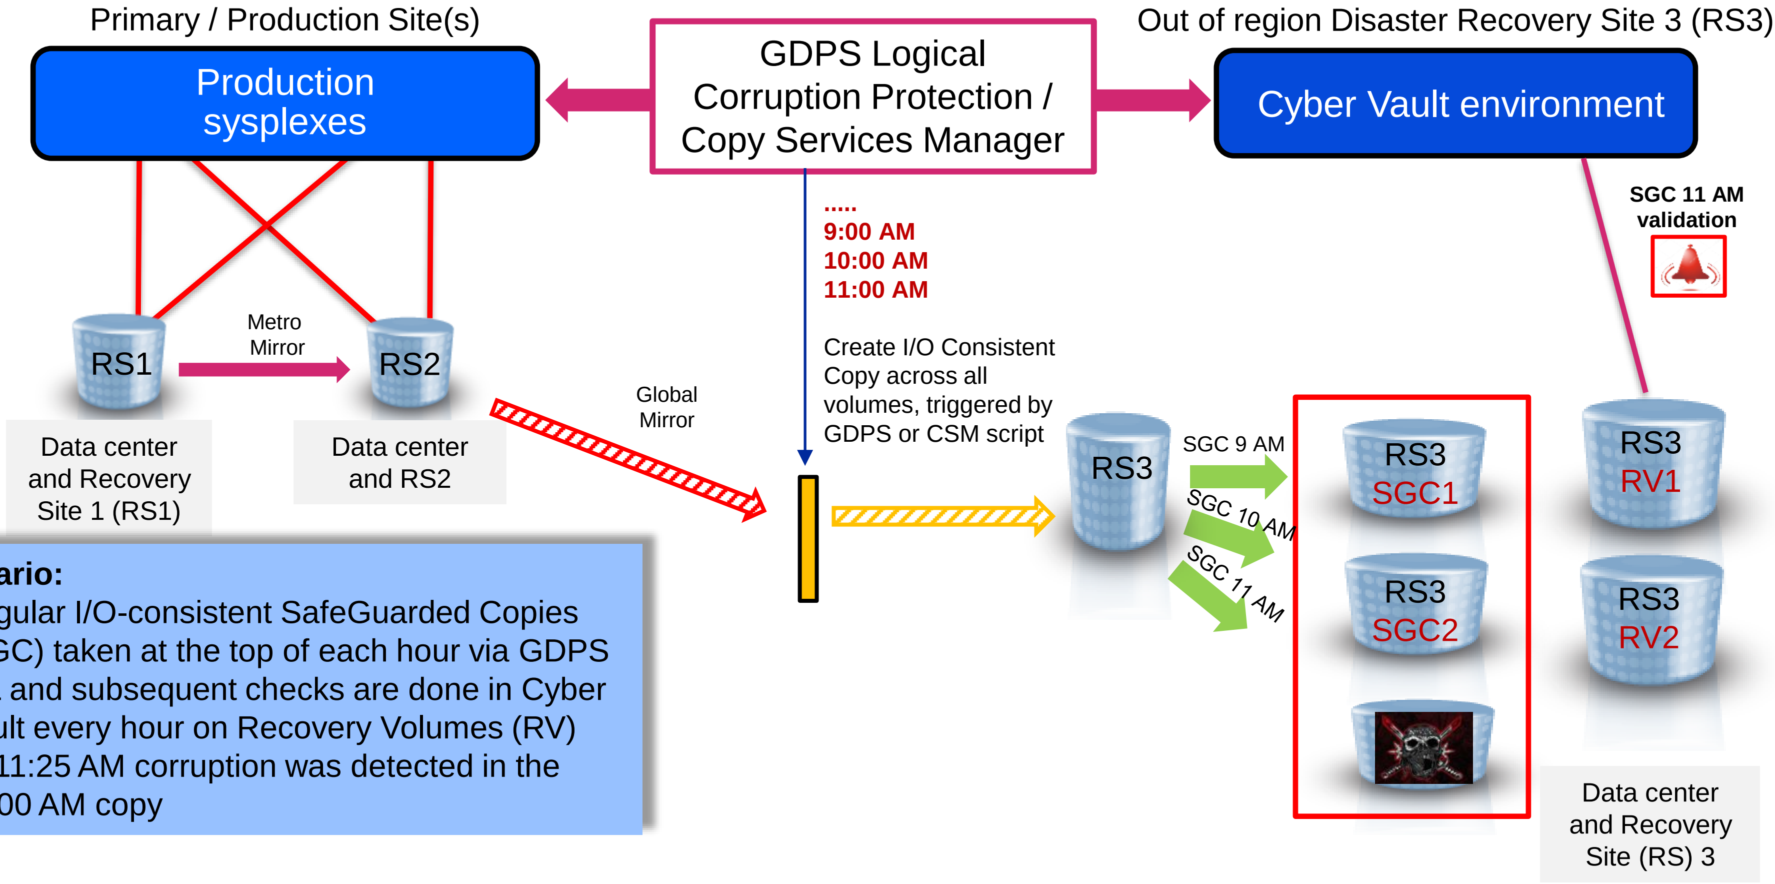
Primary / Production Site(s)
Production sysplexes
GDPS Logical Corruption Protection / Copy Services Manager
Out of region Disaster Recovery Site 3 (RS3)
Cyber Vault environment
SGC 11 AM validation
.....
9:00 AM
10:00 AM
11:00 AM
Create I/O Consistent Copy across all volumes, triggered by GDPS or CSM script
RS1
Metro Mirror
RS2
Global Mirror
Data center and Recovery Site 1 (RS1)
Data center and RS2
RS3
SGC 9 AM
SGC 10 AM
SGC 11 AM
RS3 SGC1
RS3 SGC2
RS3 RV1
RS3 RV2
Data center and Recovery Site (RS) 3

