Redbooks
ibm.com/redbooks

Getting Started with IBM Z Cyber Vault

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

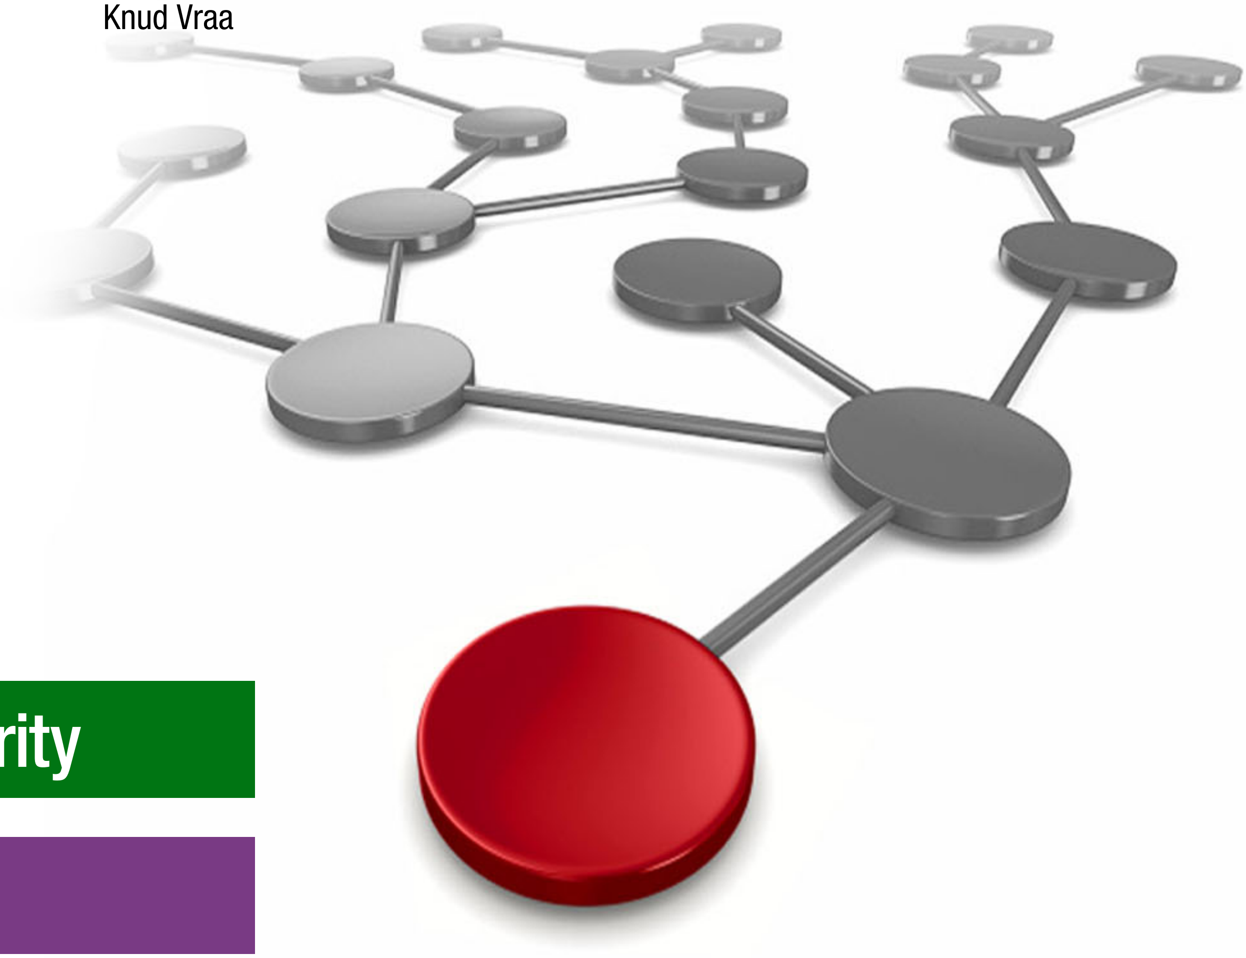

Security

IBM Z

IBM
Redbooks

IBM Redbooks

Getting Started with IBM Z Cyber Vault

November 2021

SG24-8511-00

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

First Edition (November 2021)

This edition applies to the IBM Z Cyber Vault solution.

Contents

Notices vii
Trademarks viii

Preface ix
Authors. ix
Now you can become a published author, too! xii
Comments welcome. xiii
Stay connected to IBM Redbooks xiii

Chapter 1. Protecting against key cyberthreats. 1
1.1 Common cybersecurity threats 2
1.1.1 Main factors that determine favorable outcomes 3
1.1.2 What is needed in a cyber resiliency strategy. 3
1.1.3 Frameworks for IT cyber resiliency 4
1.2 Why cyber resiliency is important 5
1.2.1 Backup and recovery 5
1.2.2 Achieving enterprise-grade cyber resiliency 5
1.3 IBM Z Cyber Vault solution 6
1.3.1 Cyber resiliency capabilities 7
1.3.2 IBM Z Cyber Vault solution technical summary 8
1.3.3 Common IBM Z Cyber Vault configurations 8
1.3.4 Evaluating an IBM Z Cyber Vault solution. 10
1.4 Use cases for the IBM Z Cyber Vault solution 11
1.4.1 Validation use cases. 11
1.4.2 Recovery use cases 13
1.4.3 Extra protection layer by using offline copy. 13
1.5 Data protection technologies for a zero trust security model 14
1.5.1 IBM Z pervasive encryption. 14
1.5.2 IBM Z Multi-Factor Authentication. 15
1.5.3 **IBM Hyper Protect Data Controller 15**
1.6 IBM Z Cyber Vault service offerings 16
1.6.1 IBM Z Cyber Vault Discover and Architecture Workshop 16
1.6.2 IBM Z Cyber Vault Installation and Configuration Technical Assistance 16
1.6.3 IBM Z Cyber Vault Data Recovery System Validation Technical Assistance 17

Chapter 2. Planning and designing the IBM Z Cyber Vault environment 19
2.1 Planning approach for the IBM Z Cyber Vault environment 20
2.1.1 Implementation phases. 20
2.1.2 IBM Z Cyber Vault environment operational roles 23
2.2 IBM Z Cyber Vault reference architecture. 24
2.2.1 Production environment 25
2.2.2 IBM Z Cyber Vault recovery system environment. 26
2.2.3 IBM Z Cyber Vault storage environment. 28
2.2.4 Isolating the IBM Z Cyber Vault environment 32
2.2.5 Network isolation for the recovery system 34
2.3 IBM Z Cyber Vault infrastructure design 34
2.3.1 Replication management 34
2.3.2 IBM GDPS capabilities and supported models 35
2.4 Prerequisites and considerations for implementing Safeguarded Copy 41

