z/OS Expertenforum in Vitznau:

IBM Z CyberVault

Walter Kläy
walter.klaey@ch.ibm.com

19. October 2021

IBM

The question is not IF you will be attacked but WHEN

$53 Billion

Predicted economic losses of the next global cyber attack

$3.86 Million

Average total cost of a data breach*

$8 Billion

Estimated global cost of WannaCry attack*

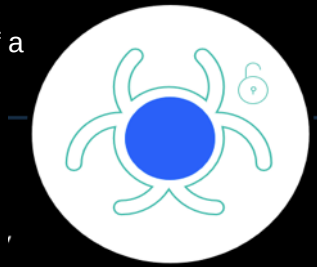

$??? Million

GDPR fine for one data breach*

280 Days

Average amount of time hackers spend inside IT environments before discovery*

- Cost of a Data Breach Report 2020, Ponemon Institute
- Riensirce news May 23 2017

HONDA

Honda Hackers May Have Used Tools Favored by Countries

The New York Times

REUTERS®

'Payment sent' - travel giant CWT pays $4.5 million ransom to cyber criminals

GARMIN

The Garmin Hack Was a Warning

As ransomware groups turn their attention to bigger game, expect more high-profile targets to fall.

WIRED

The Register®

UBS logic bomber jailed for eight years

Real-life BOFH ordered to pay $3.1m restitution

Switzerland: Comparis, various Swiss Communities

A cyber attack can quickly put you out of business

In an age of ongoing digital transformation, cybercrime has quickly become today's fastest growing form of criminal activity.

Sixty percent of business who are victimized by a cyber attack go out of business within six months.[1]

Increase in data breach observed due to remote work during COVID-19.

Mainframes becoming a target

A large amount
of all active code
runs on the mainframe

A large amount
of enterprise data is
housed on the mainframe

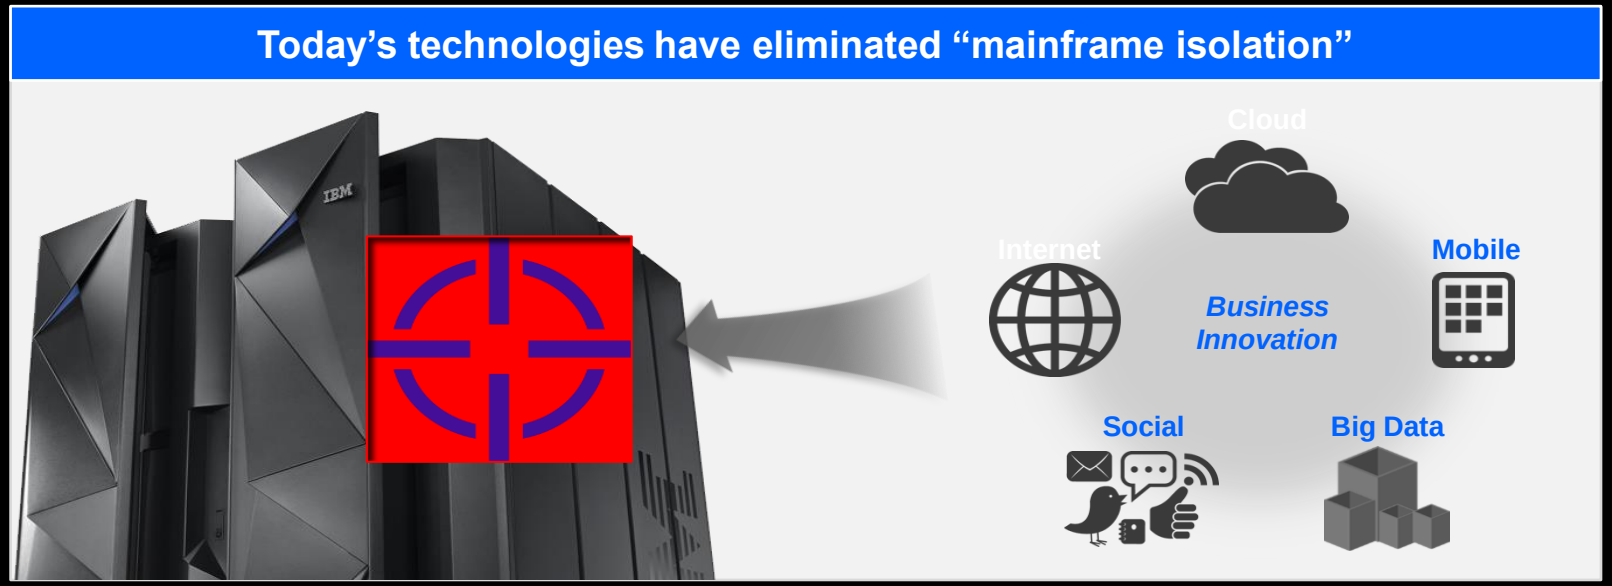

Attack vectors

Definition:

- **Malware** is any software intentionally designed to cause damage to a computer, server, client, or computer network. A wide variety of malware types exist, including computer viruses, worms, Trojan horses, ransomware, spyware, adware, rogue software, wiper and scareware.
- **Ransomware** as a type of malware from cryptovirology that threatens to publish the victim's data or perpetually block access to it unless a ransom is paid.