Scenario:

- Regular I/O-consistent SafeGuarded Copies (SGC) taken at the top of each hour via GDPS
- IPL and subsequent checks are done in Cyber Vault every hour on Recovery Volumes (RV)
- At 11:25 AM corruption was detected in the 11:00 AM copy

Choosing a recovery scenario

Recovery scenario 1

Image copies and logs for applications A3 and A4 are available at the production site

Image copies were taken after batch end at 5:10 AM
(before the 10:50 AM corruption)

Recovery scenario 2

Image copies and logs in production have been made unusable by malicious activity

Image copies from 5:10 AM are on disk in every SafeGuarded Copy (SGC) since 6:00 AM

Logs are available on disk in every SafeGuarded Copy

Recovery scenario 3

Image copies created directly on tape and corrupted in production environment

No image copies exist in the Cyber Vault environment

A good copy of the database is found in the 10:00 AM SafeGuarded Copy (SGC) and the database logs are found in the 11:00 AM SGC

Surgical recovery - scenario 1

Image copies and logs available in production

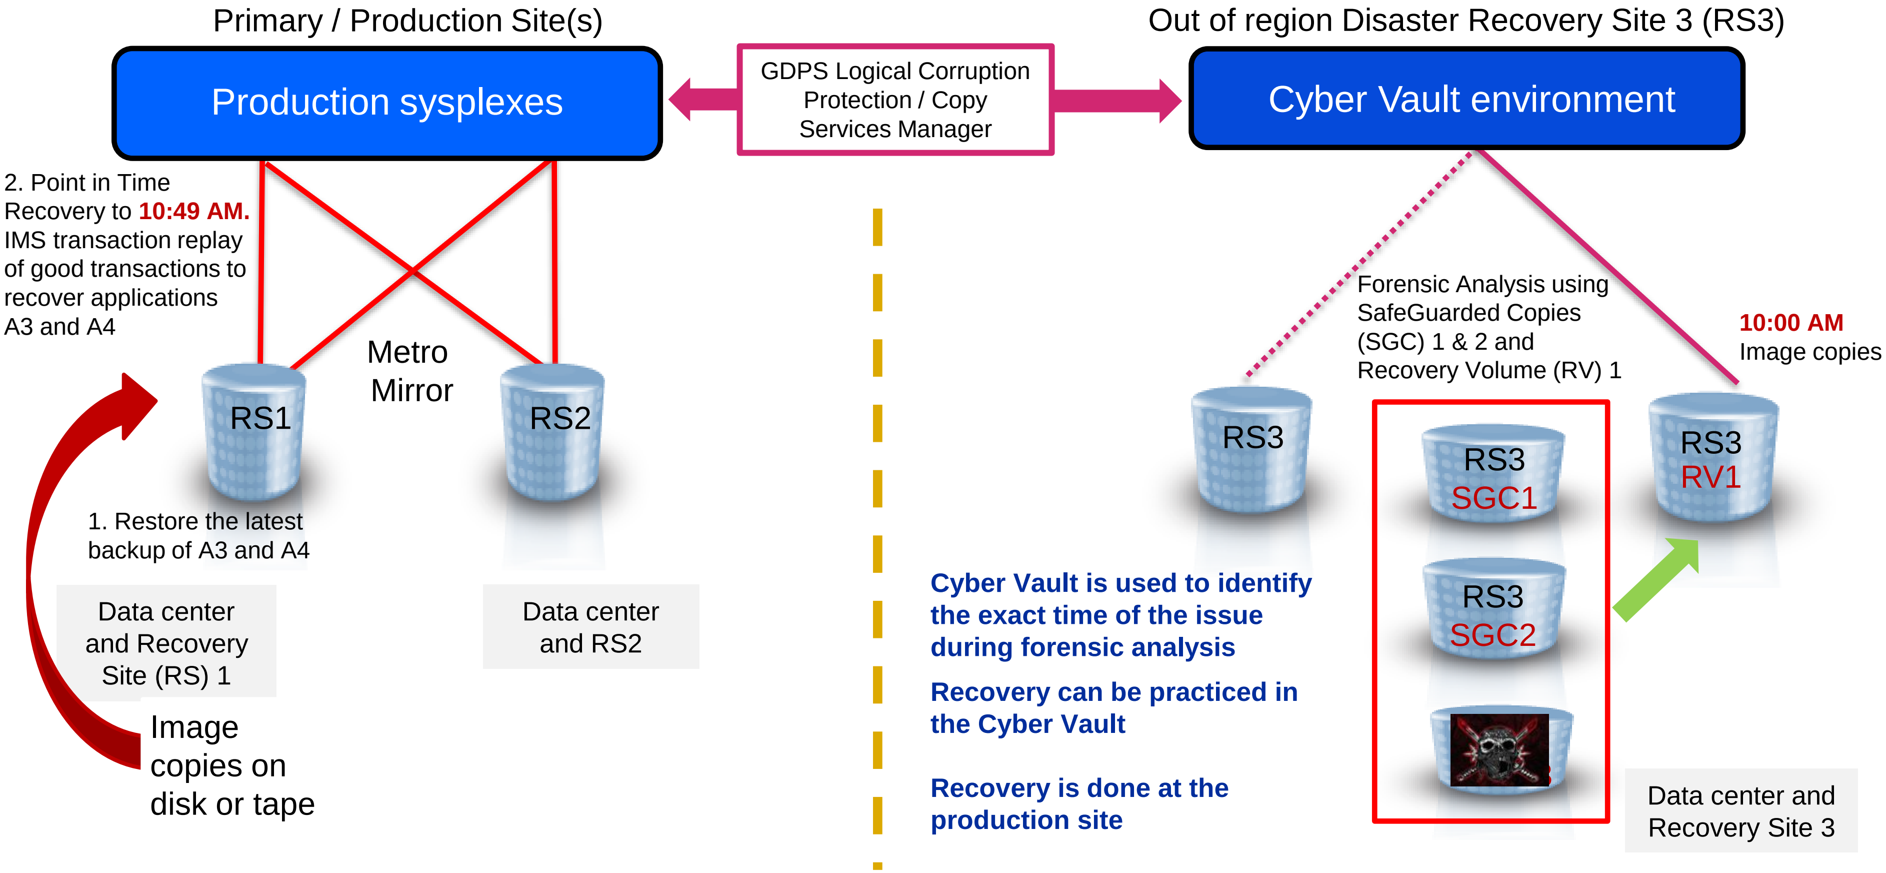

Primary / Production Site(s)

Production sysplexes

GDPS Logical Corruption Protection / Copy Services Manager

Out of region Disaster Recovery Site 3 (RS3)

Cyber Vault environment

2. Point in Time Recovery to **10:49 AM.** IMS transaction replay of good transactions to recover applications A3 and A4