2.4.1 Safeguarded Copy prerequisites 41
2.4.2 Safeguarded Copy considerations 41
2.4.3 Sizing extra storage for Safeguarded Copy 42
2.4.4 IBM Z Cyber Vault storage sizing 45
2.5 Key operational considerations for Safeguarded Copy. 47
2.5.1 Frequency of Safeguarded Copy backups 47
2.5.2 Retention period of Safeguarded Copy backups 48
2.6 IBM Z platform considerations 48
2.6.1 IBM Z System Recovery Boost and IBM Z Cyber Vault 49
2.6.2 Sizing the environment 49
2.7 IBM z/OS considerations. 50
2.7.1 Non-database managed files 52
2.7.2 IBM z/OS security 53
2.8 Application considerations 54
2.8.1 Db2 subsystem 54
2.8.2 IMS subsystem 56
2.8.3 CICS subsystem 57
2.8.4 Application dependency 58
2.9 Validation, forensic analysis, and recovery considerations. 59
2.9.1 Validation 59
2.9.2 Forensic analysis 62
2.9.3 Recovery. 64

Chapter 3. Deploying the IBM Z Cyber Vault environment 69
3.1 Preparing to deploy the IBM Z Cyber Vault environment 70
3.2 Description of our IBM Z Cyber Vault environment. 71
3.3 Setting up IBM Z Cyber Vault storage environment 74
3.3.1 Metro secondary and Safeguarded Copy 74
3.3.2 Recovery copy 75
3.3.3 Staging and persistent volumes 76
3.3.4 Recovery in case of full data corruption 77
3.3.5 IODF definitions for the IBM Z Cyber Vault environment 78
3.3.6 Couple data sets for the IBM Z Cyber Vault environment. 78
3.3.7 Considerations for pre-allocated WLM CDSs 79
3.3.8 Considerations for pre-allocated XCF CDSs. 79
3.4 Setting up Safeguarded Copy in IBM Z Cyber Vault storage 80
3.4.1 Automating Safeguarded Copy captures 85
3.4.2 Using digital certificates to authenticate to the GDPS RESTful API 89
3.4.3 Automating the process 90
3.5 Preparing the IBM Z Cyber Vault environment for validation 92
3.5.1 Setting up CFs and LPARs for the IBM Z Cyber Vault environment. 92
3.5.2 Network isolation considerations for the IBM Z Cyber Vault environment 95
3.5.3 Security access considerations for the IBM Z Cyber Vault environment 95
3.5.4 IODF requirements for the IBM Z Cyber Vault environment. 96
3.5.5 IPL changes for the IBM Z Cyber Vault environment 96
3.5.6 PARMLIB changes 97
3.5.7 CFRM changes 99
3.5.8 TCP/IP configuration changes 100
3.5.9 IBM Z Multi-Factor Authentication consideration 101
3.5.10 Pervasive encryption considerations. 101
3.5.11 Determining tasks to start on IBM Z Cyber Vault 101
3.5.12 Considerations for Db2 data sharing. 101
3.5.13 Automating your IBM Z Cyber Vault IPL procedures 101

Chapter 4. Establishing a validation framework 103
4.1 Creating a validation framework 104
4.1.1 Data validation assumptions 104
4.1.2 IBM Z Cyber Vault validation terminology 105
4.1.3 Constructing data structure validation jobs 106
4.1.4 IBM Z Cyber Vault validation framework 106
4.1.5 IBM Z Cyber Vault job naming conventions 107
4.1.6 IBM Z Cyber Vault validation data set naming conventions 108
4.1.7 Creating an IBM Z Cyber Vault validation monitoring task 109
4.2 Type 1 validation 111
4.2.1 Collecting feedback from a Type 1 validation 111
4.3 Type 2 validation 114
4.3.1 Description of the construction of a simple validation job 115
4.3.2 Detailed description of the Db2 data structure validation jobs 116
4.3.3 Automating data structure validation 122
4.3.4 Collecting feedback from data structure validation 127
4.4 Type 3 validation 140
4.4.1 Constructing data content validation jobs 140
4.4.2 Automating data content validation 140
4.4.3 Collecting feedback from data content validation 141

Abbreviations and acronyms 143

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 http://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®
CICS®
Db2®
DS8000®
FICON®
FlashCopy®
GDPS®
HyperSwap®
IBM®
IBM FlashSystem®
IBM Garage™
IBM Security™
IBM Services®
IBM Z®
IBM z14®
Parallel Sysplex®
RACF®
Redbooks®
Redbooks (logo) ®
Tivoli®
VTAM®
X-Force®
z/Architecture®
z/OS®
z/VM®
z15™

The following terms are trademarks of other companies:

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.

OpenShift, are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

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

With cyberattacks on the raise, cyber resiliency is becoming more important than ever. Cyber resiliency can provide the required capability against cyberattacks and prevent significant impact if the attacks are detected early. A cybersecurity strategy might minimize the risk of attacks getting through to systems, applications, and data, but a cyber resiliency strategy is needed to minimize the impact of a cyberattack by recovering quickly. Preparing for, responding to, and recovering from a cyberattack is not something that just happens, but must be thoroughly designed, planned for, and tested.

This IBM® Redbooks® publication looks at some common cyberthreats and introduces a cyber resiliency solution that is called IBM Z® Cyber Vault. This book describes the technology and cyber resiliency capabilities of the solution at various hardware, software, and operational levels, and describes what to consider when pursuing higher cyber resiliency goals.

Guidance and step-by-step examples for the deployment of IBM Z Cyber Vault are also included, and a suggested framework with advice for conducting basic data validation and samples that can be tailored to individual business needs, priorities, and IT configurations.

This publication is intended for IT managers, line of business (LOB) managers, IT architects, system programmers, storage administrators, security administrators, database administrators (DBAs), and system operations professionals.

Authors

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

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

Matthias Bangert is an Executive IT Specialist at IBM, working in the worldwide technical sales organization and focusing on IBM Z resiliency and tailored fit pricing. Matthias started in IT 35 years ago as an MVS system programmer. He moved to MVS/SP and MVS/XA while working for a customer in Frankfurt/Mainz, Germany. During his career, he has focused on the entire IBM Z stack and led many disaster recovery (DR) projects, sysplex merges, and data center moves. When he started at IBM 22 years ago, Matthias established an IBM Geographically Dispersed Parallel Sysplex (IBM GDPS®) community in Germany and focused on IBM Z technical sales support for large financial and public clients in Germany. After his role as Technical Sales Manager for Europe, he advanced his career with worldwide responsibility for resiliency topics. In this role, Matthias promoted the idea of IBM Z Cyber Vault together with colleagues from IBM Storage and GDPS.

Cyril Armand is an IBM Z Specialist working in IBM Garage™ for Systems in Montpellier, France. At the end of 2017, and after two years working on the mainframe infrastructure area, he joined the IBM Z Benchmark Center of Montpellier with a specific focus on IBM z/OS®. His range of missions includes proofs of concept (PoCs) regarding z/OS new features or leading large performance benchmarks to support IBM Z hardware sales.