Attack vectors:

- **Malware**
- **External intrusion**

but also:

- **Unintentionally** corrupted files / databases by SW bugs or handling errors (human errors)
- **Internal attacks**

Cyber Resiliency - High Level View

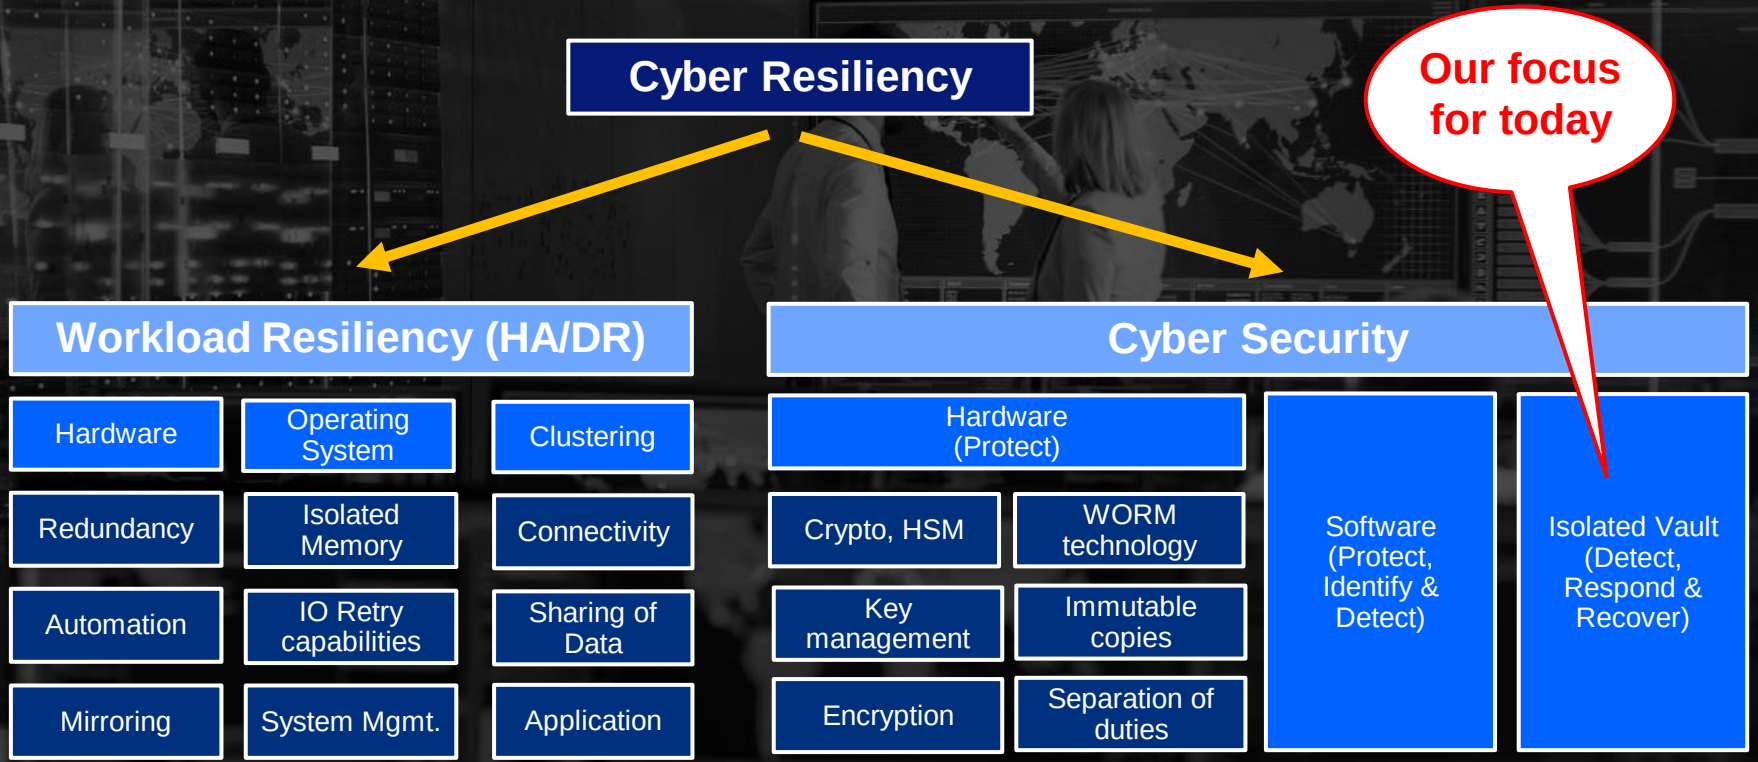

Cyber Resiliency lifecycle (based on "*NIST Cybersecurity Framework*")

IDENTIFY

Risks and vulnerabilities

PROTECT

Your critical applications and data

DETECT

Data corruption and configuration anomalies

RESPOND

To changes in configuration and data

RECOVER

Access to critical applications and data

NIST = National Institute of Standards and Technology of the USA Government

Cyber Resiliency lifecycle (based on "*NIST Cybersecurity Framework*")

IDENTIFY

Risks and vulnerabilities

PROTECT

Your critical applications and data

DETECT

Data corruption and configuration anomalies

DISA-Stig

(Checking a companies security setup against best practises)

Proactive actions required

(Immutable disk copies, Logical Worm, separation of duties)

Software Tools

(Access patterns, write activity, ...)

