

# Getting Started with IBM Z Cyber Vault

Bill White

Karen Smolar

Dino Amarini

John Thompson

Diego Bessone

Paolo Vitali

Tom Bish

Joseph Welsh II

Nathan Brice

Richard Cairns

Giovanni Cerquone

Nick Clayton

Michael Frankenberg

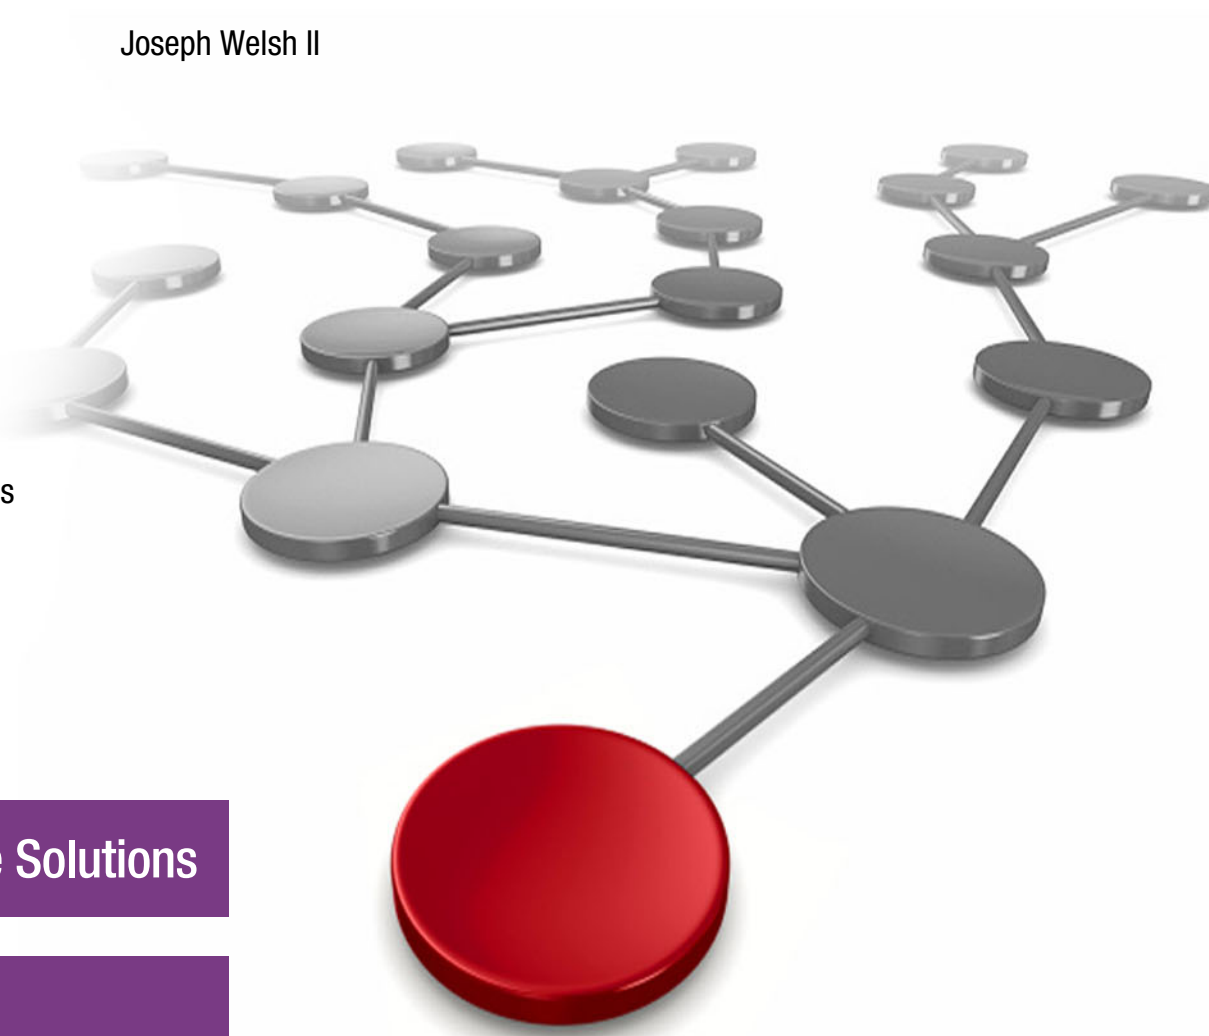
Nathan Gurley

Maryellen Kliethermes

David Matoe

Kevin Miner

Nadim Shehab



Infrastructure Solutions

IBM Z





IBM Redbooks

**Getting Started with IBM Z Cyber Vault**

January 2026

**Note:** Before using this information and the product it supports, read the information in “Notices” on page vii.

**Second Edition (January 2026)**

This edition applies to the IBM Z Cyber Vault solution.

© Copyright International Business Machines Corporation 2021, 2026. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

# Contents

<b>Notices</b> .....	vii
Trademarks .....	viii
<b>Preface</b> .....	ix
Authors .....	ix
Now you can become a published author, too! .....	xiii
Comments welcome .....	xiii
Stay connected to IBM Redbooks .....	xiii
<b>Chapter 1. Business resiliency: Proactive analysis and expedited recovery</b> .....	1
1.1 Common data corruption threats and remediation .....	2
1.1.1 Main factors that determine favorable outcomes .....	3
1.1.2 What is needed in a cyber resiliency strategy .....	4
1.1.3 Frameworks for IT cyber resiliency .....	4
1.2 What is cyber resiliency .....	5
1.2.1 Traditional resiliency versus cyber resiliency .....	5
1.2.2 Cybersecurity versus cyber resiliency .....	6
1.3 Capabilities for an effective cyber resiliency strategy .....	7
1.3.1 Characteristics of an ideal cyber resiliency solution .....	7
1.3.2 Capture, analysis, recovery, and restore characteristics .....	8
1.3.3 Data validation characteristics .....	8
1.3.4 Backup characteristics .....	8
1.4 IBM Z Cyber Vault solution and key capabilities .....	9
1.4.1 IBM Z full stack resiliency .....	11
1.4.2 IBM Z Cyber Vault capabilities .....	11
<b>Chapter 2. Planning and designing the IBM Z Cyber Vault solution</b> .....	15
2.1 Planning approach for the IBM Z Cyber Vault solution .....	16
2.1.1 Defining the architecture .....	17
2.1.2 Implementation considerations .....	18
2.1.3 Testing considerations .....	19
2.1.4 Operational considerations .....	19
2.2 IBM Z Cyber Vault solution reference architecture .....	20
2.2.1 Production environment .....	22
2.2.2 IBM Z Cyber Vault Storage .....	23
2.2.3 IBM Z Cyber Vault environment .....	23
2.2.4 IBM Z Cyber Vault automation .....	24
2.2.5 Architectural decisions .....	25
2.3 IBM Z Cyber Vault Storage considerations .....	27
2.3.1 DS8000 Safeguarded Copy concepts .....	28
2.3.2 Safeguarded Copy prerequisites .....	31
2.3.3 Safeguarded Copy capacity sizing considerations .....	32
2.3.4 DS8000 performance considerations for Safeguarded Copy backups .....	36
2.3.5 DS8000 Safeguarded Copy backup operational considerations .....	37
2.3.6 TS7700 and virtual tape solution considerations for IBM Z Cyber Vault .....	40
2.3.7 TS7700 and virtual tape best practices for IBM Z Cyber Vault .....	42
2.4 IBM Z Cyber Vault environment considerations .....	43
2.4.1 z/OS recovery system configuration .....	43
2.4.2 Network isolation .....	44

2.4.3	z/OS security . . . . .	44
2.4.4	Database and middleware consideration . . . . .	45
2.4.5	Offensive security environment . . . . .	46
2.5	IBM Z Cyber Vault automation considerations . . . . .	47
2.5.1	Safeguarded Copy and related storage processing . . . . .	47
2.5.2	Recovery system operation . . . . .	48
2.5.3	Data validation . . . . .	49
<b>Chapter 3. IBM Z Cyber Vault capabilities . . . . .</b>		<b>51</b>
3.1	Key design considerations . . . . .	52
3.1.1	Applying business and technical requirements . . . . .	52
3.1.2	IBM Z Cyber Vault environment design . . . . .	52
3.1.3	Developing and testing validation processes . . . . .	53
3.1.4	Selecting the tools and utilities for data validation and recovery . . . . .	53
3.1.5	Testing the IBM Z Cyber Vault environment . . . . .	54
3.1.6	Production cut-over and ongoing monitoring . . . . .	54
3.2	IBM Z Cyber Vault environment operational roles . . . . .	54
3.3	Environment setup . . . . .	56
3.3.1	Networking best practices . . . . .	56
3.3.2	Required software and tools . . . . .	57
3.3.3	Utilities for data validation by subsystem . . . . .	57
3.4	Validation, forensic analysis, and recovery processes . . . . .	58
3.4.1	Validation . . . . .	59
3.4.2	Forensic analysis . . . . .	62
3.4.3	Recovery process . . . . .	66
3.5	Extending the air gap . . . . .	75
3.6	Enhancing the security posture . . . . .	75
<b>Chapter 4. Deploying the IBM Z Cyber Vault environment . . . . .</b>		<b>77</b>
4.1	Preparing for IBM Z Cyber Vault deployment . . . . .	78
4.2	Setting up IBM Z Cyber Vault Storage . . . . .	80
4.2.1	Global Mirror secondary and Safeguarded Copy . . . . .	81
4.2.2	Recovery volumes . . . . .	82
4.2.3	Staging and persistent volumes . . . . .	82
4.2.4	Recovery for full data corruption . . . . .	83
4.3	Setting up Safeguarded Copy for IBM Z Cyber Vault Storage . . . . .	85
4.4	Setting up IBM Z Cyber Vault automation . . . . .	90
4.4.1	Automating Safeguarded Copy management . . . . .	90
4.4.2	Automating the IBM Z Cyber Vault IPL process . . . . .	91
4.5	Preparing the IBM Z Cyber Vault environment for validation . . . . .	92
4.5.1	Setting up CFs and LPARs for the IBM Z Cyber Vault environment . . . . .	93
4.5.2	Network isolation considerations for the IBM Z Cyber Vault environment . . . . .	95
4.5.3	Security considerations for the IBM Z Cyber Vault environment . . . . .	95
4.5.4	IODF and HCD definitions for the IBM Z Cyber Vault environment . . . . .	96
4.5.5	Couple Data Sets for IBM Z Cyber Vault environment . . . . .	97
4.5.6	IPL changes for the IBM Z Cyber Vault environment . . . . .	98
4.5.7	Parmlib changes for the IBM Z Cyber Vault environment . . . . .	99
4.5.8	CFRM changes . . . . .	100
4.5.9	TCP/IP and VTAM configuration considerations . . . . .	100
4.5.10	IBM Multi-Factor Authentication considerations . . . . .	100
4.5.11	z/OS dataset encryption considerations . . . . .	101
4.6	Establishing a validation framework . . . . .	101
4.6.1	Creating a validation framework . . . . .	102

4.6.2	Validation framework assumptions . . . . .	102
4.6.3	IBM Z Cyber Vault validation terminology . . . . .	103
4.6.4	Constructing data structure validation jobs . . . . .	104
4.6.5	IBM Z Cyber Vault job naming conventions . . . . .	105
4.6.6	IBM Z Cyber Vault datasets . . . . .	106
4.6.7	IBM Z Cyber Vault validation monitoring task . . . . .	106
4.7	Cyber Vault Data Validation for IBM Z Asset . . . . .	109
<b>Chapter 5. IBM Z Cyber Vault: Building on the foundations . . . . .</b>		<b>111</b>
5.1	The evolution of the IBM Z Cyber Vault solution . . . . .	112
5.1.1	Roadmap for the IBM Z Cyber Vault solution . . . . .	112
5.1.2	IBM Threat Detection for z/OS: Integration with IBM Z Cyber Vault . . . . .	113
5.1.3	Enhancing IBM Z Backup Resiliency for surgical recovery . . . . .	114
5.2	Db2 for z/OS subsystem rollforward recovery . . . . .	114
5.2.1	IBM Db2 Recovery Expert Pro for z/OS: Subsystem rollforward recovery . . . . .	116
5.3	Advanced cyber resilience for IBM Z . . . . .	117
5.3.1	Offline backups . . . . .	117
5.3.2	Offensive security . . . . .	118
5.4	Operational practices for enhancing integrity . . . . .	120
5.4.1	Case for monitoring the IBM Cyber Vault environment . . . . .	120
5.4.2	Ensuring that the environment is operating as expected . . . . .	121
5.4.3	Fortifying the network . . . . .	121
5.4.4	Monitoring file integrity . . . . .	122
5.4.5	Reporting on the state . . . . .	122
5.4.6	Security considerations . . . . .	125
<b>Appendix A. Monitoring the IBM Z Cyber Vault environment . . . . .</b>		<b>127</b>
	Monitoring: Key terms and definitions explained . . . . .	128
	IBM Security zSecure Alert capabilities . . . . .	129
	A closer look at FIM . . . . .	131
	Monitoring matters . . . . .	135
<b>Abbreviations and acronyms . . . . .</b>		<b>137</b>
<b>Related publications . . . . .</b>		<b>139</b>
	IBM Redbooks . . . . .	139
	Other publications . . . . .	139
	Online resources . . . . .	139
	Help from IBM . . . . .	140



# Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of 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.

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

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

## COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

# Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at <https://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

AIX®	HyperSwap®	Redbooks®
CICS®	IBM®	Redbooks (logo)  ®
Db2®	IBM Security®	Tivoli®
DS8000®	IBM Z®	VTAM®
FICON®	OMEGAMON®	z/OS®
FlashCopy®	Parallel Sysplex®	z17™
GDPS®	RACF®	

The following terms are trademarks of other companies:

ITIL is a Registered Trade Mark of AXELOS Limited.

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.

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

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

Other company, product, or service names may be trademarks or service marks of others.

# Preface

Cyberattacks and data corruption threats continue to escalate, which makes cyber resiliency a critical component of an enterprise IT strategy. Although traditional disaster recovery (DR) focuses on physical failures, cyber resiliency addresses logical corruption that is caused by ransomware, malware, or insider threats. This IBM Redbooks® publication introduces the IBM Z® Cyber Vault solution, which is a comprehensive approach to safeguarding IBM Z environments through immutable backups, isolated recovery systems, and automated validation processes.

The IBM Z Cyber Vault solution combines three core domains:

- ▶ Storage, which uses IBM DS8000® Safeguarded Copy for point-in-time, tamper-proof backups
- ▶ Automation, which uses Geographically Dispersed Parallel Sysplex® (IBM GDPS®) Logical Corruption Protection (LCP) Manager and IBM Copy Services Manager (CSM) to orchestrate capture, validation, and recovery
- ▶ Environment, which is an air-gapped IBM Z system for forensic analysis and surgical or catastrophic recovery

The solution supports proactive data validation, forensic investigation, and recovery strategies that minimize downtime and data loss, and integrate with advanced tools such as IBM Z Backup Resiliency (IZBR) and IBM Threat Detection for z/OS (TDz).

This publication provides practical guidance for planning, designing, and deploying the IBM Z Cyber Vault solution, including architecture decisions, storage sizing, automation scripts, and validation frameworks. It also explores advanced capabilities such as offline backups, offensive security (OffSec) testing, and IBM Db2® rollforward recovery.

Intended for IT managers, architects, system programmers, and security professionals, this publication equips organizations to strengthen cyber resiliency and help ensure business continuity (BC) in the face of evolving threats.

## Authors

This book was produced by a team of specialists from around the world working with the IBM Redbooks team, Poughkeepsie Center.

**Bill White** is an IBM Redbooks Project Leader and Senior IT Infrastructure Specialist at IBM Redbooks, Poughkeepsie Center.

**Dino Amarini** is a seasoned Storage Administrator and subject matter expert (SME) with more than 42 years of experience in IBM® mainframe software and hardware, and mainframe and distributed DASD and virtual tape storage subsystems. He joined IBM US in 2020 and initially worked with IBM Technology Expert Labs (TEL), where he developed and maintained a storage infrastructures and the IBM LABPLEX environment to support client engagements worldwide.

Before joining IBM, Dino served for more than 20 years as Director of Mainframe and Distributed Storage at AXA, a global insurance company, where he led enterprise storage strategy and operations. Dino has been a key contributor to IBM Z Cyber Vault Data Validation initiatives since 2025. He supports deployments across the US and worldwide, with a primary focus on implementing and optimizing IBM Z Cyber Vault Data Validation topologies. Dino holds a Bachelor of Science in Management Information Systems degree with a concentration in computer science.

**Diego Bessone** is a Global Software Sales Director for IBM Z. He entered the mainframe field in 1987 when he started working at the data center of Aerolíneas Argentinas, the flag carrier of Argentina. After 10 years, he continued to use his experience as a technical consultant for the IT departments of multinational companies, financial institutions, utility providers, and government agencies. Diego joined IBM in 1998 in Argentina, where he led the sales strategy for the IBM mainframe systems management portfolio. He moved to the US in 2001 to lead the IBM Tivoli® S/390 Latin America team, and added business operations management to his role. He began leading IBM Z middleware software sales at a worldwide level in 2009.

In July 2022, Diego was promoted to an executive role as IBM Z Sales Director, managing global sales of IBM Z high-end platforms. Since 2024, he has led the IBM Z Software sales strategy for pricing, and strategic initiatives such as IBM Z Cyber Vault and IBM Z Cyber Security. He has written extensively on data management, systems management, mission-critical systems, and resiliency.

**Tom Bish** has been with IBM for 34 years and is an IBM Senior Technical Staff Member, Principal Architect, and Master Inventor. Tom worked in IBM storage product development and architecture for 25 years and then joined IBM Global Technology Services to define their software-defined storage strategy. Tom is part of the IBM Advanced Technology Group (ATG) within technical sales, where he helps clients and sellers use IBM storage technologies effectively.

**Nathan Brice** Nathan Brice is the GDPS and IBM Z Cyber Vault Product Manager. He is based in the United Kingdom and has worked for IBM in many roles for more than 25 years. Over the last 10 years, he has held Product Manager roles in GDPS, IBM Customer Information Control System (IBM CICS®), and, during an assignment in North Carolina, in the IBM AIOps team. During this time, he presented at conferences around the world and wrote many articles on IBM products and technology.

**Richard Cairns** is the Worldwide IBM Z Security and Cyber Resiliency Software Sales Leader. He is based in the United Kingdom. He has more than 40 years of experience working in sales, technical sales, and various leadership roles within the IBM Z business unit, with assignments in South Africa and Central and Eastern Europe. Richard's current area of expertise includes the IBM Z Cyber Vault solution.

**Giovanni Cerquone** is an IBM Certified IT System Management Specialist working for IBM Expert Labs, supporting multiple clients with a focus on IBM z/OS® security. He has more than 35 years of experience with IBM mainframe technologies and joined IBM in 2007. Giovanni holds a degree in computer science from Central de Venezuela University. His areas of expertise include z/OS security, IBM CICS TS, IBM RACF®, and the IBM zSecure suite of products.

**Nick Clayton** works in IBM Infrastructure development as the Solution Architect for the DS8000 Storage System and focuses on solution integration and product strategy. His specific interests include storage performance, replication technology, and business resilience. He is a member of the GDPS design team. He also works with clients on their storage strategy and advises them on their deployment of IBM Storage technology. Nick is an author and co-author of many patents, articles, white papers, and IBM Redbooks publications on IBM storage technology. He is a regular presenter at storage conferences. Nick graduated from Trinity College, Cambridge, in 1994 with a degree in mathematics and has held previous roles both within and outside IBM in the areas of IBM Parallel Sysplex®, performance, and enterprise storage.

**Michael Frankenberg** is a Certified IT Specialist in Germany with more than 25 years of experience in high-end storage. He works in IBM Technical Sales Support, EMEA, and his areas of expertise include performance analysis, establishing high availability and disaster recovery (HADR) solutions, and implementing IBM Storage Systems. Michael supports the introduction of new IBM storage products and provides advice to clients, IBM Business Partners, and IBM Technical Sales personnel. He holds a degree in electrical engineering and information technology from the University of Applied Sciences Bochum in Germany.

**Nathan Gurley** is a GDPS developer who is based in the US. He has 5 years of experience in the disaster recovery (DR) and corruption protection fields. He holds a degree in computer science from Miami University. His areas of expertise include GDPS Logical Corruption Protection (LCP) and security on z/OS. He has presented at conferences in the US and Europe about role-based and dual control security within GDPS and enhancing the capabilities of the IBM Z Cyber Vault solution.

**Maryellen Kliethermes** is a Senior Project Manager who is based in the US with more than 30 years of experience in information technology. She held roles in the federal government and the utility industry as a developer, Manager of IT Software Development, and Manager of the IT Project Management Office. She holds a Bachelor of Science degree in computer information systems, a Master of International Business degree from St. Louis University, and certifications in PMP, Agile, and ITIL. Most recently, she served as Lead Project Manager at IBM for deploying the IBM Z Cyber Vault solution and IBM Z pervasive encryption solutions at clients worldwide.

**David Matoe** is a New Zealand-born GDPS practitioner who has worked for IBM in the United Kingdom, Poland, and Australia. He has more than 40 years of experience with IBM mainframe software and hardware and 25 years of experience with NetView, GDPS, and System Automation. David first joined IBM in 1995, working for IBM Global Networks, where he developed and maintained SNA Automation for European clients before moving into the newly formed GDPS team in the United Kingdom. David eventually moved to IBM Perth, Australia, where he supported the z/OS feeds into IBM Tivoli Business Service Manager and provided GDPS consulting services to Australian clients. He is currently a GDPS practitioner who assists clients with the deployment of all GDPS topologies, including GDPS Continuous Availability. David graduated with a New Zealand Certificate in Civil Engineering.

**Nadim Shehab** is a Senior GDPS developer who is based in the US and has 20 years of experience in the D and corruption protection fields. His areas of expertise include the development of GDPS LCP and CS. Nadim has presented at conferences around the world regarding LCP and CS. He holds a Bachelor of Science degree in computer engineering from the University of Arizona.

**Karen Smolar** is a Principal Solution Architect who is based in the US. She has 36 years of experience in the information technology field. She holds a degree in computer science from Clarkson University. Her areas of expertise include helping clients achieve success by designing infrastructure and application solutions that meet their business requirements. She has written extensively on all aspects of resiliency and BC planning.

**John Thompson** is a Senior Technical Staff Member at IBM who focuses on IBM Z resilience. He has worked in mainframe storage software at IBM for more than 36 years, and his current role is to design solutions for the GDPS LCP Manager and IBM Z Cyber Vault. John works with clients to improve and modernize their mainframe resilience position. He also presents GDPS and IBM Z Cyber Vault topics at conferences and councils around the globe.

**Paolo Vitali** is a Thought Leader IT Specialist who has worked at IBM since 1985 in the mainframe area. He spent 26 years at GTS Italy working with financial sector enterprises on z/OS Parallel Sysplex implementation and tuning. From 2009 to 2012 at the Benchmark Center in Montpellier, France, he performed several benchmarks and proofs of concept (PoCs) for customers worldwide. Paolo joined the GDPS Solution Test Team in Montpellier in 2013, where he focuses on 3-site, 4-site, and 6-site GDPS solutions with GDPS Metro Global Mirror (MGM). More recently, he worked in GDPS LCP design and testing and IBM Z Cyber Vault.

**Joseph Welsh II** is an IBM Senior Management IT Consultant and Certified IT Network Specialist who is based in the US. He joined IBM in 1988 with a Bachelor's degree in Computer Science from Transylvania University. He has held the roles of software developer, designer, and tester for IBM z/OS Communications Server (IBM VTAM® and TCP/IP). Since 1998, he has performed IT consulting services engagements that are focused on SNA, Advanced Peer-to-Peer Networking (APPN), high-performance routing (HPR), Enterprise Extender, and networking security for the Communications Server for IBM AIX®, Linux, Microsoft Windows, and z/OS at Fortune 500 companies around the world.

His work includes SNA, APPN, HPR, TCP/IP, and IP Security education and training, developing network designs and migrations, defining strategy and product direction, performing problem determination, and delivering implementation and installation and migration assistance. In recent years, he has provided IBM Z Networking Security consulting services and deployed IBM Z Cyber Vault solutions to clients worldwide.

A special thanks to the following people for their contributions to this project:

Anthony Ciabattoni  
**IBM Db2 for z/OS SWAT Team Leader**

Dave Clitherow  
**IBM GDPS Product Manager (Retired)**

Dave Petersen  
**IBM Distinguished Engineer**

Vijay Radhakrishnan  
**IBM GDPS Solution Architect**

Thanks to the authors of the first edition, "Getting Started with IBM Z Cyber Vault", published in July 2021:

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, and Knud Vraa

## Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an IBM Redbooks residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at: [ibm.com/redbooks/residencies.html](http://ibm.com/redbooks/residencies.html)

## Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks form found at:  
[ibm.com/redbooks](http://ibm.com/redbooks)
- ▶ Send your comments in an email to:  
[redbooks@us.ibm.com](mailto:redbooks@us.ibm.com)
- ▶ Mail your comments to:  
IBM Corporation, IBM Redbooks  
Dept. HYTD Mail Station P099  
2455 South Road  
Poughkeepsie, NY 12601-5400

## Stay connected to IBM Redbooks

- ▶ Find us on LinkedIn:  
<https://www.linkedin.com/groups/2130806>
- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:  
<https://www.redbooks.ibm.com/subscribe>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:  
<https://www.redbooks.ibm.com/rss.html>





# Business resiliency: Proactive analysis and expedited recovery

Businesses, organizations, and individuals face information technology threats of greater severity and cost than ever before. New forms of data corruption and malware (ransomware) are a constant threat that target computer systems and impede service delivery.

Victims of data corruption typically suffer in the following ways:

- ▶ Loss of trust and reputation both individually and for a company's brand
- ▶ Inability to provide goods and services to their customers or citizens
- ▶ Large financial costs to investigate, remediate, and recover from the attack

The increased risk demands that precautions be taken against data corruption and at the same time comply with new and updated regulations, which includes reporting incidents and attacks, both privately to insurance, government, or regulatory agencies, and publicly, and demonstrating a thorough and tested plan to prevent, detect, and recover from data corruption.

The following topics are covered in this chapter:

- ▶ 1.1, “Common data corruption threats and remediation” on page 2
- ▶ 1.2, “What is cyber resiliency” on page 5
- ▶ 1.3, “Capabilities for an effective cyber resiliency strategy” on page 7
- ▶ 1.4, “IBM Z Cyber Vault solution and key capabilities” on page 9

## 1.1 Common data corruption threats and remediation

Cyberattacks are constantly evolving, and new variants of malware and ransomware are increasing in number and financial severity. Malware and ransomware are often used as general terms, but they are different. *Malware* is software that attackers use to gain unauthorized access to IT systems and steal data or disrupt services. *Ransomware* is a form of malware, but it fences off specified data or IT systems with a demand for payment to remove the restrictions.

Attackers might use malware to implant ransomware weapons and programs into IT infrastructure assets. Malware might also use methods such as intentional deletion or erasure of data.

External entities typically launch ransomware and malware attacks after they successfully attack a business by using phishing, identity theft, or other scams. The most common attacks start with an email with infected attachments, a smartphone message, or multimedia. Another common attack vector is infected websites that exploit unpatched browser vulnerabilities, which, although less common than email attacks, are still a successful way to gain entry to IT environments.

Businesses and organizations that detect and respond to these cyberattacks must now consider the following scenarios:

- ▶ Unusable data (through unauthorized encryption)
- ▶ Loss of data in a file
- ▶ Loss of files, file systems, and databases

To minimize the risk these cyberattacks pose, protection is required at all levels of the IT environment:

- ▶ Computer systems and infrastructure, such as operating systems, clustering technology, storage systems, and disaster recovery replication
- ▶ Application middleware and run times, database servers, and core data file systems
- ▶ Application and business data that is stored in databases and file systems

Insider attacks are also on the rise. An *insider attack* is a data breach that occurs when an employee or contractor within an organization intentionally or accidentally exposes sensitive data. Insider threats originate with authorized users, such as employees, contractors, and business partners, who misuse their legitimate access or have their accounts hijacked by cybercriminals.

Another example of unintentional data corruption is when bad logic is introduced into an application code update. Subsequent fixes that attempt to fix the issue sometimes exacerbate the problem and lead to further data loss. Alternatively, simple human error in commands or processes, or inadequate testing and backup practices, can also be factors in this type of data error.

Although external threats are more common and grab the biggest cyberattack headlines, insider threats, whether malicious or the result of negligence, can be more costly and dangerous. According to the [Cost of a Data Breach Report](#), data breaches that are initiated by malicious insiders were the costliest, at USD 4.99 million on average.<sup>1</sup>

---

<sup>1</sup> Jointly produced by the Ponemon Institute and IBM Security®.

Another increasingly common attack vector involves accessing systems with stolen credentials. Sophisticated, targeted spear phishing attacks on individuals are becoming far more common. The [IBM X-Force Threat Intelligence Index report](#) finds that there has been a 71% year-over-year increase in cyberattacks that used stolen or compromised credentials.

As the IBM Z platform is a critical part of the data infrastructure for many organizations, particularly in industries such as finance, healthcare, and insurance, it is now more than ever the focus of malicious attacks and attempts to steal or corrupt data. A simple online search makes it simple to find publicly available exploits that might surface unwanted exposures and exploit insufficient security practices. Although IBM Z might be the most securable platform, data never rests or stays in one place. IBM Z is still a server, and no server is immune to attacks because they are all at risk. Although it can be argued that the likelihood of a breach occurring on IBM Z is reduced compared to other types of servers, it can also be argued that the impact of a breach on an IBM Z platform is greater than that of other types of servers because of the critical role IBM Z platforms play. As a result, the risk is a real and significant concern.

### 1.1.1 Main factors that determine favorable outcomes

Organizations must plan and prepare to evaluate the impact on business reputation, the implications of regulatory compliance, and the real risk of financial loss from a cyberattack.

A strong correlation exists between the consequences of cyberattacks and the consequences of data breaches, especially because ransomware has evolved and most likely includes a data theft element.

The [Cost of a Data Breach Report](#) also found that the following actions reduced the financial and brand impacts of a data breach while also reducing the time to detect, respond to, and recover from the breach:

- ▶ Know your information landscape.
- ▶ Strengthen prevention strategies with AI and automation.
- ▶ Take a security-first approach to gen AI adoption.
- ▶ Level up your cyber response training.

These actions build on well-established guidelines:

- ▶ Invest in security orchestration, automation, and response.
- ▶ Adopt a zero-trust security model to help prevent unauthorized access to data.
- ▶ Stress test your incident response plan to increase cyber resilience.
- ▶ Use tools to protect and monitor endpoints and remote employees.
- ▶ Invest in governance, risk management, and compliance programs.
- ▶ Minimize the complexity of IT and security environments.

Although a perfect cybersecurity solution does not exist, a thought-out, designed, and tested cybersecurity strategy can minimize the risk of attacks. Most importantly, a comprehensive cyber resilience strategy can minimize the impact of a cyberattack by enabling rapid recovery.

## 1.1.2 What is needed in a cyber resiliency strategy

A comprehensive cyber resiliency strategy always includes the following elements:

- ▶ Identification

This function focuses on the need for an enterprise to understand its most important assets and resources. These assets include categories such as asset management, business environment, governance, risk assessment, risk management strategy, and supply chain risk management. IT resources are a subset of these high-level, business-oriented categories.

- ▶ Protection

This function covers the technical and physical security controls that support the development and implementation of appropriate safeguards for protecting critical infrastructure. These categories are identity management and access control, awareness and training, data security, information protection processes and procedures, maintenance, and protective technology.

- ▶ Detection

This function focuses on measures that alert an organization to cyberattacks. These measures might include anomalies and events, continuous security monitoring, and early detection processes.

- ▶ Incident response

This function helps ensure an appropriate response to cyberattacks and other cybersecurity events. Categories include response planning, communications, analysis, mitigation, and improvements.

- ▶ Recovery

This function covers the implementation of plans for cyber resilience to help ensure business continuity (BC) if there is a cyberattack, security breach, or another cybersecurity event. The recovery functions are recovery planning, improvements, and communications.

These key elements are based on information in the [NIST Computer Security Incident Handling Guide](#). For more information, see [Data Integrity Detecting and Responding to Ransomware and Other Destructive Events](#).

Several frameworks are also available that organizations can use to improve readiness while building and enhancing a cyber resiliency strategy.

## 1.1.3 Frameworks for IT cyber resiliency

Specific regulations and frameworks vary by country or region, although it is worthwhile to evaluate multiple frameworks even if they are created by another country or entity. One commonly cited framework that was released in 2013 and updated in 2018 by the National Institute of Standards and Technology is the [Framework for Improving Critical Infrastructure Cybersecurity](#). Another framework is the [ICT risk management framework](#) from the Digital Operational Resilience Act, which organizations doing business in the European Union should evaluate.

Complementary [International Organization for Standardization \(ISO\)](#) documents ISO 31000 and ISO 27005 can map to the guidelines in the [Framework for Improving Critical Infrastructure Cybersecurity](#) or to other international and industry regulations.

The *Framework for Improving Critical Infrastructure Cybersecurity* is a comprehensive document that can guide an entity from initial risk evaluation and planning through steps to respond to and evaluate plans for future events. Here are some examples.

- ▶ PR.IP-4  
Backups of information are conducted, maintained, and tested.
- ▶ PR.IP-7  
Protection processes are improved.
- ▶ PR.IP-10  
Response and recovery plans are tested.
- ▶ DE.AE-2  
Detected events are analyzed to understand attack targets and methods.
- ▶ RS.RP-1  
A response plan is run during or after an incident.
- ▶ RS.AN-3  
Forensics are performed.
- ▶ RC.RP-1  
A recovery plan runs during or after a cybersecurity incident.

## 1.2 What is cyber resiliency

Cyber resiliency is an extension of traditional DR and BC solutions that many IBM Z clients have adopted. Cyber resiliency builds on the traditional DR building blocks of redundant systems, multiple copies of data, and replicating data to multiple locations.

A cyber resiliency strategy is vital for business continuity. It can provide benefits beyond increasing an enterprise's security posture and reducing the risk of exposure to its critical infrastructure. Cyber resiliency also helps reduce financial loss and reputational damage if there is a cyberattack.

### 1.2.1 Traditional resiliency versus cyber resiliency

A key difference between designing and deploying a traditional resiliency solution versus a cyber resiliency solution is that a traditional resiliency solution must protect against situations in which data is in its normal state but might need to be synchronized to a single point in time. A cyber resiliency solution must protect against and respond to situations in which systems and data are intentionally corrupted, erased, or encrypted.

For example, a traditional resiliency solution aims to provide protection against power failure or other physical issues. Examples include a weather event such as a hurricane that causes localized flooding and takes a production data center offline. In all cases, the primary goal of traditional resiliency solutions is to manage multiple copies of data so that they remain as consistent with the production data as possible. That protection might use a synchronized write to a second copy of data with complete consistency or an asynchronous write in which the secondary copy of data is several seconds behind the production copy.

If there is a cyberattack, any corruption, encryption, or erasure of data in production replicates rapidly to other copies of data and corrupts those additional copies. For cyber resiliency, organizations require protection against logical data corruption rather than physical issues, which is why a cyber resiliency solution takes a fundamentally different approach from a traditional resiliency solution and must complement an existing resiliency solution rather than replace it.

Another key differentiator involves how a team decides to run a recovery procedure. In an infrastructure-related incident, the system operations team has clear instructions and procedures in place to revert to a specific point in time or move the production site to an alternative location, depending on the issue that is flagged by monitoring systems. Most of these procedures are automated because they address known situations that are related to system or infrastructure malfunctions. However, when a data corruption event happens and the nature and extent of the incident are not known, the operations team must work with the applications and line of business (LOB) teams to determine the best approach to recovering corrupted data. This work requires more analysis and resources that must run in a clean-room environment with sufficient available capacity and tools to support the analysis and recovery processes.

## 1.2.2 Cybersecurity versus cyber resiliency

It is important to understand the difference between cybersecurity and cyber resiliency. The best way to understand this distinction is that cybersecurity focuses on preventing cyberattacks, and cyber resiliency focuses on minimizing the impact that a cyberattack or data breach might have on service delivery.

In today's environment, a hardened DMZ with physical security is no longer an option. Requiring all employees to access a company network only from a physical terminal in a secure office building is not viable. Working from home and mobile access are critical. To address this need, many businesses and organizations are adopting a zero-trust approach with three core principles for cybersecurity.

- ▶ Continuous monitoring and validation

Zero trust makes all network assets inaccessible by default. Users, devices, and workloads must pass continuous, contextual authentication and validation to access any resource, and they must pass these checks every time that they request a connection. Dynamic access control policies determine whether to approve requests based on data points such as user privileges, physical location, device health status, threat intelligence, and unusual behavior. Connections are continuously monitored and must be periodically reauthenticated to continue the session.

- ▶ The principle of least privilege

In a zero-trust environment, users and devices have least-privilege access to resources. This approach provides the minimum level of permission that is required to complete a task or fulfill a role. Permissions are revoked when the session is over. Managing permissions in this way limits the ability of threat actors to gain access to other areas of the network.

- ▶ Assume that a breach occurred

In a zero-trust enterprise, security teams assume that attackers have already breached network resources. Actions that security teams often use to mitigate an ongoing cyberattack become standard operating procedure. These actions include network segmentation to limit the scope of an attack; monitoring every asset, user, device, and process across the network; and responding to unusual user or device behaviors in real time.

However, even with strong cybersecurity protection in place, it is impossible to prevent all potential attacks. A successful cyber resiliency strategy considers how to minimize the impact of a successful cyberattack or data breach.

For example, after a ransomware attack, a business typically wants to restart its production environment from a recent, known good state to resume operations quickly. This capability can help avoid paying a ransom to unlock data, even assuming that the decryption capability works as expected. Sometimes, even after paying to receive the decryption key, a company might not return to full operations for an extended period.

The subsequent sections examine the capabilities of an ideal cyber resiliency strategy and IBM's solution, IBM Z Cyber Vault.

## **1.3 Capabilities for an effective cyber resiliency strategy**

A cyber resiliency solution must provide core capabilities in addition to standard backup, recovery of data or systems, and existing DR solutions. A full evaluation of the risks from cyberattacks demonstrates enterprise readiness with a plan in place to realize its value.

A significant challenge after any cyberattack is determining which systems are compromised.

Full cyber resiliency requires intrusion detection, monitoring for unusual behavior by individuals, programs, and systems, and reporting and dashboards to alert teams to this unusual behavior. All employees, contractors, and other people working with IT tools or systems must continue to receive education about how to prevent common attack points, such as phishing, smishing, vishing, or social engineering. They must also be trained to recognize and report unusual behavior. Investment and dedication to proper technologies, tools, processes, monitoring, education, and communication are critical before an incident occurs. These items are key to achieving enterprise-grade cybersecurity and cyber resiliency.

Early detection is only the first step of an effective cyber resiliency strategy. When a breach is identified, it is critical that the business quickly identifies which data is affected and to what extent, and determines the best approach to recover, replace, or re-create the corrupted data.

### **1.3.1 Characteristics of an ideal cyber resiliency solution**

An ideal cyber resiliency solution must protect against the unique challenges of a cyberattack and provide the capacity, tools, and resources that are required to address it.

The most critical characteristic is the ability to take regular, immutable copies of the production environment that attackers cannot corrupt or erase. These copies should be stored in a secure location, which is separated from production, with access governed separately from the production environment. By taking regular point-in-time copies, as opposed to continuous replication, it is possible to return to a known good copy of data that was captured before the corruption event.

The second key characteristic is that the system processes should be fully automated and not require any manual intervention. The goal is to take backups as regularly and as often as possible to reduce the window of time between backups. This result is possible only when this process is fully automated.

The third characteristic is having a dedicated vault or clean-room environment that exists entirely separate from production. This environment provides a secure location where backups can be restored and investigated. Access to this environment should be limited and managed separately from the production environment.

### **1.3.2 Capture, analysis, recovery, and restore characteristics**

If there is a data breach, this capability enables the IT staff to recover a specific point-in-time backup into the vault environment to support the analysis that is required to assess the incident. This work consists of formulating optimal recovery strategies and options, and determining the scope of recovery, which might include files, databases, or entire systems.

If only part of the production environment is corrupted, it might be possible to extract the required data from a point-in-time backup and use that data to fix the corruption in the running production environment. Alternatively, if the corruption in production is extensive, an entire point-in-time backup might be required as the basis for recovery. Both capabilities must be present in any cyber resiliency solution.

This capability provides the applications and LOB teams the tools that are required to make the best decision to restore their business to the best possible state after a serious cyber event.

### **1.3.3 Data validation characteristics**

Automation should continuously recover, start, and validate these backups in advance in the vault environment. The benefits of proactive, automated validations are as follows.

- ▶ They provide confidence that if a specific backup is needed, it is already validated as “good” and usable. If there is a data breach, there is pressure to restore production as quickly as possible. If the backup candidate is already validated, the IT staff can focus on the recovery strategy without needing more tests to determine whether this specific backup is a suitable candidate for the restore into production.
- ▶ Automating and regularly validating backups means that any validation failure might provide an early warning that production is corrupted. This corruption might not yet be detected in production, but a validation failure might trigger further investigation and potentially reduce the time that is required to identify certain types of data breaches.

### **1.3.4 Backup characteristics**

It is important to understand that there is a tradeoff between the capture interval and the retention duration for point-in-time backups. Ideally, the capture rate should be as small as possible, with the retention duration as large as possible. However, moving these values in different directions results in exponential growth in storage requirements. To address this tradeoff, different types of storage might be used across the solution.

Ideally, organizations should use disk for the small-interval initial captures. This capture rate should be in the single-digit-hours range, for example, every 1 - 4 hours. Each backup is then retained on disk for between 7 - 28 days to help ensure that, if needed, these backups can be rapidly restored.

However, it is possible that the system was compromised further back in time, and it might be necessary to go back to a backup that is several months old. This requirement can be met by archiving older backups on slower-access but more cost-effective storage, such as tape. By combining disk, virtual tape devices, and physical tape devices, it is possible to provide rapid access for frequently needed data while also enabling long-duration retention for data that is needed less often.

## 1.4 IBM Z Cyber Vault solution and key capabilities

Businesses and organizations must be both cyber secure and cyber resilient. The IBM Z Cyber Vault solution is built on point-in-time immutable copies of data from a production environment that are stored in an isolated, secure, and clean location on which regular and proactive data analytics run to validate the infrastructure, data structures, and data content. The IBM Z Cyber Vault environment is housed on an isolated IBM Z platform to prevent contamination of the validated immutable copies of the production data. If there is logical data corruption, organizations can use the IBM Z Cyber Vault solution to perform forensic, surgical, or catastrophic recoveries to the production environment.

Anticipating and responding to an event are increasingly complex, and the speed and precision of the response are increasingly critical.

Figure 1-1 shows the areas of concern when dealing with an event.

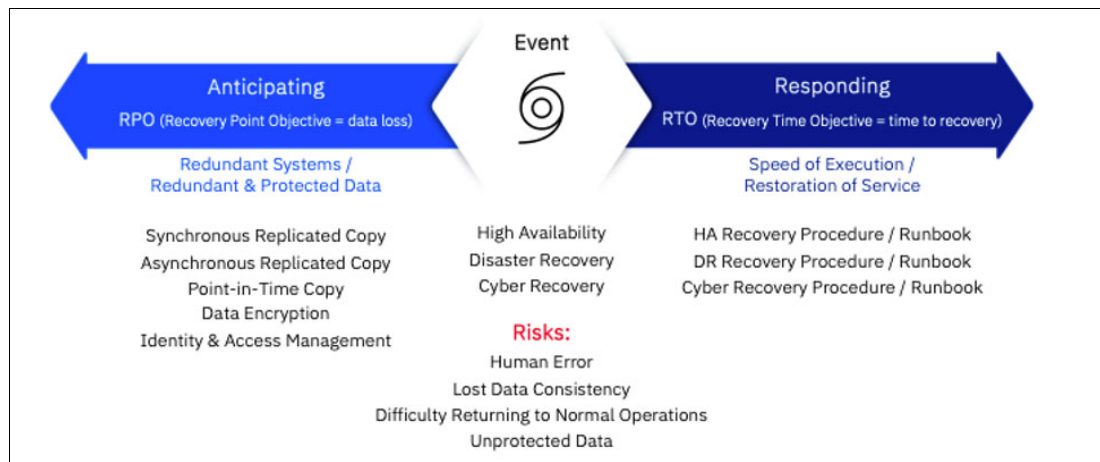


Figure 1-1 Anticipating and responding to an event

IBM learned a great deal by collaborating with businesses and organizations worldwide across every industry and by using best practices for systems and data center resiliency. Through that experience, IBM identified three areas or domains that any cyber resiliency solution should include. Only when organizations address these three domains with the right technology are business services better protected and capable of being restored to a point-in-time that makes it possible to restart business services more quickly and with the correct data.