Roger Bales is an IBM Consulting Practice Leader who is responsible for leading the global IBM Z and LinuxONE Lab Services practice. He has over 40 years of experience working with the IBM Z family of systems and software as a Systems Programmer, Networking Specialist, Computer Operator, and Data Center Manager. As a Practice Leader, he has overall responsibility for the strategic direction and success of the global practice. He also serves as an Executive Consultant specializing in business and IT strategic alignment. The scope of Roger's advisory services includes industry-specific business strategies, IT infrastructure, Operations, Applications, Data Management, and Analytics. Previously, he led global sales, marketing, and business development teams at IBM. Before coming to IBM, Roger held various technical and management positions at AT&T Long Lines and served as General Manager for a global ISV.

Diego Bessone has been part of the mainframe world since 1987, when he worked at the data center of Aerolíneas Argentinas, an airline carrier of Argentina. For 10 years, he worked on various mainframe systems (including DOS/VSE, VM/ESA, and MVS), ran IBM Db2® and IBM Information Management System (IMS) applications, and supported the largest airline reservation and check-in system in Latin America. He also worked as a technical consultant for the IT departments of multinational companies, finance organizations, utilities, and government agencies in Argentina. Diego joined IBM in 1998 and is an IBM Global Sales Business Unit Executive for IBM Z Software. During his career at IBM, he led the sales strategy for the IBM Systems management portfolio in Argentina. He moved to Austin, TX in 2001 to lead the Latin America team. He has worked in business operations management when he moved to Miami, FL in 2003, and led sales at a worldwide level since 2009. He managed the IBM Z Application and Platform software portfolio and grew the market share for the platform while helping IBM customers maximize the value of their IBM mainframes.

Anthony Ciabattoni is an IBM Certified Executive IT specialist and a current member of the Db2 for z/OS Development SWAT team that is based out of the Silicon Valley Lab. He has worked directly with Db2 for z/OS for over 27 years. His experience includes being a customer lead Db2 architect and working in the Db2 for z/OS Tools development lab. He is a recognized Db2 for z/OS subject matter expert (SME); has been a member of the IBM Db2 for z/OS Customer Advisory Council for over 12 years; and was elected as an IBM Information Champion for 9 consecutive years before joining IBM. In Anthony's current role as a senior technical advocate for the Db2 for z/OS Development Lab, he conducts consulting and educational workshops, Db2 health checks, and 360 Continuous Availability studies for worldwide Db2 for z/OS customers. He also works as a lab advocate to some of largest North American customers of IBM. He has presented on Db2 topics at IBM worldwide conferences, IDUG (NA, SA, and EMEA), and SHARE Conferences (US). He has co-authored numerous IBM Redbooks publications, and white papers.

Michael Frankenberg is a Certified IT Specialist in Germany and joined IBM in 1995. With more than 20 years of experience in high-end storage, he works in Technical Sales Support at the EMEA Storage Competence Center. His area of expertise includes performance analysis, establishing high availability (HA) and disaster recovery (DR) solutions, and implementation of storage systems. He supports the introduction of new products and provides advice for IBM Business Partners, Technical Sales, and customers. He holds a degree in Electrical Engineering / Information Technology from University of Applied Sciences Bochum, Germany.

Debra Hallen is an IBM Z IT Specialist working in the IBM Systems Lab Services group. She has over 30 years of professional computer technical and management experience that encompasses numerous facets of mainframe technical development software, strategic planning, project management, analysis and programming, and DR. Debra has extensive experience in z/OS concepts and facilities, storage management, performance and capacity management, and automation by using various vendor tool sets. She joined IBM Systems Lab Services in 2013 with a specific focus on projects encompassing z/OS Sysplex Splits, Data Center Migrations, z/OS hardware and software upgrades, DFSMShsm Health Checks, planning for exploitation features, Virtual Storage Access Method (VSAM) record-level sharing (RLS) implementation planning and Health Checks, z/OSMF implementation and reviews, and z/OS Container Extensions (zCX) planning and education.

DeWayne Hughes is an Executive IBM Z Client Architect covering financial services clients in the US. He has worked for IBM since 1989 and has been an IT Architect since 2002. DeWayne speaks at SHARE and IBM Systems TechU events. He is working on IBM Z projects in the areas of DevOps, OpenShift, containers, AI, machine learning, and security solutions and resiliency. DeWayne works to integrate mainframes into hybrid cloud environments.

Vinod Kanwal is a senior Storage Lab Consultant assisting clients nationally and worldwide with implementing replication, IBM HyperSwap®, and migration from old to new IBM storage. He started his career in 1980 as a designer of an SQL for a vendor, three years before IBM announced IBM Db2. He started working for IBM in 1982 as a z/OS Systems Programmer and performed his first data center migration in 1984. He worked in the Financial and Telecom industries in the US for almost 20 years as developer, support, architect, and ma License Compliance Manager (LCM) architect and later moved to the IBM Copy Services Manager (CSM) development team for couple years, where he helped clients with Spectrum Control and TotalStorage Productivity Center - Replication (TPC-R) (now known as IBM Tivoli® Storage Productivity Center) installations, set up and use HyperSwap, and helped with various replication technologies. Having worked with over 50 clients worldwide, he has a wide perspective of client needs and wants. He holds master’s degrees in Mathematics and Computer Science.

Karen Smolar is an IBM Certified Executive Architect. She has 33 years of experience designing and building infrastructure and application solutions that meet complex business objectives and maximize the value from IBM Systems and IBM Storage products. Karen works with clients worldwide and across every industry to help transform their infrastructure solutions by aligning IT technologies to their business initiatives. She specializes in identifying and solving complex resiliency and cyber resiliency challenges, which include IBM Z, IBM Storage, IBM Parallel Sysplex®, and GDPS. Karen' s application development experience and IT service management background allow her to bring a unique end-to-end, holistic perspective to each challenge and deliver complete solutions.

Jean-Marc Vandon is an IBM certified IT specialist with 20 years of IBM Z experience. He is responsible for Technical Sales of IBM GDPS High Availability and Disaster Recovery (HADR) solution offerings in Europe, Asia Pacific, and Africa. Jean-Marc directly engages with customers when a GDPS project is initiated to understand their requirements; explain how the solution can achieve expectations; and size the implementation project, both for the asset purchase and implementation services. Jean-Marc is a regular speaker at large IBM Z customer events in Europe, China, Japan, and several countries in the Middle East. In addition, he acts as an advisor to customers; is a preferred technical communications focal point; and a customer advocate for Lloyds Banking Group and UBS for GDPS HADR solutions.

Paolo Vitali is a Thought Leader IT Specialist who has worked in IBM since 1985 in the mainframe area. He spent 26 years at GTS Italy with Financial Sectors Enterprises z/OS Parallel Sysplexes implementation and tuning. Between 2009 and 2012 at the Worldwide Benchmark Center in Montpellier, France, he realized several benchmarks and PoCs for customers worldwide. He joined the GDPS Solution Test Team in Montpellier in 2013, where he focuses on three- and four-site GDPS solutions with GDPS Metro Global - XRC (GDPS MzGM), Metro/Global Mirror (MGM), and Multi-Target Metro Mirror (MTMM). More recently, he worked in GDPS Logical Corruption Protection design and testing.