All three components are important and need to run in production. However, the "DETECTION" is the most critical - and complicated part.

There are no "error messages" you can use as trigger as a result of a logical corruption.

Access pattern and write activity analysis can easily lead to "false alarms" - how many of those do you tolerate ?

This is clearly a future AI topic .....

From a z/OS point of view

- **z/OS is the most securable platform in the industry, but security does not come by default. To prevent from any incident, do you…**
  - Constantly check for vulnerabilities ?
  - Enforce the security policies ?
  - Monitor security events and incidents ?
  - Data base activity monitoring in place ?
  - … ?

  - Are you ready for an IBM z/OS Security Health Check ? IBM z/OS Security Health Check

Cyber Resiliency lifecycle (based on “*NIST Cybersecurity Framework*”)

It is maybe not possible to run all checks regularly in production - so why not using the immutable copies for advanced checking ?

“RESPOND” also includes, that you are able to identify where the corruption started - to avoid a reoccurrence.

The “RESPOND” and “RECOVER” phase must be tested and prepared in an isolated environment.

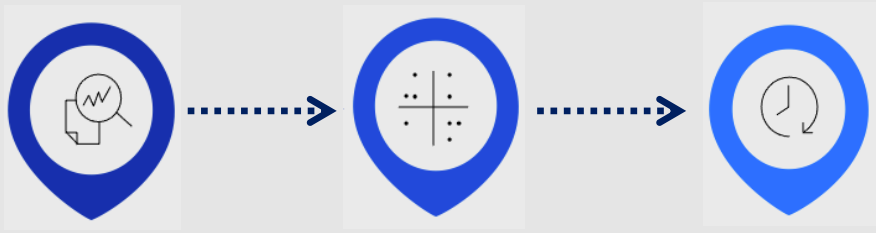

DETECT

Data corruption and configuration anomalies

RESPOND

To changes in configuration and data

RECOVER

Access to critical applications and data

Isolated Environment

(Advanced checking, analysis of data outside production, testing recovery actions)

Early detection is key to significantly reduce the impact of breaches

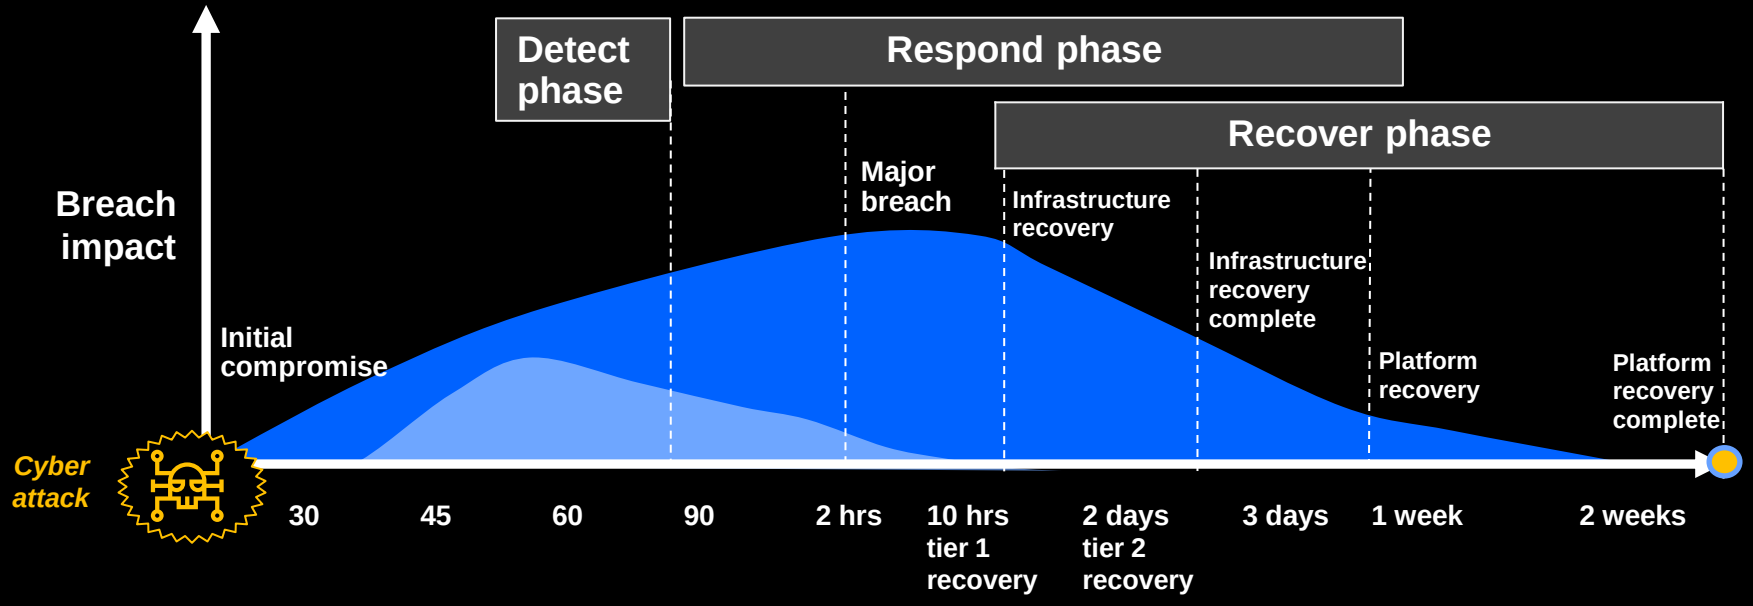

(Remark: timeline is an example, however based on real client experience)

“If we are a victim of logical corruption ....
.... how could we recover our 300 most critical services within 24 hours?”