These three domains can also serve as entry points to a full implementation of the IBM Z Cyber Vault solution. Organizations can decide where to start while they put together a deployment roadmap that ultimately implements all the provided capabilities.

The three domains of the IBM Z Cyber Vault solution include the following items.

► IBM Z Cyber Vault Storage

IBM Storage DS8000 is the latest innovation in enterprise-class storage for IBM Z. It is designed to help ensure the availability of critical business workloads and to reduce the business risk of outages while safeguarding sensitive data and meeting regulatory compliance requirements with advanced security features such as enhanced encryption, access controls, and data protection measures.

As the cornerstone of the IBM Z Cyber Vault solution, it prevents data from being modified or deleted because of user errors, malicious destruction, or malware or ransomware attacks by maintaining hundreds of immutable copies per volume. These copies can serve as a trusted source for surgical or full recovery of a production environment.

The Safeguarded Copy function supports the creation of cyber resilient point-in-time copies of volumes that attackers or users cannot change or delete. The DS8000 system integrates with IBM Copy Services Manager (CSM) to provide automated backup copies and data recovery.

► IBM Z Cyber Vault automation

The preferred technology for the IBM Z Cyber Vault automation domain is IBM Geographically Dispersed Parallel Sysplex (IBM GDPS), which is a family of DR and resiliency software for IBM Z. It manages the storage subsystem and remote copy configuration across heterogeneous platforms, automates IBM Parallel Sysplex operational tasks, and performs failure recovery from a single point of control.

GDPS Logical Corruption Protection (LCP) Manager is a core component of the IBM Z Cyber Vault solution and is integrated into the backbone of the IBM Z infrastructure. Administrators can manage LCP configurations, initiate data captures, and perform targeted data recovery through a dedicated interface within the GDPS management console.

IBM Technology Expert Labs (TEL) has developed special assets that use GDPS and LCP to drive data validation in the IBM Z Cyber Vault environment.

► IBM Z Cyber Vault environment

To stay compliant, protected, and prepared for any data corruption event, an isolated zero-trust clean room with an immutable data vault and rapid recoverability capabilities is required.

This isolated (air-gapped) environment can be physically or virtually separated from production, development, and test environments. It implements network and security safeguards that preserve the integrity of the clean room.

The IBM Z Cyber Vault environment must be sized according to the amount of data that it manages and must include the following components:

- IBM Z Cyber Vault environment Licensing (5770-ZCV), which provides the full IBM Z software production stack that is licensed for running in the environment.
- IBM Z software tools, which are used for validation, analysis, and recovery. These tools are also installed and available in the production environment because they typically collect data that is later used in the vault.

For more information about the IBM Z Cyber Vault solution, see Chapter 2, “Planning and designing the IBM Z Cyber Vault solution” on page 15.

### 1.4.1 IBM Z full stack resiliency

Designed for continuous availability and rapid DR, IBM Z provides industry-leading resiliency to protect a business from downtime, as shown in Figure 1-2. Implementing an IBM Z Cyber Vault solution helps optimize availability, keep systems running, detect problems in advance, and recover critical data.

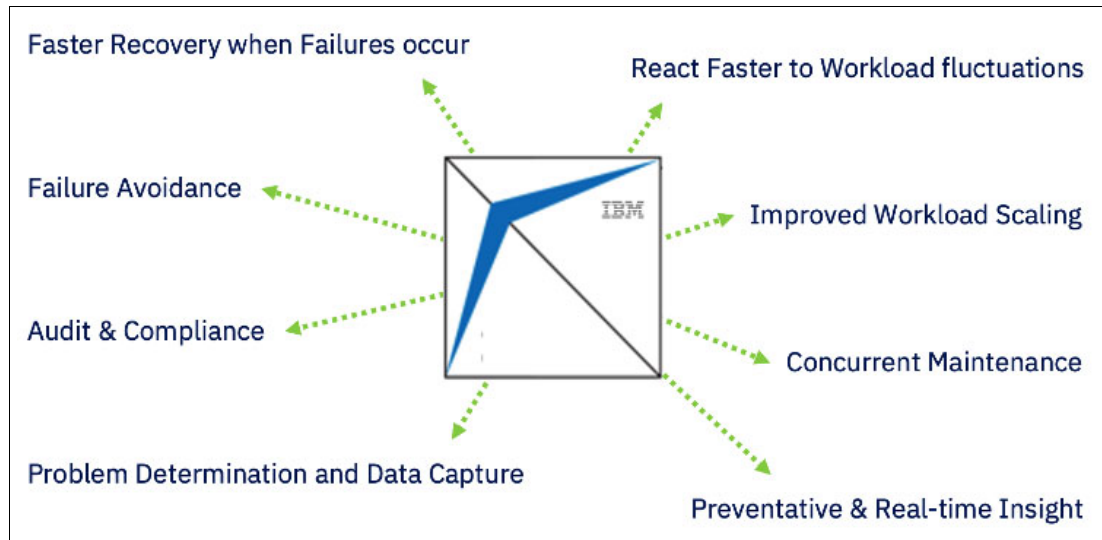


Figure 1-2 IBM Z continuous availability and rapid disaster recovery capabilities

### 1.4.2 IBM Z Cyber Vault capabilities

The IBM Z Cyber Vault solution delivers five key capabilities (see Figure 1-3) that businesses can use to prepare to work through any type of data corruption event.



Figure 1-3 IBM Z Cyber Vault - key capabilities

#### Data validation

Data validation is the process of running early and regular analytics to identify a data corruption situation. This analysis must be performed on all data because unless data is accessed, its status cannot be verified. Performing comprehensive data corruption detection and validation processes against a copy of production data is more practical and cost-effective than doing it in a live environment.

There are three types of validations that can be performed against the Safeguarded Copy backup in the IBM Z Cyber Vault environment:

► Type 1

Infrastructure validation, which consists of recovering a point-in-time copy of the production environment data into the IBM Z Cyber Vault environment and starting the system and its subsystems from this recovered copy. If the system starts correctly and resolves all inconsistencies when restarting the middleware (transactions in flight and other items), you have a working replica of production, which means that the infrastructure to run the business applications is working.

► Type 2

Data structure validation, where data structures are actively examined to identify inconsistencies or errors, often requiring specialized software to analyze data integrity. Some of the data structures and software solutions that help with validation are as follows:

- Db2, which uses Db2 utilities, **CHECK DATA/INDEX**, and log analysis
- Information Management System (IMS), which uses IMS utilities such as Pointer Checker
- z/OS Catalog, which uses Integrated Data Cluster Access Method Services (IDCAMS) and system management tools from IBM or other software vendors
- Virtual Storage Access Method (VSAM), which uses Indexcheck and Datacheck
- DFSMSHsm Control Data Sets and the DFSMSrmm catalog, which use z/OS and other IBM tools
- RACF (IRRUT200) database, which uses IBM zSecure Audit
- ISV software

► Type 3

Data content validation, where application data is analyzed to see whether the actual business data makes sense. This validation can be accomplished by special validation programs and procedures that must be developed by the application development teams because these teams are the ones that understand the business data and can determine whether it is valid or not. The IBM Z Cyber Vault solution provides the framework to automate the running and verification of these validation processes.

## Forensic analysis

Forensic analysis after a data corruption event involves using specialized software to examine a corrupted storage device, database, or file to develop a plan to recover as much usable data as possible and to document the process to preserve evidence, even when the data is partially damaged or overwritten, with the goal of identifying the cause of the corruption and potentially recovering valuable information.

## Surgical or catastrophic recovery

When the nature and extent of the data corruption event is understood, and a plan to recover lost data is determined, start the recovery process. Usually, the recovery should be a surgical recovery so that the user can restore only the affected data to production, usually by recovering data from one or more clean backups and rerunning processes to synchronize data across applications to achieve an application consistency point. Transactions for forward recovery are also part of this process if all the required logs and data are available and free of data corruption.

In a worst-case scenario, a catastrophic recovery is required, which involves taking the most recent, validated, and clean Safeguarded Copy, and restoring it into the production environment.

### **Offline backup**

Offline backups add a second layer of protection. When a recovered Safeguarded Copy backup is validated and determined to be clean, it can be stored in an isolated and offline environment that is physically separated from other environments to protect it from attacks or cyberthreats. This approach helps ensure a clean and uncompromised data recovery point if there is a major security breach. Magnetic tapes serve this function by storing the validated environment in an offline and physically isolated manner in a secure offsite facility, which creates an *air gap* that protects the data from cyberthreats by making it inaccessible to online systems. This separation acts as a last line of defense against cyberattacks by keeping backups disconnected from the network.

### **Offensive security**

Offensive security (OffSec) uses adversarial tactics to strengthen network security rather than compromise it. Ethical hackers typically conduct OffSec. These cybersecurity professionals use hacking skills to detect and fix IT system flaws and to identify security risks and vulnerabilities in the way that users respond to attacks.

OffSec measures that can help strengthen insider threat programs include phishing simulations and red teaming. In red teaming, a team of ethical hackers starts a simulated and goal-oriented cyberattack on the organization.

The IBM Z Cyber Vault environment provides an exact replica of the production environment to perform these exercises, and because it is isolated, there is no risk of harming other parts of the system.

For more information about using the IBM Z Cyber Vault capabilities, see Chapter 3, “IBM Z Cyber Vault capabilities” on page 51.





# Planning and designing the IBM Z Cyber Vault solution

This chapter describes a planning approach and the components that are required to implement the IBM Z Cyber Vault solution. The chapter describes planning and design considerations and explains the prerequisites that you need for the deployment of an IBM Z Cyber Vault solution.

The following topics are covered:

- ▶ 2.1, “Planning approach for the IBM Z Cyber Vault solution” on page 16
- ▶ 2.2, “IBM Z Cyber Vault solution reference architecture” on page 20
- ▶ 2.3, “IBM Z Cyber Vault Storage considerations” on page 27
- ▶ 2.4, “IBM Z Cyber Vault environment considerations” on page 43
- ▶ 2.5, “IBM Z Cyber Vault automation considerations” on page 47

## 2.1 Planning approach for the IBM Z Cyber Vault solution

This section provides the information that you need to start your cyber resiliency journey with the IBM Z Cyber Vault solution. It helps you understand each component of the solution, the capabilities that the components provide, and what you must consider as you design your IBM Z Cyber Vault solution.

Before you begin your design and implementation, consider your business requirements and IT objectives. You might not have all the answers right now. If you do not, using this list helps you start the cyber resiliency conversation with your teams and helps ensure that your solution satisfies your business requirements. The answers drive the design decisions for your solution.

- ▶ What data do you need to protect? What applications own the most critical data?
- ▶ How isolated does your solution need to be from your production environment?
- ▶ Do you have an enterprise cyber resiliency direction that you need to align with?
- ▶ How much data loss can the business tolerate?
- ▶ If you must recover to a previous point in time, do you have a requirement to replay transactions or updates that took place after that time?
- ▶ How long do you need to retain the data that you capture?
- ▶ Are there any regulations related to resiliency or cyber resiliency that need to be considered?
- ▶ What are your Recovery Time Objectives (RTOs)?

**Terminology:** Throughout this document, the text refers to logical partitions (LPARs), sysplex, and Parallel Sysplex. The environments that you are trying to protect and the related recovery environments can be any of these items. Because it is more common to have sysplexes or Parallel Sysplexes, much of this document uses those terms. If that consideration does not apply to your environment, you can substitute LPARs usually.

Because your IBM Z Cyber Vault solution must be integrated into your existing environment, understanding that environment is critical. As you answer these questions about your situation, also ask yourself whether there are strategic plans that might change your answers during the time you plan to implement your IBM Z Cyber Vault solution. If there are changes that are planned, design your solution based on the new architecture.

- ▶ How many physical data centers do you have and what is the purpose of each one?
- ▶ How many copies of data do you have and what combination of synchronous and asynchronous replication do you use?
- ▶ What is the scope of your replication? Is it the sysplex? If you do not have a sysplex, is it the LPAR?
- ▶ Do you failover and run production workload in your alternative data center? If so, for how long?
- ▶ How many sysplexes do you have? How many LPARs are in each sysplex? What runs in each one?
- ▶ What is your current backup and restore solution?
- ▶ How do you use virtual tape or other long-term storage solutions?

When you understand these items, you are ready to define your IBM Z Cyber Vault solution architecture. IBM can help you with this process. Contact your IBM team and ask about the “IBM Z Cyber Vault Discovery and Architecture Workshop”. The workshop helps you with the initial planning, defines the architecture, and provides an implementation roadmap.

### 2.1.1 Defining the architecture

The IBM Z Cyber Vault solution is based on a framework that includes three domains that work together to meet your cyber resiliency needs (see Figure 2-1). This book describes considerations, options, and implications for each of these domains in detail. This section introduces these components and explains how your requirements and IT environment influence the architecture that you define and ultimately implement.

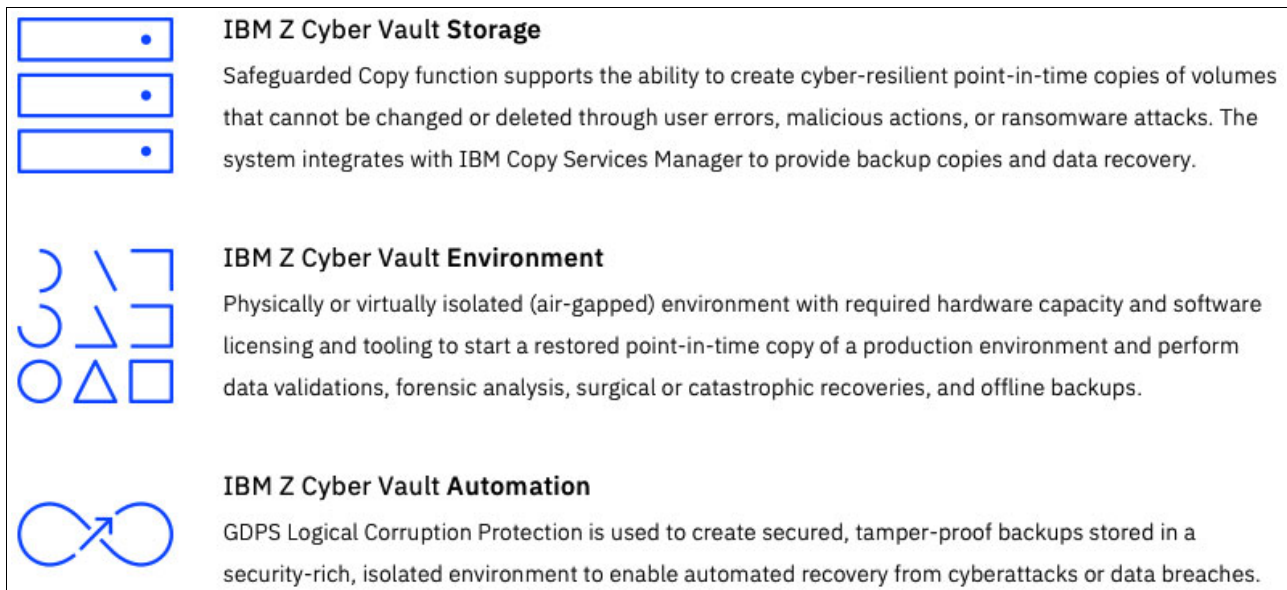


Figure 2-1 IBM Z Cyber Vault solution framework

The details in the rest of this chapter provide the information and guidance that you need to make these architectural decisions:

- ▶ Where will your IBM Z Cyber Vault solution be located? Will it be one of your existing data centers or a new location?
- ▶ Will you implement one of the virtual isolation topologies and use existing storage devices, or opt for physical isolation on dedicated storage devices?
- ▶ Will you use Geographically Dispersed Parallel Sysplex (IBM GDPS) Logical Corruption Protection (LCP) Manager or IBM Copy Services Manager (CSM) for storage management? Will you use GDPS LCP Manager for automation?
- ▶ Will you use GDPS LCP Manager for validation and recovery automation, or will you develop your own automation for these items?
- ▶ Will your IBM Z Cyber Vault environment be on IBM Z hardware that is also used for other workloads, or will it be dedicated to IBM Z Cyber Vault processing?
- ▶ Will you have access to any virtual tape data from your IBM Z Cyber Vault environment? If so, how will that data be accessed and protected?

## 2.1.2 Implementation considerations

When your architecture is defined, you know where your IBM Z Cyber Vault solution is, the source of the backups, whether the storage topology is virtually or physically isolated, and whether the IBM Z environment is on dedicated hardware or shares existing hardware. Now, consider what it takes to build, test, and operate the solution.

### Considerations for an IBM Z Cyber Vault Storage implementation

From a storage perspective, you have source volumes, Safeguarded Copy storage, recovery volumes, and persistent data storage. Your IBM Z Cyber Vault environment needs access to the recovery volumes and any persistent resource definitions that are used for maintaining validation scripts, logs, or event recordings. Based on your requirements, you also need to set up CSM or GDPS LCP Manager policies and scripts for capture frequency and retention.

### Considerations for an IBM Z Cyber Vault environment implementation

This task involves defining the specific configuration for the IBM Z Cyber Vault environment that is used for validation, forensic analysis, and recovery planning. This section briefly highlights topics that you need to consider.

- ▶ Network isolation

Your IBM Z Cyber Vault environment must be network-isolated. You do not want your IBM Z Cyber Vault processing to impact your production workload. Limit access to your IBM Z Cyber Vault environment. Your network team can help determine the best way to achieve that isolation based on your current network architecture.

- ▶ Security

Limit user access to your IBM Z Cyber Vault environment. Your Safeguarded Copy backup capacity is immutable, but ensure that the GDPS LCP or CSM policies that you use are secured properly so that bad actors cannot change those policies. You also need to consider the implications of user access to the IBM Z Cyber Vault environment.

- ▶ I/O connectivity

Your IBM Z Cyber Vault environment needs access to the recovery volumes where the Safeguarded Copy backups are recovered and the persistent volumes that contain scripts and reporting data.

- ▶ Data validation approach

Decide about how often you plan to validate environments, whether you want to use an IBM FlashCopy® or a recovered Safeguarded Copy backup, what types of data corruption that you want to detect, and how you want to report the results. These decisions are reflected in the automation scripts.

- ▶ z/OS configuration

A best practice is to duplicate the primary environments that you are protecting, minus the GDPS K-Systems if you have GDPS. This approach means that you should have the same general configuration from an LPAR and initial program load (IPL) perspective. For example, if you have an 8-way sysplex that you are protecting, define an 8-way sysplex in your IBM Z Cyber Vault environment. If you are protecting individual LPARs, define those LPARs in your IBM Z Cyber Vault environment. Resources such as processor capacity, memory, and coupling facilities (CFs) are likely fewer. How much less depends on factors such as the size of the environment and the complexity of the validations.

- ▶ **GDPS**  
If you use GDPS and GDPS LCP, consider how the configuration must be set based on the topology that you choose and the security requirements that you have.
- ▶ **Access to tape**  
If tape mounts are needed from your IBM Z Cyber Vault environment for validation or forensic analysis, the environment needs access to the tape storage environment.

### **Considerations for IBM Z Cyber Vault automation implementation**

Automation is critical for helping ensure that your IBM Z Cyber Vault solution provides consistent copies, regularly validates your data, and can efficiently restore services. GDPS LCP provides the framework and scripts for that automation.

If you do not use GDPS and GDPS LCP, you still must define processes and build automation scripts that can be used in your environment. CSM provides the Safeguarded Copy management. You need a solution for automating the validation processing. This approach needs to include an automated IPL process.

## **2.1.3 Testing considerations**

There are several types of testing that you must perform for your IBM Z Cyber Vault solution. Each type has a different objective.

- ▶ **Implementation verification testing** is traditional fit-for-purpose testing. This testing validates that the infrastructure is in place and working properly. Test cases include ensuring that the captures are taken, that the captures can be recovered to the recovery volumes, that the IBM Z recovery LPARs can be started, that validations run successfully, and that the automation works as expected. The results determine whether the solution can be considered production ready.
- ▶ **Operational procedures testing** is fit-for-use testing. Test cases include ensuring that the processes are defined and operational, that monitoring is in place, that operators are trained, and that everyone knows how to respond. Verify that the automation is in place and working as intended. Ensure that manual tasks are defined and documented and that operators know what to do if the automation fails.
- ▶ **Recovery testing** includes both functional and operational components. Test the recovery processes and automation that you put in place for both surgical and catastrophic recovery. Surgical recovery is unique to your environment. The infrastructure, processes, and automation that move data and apply it in production vary based on where your IBM Z Cyber Vault solution is. If the environments are in the same data center, you can copy the data to shared volumes. Otherwise, you need a network infrastructure in place to send the data between locations. Recovery processing might also require new processes and scripts in your production environments to apply the data. All these items need to be tested, for example:
  - Functionally tested to ensure that they work as expected
  - Regularly exercised as part of your business continuity (BC) testing plan

## **2.1.4 Operational considerations**

Ideally, your IBM Z Cyber Vault solution does not require much manual intervention. GDPS LCP automates the storage management policies and initiates the daily validations. To minimize the resources that are needed to manage your IBM Z Cyber Vault solution, implement monitoring and alerts for it.

Documenting operational procedures is critical, including expected outcomes and how to achieve them. Consider the following situations:

- ▶ Monitoring detects infrastructure issues, such as storage filling up or LPARs not starting properly.
- ▶ Validation processes fail and alerts are created.
- ▶ Response time or performance problems are identified.
- ▶ Corruption is suspected or identified.
- ▶ A security breach is detected.

Your environment is network-isolated, and you need to limit user access to it. It is important to understand the roles that are needed daily and to provide a mechanism for granting additional access only when needed. Daily operations might require access for a limited number of storage administrators, system programmers, network administrators, and monitoring or operations staff. There are two operational objectives:

- ▶ Ensure that the Safeguarded Copy captures are working and that there is no risk of running out of space.
- ▶ Ensure that validations are running as designed and, if alerts are generated, that they are investigated and addressed quickly.

If alerts are generated from the validation processes or if corruption is suspected, the objective changes to identifying the corruption, restoring services, and protecting the backup copies. This situation might require suspending captures until the problem is understood and corrected.

If there is corruption or other issues, other teams or team members might need access to the environment. For example, database administrators (DBAs) and application support teams might need access to perform forensic analysis and define recovery processes. You need a process to grant temporary access.

## 2.2 IBM Z Cyber Vault solution reference architecture

This section introduces the infrastructure and the underlying components that make up the IBM Z Cyber Vault solution. The section also explores several high-level architectural decisions that must be made when designing and planning for the implementation of the IBM Z Cyber Vault solution.

Figure 2-2 on page 21 depicts the reference architecture for the IBM Z Cyber Vault solution. This reference architecture helps you become familiar with the terminology and find the various components that are addressed in this chapter. The IBM Z Cyber Vault solution consists of three domains: IBM Z Cyber Vault Storage, IBM Z Cyber Vault environment, and IBM Z Cyber Vault automation.

For simplicity, the figure shows only the primary volume, IBM Z Cyber Vault source volume, and IBM Z Cyber Vault recovery volume in the data capture process. In most configurations, extra volumes are created by replication tools to support BC requirements. You might have up to five copies of production data, with any of those copies serving as the source for the IBM Z Cyber Vault volumes.

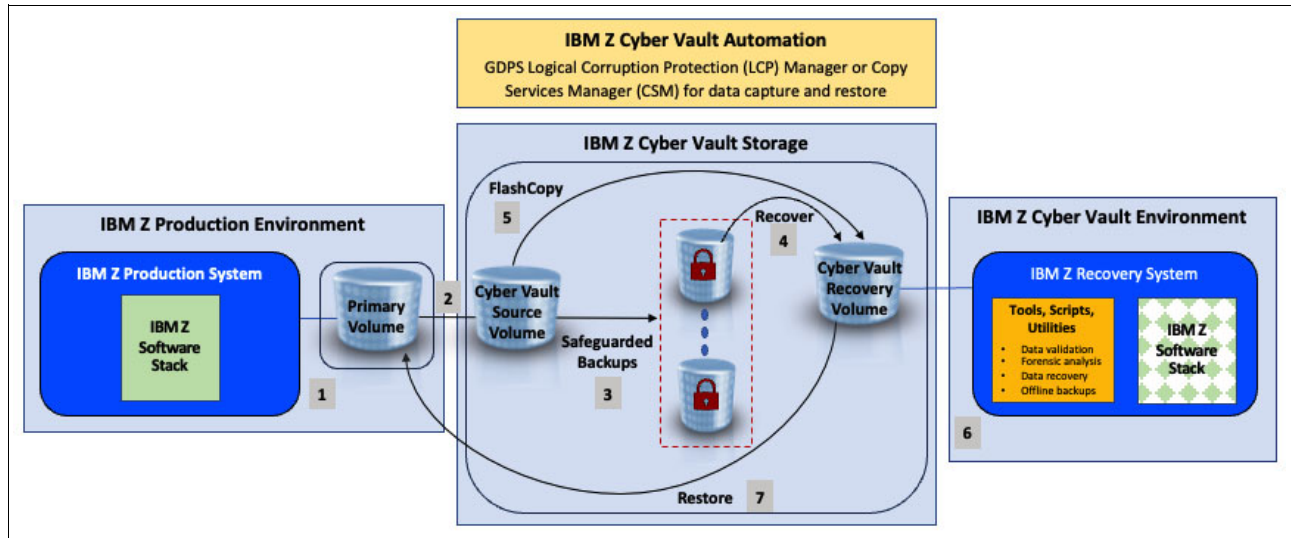


Figure 2-2 IBM Z Cyber Vault reference architecture

At its core, the IBM Z Cyber Vault solution uses isolated and immutable copies of the production data that are taken at multiple points in time, with recovery and restore functions. The solution also helps identify data corruption through a validation process. The IBM Z Cyber Vault solution capabilities are described in Chapter 3, “IBM Z Cyber Vault capabilities” on page 51.

As shown in Figure 2-2, the overall IBM Z Cyber Vault solution process works as follows:

1. Workloads running in an IBM Z production system read/write production data on the primary volume.
2. An asynchronous disk replication tool, such as GDPS Global Mirror (GM), mirrors that production data to the IBM Z Cyber Vault source volume.
3. Safeguarded Copy captures the production data from the IBM Z Cyber Vault source volume to create immutable copies.

CSM or GDPS LCP Manager, running in either the production environment or the IBM Z Cyber Vault environment, creates and manages the Safeguarded Copy backups.

Those Safeguarded Copy backups are used for data validation, analysis, recovery, or offline backup as needed.

4. The Safeguarded Copy backup is recovered to the IBM Z Cyber Vault recovery volume to perform data validation, analysis, recovery, or offline backup.
5. Optionally, when validation time is a consideration, you can use FlashCopy for validation instead of recovering a Safeguarded Copy backup. FlashCopy is faster than a recovery action and might provide performance advantages.

Immediately after capturing the Safeguarded Copy backups, FlashCopy duplicates the production data from the IBM Z Cyber Vault source volume to the IBM Z Cyber Vault recovery volume to perform an IPL of the IBM Z Cyber Vault environment and validate the current version of the production data.

GDPS LCP Manager or CSM initiates FlashCopy to create the IBM Z Cyber Vault recovery volume for data validation in the IBM Z recovery system.

6. Data validation, analysis, recovery, or offline backup takes place in the IBM Z Cyber Vault environment by using various tools, scripts, and utilities. The IBM Z Cyber Vault environment is a network-isolated and secured system that prevents contamination of the validated copy and protects the production environment.
7. Corrupted data in the IBM Z production system can be restored from the IBM Z Cyber Vault solution by using Safeguarded Copy backups. If the event is targeted, you can restore only the affected data. If the event is widespread, you can restore the entire production environment.

**Note:** When you use FlashCopy and Safeguarded Copy for IBM Z Cyber Vault, the IBM Z Cyber Vault Source and IBM Cyber Vault Recovery volumes must be in the same IBM DS8000 Storage System.

The IBM Z Cyber Vault solution is a combination of various components from the IBM Z platform, IBM Z software, and IBM Storage, including the following components:

- ▶ IBM DS8000 Storage with Safeguarded Copy, which enables the creation of point-in-time copies of production volumes that cannot be changed or deleted (immutable copies).
- ▶ IBM software for data validation, forensic analysis, and surgical recovery.
- ▶ IBM GDPS LCP Manager or IBM CSM, which creates secured and tamper-proof backups that are stored in an isolated environment.
- ▶ IBM Z (isolated LPARs), which restores point-in-time copies of the production environment and performs data validations, forensic analysis, surgical or catastrophic recoveries, and offline backups when needed.

## 2.2.1 Production environment

In the context of the IBM Z Cyber Vault solution, the production environment is composed of your existing business-critical workloads, infrastructure, and systems, plus other software products and security policies that are implemented to protect the production environment from cyberattacks and data corruption.

Because the business-critical data is in the production environment, it is where the implementation of your cyber resiliency strategy begins. The purpose is twofold: securing your production systems against cyberattacks and making them more resilient to endure cyberattacks without impacting expected service levels. The IBM Z Cyber Vault solution plays a key role by providing the ability to survive data corruption events despite severe impact. It helps you be better prepared to recover faster and minimizes the impact of possible downtime.

Preparing the production environment for the IBM Z Cyber Vault solution involves adding tools and using scripts and utilities. Some tools might need to be installed in the production environment but not be active because they are activated and used only in the IBM Z Cyber Vault environment. Other tools must be active in production because they collect and create data and metadata that is used in the IBM Z Cyber Vault environment.

Licensed software products are not installed or configured directly in the IBM Z Cyber Vault environment. Even the product preparations and configurations that are specific for the IBM Z Cyber Vault environment must be done in the production environment. This work includes creating the special IPL procedures that select the products to be active in the IBM Z Cyber Vault environment. The resulting software architecture is specific to each production environment. Existing tools and practices, and how new tools are integrated, must all be considered.

Regular validation of the production data in your IBM Z Cyber Vault environment provides a greater degree of confidence in your Safeguarded Copy backups. Those validations are not done in real time. Consider having real-time detection capability in the production environment. Real-time detection can provide immediate or earlier warning of a cyberattack or data corruption taking place.

## 2.2.2 IBM Z Cyber Vault Storage

Secure data replication technologies and immutable point-in-time copies are the underpinning of the IBM Z Cyber Vault solution. Together, these technologies create a powerful approach for protecting data and responding to various types of data corruption in an expedited manner.

The IBM DS8000 supports both technologies. The data replication technologies are Metro Mirror (synchronous mirroring) and GM (asynchronous mirroring), and they are supported by the DS8000. Metro Mirror and GM are hardware solutions that provide consistent updates across storage platforms and help ensure that those updates are applied in time sequence with a high degree of data integrity. The Safeguarded Copy function of the IBM DS8000 provides immutable point-in-time copies that can be restored if corruption of the primary and replicated copies occurs.

With Safeguarded Copy, you can have up to 500 consistent point-in-time copies per volume. To manage, create, recover, and expire Safeguarded Copy backups, CSM or GDPS LCP Manager is required. The point-in-time copies are not accessible by systems or applications because they do not have logical control units (LCUs) or volume IDs that are associated with them. They cannot be erased or deleted by using the DS8000 native interfaces. To access a Safeguarded Copy backup, a recovery action to an IBM Z Cyber Vault recovery volume is necessary so that the production data can be accessed from the IBM Z Cyber Vault environment (see Figure 2-2 on page 21).

Because FlashCopy and Safeguarded Copy functions are restricted to the same physical IBM DS8000, Metro Mirror or GM is required in configurations that have point-in-time copies that are spread across physically separate storage systems (see 2.2.5, “Architectural decisions” on page 25).

For more information about IBM Z Cyber Vault Storage, see 2.3, “IBM Z Cyber Vault Storage considerations” on page 27.

## 2.2.3 IBM Z Cyber Vault environment

The IBM Z Cyber Vault environment is where the IBM Z Cyber Vault solution runs. The solution’s capabilities include data validation, forensic analysis, recovery, any offline backups that are necessary to satisfy your data retention requirements, and offensive security (OffSec) procedures.

The IBM Z Cyber Vault environment consists of one or more recovery systems (LPARs) that start empty and remain isolated from the production environment. A recovery system starts from a point-in-time image of a production system with parameters that restrict the network to the IBM Z Cyber Vault environment and limit user access.

When Safeguarded Copy backups are selected, systems first restore them to the IBM Z Cyber Vault Recovery volumes. These volumes are then used to perform an IPL of the IBM Z Cyber Vault environment. Safeguarded Copy backups are exact replicas of the production environment. They include all production system volumes to help ensure broader protection.

All licensed software that is required to perform data validation, forensic analysis, and any type of recovery or security exercise must be installed in the production environment. Any software that is used only in the IBM Z Cyber Vault environment can remain inactive until it is started and used in the recovery system.

After all activities in the IBM Z Cyber Vault environment are complete, the recovery system or LPARs can be shut down, which prepares them for the next set of Safeguarded Copy backups. This mechanism keeps the IBM Z Cyber Vault environment isolated and helps ensure that Safeguarded Copy backups remain immutable.

The recovery system in the IBM Z Cyber Vault environment runs in an isolated network to help ensure that Safeguarded Copy backups cannot be accessed from the production environment or any other system. You can achieve this isolation by using dedicated Open Systems Adapter-Express (OSA-Express) features in the recovery system to provide physical isolation. Alternatively, when the IBM Z Cyber Vault environment uses virtual isolation, shared OSA-Express features can provide logical separation by using a separate virtual LAN (VLAN) and IP subnetwork with optional firewalls.

Both the recovery system environment and the production system environment can use built-in z/OS Communications Server security features.

- ▶ IP filtering blocks IP traffic that the system does not explicitly permit in its defined IP Filter Policy.
- ▶ IDS protects the system's services against various types of attacks.
- ▶ AT-TLS provides SSL/TLS encryption services at the TCP transport layer to protect sensitive application data in transit. AT-TLS is transparent to the application.

Because a recovery system starts by using a replica of a production system, including its configuration, the production environment must also implement all security features.

For more information, see 2.4, "IBM Z Cyber Vault environment considerations" on page 43.

## 2.2.4 IBM Z Cyber Vault automation

Automation in the running and monitoring of the data validation and recovery process can help reduce human intervention and minimize risk. Management software (such as GDPS LCP Manager) can be used to coordinate and perform repeatable steps with a high degree of confidence and lessen burden on operations.

GDPS LCP Manager is a core component of the IBM Z Cyber Vault solution that is integrated into the IBM Z infrastructure. It provides copy services management (as does CSM), and orchestration and automation for production systems and IBM Z Cyber Vault environment (recovery system LPARs) for data validations. GDPS is an IBM Z centric family of solutions that support high availability and disaster recovery (HADR) goals. There are different supported topologies that are based on the precise requirements for an environment. The common theme across these different solutions is the management of data replication, and management of system resources in the production and DR sites, with the correct management of system resources and orchestration of workflows to drive recovery actions when needed.

A key part of any GDPS LCP Manager is to help ensure that there is always a point of consistency for the production data that can be used to restore service from a previous point in time. This point is often referred to as a *crash or power-fail consistent copy* of data.

GDPS LCP Manager also provides automation capabilities and procedures that support the IBM Z Cyber Vault environment and simplify operations for data validation (for more information, see 2.4, “IBM Z Cyber Vault environment considerations” on page 43).

## 2.2.5 Architectural decisions

This section describes several high-level architectural considerations that influence the design and implementation of the IBM Z Cyber Vault solution.

One of the key architectural choices is how the required isolation for your IBM Z Cyber Vault solution is implemented. The components must be isolated from the production environment and have restricted access. This isolation can be done physically by implementing the solution components on separate, dedicated storage and IBM Z hardware or it can be done virtually on shared hardware. Physical storage isolation provides what is commonly called a physical air gap.

Figure 2-3 represents a physically isolated topology. Hardware-based replication is used to copy the production data into IBM Z Cyber Vault Storage environment through a secondary volume. An asynchronous replication method that is called GM or Global Copy (GC) is used for the replication into the IBM Z Cyber Vault Storage Source volume. The distance between the production environment and IBM Z Cyber Vault Storage can potentially be hundreds or thousands of kilometers, or it can be next to either the primary or secondary volumes.

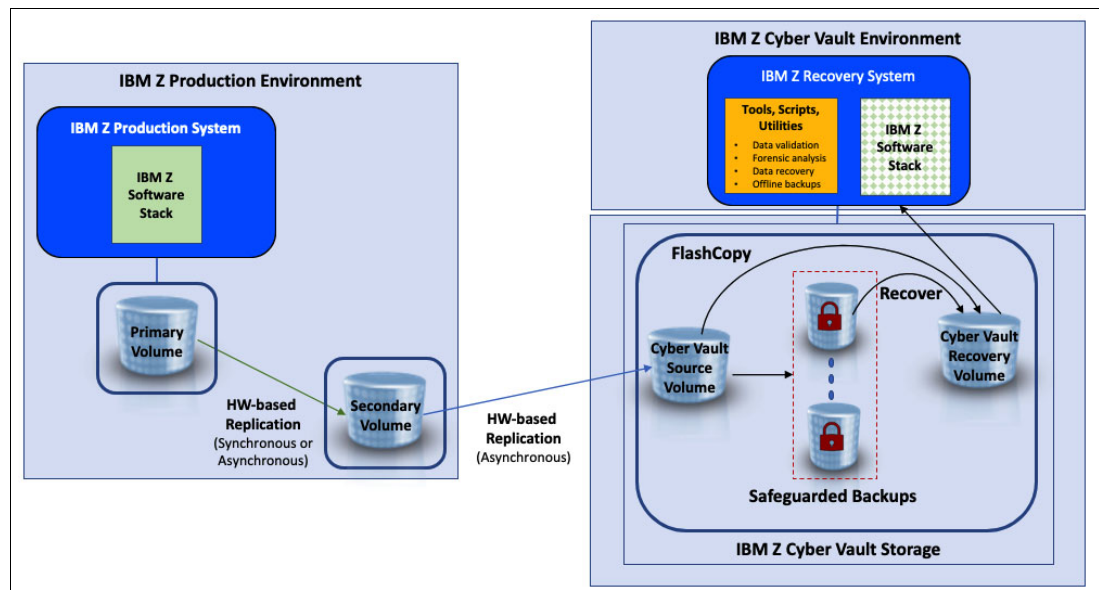


Figure 2-3 IBM Z Cyber Vault solution reference architecture (physically isolated)

The immutable copies (Safeguarded Copy backups) are in an isolated storage environment and inaccessible from all systems. Using IBM Z Cyber Vault automation, the Safeguarded Copy backups are recovered from a combination of the IBM Z Cyber Vault source volume and the immutable backups onto an isolated pool of IBM Z Cyber Vault recovery volumes. The recovery systems perform an IPL as needed to support the relevant IBM Z Cyber Vault solution capabilities.

Figure 2-3 on page 25 contrasts the physical isolation depiction with a virtual isolation topology, as shown in Figure 2-4.

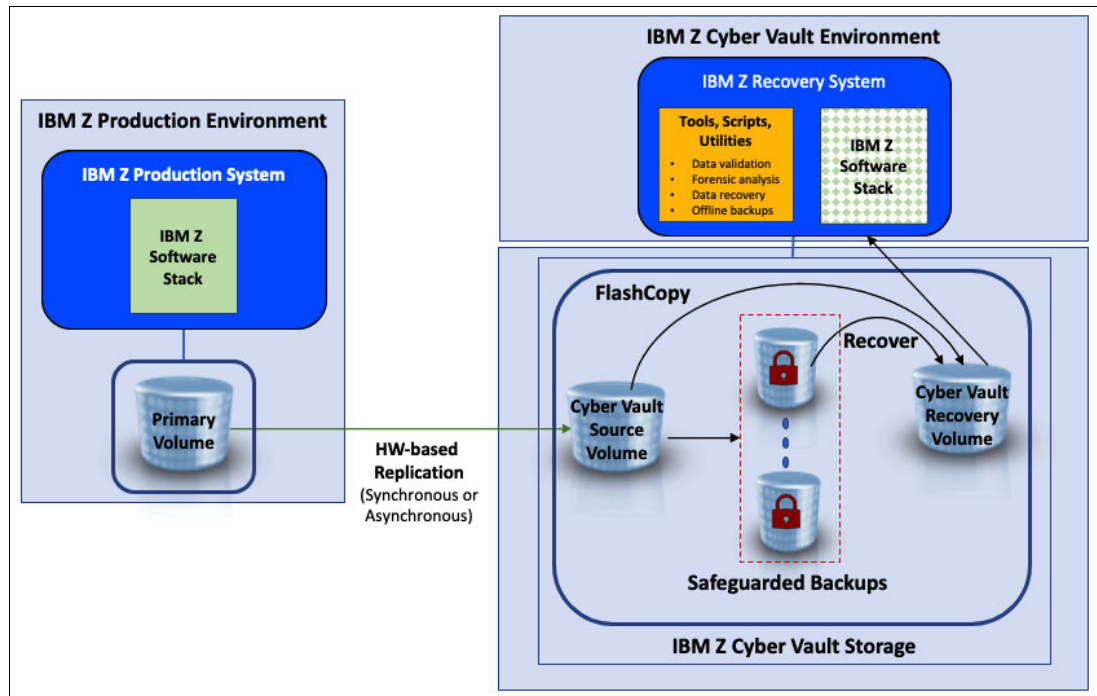


Figure 2-4 IBM Z Cyber Vault solution reference architecture (virtually isolated)

The key difference in the virtual isolation topology is that the IBM Z Cyber Vault source volume is one of the primary volumes that is created by the replication technology. The captured immutable copies are inaccessible from production systems and must be recovered to the recovery volume that is accessible to the IBM Z Cyber Vault environment before they can be validated by the recovery systems.

One characteristic to consider is the mechanism for capturing a consistent, point-in-time copy of the production data. When capturing production data, the I/O traffic to the primary volume must be paused across all volumes in the consistency group (CG).<sup>1</sup> Physical isolation pauses replication to the IBM Z Cyber Vault Storage environment in a consistent state, and this action does not directly impact production I/O traffic. However, with virtual isolation, such as when production data is copied directly from the primary volume where synchronous replication (Metro Mirror) is in use, I/O traffic to all volumes must be paused while the point-in-time consistent backup is taken. Because the impact is similar to a consistent FlashCopy operation across all devices in production, you might use this behavior to determine the expected impact.

A second consideration is where to provision capacity in the IBM Z Cyber Vault environment for the validation process. You can provision this capacity in the same site as the production environment or in a remote site. The recovery system requires connectivity to the IBM Z Cyber Vault recovery volumes in the IBM Z Cyber Vault environment.

<sup>1</sup> A CG can manage the independent relationships for all volumes that are related to an application, helping ensure consistent and synchronized data.

IBM Z production environments vary in the types of applications that are deployed, the methods by which the data is stored and backed up, and the usage of data mirroring for BC. The production environment might consist of a single IBM Z system with a few production LPARs or multiple IBM Z systems running many production LPARs in a single data center or multiple data centers.

Another characteristic to consider is the separation from a management environment perspective. With virtual isolation that uses GDPS LCP Manager, the control point for operations such as taking a backup, recovering a backup, or releasing a backup uses the same control point (or GDPS controlling system) that manages day-to-day production data replication.

GDPS LCP Manager supports both virtual isolation or physical isolation of the IBM Z Cyber Vault Storage environment. Depending on the selected isolation method, the implementation approach is different. For more information, contact [GDPS@us.ibm.com](mailto:GDPS@us.ibm.com).

With physical isolation, a separate control point is used, which enables another degree of separation from the production environment and supports the usage of a different security database with rules that are tailored for the IBM Z Cyber Vault environment. This consideration borders on the interface between cyber resilience, which is where the IBM Z Cyber Vault solution fits, and cybersecurity, as described in 1.2.2, “Cybersecurity versus cyber resiliency” on page 6.