Knud Vraa has been with IBM Denmark for 23 years working with z/OS, IBM z/VM®, and related IBM Z software. Specifically, storage administration, GDPS, and Parallel Sysplex configurations are the key areas that he has supported for the last 5 years. Knud currently holds a position as an IBM Z Client Architect in the Nordics with a main focus on IBM Z resiliency and IBM Z Cyber Vault.

Thanks to the following people for their contributions to this project:

Wendi Gervis Watson
IBM Redbooks, Raleigh Center

Frederic Bonvalet
IBM France

Francisco Duarte
IBM Brazil

Nick Clayton and Dave Clitherow
IBM UK

Javier Chavez, Dave Petersen, Joe Welsh
IBM US

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

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

- Send your comments in an email to:

  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:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

1

Protecting against key cyberthreats

Businesses, nations, and individuals face information technology threats of greater severity and cost than ever before. New forms of malware (or ransomware) exist that can attack computer systems, and the results can impede service delivery. The recent pandemic and increased digital transformation accelerated the incentives for bad actors and criminals to steal or disable data and computer systems.

Companies, governments, and people who are victims of a ransomware attack might suffer in the following ways:

- The loss of trust and reputation both individually or for a company's brand
- The inability to provide goods and services to their clients or citizens
- A large financial cost to recover data
- Even larger financial costs to investigate, remediate, and recover from the attack

One result of this risk is the increasing demands on companies to take proper precautions against cyberattacks, which includes reporting attacks (both privately to insurance, government, or regulatory agencies, and publicly), and demonstrating a thoughtful and tested plan to prevent or address cyberattacks.

As a further matter, not only are cyberattacks happening more frequently, but ransomware victims are suffering repeat attacks. According to a recent report, approximately 25% of businesses that were surveyed eventually shut down and 29% were forced to eliminate jobs.

The following topics are covered in this chapter:

- Common cybersecurity threats
- Why cyber resiliency is important
- IBM Z Cyber Vault solution
- Use cases for the IBM Z Cyber Vault solution
- Data protection technologies for a zero trust security model
- IBM Z Cyber Vault service offerings

1.1 Common cybersecurity threats

Cyberattacks are constantly evolving, and new variants of malware and ransomware are rising in number and financial severity. Some industries have witnessed over twice the increase in the number of attacks during the COVID-19 pandemic of 2020. To unlock the data, ransoms increased from hundreds or thousands of dollars up to tens of millions of dollars. These numbers are expected to continue to rise.

A key finding by IBM Security™ X-Force® and other security experts is that ransomware typically attacks application and business data first rather than computer systems data. The goal is to keep computers operating to encrypt more data and to eventually share the ransom demand to a user. These attacks require protection at all levels of the IT environment, including:

- Computer systems and infrastructure itself, such as operating systems, clustering technology, storage systems, and disaster recovery (DR) 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

Malware is closely related to ransomware and might be used to implant ransomware weapons and programs into IT infrastructure assets. Malware might also use methods such as intentional deletion or erasure of data.

How businesses detect and respond to these attacks must now consider the following scenarios:

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

Ransomware and malware attacks generally are launched by external entities who have successfully attacked a business by using phishing, identity theft, or other scams. The most common attacks start with infected attachments that are sent by email, smartphone message, or multimedia. The other common attack vector is infected websites that exploit unpatched browser vulnerabilities, which, although less common than email attacks, still is a successful way to gain entry into the IT environment.

Another reason for loss of data is unintentional data corruption, which happens when bad logic is introduced into an application code update, a patch to fix one problem causes other problems like data loss, or simple human error (either in commands or a process). Inadequate testing and backup practices can also be factors in this type of data loss. None of these changes are wanted, yet many people are impacted by the fallout from these errors.

The act of intentional data loss or destruction occurs when an internal employee, contractor, or trusted entity who has access and permission to internal systems, networks, or data deliberately modifies, destroys, or steals data. They usually want to seek revenge on the company and cause harm by corrupting or erasing data. This type of attack might be one of the most difficult to detect if it involves the scrambling of valid data. For example, an attack can be to switch the values of valid Social Security or government identification numbers with the wrong people in a database. All the actual numbers are still accurate, but now are associated with the wrong person. However, methods that are used to detect whether data has been encrypted or erased might not detect this kind of activity. Both intentional or human error attacks, and malicious attacks or sabotage, can come from people inside the company.

1.1.1 Main factors that determine favorable outcomes

Careful planning and preparation should be given to evaluating the impact on business reputation, implications of regulatory compliance, and risk of financial loss from a cyberattack.

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

The 2021 Cost of a Data Breach Report jointly produced by the Ponemon Institute and IBM Security found the following actions reduced the financial and brand impacts of a data breach while reducing the time to detect and respond to the breach:

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

A perfect cybersecurity solution does not exist for protecting against every possible form of cyberattack. However, a well thought-out cybersecurity strategy can minimize the risk of attacks getting through to computer systems and data. Also, a cyber resilience strategy can minimize the impact of a cyberattack by recovering quickly.

1.1.2 What is needed in a cyber resiliency strategy

A comprehensive cyber resiliency strategy typically includes the following elements:

- Identification: This function focuses on the need of a business to understand its most important assets and resources, which 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 much of the technical and physical security controls for developing and implementing appropriate safeguards and 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, which might include anomalies and events, continuous security monitoring, and early detection processes.

- Incident response: This function ensures an appropriate response to cyberattacks and other cybersecurity events. Categories include response planning, communications, analysis, mitigation, and improvements.

- Recovery: This function covers the actual implementation of plans for cyber resilience to ensure business continuity (BC) in the event of a cyberattack, security breach, or another cybersecurity event. The recovery functions are recovery planning improvements and communications.

There are several frameworks that are available that can be used to build and enhance cybersecurity and resiliency plans, and also improve readiness.

1.1.3 Frameworks for IT cyber resiliency

Specific regulations and frameworks vary by country or region of the world, yet it is worthwhile to evaluate multiple frameworks even if they are created by another country or entity. One commonly cited framework was released in 2013 and updated in 2018 by the National Institute for Standards and Technology (NIST) is the *Framework for Improving Critical Infrastructure Cybersecurity*.

Complementary International Organization for Standardization (ISO) documents are ISO 31000 and ISO 27005, which are found at ISO. These documents can be used to 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 lead an entity from initial risk evaluation and planning through steps to respond and evaluate plans for future events:

- 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: Response plan is run during or after an incident.
- RS.AN-3: Forensics are performed.
- RC.RP-1: Recovery plan is run during or after a cybersecurity incident.

Several documents are available for study and reference on cybersecurity and cyber resilience.

For more information about ransomware by the IBM Security X-Force team, see *The definitive guide to ransomeware: Readiness, response, and remediation*.