Metro Mirror

RS1

RS2

1. Restore the latest backup of A3 and A4

Data center and Recovery Site (RS) 1

Data center and RS2

Image copies on disk or tape

Forensic Analysis using SafeGuarded Copies (SGC) 1 & 2 and Recovery Volume (RV) 1

10:00 AM
Image copies

RS3

RS3 SGC1

RS3 SGC2

RS3 RV1

Cyber Vault is used to identify the exact time of the issue during forensic analysis

Recovery can be practiced in the Cyber Vault

Recovery is done at the production site

Data center and Recovery Site 3

Detailed description - **scenario 1**

Assumption:

- Uncorrupted image copies available in production
- Log files are usable (not corrupted)
- Tape and disk are accessible

Recovery approach:

- Identify the malicious transaction in the Cyber Vault environment
- Perform Point in Time Recovery of the database from production site image copies to a point before the malicious transaction (10:49 AM)
  - Includes log forward apply
  - Business as Usual
- Use IMS Queue Control Facility to replay “good” transactions after 10:50 AM malicious transaction

Consequence:

- At a minimum, database recovered to consistent point before malicious activity at 10:50 AM
- Additional good transactions starting from the time of the malicious activity (10:50 AM in our example) are recovered if you are able to identify and replay them without jeopardizing consistency