In addition to the IBM Z Cyber Vault Storage environment, teams must determine whether physical or virtual isolation best meets the requirements, and they must consider several factors for the IBM Z Cyber Vault environment. Teams must also decide whether to use dedicated IBM Z hardware for the IBM Z Cyber Vault environment. Both options are technically viable. However, this decision affects the current configuration and software licensing. Organizations must discuss options with the IBM team and any vendors that supply the software products.

## 2.3 IBM Z Cyber Vault Storage considerations

Secure data storage is the cornerstone of the IBM Z Cyber Vault solution, including disk storage and virtual tape. This section outlines the concepts of the key IBM DS8000 and IBM TS7700 functions that are relevant to the IBM Z Cyber Vault solution. It also describes the capacity, performance, operational considerations, and the prerequisites for those functions.

The DS8000 Safeguarded Copy function is the foundation of the IBM Z Cyber Vault solution because it provides the crash-consistent point-in-time copy that is required for effective recovery of a z/OS production environment.

The use of virtual tape, such as the TS7700 Virtual Tape Servers, includes functions such as DFSMSHsm offload of data from the primary disk to tape. This situation is common for traditional backup and restore processing and is typically done when there is a need for longer-term retention of data past the disk retention period. For the IBM Z Cyber Vault solution, virtual tape can be used to back up the recovery volumes that must be retained longer than the Safeguarded Copy backup that is held by the DS8000.

## 2.3.1 DS8000 Safeguarded Copy concepts

Safeguarded Copy provides immutable point-in-time copies that are not accessible to systems or applications because they do not have LCUs or volume IDs that are associated with them. Safeguarded Copy source volumes and their backups cannot be erased or deleted through the DS8000 native interfaces (DS Command-line Interface (DS CLI) or DS Graphical User Interface (DS GUI)). To access a Safeguarded Copy backup, a recovery action to an IBM Z Cyber Vault recovery volume is necessary so that the data can be accessed from an IBM Z Cyber Vault recovery system (see Figure 2-5).

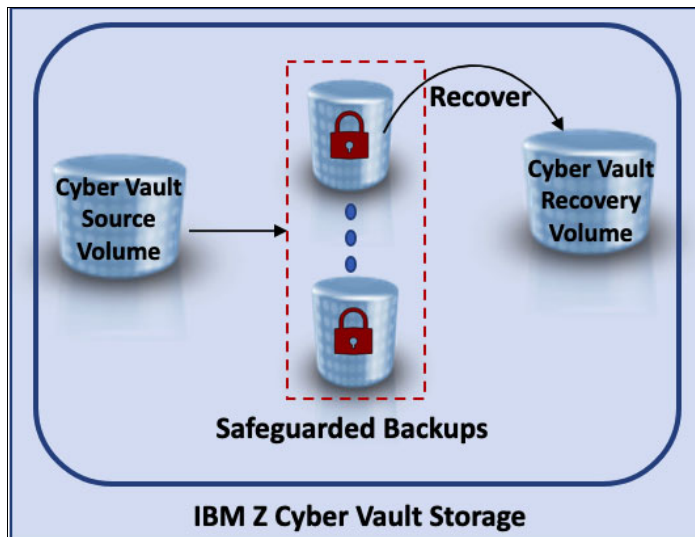


Figure 2-5 IBM Z Cyber Vault Storage

Safeguarded Copy supports up to 500 crash-consistent point-in-time copies per volume. To optimize capacity usage, the Safeguarded Copy backups depend on each other, and changed tracks are stored only once per backup in the Safeguarded Copy backup capacity. The Safeguarded Copy backup capacity is similar to an undo log.<sup>2</sup> The Safeguarded Copy capability integrates with different high availability (HA) or disaster recovery (DR) configurations.

Either GDPS LCP Manager or CSM is required to manage Safeguarded Copy backups. For more information about management software options, see 2.2.3, "IBM Z Cyber Vault environment" on page 23.

Before you can use Safeguarded Copy, you must specify the amount of space for each volume that you want to use for backups in the DS8000. The required capacity depends on the data change rate, the number of backups (frequency), the amount of time you want to keep the backups (retention period) and the threat (such as ransomware attack) you want to protect your data against. Therefore, a Backup Capacity Multiplier per volume for Safeguarded Copy needs to be defined.

To store the data changes, you must have storage capacity in the DS8000. Without a Safeguarded Copy backup, the capacity is pure virtual capacity; physical capacity is allocated as you capture or create backups; and data that is overwritten in the Safeguarded Copy source volume is saved in backups. Backup data is saved in contiguous tracks, which leads to better efficiency than with FlashCopy, which can allocate an entire extent for a single changed track.

<sup>2</sup> When a change is made to data, the system logs an undo record that describes the change so that a transaction can be rolled back if needed, which reverses the changes.

To recover previous Safeguarded Copy backups to the IBM Z Cyber Vault recovery volumes, they must be configured in the DS8000. These recovery volumes have a one-to-one relationship with Safeguarded Copy source volumes. They must have the same volume size and have the same DS8000 system affinity as your Safeguarded Copy source volumes. Usually, they are defined as thin-provisioned volumes and allocate only physical capacity during a recovery action.

Once Safeguarded Copy backup capacity is defined, you must have physical storage capacity in the DS8000 to store the data changes. Without a Safeguarded Copy backup, the backup capacity is pure virtual capacity; physical capacity is allocated as you use your management software (GDPS LCP Manager or CSM) to set up and manage Safeguarded Copy. The management software is used to capture, recover, restore, and expire Safeguarded Copy backups and contains the policy to control them.

## Creating and capturing Safeguarded Copy backups

When you create and capture a Safeguarded Copy backup with your management software, a CG is created across all involved volumes and DS8000 systems. The DS8000 system sets up metadata and bitmaps to track updates to the Safeguarded Copy source volume. After the backup is set up, the storage system continually copy data that was overwritten by host I/O in the Safeguarded Copy source volume to the CG log within the Safeguarded Copy backup capacity.

When you capture the next backup, the DS8000 system closes the previous backup and creates one. Therefore, it does not have to maintain each backup individually.

To minimize the impact during Safeguarded Copy backup creation, the process consists of three steps:

### 1. Reservation

The DS8000 sets up the required bitmap and prepares the metadata in the Safeguarded Copy backup capacity. It also ensures that all changed data from the previous backup is stored. After all preparations are done, the CG formation can take place.

### 2. Check in

To create a consistent backup, the DS8000 system must stop updates for all volumes within the session for a short period. It does this task by setting an Extended Long Busy (ELB) state and then performing the final tasks to create the backup. When the check-in is complete for all volumes and all volumes in ELB, the backup is consistent.

### 3. Check out or completion

The DS8000 lifts the ELB and write operations can continue. From now on, the DS8000 system writes further backup data into the CG logs of the new backup.

The management software (either GDPS LCP Manager or CSM) coordinates and performs these steps automatically and with minimal impact to the host operations. However, you should still consider the ELB time for your host write operations to determine the impact to your existing workload or your Recovery Point Objective (RPO).

## Recovering Safeguarded Copy backups

To access a Safeguarded Copy backup, use a recover action. Recovering Safeguarded Copy backups requires recovery volumes that you must specify while establishing the Safeguarded Copy environment. The recovery volumes must have the same capacity as the Safeguarded Copy source volumes, and they can be thin-provisioned. You can perform the recover action with background copy or **NOCOPY**. Typically, you specify **NOCOPY** if you need the recovered data only for a limited period, and you specify **COPY** if you intend to use it for longer time.

To recover a specific backup, the DS8000 system needs all backups that are newer than the one that will be recovered. During the recover action, the DS8000 system creates a relationship between the Safeguarded Copy source volume and the recovery volume and then the DS8000 system creates a recovery bitmap that indicates all the data that changed since the backup that should be recovered and must be referenced from the CG logs of the newer backups rather than from the Safeguarded Copy source volume.

Now, you have read/write access to the recovery volumes that contain pointers to the data of the selected backup. If the recovery system reads data from the recovery volume, the DS8000 examines the recovery bitmap and decides whether it must fetch the requested data from the source volume or from one of the Safeguarded Copy backups.

**Note:** The recovery action provides access to only a selected backup on recovery volumes, which allows you to do data validation and forensic analysis. Then, the data on the recovery volumes can be used to restore that data to the production volumes.

### Expiring Safeguarded Copy backups

With management software, you can expire Safeguarded Copy backups manually or automatically after the retention period for those backups ends. Because the backups depend on each other, expiring a backup also expires all older backups.

With the default setting, the DS8000 system forces roll-offs of the oldest backups on a volume basis if the following situations occur:

- ▶ The system reaches 500 Safeguarded Copy backups per volume.
- ▶ A DS8000 storage pool runs out of physical space.
- ▶ The specified Safeguarded Copy backup capacity (backup multiplier) for a particular volume is too small.

These DS8000 mechanisms help ensure that all host I/O requests can be fulfilled, and that no production impact happens because of Safeguarded Copy.

**Note:** The default behavior for roll-off backups can be modified in a DS8900F system with release 9.3.2 and later. For more information, see *IBM Storage DS8000 Safeguarded Copy: Updated for DS8000 Release 10.1*, REDP-5506.

### Safeguarded Copy backup restore- to-production

If there is a wide-spread corruption event, you might need to perform a catastrophic recovery. If you determine that you must restore all of your system and application data back to a consistent point-in-time, use the restore-to-production function.

**Note:** The restore-to-production function is for catastrophic recovery only. For a partial (surgical) recovery, use other mechanisms. For example, you might be able to define shared volumes if your recovery environment is close enough to be accessed directly. Alternatively, you might use a GC session to move the data back to the production environment.

Both GDPS LCP Manager and CSM support the restore-to-production function. During this process, the Safeguarded Copy backup that was recovered to recovery volumes is restored to the production volumes.

The production volumes are on a remote DS8000 system that has a replication relationship (Metro Mirror, GM, or GC) with the DS8000 system that is running Safeguarded Copy. This process is doing an incremental copy of a selected backup from recovery volumes to the production volumes in the remote DS8000 system. With that process, all production volumes contain the data from the point-in-time of the selected Safeguarded Copy backup.

The process is managed and controlled with the management software.

To do this restore-to-production process, complete the following steps:

1. Recover the selected backup to the recovery volumes by using your management software.
2. Validate data on the recovery volumes.
3. Perform forensic analysis to check whether a catastrophic recovery is required.
4. If you have not done so, stop production and the applications.
5. Suspend the replication relationship.
6. Make sure that a replication relationship does not exist on the recovery volumes.
7. Start the restore-to-production by using your management software.
8. After all data is copied by using the incremental copy process, validate data on the production system.
9. Reestablish the suspended replication relationship.
10. Restart the production environment.

For more information about the restore-to-productions process, see your management software documentation and *IBM Storage DS8000 Safeguarded Copy: Updated for DS8000 Release 10.1*, REDP-5506.

### 2.3.2 Safeguarded Copy prerequisites

The Safeguarded Copy function is integrated into the IBM DS8000 storage system models with microcode release 8.5 or later.

To start using Safeguarded Copy, you must have a Copy Services (CS) license that is installed on the DS8000 system. The CS licenses bundle is based on usable capacity and on actual usage. For example, if you must protect 200 TB of your production data with Safeguarded Copy, then 200 TB of DS8000 CS license is required.

For managing Safeguarded Copy, either a fully licensed CSM 6.2.3.1 or later or GDPS LCP Manager 4.2 SP2 or later is required. You cannot use the DS8000 interfaces DS GUI or DS CLI or a z/OS interface to capture, recover, expire, or restore a Safeguarded Copy backup.

Consider the following items when planning for a Safeguarded Copy implementation:

- ▶ Safeguarded Copy operates at the volume level.
- ▶ The DS8000 system maintains a maximum number of 500 backups per volume.
- ▶ If you intend to use a backup frequency of less than 10 minutes, you must submit a Request Price Quotation (RPQ) to IBM for approval and support.
- ▶ The maximum Safeguarded Copy backup capacity for a volume is 14.6 TiB for CKD volumes.

- ▶ A Safeguarded Copy source can be a FlashCopy target if DS8900F microcode release 9.3 or later is installed.
- ▶ The source volume and recovery volume must be managed by the same DS8000 internal server. Therefore, they must both be either in an even or odd logical subsystem (LSS).
- ▶ Only a single Safeguarded Copy backup for a volume can be recovered concurrently.
- ▶ During a Safeguarded Copy recovery action (with the **NOCOPY** option), a cascaded FlashCopy from the recovery volume to another volume is not possible.
- ▶ DS8000 Dynamic Volume Expansion (DVE) is not supported for Safeguarded Copy source volumes.
- ▶ Space Release for a volume that is in a Safeguarded Copy relationship is supported by DS8900F microcode release 9.1 or later.

For more information about Safeguarded Copy planning considerations and how to implement Safeguarded Copy with CSM, see *IBM Storage DS8000 Safeguarded Copy: Updated for DS8000 Release 10.1*, REDP-5506.

In addition to the previous hardware and software requirements, extra physical storage capacity in the DS8000 system is required for the following items:

- ▶ The changed data that is stored in Safeguarded Copy backup capacity over the retention period
- ▶ The small Safeguarded Copy impact for each backup
- ▶ Recovery volumes
- ▶ Safeguarded Copy source volumes (for physical isolation)
- ▶ Provisioning of GM FC Journal volumes if GM is used in a physical isolated solution

### 2.3.3 Safeguarded Copy capacity sizing considerations

During the implementation of an IBM Z Cyber Vault environment, consider how often you create Safeguarded Copy backups and how long you keep these backups in the DS8000. These considerations might depend on regulatory or business requirements. How often you validate data in your IBM Z Cyber Vault environment, and how long that validation takes can have implications on the physical capacity that is required for recovery volumes.

A higher backup (capture) frequency can reduce data loss. A longer retention provides more recovery options. However, the Safeguarded Copy backup frequency combined with the backup retention period of your backups and the data change rate are the key factors that influence how much capacity that you need to store the backups. You also need capacity for the recovery volumes on which you are doing the data validation.

This section describes the capacity sizing for Safeguarded Copy.

#### Safeguarded Copy sizing overview

It is crucial to do an accurate Safeguarded Copy capacity sizing. It is a best practice to use small extents and thin-provisioned volumes in a DS8000 system, and the DS8000 physical and virtual capacity limits should not be reached.

For sizing of a Safeguarded Copy solution, the following actions are required:

- ▶ Understand the topology, whether it is virtual or physical isolation.
- ▶ Determine the requirements for backup retention and frequency.
- ▶ Understand how the recovery volumes are used in different use cases.

- ▶ Size the Safeguarded Copy recovery and source volumes physical and virtual capacity.
- ▶ Size the Safeguarded Copy backup physical and virtual capacity.
- ▶ Model the performance of the new or upgraded storage systems.

The capacity limits of a DS8000 system depend mainly on the cache size of the system, The limits for DS8000 systems with less than a 1 TB cache is only about 25% of the maximum capacity limit. For more information about capacity limits, see the DS8000 product documentation for your DS8000 model.

Estimate the physical and virtual capacity of the following components in the DS8000 storage system:

- ▶ Safeguarded Copy backup capacity
- ▶ Recovery volume
- ▶ Safeguarded Copy source volume if a physical isolation approach is implemented

Physical capacity estimation is necessary to determine how much capacity is required to store all changed data within the Safeguarded Copy backup capacity. The virtual capacity limit of the DS8000 system is based on its cache size, so to determine whether that limit will be exceeded, the virtual capacity for all volumes within the DS8000 must be estimated. For each Safeguarded Copy source volume, you must calculate the required Safeguarded Copy virtual capacity to estimate the Backup Capacity Multiplier.

Both the required Safeguarded Copy backup capacity and the Safeguarded Copy virtual capacity depend on the data change rate and the following backup management policies:

- ▶ Frequency of Safeguarded Copy backups that are taken
- ▶ Retention period for the backups

### **Frequency of Safeguarded Copy backups**

When you have a high frequency of Safeguarded Copy captures, you have more recovery points and potentially reduce the amount of data that is lost as a result of the corruption. You might be able to recover to a time closer to the point of corruption. However, you might not always be able to restore to the last backup. The recovery point for a particular event might be several captures older than the latest Safeguarded Copy backup. More frequent captures are often required and preferred but can impact the capacity that is required and make it more challenging to find the best copy.

There are also implications based on the source of your backups. When backups are taken in a Metro Mirror environment, you must “freeze” all write I/O to the volumes being backed up. A higher backup frequency results in a freeze frequently impacting production. However, if the Safeguarded Copy backups are taken on a GM DR DS8000 system or on an isolated third or fourth site, such a freeze does not impact production, which enables more frequent backups.

The DS8000 supports a frequency of every 10 minutes. However, 4 - 6 hours is the average frequency with some clients opting for more frequent copies to minimize data loss and some doing only one or two a day.

### **Retention period of Safeguarded Copy backups**

In addition to the backup frequency, you must decide how long you want to keep Safeguarded Copy backups. The longer your retention period is, the more capacity is required to store the changes in the DS8000 system.

The following considerations help you define the best retention period to meet your needs:

- ▶ Do you have regulatory or business requirements that define how long you must keep the backups?
- ▶ Is it helpful to restore a backup that is 14 days old? Is it acceptable for your business?
- ▶ How long does it take to detect that logical corruption occurred?

Today, the most common retention period is 7 - 14 days. Copying validated data from the DS8000 system to logical WORM (LWORM) tapes is a good option for long-term retention on reduced cost media, which also reduces the amount of capacity that is required in the DS8000 system.

Define the retention period based on your requirements. Ensure that you do Safeguarded Copy backup capacity sizing to reflect your retention period and backup frequency requirements.

### **Safeguarded Copy backup capacity sizing**

As an example, if you are creating a backup every 6 hours (four backups per day) and retaining it for 7 days, you must know the data change rate of 28 backup intervals over a 7-day retention period to estimate the Safeguarded Copy backup capacity.

In addition to the Safeguarded Copy backup capacity, the Safeguarded Copy recovery volumes must be sized too. The required physical capacity for recovery volumes depends on how long you intend to keep the recovery volume copy relationship active, and how much the Safeguarded Copy source volumes and recovery volumes change while the relationship exists.

**Best practice:** Provision about 20% of the physical source volume capacity for the recovery volumes so that the volumes can be used for data validations during normal operations.

The sizing for the Safeguarded Copy backup capacity (physical and virtual) and recovery volumes can be done by using the methods that are described in this section. The methods determine the data change or destage rate in tracks. Then, this absolute number is used for converting to actual GiB (or TiB) capacity.

To calculate the required capacity based on the data change rate or destage rate, use a sliding sum approach to estimate the peak capacity. Add the data change rate or destage rate in GiB (or TiB) per backup interval for as many intervals as the length of the retention period. Do this task for each DS8000 system to calculate the required physical capacity for the Safeguarded Copy backups.

Do the same for each Safeguarded Copy source volume to estimate the Backup Capacity Multiplier if you cannot use the simple approach of using the number of backups in a retention period as the Backup Capacity Multiplier for each volume.

**Tip:** Use the simple approach to determine the Backup Capacity Multiplier per volume only if a few Safeguarded Copy backups are required and the source volume capacity is small. Only then will you not reach a DS8000 capacity limit.

Different methods are available to do a capacity sizing for Safeguarded Copy backup. The most common methods are the following ones:

- ▶ Analyzing the DS8000 Write Monitoring Bitmap, which is the preferred method for existing DS8880 and DS8900F systems.

For more information about this sizing method by using the CSM ESESizer session, see [DS8000 Safeguarded Copy and Extent Space Efficient \(ESE\) FlashCopy capacity sizing by using the new CSM ESESizer or Safeguarded Copiesizer functions](#).

- ▶ Analyzing performance data, such as Resource Measurement Facility (RMF) data in z/OS or IBM Storage Insights.
- ▶ IBM Storage pre-sales and IBM Business Partners can use IBM Storage Modeller to do a sizing based on RMF data.

Estimating the virtual capacity and assigning a correct Backup Capacity Multiplier is important. It is possible to run out of virtual capacity for volumes or reach the DS8000 virtual capacity limit. To avoid this outcome, a best practice is to assign the same Backup Capacity Multiplier to all volumes that belong to the same z/OS DFSMS Storage Group after you estimate the virtual capacity per volume. Assign the highest Backup Capacity Multiplier estimate for a DFSMS Storage Group. To simplify storage management, reduce the number of different Backup Capacity Multiplier to three or four per environment. This approach increases the required virtual capacity, but it avoids losing a Safeguarded Copy backup.

Performance data provides only an estimate of the capacity for the Safeguarded Copy backups. This method might overestimate the Safeguarded Capacity because it does not consider tracks that are destaged multiple times within one Safeguarded Copy backup interval. Therefore, this method tends to be most accurate for configurations where there is a shorter period between backups.

For more information, see *IBM Storage DS8000 Safeguarded Copy: Updated for DS8000 Release 10.1*, REDP-5506.

A best practice is to gather data change rate information during a peak workload period, for example, in banking environments a month-end or quarter-end period is usually a peak period. Gather data for at least a full week. If you plan a longer retention period, gather data for a whole retention period. In addition, make sure that things like *DB reorgs* and maybe burst FlashCopy operations are included during the period of data gathering.

## Extra storage for the IBM Z Cyber Vault solution

Extra storage capacity is required for the IBM Z Cyber Vault environment beyond the backup capacity and recovery volume capacity. For example, storage volumes are needed for surgical recovery. These extra volumes are called *staging volumes*.

During a surgical recovery action, you copy the data that is needed for recovery from the recovery volume to the staging volumes. In a surgical recovery, the restoration to the production environment happens mainly from these staging volumes by either bringing the volumes online directly in your production environment (if the staging volumes are in a Metro Mirror secondary (virtually isolated)) or by using the DS8000 Global Copy function to copy the data to another set of staging volumes in your production environment.

You must have some persistent volumes in your IBM Z Cyber Vault environment to store reports and other historical datasets that are created by the data validation process.

An example of an IBM Z Cyber Vault Storage architecture is shown in Figure 2-6.

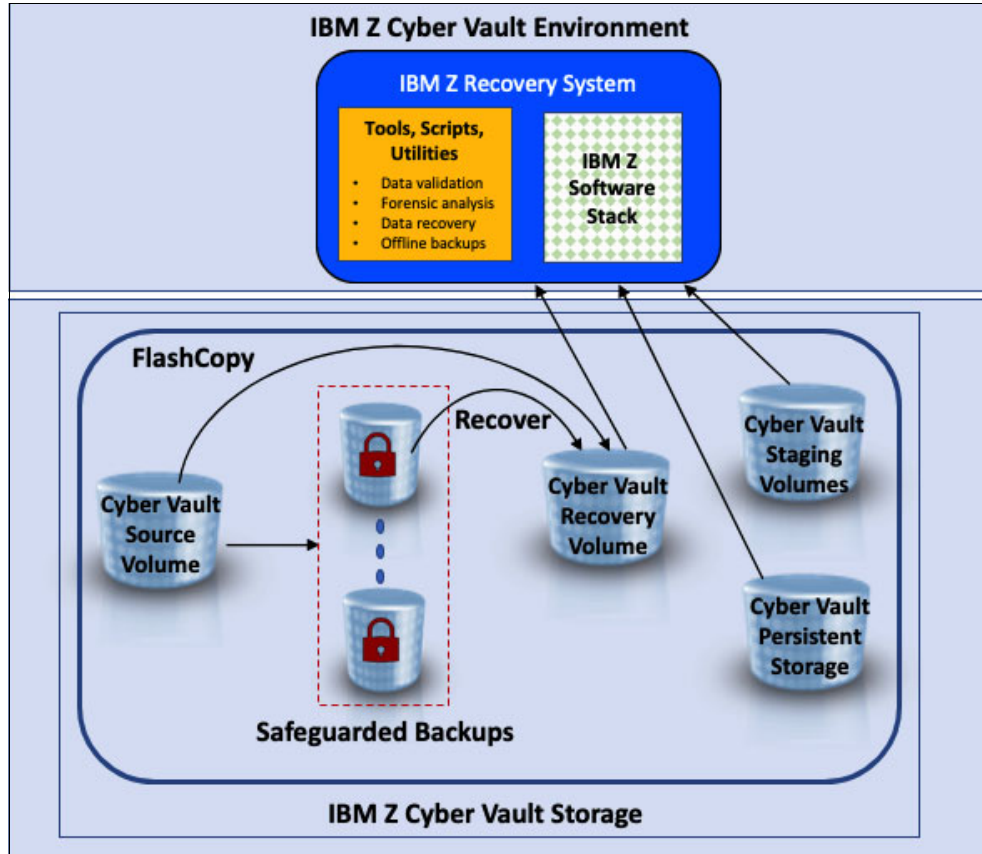


Figure 2-6 IBM Z Cyber Vault Storage architecture

The IBM Z Cyber Vault recovery volume is used for the data validation of a recovered Safeguarded Copy backup.

In addition, a smaller set of volumes that are called staging volumes are used for surgical recovery, and a few persistent volumes are used to store reports and other historical files. *Persistent* means that these volumes must not be lost between different validation runs, and they can be System Management Facility (SMF) records or historical information that are used to identify the last validated copy before corruption.

You can also use a second set of recovery volumes so that regular tasks such as data validation, which are performed at the same time as recovering a Safeguarded Copy backup for surgical recovery or forensic analysis, can be completed.

### 2.3.4 DS8000 performance considerations for Safeguarded Copy backups

In addition to Safeguarded Copy backup capacity sizing, another important aspect is to model the performance of the new or upgraded IBM DS8000 systems by contacting your IBM representative or IBM Business Partner.

There are many different workloads running on the storage systems that are used by your IBM Z Cyber Vault environment.

- ▶ The production write workload that updates the Safeguarded Copy source volumes
- ▶ The saving of data into the Safeguarded Copy backups

- ▶ FlashCopy activity for the recovery copy (RC) that is used for data validation
- ▶ FlashCopy activity for the GM journals (if you are using GM for replication into IBM Z Cyber Vault Storage)
- ▶ The data validation workload, which is typically a read-intensive workload on the recovery volume

Safeguarded Copy and FlashCopy can each be considered, in the worst case, to generate an extra back-end read/write I/O for every write to the Safeguarded Copy source volume. For sizing purposes, teams might consider the worst-case scenario.

When running a Safeguarded Copy backups on a GM secondary volume, consider using FlashCopy with GM. By using Remote Pair FlashCopy (RPFC) in the Metro Mirror environment, this function can help keep replication relationships in full duplex even though a FlashCopy operation is done on a Metro Mirror primary volume.

However, if RPFC is used in a Multi-Target Metro Mirror or in a Metro Global Mirror (MGM) environment, the data must be copied by using the PPRC link to either the second Metro Mirror leg or the GM leg, which means that there is a burst write activity with a 100% data change rate for all the involved volumes or datasets. During this time, the second Metro Mirror leg is in copy pending, and a GM leg might be unable to create a GM CG.

If you use RPFC for many volumes, you must coordinate RPFC and Safeguarded Copy capture. Otherwise, you might not be able to create a Safeguarded Copy backup because the GM RPO is too high to successfully pause GM.

A best practice approach is to do all RPFC at the same time and create a Safeguarded Copy backup after a second Metro Mirror leg is in full duplex or after GM created a CG.

Safeguarded Copy backup performance and capacity sizing is a crucial part of a Safeguarded Copy backup implementation. For the Safeguarded Copy backup sizing, contact your IBM Storage Technical Sales or your IBM Business Partner for support.

### 2.3.5 DS8000 Safeguarded Copy backup operational considerations

For Safeguarded Copy environments, you have more operational considerations. This section describes operational changes like the impact of capturing a Safeguarded Copy backup, adding and removing volumes, and changing the Safeguarded Copy Volume Backup Multiplier when adding volumes to an Storage Management Subsystem Storage Group. It also describes monitoring and alerting, and DS8000 security considerations.

#### **The impact of capturing Safeguarded Copy backups**

A DS8000 Safeguarded Copy backup is a protected, crash-consistent point-in-time copy. To create consistency, the DS8000 pauses all updates to volumes within the CG for a short period by presenting an ELB state. If multiple DS8000 systems are involved in a Safeguarded Copy CG, then the management software GDPS LCP Manager or CSM helps ensure that the backup is crash-consistent across all volumes. The fewer volumes, the smaller the impact is.

In a Metro Mirror environment where either Safeguarded Copy is running on a Metro Mirror primary or a Metro Mirror secondary, the z/OS LPARs notice this short pause, which might be 100 - 200 ms for an environment with only a few thousand volumes up to perhaps 1 - 2 seconds for environments with 10,000 or more volumes.

For a CSM-managed environment, ensure that you implemented the z/OS IOS enhancement APAR OA59561 and established a secure IP connection to the z/OS IBM HyperSwap® Address space to reduce the freeze impact. For more information, see *IBM Storage DS8000 Safeguarded Copy: Updated for DS8000 Release 10.1*, REDP-5506. The need to avoid an ELB in production is one reason clients implement a physically isolated solution with a GM relationship to a separated system that is running Safeguarded Copy.

In a GM environment where Safeguarded Copy is running on a GM secondary, pause the GM with consistency before capturing or creating a Safeguarded Copy backup. This action creates a consistent Safeguarded Copy backup from the GM secondary. This GM pause increases the GM RPO. The GM pause and creating a Safeguarded Copy backup can be automated by the management software GDPS LCP Manager.

## Monitoring and alerting considerations

Safeguarded Copy protects your data, so it is important that your Safeguarded Copy environment works as expected, and that all backups can be kept until the retention period is over. Therefore, monitoring a Safeguarded Copy environment is another important aspect. The DS8000 and the management software (GDPS LCP Manager or CSM) provide messages that are related to Safeguarded Copy that can be used for monitoring, automation, and alerting.

It is critical to monitor DS8000 capacity. No matter how careful you are with Safeguarded Copy Backup Capacity sizing, you might still exhaust your storage pools or the virtual capacity. Running out of capacity might lead to losing Safeguarded Copy backups or impacting production or the GM relationship.

The DS8000 sends out a warning message if a storage pool physical capacity reaches a certain threshold or if the Safeguarded Copy backup virtual capacity of a volume exhausts or reaches the warning level. These messages should be monitored and alerts should be generated to allow immediate action.

The DS8000 provides different interfaces to monitor these events. It uses the DS8000 Event log or DS8000 SNMP traps, or it sends messages to z/OS or to a syslog server. The important z/OS messages are as follows.

- ▶ IEC817I provides a warning for a virtual capacity out-of-space situation.
- ▶ IEA499E provides a warning that is related to the physical capacity of the storage pools.

GDPS LCP Manager also provides a message that enables monitoring and alerting.

It is a best practice to establish a capacity alert depending on your capabilities and preferences. In addition, you can establish automation to increase the Safeguarded Copy Backup Capacity if a certain usage is reached for a volume. For that purpose, you can use the DS8000 REST API.

## Changing the volume backup multipliers

During the lifetime of a Safeguarded Copy environment, there might be circumstances that require changes to Volume Backup Multipliers, for example, adding a volume to the Storage Management Subsystem Storage Group or because the workload changed. Requirements might change over time and a higher frequency or retention period might be required.

The DS8000 provides dynamic expansion of the Safeguarded Copy backup capacity. You can increase the volume backup multiplier in the DS8000 by using the DS CLI or the DS GUI. However, an expansion might be delayed until the retention period ends or until backups that block an expansion are manually deleted.

Unfortunately, it is not possible to predict whether a backup will block an expansion. If an expansion is not possible, the volume remains in an expanding state for a longer period, and both the CSM Safeguarded Copy session and GDPS LCP Manager indicate which backups are blocking the expansion. At that point, you can decide whether to delete backups manually or wait until the retention period for those backups ends.

If you get a “virtual capacity” warning message like IEC817I, you must react quickly to avoid losing backups and expand the Safeguarded Copy Backup Capacity of the volume that is mentioned in the message.

If your requirements change, and you must increase the backup frequency or the retention period, perform a Safeguarded Copy capacity sizing by using the ESESizer or Safeguarded Copiesizer tool, compare the results with your current configuration, and adjust the volume backup multipliers.

### **DS8000 security aspects that are related to Safeguarded Copy**

The Safeguarded Copy feature provides immutable point-in-time copies to protect your data if there is a logical corruption event. To help ensure that those copies are protected, consider increasing security in a Safeguarded Copy environment. Ensure that a single user cannot modify the policies that control the Safeguarded Copy backup creation and expiration and damage the DS8000 configuration.

Follow these best practices:

- ▶ Implement a separation of duty policy so that a single user does not have access to both the management software GDPS LCP or CSM and the DS8000 system.
- ▶ Restrict DS8000 user access and limit user rights in the management software. There are different functions that are available to restrict access in the DS8000 and in the management software.
- ▶ Use Multi-factor Authentication (MFA) in the DS8000 to improve security for the login process and implement custom user roles that enable a granular definition of user rights. For more information, see [Security](#).

Both GDPS LCP Manager and CSM provide role-based security and dual-control functions for Safeguarded Copy. With role-based security, you can restrict user rights in the management software by assigning specific roles. The dual-control function provides extra protection for certain tasks, such as expiring Safeguarded Copy backups or changing GDPS LCP management profiles or CSM properties. After dual control is enabled, a second user with the same privileges must approve the task before it runs. For more information about GDPS LCP Manager dual control, see the GDPS LCP Manager documentation. For the CSM dual control function, see [Dual control](#).

You might also consider integrating the Safeguarded Copy and IBM Z Cyber Vault environment into a security information and event management (SIEM) solution to detect unauthorized actions in this environment. You can forward audit logs and messages from the DS8000 and the management software into a SIEM solution. Depending on log entries or messages, certain actions, such as sending a notification or taking another Safeguarded Copy backup, can be initiated.

## 2.3.6 TS7700 and virtual tape solution considerations for IBM Z Cyber Vault

Most z/OS environments include the use of the IBM TS7700 or virtual tape solutions from other vendors. The information in this section focuses on the TS7700 and its capabilities, although similar parallels might be available with other vendor solutions. Virtual tape is a widely used technology to provide reduced cost media for typical backup and restore use cases, and it is also key in supporting an *operational tape* function for primary production operations.

Examples of operational tape use might include the following items:

- ▶ DFSMSHsm migrated data
- ▶ Operations, Administration, and Maintenance (OAM) object data
- ▶ Batch input and output datasets
- ▶ Reporting data

Some of this data is no longer stored on the primary storage system, and as a result, becomes the *master data*, and it must be protected beyond normal HADR best practices by using extra protections against compromised environments as part of the overall IBM Z Cyber Vault solution.

There are several security and data gapping capabilities that should be considered for improving the security of a TS7700 environment across all facets of the solution.

- ▶ Expire hold
- ▶ LWORM and LWORM retention
- ▶ Secure Data Transfer
- ▶ Data at rest encryption
- ▶ Event and rsyslog logging
- ▶ Dual admin authorization
- ▶ Physical tape and copy export
- ▶ Cloud tiering and cloud export
- ▶ Selective Device Access Controls (SDAC)

For more information about these features, see *IBM TS7700 Release 6.0 Guide*, SG24-8464.

The use of virtual tape for the IBM Z Cyber Vault solution falls into two primary functions:

- ▶ Providing backups of disk volumes and application data, which is a common capability of using tools like DFSMSdss, ImageCopy, others to create a point-in-time dump of critical data to tape volumes.

For IBM Z Cyber Vault, this function includes backing up the disk volumes that are associated with a point-in-time recovery volume that can be retained for a longer duration than the Safeguarded Copy backup restore depth that is held by the DS8000.

- ▶ Accessing the tape volumes from an IBM Z Cyber Vault environment that might be needed to recall data from tape as part of validation, forensic analysis, or recovery efforts. This function focuses on concepts and considerations for when a z/OS sysplex undergoes an IPL from a recovery volume set, potentially from a considerably older point-in-time, and needs access to tape volumes that are held by the tape environment that is at the “current” view of the tape data and tape catalog state.

Figure 2-7 on page 41 shows the use of tape storage in the IBM Z Cyber Vault solution with the TS7700.

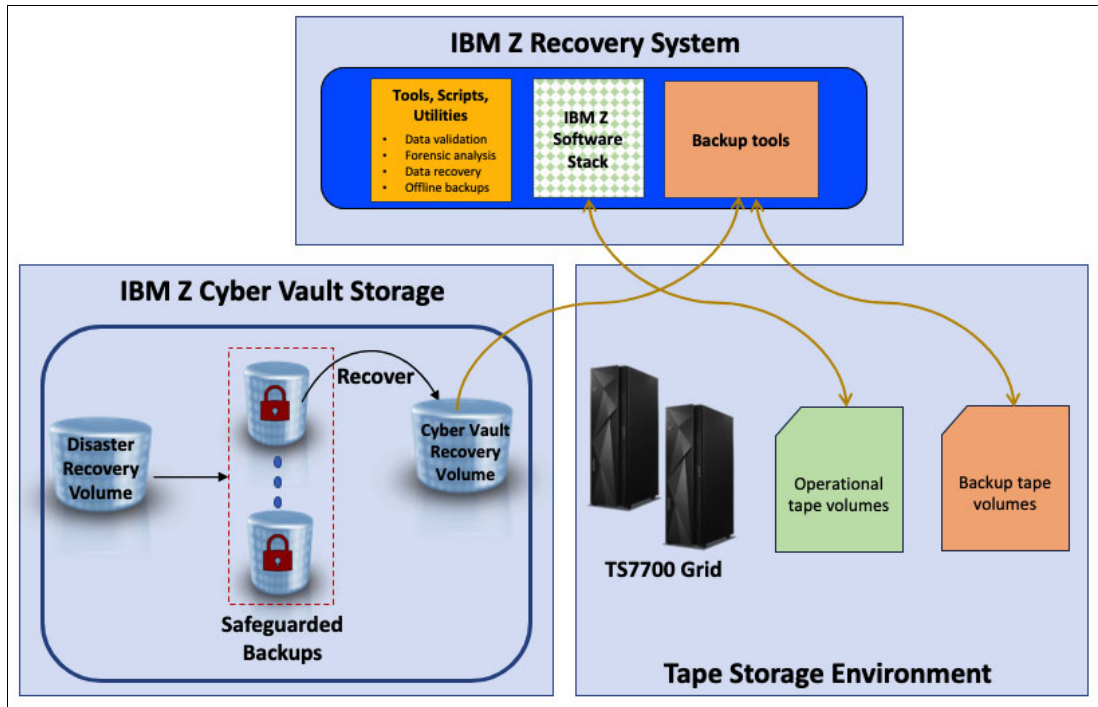


Figure 2-7 Tape storage for the IBM Z Cyber Vault solution

## Protecting data with logical WORM

Although the deletion of tape data is different from disk because tape volumes are moved to a scratch category, which makes them eligible for future use and overwrite rather than causing an immediate impact, a knowledgeable attacker might know how to accelerate this process. These protections must be considered for backup and restore operations, and for operational tape volumes such as DFSMSHsm.

To protect data that is stored within the TS7700 grid, use a combination of LWORM, LWORM Retention, and the Expire Hold function.

LWORM prevents any overwrite of existing data on TS7700 volumes and prevents attempts to move the volume to scratch until the defined LWORM retention period elapses. The LWORM protection and retention duration are enforced by the TS7700 and are established at the write from the beginning of tape (BOT) when the tape volume is first created. These values are ingrained in the definition of the logical tape volume itself and are retained when the logical tape volume is moved to physical tape, cloud tiering, cloud export, or other media types.

**Note:** The LWORM retention duration is defined for a tape volume when the volume is created and does not change for that volume even if the retention period is modified for future volumes. For example, if the LWORM retention is set to 3 months when the tape volume VOL001 is written from BOT, this 3-month value remains with that tape volume regardless of whether the volume is copy-exported, cloud-exported, or whether the retention period is changed for the associated data class. Any new retention period applies to future tape volumes that are written from BOT.

Expire Hold prevents any reuse of volumes that are returned to scratch during the hold period, which creates a defined grace period from the time a tape volume is returned to scratch until the tape environment attempts to reuse that tape volume to overwrite it with new data.

## 2.3.7 TS7700 and virtual tape best practices for IBM Z Cyber Vault

Here are the primary considerations for augmenting this capability:

- ▶ LWORM and LWORM retention

This option prevents a compromised environment from deleting or overwriting data on tape volumes. In a virtual tape environment, there are no additional requirements other than defining the retention policies and data class constructs to instruct the TS7700 and the tape management software to provide this level of protection. However, these tape volumes can still be cataloged, visible, and accessible for read operations to a compromised environment. LWORM protects the data from alteration, but the volumes remain accessible.

- ▶ SDAC controls

These controls partition access to the tape volumes for the dumps to a specific device range that only a dedicated backup host LPAR can access. This approach should be considered in addition to LWORM to isolate the tape catalog and management system to a highly controlled LPAR.

- ▶ Physically isolating the TS7700 cluster or grid provides the ultimate separation of tape dumps to a cluster or grid that is dedicated solely to holding the dumps. This approach uses different hardware, which can have different physical controls, management access, and administrator sets.

### **Accessing operational tape from an IBM Z Cyber Vault environment**

To access a tape grid from an isolated IBM Z Cyber Vault environment, you need physical and addressable access from IBM Z Cyber Vault environment to your tape grid.

Because the IPL of an IBM Z Cyber Vault environment LPAR might be from a Safeguarded Copy backup instance that is several days old, it has a different tape management catalog view of the tape categories, tape data, and related information compared to the current TS7700 database and tape management system in production. Protections against impacts from the IPL of a different system must be put in place to help ensure that the older system does not push changes to the environment. These protections can be applied at the TS7700 hardware level or at the tape management level (RMM, CA1, and others) depending on your operational needs. This requirement is similar to considerations for a normal DR test, with the additional consideration that the tape management system view in use can be from a point back in time.

It is important to help ensure that the older system undergoing an IPL from a backup that was taken days earlier can still locate the tape data on the tape volumes it cataloged, even if those volumes were scratched by the current system. Expire Hold must be used to help ensure that tape volumes that were moved to scratch by the current production view are not overwritten until the oldest possible Safeguarded Copy recovery depth might undergo IPL to a IBM Z Cyber Vault environment.

LWORM and LWORM retention must be considered for the master data instances on tape to further protect against accidental or malicious attempts to overwrite data by using direct IBM FICON® CCW tape commands or by bypassing the tape management software.

There are multiple tradeoffs among the options that are described in this section, depending on the environment, business security or compliance requirements, and the objectives that influence the optimal configuration for that environment. These tradeoffs should be explored with storage subject matter experts (SMEs) to help evaluate the available options.

Tape management environments like DFSMSrmm added capabilities to support the IBM Z Cyber Vault solution and the IPL of an older version of a tape catalog sharing the same grid as the current production.

For more information, see *What is New in DFSMSrmm*, SG24-8529.

## 2.4 IBM Z Cyber Vault environment considerations

This section describes considerations for the design and implementation of your IBM Z Cyber Vault environment.

### 2.4.1 z/OS recovery system configuration

The purpose of the IBM Z Cyber Vault environment is to provide a clean-room system where users can safely perform data corruption activities and related actions. These activities include data validation processes, forensic analysis and recovery tasks, and security strengthening by running more offline backup procedures or running OffSec exercises.

Your planning for z/OS in the IBM Z Cyber Vault environment should consider the following items:

- ▶ As a best practice at the LPAR definition level, the recovery system should mimic the production environment that is being protected in the number of z/OS LPARs, minus the GDPS Kg/Kr systems and cryptographic hardware. This approach reduces complexity in managing the IPL and validation processes in the recovery environment. You need only a single CF LPAR.
- ▶ The size of the LPARs for the recovered systems and CFs depends on the workload and the amount of data that you intend to validate and process. When you size the recovery environment, the workload profile of the IBM Z Cyber Vault environment is likely different from the profile of the production counterparts in terms of I/O activity and CPU capacity when running potentially many thousands of concurrent structure and data validation jobs. Production workloads never run in an IBM Z Cyber Vault environment, so sizing considerations are not related to actual transaction processing in production.
- ▶ Persistent volumes are a best practice for the storage of validation scripts and associated artifacts, and for data that should not be lost between validation runs, such as historical data or SMF records. Consider defining multiple volumes for this purpose with a unique esoteric to avoid hardcoding VOLSERs when creating output files during validation runs and to reduce I/O contention during validation cycles.
- ▶ Review and update operational procedures regularly to accommodate changes in both the production and IBM Z Cyber Vault environments. Processes such as master key rotation or LPAR configuration changes must include the corresponding updates that are necessary in the IBM Z Cyber Vault environment.