This document provides guidelines about how to prepare and respond to a cyberattack. Although many of the specific examples are based on Microsoft Windows, the principles and lessons that are provided also work for UNIX and Linux systems. While IBM Z and z/OS are different in many ways from x86-based platforms, valuable ideas can be applied when building IBM Z cybersecurity and cyber resiliency plans. The document contains information about the four key elements of a ransomware incident response lifecycle strategy:

- Preparation
- Detection and analysis
- Containment, eradication, and recovery
- Post-incident activities

The 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*.

1.2 Why cyber resiliency is important

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

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 a situation where 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 where systems and data are intentionally corrupted, erased, or encrypted.

A recent report about the results of a Forbes Insights and IBM survey of 353 executives worldwide on what is a true cyber resilience had the following statement:

"True cyber resilience means having a plan that could protect data and systems as they evolve and minimizing damage and disruption that could come from new threats. It means addressing the lifecycle of data as it is created, dispersed and stored as well as building resilience into every step."[1]

1.2.1 Backup and recovery

Successful backup and recovery procedures are required to maintain established and tested plans for protecting data. These procedures usually involve various methods of copying data into storage systems and tape. Backup and recovery operational processes should run on a regular prescribed basis to meet the needs of the business. These backups are likely automated, managed on how long they are kept in storage before rolling off to tape, and either kept permanently or deleted after a period.

Time and effort are invested in recovery capability to cope with potential hardware incidents that might cause production disruption. Reasons for failures are limitless, and usually a resilient architecture based on redundancy of all individual IT equipment and data is deployed. Backup servers, duplicate switches, and replicated data converge to provide insurance that if one element fails, the infrastructure will use one remaining spare instance. Servers are being made redundant in clusters, and data is replicated at second, third, or more sites.

Although this concept shows its efficacy in protecting from loss of data or equipment, limitations exist when dealing with data corruption issues, such as:

- Replication of corrupted data throughout nonisolated systems.
- Unlikely events that trigger DR. Data corruption can go undetected for extended periods.
- Lack of granularity for short to medium term backup recovery points.
- Recovery flexibility is limited, and partial recovery is difficult.

1.2.2 Achieving enterprise-grade cyber resiliency

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

The greatest challenge of a cyberattack is detecting that systems have been compromised or an attacker has successfully gained access somewhere in the enterprise IT environment.

[1] https://www.ibm.com/account/reg/us-en/signup?formid=urx-38121

Many studies show that an attacker might be working undetected for several months before the actual attack begins, looking for weaknesses, gaining higher privileges, and planting malware to conduct the attack.

Full cyber resiliency requires intrusion detection, monitoring for unusual behavior by individuals, programs, and systems, and reporting and dashboards to alert teams of this unusual behavior. All employees, contractors, and other people working with IT tools or systems must continue to be educated on how to prevent common attack points, such as phishing, smishing, vishing, or social engineering. They must also be trained in how to recognize and report unusual behavior. It is too late if the first detection of a ransomware attack is after it has occurred. Investment and dedication to proper technologies, tools, processes, monitoring, education, and communication are critical before an incident has occurred. These items are key to achieving enterprise-grade cybersecurity and resiliency.

In addition, a cyber resiliency solution must provide capabilities to protect against the unique challenges of a cyberattack. The first unique capability is the need for airgap, immutable copies of data that cannot be corrupted or erased by the cyberattacker. Tools are also needed to continually validate data to help detect an attack and build confidence in the quality of a backup if it is needed in the future. These tools will also help IT staff perform the forensic analysis that is needed to assess the incident; formulate optimal recovery strategies and options; and determine the scope of recovery, files, databases, or entire systems. An IBM Z cyber resiliency solution should be part of a larger enterprise cybersecurity plan and framework.

For IBM Z cyber resiliency, IBM introduced the *IBM Z Cyber Vault* solution. This solution addresses key components of both the NIST Cybersecurity Framework and the *Framework for Improving Critical Infrastructure Cybersecurity* that are described in 1.1.3, “Frameworks for IT cyber resiliency” on page 4.

The IBM Z Cyber Vault solution allows for a custom deployment that best meets business and technical requirements for IBM Z cyber resiliency.

1.3 IBM Z Cyber Vault solution

The IBM Z Cyber Vault solution provides a safe, isolated environment where an exact replica of the production environment is stored. The IBM Z Cyber Vault environment does not improve production environment robustness because that is not its function. Existing tools and best practices are used to keep the z/OS production system secure and provide the appropriate level of backup and recovery capabilities.

An extra benefit of having an IBM Z Cyber Vault environment is to use this system as a sandbox to run data validation processes without affecting production workloads, which might reduce high costs, performance issues, and the risk of introducing errors into business applications. This environment is also the perfect place to conduct forensic analysis after data corruption is detected, and based on the analysis, exercise surgical recovery procedures with peace of mind that if something goes wrong with the recovery, you can always go back to the Safeguarded Copy backup that you started from.

1.3.1 Cyber resiliency capabilities

With the IBM Z Cyber Vault solution, you can develop and use the following distinct cyber resiliency capabilities:

- Data validation: Regular operational validation of the Safeguarded Copy backups to provide proactive detection of data corruption or reassurance that the copy is clean before you take further action. The data corruption detection effectiveness depends on the custom-designed processes and their implementation. (Both are the responsibility of the user.)
- Forensic analysis: Start a copy of the production system and use it to investigate the problem and determine what the recovery action is.
- Surgical recovery: Extract data from the Safeguarded Copy and logically restore it into the production environment.
- Catastrophic recovery: Recover the entire environment to the point in time of a previously validated Safeguarded Copy backup if this option is the only recovery option.
- Offline backup: Take a fresh backup of the successfully validated environment to add an additional protection layer.

These capabilities are enabled by running z/OS utilities and other tools that are identified as necessary to maximize the value of this solution (see Figure 1-1). The usage, design, and assembly of these utilities, tools, and associated automation differ based on the individual IT environment and business requirements, and they must be developed and tested as a part of building the IBM Z Cyber Vault environment.