Surgical recovery - scenario 2

Image copies only available in Cyber Vault

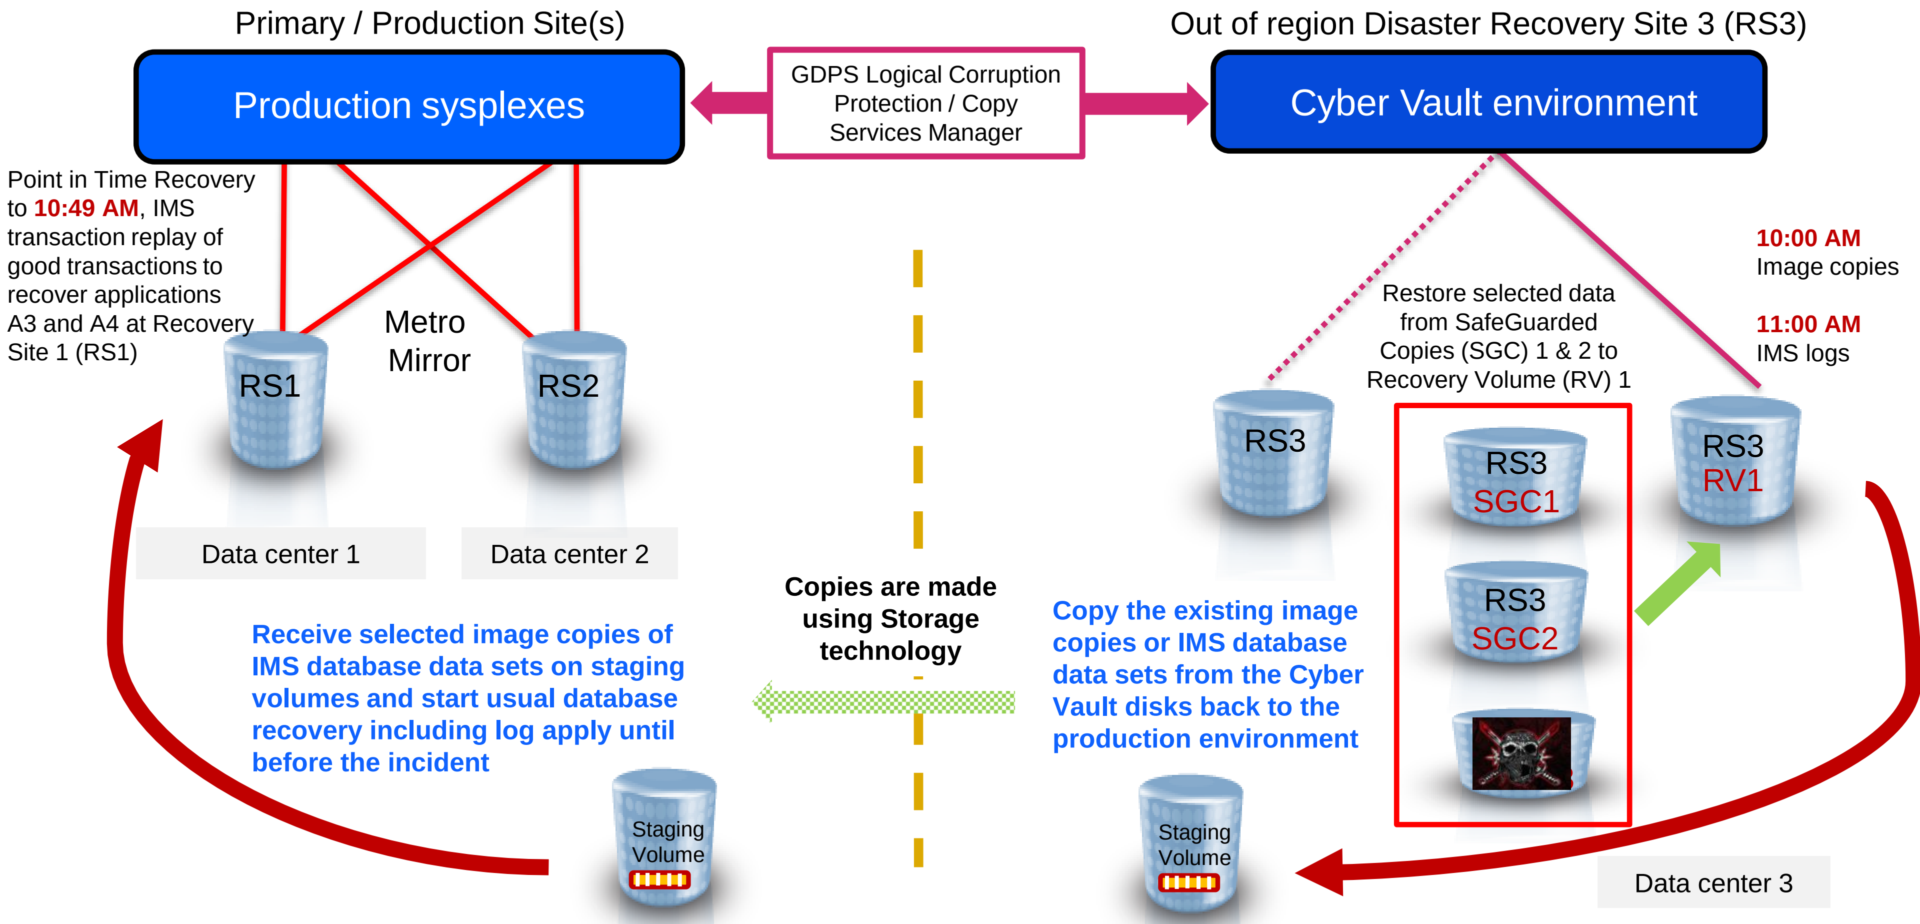

Point in Time Recovery to **10:49 AM**, IMS transaction replay of good transactions to recover applications A3 and A4 at Recovery Site 1 (RS1)

Restore selected data from SafeGuarded Copies (SGC) 1 & 2 to Recovery Volume (RV) 1

Copies are made using Storage technology

Copy the existing image copies or IMS database data sets from the Cyber Vault disks back to the production environment

Receive selected image copies of IMS database data sets on staging volumes and start usual database recovery including log apply until before the incident

Detailed description - **scenario 2**

Assumptions:

- IMS is available for most applications (including log files, etc.)
- Log files are corrupted in production
- Image copies are **not** accessible in production because either they do not exist, were corrupted by malicious activity, or the tape catalog was corrupted

Recovery approach:

- Identify the malicious transaction
- If an image copy was stored on disk before copied to tape
  - Should be found in one of the SafeGuarded Copies (SGC)
  - Image copy can be copied from Cyber Vault to production using staging volumes
- If no image copies available
  - Create image copies from the “clean” database within the Cyber Vault
  - Copy image copies from Cyber Vault to production using staging volumes
- Copy logs from 11:00 AM SGC to production using staging volumes
- Recover the database from image copies and replay good transactions – same as scenario 1
- Alternatively copy database data sets from SGC to staging volumes and replace production database data sets

Consequence:

- At a minimum, database recovered to consistent point before malicious activity at 10:50 AM
- Additional good transactions starting from the time of the malicious activity (10:50 AM in our example) are recovered if you are able to identify and replay them without jeopardizing consistency