Focus

How do we **reduce outage time** as a result of corruption or destruction of data and systems ?

Problem

Existing HA/DR protects from **physical** infrastructure outages.

Traditional backups are specific to **individual** databases and do not allow re-creation of a consistent system image.

Solution

Establish a set of **recent ‘point in time’ copies** of the entire system, with ability to **regularly inspect quality** of data in the isolated vault, and test recovery scenarios.

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
  - Restore production data
  - Test application and system changes
  - Proactive vs. reactive checks

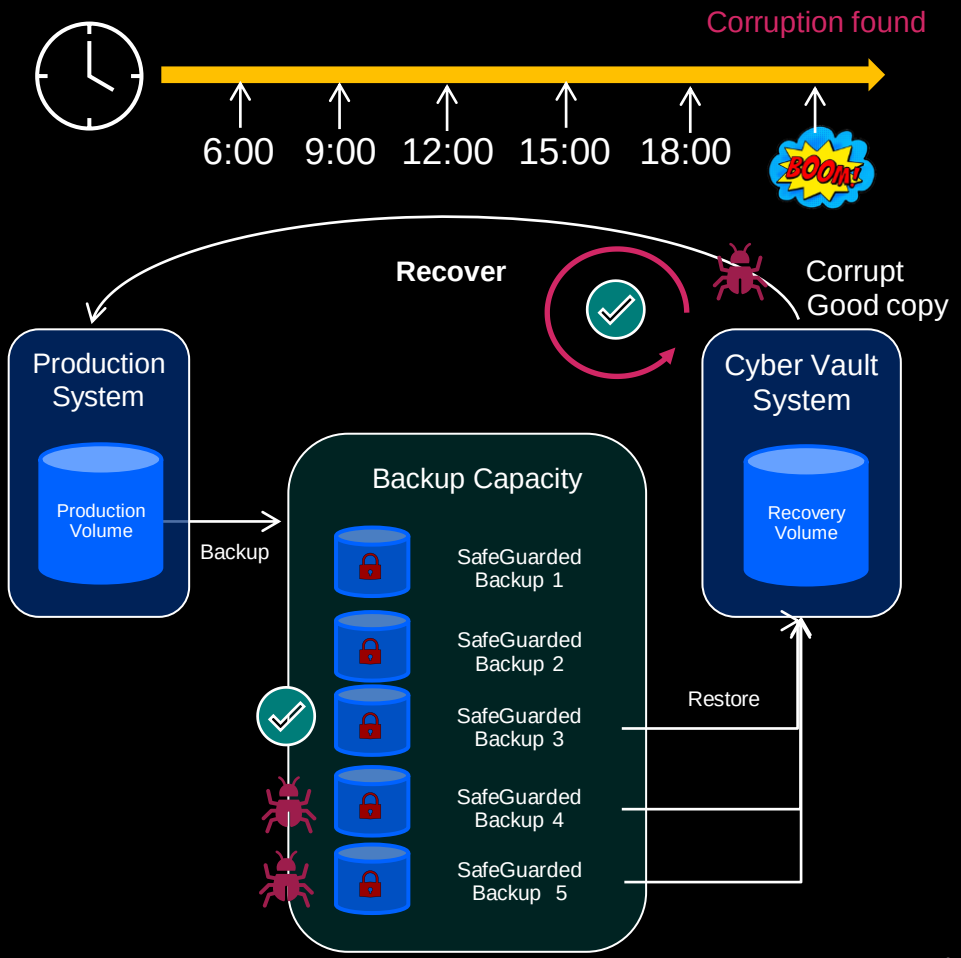

Air gap: Virtual or physical isolation of protection copies

Virtual isolation

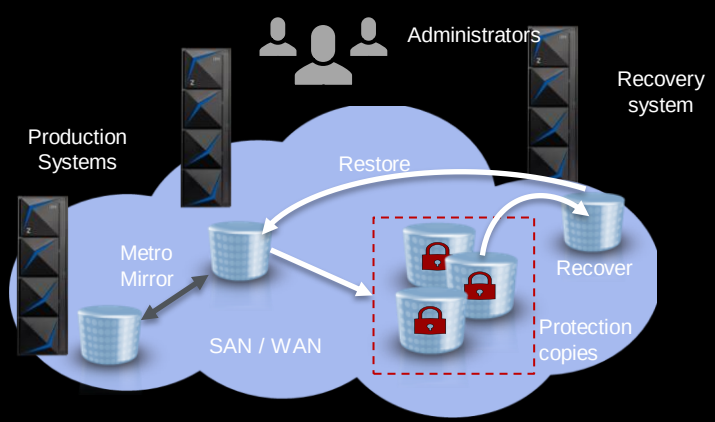

- The protection copies are created in one or more storage systems in the existing high availability and disaster recovery topology
- The storage systems are typically in the same SAN or IP network as the production environment