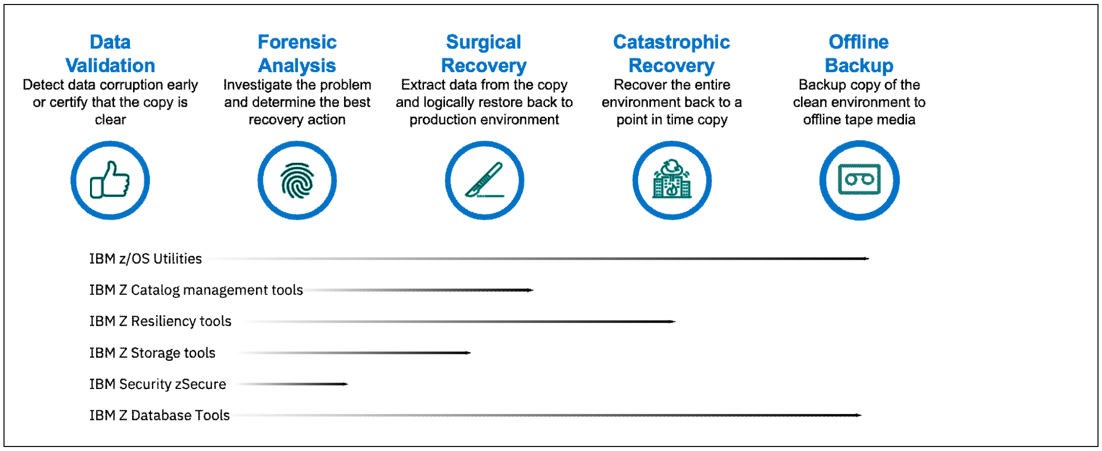
Data Validation
Detect data corruption early or certify that the copy is clear
Forensic Analysis
Investigate the problem and determine the best recovery action
Surgical Recovery
Extract data from the copy and logically restore back to production environment
Catastrophic Recovery
Recover the entire environment back to a point in time copy
Offline Backup
Backup copy of the clean environment to offline tape media
IBM z/OS Utilities
IBM Z Catalog management tools
IBM Z Resiliency tools
IBM Z Storage tools
IBM Security zSecure
IBM Z Database Tools

Figure 1-1 z/OS utilities and other tools

An extra use case for the IBM Z Cyber Vault environment, after it is deployed and tested, is to build a cyber range environment to perform penetration testing and ethical hacking of IT systems. The environment must run on systems that are isolated from production with access to the same applications and data, which are stored in the Safeguarded Copy backups and can be loaded in to the IBM Z Cyber Vault environment. This work can either alternate with ongoing validation (not necessarily recommended) or another logical partition (LPAR) can be constructed on the IBM Z Cyber Vault recovery system to do the cyber range work. This shared use increases the value that is gained from the investment that is made in the IBM Z Cyber Vault solution.

Another use case for the IBM Z Cyber Vault environment is to perform controlled and planned *chaos engineering* experiments. According to principles of chaos engineering, chaos engineering is the discipline of experimenting on a system to build confidence in the system's capability to withstand turbulent conditions in production. So, in this case we propose to proactively prepare for unpredictable data corruption to prepare your teams for failure. This experience can be used to improve service levels by revealing weak points in your production environment and then fortifying it, thus building new confidence in your system.

1.3.2 IBM Z Cyber Vault solution technical summary

IBM Z Cyber Vault solution contains a combination of various capabilities from IBM Z, IBM Storage, IBM Software, and IBM Services®. The IBM Z Cyber Vault environment is established with airgap, immutable copies, which are taken at multiple points-in-times on production data, with rapid recovery capability. In addition, this solution can enable validation of data to support testing on the validity of each captured copy.

The IBM Z Cyber Vault environment includes the following components:

- IBM Z (isolated LPARs or systems)
- IBM DS8000® Storage with Safeguarded Copy
- IBM Software for data validation, forensic analysis, and surgical recovery
- IBM Copy Services Manager (CSM) or IBM Geographically Dispersed Parallel Sysplex (GDPS) for data capture and recovery purposes
- IBM System Recovery Boost (faster recovery on an IBM z15™ or later)

The IBM Z Cyber Vault environment is isolated from the production environment. Isolation can be achieved by either fencing off the components virtually or physically, as shown in Figure 1-2.

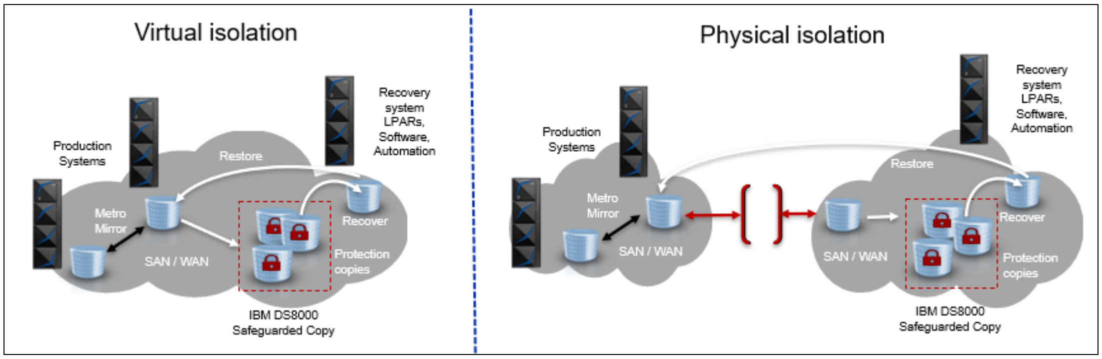

Figure 1-2 IBM Z Cyber Vault environment: Virtual versus logical isolation

1.3.3 Common IBM Z Cyber Vault configurations

IBM Z environments vary in size, capacity, topology, and configuration. They also vary in the types of applications that are deployed, methods in which the data is stored and backed up, the use of mirroring and replication for DR, and so on.

Regardless of the specific IBM Z environment, they typically fall under one of the three configuration types that are shown in Figure 1-3.

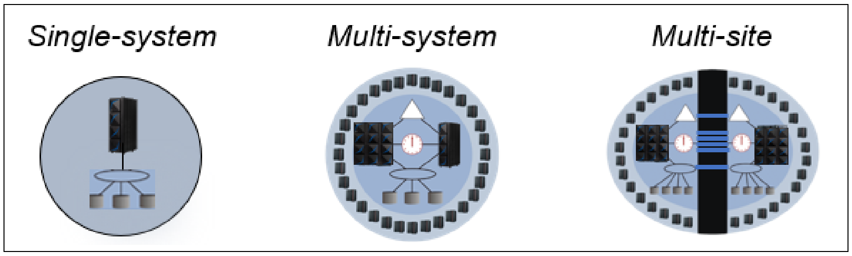

Figure 1-3 IBM Z configuration types for cyber resiliency

Each of these IBM Z configuration types, which are combined with the IBM Z Cyber Vault capabilities, can provide a high degree of cyber resiliency.

Single-system configuration

A single IBM Z platform running systems of record might have only a few production LPARs, with an extra set of development and test LPARs. Some use Virtual Storage Access Method (VSAM) or other sequential data sets to store application data, and others also use a database such as Db2 or IBM Information Management System (IMS) for z/OS. DR is likely provided on a single system in a remote location by using either tape- or disk-mirroring to copy data to the remote site.

Regardless of the size of an IBM Z environment, the IBM Z Cyber Vault components must be separate. A separate IBM DS8000 storage system should be used for Safeguarded Copy backups from the primary storage systems. A best practice might be to create an isolated LPAR with dedicated processors, memory, adapters, and so on, if it is built on the same IBM Z footprint that is already in use. However, to obtain maximum value, a fully isolated IBM Z Cyber Vault environment should exist to ensure that an attacker cannot access the Safeguarded Copy backups or the IBM Z Cyber Vault LPAR. Security to prevent unauthorized access or attack of the entire solution should also be included in the IBM Z Cyber Vault environment.