- ▶ The software stack in the IBM Z Cyber Vault environment is an exact replica of the production environment. Therefore, any tools, utilities, and procedures that manage, analyze, validate, and restore system components must be identified and installed in production. This requirement includes preparing a special IPL procedure, setting up the required software, activating tools that collect metadata and track activity, and preparing all automation and data validation tools.
- ▶ Specialized configuration items that are required to direct the IPLs of the recovery systems and subsequent processing must be present in the production environment so that they are available in the IBM Z Cyber Vault environment. These items include IBM Z Cyber Vault specific parmlib members, PROCs, automation routines, and Cross-System Coupling Facility (XCF) policy definitions for the Coupling Facility Resource Manager (CFRM) and Workload Manager.

## 2.4.2 Network isolation

The IBM Z Cyber Vault environment should be in a logical bubble with no intersection with the existing production environment. The IBM Z Cyber Vault environment should be in an isolated network to help ensure that the Safeguarded Copy backups cannot be accessed from the production environment or any other system. This isolation is achieved by using dedicated Open Systems Adapter-Express (OSA-Express) features in the recovery system to provide physical isolation. Alternatively, if the IBM Z Cyber Vault environment is virtually isolated, shared OSA-Express features define logical separation by configuring a separate VLAN and IP subnetwork and optional firewalls.

In addition, you can protect both the IBM Z Cyber Vault environment and the production environment by using built-in z/OS Communications Server security features.

- ▶ IP filtering blocks IP traffic that the system does not explicitly permit within its defined IP Filter Policy.
- ▶ IDS protects the system's services against various types of attacks.

From the router network perspective, fence off the IBM Z Cyber Vault environment from the existing production environments with limited access points. The IBM Z Cyber Vault environment should be accessible only through a few predetermined IP addresses, ports, and protocols, and it must be protected by firewalls and routers that require VPN access and encrypted flows to communicate with the IBM Z Cyber Vault environment (z/OS LPARs). The z/OS LPARs in the IBM Z Cyber Vault environment should be accessible only through the network by using Network Address Translation (NAT) IP addresses because the z/OS LPARs are activated by using the same configuration as production, and the network must avoid duplicate IP addresses.

## 2.4.3 z/OS security

As a best practice, the IBM Z Cyber Vault environment should mimic the production environment, including the External Security Managers (ESMs) that are used to secure the z/OS environment, such as IBM RACF, CA ACF2, or CA Top Secret. The security policies that are defined in production should be replicated in the IBM Z Cyber Vault environment without modification.

To perform validation within the IBM Z Cyber Vault environment, extra security access might need to be permitted, for example, to run data validation scripts and utilities within the IBM Z Cyber Vault environment.

- ▶ RACF started task user ID might require “System Special” access in the existing production environment so that the access is propagated into the IBM Z Cyber Vault environment when it is activated by using production volumes.
- ▶ Parmlib methodology activation is a best practice during RACF startup in the IBM Z Cyber Vault environment.
- ▶ The Started Task user ID that activates or owns validation scripts and utilities must be defined in RACF as “Trusted” and is restricted to operations capability within the IBM Z Cyber Vault environment.

## 2.4.4 Database and middleware consideration

Applications use different subsystems and are developed for specific purposes. Therefore, they require distinct validation and recovery procedures, which might rely on utilities that the corresponding subsystems already provide or use extra software tools and utilities that improve these processes.

### Data validation

Application subsystems enable transaction and data management, providing data and application integrity. Their software controls the creation, organization, and modification of data, and access to it. Many structures and processes are associated with data. The structures are the key component of any set of data, and the processes are the interactions that occur when applications access the data.

Data corruption can occur in any of the data structures or in the supporting metadata of the processes that provide application subsystem services, such as log records and backup catalogs. As a result, it is important to perform data validations for all these subsystems to help ensure that their inner workings are not compromised. In the IBM Z Cyber Vault environment, it is a best practice that all applications and databases are activated by system automation during the IPL process to enable comprehensive validation of the database and middleware together with the infrastructure.

For an IBM Z Cyber Vault environment configuration, the best practice database and middleware capabilities include, but are not limited to, the following items.

- ▶ Validate database structures.
- ▶ Monitor data changes through database log reporting.
- ▶ Identify and fix problems with the databases.
- ▶ Undo or roll back changes to data and database structures by using log reports.

Some examples of data structure validation utilities include but are not limited to the following items:

- ▶ Db2 Utilities (CHECK DATA/INDEX, and log analysis)
- ▶ Information Management System (IMS) utilities (Pointer checker)
- ▶ Catalog tools (IBM Tivoli, Integrated Data Cluster Access Method Services (IDCAMS), and ISV products)
- ▶ Virtual Storage Access Method (VSAM) Indexcheck and Datacheck
- ▶ DFSMSHsm and DFSMSrmm tools
- ▶ RACF (IRRUT200) and zSecure-Audit

- ▶ ISV software
- ▶ Custom-built programs

An example of IBM Customer Information Control System (IBM CICS) validation might include issuing CICS commands such as **CEMT** against CICS regions. You can also run the CICS Installation Verification Program “Catalog Manager” if it is available within the CICS regions to validate.

An example of Db2 database validation might include performing **DSN1COPY CHECK** commands against system and application tables.

An example of IMS validation might include submitting Pointer Checker jobs to the IBM Z Cyber Vault environment to validate the IMS databases.

An example of MQ for z/OS validation might include running the Installation Verification Program **CSQ4IVP1**, the Dead Letter Verification Program **CSQUDLH**, or issuing commands to an IBM MQ queue by using the **CSQUTIL** program.

### **Backup and recovery**

Backup and recovery are the most complicated areas of database management. Having the correct resources to perform a recovery is critical. Without them, there is a risk of losing key data. Database backup and recovery tasks vary from recovering a dropped object to rebuilding after a major disaster. Manual recoveries can be error-prone, time-consuming, and resource-intensive.

The backup design helps ensure that the process is repeatable and automated; the time-consistent copy is clean; and the system is operational.

The recovery design must account for surgical and catastrophic recovery. Surgical recovery applies if only a small portion of the production data is corrupted and if consistency between current production data and the restored parts can be reestablished. A catastrophic recovery applies when there is massive corruption to all or most of the data in the environment.

## **2.4.5 Offensive security environment**

An OffSec environment is a space where ethical hackers use techniques to identify and exploit vulnerabilities in systems and networks. The goal is to find security gaps before malicious actors can use them.

Here are some techniques that are used in OffSec:

- ▶ Vulnerability assessment
- ▶ Penetration testing
- ▶ Red teaming
- ▶ Social engineering
- ▶ Exploit development
- ▶ Threat hunting

The IBM Z Cyber Vault environment provides an ideal setup to perform these activities because it is an exact replica of the actual production system that is isolated. Any actions that occur on this system have no consequences outside of it. As a result of these exercises, weaknesses and vulnerabilities might be identified, which can then be remediated in the actual production system.

The benefits of regularly performing OffSec practices include the following items:

- ▶ Reduced risk of cyberattacks
- ▶ Enhanced security posture
- ▶ Improved incident response
- ▶ Increased security awareness
- ▶ Cost savings and efficiency
- ▶ Competitive advantage
- ▶ Regulatory compliance and risk management

## 2.5 IBM Z Cyber Vault automation considerations

There are essentially two levels of automation to consider. Fully automated tasks are ones that can be scheduled or triggered based on policy with no manual interaction. Although fully automating all processes is ideal, it is not always possible. Partial automation includes scripts that can be manually triggered based on decisions that are made during incident management or daily operations. For example, the process of creating images on the recovery volumes and then activating images and initiating validation is fully automated. If a recovery action and IPL must be initiated for forensic analysis, that action requires manual initiation, although the process itself is scripted.

This section describes automation considerations when designing or planning for the implementation of an IBM Z Cyber Vault solution. It describes automation for Safeguarded Copy and related storage processing, recovery system operation, and data validation.

### 2.5.1 Safeguarded Copy and related storage processing

Managing, creating, recovering, and expiring Safeguarded Copy backups requires management software like IBM CSM or GDPS LCP Manager. They coordinate and perform these steps automatically and with minimal impact to the host operations. For more information about these tools, see 2.3, “IBM Z Cyber Vault Storage considerations” on page 27.

GDPS LCP Manager minimizes the risk of errors by reducing manual operations. To simplify the administration of Safeguarded Copy captures, it defines management profiles. A *management profile* describes the management characteristics of the volume captures that are taken.

- ▶ The replication site (RS) where the captures are taken.
- ▶ The CG to capture.
- ▶ Copy sets that are assigned to this management profile.
- ▶ How long a capture should be retained before it expires and becomes eligible for release.
- ▶ How much time must elapse before a new capture can be taken.
- ▶ The maximum permissible elapsed time for the Safeguarded Reservation Scan phase.
- ▶ The maximum permissible elapsed time for the Safeguarded Check In phase.
- ▶ The maximum permissible time that is allowed for a Consistent Group Pause to complete in preparation for capture processing.

Another key capability of GDPS LCP Manager is standard GDPS scripting, which enables automated procedures and actions for an individual capture set or for all capture sets that are managed by a profile with specific statements.

► **CAPTURE**

This statement performs a consistent capture of the RS(*n*) volume set to an FC(*n*) or Safeguarded Copy(*n*) copy set.

► **RELEASE**

This statement performs a release of one or more expired captures from an FC(*n*) or Safeguarded Copy(*n*) copy set.

► **RESTORE**

This statement performs a restore from an FC(*n*) copy set to the RS volume set. When they are restored, it is then possible to access the RS(*n*) volume set for data analysis, extraction, or test purposes. The **RESTORE** operation is not supported for Safeguarded Copy backup sets.

► **RECOVER**

This statement performs a recovery to the RC set from either an FC(*n*) or a Safeguarded Copy(*n*) copy set. When recovered, it is then possible to access the RC(*n*) copy set for data analysis, extraction, or test purposes.

## 2.5.2 Recovery system operation

The *recovery system* performs all validations, forensic analysis, and recovery planning. You create the images on the recovery volumes by using the commands that are described in 2.5.1, “Safeguarded Copy and related storage processing” on page 47, and then activate systems perform an IPL of them. GDPS LCP Manager scripts can initiate these actions. For CSM environments, that automation must be created. For that reason, the following content focuses on the capabilities that are included in GDPS LCP Manager. Understanding the GDPS capabilities provides the background that is necessary to implement manual processes or custom automation where required.

Once the recovery volumes are created, you can run an IPL on the systems and run your validations. GDPS, by using the Base Control Program internal interface (BCPii), can automate the activation and load (perform an IPL of) the IBM Z Cyber Vault environment LPARs. The GDPS script includes the activation of each LPAR in the sysplex or protected environment. The IPL process uses the existing IPL automation and should include all system and subsystem components. All database and middleware subsystems must start so that they go through their normal restart and recovery processing. If this level of IPL automation is not in place, the production IPL processes must be enhanced.

The IPL process for the IBM Z Cyber Vault systems must differ slightly from the production versions and must accommodate changes in the I/O configuration, and changes in the parmlib members that facilitate the validation process. These changes are driven by the LOADPARM that is used to perform an IPL of the IBM Z Cyber Vault systems and by the introduction of a system symbol in IEASYMxx that differentiates between IPLs of the systems in the production and IBM Z Cyber Vault LPARs. The IBM Z Cyber Vault members must be present in the production environment so that they are copied to the recovery volumes that are used to perform an IPL of the IBM Z Cyber Vault environment.

The IPLs of the IBM Z Cyber Vault LPARs should complete without manual action by implementing an auto-reply policy (AUTORxx) to address write-to-operator-with-reply (WTOR), which modifies the IPL startup command lists in parmlib and uses an automation tool, such as IBM Z System Automation.

### 2.5.3 Data validation

The first step in data validation is the IPL processing (Type 1). Data structure validation (Type 2) and data content validation (Type 3) are unique to each environment. More automation scripts run after the standard IPL processing completes.

The complete validation process should include Type 1 (IPL validation), Type 2 (Data Structure validation), and Type 3 (Data Content validation) scripts, and might use REXX, z/OS utilities, and middleware-specific tools to perform validation. The startup of the Type 2 and Type 3 validation processes should be automated and performed as soon as possible after the IPL.

The validation process sends a report of the results to specific email addresses by using SMTP or forwards the report to a monitoring solution.

Once the validation process is complete, the IBM Z Cyber Vault systems can be taken offline, the LPARs deactivated, and the environment prepared for the next validation run.

For more information about data validation and the IBM Z Cyber Vault capabilities, see Chapter 3, “IBM Z Cyber Vault capabilities” on page 51.





## IBM Z Cyber Vault capabilities

The IBM Z Cyber Vault solution provides a comprehensive framework for protecting critical IBM Z data against logical corruption and cyberattacks by using isolated hardware and Safeguarded Copy backups for data validation and recovery. Therefore, some important design decisions must be made, such as identifying the data to protect, defining security requirements, setting data capture frequency, and developing operational procedures.

This chapter outlines the supported IBM Z Cyber Vault capabilities with best practice design considerations, operational roles, environment setup, validation processes, recovery methods, and security enhancements.

The following topics are covered in this chapter:

- ▶ 3.1, “Key design considerations” on page 52
- ▶ 3.2, “IBM Z Cyber Vault environment operational roles” on page 54
- ▶ 3.3, “Environment setup” on page 56
- ▶ 3.4, “Validation, forensic analysis, and recovery processes” on page 58
- ▶ 3.5, “Extending the air gap” on page 75
- ▶ 3.6, “Enhancing the security posture” on page 75

## 3.1 Key design considerations

Before you use the IBM Z Cyber Vault supported capabilities, consider the following items.

- ▶ Which data must be protected by the IBM Z Cyber Vault solution?

This information helps determine the storage and system capacity, and the location for protecting and validating the environment. The implementation of the IBM Z Cyber Vault solution can be completed in stages across the relevant systems, based on the requirements and priorities of the application environments.

- ▶ Are there security or other business requirements that prescribe the location for the IBM Z Cyber Vault environment and the necessary degrees of isolation for data and processing?

These factors drive the topology and the locations of the storage and system environments for data validation.

- ▶ What is the wanted capture frequency and retention of the data?

The frequency of the Safeguarded Copy backups, the length of time they are retained, and the measured change rate of the application storage environment determine the storage capacity that is required for the Safeguarded Copy backups.

- ▶ What is the current architecture and replication environment?

Using the current environment as a starting point, which is combined with the decisions that are needed to fulfill protection and recovery requirements, provides the next steps for planning and the acquisitions that are required to implement the blueprint.

### 3.1.1 Applying business and technical requirements

Before beginning the technical design and deployment of the IBM Z Cyber Vault solution, consider the business goals for business continuity (BC) and recovery of critical IT services. Reliance on digital technology to conduct business operations is extensive. Identify which IT services are most critical among the ones that are delivered by the targeted IBM Z environments. Identify straightforward application services and the data services that are provided by the IBM Z environment that support distributed processing or IT services that are not hosted on IBM Z.

### 3.1.2 IBM Z Cyber Vault environment design

When performed in a structured and methodical manner, the IBM Z Cyber Vault environment can be designed in a way that supports business goals and priorities.

A logical progression for the design of the IBM Z Cyber Vault environment should include the following items:

- ▶ The management method for generating Safeguarded Copy backups of the data by using either Geographically Dispersed Parallel Sysplex (IBM GDPS) Logical Corruption Protection (LCP) or IBM Copy Services Manager (CSM).
- ▶ The sizing of extra storage that is required for the Safeguarded Copy backups based on the location, frequency, retention period of the copies, and the data change rate.
- ▶ The number of IBM Z Cyber Vault logical partitions (LPARs) that are required for performing validation, analysis, and recovery of the production environments and data.

- ▶ To comply with security requirements, consider using a jump server or deploying more IBM Z Cyber Vault capabilities, such as offensive security (OffSec) exercises or offline tape backups.
- ▶ Whether the IBM Z Cyber Vault environment uses an existing IBM Z system, such as a current disaster recovery (DR) environment, or whether the IBM Z Cyber Vault environment runs stand-alone in its own IBM Z server and possibly in a different location.

As part of the IBM Z Cyber Vault environment implementation, consider scheduling an IBM Z Cyber Vault workshop, which an IBM representative or IBM Business Partner can arrange. In this workshop, subject matter experts (SMEs) assess the production environment, including storage, base z/OS, RACF, Virtual Storage Access Method (VSAM), IBM Customer Information Control System (IBM CICS), Db2, network components, Information Management System (IMS), and other subsystems to determine the most appropriate environment setup. In this context, IBM can also create the parameter files for the IBM Z Cyber Vault data validation asset as a service offering.

### 3.1.3 Developing and testing validation processes

Operational processes are important to effectively use the IBM Z Cyber Vault solution. Some of the areas to address include defining the operational sequencing of daily IBM Z Cyber Vault processes and procedures. This work includes automation strategies, scheduling regular daily operational validation, and checking validation status. The operation of the daily IBM Z Cyber Vault processes is important and can be fully automated, which is preferred, or performed by an individual or orchestrated across several operational actors. Safeguarded Copy backups are taken regularly, and validation is performed one or more times a day. The capture frequency and validation frequency are independent. It is common to capture Safeguarded Copy backups more frequently than validations are performed. Validations can be performed on recovered Safeguarded Copy backups or on a FlashCopy of the source.

As a best practice, perform as many validations as possible during a 24-hour period. The achievable number depends on several factors:

- ▶ How many environments will be validated.
- ▶ The size and complexity of the total environments to validate.
- ▶ How long each validation takes to complete.
- ▶ Whether your validations are automated or performed manually.
- ▶ Whether you want to validate from FlashCopy or recovered Safeguarded Copy backup.

### 3.1.4 Selecting the tools and utilities for data validation and recovery

The IBM Z Cyber Vault solution is flexible enough to incorporate tools and utilities that can make detection and recovery faster, consistent, and more reliable, and help reduce downtime after a logical corruption incident. Standard utilities that come with the operating system (z/OS) and IBM middleware products (such as CICS, IMS, Db2, and others) can be used for data validation and recovery purposes.

The tools and utilities that are described in this publication cover the most common scenarios. However, there might be other scenarios where you must check different databases, control files, or other elements. The IBM Z Cyber Vault solution enables the use of other tools and utilities that support your validation and recovery requirements.

### 3.1.5 Testing the IBM Z Cyber Vault environment

After the IBM Z Cyber Vault environment is defined and verified, activate it for operations. This task involves testing IBM Z Cyber Vault operational capabilities and procedures.

1. Verify the basic IBM Z Cyber Vault environment by starting and operating it.
2. Test the Type 1, Type 2, or Type 3 validation capabilities that were built for the IBM Z Cyber Vault environment (see 1.4.2, “IBM Z Cyber Vault capabilities” on page 11).
3. Check that the operational and run procedures are documented and work properly.
4. Make any adjustments that are necessary for smooth operations.

### 3.1.6 Production cut-over and ongoing monitoring

Production cut-over and monitoring occur when the IBM Z Cyber Vault environment starts and performs as a part of the IT production environment. Ongoing monitoring is important to identify and remediate any operational or technical adjustments that are needed for good operation. During this step, consider periodic audit checks of the overall IBM Z Cyber Vault environment and operation. Ideally, this auditing should be done outside of the daily IBM Z Cyber Vault operations and focused on inspecting and confirming a solid operation environment. Because IT operations are always changing, the IBM Z Cyber Vault environment can be affected by these changes. Ongoing monitoring should include proper change management to help ensure that changes to the systems, applications, or data occur.

At the end of every IBM Z Cyber Vault validation, the results should be analyzed to determine whether any course of action is required. The validation results can be forwarded by email, generate an alert to be picked up by automation, or be written to a dataset that is accessible on a staging volume.

## 3.2 IBM Z Cyber Vault environment operational roles

Operational roles are critical to planning, deploying, and effectively managing an IBM Z Cyber Vault environment. Also, an essential attribute of any cybersecurity strategy is separation of duties between administrators. More than one person is needed to complete a security-related task. This process helps avoid conflicts of interest and can help you better detect control failures that might lead to security breaches, information theft, and violations of corporate security controls and policies.

The roles and responsibilities of team members are as follows:

- ▶ IBM Z Cyber Vault solution architect

The IBM Z Cyber Vault environment requires thoughtful and comprehensive design and planning. Determining the scope of IBM Z Cyber Vault operations and specific system hardware and software configuration are key to both deployment and efficient operations. Understanding specific BC requirements and collaboration with application and line of business (LOB) teams are important to effective planning and design. This person is responsible for overall design and close collaboration with operations to help ensure efficient and effective daily operations of IBM Z Cyber Vault.

- ▶ Business continuity representative

The BC representative in this role represents business requirements and works closely with the IBM Z Cyber Vault solution architect to help ensure that BC risk management, and recovery requirements are understood. They are also involved in assessing the financial risk of business loss in helping to align the appropriate IBM Z Cyber Vault capabilities and return on investment analysis during early planning phases for IBM Z Cyber Vault.

- ▶ IBM Z systems programmers

Systems programmers play a key role in the IBM Z Cyber Vault installation and configuration activities to set up both basic IBM Z Cyber Vault environments and building and customizing Type 1 and Type 2 validation capabilities by using available IBM Z utilities and tools. They also provide ongoing technical support for the IBM Z Cyber Vault environment.

- ▶ Application specialists and owners

Application specialists and owners understand the various applications that use the IBM Z Cyber Vault environment and play a key role in a Type 3 validation, that is, designing and developing application data validation techniques to meet the unique needs of the business. Operations and automation specialists perform daily operations and monitoring of the IBM Z Cyber Vault environment. They also play a key role in designing and developing automation capabilities by using various automation tools and their knowledge of operational specifics for the IBM Z Cyber Vault environment.

- ▶ Security architects and administrators

Security architects and administrators draw from multiple teams. Mainframe security administrators who support IBM RACF or another SAF product have a key role in setting proper security on the mainframe. Network security experts set up firewall rules and other network security to help ensure privacy and limited access to IBM Z Cyber Vault components. Finally, IT security architects and people that are involved with overall enterprise cybersecurity provide direction, standards, and help implement the most secure environment. This role is critical for the solution.

- ▶ Storage administrators

Storage administrators play a key role in managing the DS8000, the overall storage, Flash Storage, mirroring, and other items. They work closely with the IBM Z systems programming team and the IBM Z security administrators to set up and operate this environment.

- ▶ Automation administrators

Automation administrators (GDPS and CSM) play a key role in managing the scripts to run the validation and use the Safeguarded Copy environment. They work closely with the IBM Z systems programming team and the IBM Z security administrators to set up and operate this environment.

- ▶ Database administrators

Database administrators (DBAs), who recover data, run forensics, forward recovery, and do other work in the IBM Z Cyber Vault environment work in the larger team to define the processes that are needed to benefit from IBM Z Cyber Vault. There might be new tools that they learn. This role is key for all IBM Z environments that use databases.

- ▶ Network administrators

The network administrator provisions and manages the network connections that are used between all IBM Z Cyber Vault components. They also play a security role, as noted in “Security architects and administrators” on page 55.

## 3.3 Environment setup

Implementing an IBM Z Cyber Vault environment requires isolated (virtually or physically) hardware and software components. You need storage systems with functions like Safeguarded Copy for capturing and storing protected copies of your data. You also need IBM Z hardware for implementing logical partitions (LPARs), where you can run data validations, forensic analysis, recovery actions, or any other use case that is related to data corruption. To optimize the IBM Z Cyber Vault capabilities, use a set of IBM Z software tools.

Software tools enable and define the level of automation and depth of analysis for the recovery of the environment, but the number of tools and the subsystems that are included in the solution determines the cost and complexity of the environment. It is a best practice that the infrastructure and application teams work together to define the best configuration that provides all the expected benefits while balancing risk and cost.

### 3.3.1 Networking best practices

The IBM Z Cyber Vault networking environment should be constructed in a logical bubble with no intersection with the existing production networking environment. The IBM Z Cyber Vault networking environment should be in an isolated network, which can be achieved by using the following methods:

- ▶ Physical isolation that uses dedicated Open Systems Adapter-Express (OSA-Express) interfaces for the IBM Z Cyber Vault LPARs.
- ▶ Virtual isolation that uses shared OSA-Express interfaces with defined logical separation that is achieved by configuring a separate virtual LAN (VLAN) and IP subnetwork for the IBM Z Cyber Vault LPARs with optional firewalls.

From the core router network perspective, the IBM Z Cyber Vault environment should be fenced off from the existing production environments with limited access points. The IBM Z Cyber Vault environment should be accessible only through a few predetermined IP addresses, ports, and protocols and protected by a firewall router that requires VPN access and encrypted flows to communicate with the IBM Z Cyber Vault z/OS LPARs. The LPARs should be accessible only through the network by using Network Address Translation (NAT) IP addresses because the LPARs are activated by using the same configurations as production and duplication of IP addresses in the network must be avoided.

Access to the IBM Z Cyber Vault environment should be controlled and managed by a VPN that allows only DNS, OSPF, Telnet, SNMP, and SMTP traffic to a predefined set of NAT IP addresses that are allocated to the LPARs while all other traffic is blocked by the firewall rules that are defined within the access control lists (ACLs). If OSPF is used by the production z/OS LPARs, OSPF flows must be permitted within the IBM Z Cyber Vault environment boundary to the VPN edge firewall router. The IBM Z Cyber Vault VPN edge routers should advertise only the NAT IP addresses representing the static VIPA of the IBM Z Cyber Vault LPARs into the core network while all other z/OS LPAR IBM Z Cyber Vault IP addresses are suppressed.

### 3.3.2 Required software and tools

You might already have a system automation tool that can improve the capabilities of the IBM Z Cyber Vault solution. This tool might include data recovery, forensic analysis, and different types of validation (including Type 3 validation) capabilities.

IBM offers automation solutions that are policy-based that provide the foundation for fully automating an IBM Z Cyber Vault environment.

The software should be installed in the production environment before you start the IBM Z Cyber Vault implementation. Because the IBM Z Cyber Vault environment is an exact replica of the production environment, the following software should be included:

- ▶ z/OS and JES (SDSF REXX and REXX Runtime)
- ▶ IBM GDPS LCP Manager or IBM CSM
- ▶ IBM Z NetView and IBM Z System Automation

### 3.3.3 Utilities for data validation by subsystem

Table 3-1 lists the subsystem utilities that you can use in the IBM Z Cyber Vault environment to validate data.

Table 3-1 Utilities to build the IBM Z Cyber Vault environment

Data type	Validation program
<b>Type 2 (Data Structure)</b>	
RACF database	IRRUT200 (INDEX MAP)
ICF catalog (basic catalog structure (BCS) only)	IDCAMS DIAGNOSE IDCAMS EXAMINE
ICF catalog (BCS + VSAM Volume Data Set (VVDS))	IDCAMS DIAGNOSE <sup>a</sup>
VSAM-key sequenced dataset (KSDS) (includes AIX and VSAM Recovery Record Data Set (VRRDS))	IDCAMS EXAMINE <sup>b</sup>
Partitioned dataset (PDS)	IEBCOPY UNLOAD
Partitioned dataset extended (PDSe)	IEBPDSE
CICS/TS	Integrated Data Cluster Access Method Services (IDCAMS)
IBM MQ	CSQ4IVP1 CSQUTIL CSQDLQH

IBM Db2	DSN1COPY DSNUTILB <sup>c</sup>
IMS	FABPMAIN
Endevor <sup>d</sup>	NDVRC1 (C1BM5000)
ADABAS <sup>d</sup>	ADAREP
<b>Type 3 (Data Content)</b>	
Application data <sup>e</sup>	Client application dependent

- a. You may run a **DIAGNOSE** with a comparison of each owned VVDS.
- b. You may run an optional **VERIFY** and optional **INDEX CHECK** system and application databases.
- c. Data type validation under development.
- d. Requires separately licensed software (like IBM IMS High Performance Pointer Checker).
- e. All Type 3 validation processes are client-developed and provided.

### 3.4 Validation, forensic analysis, and recovery processes

Section 1.4.2, “IBM Z Cyber Vault capabilities” on page 11 introduces the capabilities of the IBM Z Cyber Vault solution. You can use the capabilities in different ways depending on your unique business requirements and service-level agreements (SLAs). SLAs typically tie into recovery events that are measured by the targets that are defined in the Recovery Point Objective (RPO) and the Recovery Time Objective (RTO). Implementing IBM Z Cyber Vault capabilities and a data validation frequency can help meet RPO and RTO targets if there is a cyberattack.

Figure 3-1 on page 59 shows the IBM Z Cyber Vault capabilities as a repeatable data validation process that, if necessary, enables forensic analysis and recovery. Do data validation at intervals that align with your recovery strategy.

If no corruption is detected, you can save the environment to offline tape media for added security and longer-term retention. If data corruption is found, move to the forensic analysis phase to investigate how it happened, when, and to what extent. With that information, you can prepare a recovery plan that includes recovery actions in the IBM Z Cyber Vault and production environments.

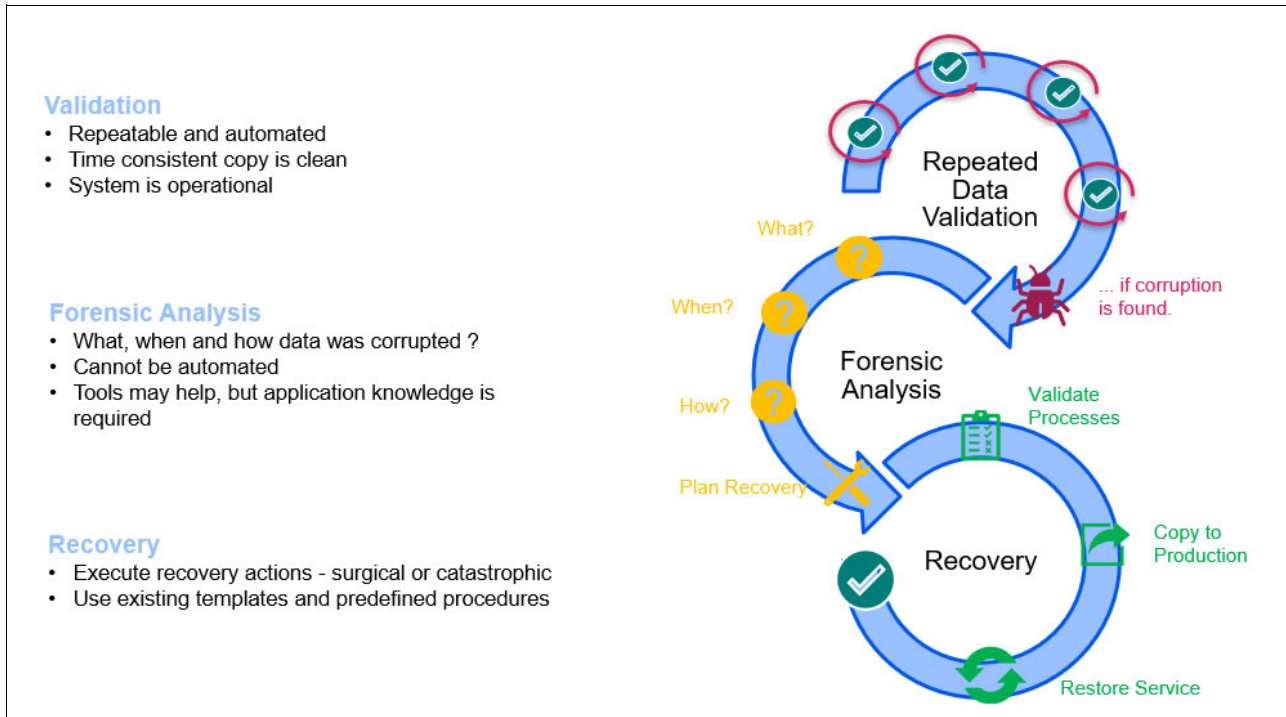


Figure 3-1 IBM Z Cyber Vault capabilities

### 3.4.1 Validation

When considering or implementing the IBM Z Cyber Vault capabilities, and specifically an IBM Z Cyber Vault environment, the software stack plays a fundamental role because it determines the strength of your system when facing cyberattacks and the resiliency level that enables your business to get back to normal in the shortest possible time while minimizing data loss.

We focused so far on the architecture that is required to have a proper IBM Z Cyber Vault environment, which has isolated hardware capacity to start an air-gapped system from an immutable point-in-time copy of your entire production storage repository.

Now, take the created point-in-time image of the production environment and validate it to ensure that there is no data corruption. Your options for validation include FlashCopy or Safeguarded Copy backup. Data corruption can affect the system's software structure or the application data.

The data validation process consists of three different types:

- ▶ *Infrastructure validation* is the process of performing an initial program load (IPL) from a Safeguarded Copy backup or FlashCopy to validate that no data corruption has taken place that would prevent a successful system IPL. This process is defined as Type 1 validation.
- ▶ *Data structure validation* is the process of checking the z/OS system to search for potential structural corruption in any of the control files, configuration libraries, catalogs, repositories, file systems, database systems, or any other component that makes your production z/OS environment run. This process is defined as Type 2 validation.
- ▶ *Data content validation* is concerned with each user's own application data, and because the lifecycle of this data is managed by business applications, only the user can provide validation processes to verify whether business-related data was corrupted. IBM Z Cyber Vault provides a safe environment where users can run their own programs and procedures to this end. This process is defined as Type 3 validation.

Figure 3-2 outlines the IBM Z Cyber Vault validation process.

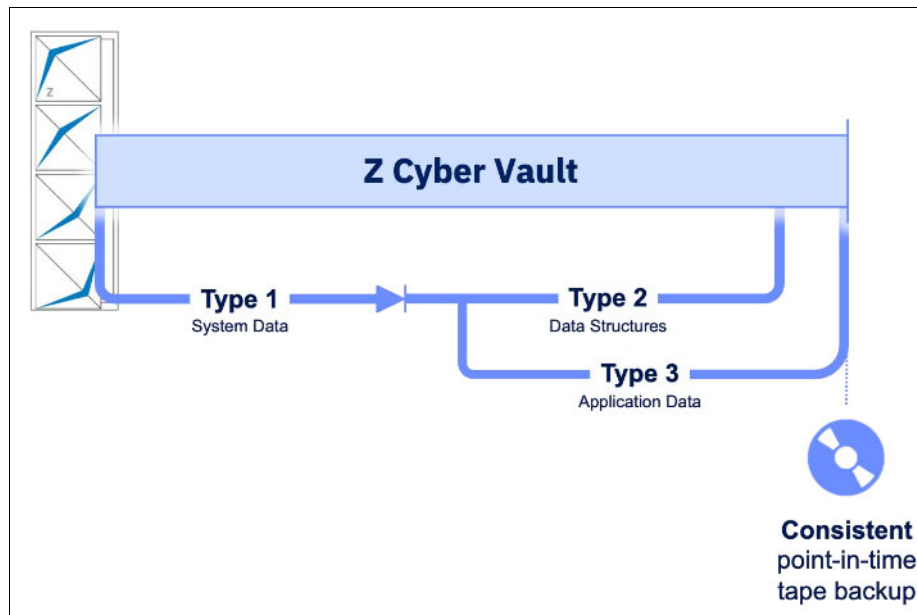


Figure 3-2 IBM Z Cyber Vault validation process

After successfully completing all the data validation procedures, you may take a tape backup copy of this consistent, validated environment as a second layer of protection (for more information, see 3.5, “Extending the air gap” on page 75).

If data corruption is encountered, you must understand how it happened and to what extent by using forensic analysis (for more information, see 3.4.2, “Forensic analysis” on page 62). This analysis also helps determine the recovery actions, which require a combination of tasks, software tools, and environments.

### Data validation frequency

Data validation is one of the most important functions in your IBM Z Cyber Vault environment because it helps ensure that a backup is valid if you need to use it. The data validation procedures include data structure validation and application data validation. By running these procedures, you can detect logical corruptions in your environment.

As a best practice, run the validation function frequently, which means that you should validate as many point-in-time copies (either FlashCopy or Safeguarded Copy) as possible within the confines of your environment. In this way, you can detect a logical corruption earlier and reduce the impact of the corruption.

To accomplish this task, you must have the appropriate tools and automation because the time that it takes to run a full validation of your complete environment ultimately determines how frequently you might start the IBM Z Cyber Vault environment. This frequency determines how often it makes sense to copy your data and how far back you must go to start a recovery if needed (see Figure 3-3).

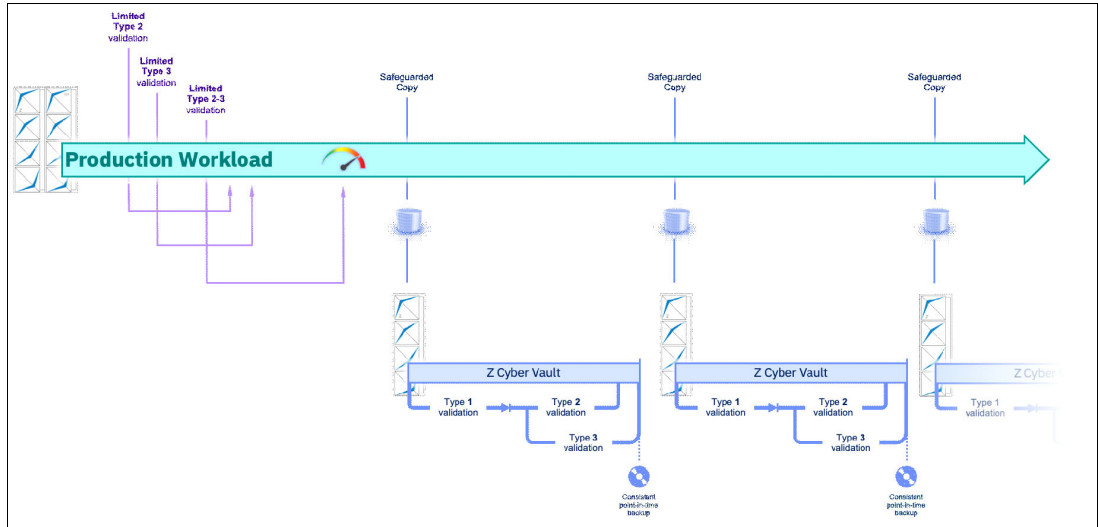


Figure 3-3 Data validation frequency

### Data validation tools

In the isolated validation environment, you use several mechanisms to identify system and data structure corruption. The first mechanism is attempting to perform an IPL of the IBM Z Cyber Vault recovery system, and then starting each of the subsystems within the z/OS image. As part of this process, it is possible to develop procedures and best practices that use system utilities and licensed programs to perform the required and recommended validations.

It is the combination of the correct tools and best practices that enable you to validate your data effectively and efficiently and be better prepared in case there is an incident that requires data analysis and recovery.

Even though z/OS and several of its subsystems provide utilities to perform system functions that aid in the validation and recovery of data, they are not designed to be used in a cyberattack scenario where a fast response and accurate actions are required. Therefore, software tools that deliver features and functions to address these requirements in a more efficient and comprehensive way should be considered.

The challenge of identifying data issues, the extent of them, and the corresponding recovery tasks can be faced only with the correct software tools, which must be readily available for use in such situations.

Software tools selection is a key step in defining the IBM Z Cyber Vault environment, the expected capabilities, and the use cases that are implemented. When selecting potential software tools to support the IBM Z Cyber Vault environment, the overall goal is always the same: reduce downtime. Everything that helps to make recovery processes faster and more comprehensive are welcome. Dealing with logical corruption, which might have spread to several applications in an enterprise, is an exceptional process that must be managed under stressed conditions. Therefore, it makes sense to think beyond the normal when selecting extra software.

### 3.4.2 Forensic analysis

Generally, corruption events can be one of two types:

- ▶ *Targeted corruption* affects a subset of data, which minimizes the amount of data that is moved. Although any data can be corrupted, here we consider two targeted corruption scenarios. One scenario addresses the case where a Db2 application data is corrupted. This scenario requires a surgical recovery of that data, which is in a capture that was taken earlier in the day, and potential steps to roll forward by using Db2 logs if they are available. The second scenario addresses the case where a sequential file that is used throughout the batch process is corrupted.
- ▶ *Widespread corruption* might not affect every byte of data, but if the impact is system-wide and requires a full system or sysplex-wide restoration, it is a catastrophic scenario, so you must restore the entire environment.

The specific steps vary based on the type of event that you had or suspect that you had.

This section provides guidance about the infrastructure, tools, and processes that are needed to identify and move the “good” data back to your production environment so you can restore services as quickly as possible. Because it is not possible to consider every potential topology and corruption scenario, this section focuses on some common ones. Although your specific situation might be different, the techniques and practices are common.

Forensic analysis is a manual activity that requires deep technical skills across the distinct IBM Z technology that is deployed in each installation. It requires fundamental IBM z/Architecture, operational understanding, and specific application and database knowledge.

One key aspect of data corruption recovery is that in a data corruption scenario the decision to restore data is not driven by the operations team and cannot be established in advance. A careful analysis of the applications that are involved and the data that is lost must be performed with the lines of business (LOBs). This process is exactly what the forensic analysis is about, and it determines what data can be recovered, what data makes sense to recover, and what data is kept.

The forensic analysis process is a vital part in the overall IBM Z Cyber Vault solution. Imagine the following situation:

- ▶ You experience a malicious activity at 10:15 AM in the morning.
- ▶ Your quick analysis shows that the cyberattack spread to numerous databases in your production environment, so you must shut down two thirds of your business.
- ▶ Your backup frequency is every hour. So, you decide to set the production environment back to 10:00 AM by restoring a validated Safeguarded Copy backup.

Everything works fine again, but where exactly did the problem start? How was it possible that someone made unauthorized changes to your system? Your production system cannot tell you because it was restored to a point before the corruption occurred. In this situation, forensic analysis is needed to find out why, when, and how something went wrong so that further attacks can be avoided.

During this analysis, consider the real impact of the cyber event. You might discover data corruption in some files, databases, or applications, and you act based on this analysis to recover what is needed. But how do you know that you discovered all the corrupted data, or even worse, that you effectively removed all the causes of data corruption? This scenario requires an analysis that goes further than the immediate recovery of known data corruption and should continue even after the recovery actions.

During the forensic analysis process, you investigate problems and check which recovery actions must be carried out. To successfully do so, you must have a system that resembles the production environment at the time of the cyberattack. You do not know which product or application introduced the data corruption operation, so potentially everything that was running in the production environment at the time of the cyberattack might be the cause.

The IBM Z Cyber Vault environment is a safe system where you can conduct all your research without worrying about affecting your production environment. You run several tools to their full potential to investigate where the data corruption started and how. Many of these tools are described throughout this publication. You might have other IBM or third-party software that you are using in your production environment that might also help. Log analysis tools are some of the key utilities that you should consider for everything that creates a log in z/OS, including z/OS itself.

In the end, the objective of forensic analysis is twofold: Find and fix the vulnerability and come up with a fast and effective recovery plan and procedure. All software tools that are required for this purpose should already be in the production environment because the IBM Z Cyber Vault environment is a replica of production, so the tools already are installed in the production environment. Also, it is necessary for some of these tools to be monitoring in real time what is happening in production, continuously collecting data that can then be used in the forensic analysis and recovery of data and applications.

To perform forensic analysis, you might need two different recovery sets depending on the DS8000 system. You can obtain these two sets, for example, with GDPS LCP Manager support.

In the Db2 scenario that is shown in Figure 3-4, the RC1 volumes contain the last clean copy of data, and the RC2 volumes contain the first set of volumes with corrupted data. Now, you can apply the Db2 log to the clean copy at the point where the corruption first started.