Physical isolation

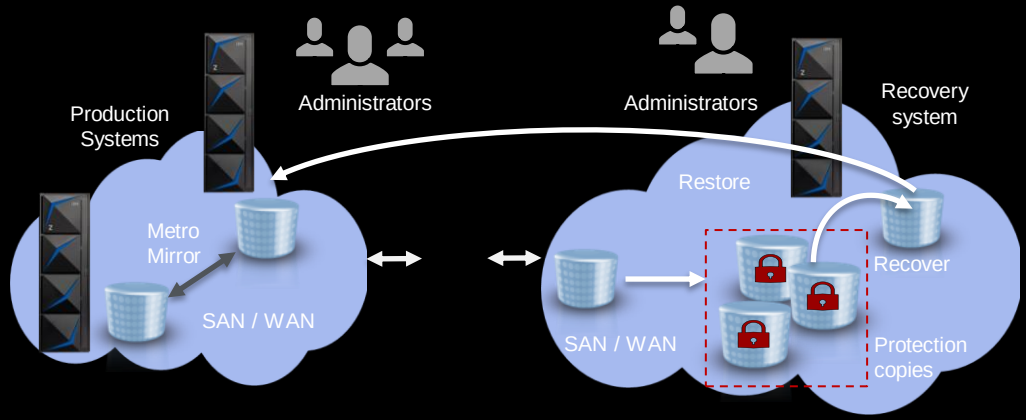

- Additional storage systems are used for the protection copies
- The storage systems are typically not on the same SAN or IP network as the production environment
- The storage systems have restricted access and even different administrators to provide separation of duties

Cyber Vault deployment example - large UK bank
Physical Airgap (and Isolation) with Global Mirror

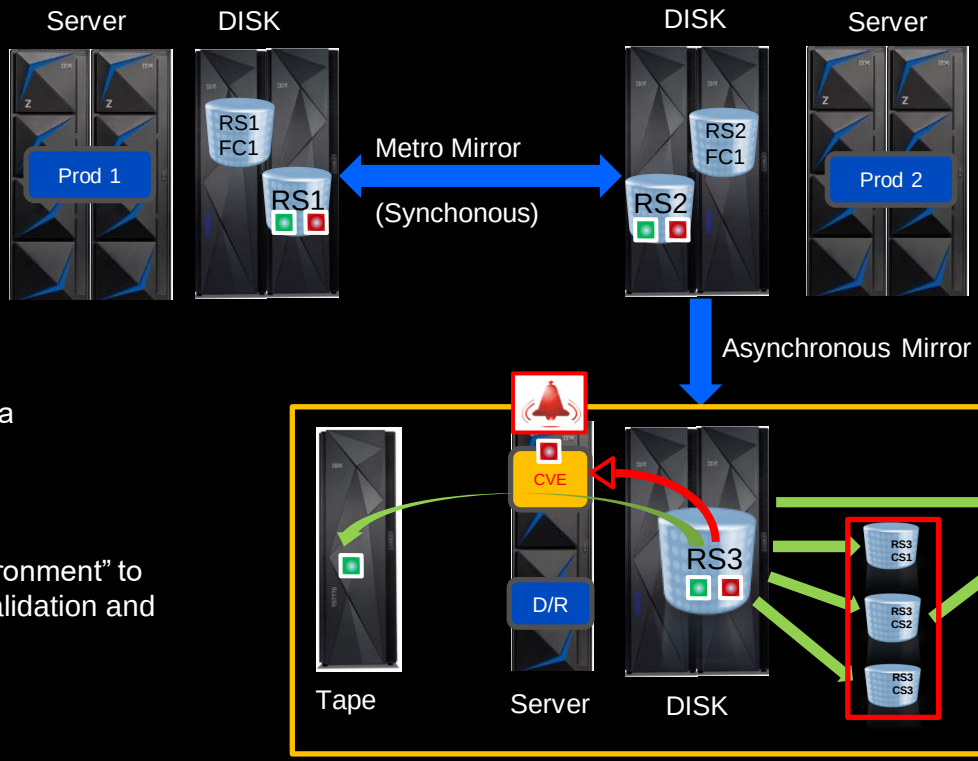

= “Good” copy of data

= “Bad” copy of data

CVE = “Cyber Vault Environment” to enable ongoing validation and recovery actions.

Alert sent in case of anomaly detection. Otherwise copy (optionally) on tape

Validation - Forensic Analysis - Surgical Recovery in a Cyber Vault

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, but also application knowledge is required