The IBM Z Cyber Vault solution in a single-system environment can provide all its core capabilities. Depending on how many applications, data sets, and databases are used, it might be simpler to perform data validation in a single-system environment than in larger IBM Z installations. The development and use of automation is important because the size of the IBM Z technical teams might be smaller in a single-system environment.

Multi-system configuration

A multi-system configuration includes at least two IBM Z platforms in the primary data center. Metro Mirror (MM) capabilities are present, and there is a moderate-to-high level of using IBM Parallel Sysplex for high availability (HA). It is more likely that databases such as Db2 store a significant percentage of system-of-record data.

This configuration should use a stand-alone system for the IBM Z Cyber Vault LPARs. As a best practice, do not use an LPAR in the production sysplex. Safeguarded Copy backups are created each day based on the number of production LPARs that exist, the amount of data sharing and HA capabilities in use, and other factors (such as the number of application instances). Development and use of automation are important because this configuration will validate as many Safeguarded Copy backups as possible per day, which will drive the need to perform validations 24x7 in the IBM Z Cyber Vault environment.

The capabilities expand for a multi-system configuration compared to a single-system configuration because there is likely a larger IT security team in place with more in-depth security tools and requirements. Governmental and industry regulations become more stringent as a company grows, so there is a greater need for tools to monitor for unusual behavior.

Multi-site configuration

In a multi-site configuration, two or more production data centers are in use. Consider whether a single IBM Z Cyber Vault implementation is sufficient for the IBM Z Cyber Vault profile.

One option for a multi-site configuration within metro distance is to have an IBM Z Cyber Vault environment in each production data center. In each data center, Safeguarded Copy backups are created once each hour at different times. In this case, the Safeguarded Copy backups can be on a physically isolated storage system. It is critical that data is not lost because it would take a long time to restore vast amounts of data to a production system. This configuration is likely to have either IBM CSM or IBM GDPS implemented. (GDPS is the preferred solution for fast recovery through automation.)

Enterprises that have multi-site configurations must focus heavily on automation to allow for as much data validation as possible to run each day. Validation of data will be more difficult and requires more effort to develop and run over time. Tight connections between Intrusion Detection Services (IDS) and processes must be built, and the IT security team should work closely with the IBM Z operations team to detect and act if a cyberattack occurs.

1.3.4 Evaluating an IBM Z Cyber Vault solution

Most enterprises run systems with a combination of IBM and other vendor hardware and software. However, security and resiliency must be achieved regardless of the provider of each technology if the functions and capabilities that are required are delivered.

An IBM Z Cyber Vault implementation will deliver the expected capabilities if it has the following attributes:

- The hardware components of the IBM Z Cyber Vault solution are airgapped.
- Storage replication technology can create immutable point-in-time copies of volumes.
- A customized system configuration for the IBM Z Cyber Vault environment was created.
- The software stack in the IBM Z Cyber Vault environment is an exact replica of the production environment.
- Tools, utilities, and procedures to manage, analyze, and restore system components are already identified and installed in production.
- Application data validation procedures are available.
- Automation tools and procedures are defined and built to create a repeatable validation process.

Different options are available in the marketplace to comply with these attributes. The correct components and tools should be chosen based on existing hardware and software configurations, internal rules and regulations, and operational best practices. As a result, every IBM Z Cyber Vault implementation will be different, but if done properly, will deliver the required capabilities.

For more information about the technologies that are used with the IBM Z Cyber Vault solution and discussions based on planning and design considerations, see Chapter 2, “Planning and designing the IBM Z Cyber Vault environment” on page 19.

1.4 Use cases for the IBM Z Cyber Vault solution

The IBM Z Cyber Vault solution provides new extensions beyond standard backup and recovery or DR solutions by supporting *validation* and *recovery* use cases.

1.4.1 Validation use cases

One of the key use cases for an IBM Z Cyber Vault environment is the validation of data that is stored and protected by the Safeguarded Copy backup process. To minimize recovery time, proactively check each Safeguarded Copy backup before an actual recovery is needed.

To validate a Safeguarded Copy backup, complete the following steps:

1. Create airgap copies of the production system data, which cannot be affected directly by data corruption. All copies that are created by replication or mirroring pass errors on the primary systems to all copies because they are continuously updated with changes on the primary volumes. However, the airgap copies that are created by Safeguarded Copy in the IBM DS8000 storage system cannot be modified after they are created.
2. Start the automated process to create recovery volumes from the Safeguarded Copy backup that is going to be validated.
3. After all the volumes in the selected copy are recovered, the IBM Z Cyber Vault recovery system is started (initial program load (IPL)) and the processes to perform the validation of the environment and data run.

There are three types of data validation that are specific to each IBM Z environment:

- Type 1 validation – System data

  Validate whether an LPAR can fully perform an IPL from the restored volumes by checking out core parts of z/OS, and enabling subsystems and logging in to them.

- Type 2 validation – Data structures

  Perform health checks, run scripts and tools to validate catalogs and other core parts of the system, and check and validate coupling facility (CF) structures. Validate that key middleware, databases, and run times are operational. These data structure validations of Customer Information Control System (IBM CICS®), IBM MQ, Db2, IMS, and batch environments ensure that the z/OS image can run applications, handle transactions, and process data. All these validations are needed to know that a system is fully operational.

- Type 3 validation – Application data

  This validation is to ensure that application and user data that is stored in data sets, databases, or other subsystems are valid. This validation is the final step to ensure that a copy is not corrupted by malware, ransomware, or any other source of intentional or unintentional data corruption, and therefore able to be trusted. These validations can be done by running numerous database queries, running batch programs and online transactions, and running other application tests to prove that data is available. It is the responsibility of the application, database, and technology teams to provide the appropriate tools and scripts to run these tests, which will be incorporated into the IBM Z Cyber Vault automation framework to be run.

The IBM Z Cyber Vault recovery volumes will be eligible (when the actual validation performed) if one of the following conditions exists:

- There are no errors.
- The errors that are identified are deemed acceptable because they either represent a known validation error or they were flagged in the baseline as pre-existent.

When the data validation processes successfully complete and satisfy the requirements to deem the IBM Z Cyber Vault environment free of logical data corruption, it can be backed up to a virtual tape library. This consistent point-in-time copy of the IBM Z Cyber Vault environment can then be offloaded in to physical tapes and stored in a safe location, providing an extra recovery option.

Figure 1-4 depicts the IBM Z Cyber Vault validation process.