Surgical recovery - scenario 3, phase 1

No image copies in production nor Cyber Vault

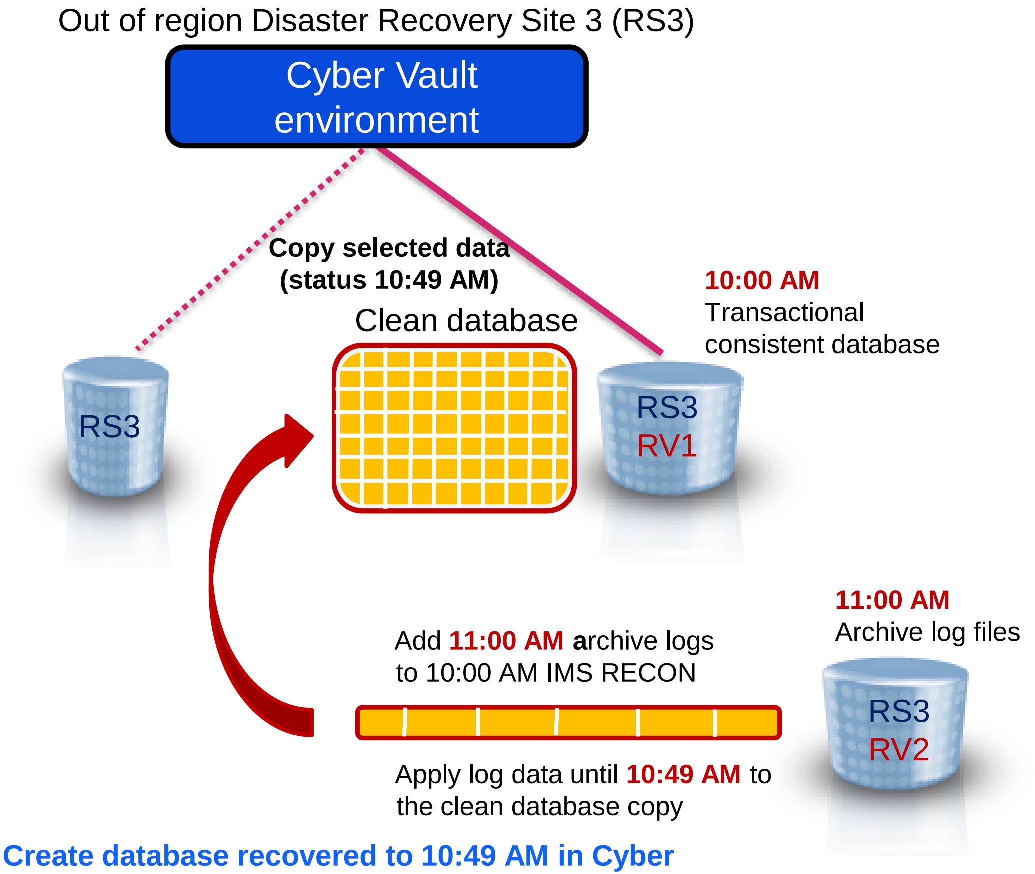

Out of region Disaster Recovery Site 3 (RS3)
Cyber Vault environment
Copy selected data (status 10:49 AM)
Clean database
10:00 AM
Transactional consistent database
RS3
RS3
RV1
Add 11:00 AM archive logs to 10:00 AM IMS RECON
11:00 AM
Archive log files
RS3
RV2
Apply log data until 10:49 AM to the clean database copy

Create database recovered to 10:49 AM in Cyber Vault environment on Recovery Volume (RV) 1

Identify base for recovery (after 11:25 AM data validation):

- Perform forensic analysis using database log files and tools to determine the time and source of the problem
- Pick a copy at Recovery Site 3 (RS3) which serves as base for recovery, 10:00 AM copy IPLed on Recovery Volume 1 (RV1) in this example
- Copy the logs and RECON from a later copy (11:00 AM in this example) over to the ”base of recovery” copy (10:00 AM)

Recovery

- Notify IMS that the 10:00 AM copy will be used for recovery with NOTIFY.UIC command
- Execute Point in Time Recovery (PITR) with IMS Database Recovery Facility specifying USEUICTIME
  - Applies database log content to the database for a status just before the corruption occurred
- Replay the “good” transactions from 10:50 AM to 11:00 AM using IMS Queue Control Facility

Checking

- Start the applications in the Cyber Vault environment (separated network - no access from outside) and check status