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

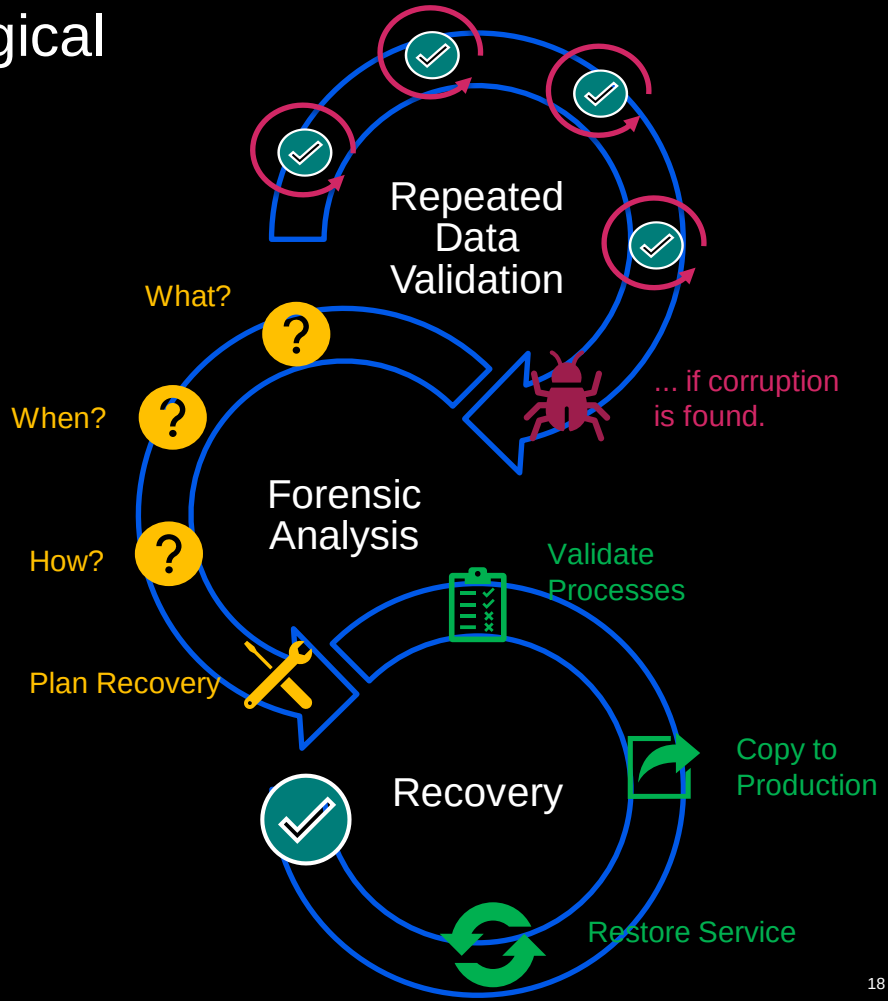

Data validation details
Identify problems

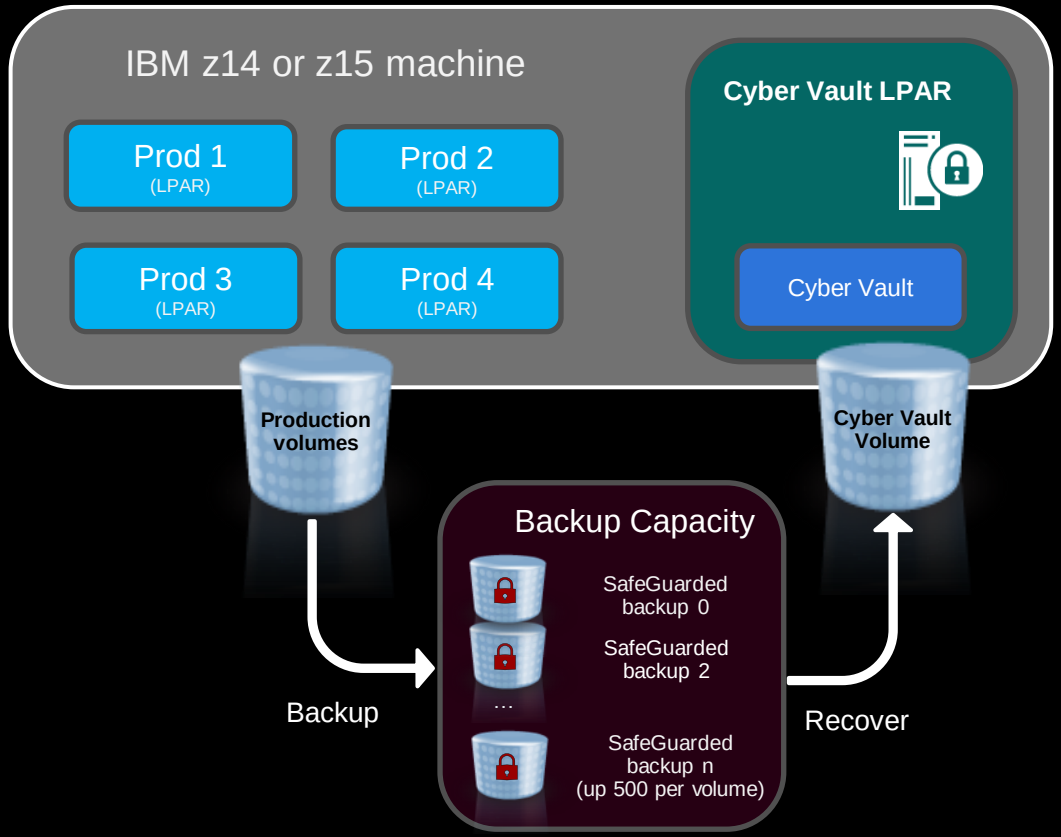

At least once a day…

Phase 1: IPL with production image

At least one LPAR per Sysplex is required

- System Recovery Boost upgrade record used for one IPL per day
- Check Sysplex infrastructure

Phase 2: Data structure validation

- Db2 restart (all data sharing group members), Utilities, Log analysis
- IMS restart, Utilities
- Catalog tools (Tivoli, IDCAMS, ISV products)
- VSAM Indexcheck, Datacheck
- SMShsm, SMSrmm tools
- RACF (IRRUT200), zSecure-Audit
- ISV software (CA1, CA7, ...)