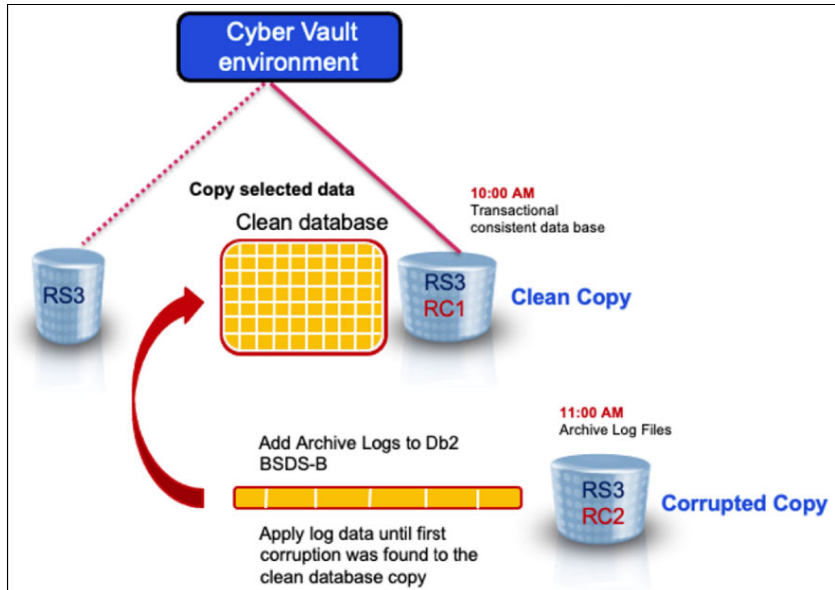


Figure 3-4 Forensic analysis in a Db2 scenario

You can either establish a FlashCopy relationship directly from the RS3 set of volumes to the recovery set of volumes (RC1), or do a recovery action of one of your Safeguarded Copy backups (captures) to your RC1 set of volumes.

For forensic analysis, a **NOCOPY** operation is suitable. The RC1 set of volumes are Extent Space Efficient (ESE) volumes (thin-provisioned). The data is copied from RS3 to RC1 only when changed to a track on a volume on RS3.

The amount of physical space that you plan to set aside for RC1 depends on your change rate and how long you use the RC1 volumes during the forensic analysis phase.

In general, you do sizing for a catastrophic recovery, which should suffice for forensic analysis and surgical recovery.

### Forensic analysis with IZBR

IBM Z Backup Resiliency (IZBR) can be used to simplify and speed up forensic analysis so that you can effectively and efficiently determine what data to move and how to best move it. Without IZBR, this process is a manual and time-consuming one.

IZBR continuously gathers System Management Facility (SMF) data in your production environment. Because that data is captured in your IBM Z Cyber Vault environment, you can use IZBR in the IBM Z Cyber Vault environment to simplify forensic analysis and speed up recovery.

IZBR has a rich set of capabilities that facilitate recovery from disaster events, application failures, and logical corruption events. For the purposes of this document, we focus on the capabilities that are specific to understanding and recovering from logical corruption events. For more information about IZBR, see the [IBM Z Backup Resiliency documentation](#).

This section highlights the features that are most interesting from an IBM Z Cyber Vault solution perspective. The focus is on three specific capabilities:

1. IZBR Cyber Vault Health Check report
2. 3D Virtual Katalog
3. TimeLiner reports

The fundamental value is that IZBR provides information about the status of the non-database managed data at the time of the Safeguarded Copy backup. Although the database management tools can process “fuzzy” backups, which are backups of open files, non-database managed datasets are considered unreliable if they are open for output at the time of a backup. Because the Safeguarded Copy backup is done at a volume level, there is no consideration of the state of the datasets, so it is likely that at least some of the non-database managed datasets are open for output.

Here are the components of IZBR that you can use to perform forensic analysis:

- ▶ The IZBR Cyber Vault Health Check report provides a list of all datasets that are open for output at the time of the Safeguarded Copy backup. This report can be used to identify the best Safeguarded Copy backup to use (the one with the fewest datasets open for write) or to find alternative Safeguarded Copy backups for selected datasets that are deemed unusable from the primary Safeguarded Copy backup.
- ▶ The 3D Virtual Katalog tracks which volumes that a particular dataset was on at a point-in-time. This feature is invaluable because a dataset might not always be on the same volume that it was originally written to. This tool saves the analyst from hunting for a dataset when the dataset is no longer on the same volume in past Safeguarded Copy backups as the catalog says it is on today. IZBR displays which Safeguarded Copy backups contain that dataset and whether the dataset was open or closed.
- ▶ You can use the IZBR TimeLiner Reverse Cascade report to look backwards for a dataset from a particular point-in-time to see all the jobs that used or influenced this dataset. The report displays these jobs and shows for each job the datasets that are used as input or output. This information provides analytics to help forensically identify the points where a dataset might be corrupted.
- ▶ You can use the IZBR TimeLiner Forward Cascade report to create a recovery plan for a dataset that you are restoring. It can be run from the point-in-time from which the dataset is being restored, and it shows everything that you must rerun to bring the application data to its current state. Without this capability, the alternative is to use the scheduler to rerun all jobs, which might cause unnecessary work (jobs that had nothing to do with that dataset) or work might be missed because the scheduler does not know of the relationship of a dataset to an application. If multiple applications depend on the dataset, secondary applications might be out of sync with the restored dataset.

### **Using IZBR: An example**

In this example, the application team implemented a software enhancement. Unfortunately, there was a defect in the code. Several days after it was deployed to production, the defect affected the batch cycles. The corrupted data was extensive. Now, you must consider the steps that are needed to understand the corruption and determine how to recover. IZBR can help identify where the good backups are and initiate those recoveries.

Here are the actions what you must take to recover your data:

- ▶ Recover from a Safeguarded Copy backup in your IBM Z Cyber Vault environment.  
Selecting the appropriate copy depends on several things. If you do not know when the corruption occurred, you might want to start with the most recent capture and work back from there. If you have some idea of when the corruption occurred, you can recover a copy from a point-in-time closer to that event. If you must safely analyze the current data or have it available for comparisons, you might take a new capture and start with it.
- ▶ Determine when the datasets were updated and which jobs made those updates.  
Use the IZBR TimeLiner Reverse Cascade report to help identify preceding jobs that updated the corrupted datasets and when they were updated.
- ▶ Determine which datasets were open and not usable.  
Use the IZBR Cyber Vault Health Check report to identify any open datasets that might need to be restored from another source to help ensure that the data is valid.
- ▶ Identify what jobs must be rerun.  
Use the IZBR TimeLiner Forward Cascade report to identify all downstream jobs from any application that must be rerun after the restore to complete the recovery.
- ▶ Recover critical datasets from valid backups.  
Use the IZBR TimeLiner function to surgically restore the datasets. IZBR automatically generates the Job Control Language (JCL).
- ▶ Recover open tape datasets.  
Use IZBR to recover the tape datasets.

If this attack was a widespread ransomware attack instead of an application defect, the IZBR Audit report can also help identify jobs that run outside the scheduler that update application data.

### 3.4.3 Recovery process

Recovery includes all the processes and procedures, starting with understanding the corruption and ending with re-establishing business as usual. Although the specific steps vary based on things like the event itself, the production architecture, and the IBM Z Cyber Vault architecture, the approach is consistent.

Start with a recovery checklist. Regardless of your specific situation, the following steps should be part of your recovery process:

- ▶ Determine what happened (forensic analysis).
  - What data is corrupted (targeted or widespread)?
  - How was the data corrupted (intentional or accidental)?
  - Determine when the data was corrupted. How far back is the good data?
- ▶ Distinguish where the “clean” data is, that is, traditional backups, previous Safeguarded Copy backups, or somewhere else.
- ▶ Identify dependencies that need to be resolved. Did corrupted data also get shared with other applications or pushed into other data sources?

- ▶ Establish the impact of going back in time.
  - Will you accept data loss?
  - Can you recover forward by using system tools?
  - Must the application rerun transactions or catch up in some way?
  - How will this choice affect upstream and downstream systems?
- ▶ Confirm whether the components that caused the corruption, such as defective code or compromised user ID, must be removed or repaired. Can the cause be identified and immediately fixed to prevent future impacts?
- ▶ Ensure that the infrastructure and processes are in place for moving smaller amounts of data during a surgical recovery and the full system restore for catastrophic.
- ▶ Practice everything in a sandbox or test environment.

Depending on the amount of data that was affected by the logical corruption event and the implemented production environment topology, the process to restore the validated data might differ. There are many ways to restore data back to production.

For simplicity, this section makes the following assumptions:

- ▶ The term “recovery” refers to all the steps that are needed to restore business operations to business as usual.
- ▶ The phrase “recover action” refers to the process of creating the recovery volumes that are used for all other activities.
- ▶ Safeguarded captures can be recovered (by using the “recover action”) to recovery volumes and Type 1 Validation can be performed.
- ▶ The recovery steps start after the Safeguarded Copy recover action and Type 1 validation.
- ▶ System administrators can log on to systems in the IBM Z Cyber Vault environment. These administrators might be the administrators who have access to the production systems or they might be different administrators if separation of duty is wanted.

### **Surgical recovery**

Surgical recovery is simply moving specific data from the recovery volumes back to your production environment and then applying those changes there. That data might be system files, batch files, databases, or any other type of application data. This process might require re-creating data manually or replaying transactions.

Surgical recovery can be done at a z/OS system level (see “System-level surgical recovery” on page 68) and for each of the transaction and data management subsystems in use. In this example, we select Db2 to describe how surgical recovery works for it (see “Db2 subsystem surgical recovery” on page 69). Similar strategies can be designed for other subsystems.

Like forensic analysis, surgical recovery is a manual activity that requires deep technical skills across the IBM Z technology that is deployed in each installation. It requires fundamental z/Architecture and operational understanding, and potentially specific application and database knowledge. The extent to which data can be surgically restored also depends on the tools and utilities in use, and the design of the IBM Z Cyber Vault environment.

### System-level surgical recovery

In an IBM Z Cyber Vault environment that is in the same data center as the primary production data, system-level surgical recovery can be done by copying the data from the recovery volumes to a set of staging volumes (see Figure 3-5) and making the staging volumes available to the production system LPARs. Use your existing processes and procedures to copy the required data from the staging volumes to the production volumes.

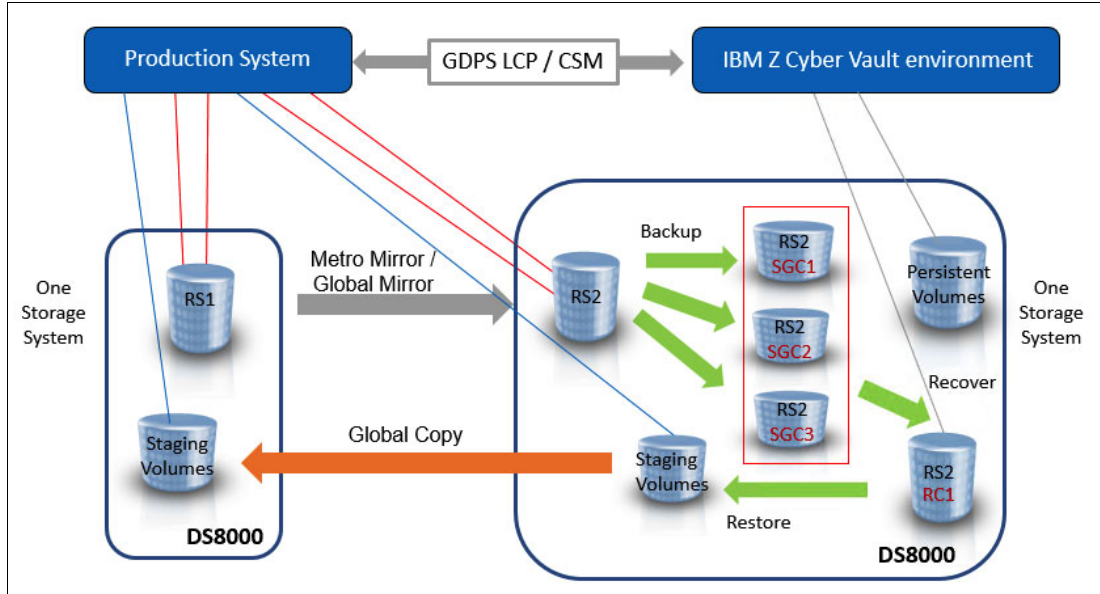


Figure 3-5 Using staging volumes for surgical recovery and partial restore to production

In an IBM Z Cyber Vault environment that is in a remote data center, copy the data from the recovery volume (RC1) to a set of staging volumes in the IBM Z Cyber Vault environment, and copy the staging volumes in the IBM Z Cyber Vault environment to a set of staging volumes in the production sysplex. When 100% of the data is copied, you can bring the staging volumes online in your production environment and copy the data to the production volumes by using standard operating system methods (see Figure 3-6).

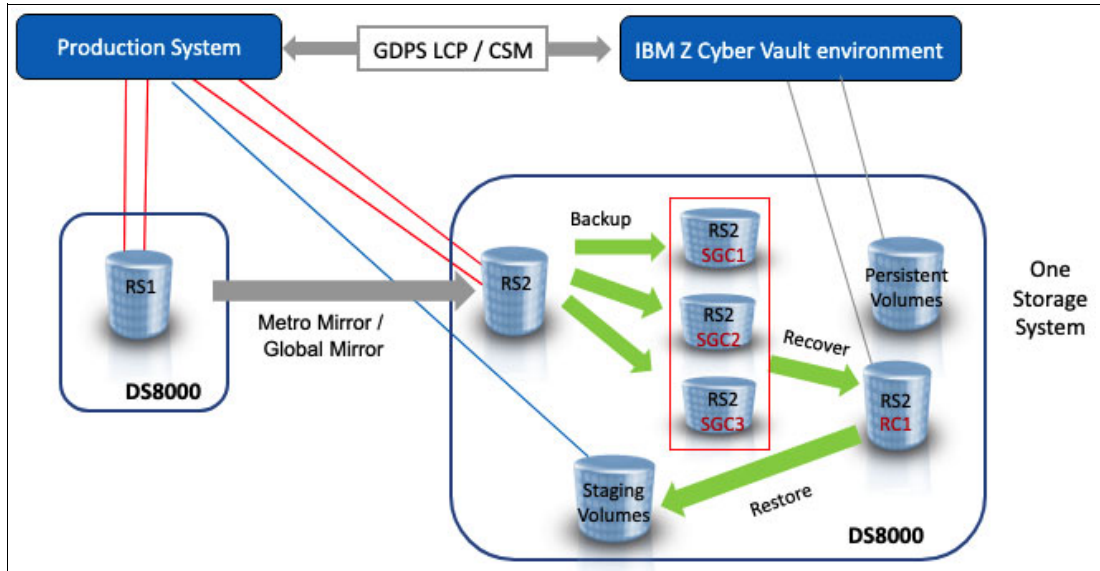


Figure 3-6 Using staging volumes for surgical recovery and partial restore to production

### ***Db2 subsystem surgical recovery***

With Db2 and other log-managed subsystems, you can use log data to roll forward the databases and closer to the point of corruption. For more information about new capabilities that simplify the current, more manual processes, see 5.2, “Db2 for z/OS subsystem rollforward recovery” on page 114.

Here are some considerations for implementing a roll-forward process today:

- ▶ Uncorrupted logs are required. If the corruption event leaves the log files in an unusable state, your ability to roll forward is limited. Consider what your log management processes look like to help ensure that you have the log data that you might need. You need Db2 archive logs, the BSDS that was captured as part of the most recent archive log, and the image copies that are needed for a LOG NO event.
- ▶ The roll-forward process is the same process that you use in your production environment. The difference is where you get the assets from. Ensure that you know where the latest good data might potentially be. If the corruption is limited, the best data might be on your production system. It might potentially be in recent traditional backups. It might be in one of your Safeguarded Copy captures.
- ▶ The specific roll-forward process depends on the version of Db2 that you are using. Use existing documentation, resources, and tools to define that process.

This section describes typical surgical recovery scenarios for Db2 that you can use as an example and reference for what must be considered when designing an IBM Z Cyber Vault environment.

There are three possible scenarios after you identify data corruption in Db2:

- ▶ Scenario 1: Your database system log files, image copies, and any incremental copies are available in production. These files and copies are also accessible through standard access methods because the disk and tape catalogs are not corrupted.
- ▶ Scenario 2: Your database system log files are available, but your image copies of the corrupted database are not available. This situation might happen if the cyberattack targeted the backup data of the database or the ICF or tape catalogs. Assume that you cannot use standard production procedures to recover the database, but your image copies were stored on recovery volumes before being migrated to tape, and so they exist in the IBM Z Cyber Vault environment.
- ▶ Scenario 3: Like Scenario 2, but no image copies exist in the IBM Z Cyber Vault environment because the database image copies are written directly to tape and are no longer accessible due to the cyberattack.

What can you do in each of these situations?

- ▶ Scenario 1: The image copies are accessible in production.  
This scenario is the simplest because data can be recovered by following standard recovery procedures in the production environment. In this case, the IBM Z Cyber Vault environment is used only for validation and forensic analysis.
- ▶ Scenario 2: The image copies are not accessible in production, but they do exist in the IBM Z Cyber Vault environment.

After data corruption is identified through validation and forensic analysis in the IBM Z Cyber Vault environment, you can determine which is the last valid database image copy, which would be on disk and is now available in one of the Safeguarded Copy backups.

The next step is to bring back into production this valid image copy by using the staging volumes and standard z/OS Copy Services (CS) tools. After the image copy of the database is available in production, the regular database recovery procedures can be used to repair the corrupted database.

- ▶ Scenario 3: The image copies of the corrupted database are not accessible anywhere.

This scenario is the most complex because you have only a clean copy of your database in a Safeguarded Copy backup, but no image copy to recover from is available. In this case, you must apply the available log data to the last clean version of the corrupted database to minimize the loss of data.

If you do not use any specialized tools to perform database recovery, here are the steps that you should follow:

1. Recover the most recent Safeguarded Copy backups that contain a copy of the clean database into the RC1 volume.
2. Recover the Safeguarded Copy backups that contain the database log files into the RC2 volumes. Both sets of volumes, RC1 and RC2, must be online to the IBM Z Cyber Vault recovery system.
3. Copy the logs and BSDSs from the RC2 to the RC1 volumes.
4. Reassemble the Db2 zParms with **DEFER ALL**.
5. Apply the database log records to the database up to the point right before the data corruption happened (Db2 **LOGONLY** recovery). This recovery point already was identified through the forensic analysis.
6. Determine the Log Record Sequence Number (LRSN) corresponding to the recovery point.
7. Use the **DSNJ003** utility to update the ENDLRSN on all Db2 member BSDSs.
8. Run Db2 object **LOGONLY** Recovery (Cat/Dir and user objects).

If you use tools such as Db2 Log Analysis Tool for z/OS and IBM Db2 Recovery Expert Pro for z/OS, there is no need to use two recovery sets of volumes. Only the latest Safeguarded Copy backups are required: After you identify the malicious transactions, you can use these Db2 tools to delete the transactions from the database.

After you recovered the database to the most recent status, start the application in the IBM Z Cyber Vault environment and check the status of the recovered database.

When you are ready to bring the recovered database back into production, you can use the staging volumes to do so by first copying the database into these volumes, and then moving it into production (after making the staging volumes available in the production environment) by using Db2 recover **NOSYSCOPY**.

As you can see through this example, the use of specialized tools is a best practice due to the consistency and speed that they enable in the recovery process.

## Catastrophic recovery

Organizations use an incremental restore to production for widespread or catastrophic recovery processes. Safeguarded Incremental Restore to Production implements a process in which the function restores a Safeguarded capture on a PPRC volume to its peer PPRC volume incrementally. The process includes steps that recover the Safeguarded capture to a recovery copy (RC) set and then incrementally restore the RC set to the PPRC production volume.

The location of the peer PPRC volume depends on the topology that you select for the IBM Z Cyber Vault architecture. This volume is often described as “the source of the source.” To demonstrate where the data is restored, this section uses three topology examples.

The following assumptions apply to all the examples:

- ▶ The recover action, the IPLs, and the forensic analysis are complete.
- ▶ The corruption is widespread.
- ▶ All data that must be restored is on the recovery target volume, which becomes the recovery source volume. Any roll forward or data changes are complete.
- ▶ The GDPS scripts are written and tested.

### **Physical isolation with GDPS Metro Mirror and Global Mirror**

The first topology is physical isolation with GDPS Metro Mirror and Global Mirror (GM) (see Figure 3-7).

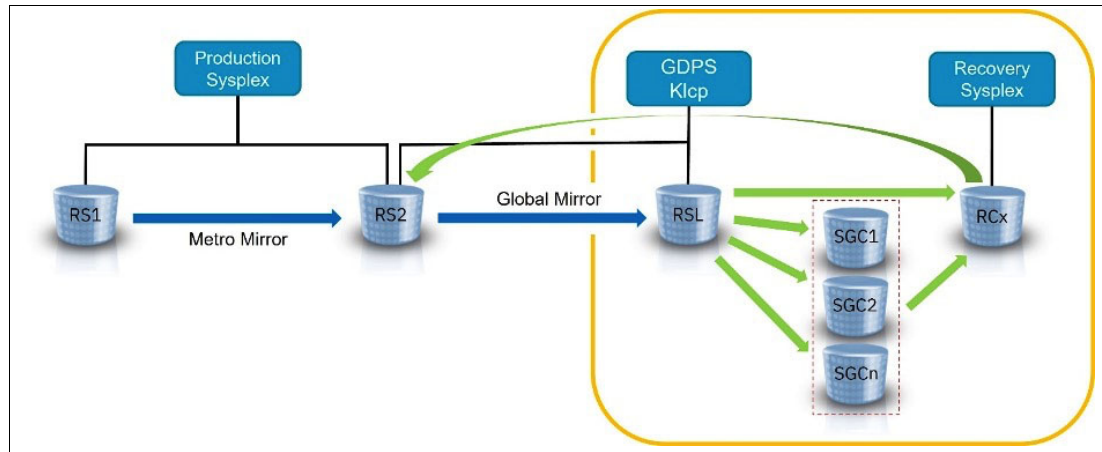


Figure 3-7 Physical isolation with GDPS Metro Mirror and Global Mirror

In this topology, RS1 is the primary production storage, and RS2 is the MM secondary. RSL is the source of the Safeguarded Copy backups, and it is a GM volume off RS2. RCx is the recovery volume, and it is the source of the data to be restored.

To help illustrate the path that the data takes by using the “source of the source” approach, RSL is the source of the Safeguarded Copy backup, and RS2 is the source of RSL. This relationship means that RS2 is the volume to which the incremental restore returns the data.

During normal operations, the GM replication path typically cascades through the Metro Mirror replication path. However, to allow incremental restoration of a Safeguarded capture, the GM primary site cannot also serve as a Metro Mirror secondary.

To address this requirement, the Metro Mirror primary is swapped to create a Multi-Target configuration. After this swap, the environment allows an incremental restore to the production environment.

In a Multi-Target setup, the Safeguarded capture is first recovered to an RC set. From there, the environment incrementally restores the data to the GM primary site.

Metro Mirror is reversed by the HyperSwap process. After the restore is complete, Metro Mirror is restarted from RS2 to RS1 so that RS1 is brought back into sync.

### Virtual isolation with GDPS GM

The second topology is virtual isolation with GDPS GM (see Figure 3-8).

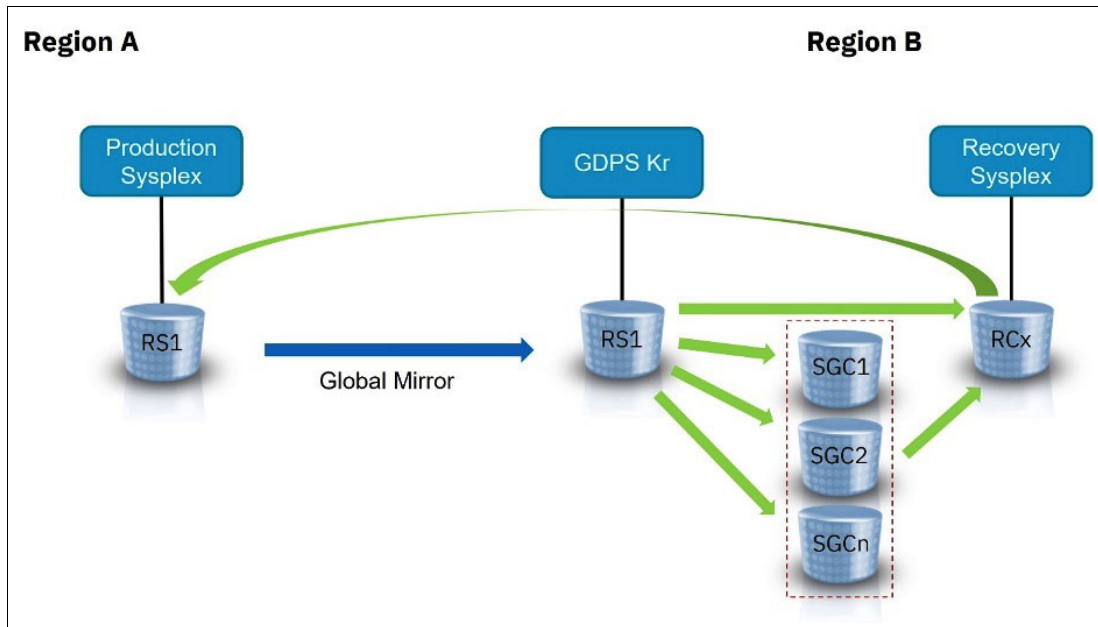


Figure 3-8 Virtual isolation with GDPS Global Mirror

In this topology, A.RS1 is the primary production storage. B.RS1 is the GM secondary, and it is the source of the Safeguarded Copy backups. RCx is the recovery volume, and it is the source of the data that is restored.

Because B.RS1 is the source of the Safeguarded Copy backups and A.RS1 is the source of B.RS1, the environment incrementally restores the data from RCx to A.RS1.

### Virtual isolation in both regions with GDPS Metro Global Mirror (MGM) 4-site

The third topology is virtual isolation in both regions with GDPS Metro Global Mirror (MGM) 4-site (see Figure 3-9 on page 73).

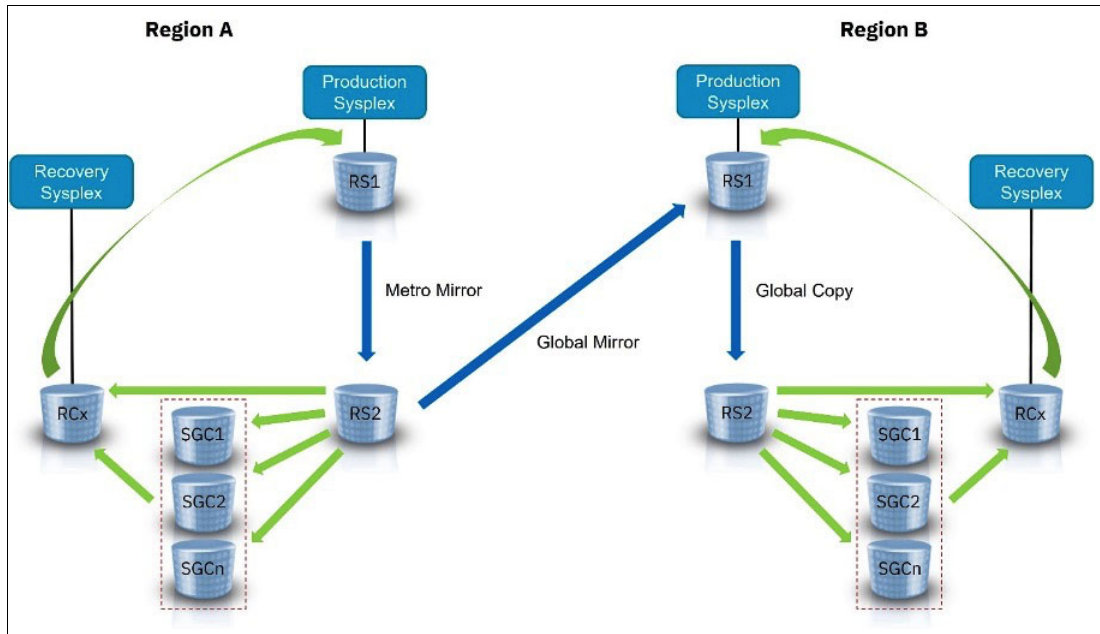


Figure 3-9 Virtual isolation in both regions with GDPS MGM 4-site

In this topology, there are two active vaults. The environment can restore data in either data center. In the primary region, A.RS1 is the primary production storage. A.RS2 is the Metro Mirror secondary and the source of the Safeguarded Copy backups. This relationship makes A.RS1 the source of the source. For incremental restore to production, the environment moves data from the A.RCx recovery volumes to A.RS1.

In the recovery region, B.RS1 is the GM secondary. B.RS2 is the Global Copy volume and the source of the Safeguarded Copy backups. This relationship makes B.RS1 the source of the source. For incremental restore to production in the recovery region, the environment moves data from the B.RCx recovery volumes to B.RS1.

If the restore is performed from the vault in Region A, the Metro Mirror leg must first be suspended. Then, the restore is directed to A.RS1, and Metro Mirror is resynchronized after the restore is complete.

If the restore is performed from the vault in Region B, a region switch must occur first. This action results in a B.RS1 > B.RS2 > A.RS1 > A.RS2 configuration. Then, the restore is directed to B.RS1, and the systems are started from that location. After the restore is complete, all mirroring relationships can be resumed to bring the remaining copies into sync. When required, a planned region switch back to Region A can occur.

### ***Restoring a valid Safeguarded Copy backup to your production environment***

Regardless of your exact topology, the following general steps are required to restore a valid Safeguarded Copy backup to your production environment:

1. Select a validated Safeguarded Copy backup and recover it with your management software GDPS LCP or CSM to the recovery volumes (RC1).

**Note:** A recovery action with the **NOCOPY** option is sufficient to restore data to production. If you have not done a validation for a Safeguarded Copy backup, you must do the validation process before you start the restore to production (see 4.6, “Establishing a validation framework” on page 101) before starting the restore to production.

2. Stop production by shutting down your production systems (LPARs).
3. Suspend all replication relationships from your production volumes (RS1) in your production environment and convert all synchronous replication (metro mirror) to GC.
4. Establish the GC relationship between your recovery volumes (RC1) to the production volumes (RS1) with the interface of your choice (DS Command-Line Interface (DS CLI), GDPS, or CSM).
5. Pause the Safeguarded Copy backups until the data is copied and verified in production. Starting new backups during this process might expire older ones, which you might need later.
6. Wait until 100% of the data is copied to the production volumes (RS1), and then suspend the relationship between RC1 and RS1. If you decided to run the production on RC1 to reduce the RTO, you must now shut down the LPARs that are running on the RC1 volumes and wait until all data is copied over to RS1. When all data is copied, which is shown by out-of-sync tracks being zero, remove the GC relationship. You cannot perform an IPL of the production system if the RS1 volumes are the GC target.
7. Start resynchronization of your production replication relationships.

**Note:** While recovering from RC1, every I/O operation goes through the Safeguarded Copy metadata, which impacts the response time. You do not get the same I/O performance for running the production workload on RC1 as on RS1 or RS2.

8. Perform an IPL of your production systems (LPARs) by using RS1. Check the environment and start your application.
9. Re-establish your Safeguarded Copy environment if you stopped it.

**Note:** Even in a virtual isolated environment, you cannot use FlashCopy to copy the data from the recovery volumes (RC1) to the Safeguarded Copy source volumes (RS2) because the RC1 set of volumes are copied from RS2 and set of Safeguarded Copy backup volumes. At the time of writing, the DS8000 microcode does not support FlashCopy from RC1, which is why GC is the best practice to do a restore to production.

## 3.5 Extending the air gap

Tapes are durable and can remain viable for 20 - 30 years, which makes them suitable for archiving data that requires long-term preservation. Also, they provide an air gap between the network and the data, which helps safeguard the data from ransomware and malware. When a validated and consistent point-in-time copy of the production environment is offloaded to tape and stored in a secure, off-site physical vault, the approach provides another layer of protection that reduces exposure to cyberattacks.

For more information, see 5.3.1, “Offline backups” on page 117.

## 3.6 Enhancing the security posture

To enhance an organization’s security posture, a comprehensive approach must include continuous vulnerability assessments, employee training, and robust incident response plans. An air-gapped environment that replicates a real production system serves as an environment for conducting these assessments, analyzing the results, and determining the actions that are required to safeguard and recover applications.

OffSec encompasses a range of proactive security strategies that use tactics similar to the ones that are used by malicious actors in real-world attacks. Common OffSec methods include red teaming, penetration testing, and vulnerability assessment. These operations are typically performed by ethical hackers, who are cybersecurity professionals who apply their skills to identify and address IT system vulnerabilities.

OffSec complements defensive security by enabling security teams to uncover and respond to unknown attack vectors that might otherwise go unnoticed. It also provides a proactive approach compared to defensive security because OffSec measures identify and address flaws before attackers can exploit them.

In essence, OffSec offers insights that improve the effectiveness of defensive security measures and reduces the workload on security teams. As a result, OffSec has become a standard practice in certain highly regulated sectors.

For more information, see 5.3.2, “Offensive security” on page 118.





# Deploying the IBM Z Cyber Vault environment

IBM Z Cyber Vault is a solution that is composed of many components and functional areas that work together to create a more robust cyber resilient environment. The preceding chapters outline key capabilities and components of the IBM Z Cyber Vault environment, and the planning and design considerations that support cyber resiliency objectives.

Several actions are required to deploy an IBM Z Cyber Vault environment, including preparing the storage configuration, Safeguarded Copy setup, automation, isolation and security, and implementing a validation framework.

This chapter describes the tasks that are required to deploy the IBM Z Cyber Vault environment and use its key capabilities.

The following topics are covered in this chapter:

- ▶ 4.1, “Preparing for IBM Z Cyber Vault deployment” on page 78
- ▶ 4.2, “Setting up IBM Z Cyber Vault Storage” on page 80
- ▶ 4.3, “Setting up Safeguarded Copy for IBM Z Cyber Vault Storage” on page 85
- ▶ 4.4, “Setting up IBM Z Cyber Vault automation” on page 90
- ▶ 4.5, “Preparing the IBM Z Cyber Vault environment for validation” on page 92
- ▶ 4.6, “Establishing a validation framework” on page 101
- ▶ 4.7, “Cyber Vault Data Validation for IBM Z Asset” on page 109

## 4.1 Preparing for IBM Z Cyber Vault deployment

Deploying the IBM Z Cyber Vault solution begins with a clear understanding of the environment topology and configuration. This section uses a Geographically Dispersed Parallel Sysplex (IBM GDPS) Global Mirror (GM) multisite configuration as a reference model and describes how GDPS Logical Corruption Protection (LCP) Manager integrates into the deployment (see Figure 4-1).

GDPS LCP Manager supports several multisite topologies. The IBM Z Cyber Vault solution supports all GDPS topologies, but the recovery steps for each topology differ based on the installation.

GDPS LCP Manager is implemented in one of the existing controlling systems of the GDPS environment. Extra controlling systems are not required for GDPS LCP Manager. The primary factor that determines where GDPS LCP Manager is implemented is whether the environment uses a virtual or physical air-gap configuration.

In a virtual air-gap configuration, the protected copies (Safeguarded Copy backups, also known as captures in GDPS) are in the same storage system as one of the existing production or disaster recovery (DR) copies in the GDPS environment.

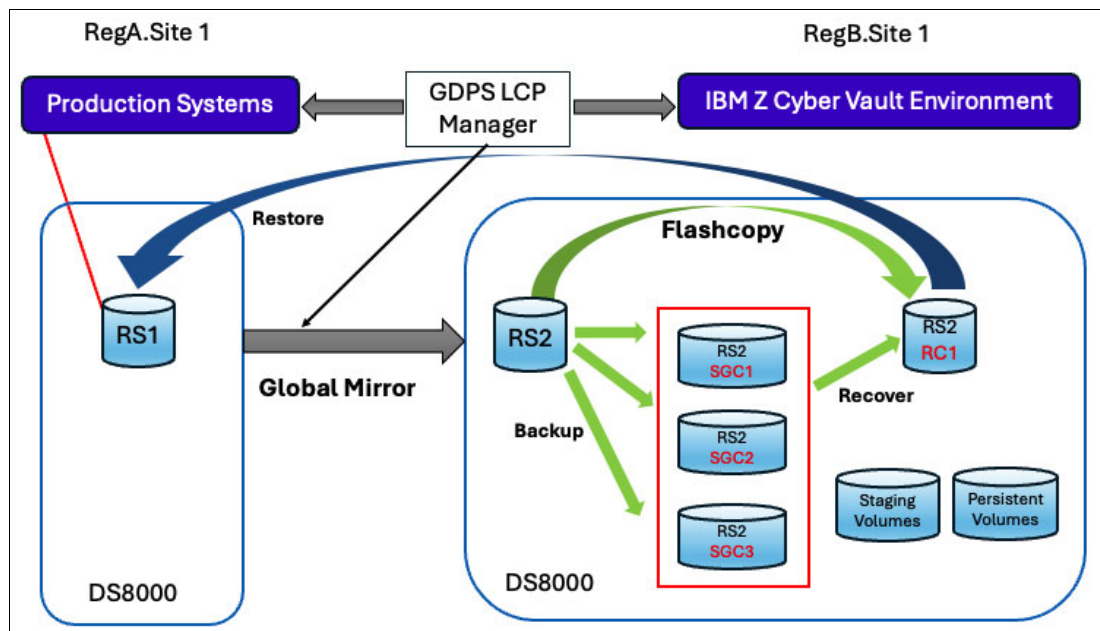


Figure 4-1 IBM Z Cyber Vault replication in a Global Mirror 2-Site configuration

The example configuration includes a single asynchronous GM session that is called CVDASD. The primary and secondary copies of data are referred to as replication sites (RSs), designated RS1 and RS2. The primary copy in RegA.Site 1 is referred to as CVDASD.RS1.

The details of the example environment, which is shown in Figure 4-2 on page 79, are as follows:

- ▶ All logical partitions (LPARs) are built on IBM z17™ platforms.
- ▶ The applications run in a multisite data-sharing Parallel Sysplex with two production systems, PRD1 and PRD2. The primary GDPS GM Control LPAR (Kg) runs in KSYSG by using its isolated disk, and its GDPS Recovery LPAR (Kr) runs on KSYSKR.

- ▶ In the IBM Z Cyber Vault environment (RegB.Site 1), two spare LPARs are available to perform an initial program load (IPL) of the recovered systems.
- ▶ The GDPS GM DS8000 systems (not shown) contain two logical subsystems (LSSs) with defined base addresses. All volumes may be on any 3390 DASD model type.
- ▶ The production environment is in a sysplex with a coupling facility (CF) (CF1). The IBM Z Cyber Vault environment is also in a sysplex that uses a CF (CV-CF1).

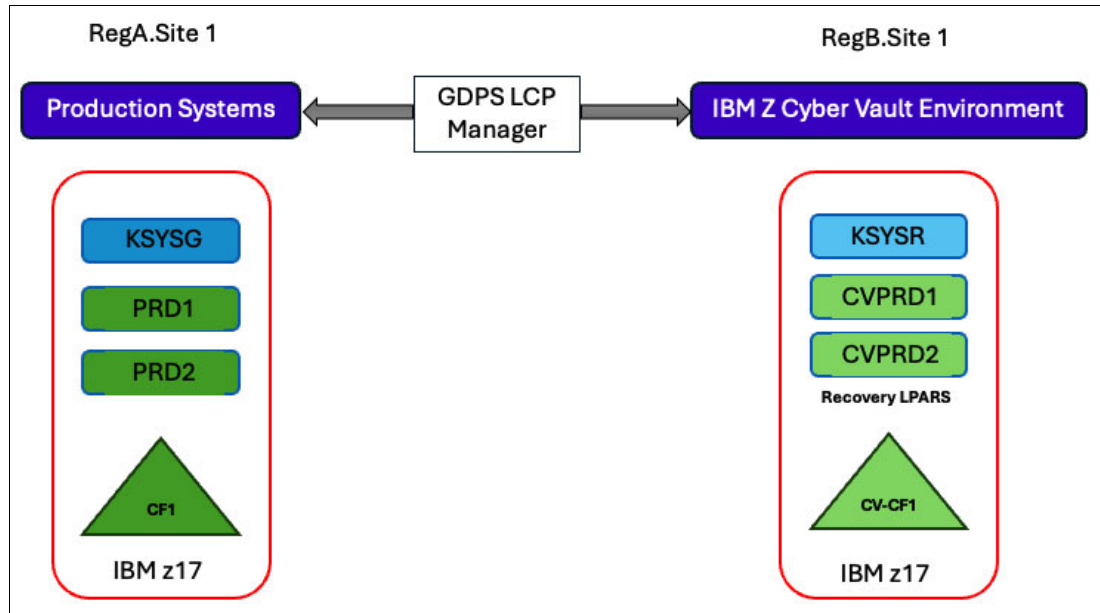


Figure 4-2 LPAR placement in the example environment

For more information about supported GDPS configurations, see *IBM GDPS: An Introduction to Concepts and Capabilities*, SG24-6374.

To help move from design to an operational IBM Z Cyber Vault environment in a structured and auditable way, a deployment checklist aligns technical tasks, security requirements, and operational readiness into an actionable plan:

- ▶ Prerequisites and decisions. Select the topology and air-gap type (virtual or physical).
- ▶ Set validation. Establish the cadence and retention goals to help determine the required storage capacity.
- ▶ Storage. Size the DS8000 Safeguarded Copy capacity (virtual multiplier), define Safeguarded Copy backups, provision RC1, and allocate Staging and Persistent volumes.
- ▶ GDPS LCP Manager. Create the CV\_TEST\_Safeguarded Copy profile, enable capture and monitor schedules, and provide suitable scripts for **CAPTURE**, **RELEASE**, and **REFRESH** operations in the IBM Z Cyber Vault environment.
- ▶ Input/output definition file (IODF) and Hardware Configuration Definition (HCD). Define CVPRD1, CVPRD2, and CV-CF1; set **IPLTYPE** and **IPLMODE RC1**; and preallocate Couple Data Sets (CDSs).
- ▶ Security and Base Control Program internal interface (BCPii). Configure Hardware Management Console (HMC) permissions, IBM RACF parity, and Integrated Cryptographic Service Facility (ICSF) and cryptographic access.
- ▶ Network isolation. Implement segmentation between the IBM Z Cyber Vault and the production environment by using industry-standard technologies such as firewalls, Network Address Translation (NAT), jump servers, and alternative IP addressing.

- ▶ Parmlib, Coupling Facility Resource Manager (CFRM), and TCP/IP. Define **&CYBERID** symbols; tailor **COMMND**, **AUTOR**, and System Management Facility (SMF); specify CFRM **PRELIST** to CV-CF1; and augment TCP/IP profiles with IP addressing as needed.
- ▶ Automation and IPL. Implement cleanup, refresh, and **LOAD** scripts for FlashCopy and Safeguarded Copy paths.
- ▶ Validation framework. Establish build and runtime libraries, job naming, Type 1 - 3 run times, monitor tasks, and email configuration.
- ▶ Operations and restore. Implement the daily cycle (capture, IPL, validate, and report), weekly Safeguarded Copy **RECOVER** test, and incident **RESTORE PREPARECOPY** flow.

The prerequisites for the IBM Z Cyber Vault solution and design considerations for the supported GDPS configurations are described in Chapter 2, “Planning and designing the IBM Z Cyber Vault solution” on page 15. Considerations regarding validation cadence and retention goals are described in Chapter 3, “IBM Z Cyber Vault capabilities” on page 51.

The subsequent sections walk through the details of the deployment checklist.

## 4.2 Setting up IBM Z Cyber Vault Storage

Storage configuration is a critical component of the IBM Z Cyber Vault architecture. This section describes how to define and allocate the various volume types, including Safeguarded Copy, recovery, staging, and persistent, which are used throughout the recovery and validation processes. It explains how these volumes interact across production and recovery systems, and outlines best practices for sizing, accessibility, and connectivity. Proper storage setup helps ensure that the IBM Z Cyber Vault solution supports both surgical and full recovery scenarios.

This section describes the steps that are needed to set up the storage environment for the IBM Z Cyber Vault environment. These steps include defining the Safeguarded Copy storage environment and the staging volumes that are used when a surgical or catastrophic recovery is required. The persistent volumes contain information that is produced by the IBM Z Cyber Vault validation process.

- ▶ The IODF operating system configuration (OSCONFIG) definitions that are defined for the IBM Z Cyber Vault environment can be summarized as follows:
  - The Safeguarded Copy source volumes are accessible only from the recovered production systems because the production systems run in RegA.
  - The recovery volumes originate from a Safeguarded Copy backup and are defined online only to the IBM Z Cyber Vault recovery system OCONFIG.
  - The staging volumes are accessible to both the recovered production and IBM Z Cyber Vault recovery systems but are normally offline.
  - The persistent volumes are accessible from the IBM Z Cyber Vault recovery systems. These volumes store the IBM Z Cyber Vault assets and are where the validation process writes validation reports.