Surgical recovery - **scenario 3, phase 2**

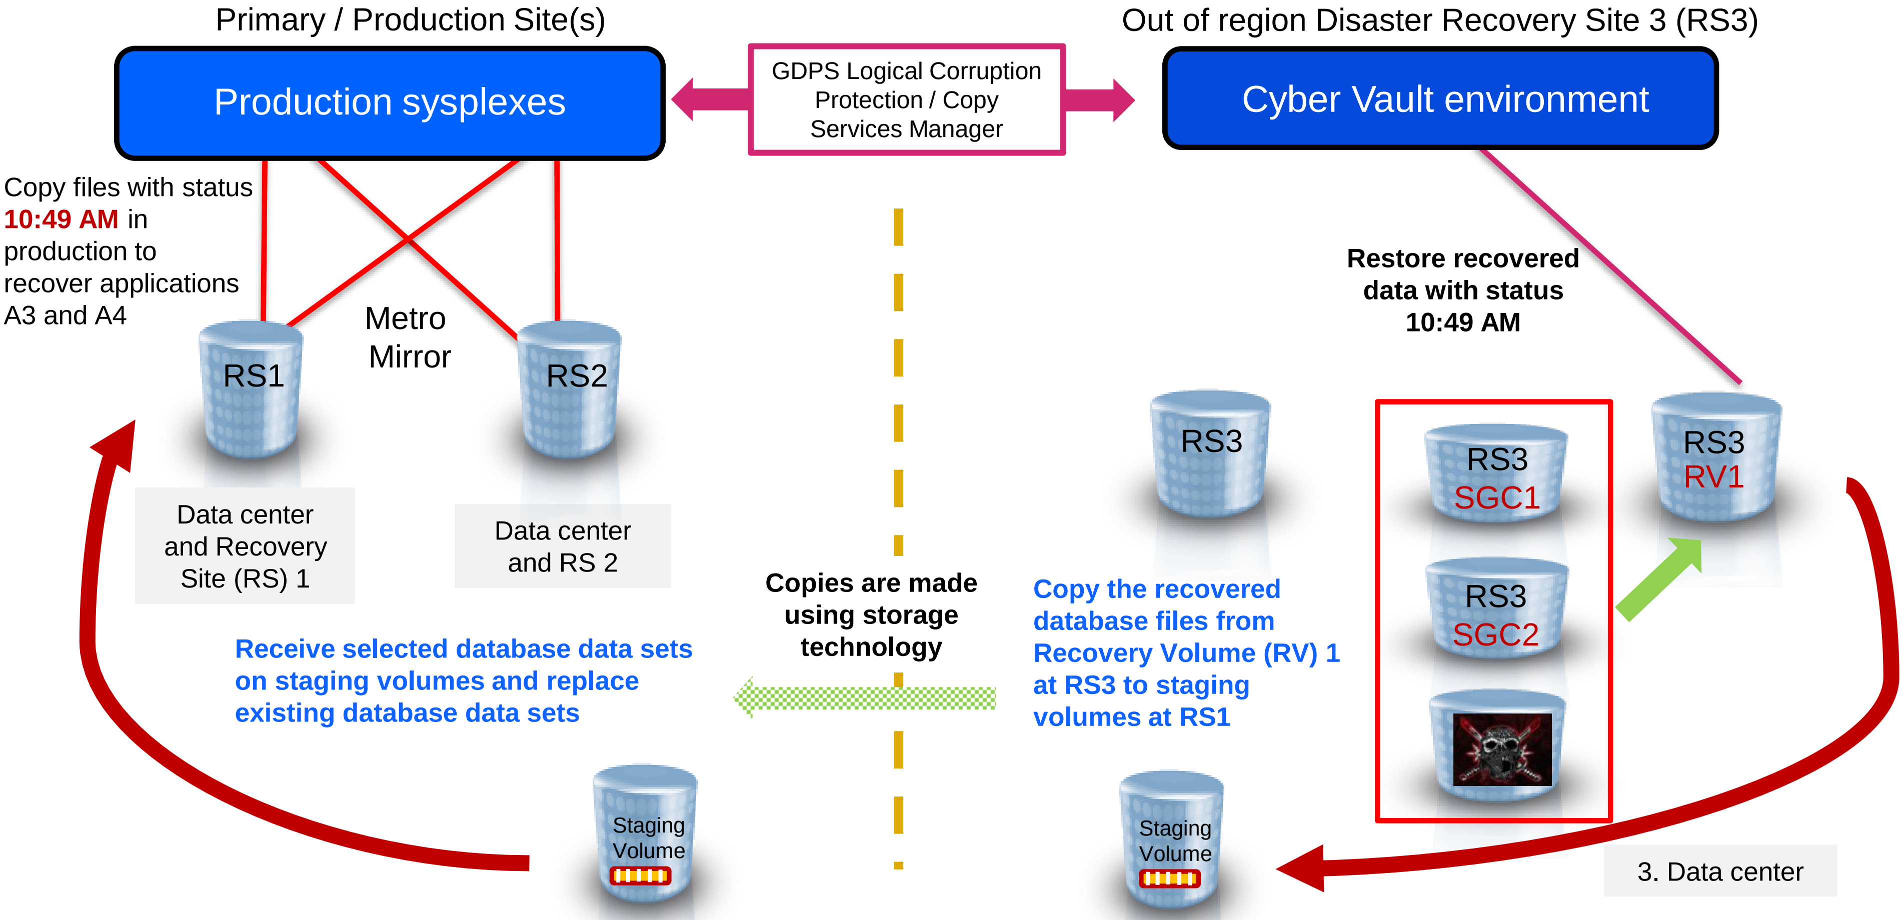

Primary / Production Site(s)

Out of region Disaster Recovery Site 3 (RS3)

Production sysplexes

GDPS Logical Corruption Protection / Copy Services Manager

Cyber Vault environment

Copy files with status **10:49 AM** in production to recover applications A3 and A4

Metro Mirror

Restore recovered data with status 10:49 AM

Copies are made using storage technology

Receive selected database data sets on staging volumes and replace existing database data sets

Copy the recovered database files from Recovery Volume (RV) 1 at RS3 to staging volumes at RS1

Detailed description – scenario 3

Assumption:

- IMS is available for most applications (including log files, etc.)
- Log files are **not** usable for the corrupted databases in production
- Image copies are not available in the Cyber Vault because they are created only on tape in production

Recovery approach:

- In Cyber Vault environment
  - Identify the malicious transaction
  - Obtain last clean copy of the database from 10:00 AM SafeGuarded Copy (SGC)
  - Obtain logs from 11:00 AM SGC
  - Execute Point in Time Recovery to 10:49 AM with IMS Database Recovery Facility
    - “USEUICLAST” option
  - Use IMS Queue Control Facility to replay “good” transactions after 10:50 AM malicious transaction
- Copy recovered database from Cyber Vault to production using staging volumes
- In production environment
  - Issue NOTIFY.RECOV in production environment to notify IMS about timestamp of recovered files
  - Perform required image copy of recovered files

Consequence:

- At a minimum, database recovered to consistent point before malicious activity at 10:50 AM
- Additional good transactions starting from the time of the malicious activity (10:50 AM in our example) are recovered if you are able to identify and replay them without jeopardizing consistency

For all recovery scenarios

The creation of regular image copy backups from all databases is vitally important.