Phase 3: Data content validation

- Customer application program(s)

If no issue found: Create tape copy

Automation is key…

IBM Z Cyber Vault systems image validation

Step 1
z/OS System

Step 2
z/OS Subsystems
& Data Structure

Step 3
Application Data

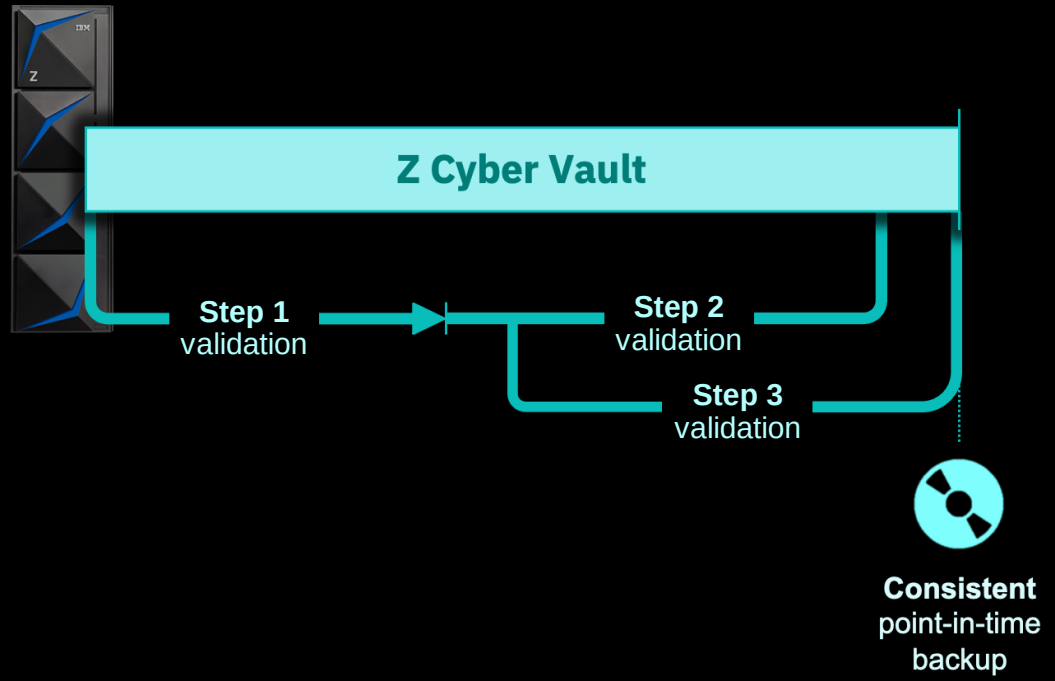

IBM Z Cyber Vault software selection

This is an initial list of software required to provide Cyber Vault capabilities.
ISVs can also be part of this list

z/OS	
Catalog / VTOC / VVDS	
Tivoli Advanced Catalog Management for z/OS	Enhanced checks on catalog integrity, runs faster than standard IDCAMS utilities, has extensive repair capabilities.
DFSMS	
Advanced Audit for DFSMShsm	Faster than standard HSM audit, has a connection to RMM to cross check tapes.
Security	
zSecure Audit	Can check unauthorized update of static libraries (load libraries, JCL, ...).
Guardium S-TAP (Db2 & IMS)	Potential identification of malicious database activities to help identify starting point of corruption.
Datasets	
IZBR	Fast identification and recovery of datasets being open during backup.
Db2 Subsystem	
Db2 Log Analysis Tool	Creation of data value change reports and fast recovery of changes recorded in the Db2 log.
Db2 Recovery Expert	Data structure check for all Db2 databases.
IMS Subsystem	
IMS High Performance Pointer Checker	Data structure check for all IMS databases.
IMS Recovery Expert	Speeds up backup and recovery processes.

IBM Z Cyber Vault software

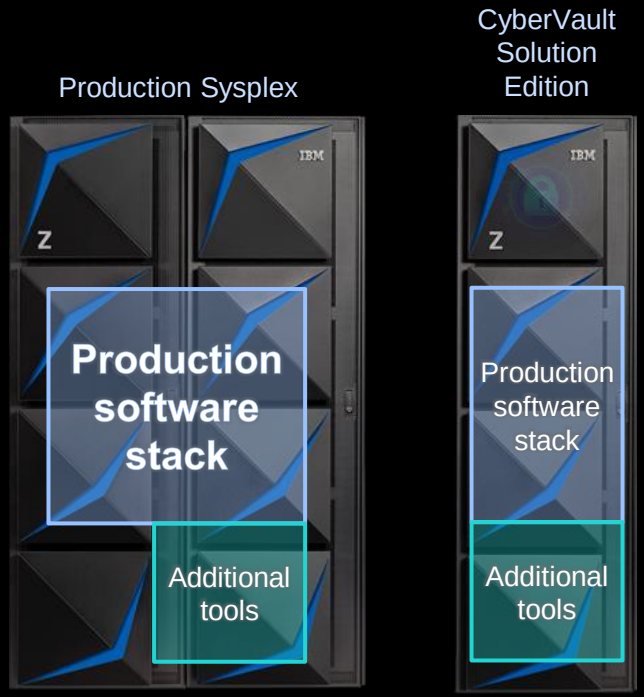

IBM Z Cyber Vault provides air gapped data corruption protection and tools to detect issues and speed recovery. This isolated environment requires hardware and software that will be configured and priced as a “Solution Edition” in order to provide the best value.