Table 4-1 on page 81 summarizes the IBM Z Cyber Vault IODF OCONFIG volumes that are defined as “Not configured,” “Offline,” or “Online” at IPL time.

Table 4-1 IODF OSCONFIG Matrix

OSCONFIG volumes	CYBER production sysplex	CYBER1 recovered sysplex on RC1
Safeguarded Copy source volumes	Online	Not configured
RC1	Not configured	Online
Staging volumes	Offline	Offline
Persistent volumes	Not configured	Online

- ▶ In the CV environment, the Recovery LPAR requires FICON connectivity to the Recovery Volumes. It also requires connectivity to the recovered sysplex CF.
- ▶ The recovery process copies the Cross-System Coupling Facility (XCF) CDSs from the production environment to the recovery environment as part of the restoration from the Safeguarded Copy backup. The CDSs must be preallocated for both the recovered production systems and the IBM Z Cyber Vault systems.
  - Use the same CDS names that are used in the production environment. Catalog these CDSs to the same VOLSERS that are used in the production environment. The difference is that the recovered sysplex CDSs are on different device addresses.
  - Another option is to use different CDS names and VOLSER values that are referenced by a COUPLExx member that differs from the one that is used by the recovered production systems.

## 4.2.1 Global Mirror secondary and Safeguarded Copy

The GDPS GM 2-Site (GM2Site) secondary volumes might be at an extended distance from the primary volumes. Each GM secondary must have a Safeguarded Copy backup that is defined, and each GM secondary requires a corresponding RC1 Recovery Volume.

**Note:** To perform multiple activities in a recovery environment, such as validation and forensic analysis, you need multiple sets of recovery volumes (RCx).

Replicated volumes, whether primary or secondary, have a definition for Safeguarded Copy backup capacity in the DS8000 system, as described in 2.3.2, “Safeguarded Copy prerequisites” on page 31. The physical and virtual capacity depends on the retention period, the capture or backup frequency, and the data change rate.

Define the virtual capacity multiplier for each volume. The smallest possible value is 1.5 times the original size. For more information, see “Safeguarded Copy sizing overview” on page 32.

Set up the Safeguarded Copy virtual capacity in the DS8000 system by using either the GUI or the `dsc1i` commands.

**Tip:** Define all the volumes and Safeguarded Copy space in the same storage extent pools of the DS8000 system.

The virtual capacity multiplier can be dynamically increased, but not decreased.

## 4.2.2 Recovery volumes

In the IBM Z Cyber Vault environment, you can define one or more sets of recovery volumes for RC1 for all the consistency groups (CGs). You can define these volumes as Extent Space Efficient (ESE) volumes.

A best practice for the recovery volumes is to maintain at least one set of recovery copy (RC) disks that provides capacity that is equivalent to the production environment. This configuration enables the migration of the recovery to a full copy by using the **StartCopy** command.

These OSCONFIGs are used to perform an IPL of the recovered systems in the IBM Z Cyber Vault environment. Both values are defined in the OSCONFIG of the GDPS LCP Manager, as shown in Table 4-2.

Table 4-2 IODF definitions

Type	LPARS	GDPS GM			Persistent	Staging	Safeguarded Copy	RC1
		Primary	Secondary	Journals				
					Simplex	Simplex		
PROD Live	PRD1 PRD2	ONLINE	OFFLINE	Not defined	OFFLINE	ONLINE	Not defined	
KG1	KSYSG	OFFLINE		Not defined			Not defined	
KR1	KSYSR		OFFLINE	Not defined			Not defined	
CV	CVPRD1 CVPRD2			Not defined	ONLINE	ONLINE	Not defined	ONLINE

## 4.2.3 Staging and persistent volumes

If you are doing a surgical recovery, you can use a set of staging volumes in the site that is hosting the GDPS LCP Manager.

### Staging volumes

A volume that contains datasets, tables, or any other objects that must be restored to the production environment cannot be copied to the staging volumes from RC1 when the recovery uses the No Copy option. When the recovery uses the Copy option, you can run a FlashCopy operation after the background copy is complete. When the recovery uses the No Copy option, you must run a Global Copy (GC) operation to copy the data to the staging volumes. Also, you can use a FlashCopy command when RC1 was established by using the **FCESTABLISH** script command.

In a GM virtual isolated environment with a physically isolated GDPS LCP Manager, the primary volumes are asynchronously mirrored to the secondary volumes, and you must copy the staging volumes back to the production site by using GC.

The staging volumes must not contain data that is generated by validation runs. Instead, the volumes are maintained for historical information or for subsequent processes because the staging volumes might be overwritten by the next validation cycle.

SMF records are saved to staging volumes in RegB.Site 1 in MSS0 in the OSCONFIG of the recovered systems.

## Persistent volumes

Historical files from a validation run are stored on persistent volumes. These volumes must not be overlaid during the recovery process. These volumes are online and available for each production validation run. In asynchronous mirroring solutions, these persistent volumes might need to be copied back to the production site by using GC.

**Note:** In a multi-site configuration, the staging and persistent volumes can be provisioned in the production region.

One of the persistent volumes contains a user catalog with a high-level qualifier (HLQ) of CYBERV, which is defined with a single alias that is connected to it. The CYBERV user catalog contains validation REXX and Job Control Language (JCL) commands, and it is where validation output is written. The user catalog is also connected to the production sysplex. Connecting to the production system enables operators to view the output that is produced during the validation process in the production environment.

Also define these persistent volumes to the recovered production systems.

### 4.2.4 Recovery for full data corruption

For a full recovery in an asynchronous solution, copy the validated RC1 volumes back to the RS1 production volumes by using the GDPS LCP Manager **RESTORE** scripts. In this configuration, it might be convenient to restart the application systems in the DR region from the RC1 volumes because an incremental copy might require more time than the restart.

In a GDPS GM environment, operators can incrementally recover the production environment from a tagged copy. The only data that is restored to the restored target is the delta between the current state of RS1 and the content of the tagged Safeguarded Copy capture, plus any additional changes that are made to RC1 after operators perform the **PREPARECOPY** operation. This process does not perform a full copy back to the restored target.

Figure 4-3 shows that the Safeguarded Copy backup (capture) that is used as part of the restoration process at the GM secondary site. This capture is recovered to a recovery set and then restored incrementally to the GM primary site.

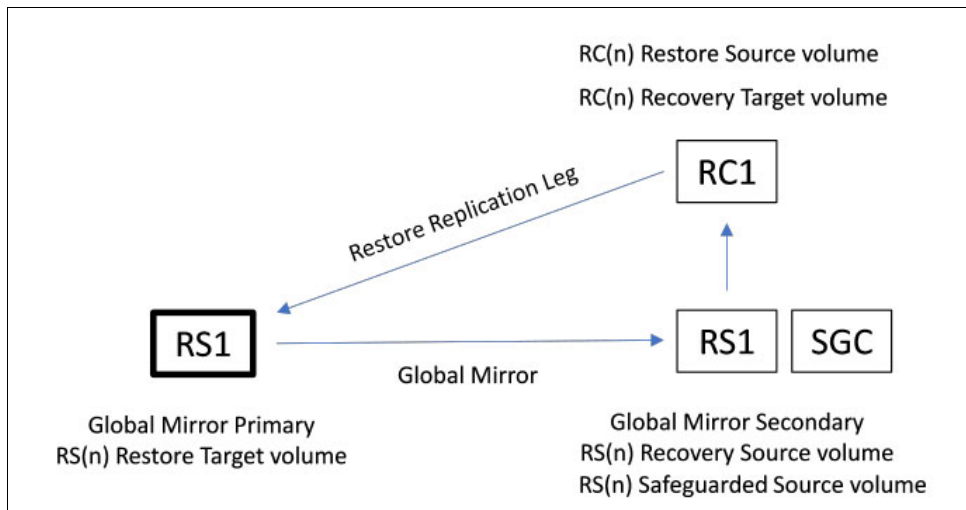


Figure 4-3 Flow of the incremental restore operations for a Global Mirror configuration

Before you run a **RESTORE** operation, you must stop the production systems in RegA.Site 1.

The steps to incrementally recover back to production are as follows:

1. Run **SYSPLEX STOP SYSTEM (system1, system2,...,systemn)** to stop all production systems at RegA.Site 1.
2. Run **GDASD STOP SECONDARY**.

Ensure that replication from the GM primary to the GM secondary is suspended. The GDPS LCP Manager rejects an attempt to run a **RESTORE** operation when the mirroring status is OK.

3. Tag the safeguarded capture to restore.

Identify the Safeguarded Copy capture that will be incrementally restored to production. You can accomplish this task by using the existing tag operations that are available in the Safeguarded Captures list, as shown in Figure 4-4.

VPCPLCSO Logical Corruption Protection Safeguarded Captures						
KSYSR			Capture Type: SAFEGUARD			
Management Profile: CV_TEST_SGC			Latest Capture: 2025/07/14.16:00:01			
Consistency Group: CVDASD			Captured/Expired: 168/0			
Replication Site: 1						
Actions: V olumes T ag U ntag						
Copy	Set	Capture	UTC TimeStamp	Volume Count	Flags	
	001	68752981	2025/07/14.16:00:01	224	NNNNN...	
I	001	68751B71	2025/07/14.15:00:01	224	NNYNN...	
	001	68750D61	2025/07/14.14:00:01	224	NNNNN...	
	001	6874FF51	2025/07/14.13:00:01	224	NNNNN...	
	001	6874F141	2025/07/14.12:00:01	224	NNNNN...	
	001	6874E331	2025/07/14.11:00:01	224	NNNNN...	
	001	6874D521	2025/07/14.10:00:01	224	NNNNN...	
	001	6874C711	2025/07/14.09:00:01	224	NNNNN...	
	001	6874B901	2025/07/14.08:00:01	224	NNNNN...	
	001	6874AAF1	2025/07/14.07:00:01	224	NNNNN...	
	001	68749CE1	2025/07/14.06:00:01	224	NNNNN...	
Sequence number 68751B71 in copy set SGC(1) is now tagged						
Command/Filter ==>						Row 1 of 168
F1=Help F3=Return F4=MonData F5=Refresh F6=Roll F7=Up F8=Down						

Figure 4-4 Tagging a capture

4. Run **LCP RESTORE PROFILE(CV\_TEST\_SGC) PREPARECOPY**.

Run the **LCP=RESTORE** script and specify the **PREPARECOPY** parameter. This parameter indicates that the request is not a standard LCP restore operation, but instead a request to prepare the replication environment and create an RC in readiness for a subsequent incremental restore from the tagged Safeguarded Copy capture. The **PREPARECOPY** operation performs all actions that are required to prepare the RCn Restore Source volumes, which are also the Recovery Target volumes, for an incremental copy back to the RS1 Restore Target volumes.

- Verify the data by using the RC set.

When the **RESTORE PREPARECOPY** script completes successfully, you can optionally perform an IPL from the recovered RC set to verify that the data in the tagged Safeguarded Copy capture is the data that you want to restore to production. Any updates that are made to the RC(*n*) volumes form part of the data that is then restored to the production volumes in step 6. If the updated data is not required, you must first run **RECOVER ENDCOPY** on the RC(*n*) RC and then repeat the **PREPARECOPY** operation.

- Run **LCP=RESTORE PROFILE(CV\_TEST\_SGC)**.

If the data that is recovered from the tagged Safeguarded Copy capture is not the required data, you can run **ENDCOPY** on the RC(*n*) RC by using the **RECOVERY ENDCOPY** script, tag a different capture, and then repeat the **RESTORE PREPARECOPY** script. When you are satisfied with the recovered data, you can run **LCP=RESTORE**, which incrementally restores the recovered copy to production.

### 4.3 Setting up Safeguarded Copy for IBM Z Cyber Vault Storage

Safeguarded Copy is the cornerstone of data protection in the IBM Z Cyber Vault solution. This section guides you through the configuration of Safeguarded Copy within the GDPS LCP framework, including how to define management profiles, capture intervals, retention policies, and monitoring schedules. It also introduces automation scripts and operational panels that streamline the creation and management of secure, immutable backups. These backups form the basis for reliable recovery and validation.

This section provides the information that you need to set up the Safeguarded Copy environment for the GDPS GM virtual air-gap solution. For more information about the GDPS LCP setup, see *GDPS Logical Corruption Protections Manager Installation and Customization Guide, ZG24-6763*.

**Note:** All IBM GDPS and LCP documentation are available only to users with a valid license.

The management profiles panel appears when you first enter the GDPS LCP Manager panel from the main GDPS panel by using the L option, as shown in Figure 4-5 (the VPCPMP00 panel).

VPCPMP00 Logical Corruption Protection Management Profiles									
KSYSR									
Actions:		Safeguard	Modify	Delete	Info				
		Captures	Backups	Volumes	Pools	Quiesce	Resume		
Management Profile			Volume Type	Copy Count	Captures Sets	Retention Tot	Flags Exp	QLMCRF...	
CVDASD.RS1									
—	CV_TEST_SGC	SGC	224	1	168	0	7D	0	NNYAE...
—	RECOVERY	RC	224	1	1	0	0M	0	
Command/Filter ==>								Row 1 of 3	
F1=Help	F2=DualCTL	F3=Return	F4=Schedule	F5=Refresh	F6=Roll	F7=Up			
F8=Down									

Figure 4-5 Logical Corruption Protection panel after creating the LCP Management Profile

In Figure 4-5 on page 85, you see one CG (CG CVDASD.RS1), the GM secondary disk, and the recovery set (RC) with 224 volumes.

To set up Safeguarded Copy in IBM Z Cyber Vault, complete the following steps:

1. Type S (SAFEGUARD) on the CVDASD.RS1 Consistency Group line.

To define a Safeguarded Copy management profile (see Figure 4-6), we used the following definitions:

- Management Profile: CV\_TEST\_SGC
- Copy Set:1
- Minimum Interval: MINUTE(50)
- Retention Period: DAY(7)
- Retention Minimum:168
- Automatic Release: EXPIRED
- Capture Interval: HOUR(1)
- Capture Start Time: 00:00:00
- Monitor Interval: HOUR(6)
- Monitor Start Time: IMMED
- Reservation Time: 0600
- Check In Time: 010
- CG Pause Time: N/A

```

VPCPMPMS          Modify an LCP Management Profile          LABK
      Capture Type: SAFEGUARD          SafeGuarded capture profile
Consistency Group: CVDASD          Consistency Group name
Replication Site: 1          Replication site number
Management Profile: CV_TEST_SGC          Profile name
      Copy Set: 1          Copy set assigned to this profile
Minimum Interval: MINUTE(50)          Minimum interval between captures
Retention Period: DAY(7)          Retention period for all captures
Retention Minimum: 168          Minimum number of captures to retain
Automatic Release: EXPIRED          Automatically release captures
Capture Interval: HOUR(1)          Enable automated captures
Capture Start Time: 00:00:00 (AUTO)          Start UTC time of capture schedule
Monitor Interval: HOUR(6)          Enable SGC monitoring
Monitor Start Time: IMMED          Start UTC time of monitor schedule
Enter NO to cancel or YES to proceed with the profile modification
Selection ==> YES          Page 1 of 2
F1=Help F3=Return F7=Up F8=Down
  
```

Figure 4-6 Safeguarded Copy management profile definitions in the first panel (VPCPMPMS)

2. Change only the default Reservation Time, Check In Time, and CG Pause Time after consulting with an IBM GDPS delivery professional. You can view these settings when you scroll down by using PF8 from the VPCPMPMS panels. Type YES and update the profile so that it matches the format that is shown in Figure 4-7 on page 87.

```

VPCPMPM2                Modify an LCP Management Profile                LABK
  Reservation Time: 0600                Maximum Reservation Scan elapsed
time
    Check In Time: 240                Maximum Check In elapsed time
    CG Pause Time: 0300                Maximum CG Pause elapsed time (GM)
Enter NO to cancel or YES to proceed with the profile modification
Selection ==>
F1=Help  F3=Return  F7=Up  F8=Down
Page 2 of 2

```

Figure 4-7 Safeguarded Copy Management profile definitions second panel (VPCPMPM2)

3. In this IBM Z Cyber Vault use case, we plan to automatically take one Safeguarded Copy backup (capture) every hour and keep it for 7 days, when it expires and is released. The retention minimum could be set to 168 (24 x 7) to retain at least all the Safeguarded Copy backups that are captured every hour for 7 days.
4. The Capture Interval and Monitor Interval that are defined in the Management profile create the Capture and Monitor Schedule based on the values that you provide. You can see the schedules by selecting PF4=Schedule.

Figure 4-8 shows the VPCPMPSS panel, which displays the Captures and Monitor Schedules.

```

VPCPMPSS                LCP Automation Schedule                LABK
  Next 7 days as of 2025/04/23.08:24:21 UTC
  Capture Engine Sites: RS1                Capture Scheduler: AUTOMATIC
  SGC Monitor Sites: RS1                Capture Delay Timeout: 00:10:00
  Global Offset: 00:00:00
Event      Event      Management      Consistency  Event      Flags
UTC TimeStamp  Type      Profile      Group      Window
2025/04/23.09:00:00  CAPTURE  CV_TEST_SGC  CVDASD.RS1  000.01:00:00  YAN..
2025/04/23.10:00:00  CAPTURE  CV_TEST_SGC  CVDASD.RS1  000.01:00:00  YAN..
2025/04/23.11:00:00  CAPTURE  CV_TEST_SGC  CVDASD.RS1  000.01:00:00  YAN..
2025/04/23.12:00:00  CAPTURE  CV_TEST_SGC  CVDASD.RS1  000.01:00:00  YAN..
2025/04/23.13:00:00  CAPTURE  CV_TEST_SGC  CVDASD.RS1  000.01:00:00  YAN..
2025/04/23.14:00:00  CAPTURE  CV_TEST_SGC  CVDASD.RS1  000.01:00:00  YAN..
2025/04/23.14:22:09  MONITOR  CV_TEST_SGC  CVDASD.RS1  000.06:00:00  YAN..
2025/04/23.15:00:00  CAPTURE  CV_TEST_SGC  CVDASD.RS1  000.01:00:00  YAN..
2025/04/23.16:00:00  CAPTURE  CV_TEST_SGC  CVDASD.RS1  000.01:00:00  YAN..
2025/04/23.17:00:00  CAPTURE  CV_TEST_SGC  CVDASD.RS1  000.01:00:00  YAN..
2025/04/23.18:00:00  CAPTURE  CV_TEST_SGC  CVDASD.RS1  000.01:00:00  YAN..
2025/04/23.19:00:00  CAPTURE  CV_TEST_SGC  CVDASD.RS1  000.01:00:00  YAN..
2025/04/23.20:00:00  CAPTURE  CV_TEST_SGC  CVDASD.RS1  000.01:00:00  YAN..
Command/Filter ==>
F1=Help  F3=Return  F4=Modify  F5=Refresh  F6=Roll  F7=Up  F8=Down  F9=Restart
Row 1 of 196

```

Figure 4-8 VPCPMPSS LCP Capture & Monitor Schedules

- If you want to run an ad hoc capture between the scheduled times, define a GDPS script like the one that is shown in Figure 4-9.

```

Script Name LCP_CAPTURE
Nbr  Script statement
1    COMM=CREATE A NEW SGC & EXPIRE OLD ONES
2    MESSAGE=CV0100I LCP_CAPTURE STARTED
3    LCP=CAPTURE PROFILE(CV_TEST_SGC)
4    MESSAGE=CV0101I LCP_CAPTURE ENDED

```

Figure 4-9 Sample script to take a Safeguarded Copy capture

**Note:** The MESSAGE statement in each script is needed only when you intend to use automation to check the process, as described in 4.4.1, “Automating Safeguarded Copy management” on page 90. Client-based message numbers are used to make each script unique. In this example, the script begins with an even number and ends with an incremented odd number. This structure enables automation to react to the GEO45I message and then filter to the script that is running.

The automation can react to the following messages:

```

GEO045I MESSAGE from GDPS SCRIPT: CV0100I LCP_CAPTURE STARTED. (GDPSR)
GEO045I MESSAGE from GDPS SCRIPT: CV0101I LCP_CAPTURE ENDED. (GDPSR)

```

Figure 4-10 shows the Status Display Facility (SDF) trace entries that the scheduled GDPS LCP capture generates.

```

Time      Text
05:00:00  PERFORMING AUTOMATED LCP CAPTURE FOR MANAGEMENT PROFILE CV_TEST_SGC
05:00:00  EXCLUSIVE LCP FLAGSET SERIALIZATION OBTAINED BY LCP CAPTURE
05:00:02  SEQUENCE NUMBER 6808AC11 HAS BEEN GENERATED FOR THIS SAFEGUARDED CAPTURE
05:00:02  GEO2772I SAFEGUARD CAPTURE PHASE 1 RESERVATION STARTED
05:00:02  GEO2773I SAFEGUARD CAPTURE PHASE 1 RESERVATION ENDED SUCCESSFULLY
05:00:02  GEO2772I SAFEGUARD CAPTURE PHASE 2 RESERVATION SCAN STARTED
05:00:05  GEO2773I SAFEGUARD CAPTURE PHASE 2 RESERVATION SCAN ENDED SUCCESSFULLY
05:00:05  GEO2949I SCHEDULING CGPAUSE FOR SESSION CVDASD
05:00:05  GEO2950I MONITORING COMPLETION OF CGPAUSE FOR SESSION CVDASD
05:00:06  GEO2951I SESSION CVDASD HAS BEEN SUCCESSFULLY CGPAUSED
05:00:06  GEO2772I SAFEGUARD CAPTURE PHASE 3 CHECKIN STARTED
05:00:07  GEO2957I THE CHECKIN TIME FOR THIS SAFEGUARDED CAPTURE WAS 0.156 SECONDS
05:00:07  GEO2773I SAFEGUARD CAPTURE PHASE 3 CHECKIN ENDED SUCCESSFULLY
05:00:07  GEO2949I SCHEDULING RESUME FOR SESSION CVDASD
05:00:09  MONITORING SESSION CVDASD FOR 'RUNNING' STATE
05:00:10  SESSION CVDASD WAS RETURNED TO 'RUNNING' STATE IN 3.025 SECONDS
05:00:10  GEO2951I SESSION CVDASD HAS BEEN SUCCESSFULLY RESUMED
05:00:10  RPO EXPOSURE FOR SAFEGUARD CAPTURE WAS 5.491 SECONDS
05:00:12  GEO2775I LCP SAFEGUARD CAPTURE ENDED SUCCESSFULLY
05:00:12  AUTORELEASE(EXPIRED) ENABLED FOR MANAGEMENT PROFILE CV_TEST_SGC
05:00:12  GEO2941I ZERO EXPIRED SAFEGUARD CAPTURES WERE FOUND IN MANAGEMENT PROFILE
CV_TEST_SGC
05:00:12  GEO2942W ZERO SAFEGUARD CAPTURES WERE RELEASED IN MANAGEMENT PROFILE CV_TEST_SGC
05:00:12  EXCLUSIVE LCP FLAGSET SERIALIZATION RELEASED

```

Figure 4-10 Sample SDF trace entries when capturing a Safeguarded Copy



8. Concurrently with one of the Safeguarded Copy backups (captures), a direct FlashCopy from RS2 to RC1 is taken. The validation process runs on this direct RC1 FlashCopy to help ensure that data corruption does not exist up to that point in time.

The reasons for creating a recovered Safeguarded Copy backup include the following items:

- Only a single Safeguarded Recovery relationship is supported. With this single relationship, you can perform a Safeguarded Recovery, if needed, without interrupting the current data validation process.
  - The read I/O on the recovered relationships goes through the Safeguarded Copy metadata for every read, which negatively affects the overall DS8000 workload and the speed of the validation process. When you take more backups after the recovery, the read performance can degrade over time because increasing amounts of metadata must be scanned for each read.
9. To take a direct FlashCopy from RS1 to RC1, use **GDASD \*ALL TESTCOPY CAPTURE**, as shown in Figure 4-13.

Nbr	Script statement
1	COMM=REFRESH RC1 WITH FC WITH TESTCOPY
2	GDASD=*ALL FCWITHDRAW RC1
3	GDASD=*ALL TESTCOPY CAPTURE RC1 NOCOPY

Figure 4-13 Sample Script CV\_RC1\_EST\_NOCOPY\_TESTCOPY

## 4.4 Setting up IBM Z Cyber Vault automation

Automation is essential to help ensure consistent, repeatable, and error-free operations within the IBM Z Cyber Vault environment. This section describes how to automate key tasks such as Safeguarded Copy management, IPL processes, and environment refreshes by using GDPS scripts. It emphasizes the importance of coordinating automation with business schedules and validation cycles and includes examples of scripts that you can customize for different recovery scenarios. Automation reduces manual intervention and enhances resilience to help achieve a streamlined operation.

### 4.4.1 Automating Safeguarded Copy management

Management of Safeguarded Copy backups, such as taking a new Safeguarded Copy capture or releasing expired copies, is performed by running a GDPS script. You can perform this task manually by using K-System or enable the automated capture interval in the Safeguarded Copy profile. In an IBM Z Cyber Vault environment, you must automate these tasks by using GDPS LCP Manager and schedule them for the times of day that are most appropriate for the business operations.

To do so, start a GDPS script either by using the LCP Capture Scheduler or by using batch JCL, which starts the GDPS RESTful API. Previous versions of GDPS also used IBM Z System Automation or IBM NetView timers.

You may create GDPS scripts to perform various functions:

<b>LCP_CAPTURE</b>	CREATE A NEW SGC AND EXPIRE OLD ONES
<b>LCP_FC_TO_RC1</b>	FLASHCOPY RSL TO RC1
<b>LCP_RELEASE</b>	RELEASE EXPIRED SGC SETS

## 4.4.2 Automating the IBM Z Cyber Vault IPL process

You can also automate the IPL process for the IBM Z Cyber Vault LPARs by using GDPS scripts, and this automation should be coordinated with automated Safeguarded Copy capture processes to quickly initiate the validation cycle when a Safeguarded Copy capture is complete. The IPL process should include deactivating and activating the IBM Z Cyber Vault LPARs, including the CF LPARs, loading the most recent Safeguarded Copy onto the IBM Z Cyber Vault recovery set of volumes (RC1) or, alternatively, running a **FLASHCOPY NOCOPY** operation, and loading the z/OS LPARs from the IBM Z Cyber Vault IPL volumes.

Here are the GDPS scripts to automate the IBM Z Cyber Vault environment IPL process:

<b>CV_CLEANUP_FC</b>	CLEANUP CV FC ENVIRONMENT
<b>CV_CLEANUP_SGC</b>	CLEANUP CV SGC ENVIRONMENT
<b>CV_RC1_REFRESH_IPL_FC</b>	REFRESH DATA FOR CV ENVIRONMENT & IPL CV WITH FC
<b>CV_RC1_REFRESH_IPL_SGC</b>	REFRESH DATA FOR CV ENVIRONMENT & IPL CV WITH SGC

Figure 4-14 shows examples of these scripts for **TESTCOPY**.

```
Nbr  Script statement
1    COMM=REFRESH DATA FOR CV ENVIRONMENT & IPL CV WITH FC
2    MESSAGE=CV0005I CV_RC1_REFRESH_IPL_FC STARTED
3    SYSPLEX=DEACTIVATE SYSTEM(CVPRD1,CVPRD2,CV-CF1)
4    GDASD=*ALL FCWITHDRAW RC1
5    GDASD=*ALL TESTCOPY CAPTURE RC1 NOCOPY
6    SYSPLEX=ACTIVATE SYSTEM(CV-CF1,CVPRD1,CVPRD2)
7    SYSPLEX=LOAD SYSTEM(CVPRD1)
8    USERPROC=CVWAIT 60
9    SYSPLEX=LOAD SYSTEM(CVPRD2)
10   MESSAGE=CV0006I CV_RC1_REFRESH_IPL_FC ENDED
```

Figure 4-14 Script example of automating the IBM Z Cyber Vault IPL process by using **TESTCOPY**

Figure 4-15 shows examples of these scripts for Safeguarded Copy.

```
Nbr  Script statement
1    COMM=REFRESH DATA FOR CV ENVIRONMENT & IPL CV WITH SGC
2    MESSAGE=CV0010I CV_RC1_REFRESH_IPL_SGC STARTED
3    SYSPLEX=DEACTIVATE SYSTEM(CVPRD1,CVPRD2,CV-CF1)
4    LCP=RECOVER PROFILE(CV_TEST_SGC) RC(1) ENDCOPY UNLOCK
5    LCP=RECOVER PROFILE(CV_TEST_SGC) RC(1)
6    SYSPLEX=ACTIVATE SYSTEM(CV-CF1,CVPRD1,CVPRD2)
7    SYSPLEX=LOAD SYSTEM(CVPRD1)
8    USERPROC=CVWAIT 180
9    SYSPLEX=LOAD SYSTEM(CVPRD2)
10   MESSAGE=CV0011I CV_RC1_REFRESH_IPL_SGC ENDED
```

Figure 4-15 Script example of automating the IBM Z Cyber Vault IPL process by using Safeguarded Copy

Both scripts do the same job but one populates the RC1 disk by using a direct FlashCopy from the source volumes, and the other one populates the RC1 disk by using a **RECOVER** operation from the Safeguarded Copy backups. You can see the differences in the breakdown of each step:

1. A comment about the script.
2. Unique message for automation purposes (indicates the start of the script).
3. Deactivates the IBM Z Cyber Vault environment.
4. Stops previous sessions.
  - a. For FC. Withdraws any previous FC sessions.
  - b. For Safeguarded Copy. Ends the relationship and unlocks the capture.
5. Starts the new replication session.
  - a. For FC. Runs **TESTCOPY CAPTURE** FC data to the RC1 disk.
  - b. For Safeguarded Copy. Runs **RECOVER** for the Safeguarded Copy data to the RC1 disk.
6. Activates the IBM Z Cyber Vault environment.
7. Loads the first LPAR into the IBM Z Cyber Vault environment.
8. A wait of 1 - 2 minutes for the environment to load and complete NIP.
9. Loads the second LPAR into the IBM Z Cyber Vault environment.
10. Unique message for automation purposes that indicates the end of the script.

Managing target system hardware for the IBM Z Cyber Vault LPARs usually requires full control, from **ACTIVATE** and IPL to **DEACTIVATE**. Both IBM GDPS and IBM System Automation Processor Operations (SA ProcOps) provide this capability and must be linked to BCPii so that SA ProcOps can control the LPARs and processors.

SA ProcOps is a best practice tool when GDPS is not available and you want to fully manage and control all target system hardware across multiple central processor complexes (CPCs) and physical sysplexes. If GDPS or SA ProcOps is not available, check with your automation product vendor to discover tools that can schedule and initiate an unattended automated IPL of a system.

When you automate the IBM Z Cyber Vault validation process, use FlashCopy to copy from the production source volumes to the RC1 volumes, and then start the IBM Z Cyber Vault environment and validate the current version of the production data. This approach helps improve performance when recovering data.

For more information about the end-to-end process for capturing, recovering, and restoring production data, see 2.2, “IBM Z Cyber Vault solution reference architecture” on page 20.

## 4.5 Preparing the IBM Z Cyber Vault environment for validation

Before you start the validation process, you must configure and isolate the IBM Z Cyber Vault environment. This section describes the setup of LPARs, CFs, network isolation, security controls, and IPL parameters. It also addresses considerations for parmlib, TCP/IP, VTAM, Multi-Factor Authentication (MFA), and dataset encryption. These preparations help ensure that the environment is secure, operationally ready, and capable of supporting comprehensive validation activities without affecting the production systems.

For more information about validation cadence and retention goals, see Chapter 3, “IBM Z Cyber Vault capabilities” on page 51.

### 4.5.1 Setting up CFs and LPARs for the IBM Z Cyber Vault environment

This section describes the implementation of the data corruption protection configuration by using GDPS K-sys and the available GDPS panels and scripts.

One requirement for an IBM Z Cyber Vault solution is to create LPARs for regular testing of performing an IPL from a recovered copy of the production system. To perform this testing, build an isolated environment to support the IBM Z Cyber Vault IPL and validation.

The size of the LPARs for the recovered application systems and CFs depends on the workload that you intend to run on them for the validation process and for eventual testing.

Define the new CFs and system LPARs for the IBM Z Cyber Vault environment in the GEOGROUP file, which results in a site table, as shown in Figure 4-16.

VPCSTD1										Standard Actions		KSYSR	
Actions: M Modify													
		L Load	X Reset	A Activate	D Deactivate								
Sysname	IND	Status	IPLtype	LPAR	IPLmode	Auto	L-addr	Loadparm					
— KSYSR	C	PRIMARY	NORMAL			YN							
— CVPRD1		MANUAL	RC1	FOSP59	RS2	NN	143A	1510C1M1					
— CVPRD2		MANUAL	RC1	FOSP5A	RS2	NN	1438	1510C1M1					
— CV-CF1		MANUAL	NORMAL	FOSP34		NN							
1 CPC Ops			3 LPAR View										
Selection ==>													
F1=Help F3=Return F6=Roll F7=Up F8=Down													

Figure 4-16 GDPS site table

In Figure 4-16, note the following items:

- ▶ KSYSR is the GDPS GM Recovery LPAR that runs the GDPS LCP Manager Code.
- ▶ CVPRD1 and CVPRD2 are the two recovery LPARS that can be loaded from the RC1 disks.
- ▶ CV-CF1 is the CF that you use in the IBM Z Cyber Vault environment recovered sysplex.



The recovered systems that are started from RC1 do not join the production sysplex, so XCF and Tivoli NetView for z/OS cannot communicate with the GDPS system or the GDPS LCP Manager, or receive acknowledgment of the completed IPL.

The initiation of the IPL of the IBM Z Cyber Vault environment is complete.

## 4.5.2 Network isolation considerations for the IBM Z Cyber Vault environment

Construct the IBM Z Cyber Vault environment in a logical bubble that has no intersection with the existing production environment. The IBM Z Cyber Vault environment should be in an isolated network to help ensure that the Safeguarded Copy backups cannot be accessed from the production environment or any other system. You can achieve this isolation by using dedicated Open Systems Adapter-Express (OSA-Express) features in the recovery system to provide physical isolation. Alternatively, if the IBM Z Cyber Vault environment is virtually isolated, you can use shared OSA-Express features to define logical separation by configuring a separate virtual LAN (VLAN) and IP subnetwork with optional firewalls.

In addition, you can protect both the IBM Z Cyber Vault environment and the production environment by using built-in z/OS Communications Server security features, such as the following items:

- ▶ IP filtering blocks all IP traffic that the system does not explicitly permit within its defined IP filter policy.
- ▶ Intrusion Detection Services (IDS) protect against various types of attacks on the system services.

From the router network perspective, the IBM Z Cyber Vault environment should be isolated from the existing production environments with a minimal number of access points. The IBM Z Cyber Vault environment should be accessible only through few predetermined IP addresses, ports, and protocols, and it must be protected by firewalls and routers that require VPN access and encrypted flows for communication with the IBM Z Cyber Vault environment (z/OS LPARs). The z/OS LPARs in the IBM Z Cyber Vault environment should be accessible only by using NAT IP addresses because the IBM Z Cyber Vault environment (z/OS LPARs) is activated by using the same configurations as production, and duplication of IP addresses on the network must be avoided.

## 4.5.3 Security considerations for the IBM Z Cyber Vault environment

The IBM Z Cyber Vault environment should be secured to limit access and to perform all tasks by using automation. From a hardware perspective, configure the IBM Z Cyber Vault LPARs in the HMC to enable the automation of IPLs.

### BCPii security

Set up BCPii permissions in the HMC so that GDPS may send **LOAD**, **RESET**, **DEACTIVE**, **ACTIVATE**, and **CAPACITY** commands to the LPAR objects on the IBM Z Support Element. To do so, complete the following steps:

1. Enable the SNMP APIs.
2. Enable Capacity Change API requests.
3. Set the Community name for SNMP Connections.
4. Set the Cross Partition Flags.
  - a. Select System Management.
  - b. Select CPC Operational Customization.
  - c. Select Change LPAR Security.
  - d. Select Cross Partition Authority.

5. Select System BCpii permissions.
6. Select LPAR BCpii permissions.

In the Change LPAR Security panel in the HMC, change BCpii Permission from the default Disabled to at a minimum of Receive to enable automatic hardware activation and the loading of the IBM Z Cyber Vault LPARs with BCpii commands that are sent from the production environment.

### RACF security

From a z/OS perspective, the RACF security database is the same as the production security database. All users have the same security entitlements that they have in the production environment. The security of the IBM Z Cyber Vault LPARs must be aligned with the enterprise security, audit, and compliance policies. Access to the IBM Z Cyber Vault environment is limited to users who have access to the isolated VLAN that is set up for this environment, and their activities must be audited.

The user ID for the data validation jobs must be set up with all required access so that it is ready for use when an IPL is performed in the IBM Z Cyber Vault environment. This user ID must be configured so that all required functions can be performed without failures that are related to resource access. Required access might include access to cryptographic functions and Db2 database administrator (DBA) authorities.

## 4.5.4 IODF and HCD definitions for the IBM Z Cyber Vault environment

GDPS LCP Manager creates a copy of the production storage environment. The IBM Z Cyber Vault environment requires unique CF resources that are separate from the production environment (see Table 4-3).

Table 4-3 CF accessibility in this environment

Coupling facilities	CYBER Production sysplex	CYBER1 Recovered sysplex on RC1	CYBER2 Recovered sysplex on RC2
Production CF (CF1)	Accessible	Not accessible	Not accessible
Recovery CF (CV-CF1)	Not accessible	Accessible	Accessible

To retain the information that is collected during the validation process, define a minimal number of persistent and staging volumes in the IBM Z Cyber Vault environment that are used across IPLs. You do not need to define a unique esoteric for these IBM Z Cyber Vault persistent and staging volumes.

Persistent and stage volumes may be defined as follows:

- ▶ A Systems Managed Storage (SMS) volume can be assigned to a storage pool under SMS control with other SMS managed volumes to manage datasets automatically.
- ▶ A non-SMS volume is not under the control of SMS and will not be assigned to a storage group or manage datasets automatically.

You can define persistent volumes (for example, ZCVP01-ZCVP04) as SMS managed volumes. This way, only the persistent volumes can be varied online to the IBM Z Cyber Vault environment (recovered sysplex).

## 4.5.5 Couple Data Sets for IBM Z Cyber Vault environment

Account for special considerations for the CDSs in the IBM Z Cyber Vault environment. In a HyperSwap setup, the CDSs of the production sysplex are not mirrored, but the LOGR CDSs are. Therefore, these CDSs are not included in the GDPS GEOPARM, and there is no corresponding Safeguarded Copy or RC1 and RC2 copy set.

To restart the sysplex from the RC1 or RC2 set of volumes, you need more CDSs. Two options are available:

- ▶ Preallocate the CDS on dedicated volumes that are defined in MSS0 and are online at IPL to the CYBER1 and CYBER2 OSCONFIG in the recovery region, except for the LOGR CDS.
- ▶ Preallocate all CDS except the LOGR CDS on mirrored volumes in Site 1. These CDSs are GDPS Metro mirrored to Site 2, included in the Safeguarded Copy captures, and included in the FlashCopy operation to RC1.

These CDSs are referenced in a specific COUPLExx member that is used when performing an IPL from RC1 or RC2.

**Tip:** Because the XCF CDS is newly allocated and never activated, it does not contain the name of the last CFRM policy that was used, so the CFRM policy name must be specified in the COUPLExx member.

Prime these pre-allocated CDSs with the current policies that are used in production by specifying the CDS name in the **IXCMIAPU** utility JCL for the ARM, CFRM, and SFM policy, but not for the WLM policy (see “Considerations for pre-allocated WLM CDS”).

**Important:** Update your day-to-day CDS management procedure to include the steps to refresh any changes to the production policies that are also in the pre-allocated CDS.

### Considerations for pre-allocated WLM CDS

The WLM policy cannot be primed in the CDS by using batch JCL. It must be imported as an exported partitioned dataset (PDS) from the 3270 or z/OSMF panels, and this process cannot be automated.

A default WLM policy is activated when you use preallocated non-primed CDS. In an IBM Z Cyber Vault environment, the default WLM policy is used.

If this configuration is sufficient to run the validation process, you can preallocate the CDS on mirrored volumes in the same way as the other CDS. Otherwise, the CDS should be preallocated on volumes in the same site as RC1, similar to the permanent volumes, and defined online to the recovery OSCONFIG. The policy (which does not need to match the production policy) must be imported during the first IPL and refreshed when changes are required.

### Considerations for preallocated XCF CDSs

From an XCF perspective, preallocated CDSs are sufficient. However, other sysplex consumers store information in the XCF CDSs, and these consumers expect this content to persist. For example, Virtual Storage Access Method (VSAM) RLS stores the name of the SHCDS datasets, and IBM MQ registers information about queue-sharing groups in the XCF CDSs.

Such information cannot be re-created in a preallocated XCF CDS until the CDS is actively used in the sysplex. After you perform an IPL of the first production system into the sysplex by using the preallocated XCF CDS, when the SMSVSAM address space is started, it issues a prompt because it does not find any SHCDS information in the XCF CDS.

When SMSVSAM is initialized in a sysplex environment that uses preallocated XCF CDSs, the information in the SHCDSs is lost when you use the **V SMS,SHCDS(dsn),NEW** command. This action deletes lost lock data and other important information in the SHCDS.

To remedy this issue, an option was added to the SMS **SHCDS** command: **V SMS,SHCDS(dsn),OLD**. This command enables an existing SHCDS to be made ACTIVE without formatting its contents.

Issue commands such as the following examples to define the SHCDS to be stored in the XCF CDS:

```
V SMS,SHCDS(PRIMARY.Z1SHC1),OLD
V SMS,SHCDS(SECONDARY.Z1SHC2),OLD
V SMS,SHCDS(SPARE.Z1SHCS),NEWSPARE
```

The interface between SMSVSAM and IBM Customer Information Control System (IBM CICS) enables CICS to understand and resolve these lost locks, and all that is required is to start the regions. These regions communicate with SMSVSAM and work to resolve any outstanding transactions based on the CICS logging.

Review the DFSMSdfp Storage Administration and the MVS System Commands manuals for the details of this command. Also, see [APAR OA58064](#), which describes this function in detail.

Before any workload that requires IBM MQ services is started, use the IBM 0MQ **CSQ5PQSG** utility to re-register the IBM MQ queue sharing groups in the XCF CDS.

For more information about this utility, see the IBM MQ publications for your release of IBM MQ. For IBM MQ 8, see [the queue-sharing group utility \(CSQ5PQSG\)](#).

Both the **SMSVSAM** commands and the IBM MQ JCL submission can be automated at IPL time from RC1 during the validation process.

## 4.5.6 IPL changes for the IBM Z Cyber Vault environment

When the IBM Z Cyber Vault environment is deployed, one or more LPARs are created on the IBM Z Cyber Vault recovery system. Because the IBM Z Cyber Vault LPAR that undergoes the IPL has the same system definitions as the main recovered LPAR, some changes must be made to the LOADPARM.

Based on the z/OS configuration in the IODF, another OS group ID is used to enable the IBM Z Cyber Vault LPARs to perform an IPL by using the RC1 or RCx volumes.

Update the IEASYMxx parameter so that changes can be made in the system symbol definitions, such as a CYBERID symbol. The IEASYMxx member that is used for the IPL can be set differently depending on the LOADPARM that is used.

For example, here are excerpts from different SYS1.IPLPARM LOADxx members:

- ▶ SYS1.IPLPARM(LOADZ1) (production environment)
  - SYSPARM (Z1,L)
  - IEASYM (Z1,SS,L)
- ▶ SYS1.IPLPARM(LOADC1) (IBM Z Cyber Vault environment)
  - SYSPARM (Z1,L)
  - IEASYM (C1,SS,L) (Note the difference.)

## 4.5.7 Parmlib changes for the IBM Z Cyber Vault environment

In the z/OS parmlib concatenation, as a best practice, use system symbols for simplifying the deployment of the IBM Z Cyber Vault environment. For example, as described in 4.5.6, “IPL changes for the IBM Z Cyber Vault environment” on page 98, two IEASYMxx members are used: one for the production system and one for the IBM Z Cyber Vault environment.