The image copy creation frequency and Cyber Vault SafeGuarded Copy (SGC) retention periods need to be aligned with each other. For example, if image copies are taken once every week, but the SGC retention is 2 days, then there are 5 days without image copies in the Cyber Vault. This must be avoided to guarantee easy recovery.

In the example above, the retention of the SGC set must be at least one week - not shorter. For shorter SGC retentions, the image copy frequency from the production databases must be adjusted accordingly.

IMS recovery

Repair specific segments

- Valuable when small number of data points are impacted
- No need for full recovery
- IBM solution: IMS Database Repair Facility in:
  - ➢ **IMS High Performance Pointer Checker** for full function databases
  - ➢ **IMS Fast Path Solution Pack** for fast path databases

Recover databases

- Recover databases or IMS systems to a consistent point in time prior to the corruption
- Synchronize recovery of IMS and Db2 databases
- IBM solution
  - ➢ **IMS Recovery Solution Pack** for point in time recovery of databases or IMS systems

Replay valid transactions

- Replay valid transactions beyond recovery point
- Skip invalid transactions (based on forensic analysis)
- IBM solution
  - ➢ **IMS Queue Control Facility** to recover and replay specific transactions

IBM Z Cyber Vault software selection – summary for IMS

These are the products that, following IBM Best Practices, provide resiliency capabilities to your IMS database and transaction processing subsystems.

Solution	P	CV	Capability
IBM IMS High Performance Pointer Checker Pointer checking IMS full function databases	✗	✓	Data Validation
IBM IMS Fast Path Solution Pack Pointer checker function for Fast Path databases, aka DEDBs	✗	✓	
IBM IMS Recovery Solution Pack Database Recovery Facility component to validate all assets needed for recovery are available and can get to all of them	✗	✓	
IBM IMS Connect Extensions Collect and write data about IMS transactions coming in through IMS Connect	✓	✗	Forensic Analysis
IBM IMS Problem Investigator Deep dive analysis of IMS logs and IMS Connect Extensions journals	✗	✓	
IBM IMS Performance Analyzer Report on transactions that occurred during a specified period	✗	✓	
IBM IMS Recovery Solution Pack Recover specific IMS systems or databases based on the volume level backups	✓	✓	Surgical Recovery
IBM IMS High Performance Pointer Checker Repair specific segments in IMS full function databases without requiring full recovery	✗	✓	
IBM IMS Fast Path Solution Pack Repair specific segments in IMS Fast Path databases without requiring full recovery	✗	✓	
IBM IMS Queue Control Facility Recover and/or replay specific transactions	✓	✓	