- The IBM Z Cyber Vault will be started from a copy of the existing production environment that has already been replicated. This means that all the current software in production needs to be available (and licensed) in the IBM Z Cyber Vault. This includes all IBM and non-IBM software.

- For enhanced diagnostics and recovery, additional IBM Software could be required in both the production Sysplex and the Cyber Vault.
- IBM will work with you to identify your specific needs and requirements in a discovery and architecture workshop to define a final software list.
- If the required software is not currently licensed, it can be provided through the IBM Z Cyber Vault Solution Edition contract as a limited use license.

IBM Z Cyber Vault
Challenge

Prepare to Recover

- IBM Z Cyber Vault is an extension of the current DR concept
- RPO > 0 and unpredictable RTO are challenging for every business
- Analysis and recovery of such an incident is very time critical for the business
- A large variety of recovery plans as they depend on the type of attack.
- As we are used from the current DR preparation:
  Management and employees need to be aware, educated and regular training is essential

IBM Z Cyber Vault is a solution (not a product) based on IBM technology, products & services, help to minimize the impact of such an incident.

This end up typically a project (with a larger scope)

Implementation is typically step by step approach

IBM Z Cyber Vault Solution

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 Cyber Vault for IBM 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

IBM Z Cyber Vault

A secure, air-gapped solution to ensure you can recover your data to a good, known version –

No matter what!

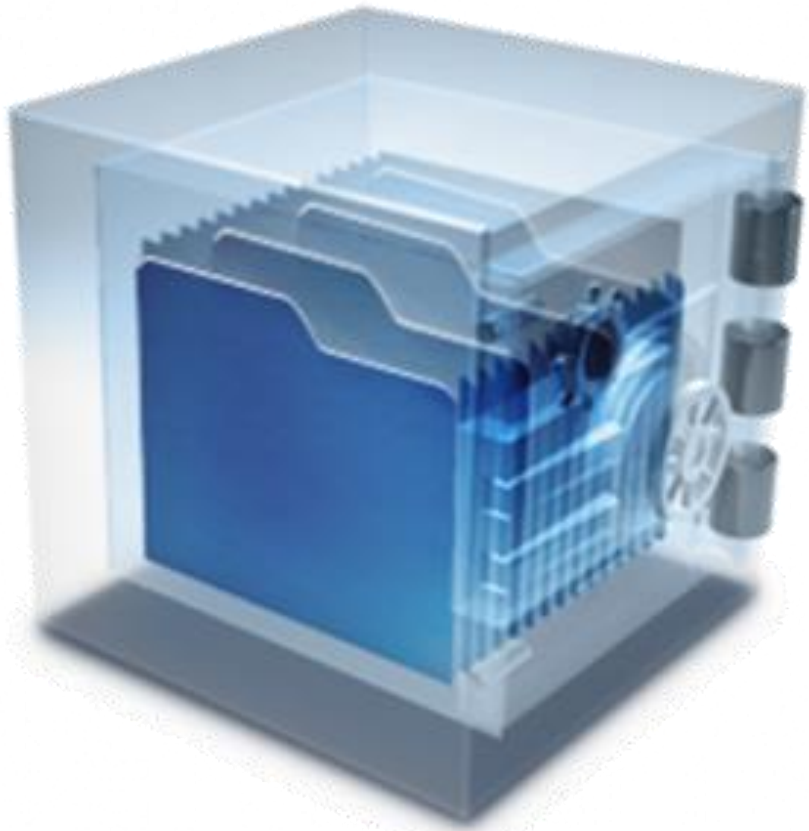


Questions?

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

*** Registered trademarks of IBM Corporation**

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

IT Infrastructure Library is a Registered Trade Mark of AXELOS Limited.

ITIL is a Registered Trade Mark of AXELOS Limited.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

OpenStack is a trademark of OpenStack LLC. The OpenStack trademark policy is available on the OpenStack website.

Red Hat®, JBoss®, OpenShift®, Fedora®, Hibernate®, Ansible®, CloudForms®, RHCA®, RHCE®, RHCSA®, Ceph®, and Gluster® are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

RStudio®, the RStudio logo and Shiny® are registered trademarks of RStudio, Inc.

UNIX is a registered trademark of The Open Group in the United States and other countries.

VMware, the VMware logo, VMware Cloud Foundation, VMware Cloud Foundation Service, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

Zowe™, the Zowe™ logo and the Open Mainframe Project™ are trademarks of The Linux Foundation.

Other product and service names might be trademarks of IBM or other companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

This information provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g, zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines" provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html ("AUT"). No other workload processing is authorized for execution on an SE. IBM offers SE at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types a