The &CYBERID system symbol is defined in each IEASYMxx member to represent the configuration of the environment that undergoes an IPL (production or IBM Z Cyber Vault). It is defined as follows:

```
&CYBERID          0 for the production system (in IEASYM00)
&CYBERID          R for the IBM Z Cyber Vault environment (in IEASYM0R)
```

The production and IBM Z Cyber Vault environment can use the same IEASYSxx member by using symbols to select different parmlib members during the IPL process to accommodate each environment:

- ▶ IEASYMxx

You may add symbolics that can be used as parameters in IEASYSxx. The CV-LPARs use the same IEASYSxx as the production counterpart. As a best practice, use the **SYSDEF LPARNAME()** keyword to set symbols for the CV-LPARs.

- ▶ IEASYSxx

These parameters, which may use different parmlib members.

- **ALLOC**

Cancel any job that encounters a volume enqueue or requires a volume that is not online.

- **AUTOR**

The IEASYSxx default for the **AUTOR** parameter is 00. If the source system has an **AUTOR00** in the system parmlib concatenation, then the **AUTOR=** parameter must reference multiple members.

A best practice is to code the IEASYSx **AUTOR** keyword as **AUTOR=(&AUTOR)**. This setting enables the &AUTOR symbol to point to one or multiple members in IEASYMxx without having to deal with parentheses.

- **COMMND**

Used to add or skip starting tasks on the CV-LPARs.

- **SMF**

For convenience of use, consider using MAN datasets to collect SMF records.

## 4.5.8 CFRM changes

The CFRM CDS that is allocated with the &CYBERID system symbol in the IBM Z Cyber Vault environment must be defined by using the same policy that is used for the production environment, and the PREFLIST of the structures must reference the CF in the IBM Z Cyber Vault environment. The hardware connectivity determines where the structures are physically allocated. Therefore, the CFs in the IBM Z Cyber Vault environment must not be connected to the production LPARs, and the production CFs must not be connected to the IBM Z Cyber Vault LPARs. The PREFLIST links to the unique CFs that are defined in the IODF. This configuration makes it possible to use a single CFRM policy that works in both environments.

As an example, a CFRM policy is defined with four CFs: two for the production sysplex and two for the IBM Z Cyber Vault environment. All four CFs are defined in the CFRM policy and are listed in the correct order in the PREFLIST, which is the ordered list of CFs where structures are preferably allocated, as shown in Example 4-1.

*Example 4-1 XCMIAPU SYSIN for the IBM Z Cyber Vault CFRM CDS*

---

```
DEFINE POLICY NAME(CFPOLC)
  CF NAME(CF1) DUMPSPACE(5000K) PARTITION(53) CPCID(00)
    TYPE(008561) MFG(IBM) PLANT(02) SEQUENCE(0000000xxxxx)
  CF NAME(CVCF1) DUMPSPACE(5000K) PARTITION(34) CPCID(00)
    TYPE(008561) MFG(IBM) PLANT(02) SEQUENCE(0000000xxxxx)
  STRUCTURE NAME(ISGLOCK) SIZE(33000K)
    REBUILDPERCENT(1)
  PREFLIST(CVCF1, CF1)
```

---

## 4.5.9 TCP/IP and VTAM configuration considerations

You can manage the TCP/IP and IBM VTAM configuration of the IBM Z Cyber Vault environment in one of the following ways, depending on the networking infrastructure.

If the IBM Z Cyber Vault environment is isolated and separated from the production network, and access into or out of it is controlled by using NAT, the production TCP/IP configurations can be activated in the IBM Z Cyber Vault environment as-is without modifications or the use of system symbols.

Alternatively, you can use a system symbol such as &CYBERID in the TCP/IP procedure to select either the production environment TCP/IP configurations or the IBM Z Cyber Vault environment TCP/IP configurations. With this method, two different configurations are available for isolation. Use the same symbol methodology with IPNODES to configure the resolver.

For VTAM, you must define the IBM Z Cyber Vault devices, such as **TRL**, **PORTNAME** definitions. If the IBM Z Cyber Vault environment is a mirror of production, no changes are required. If there are differences in the device configurations, you can use a system symbol such as &CYBERID to activate the appropriate devices for the specific environment.

## 4.5.10 IBM Multi-Factor Authentication considerations

If you use MFA, you must configure it in the same way that it is configured for the production environment, including considerations for out-of-band authentication, which requires a network configuration that enables access to the web login page. Some considerations are required for an IBM Z Cyber Vault environment to provide extra security, such as firewall configuration.

### 4.5.11 z/OS dataset encryption considerations

If you implement IBM Z dataset encryption in your environment, ensure that the following tasks are complete:

- ▶ The ICSF component is properly configured and active.
- ▶ The master keys are loaded into the crypto-card domains for the IBM Z Cyber Vault LPAR.
- ▶ The master keys are rotated in the production environment accounts for the IBM Z Cyber Vault environment.
- ▶ You need access to RACF security rules that are related to dataset encryption, including encryption controls, cryptographic functions, key labels, and datasets.

## 4.6 Establishing a validation framework

*Validation* is the process of confirming the integrity and usability of recovered data. It is a key capability of the IBM Z Cyber Vault solution. Data validation requirements and infrastructure design are unique to each environment and require investment to build, customize, and maintain a robust validation process. The goal is to automate and audit validation activities to help ensure that recovery copies are trustworthy and free from corruption.

This section provides a best practice validation framework and process flow with advice for conducting basic data validation. It also introduces Cyber Vault Data Validation for IBM Z Asset and the related implementation services that are delivered by IBM Technology Expert Labs (TEL), which provide a prebuilt framework of validation routines that is tailored and implemented according to individual configurations. For more information, see 4.7, “Cyber Vault Data Validation for IBM Z Asset” on page 109.

The three basic types of data validation that can be regularly performed in an IBM Z Cyber Vault environment include the following items:

- ▶ Type 1 validation (infrastructure validation)  
Verify that an LPAR can fully perform an IPL from the restored volumes and that required system tasks and middleware subsystems, such as CICS and Db2, initialize successfully.
- ▶ Type 2 validation (data structure validation)  
Run scripts and tools to verify the data structure integrity of key middleware databases, such as ADABAS, Db2, CICS, Information Management System (IMS), and IBM MQ, and critical system components, such as ICF catalogs, VSAM Volume Data Set (VVDS), PDS, PDSE, VSAM, and the RACF database to help ensure that the system is fully operational.
- ▶ Type 3 validation (data content validation)  
This validation helps ensure that application and user data that is stored in datasets, databases, or other subsystems is valid. This validation is the final step to help ensure that a copy is not corrupted by malware, ransomware, or any other type of intentional or unintentional data corruption and can be trusted. These validations can be performed by running database queries, batch programs, online transactions, and other application tests to verify that data is available. The application, database, and technology teams are responsible for providing the appropriate tools and scripts to run these tests, which are incorporated into the IBM Z Cyber Vault automation framework.

In Type 1 validation, much of the automation that is required to set up and perform an IPL on the IBM Z Cyber Vault environment is managed and tracked by the GDPS K-sys. However, GDPS can initiate the IPL only. It cannot manage or track activities after the IPL because the IBM Z Cyber Vault recovery system is in a network-isolated sysplex. Therefore, the validation framework that you implement should ensure that validation can be tracked, monitored, and reported on in an automated fashion within the IBM Z Cyber Vault environment.

The methods that are used to define jobs or processes for validation are flexible. Different vendor z/OS automation and z/OS batch scheduling products can be used to manage the validation process.

### 4.6.1 Creating a validation framework

The basic assumptions that are used to design a framework for performing data validation are listed in this section. The framework is described in terms of a general process flow that outlines the creation and running of a monitoring task that controls the sequencing, monitoring, and reporting of the results of validation jobs for a validation run.

### 4.6.2 Validation framework assumptions

Here are some considerations for designing a validation framework.

- ▶ The framework must have a high degree of automation in the construction and monitoring of IBM Z Cyber Vault data validation jobs in a closed network environment that is unattended.
- ▶ The framework must have a high degree of self-discovery for input to IBM Z Cyber Vault data structure validation for items that might naturally vary, such as new VSAM files that are added or deleted, during the construction of data structure validation jobs.
- ▶ The framework must identify and document errors that are encountered for IBM Z Cyber Vault validation jobs.
- ▶ The framework must allow for the exclusion of certain conditions during data validation checking. However, the excluded conditions must be documented and auditable.
- ▶ The framework must avoid any write to operator with reply (WTOR) messages from being generated by a validation job, such as a reference to a dataset that is cataloged to an offline volume.
- ▶ The framework must have a mechanism for capturing IBM Z Cyber Vault validation output and delivering validation process results to a monitored destination.
- ▶ The framework must maintain an auditable trail of all validation jobs that run.
- ▶ The framework must include a mechanism that allows data content (Type 3) validation jobs to fit into the IBM Z Cyber Vault validation framework.
- ▶ The framework must have the necessary security requirements that are documented for an automated process to perform validation, including access to the files to validate.

### 4.6.3 IBM Z Cyber Vault validation terminology

The following terminology is used when describing IBM Z Cyber Vault validation.

▶ Subsystem

Refers to a major category of data-management product or an OEM product.

▶ Data content validation

Refers to the validation of the content of application data that is constructed or maintained by a business. Validation of the functional product that is used for a specific purpose is called Type 3 validation and requires the owning application team to construct these jobs.

▶ Product

Refers to a specific product in use, which might be one of the following items:

- A product that is provided as part of the base z/OS operating system by IBM.
- A product that is available as open source for z/OS, either at no charge or with optional fee-based service plans, and provided by IBM or other vendors.
- A purchasable product for z/OS that is provided by IBM or other vendors.

▶ Metadata

Refers to information that describes the characteristics of data. For example, structural metadata describes how data is organized. Metadata does not describe data content. Many products use published metadata schemas or constructs to define, manage, or access the data content that is maintained by the product.

▶ Data structure validation

Refers to validations that can be performed on major products to detect anomalies in the metadata that is used or maintained by that product. Sometimes, the terms products and subsystems are used interchangeably, such as Db2, IBM IMS, and IBM CICS. Several products might be installed in an environment: Some products relate specifically to the management and enhanced functions of data-management subsystems, and others relate to enhanced functions that are not specifically tied to a subsystem. Generally, data-structure validation utilities are provided and documented by the vendor of the data-management subsystem. This type of validation is also referred to as Type 2 validation.

▶ Data validation

Refers to general validation concepts or activities, which might include data-structure validation and data-content validation. Data-content validation is performed by user-supplied jobs and processes that validate the content of their data. This type of validation is also referred to as Type 3 validation.

▶ Validation object

Refers to a single entity or object to be validated, for example, a single catalog, a dataset, or a unique Db2 table space.

▶ Validation build library

Refers to a partitioned library that contains the elements that are needed to run the validation processes and determines the following items:

- Which processes to run for each validation type (Types 1 - 3).
- Which processes to run to collect the validation output and where to send the output.
- The basic JOB card to use for the IBM Z Cyber Vault environment.
- Control cards that define the scope of the objects to be validated.

- JCL processes to invoke for the following reasons:
  - Submitting a simple validation.
  - Submitting a more complex validation process that requires written code, such as REXX or another language, which dynamically composes JCL in a runtime library and then submits the JCL.
  - Monitoring more complex validation that requires a specific sequencing of validation for a product (for data-structure validation) or for an application (for data-content validation).
- Written code members that are invoked by JCL in the validation build library. Depending on the language that is used, these members might be in a separate code-only library.
- ▶ Validation runtime library
 

Refers to a partitioned library that is newly built with each validation run. It contains JCL members that are dynamically created and submitted during the validation run.

#### 4.6.4 Constructing data structure validation jobs

The validation framework can be constructed by using various jobs, processes, and code that enable the following functions:

- ▶ Auto-generation of validation jobs
- ▶ Sequencing of validation jobs when required
- ▶ Monitoring of the jobs to help ensure completion
- ▶ Collecting validation information for reporting and historical archiving
- ▶ Reporting results by using email

Validation processes should be built to perform structural validation for a range of key system components and middleware objects.

- ▶ Integrated Data Cluster Access Method Services (IDCAMS) **DIAGNOSE** and **EXAMINE** activities for basic catalog structure (BCS) datasets
- ▶ BCS–VSAM Volume Dataset (BCS–VVDS) IDCAMS **DIAGNOSE** activities
- ▶ VSAM key sequenced dataset (KSDS) **EXAMINE** activities
- ▶ IBM MQ and Endeavor data structure validation
- ▶ RACF database data structure validation
- ▶ ADABAS database data structure validation
- ▶ Db2 database data structure validation for a stand-alone Db2 subsystem
- ▶ Db2 database data structure validation for a data sharing group Db2 subsystem
- ▶ IMS database data structure validation
- ▶ PDS and partitioned dataset extended (PDSE) data structure validation

There are three variations for building validation jobs based on the number of objects to be validated for a specific component, such as RACF, catalogs, or Db2, and whether the objects to be validated are provided from a predetermined input list or must be self-discovered at the time that the validation is run.

► Basic validation

This validation is used for a single object and is based on running a single job. The job writes the output to a predetermined error file name. If the return code is successful, the job deletes the error file. Otherwise, as a final step in the validation process, the error file is retained and noted when the results of all validation jobs are collected.

An example of basic validation is the data-structure validation job for the RACF database.

► Controlled validation

This type of validation uses an input file that contains a list of objects that can be included or excluded for validation. The validation process builds individual jobs for each object. The validation process verifies the existence of each object to prevent JCL failures. Nonexistent objects are reported as runtime errors in the final report. For objects that pass the verification process, all jobs are concatenated into a single validation runtime library member that is submitted at the end of the build process.

► Self-discovery

In this type of validation, the objects to be validated are automatically and dynamically discovered during each validation cycle. Objects of a specific object type are identified at run time, verified for existence and accessibility, and a validation job is built for each object and concatenated into a single validation runtime library member to be submitted at the conclusion of the build process.

An example of this variation is the process that is built to perform structural validation for VSAM KSDS datasets, where the list of objects to be validated varies between validation runs because objects are deleted or new objects are created in the production environment. In this example, the list of VSAM objects to be validated is reproduced with each validation cycle, and each VSAM object is verified for existence and accessibility, for example, whether it is cataloged properly and physically is on an online DASD volume, before a validation job is built and concatenated into the validation runtime library member and then submitted for running. Any discrepancies or anomalies that are detected during the self-discovery verification process and that prevent a validation job from being built must be noted and reported so that appropriate clean-up activities can be performed in the production environment.

When performing multiple types of validations for a product but they must be performed in a specific sequence, use the following steps:

1. Define separate, unique job prefixes and validation runtime library members for each sequence of jobs to run.
2. Construct a validation monitoring job to help ensure that each sequence is completed before triggering the next sequence.

#### **4.6.5 IBM Z Cyber Vault job naming conventions**

When constructing validation jobs, consideration proper job naming conventions because this best practice is important when monitoring job progress and the status of the overall validation cycle. Validation jobs should be assigned a two-character job name prefix that is unique and used exclusively within the IBM Z Cyber Vault environment. No other jobs or tasks should use this prefix or be inadvertently included in the monitoring process, such as CV.

Also, validation job names should contain an identifier that denotes the data-structure type of validation being performed, for example, C for catalog **DIAGNOSE** and **EXAMINE** validation jobs or P for PDS and PDSE partitioned datasets. Because there can be many thousands of objects of a specific data type and many thousands of associated validation jobs, the job names should also include an incremental value that uniquely associates the job name with the validation object. For example, job CVC00007 might be used to perform data-structure validation, by using IDCAMS **DIAGNOSE** and **EXAMINE**, for the seventh entry in a list of ICF user catalogs to be validated.

#### **4.6.6 IBM Z Cyber Vault datasets**

The validation framework requires several permanent datasets to support the validation process, and each running of the validation cycle produces more datasets, many of which might need to be retained across validation cycles. As with any application, plan for dataset naming conventions and help ensure that adequate DASD space is available for both temporary and permanent datasets. Persistent volumes that are defined in the IBM Z Cyber Vault environment should be used to store the results of validation cycles, reports of detected errors, detailed and summary reports of validated objects, and historical backups of important validation cycle artifacts. Audit trail data, such as SMF data that is produced during the validation cycle, might also need to be stored on the staging volumes to allow that data to be transmitted to and processed on the production LPAR.

#### **4.6.7 IBM Z Cyber Vault validation monitoring task**

The key to creating an automated and unattended validation process is to use a monitoring task to control the sequence of validation processes and monitor the completion of one type of validation before moving on to the next type. For example, during Type 2 data-structure validation, automated self-discovery identifies and creates a list of objects to be validated and dynamically builds the JCL by using the job-naming conventions that are described in 4.6.5, “IBM Z Cyber Vault job naming conventions” on page 105. The generated jobs are submitted in batch and monitored for completion before the framework proceeds to create and submit the Type 3 data-content validation jobs.

Figure 4-18 on page 107 shows the simple sequence of steps that is performed by this monitoring task. You can build dependencies so that Type 3 data content validation can be submitted when a subset of Type 2 data-structure validation jobs completes.

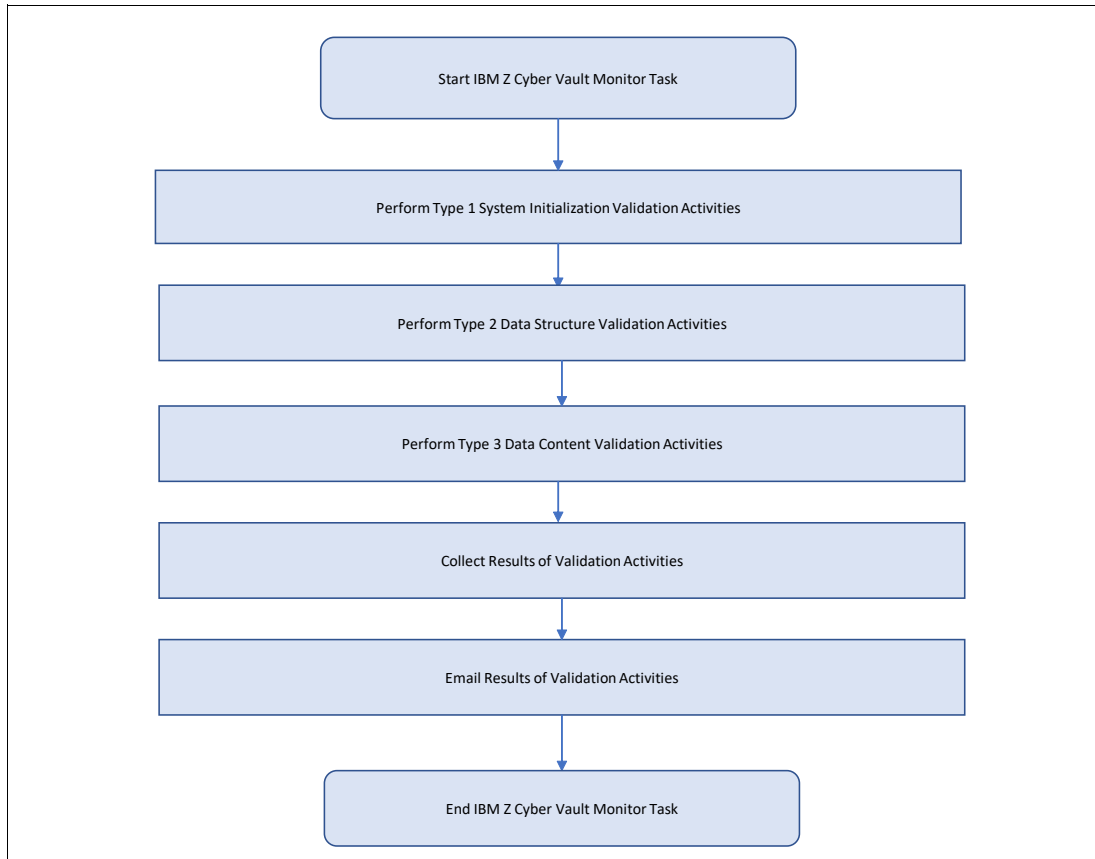


Figure 4-18 Process flow for the IBM Z Cyber Vault monitoring task

You can manage the order and monitoring of validation jobs by using a combination of the batch scheduling product, the automation product, available JES facilities, or applications that are written to use the SDSF REXX APIs to query the output of jobs or tasks and monitor their status in the JES queues. The collection of validation results is performed after all submitted validation activities are completed.

After the results of validation activities are collected, communicate those results to the correct personnel in a timely manner. This communication depends on-site requirements. As a best practice, communicate the validation cycle results through email by using a collective email address or by sending the summary file to a staging volume, where it can be processed on a production LPAR.

The processes that are involved in the IBM Z Cyber Vault monitoring tasks are described in the following sections.

### Type 1 (infrastructure) validation

Type 1 infrastructure validation includes verifying a successful IPL of the IBM Z Cyber Vault systems. This verification includes confirming that critical system tasks and interactive subsystems are active and, when required, confirming that tasks issued the expected initialization-complete type messages. Each system might have its own unique list of required tasks and associated messages.

As a basic process, Type 1 validation routines perform the following actions:

- ▶ Process a list of task names and, optionally, their respective initialization-successful type messages.
- ▶ Verify that each task is active.
- ▶ Optionally, verify that the expected initialization-successful type message was issued by each task.
- ▶ Summarize the results of the verification.

The IBM Z Cyber Vault monitoring task initiates and monitors the Type 1 validation activities and, when they are complete, proceeds with Type 2 validation.

### **Type 2 (data structure) validation**

Type 2 data-structure validation is normally the most time-consuming component of the overall validation cycle. The combination of the type and volume of objects, such as Db2, VSAM, and other subsystems, the wanted frequency of the validation cycle, and the resources, such as CPU and memory, that are available to the IBM Z Cyber Vault environment can result in thousands, or even tens of thousands, of validation jobs that must be generated, submitted, and monitored for completion. The results of each job must be collected for later reporting.

For each object type, such as Db2, VSAM, IBM MQ, PDS, IMS, and others, the Type 2 validation process can consist of the following activities:

- ▶ Runs self-discovery routines to generate a list of all objects that are eligible for validation, such as all VSAM KSDS clusters, or processing a static list of objects, such as a single RACF database.
- ▶ Dynamically builds the JCL job to run the appropriate validation utility for each object.
- ▶ Submits and monitors each validation job for completion.
- ▶ Summarizes the results of the validation.

As with Type 1 validation, the IBM Z Cyber Vault monitoring task is responsible for initiating and monitoring all Type 2 activities and then proceeding with Type 3 activities when the Type 2 validation completes.

### **Type 3 (data content) validation**

Type 3 data-content validation is specific to an application, and constructing validation jobs requires in-depth knowledge of the application and its internal data and associated relationships. However, the general process for Type 3 validation resembles Type 2 validation process regarding the following activities:

- ▶ Builds the JCL job to run the appropriate validation utility for each application object.
- ▶ Submits and monitors each validation job for completion.
- ▶ Summarizes the results of the validation.

Completing of the Type 3 validation activities triggers the IBM Z Cyber Vault monitoring task to continue with the collection and reporting of all types of validation activities.

## Collecting and reporting the results of validation activities

At its most basic level, each validation job or process, regardless of type, should provide a detailed record of its processing. This record should include the identity of the object, the results of its validation, whether the validation is successful or unsuccessful, and any error messages that are generated during the validation process. At a higher level, the results of the overall validation cycle, such as the number of objects that are validated for each component, the number of validation errors that are identified, and statistical information about run times, must be recorded. An audit trail of any objects that are bypassed or excluded from validation activities must also be retained. All such information must be archived to persistent storage for historical purposes at the end of each validation cycle.

The IBM Z Cyber Vault monitoring task concludes its processing by collecting and summarizing the results of all validation activities and producing a final report, which is delivered to the personnel who are responsible for monitoring the IBM Z Cyber Vault environment.

## 4.7 Cyber Vault Data Validation for IBM Z Asset

The IBM Z Cyber Vault validation process is complex. Developing, implementing, and automating a comprehensive validation framework can be a large task and require considerable time to complete. To simplify and accelerate the implementation of a validation framework, IBM offers the Cyber Vault Data Validation for IBM Z Asset.

The Cyber Vault Data Validation for IBM Z Asset is a data validation framework with supporting software for the IBM Z Cyber Vault solution. It is designed to automate and streamline the process of performing all types of validation in an IBM Z Cyber Vault environment and to report validation cycle results.

The asset is licensed software that is available through IBM Technology Expert Labs (TEL) as part of an IBM Z Cyber Vault services engagement. It cannot be purchased or downloaded from IBM Shopz and is not part of any software catalog. The asset is priced and licensed per IBM Z sysplex because it validates a production sysplex environment. The asset requires IBM Storage (DS8000), JES2, SDSF, and REXX, and it provides built-in support for validation of the following subsystems and data types: RACF, ICF catalogs, VSAM, PDS, PDSE, CICS/TS, IBM MQ, Db2, and IMS.

The asset can be extended to include more subsystem types and client-customized validation processes.

For more information, contact [systems-expert-labs](https://www.ibm.com/support/entry.do?entry=ibmcom+systems-expert-labs).





# IBM Z Cyber Vault: Building on the foundations

This chapter builds on the foundational concepts of the IBM Z Cyber Vault solution that were introduced earlier in this publication. It highlights advanced capabilities, operational best practices, and future enhancements that are designed to strengthen cyber resiliency. Key focus areas include integration with offensive security (OffSec) techniques, improved monitoring and validation tools, and recovery processes at the subsystem level.

More resources and references are provided throughout this chapter to guide deeper exploration of specific topics.

The following topics are covered in this chapter:

- ▶ 5.1, “The evolution of the IBM Z Cyber Vault solution” on page 112
- ▶ 5.2, “Db2 for z/OS subsystem rollforward recovery” on page 114
- ▶ 5.3, “Advanced cyber resilience for IBM Z” on page 117
- ▶ 5.4, “Operational practices for enhancing integrity” on page 120

## 5.1 The evolution of the IBM Z Cyber Vault solution

The IBM Z Cyber Vault solution is part of an ongoing innovation journey. Although it already delivers value today, designers created it with extensibility in mind. IBM plans future enhancements to further strengthen its capabilities. These updates continue to align with emerging industry needs and evolving client requirements. For example, see the Geographically Dispersed Parallel Sysplex (IBM GDPS) Statements of Direction in [GDPS 4.8 What's New](#).

### 5.1.1 Roadmap for the IBM Z Cyber Vault solution

Before you review specific technologies, this section introduces general trends. Later sections examine how these trends apply to the IBM Z Cyber Vault solution. This roadmap groups the focus areas into three main domains:

- ▶ IBM Z Cyber Vault storage
  - Covers data protection, backup capabilities, and secure storage mechanisms.
- ▶ IBM Z Cyber Vault automation
  - Focuses on orchestration, recovery workflows, and automated threat response.
- ▶ IBM Z Cyber Vault environment
  - Encompasses the infrastructure, configurations, and operational context that are needed to support IBM Z Cyber Vault.

A key focus in the IBM Z Cyber Vault storage domain is the number of Safeguarded Copy backups. The previous limit was 500, which allowed for 83 days of backups at 4-hour intervals. However, with growing demand for shorter intervals, such as 1 hour, this configuration covers only about 21 days. IBM recently increased the maximum limit for Safeguarded Copy backups to 1,024 in IBM DS8000 R10.1. By raising the backup limit, organizations can capture backups more frequently or retain them for longer durations.

Another focus area in the IBM Z Cyber Vault storage domain is extending inline data corruption detection ([available in IBM Storage FlashSystem](#)) to IBM DS8000 through Flash Core Modules 4 (FCM4). This technology uses machine learning to continuously analyze I/O statistics and detect anomalies at the block level. Similar to the planned integration that is described in 5.1.2, “IBM Threat Detection for z/OS: Integration with IBM Z Cyber Vault” on page 113, it adds another layer of intelligence for identifying unusual activity. When linked to the automation domain, it enables policy-based automated actions in response to detected threats.

Within the IBM Z Cyber Vault automation domain, The stand-alone Logical Corruption Protection (LCP) Manager that was recently announced extends GDPS LCP Manager integration to environments that use IBM Copy Services Manager (CSM) in addition to GDPS. Although GDPS remains the leading solution for high-end enterprise high availability and disaster recovery (HADR) deployments, many IBM Z installations rely on CSM to meet their HADR requirements. Enabling LCP Manager functions in CSM-based infrastructures facilitates broader adoption of IBM Z Cyber Vault capabilities without requiring a migration away from existing HADR tools.

Another area of development for the IBM Z Cyber Vault automation domain is to introduce the current validation services offering into GDPS LCP Manager. This change enables more clients to access this capability in a way that is fully supported by IBM.

## 5.1.2 IBM Threat Detection for z/OS: Integration with IBM Z Cyber Vault

IBM Threat Detection for z/OS (TDz) is an AI-powered solution that continuously monitors z/OS systems for suspicious or unauthorized data access. When TDz detects a potential threat, TDz generates a detailed notification that includes the affected datasets, the user ID involved, and the type of access, such as read or write.

The IBM Z Cyber Vault solution is designed to act on TDz threat notifications by triggering a set of user-defined policies. GDPS runs these policies and uses its REST API to notify GDPS LCP Manager. When LCP Manager receives the alert, it activates IBM Z Cyber Vault policies, which are automated actions that are tailored to the nature and severity of the threat.

These policies are also foundational for future enhancements, including IBM Z Cyber Vault validation updates and expanded threat detection capabilities. Each policy has four distinct actions that can run automatically when a threat is identified, as shown in Table 5-1.

Table 5-1 Distinct actions

Action	Description
ALERTS	Issues an Status Display Facility (SDF) alert or WTO message to signal the event.
CAPTURE	Initiates a new data capture for forensic analysis.
QUIESCE	Pauses all LCP capture and release operations.
SUSPEND	Halts replication into the vault to prevent DS8000 rollofs due to space exhaustion.

These actions provide a flexible and scalable response mechanism that allows organizations to tailor threat response strategies. Within GDPS LCP management profiles, administrators can configure which IBM Z Cyber Vault policy to run based on the type of event, such as read or write.

When TDz detects a threat, the following actions occur:

1. LCP Manager is notified through the REST API.
2. Each management profile is evaluated to determine the appropriate response.
3. The associated IBM Z Cyber Vault policy runs automatically and applies the defined actions that are relevant to the event type and profile configuration.

This integration of TDz and IBM Z Cyber Vault automation (LCP Manager) helps ensure a seamless and intelligent response to threats and enhances the overall resilience of the IBM Z environment.

### 5.1.3 Enhancing IBM Z Backup Resiliency for surgical recovery

IBM Z Backup Resiliency (IZBR) provides critical capabilities for forensic analysis and targeted recovery when data corruption is detected. In cases where the corruption is limited in scope, IZBR enables surgical recovery, which restores only the affected datasets rather than the entire system.

To improve the effectiveness of surgical recovery operations, IBM is planning several key enhancements.

- ▶ Real-time backup inventory

IZBR maintains a real-time inventory of available backups within IBM Z Cyber Vault. This enhancement includes integration with GDPS LCP Manager, which sends real-time notifications to IZBR when events occur that affect the state of backup captures. These events include capture creation, capture expiration and release, and DS8000 roll-off events to free Safeguarded Copy backup capacity. With this real-time visibility, IZBR can present users with a curated list of viable recovery points, which reduces decision time and improves accuracy.

- ▶ Increased automation of recovery operations

To streamline the recovery process, IZBR introduces greater automation. After a user selects the datasets to recover and the corresponding backup, the system automatically performs the following actions:

- Retrieves the selected datasets from the backup.
- Copies them to staging volumes.
- Completes the operation without requiring additional user intervention.

These enhancements aim to reduce manual effort, minimize recovery time, and improve reliability, especially during high-pressure scenarios such as active cyber incidents.

## 5.2 Db2 for z/OS subsystem rollforward recovery

By default, a production z/OS logical partition (LPAR) can be restored only to a previously created Safeguarded Copy backup. Because these copies are typically created every x hour, restoring a production environment to the point-in-time of a Safeguarded Copy backup almost always results in the loss of valid transactions and updates to Db2 for z/OS, assuming that a corruption event occurred between the creation of two Safeguarded Copy backups.

For example, if Safeguarded Copy backups are created at 9:00 AM and midnight and a corruption event occurs at 11:52 AM, restoring the production environment to the 9:00 AM Safeguarded Copy backup does not account for valid transactions that occurred between 9:00 AM and 11:52 AM. All transactions and updates that were made between 9:00 AM and 11:52 AM in Db2 for z/OS are lost (see Figure 5-1 on page 115).




Figure 5-1 Capture interval and corruption event detection

The goal of a Db2 for z/OS subsystem rollforward recovery is to avoid losing transactions that successfully completed between the creation of a Safeguarded Copy backup and the time a corruption event starts.

**Note:** Only the customer can determine the exact point-in-time to which a Db2 for z/OS subsystem should be recovered based on forensic analysis. For more information, see 3.4.2, “Forensic analysis” on page 62.

To reduce the data loss window through a Db2 for z/OS rollforward recovery from multiple hours to minutes or less, Db2 for z/OS rollforward recovery requires a collection of the following assets:

- ▶ Copies of the bootstrap dataset
- ▶ Copies of new archive logs
- ▶ Copies of image copies that were created after LOG NO events, such as **LOAD REPLACE LOG NO**

**Note:** If image copies are not created after LOG NO events, affected objects might not be recoverable. To protect these assets from potential corruption during a cyber event, store these data assets on immutable storage, similar to Safeguarded Copy backups.

If the data assets are available in an IBM Z Cyber Vault recovery LPAR, the following steps support a rollforward recovery of Db2 for z/OS.

- ▶ Preparation at the production system.
  - Set the archive log frequency based on the asset collection interval.
  - Periodically save recovery assets to immutable storage that is accessible at the recovery system.
    - Archive logs.
    - Bootstrap dataset.
    - Image copies for table spaces and indexes that were created after LOG NO events.
- ▶ Preparation at the Cyber Vault recovery system after a Safeguarded Copy is restored.
  - From immutable storage, retrieve the archive logs, BSDS, and image copies that are more recent than the restored Safeguarded Copy.
  - For each Db2 member, do the following actions:
    - Replace the BSDS with the most recent BSDS.
    - Run **DSN1LOGP** on the most recent archive log dataset.

- Identify the smallest Log Record Sequence Number (LRSN) across the members from the DSN1LOGP output, which is the ending point for the log apply.
- Adjust the SYSPITR system recovery point.
- ▶ Recovery rollforward at the Cyber Vault recovery system.
  - Run a conditional restart with **SYSPITR**.
  - Run **RESTORE SYSTEM LOGONLY**.
  - Stop and start Db2.
  - Stop any outstanding utilities.
  - Run the **RECOVER** utility on table spaces and indexes in RECOVER-pending status.
  - Run **REBUILD INDEX** on indexes in REBUILD-pending status.
  - Validate data.

Because collecting these additional recovery assets and running these steps require manual effort and can be error-prone, especially in high-pressure situations such as an ongoing cyberattack, IBM has a solution that reduces the complexity of these tasks to support the fastest possible recovery of the production environment.

## 5.2.1 IBM Db2 Recovery Expert Pro for z/OS: Subsystem rollforward recovery

[IBM Db2 Recovery Expert Pro for z/OS V1.1](#) supports Db2 for z/OS rollforward recovery in an IBM Z Cyber Vault recovery environment. With this solution, you can collect the required recovery assets at pre-defined intervals (in minutes) and store them on immutable storage.

**Note:** The solution requires that image copies that are created after LOG NO events are recorded in SYSIBM.SYSCOPY. Multiple architectural options are available to store data assets on immutable storage to support individual customer requirements. Because this area continues to evolve, contact your IBM representative to understand the latest available supported architectural options.

In Figure 5-1 on page 115, assuming that recovery assets are collected and stored on immutable storage every 4 minutes, Db2 for z/OS can be rolled forward to 11:48 AM. This change reduces the Recovery Point Objective (RPO) to 4 minutes before the cyber event occurs, instead of 2 hours and 52 minutes when using the traditional IBM Z Cyber Vault recovery approach. The solution fully automates all the manual steps that are outlined in 3.4.3, “Recovery process” on page 66, which reduces the time that is needed to resume normal business operations.

**Note:** Although Db2 is used in the following example, similar techniques might apply to other subsystems.

## 5.3 Advanced cyber resilience for IBM Z

As cyberthreats grow in sophistication and frequency, IBM Z environments must evolve beyond traditional security and recovery practices. Two key strategies, offline backup and offensive security (OffSec), play complementary roles in strengthening cyber resilience.

Together, offline backup and OffSec form a defense-in-depth approach:

- ▶ Offline backups help ensure data survivability and integrity.
- ▶ OffSec helps ensure system hardening and proactive threat mitigation.

By integrating both into the IBM Z Cyber Vault environment, organizations can better prepare for, withstand, and recover from cyber incidents, whether they are caused by external attacks or internal failures.

### 5.3.1 Offline backups

Offline backup is already a feature of the IBM Z Cyber Vault solution and is implemented after the initial setup phase. This process moves primary data to a storage tier that is less accessible if there is data corruption.

Tape and virtual tape are the most commonly used technologies for offline backups, although alternatives such as object storage can also be considered for this purpose.

- ▶ Tape: Best for long-term, infrequent access and high-security needs.
- ▶ Virtual tape: Supports faster recovery and integration with existing backup software.
- ▶ Object storage: Suitable for scalable, cloud-native environments with frequent access requirements.

Unlike Safeguarded Copy backups on the DS8000, which keep data within the same storage system, offline backups transfer data to a separate storage tier. This added step increases the time and complexity that are required to retrieve and restore data if it is needed.

Some key attributes of this capability that are wanted for this role include the following items:

- ▶ Data that is held is less accessible for modification, deletion, or disruption from a compromised environment. This capability can include logical and physical separation to address requirements.
- ▶ Administrative separation of management, storage administration roles, and access profiles layers provide more safeguards against destructive actions that are performed through administrative operations.
- ▶ Ability to use technology-specific controls to protect data.
  - For TS7700, this capability includes logical WORM (LWORM), LWORM retention, expire hold, and related functions.
  - For object storage technologies, this capability includes hold policies on objects, retention controls, and related mechanisms.

The best offline backup approach and technologies depend on the requirements of the environment and the operational processes that are needed by the clients. There is no single best approach; the ideal approach is tailored to each installation.

Offline backup uses the traditional backup and restore model in which data is transferred from the primary disk system to the auxiliary storage to protect it. Data is transferred back from the auxiliary storage to the primary disk system to restore it. Safeguarded Copy backups on disk are always faster in the protect and restore operations because the data is immediately available rather than requiring completion of data transfers.

Offline backup should be considered to extend the depth of protection beyond what is practical to hold on primary disk for cost or capacity reasons, with the understanding that this capability focuses on retrieving specific, targeted data and that restoring from it should be a last resort.

At the time of writing, offline backup is performed by host tools such as ImageCopy and DFSMSdss to coordinate the backup and restore processes. The transport mechanism can use transparent cloud tiering (TCT) for direct storage-to-storage transfers as objects or FICON data transfer to a virtual tape device.

Storage systems such as the IBM TS7700 can further copy or protect data with tiering behind themselves, such as physical tape or object storage transfers, to provide more isolation and protection options.

## 5.3.2 Offensive security

*OffSec*, also known as ethical hacking or penetration testing, involves simulating cyberattacks to identify vulnerabilities and weaknesses in a system before malicious actors can exploit them. This proactive approach to security is crucial for IBM Z environments because they often handle critical workloads and sensitive data.

### Why offensive security is important for IBM Z

Although IBM Z systems are renowned for their reliability and security, the evolving threat landscape and increasing connectivity of mainframes make OffSec an indispensable aspect of a robust security strategy. Some major benefits of this strategy include the following items:

- ▶ Identifies unknown vulnerabilities.

OffSec helps uncover vulnerabilities that traditional security practices might not identify.

- ▶ Mitigates threats proactively.

By mimicking real-world attacks, penetration tests and red teaming exercises allow organizations to identify and address weaknesses before attackers can exploit them.

- ▶ Strengthens defenses.

The information that is gathered from OffSec operations provides insights that enhance defensive security measures and improve incident response capabilities.

- ▶ Helps ensure compliance.

Many regulatory frameworks and standards, such as Payment Card Industry Data Security Standard (PCI DSS) and the International Organization for Standardization (ISO), require regular penetration testing to demonstrate compliance and reduce the risk of non-compliance penalties.

- ▶ Addresses misconceptions

Despite their inherent security features, mainframes are not “bulletproof”. A misconception about mainframe security can lead organizations to overlook the need for rigorous security measures, including penetration testing. This oversight can expose mainframes to threats such as insider attacks, code-based vulnerabilities, human error, and ransomware.

## Key elements of offensive security for IBM Z

Several methods and tools are employed in OffSec assessments for IBM Z environments:

- ▶ Vulnerability scanning  
Automated processes scan for known vulnerabilities in mainframe systems, applications, and configurations.
- ▶ Penetration testing  
Ethical hackers simulate real-world attacks to identify exploitable vulnerabilities and assess the effectiveness of security controls. This activity includes examining configurations, applications, and access controls within the mainframe environment.
- ▶ Red teaming  
Adversarial simulations in which expert teams emulate the tactics, techniques, and procedures of real-world attackers to test organizational defenses and incident response capabilities.
- ▶ Social engineering  
Testing human vulnerabilities through tactics such as phishing simulations to assess employee awareness and susceptibility to attacks.
- ▶ Exploit development  
Creating and automating exploits to identify and address vulnerabilities within IT systems.

## IBM Z and offensive security tools

Although IBM Z systems are securable by design, incorporating OffSec practices requires the use of specialized tools and techniques that are tailored for the mainframe environment. These tools and techniques include the following items:

- ▶ Open-source tools  
You can use tools such as Nmap, Python, Kali Linux, and Metasploit to scan, analyze, and exploit vulnerabilities on the mainframe.
- ▶ Proprietary software  
Vendors such as IBM and Broadcom offer mainframe security solutions, such as TDz, IBM zSecure, and SDS IronSphere, to address vulnerabilities and enhance security postures.
- ▶ Ethical hacking expertise  
Professionals with a deep understanding of mainframe systems, security protocols, and ethical hacking techniques are crucial for conducting OffSec assessments effectively.

Embracing OffSec practices for IBM Z involves a continuous process of vulnerability management, penetration testing, and integration with modern security operations centers to proactively address threats and help ensure regulatory compliance. The IBM Z Cyber Vault environment supports investment in ethical hacking expertise, the use of appropriate tools, and continuous review of the security posture to protect critical mainframe assets in the face of evolving cyberthreats.

## 5.4 Operational practices for enhancing integrity

The principles and practices that are described in this section focus on improving the integrity of the IBM Z Cyber Vault solution. Each subsection describes the monitoring, reporting, and security tools and methods that enhance the overall integrity of these solutions. Many organizations have already taken steps to enhance their cybersecurity and resilience, and the insights they have gained have helped IBM identify best practices for operating the IBM Z Cyber Vault solution more effectively. Every organization's journey is unique, and it is important to recognize that the path to becoming cyber resilient is a marathon rather than a sprint. The items that are outlined in this section do not need to be implemented all at once; their adoption should align with business priorities and timelines. These operational practices provide a strong foundation for building a secure and reliable IBM Z Cyber Vault environment.

### 5.4.1 Case for monitoring the IBM Cyber Vault environment

Controls around the operation of the IBM Z Cyber Vault environment can be enhanced beyond the standard validation process by using extra products and components as part of the IBM Z Cyber Vault architecture. In particular, organizations can add and use products and components from the IBM Security zSecure Suite (zSecure) to implement monitoring of critical environment components, such as network activity, and application-related controls such as data integrity. When these extra controls are combined with the standard validation process, the IBM Z Cyber Vault environment provides a more robust and reliable operational environment.

In the following sections, specific monitoring controls that enhance the operation of the IBM Z environment are described, and the zSecure products for implementing these controls are introduced. Before that, this section reviews several concepts and terminology that support the case for monitoring the IBM Z Cyber Vault environment.

The zero trust principle “never trust, always verify” should be considered a first line of defense for rigorously managing security entitlements and policies where data copies are stored. This principle includes identification and authentication of personnel who manage the IBM Z Cyber Vault environment and pre-approved, business-justified access control to the components and resources in this environment. Continuous monitoring of all activities that are performed by authorized personnel is required to help ensure compliance with security policies, along with real-time notifications of suspicious activities and actionable capabilities to stop any access that is deemed unauthorized. If organizations must use the IBM Z Cyber Vault environment, it is because an unexpected event affecting the availability, security, or reliability of the primary production facility occurred. Therefore, the “never trust, always verify” principle should be thoroughly followed in the IBM Z Cyber Vault environment.