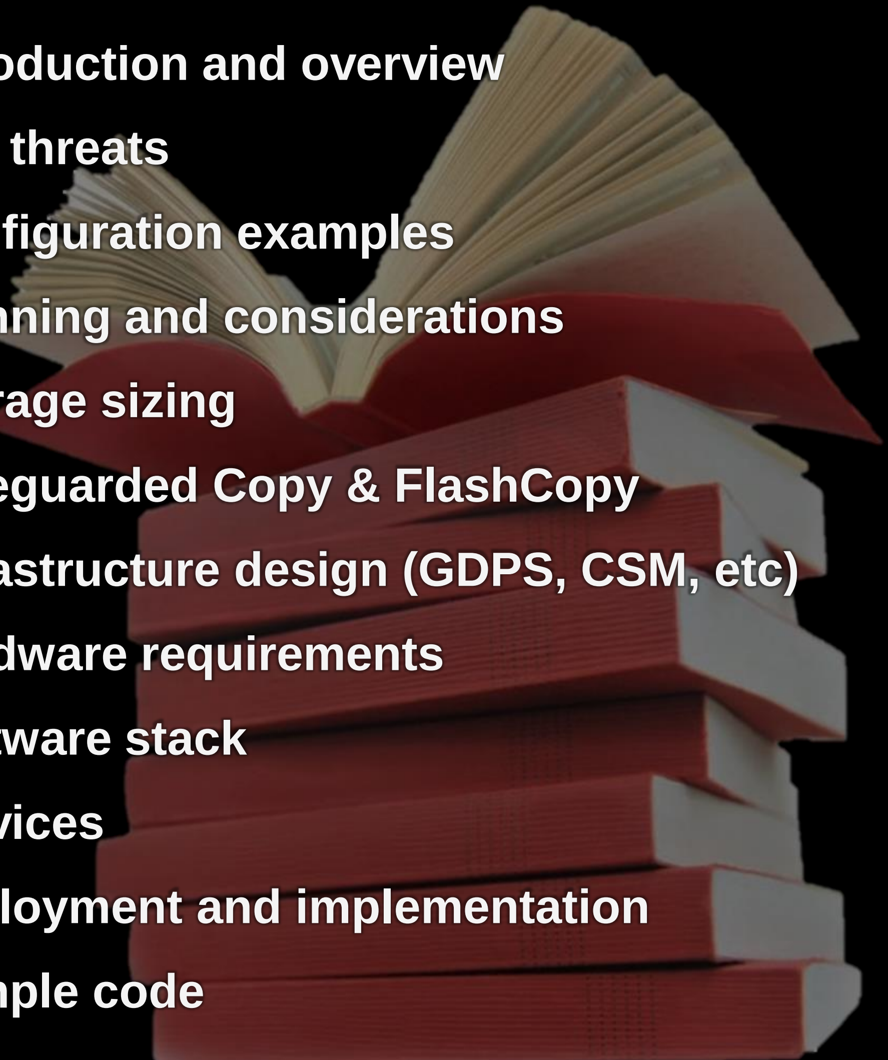

- ✓ Introduction and overview
- ✓ Key threats
- ✓ Configuration examples
- ✓ Planning and considerations
- ✓ Storage sizing
- ✓ Safeguarded Copy & FlashCopy
- ✓ Infrastructure design (GDPS, CSM, etc)
- ✓ Hardware requirements
- ✓ Software stack
- ✓ Services
- ✓ Deployment and implementation
- ✓ Sample code

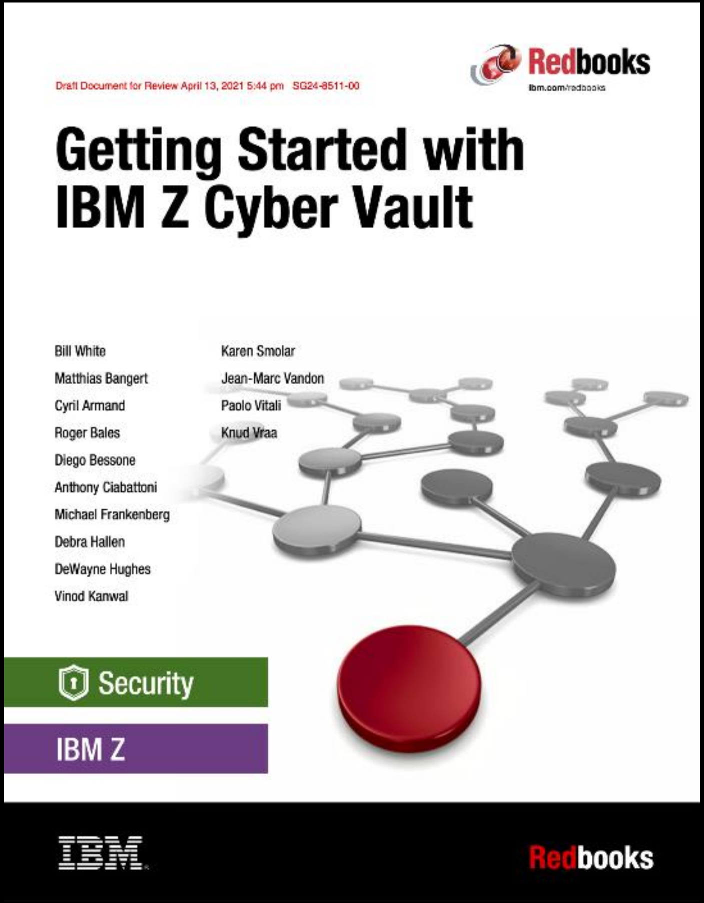

Draft Document for Review April 13, 2021 5:44 pm SG24-8511-00

Redbooks
ibm.com/redbooks

Getting Started with IBM Z Cyber Vault

Bill White
Matthias Bangert
Cyril Armand
Roger Bales
Diego Bessone
Anthony Ciabattoni
Michael Frankenberg
Debra Hallen
DeWayne Hughes
Vinod Kanwal
Karen Smolar
Jean-Marc Vandon
Paolo Vitali
Knud Vraa

Security

IBM Z

IBM

Redbooks

Link to Redbook

Additional IMS and IMS Tools Links

IMS Tools website
www.ibm.com/it-infrastructure/z/ims/tools

IBM Z Software Newsletter, Operations and Management
http://ibm.biz/zITSMNewsletterSubscribe

IMS Tools Product Documentation
www.ibm.com/support/docview.wss?uid=swg27020942

IMS listserv
http://imslistserv.bmc.com

IMS Tools new functions
www.ibm.com/support/docview.wss?uid=swg22015506

IMS Tools support for IMS V15
https://www.ibm.com/support/pages/node/6572967

IMS Tools support for Managed ACBs
www.ibm.com/support/docview.wss?uid=ibm10731745

IMS Tools support for Data Set Encryption
www.ibm.com/support/pages/ibm-ims-tools-and-data-set-encryption-support

IMS Fundamentals videos:
https://mediacenter.ibm.com/playlist/dedicated/122579632/1_b56rpdpt/1_jy8lv5f5

IMS Tools Videos on IBM MediaCenter
ibm.biz/ims-tools-mediacenter

IMS new functions
www.ibm.com/docs/en/ims/15.3.0?topic=enhancements-ims-enhancement-ptfs

THANK YOU!

QUESTIONS?

Trademarks

The information contained in this publication is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this publication, it is provided AS IS without warranty of any kind, express or implied. In addition, this information is based on IBM’s current product plans and strategy, which are subject to change by IBM without notice. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this publication or any other materials. Nothing contained in this publication is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.

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