The zero-trust architecture is defined by the National Institute of Standards and Technology (NIST)<sup>1</sup>. IBM aligns with NIST guidelines and is committed to embedding security and privacy capabilities into the design of its offerings.

---

<sup>1</sup> Source: <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf>

## 5.4.2 Ensuring that the environment is operating as expected

The IBM Z Cyber Vault solution can help ensure an operational environment in which business processes can be reliably conducted. The operational paradigm includes controls that allow the environment to be validated and the data to remain sound and not wrongly altered. In other words, this approach supports an operational environment in which the infrastructure, operating system services, business applications, and data are in a state that enables business to run with certainty and the expected outcomes.

The standard validation process includes provisions for analyzing critical system components such as catalogs, security databases, DBMS structures, and transaction processing systems. However, other system components and services that are not included in the standard validation process can also affect the stability of the operational platform. Changes to any of these components can lead to varying degrees of disruption, ranging from a single service interruption to a global outage.

IBM Z Cyber Vault is built on regular point-in-time immutable copies of a production environment (stored in an isolated and secure site) in which organizations can run regular data analytics to validate the infrastructure, data structure, and data content. Access to this site should be protected and monitored by using the same security policies, processes, and procedures that are applied to the production environment.

At least two techniques, such as fortifying the network and monitoring file integrity, can be added to enhance the validation process. When implemented correctly, these techniques improve the level of confidence and reliability of the operational platform to conduct business that, in principle, yields business results from an environment that is not only tightly protected and monitored but also validated with elements of consistency and integrity in mind.

## 5.4.3 Fortifying the network

The first technique focuses on protecting the IBM Z Cyber Vault environment from external actors that might compromise and jeopardize the stability of the operational platform. Network configuration controls are defined to prevent the following conditions:

- ▶ Unauthorized and unapproved network traffic from reaching the internal network of the IBM Z Cyber Vault environment.
- ▶ Modifications to the network configuration as defined by the network administrator, including any attempt to suppress the logging of unauthorized access that might result in a loss of accountability.

These two controls are defined as part of the IBM Z Cyber Vault network configuration. However, a monitoring mechanism is required to effectively determine the efficacy and robustness of these controls.

To be effective, a monitoring mechanism must run outside the target being monitored. For network monitoring, it should be implemented without relying on the configuration of the network. Its operation should not require network definitions and should instead run solely as an operating system service or component on top of the monitored target.

The IBM Security zSecure Alert (zSecure Alert) product is a real-time monitor for z/OS systems that are protected with IBM RACF or Broadcom ACF2, commonly referred to as an External Security Manager (ESM). zSecure Alert issues notifications for important events that are relevant to the health of the system at the time that they occur. These events intercept and use information that is recorded in the z/OS logging mechanism, the System Management Facility (SMF) component, messages that are issued by any component running on z/OS, or the results of inspecting and comparing system or component control block state content.

zSecure Alert provides more than 50 default alert notifications, and administrators can define custom alert notifications as needed. Notifications can be sent to various destinations, ranging from email and text messages to WTOs and external security information and event management (SIEM) dashboards for event aggregation and correlation.

For more information about how zSecure helps monitor the IBM Z Cyber Vault solution, see “IBM Security zSecure Alert capabilities” on page 129.

#### 5.4.4 Monitoring file integrity

The File Integrity Checking (FIC) definition from NIST covers the File Integrity Monitoring (FIM) function that is implemented by IBM zSecure. From a pragmatic standpoint, monitoring is perceived as an ongoing process, and checking is perceived as an on-demand process.

In general, FIM is suited for applications’ file validation because any data that is used by business applications must be consistent and maintain integrity. FIM provides an extra validation layer beyond the structural correctness of the file containing the data. For example, a file can be structurally correct while its data content is corrupted, which can cause different degrees of outages. Alternatively, the content might change without being corrupted, which can cause the application results to be, in principle, valid but contain incorrect information.

FIM can also be applied to infrastructure components’ configuration or control files. Some of these files are critical for the stability of the operational platform or the component. By applying FIM concepts to the infrastructure’s files, administrators can detect changes that might severely affect the stability of the operational platform, including the correct operation of components such as storage management, DBMS processing, and batch scheduling.

The IBM Security zSecure Audit product, part of the IBM zSecure Suite portfolio, provides a mechanism to implement FIM. FIM is primarily intended for integrity verification. It is implemented by using a specialized program, zSecure Collect (CKFCOLL), which can be configured to generate checksums for integrity verification purposes. These checksums can be generated for specific supported datasets and files or for all supported datasets and files of the operational platform.

In practical terms, the zSecure Collect program generates two types of checksums: anti-tamper digests and fingerprints. To learn more about what them and how they can be used in the solution, see “A closer look at FIM” on page 131.

#### 5.4.5 Reporting on the state

After operations in the IBM Z Cyber Vault environment complete, that is, when your personnel sign off from the IBM Z Cyber Vault environment, reports on all security activities that were performed in the IBM Z Cyber Vault environment should be generated. You can use these reports to demonstrate compliance with due-diligence requirements that are associated with operating the IBM Z Cyber Vault environment. You can show that provisions for the protection and monitoring of the environment are implemented, and that security activities that are related to authentication and access to resources while operating the environment are also reported. These reports help complete the cycle of secure operation of the IBM Z Cyber Vault environment: security protection, monitoring, and reporting.

You can use zSecure Access Monitor (Access Monitor), which is a component of the IBM Security zSecure Admin for RACF (zSecure Admin) product, to generate reports that are related to all authentication and access activities that are performed in the IBM Z Cyber Vault environment. These reports help validate that only approved personnel gained access to the environment and that resources and data were accessed according to the intended security definitions. The reports highlight how the security rules and access control lists were used compared to what was defined in the security database, and expose user-to-resource relationships as they occurred. This information helps validate security policies, the zero trust least-privilege access guideline, and provides decision-making input for improving the overall protection of the IBM Z Cyber Vault environment.

Reports can be generated in a format that can be interacted with by using the facilities of TSO/ISPF, or they can be generated as standard reports that can be printed or archived by a report management output tool for deferred analysis. Because of the transient nature of the IBM Z Cyber Vault environment, generating standard reports that can be printed or transferred to the primary production location for analysis is a best practice.

Three reports can be generated to analyze the security activities in the IBM Z Cyber Vault environment.

- By using the interactive report that is shown in Figure 5-2, security analysts, security auditors, or compliance officers can assess all access that is performed over the protected resources from the user ID perspective. The objective is to display activity that is performed by a user ID regardless of which resources were accessed. To analyze access events that are performed by the user ID that is associated with the z/OSMF component, select the IZUSVR user ID entry.

```

IBM Security zSecure ACCESS summary
Command ==>
All access monitor records
  Occurrence  Userid  Name                               First occurrence  Last occurrence
-----
   56 BBUCKNE  BUCKNER BOB                       15Aug2024 23:55  15Aug2024 23:55
   65 BOYDG   BOYD GREG                          15Aug2024 16:13  15Aug2024 16:15
   33 CICS1   CICS1 REGION USER ID              15Aug2024 23:03  15Aug2024 23:22
 2914 C2PSUSER                                14Aug2024 23:59  15Aug2024 23:58
   647 DANDRE  ANDRE DIDIER                       15Aug2024 11:48  15Aug2024 11:59
   20 DFHSM   DFHSM STARTED TASK                15Aug2024 08:24  15Aug2024 16:04
   144 DFLT1   CICS1 DFLT USER ID                15Aug2024 23:57  15Aug2024 23:57
   63 GCERQ01 CERQUONE GIOVANNI                 15Aug2024 10:18  15Aug2024 12:44
 11181 GEORGEN NG GEORGE                          15Aug2024 00:08  15Aug2024 18:37
   7036 GPMSERVE RMF DDS ID              15Aug2024 23:58  15Aug2024 23:58
   182 HCHECK  HEALTH CHECKER                     15Aug2024 10:22  15Aug2024 23:22
 17449 IZUSVR   ZOSMF STARTED TASK U              15Aug2024 15:22  15Aug2024 23:58
   63 OMVSKERN OPEN MVS                            15Aug2024 18:37  15Aug2024 18:37
   1824 SRGOEDD GOEDDE STEVEN                      15Aug2024 08:58  15Aug2024 15:53
 22250 STCRACF RACF STC USERID                  14Aug2024 23:59  15Aug2024 23:58
   160 SYSADM1 DB2 SYSAMIN                15Aug2024 06:29  15Aug2024 20:49
   104 TCPIP   TCPIP                              15Aug2024 23:41  15Aug2024 23:41
***** Bottom of Data *****

```

Figure 5-2 Access summary by user ID

- By using the interactive report that is shown in Figure 5-3, security analysts, security auditors, or compliance officers can assess all access that is performed over the protected resources from the resource class perspective. The objective is to see the activity on the resource regardless of who accessed it. To analyze activity over cryptographic services, select the CRYPTOZ, CSFKEYS, or CSFSERV class entries.

```

IBM Security zSecure ACCESS summary
Command ==>
All access monitor records
  Occurrence Class      First occurrence Last occurrence
-----
   6 ACCTNUM  15Aug2024 08:58 15Aug2024 18:31
   1 APPL     15Aug2024 11:22 15Aug2024 11:22
  104 CRYPTOZ 15Aug2024 23:41 15Aug2024 23:41
   11 CSFKEYS 15Aug2024 11:59 15Aug2024 11:59
 14567 CSFSERV 15Aug2024 15:22 15Aug2024 23:58
 19180 DATASET 14Aug2024 23:59 15Aug2024 23:58
   56 DSNR    15Aug2024 23:55 15Aug2024 23:55
 22266 FACILITY 15Aug2024 01:00 15Aug2024 23:58
   804 OPERCMDS 15Aug2024 00:11 15Aug2024 23:58
   8 RACFVARS 15Aug2024 21:15 15Aug2024 22:23
  1658 SDSF   15Aug2024 00:10 15Aug2024 18:34
  4222 SERVAUTH 15Aug2024 18:37 15Aug2024 23:58
   177 TCICSTRN 15Aug2024 23:03 15Aug2024 23:57
   606 TSOAUTH 15Aug2024 08:58 15Aug2024 18:36
   6 TSOPROC  15Aug2024 08:58 15Aug2024 18:31
   519 XFACILIT 15Aug2024 11:57 15Aug2024 23:58
***** Bottom of Data *****

```

Figure 5-3 Access summary by Resource Class type

- By using the interactive report that is shown in Figure 5-4, security analysts, security auditors, or compliance officers can assess all authentication mechanisms that took place in the IBM Z Cyber Vault environment from the authentication method (type) perspective. The objective is to see the authentication type that is used regardless of who uses it. To analyze all user IDs that were authenticated by using a password, select the Password method.

```

IBM Security zSecure ACCESS summary
Command ==>
All access monitor records
  Occurrence Meth      First occurrence Last occurrence
-----
  868 None             15Aug2024 01:00 15Aug2024 23:58
   73 Password         15Aug2024 08:58 15Aug2024 23:55
   17 Started          15Aug2024 00:59 15Aug2024 22:44
***** Bottom of Data *****

```

Figure 5-4 Authentication activity by authentication method

Access Monitor reports are not a replacement for detailed audit trail reports. They provide summarized reports based on events such as access activity (by user ID or by resource class) or authentication type (passwords, pass tickets, password phrase, or Multi-Factor Authentication (MFA)).

For more information about how to configure Access Monitor to generate access reports, see the following publications:

- [zSecure CARLa-Driven Components Installation and Deployment Guide](#)
- [zSecure Admin and zSecure Audit for RACF User Reference Manual](#)

## 5.4.6 Security considerations

To effectively implement a secured IBM Z Cyber Vault solution, follow the principle of least privilege by granting operators and administrators only the minimum level of access that is required to perform daily tasks and respond to cyber incidents. This approach is a cornerstone of the zero-trust security model.

Adopting a zero-trust strategy substantially reduces the risk of insider threats that might compromise an environment and lead to cyber incidents. The importance of this model cannot be overstated. It is built on six fundamental principles:

- ▶ Least privilege. Grant the least level of privilege as necessary.
- ▶ Fail securely. Minimize the harm of a failed security check.
- ▶ Separation of duties. At least two individuals are responsible for the completion of a task.
- ▶ Defense in depth. Use multiple layers of defense.
- ▶ Establish secure defaults. Default configurations should prioritize the most secure settings.
- ▶ Minimize the attack surface. Keep security simple to avoid confusing or complex situations.

These six pillars help reduce the risk of insider attacks by helping ensure that one individual does not have the power to cause widespread harm. Organizations can severely inhibit catastrophic damage by minimizing a solution’s attack vectors, granting the least necessary levels of access to operators, and separating the authority to run actions across multiple individuals. Cyber resiliency can be greatly enhanced when organizations adopt the “never trust, always verify” mindset. Within their configurations, responsibilities must be clearly divided to prevent users from accessing resources beyond their job scope.

Consider a scenario in which two users manage a GDPS solution: one is the storage administrator, and the other is a systems operator. To uphold the principles of least privilege and separation of duties, their levels of access should be carefully restricted. Therefore, a typical ACL might resemble the one that is shown in Table 5-2.

*Table 5-2 Upholding the principles of least privilege and separation of duties*

<b>Task</b>	<b>Storage administrator</b>	<b>Systems operator</b>
API	Yes	Yes
Main Menu	Yes	Yes
Disk Management	Yes	No
Standard Actions	No	Yes
Systems	No	Yes
Planned Actions	Yes	Yes
Sysplex Resource Management	No	Yes
Health Checks	Yes	Yes
Monitor Commands	Yes	No
Mon 1/2/3 <sup>a</sup>	Yes	No
NetView Command Interface	No	No

Task	Storage administrator	Systems operator
Configuration Management	Yes	No
Debug On/Off	No	No

a. These monitors are related to GDPS. Each monitor checks different aspects of the GDPS environment at varying intervals to help ensure system health and operational integrity.

In this example, each user has only enough access to do their daily tasks and can affect only specific areas of an environment on their own.

### Role-based security and dual control

GDPS LCP and CSM both offer role-based and dual-control security solutions to help enterprises meet their security needs. Users in a GDPS LCP environment can be granted read-only or alter authorities based on their access levels to certain SAF profiles that are associated with an LCP action. In addition to these general authorizations, GDPS LCP offers more granular security options by protecting actions at the sysplex, domain, and management profile levels. You can secure each parameter within a management profile through the LCP field-level security distinctions. With this extra level of security, administrators can restrict users from modifying specific fields within a management profile even when a user is granted access to the modify action. Role-based security is available for LCP starting at GDPS 4.5 and later.

Under the dual-control security model, any pervasive actions require a second user to review and approve an action before it runs. GDPS LCP Manager actions that have this protection are shown in Table 5-3.

Table 5-3 Protected GDPS LCP Manager actions

Action	Description
CREATE	Create a Safeguarded Copy or FlashCopy management profile.
MODIFY	Modifying an existing management profile.
DELETE	Delete a management profile.
FLAGSET	Include or exclude a specific volume from a management profile's future captures.
MODIFY Scheduler	Modify the parameters of the LCP Manager capture scheduler.
QUIESCE	Suspend all LCP Manager capture and release operations for a management profile or consistency group (CG).
RESUME	Resume all previously quiesced LCP Manager capture and release operations.

The scope of dual-control protected actions continues to expand as more features are incorporated into GDPS LCP. Dual-control security is available for LCP starting at GDPS 4.7 and later. For more information about the GDPS LCP role-based and dual-control security implementations, see the "LCP Security" section of the GDPS LCP Manager documentation.



# Monitoring the IBM Z Cyber Vault environment

To strengthen IBM Z Cyber Vault environments against cyberthreats, implementing robust monitoring, alerting, and integrity validation that is aligned with NIST security principles and a zero-trust architecture is essential. This appendix describes how to monitor the system security posture with a focus on privileged user actions and sensitive datasets by using real-time notifications and examples of automated responses.

The following topics are covered in this appendix:

- ▶ “Monitoring: Key terms and definitions explained” on page 128
- ▶ “IBM Security zSecure Alert capabilities” on page 129
- ▶ “A closer look at FIM” on page 131
- ▶ “Monitoring matters” on page 135

# Monitoring: Key terms and definitions explained

NIST has information assurance practices that applied by using controls that are associated with information security: confidentiality, integrity, and availability. NIST also makes provision for processes and procedures that are related to controlling the security of the resources by implementing mechanisms for the backup of enterprise resources, restore of resources to preserve their integrity, and detection and reaction to cyberthreats. These controls might be specific to each enterprise and internally defined practices in general.

NIST provides information-assurance practices through controls that are associated with information security: confidentiality, integrity, and availability. NIST also defines processes and procedures that are related to controlling the security of resources by implementing mechanisms for the backup of enterprise resources, the restoration of resources to preserve their integrity, and the detection of and reaction to cyberthreats. These controls can be specific to each enterprise and typically align with internally defined practices.

This section includes basic definitions to provide a common context around some of the concepts that are described in this appendix. The following definitions were extracted from the [glossary](#) that is provided by the NIST:

<b>Monitoring</b>	“Continual checking, supervising, critically observing or determining the status in order to identify change from the performance level required or expected.” <sup>1</sup>
<b>Integrity</b>	“Guarding against improper information modification or destruction and includes ensuring information non-repudiation and authenticity.” <sup>2</sup>
<b>File Integrity Checking (FIC)</b>	“Software that generates, stores, and compares message digests for files to detect changes made to the files.” <sup>3</sup>
<b>Zero trust</b>	“A collection of concepts and ideas designed to minimize uncertainty in enforcing accurate, least privilege per-request access decisions in information systems and services in the face of a network viewed as compromised.” <sup>4</sup>

In addition, IBM includes the following definitions (excerpts):

► **Data consistency**

Refers to the state of data in which all copies or instances are the same across all systems and databases. Consistency helps ensure that data is accurate, up-to-date, and coherent across different database systems, applications, and platforms. It plays a critical role in helping ensure that users of the data can trust the information that they access. Ways to help ensure data consistency include implementing data validation rules, using data standardization techniques, and employing data synchronization processes.

► **Data integrity**

Refers to the accuracy, completeness, and consistency of data throughout its lifecycle. It is the assurance that the data is not tampered with or altered in any unauthorized way. Therefore, data integrity helps ensure that the data remains intact, uncorrupted, and reliable. Common methods to improve data integrity include data validation techniques, implementing access controls and authentication mechanisms, and employing data backup and recovery procedures.

<sup>1</sup> Source: <https://csrc.nist.gov/glossary/term/monitoring>

<sup>2</sup> Source: <https://csrc.nist.gov/glossary/term/integrity>

<sup>3</sup> Source: [https://csrc.nist.gov/glossary/term/file\\_integrity\\_checking](https://csrc.nist.gov/glossary/term/file_integrity_checking)

<sup>4</sup> Source: [https://csrc.nist.gov/glossary/term/zero\\_trust](https://csrc.nist.gov/glossary/term/zero_trust)

- ▶ File Integrity Monitoring (FIM)

A security function that detects unauthorized changes to critical files and systems within the IBM z/OS mainframe environment. It works by comparing a file's current state to a trusted baseline by using checksums and alerting operators to any unauthorized modifications, corruption, or tampering to help ensure system integrity and support compliance with regulations such as Payment Card Industry Data Security Standard (PCI DSS).

## IBM Security zSecure Alert capabilities

IBM Security zSecure Alert (zSecure Alert) enhances the Identify domain of Cyber Security and Cyber Resilience within the NIST framework by alerting about changes to selected areas in the operational platform. IBM Z Cyber Vault environments that run zSecure Alert can use the alerting and notification mechanisms to identify changes to z/OS areas.

- ▶ User privileges
- ▶ datasets and general resource access entitlements
- ▶ Critical programs
- ▶ Security configuration and security rules and definitions

Note the alert notifications that are available for network monitoring that you can implement for the IBM Z Cyber Vault network configuration. As of zSecure Alert 3.1.0 and later, the following network-related alerts are available:

- ▶ Logon from a not allowed IP address (1124)
- ▶ IP attacks blocked by filter no longer logged (1609)
- ▶ IP attacks blocked by default filter no longer logged (1610)
- ▶ IP System Management Facility (SMF) subtype 119 no longer written (1611)
- ▶ IP filtering and IPsec tunnel support deactivated (1612)
- ▶ IP ports below 1024 no longer reserved (1613)
- ▶ IP interface security class changed (1614)
- ▶ IP filter rules changed (1615)

You can use alert 1124 (RACF) or the ACF2 equivalent (2124) to monitor the IBM Z Cyber Vault network configuration. You can configure alert 1124 (2124) to help ensure that a logon from a not allowed IP is reported when the event is detected by zSecure Alert. In general, such notifications are not issued based on how the network is configured. However, because there is always a possibility that the configuration can be accidentally or intentionally changed, implementing such a notification helps close the gap for any intentional or unintentional changes to the network configuration.

Figure A-1 shows an example of an email notification.

```
Alert: Authorized user CRMBXX2 logged on from 9.145.159.178
Logon by a userid from a not allowed IP address

Alert id          1124
Date and time     29Mar2021 13:33:08.88
User              CRMBXX2  IBM DEFAULT USER
Result           Success
Job name + id    CRMBXX2  TSU07970
System ID        8018
Source terminal   STCP0010
Source IP        9.145.159.178
```

Figure A-1 Notification through email

Alerts are identified with a four-digit numeric value that is based on the security product. This numeric value is added to the Alert ID field in the email notification. It is also added as the suffix for the message that is generated when WTO is set as the destination recipient. For alert 1124, the generated message ID is C2P1124I.

Then, the WTO message may be post-processed by automation tools to perform actions on the reported event. For alert 1124, automation can shut down a portion of the network, which in this case blocks any traffic coming from the monitored IP address that is defined as not allowed.

The other set of alerts is oriented toward maintaining the soundness of the IBM Z Cyber Vault network configuration as intended by the network administrator. These alerts can be activated for completeness and to help ensure that accountability for any actions or activities involving network management and operation is preserved. For example, alert 1611 can be activated to prevent the logging of a specific TCP/IP component from being deactivated, which allows network activity to occur without trace and ultimately with incomplete or no accountability (see Figure A-2).

```

From: C2POLICE at DINO
Subject: Alert: SMF 119 FTPCLIENT is no longer written by stack name

Alert: SMF 119 FTPCLIENT is no longer written -
audit trail incomplete in TCP/IP stack TCPIP
Alert id      1611
Changed field SMF119_FTPCLIENT(Yes->No)-
Stack        TCPIP
System ID    DINO

```

Figure A-2 Activating alert 1611

The scope of intrusion detection is defined by implementing controls that help prevent modifications to system services that might bypass or suppress security. Figure A-3 shows an alert that triggers when a RACF security option is changed (alert 1503). In this alert, supported RACF commands that are issued by user IDs with the RACF system-wide **SPECIAL** or group-level **SPECIAL** attribute are not logged unless other compensatory controls are in place. This action can compromise the security of the environment.

```

From: C2POLICE at DINO
Subject: Alert: Global security countermeasure changed by C##BNAT

Alert: Global security countermeasure changed by C##BNAT
SETROPTS command changed system security

Alert id      1503
Date and time 23Jan2003 11:51:56.01
RACF command  SETROPTS NOSAUDIT
User         C##BNAT NICK AFTERSOCK
Result       Success
Job name     C##BNAT
System id    DINO

```

Figure A-3 RACF security option were changed

Figure A-4 on page 131 shows an alert that triggers when an APF dataset is updated (alert 1204). Programs that are stored in APF datasets can run code that can bypass security, modify a critical file, and disable logging, and these actions go unnoticed. Therefore, updates to any APF-defined dataset should be continuously monitored.

```
From: C2POLICE at DINO
Subject: Alert: Update by C##ASCH on APF data set C##A.D.C##NEW.APF.LOAD

Alert: Update by C##ASCH on APF data set C##A.D.C##NEW.APF.LOAD
APF data set successfully updated

Alert id          1204
Date and time    03Feb2003 10:12:05.30
Data set         C##A.D.C##NEW.APF.LOAD
Access           ALTER
User             C##ASCH SIRAM CHRISTIAN
Result           Success
Job name         C##ASCHL
System ID        DINO
```

Figure A-4 APF dataset updated

Bypassing security and disabling logging should not be allowed. To minimize this risk, implement two monitoring actions:

- ▶ Actions that are performed by privileged users, including security and system administrators, must be monitored because these users have the authority to modify the security and configuration of the environment. Monitoring privileged users helps minimize insider threats. One way to control changes to access entitlements for authorized libraries is by allowing RACF commands that affect access to authorized libraries to be accepted only when a change control number (ticket ID) is associated with the command. You can use the Command Review component of the zSecure Admin for RACF product to enforce entering a ticket ID along with the RACF command when the command is generated from the zSecure ISPF interface.

For more information about how to configure and use Command Review, see the following resources:

- [zSecure Alert User Reference Manual](#)
- [Command logging facility](#)

- ▶ Changes to the list of APF datasets or any system list where the datasets are implicitly authorized. Restrict access to commands that modify the list of authorized libraries or restrict access to monitoring products such as IBM OMEGAMON® for z/OS that allow you to modify the list of authorized libraries.

For information about how to configure zSecure Alert to provide network monitoring alerting and other controls, see the following resources:

- [zSecure CARLa-Driven Components Installation and Deployment Guide](#)
- [zSecure Alert User Reference Manual](#)

## A closer look at FIM

To verify the integrity of datasets, the zSecure main program, CKRCARLA, compares anti-tamper digests. This field is based on the content of a dataset and various metadata, such as the dataset and member names, partitioned dataset (PDS) directory information, ISPF statistics, and Installation Data Record (IDR) and ZAP data. Because the metadata is used to compute the anti-tamper digest, it is difficult for any changes to go undetected. The anti-tamper digest is considered safer than the fingerprint and is best suited for integrity verification. Its primary use is anti-malware protection. In addition, a sensitivity type is automatically assigned to the object that is controlled by the FIM function.

Conversely, fingerprints are more generic because they are based solely on the content of the dataset or library without metadata. Their primary use is duplicate detection.

The term *checksum* is used to include both anti-tamper digests and fingerprints.

zSecure Audit enhances the Detect domain of the Cyber Security and Cyber Resilience with the NIST framework by adding data integrity validation at the content level. IBM Z Cyber Vault environments running zSecure Audit can use the FIM function to address and answer the following questions:

- ▶ Has my (critical) data being modified?
- ▶ Have my configuration or control files changed?

The business case for using FIM is to assist with intrusion detection, help with integrity validation, and comply with regulations such as PCI DSS<sup>5</sup>, GDPR, HIPAA, and others.

To implement FIM, complete the following steps:

1. Create a BASELINE environment data or configuration file with checksums that are computed for integrity verification based on what is considered a trusted or gold copy.
2. Compute new checksums for the datasets that the installation designates as critical, as captured in a CURRENT point-in-time environment data or configuration file.
3. Generate reports to help confirm that the data is not tampered with. These data or configuration files are called CKFREEZE in zSecure terminology. They contain information about everything running on the system at the time the CKFCOLL program runs, including all datasets and z/OS UNIX sensitive files.

The BASELINE file is created on the source environment that creates the IBM Z Cyber Vault environment by using replication mechanisms. The CURRENT file is created in the IBM Z Cyber Vault environment counterpart.

For the FIM process to be effective, create the BASELINE copy in the environment where business-as-usual operations are performed, usually the production environment. This BASELINE copy is replicated to the IBM Z Cyber Vault environment as part of the volume replication mechanism. Then, in the IBM Z Cyber Vault environment, the CURRENT copy is created. These two files are compared to generate various reports.

For FIM, in addition to the files for which checksums are computed, known as the files subject to integrity monitoring (which can include selected critical files as identified by the system or installation or all files in the environment), specify the algorithm that computes the checksum values must be specified, or the program can determine which algorithm to use.

The results of the compare process are integrity verification reports that you can use to identify the files that have not been tampered with. Two reports categories are available:

- ▶ Member (source and non-source programs) level
- ▶ Dataset level

These reports include dataset (file) sensitivity for system and component datasets. As of IBM zSecure 3.1.0 and later, over 700 sensitivity types for z/OS datasets, z/OS UNIX files, and other programs are defined and automatically assigned, ranging from standard types such as APF and LINKLST to more specialized types such as Db2 bootstrap datasets, and selected third-party products or components. You can also add your own sensitivity types, for example, for business applications files where the application design defines what files are considered sensitive.

---

<sup>5</sup> PCI DSS Requirement 11.5 mandates using a change-detection mechanism such as FIM to detect and alert on unauthorized modifications to critical files, including system, configuration, and content files.

A typical FIM implementation includes these events:

- ▶ Day-1 checksum compute – BASELINE
- ▶ Day-2 checksum compute – CURRENT (Same value – data integrity not compromised)
- ▶ Day-n checksum compute – CURRENT (Different value – data integrity may have been compromised)

A compare process that leads to different checksum values does not always indicate when the data was compromised. There might be necessary, one-time changes to the files, in which case different checksums are expected. This situation is typical for last-minute approved changes. Therefore, documentation about the system management actions that are required for the IBM Z Cyber Vault environment is fundamental to making the final determination of the results that are included in the integrity reports.

The key factor for an effective FIM implementation rests on how to determine the comparison criteria. In general, installation policies, system management, and change management practices must be aligned to the following items:

- ▶ Generate a trusted or golden (BASELINE) version, such as when it is created immediately after creation or published by software vendors or applications.
- ▶ Generate a CURRENT (point-in-time version) that is based on a client-defined date, which might be required by change management procedures.
- ▶ Generate a CURRENT (point-in-time version) that is required by internal policies as part of an annual audit review.
- ▶ Generate a CURRENT (point-in-time version) weekly that is driven by industry regulations, such as PCI DSS requirement 11.5.

Figure A-5 shows how you can use FIM detect changes to specific datasets.

```

z/OS data sets
Command ==>
Line 1 of 61
5 Aug 2024 16:39
Scroll==> CSR

Changes
ANTI_TAMPER_DIGEST(changed)
FINGERPRINT(changed)
Identification
Security complex name      LABPLEX
System name                LAB1 LAB1
Data set name              GCERQ01.LAB1.JCLLIB
Data set type              nvsam
Device class               DASD      Unit type      3390
Volume serial or SMS managed *SMS*    Volume is scratch  No
Start of multi-file complex File sequence number
Volume serial              TSUSR1    Data set is migrated  No
Real volume serial         TSUSR1    Volume is mounted    Yes
Real data set name         GCERQ01.LAB1.JCLLIB
DASD box serial number and id IBM-75-0000000LFZ71-0257
Catalog name               CATALOG.LABPLEX.USERCAT
Catalog volume             LABSYS

Detail information
Type of sensitive data set Site-Dsn-R Site-Dsn-U
Catalog alias              GCERQ01
  
```

Figure A-5 Changes to specific datasets

Monitor application libraries by using the Datasets report option of FIM. In this example, the library is a Job Control Language (JCL) library. The library was changed as indicated by the Changes section where both the ANTI\_TAMPER\_DIGEST and the FINGERPRINT fields show as changed.

The sensitivity is user-defined, which is indicated by the Site-Dsn prefix. The -R and -U suffixes stand for the Read (R) and Update (U) sensitivity access levels, which you can use to classify the object as 'confidential', R-sensitive or 'integrity-related', or U-sensitive.

At the dataset level, no indication of what was changed is provided. If this file is a flat file, you must investigate the installations further to discover what was changed. FIM reports only that the content of the dataset changed, which should suffice to start an investigation for the impact of such a change in the environment. However, for PDSe datasets, where members' metadata is used to detect changes, you can use Members report option of FIM to investigate further about what changed, as shown in Figure A-6.

```

Non-program members - C2PML18M - GCERQ01.LAB1.JCLLIB
Command ==> _____ Line 1 of 45
                                                                    Scroll==> CSR
                                                                    5 Aug 2024 16:28

Identification
Member name          C2PML18M
Data set name        GCERQ01.LAB1.JCLLIB
DASD box serial number and id  IBM-75-000000LFZ71-0257
Volume serial        TSUSR1
Volume serial or SMS managed *SMS*
Member of PDSE       Yes      TTR      000165
System name          LAB1      Complex name  LABPLEX
Type of sensitive resource Site-Dsn-R Site-Dsn-U

Changes
ANTI_TAMPER_DIGEST(changed)
BYTES(5600->5680)
FINGERPRINT(changed)
LAST_CHANGE( 2 Aug 2024 13:24:24.383983-> 5 Aug 2024 20:25:19.314428)
PDF_CHGDATE( 2 Aug 2024-> 5 Aug 2024)
PDF_CHGTIME(09:24:00->16:25:00)
PDF_LINES(70->71)
PDF_VERSION(01.03->01.05)

```

Figure A-6 Changes to members in a PDS

In this case, the source member C2PML18M that is indicated in the Changes section was changed. This member can be “linked” to the JCL library in Figure A-5 on page 133 because the dataset name and volume serial are the same across the two reports.

Metadata information for the source member was changed, including the size (in bytes) and the total number of lines of the member. For the latter, it is possible that lines were added, modified, or deleted, which resulted in an effective extra line (70 - 71). To discover how the source member changed, you must use other software management tools with versioning capabilities. FIM reports that the member content changed when it is compared to the BASELINE trusted copy. As it was the case for the JCL library, the sensitivity is user-defined and matches the sensitivity level for the JCL library.

Here, the program (load module) RACFCAI was changed, as indicated in the Changes section in Figure A-7 on page 135.

```

Command ==> Programs - RACFCAI - GCERQ01.T.MMO.CICSTK.LOADLIB Line 1 of 60
                                                    Scroll==> CSR
                                                    9 Aug 2024 15:09

Identification
Member name          RACFCAI
Data set name        GCERQ01.T.MMO.CICSTK.LOADLIB
DASD box serial number and id IBM-75-000000LFZ71-0257
Volume serial        TSUSR1
Volume serial or SMS managed *SMS*
Member of PDSE       Yes      TTR      000046
System name          LAB1      Complex name LABPLEX
Type of sensitive resource CICS Loadlb
Alias names for member TSSCAI
Changes
BYTES(69632->155648)
LAST_CHANGE(12 Feb 2021 15:23:07.788000-> 9 Aug 2024 19:08:30.267814)
LKEDDATE(12 Feb 2021-> 9 Aug 2024)
LKEDTIME(10:23:07->15:08:30)
STORSIZE(30668->29396)

```

Figure A-7 Changes to a specific PDS member

The metadata information for the program was changed, as reported in the fields for the last date (LKEDDATE) and last time (LKEDTIME) that the program was compiled (LKED) and its size (STORSIZE) in bytes. To discover what changed, you must use other software management tools with versioning capabilities. FIM reports that a new copy of the program was compiled when compared to the BASELINE trusted copy. If you view the assigned sensitivity, you might conclude that this program is an IBM Customer Information Control System (IBM CICS) application program (the sensitivity type is CICS Loadlb) because it is on a library that was defined to one of the CICS regions that was active at the time that the CURRENT (point in time copy) environment data or configuration file was created.

For more information about how to configure the CKFCOLL program for implementing FIM, restrictions, and how to perform the compare process and generate the integrity reports, see [zSecure Admin and zSecure Audit for RACF User Reference Manual](#).

## Monitoring matters

You can configure other products and components in the IBM Z Cyber Vault environment to provide monitoring capabilities beyond what is built in the standard IBM Z Cyber Vault solution. In 5.4.1, “Case for monitoring the IBM Cyber Vault environment” on page 120, we made the case for monitoring the IBM Z Cyber Vault environment by using some of the capabilities that are provided by the following IBM products:

- ▶ IBM Security zSecure Admin for RACF Command Review component and Access Monitor
- ▶ IBM Security zSecure Audit for RACF (or ACF2) FIM
- ▶ IBM Security zSecure Alert for RACF (or ACF2) default alerts

You may configure all or a subset of these products and components in the IBM Z Cyber Vault environment or any system (logical partition (LPAR)) that is associated with it. FIM is suitable for implementing Type 3 validations, which are primarily used for applications validation, that is, data and configuration files that make up a business application. Plan to deploy notification alerts on all systems that make up the IBM Z Cyber Vault environment to adhere to the zero trust principle.



# Abbreviations and acronyms

<b>ACL</b>	access control list	<b>IBM</b>	International Business Machines Corporation
<b>APPN</b>	Advanced Peer-to-Peer Networking	<b>ICSF</b>	Integrated Cryptographic Service Facility
<b>ATG</b>	Advanced Technology Group	<b>IDCAMS</b>	Integrated Data Cluster Access Method Services
<b>BC</b>	business continuity	<b>IDR</b>	Installation Data Record
<b>BCPii</b>	Base Control Program internal interface	<b>IDS</b>	Intrusion Detection Services
<b>BCS</b>	basic catalog structure	<b>IMS</b>	Information Management System
<b>BOT</b>	beginning of tape	<b>IODF</b>	input/output definition file
<b>CDS</b>	Couple Data Set	<b>IPL</b>	initial program load
<b>CF</b>	coupling facility	<b>ISO</b>	International Organization for Standardization
<b>CFRM</b>	Coupling Facility Resource Manager	<b>IZBR</b>	IBM Z Backup Resiliency
<b>CG</b>	consistency group	<b>JCL</b>	Job Control Language
<b>CICS</b>	Customer Information Control System	<b>KSDS</b>	key sequenced dataset
<b>CPC</b>	central processor complex	<b>LCP</b>	Logical Corruption Protection
<b>CS</b>	Copy Services	<b>LCU</b>	logical control unit
<b>CSM</b>	Copy Services Manager	<b>LOB</b>	line of business
<b>DBA</b>	database administrator	<b>LPAR</b>	logical partition
<b>DORA</b>	Digital Operational Resilience Act	<b>LRSN</b>	Log Record Sequence Number
<b>DR</b>	disaster recovery	<b>LSS</b>	logical subsystem
<b>DS CLI</b>	DS Command-Line Interface	<b>LWORM</b>	logical WORM
<b>DS GUI</b>	DS Graphical User Interface	<b>MFA</b>	Multi-Factor Authentication
<b>DVE</b>	Dynamic Volume Expansion	<b>MGM</b>	Metro Global Mirror
<b>ELB</b>	Extended Long Busy	<b>NAT</b>	Network Address Translation
<b>ESE</b>	Extent Space Efficient	<b>OAM</b>	Operations, Administration, and Maintenance
<b>ESM</b>	External Security Manager	<b>OffSec</b>	offensive security
<b>FCM4</b>	Flash Core Modules 4	<b>OSCONFIG</b>	operating system configuration
<b>FIC</b>	File Integrity Checking	<b>PCI DSS</b>	Payment Card Industry Data Security Standard
<b>FIM</b>	File Integrity Monitoring	<b>PDS</b>	partitioned dataset
<b>GC</b>	Global Copy	<b>PDS<sub>e</sub></b>	partitioned dataset extended
<b>GDPS</b>	Geographically Dispersed Parallel Sysplex	<b>RC</b>	recovery copy
<b>GM</b>	Global Mirror	<b>R</b>	Read
<b>HA</b>	high availability	<b>RMF</b>	Resource Measurement Facility
<b>HADR</b>	high availability and disaster recovery	<b>RPFC</b>	Remote Pair FlashCopy
<b>HCD</b>	Hardware Configuration Definition	<b>RPO</b>	Recovery Point Objective
<b>HLQ</b>	high-level qualifier	<b>RPQ</b>	Request Price Quotation
<b>HMC</b>	Hardware Management Console	<b>RS</b>	replication site
<b>HPR</b>	high-performance routing	<b>RTO</b>	Recovery Time Objective

<b>SA ProcOps</b>	IBM System Automation Processor Operations
<b>SDAC</b>	Selective Device Access Controls
<b>SDF</b>	Status Display Facility
<b>SIEM</b>	security information and event management
<b>SLA</b>	service-level agreement
<b>SME</b>	subject matter expert
<b>SMF</b>	System Management Facility
<b>SMS</b>	Systems Managed Storage
<b>SOAR</b>	security orchestration, automation, and response
<b>SOC</b>	security operations center
<b>TCT</b>	transparent cloud tiering
<b>TDz</b>	Threat Detection for z/OS
<b>TEL</b>	Technology Expert Labs
<b>TRL</b>	Transport Resource List
<b>TTP</b>	tactics, techniques, and procedures
<b>UCB</b>	Unit Control Block
<b>Update</b>	U
<b>VLAN</b>	virtual LAN
<b>VERRDS</b>	VSAM Recovery Record Data Set
<b>VSAM</b>	Virtual Storage Access Method
<b>VVDS</b>	VSAM Volume Data Set
<b>WTOR</b>	write-to-operator-with-reply
<b>XCF</b>	Cross-System Coupling Facility

# Related publications

The publications that are listed in this section are considered suitable for a more detailed description of the topics that are covered in this book.

## IBM Redbooks

The following IBM Redbooks publications provide more information about the topics in this document. Some publications that are referenced in this list might be available in softcopy only.

- ▶ *IBM GDPS: An Introduction to Concepts and Capabilities*, SG24-6374
- ▶ *IBM TS7700 Release 6.0 Guide*, SG24-8464
- ▶ *What is New in DFSMSrmm*, SG24-8529

You can search for, view, download, or order these documents and other Redbooks, Redpapers, web docs, drafts, and additional materials, at the following website:

[ibm.com/redbooks](https://ibm.com/redbooks)

## Other publications

These publications are also relevant as further information sources:

- ▶ *GDPS LOGICAL Corruption Protections Manager Installation and Customization Guide*, ZG24-6763

## Online resources

These websites are also relevant as further information sources:

- ▶ Command logging facility  
<https://www.ibm.com/docs/en/szs/3.1.0?topic=guide-command-logging-facility>
- ▶ Cost of a Data Breach Report 2025  
<https://www.ibm.com/reports/data-breach>
- ▶ *Data Integrity Detecting and Responding to Ransomware and Other Destructive Events*  
<https://www.nccoe.nist.gov/sites/default/files/legacy-files/di-detect-respond-nist-sp1800-26-draft.pdf>
- ▶ DS8000 Safeguarded Copy and Extent Space Efficient (ESE) FlashCopy capacity sizing by using the CSM ESESizer or SGCSizer functions  
<https://www.ibm.com/support/pages/node/6372180>
- ▶ *Framework for Improving Critical Infrastructure Cybersecurity*  
<https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.04162018.pdf>

- ▶ IBM Db2 Recovery Expert Pro for z/OS 1.1  
<https://www.ibm.com/downloads/documents/us-en/1443d5dd844f4e06>
- ▶ IBM Threat Detection for z/OS (TDz)  
<https://www.ibm.com/products/threat-detection-for-zos>
- ▶ *ICT risk management framework*  
<https://www.springlex.eu/en/packages/dora/dora-regulation/article-6/?ref=dora-info>
- ▶ International Organization for Standardization (ISO) Standards  
<https://www.iso.org/standards.html>
- ▶ Maximize the power of your lines of defense against cyberattacks with IBM Storage FlashSystem and IBM Storage Defender  
<https://www.ibm.com/new/product-blog/maximize-the-power-of-your-lines-of-defense-against-cyber-attacks-with-ibm-storage-flashsystem-and-ibm-storage-defender>
- ▶ NIST Glossary  
<https://csrc.nist.gov/glossary>
- ▶ What's new/changed in GDPS V4.8?  
<https://www.ibm.com/downloads/documents/us-en/1227c12dd4b8bc70>
- ▶ *zSecure Alert User Reference Manual*  
<https://www.ibm.com/docs/en/szs/3.1.0?topic=alert-pdf>
- ▶ *zSecure Admin and Audit for RACF User Reference Manual*  
<https://www.ibm.com/docs/en/szs/3.1.0?topic=manual-pdf>
- ▶ *zSecure CARLa-Driven Components Installation and Deployment Guide*  
<https://www.ibm.com/docs/en/szs/3.1.0?topic=deployment-pdf>

## Help from IBM

IBM Support and downloads

[ibm.com/support](https://www.ibm.com/support)

Services from IBM Consulting

[ibm.com/services](https://www.ibm.com/services)

IBM Training

[ibm.com/training](https://www.ibm.com/training)

**Redbooks**

**Getting Started with IBM Z Cyber Vault**

(0.2"spine)  
0.17"->0.473"  
90->249 pages







SG24-8511-01

ISBN 0738462403

Printed in U.S.A.

Get connected