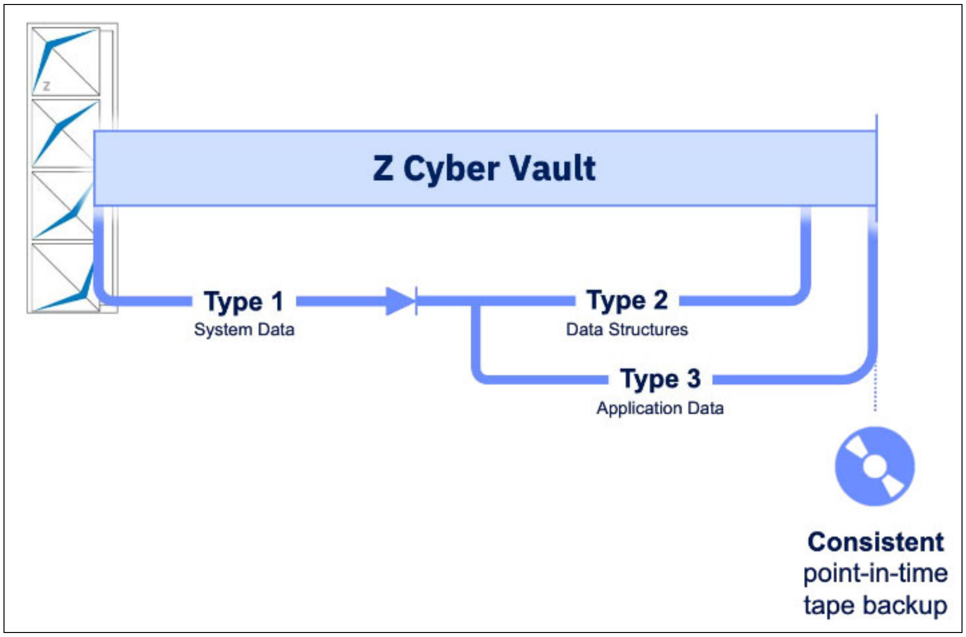

Figure 1-4 IBM Z Cyber Vault validation process

Because the validation process must be repeated for as many Safeguarded Copy backups as possible, consider automating it. If GDPS is implemented, it can be used to automate all these actions. If GDPS is not used, but some other type of automation tool (such as IBM System Automation for z/OS or an equivalent ISV solution), then that automation tool can be customized to run the procedures that are required for each of these steps.

The results of the validation processes should be checked to ensure that they are successful, and that data corruption has not been found. A validation baseline should already be established, which will be used to avoid identifying known errors as indicators of corrupted data. Error management should be automated as part of each validation type.

Some of these validations might also be performed in a more frequent but less thorough fashion in the regular production environment. Deciding where and when to run each validation process is a tradeoff between available processing capacity, performance, and risk.

For more information about using IBM Software tools and utilities to perform these validations, see 2.9.1, “Validation” on page 59 and Chapter 4, “Establishing a validation framework” on page 103.

1.4.2 Recovery use cases

All the operational validation types are typically done at prescribed intervals based on business needs and risk assessments. Recovery operations are done when a cyberattack is known to be ongoing or has occurred. With that in mind, an environment to experiment and research recovery operations is critical. The environment can be used to fully test and document how forensic analysis and various recovery tasks will be performed before the actual need ever arises. The IBM Z Cyber Vault environment can be used for this purpose to gain confidence in the effort and time that is required to perform these tasks. All recovery use cases are less procedural and require deep IBM Z skills and expertise to successfully run.

The recovery operations include:

- Forensic analysis: Identify various types of security analysis that would need to be done against z/OS and its components, core middleware, run times, databases, applications, and data. Plan what tools and procedures would be used to do forensic research to identify the cause and scope of an attack, which will vary based on what was attacked, how long it has been occurring, and other environment-specific items. This analysis must be properly trained for and tested, and then documented for possible future use.
- Surgical recovery: This operation is critical to restoring data, files, or systems back into production use if there has been an intended or unintended data loss. The tools and process that are used to perform the recovery will vary by attack vector and the attacked environment. Planning, training, testing, and documentation are critical before implementing the IBM Z Cyber Vault solution.
- Catastrophic recovery: This option is the last option that everyone hopes will never be used. The IBM Z Cyber Vault solution provides this capability, and it is a best practice to perform a full catastrophic recovery exercise against a test or development system before considering the IBM Z Cyber Vault environment, again with planning, training, testing, and documentation.

For more information about the key items to consider for these operations, see Chapter 2, “Planning and designing the IBM Z Cyber Vault environment” on page 19.

1.4.3 Extra protection layer by using offline copy

When using the IBM Z Cyber Vault environment to regularly run validations, it is valuable to take an extra step and save the validated system into an extra offline copy on a virtual tape system (such as the IBM TS7700) with no expiration or a long expiration date. This best practice adds an additional protection layer and saves valuable time if this copy must be restored as part of a catastrophic recovery procedure, for two reasons:

- The image has already been fully validated.
- Since a point-in-time copy does not quiesce transactions, a database system-level backup is a fuzzy copy, which might not contain committed data and might contain uncommitted data. After starting the IBM Z Cyber Vault recovery system and running all the data structure validation procedures, the database data is made consistent.

1.5 Data protection technologies for a zero trust security model

IBM Z provides protection for data and applications, which has been a hallmark of the platform for decades. Examples of security solutions that can help protect sensitive data on the IBM Z platform include:

- IBM Z pervasive encryption
- IBM Z Multi-Factor Authentication
- IBM Hyper Protect Data Controller

These security solutions can be used to support a zero trust security model that requires resolving the dilemma of need to access, need to know, and need to share sensitive data.

1.5.1 IBM Z pervasive encryption

The IBM z14® system introduced the concept of pervasive encryption. While encryption capabilities were offered by earlier IBM Z servers, the z14 made significant improvements in the speed and cost of performing both data-at-rest and data-in-flight encryption. In both cases, the ability to encrypt or decrypt data is transparent to applications and databases. Many companies find that enabling z/OS data set encryption for VSAM, Db2, or IMS databases adds little extra processor cost and little latency to time-sensitive transactions. This encryption is policy-based, making it easier to prevent users from accessing z/OS data because they do not have a need to know or use that data. It complements existing hardware-based encryption on storage systems.

z/OS data set encryption requires a person or application to have security permissions and access to two items before a file or database items can be viewed in the clear:

- The permission to access the file or database.
- The permission to access the operational or data encryption key that is used to protect the data.

Having only one of the two permissions results in the denial of access to the clear text data. These security permissions are proven and supported by new tools, utilities, and ISV software.

Consider implementing pervasive encryption on all z14 or later systems that have z/OS V2.3 or later with the needed APARs. Encryption is the foundation for protecting data on IBM Z. The IBM Z Cyber Vault solution works with data that is encrypted through z/OS data set encryption. Encrypted data can be encrypted multiple times, so having data that is encrypted with pervasive encryption does not provide protection from a ransomware attack.

For more information, see the following resources:

- IBM Z Pervasive Encryption content solution
- Encryption solutions

1.5.2 IBM Z Multi-Factor Authentication

IBM Z Multi-Factor Authentication (MFA) is the use of multiple factors to authenticate a person seeking access to a system. MFA authenticates new information in addition to the traditional use of a user ID and password.

The additional factors to authenticate include the following ones:

- Something that you have, such as a token or smartphone.
- Something that you are, such as a fingerprint or face or retina scan.
- Something that you know, such as a PIN code or password.

Using IBM Z MFA is best practice to reduce the ability of an attacker who has successfully completed a phishing or social engineering attack from gaining access to sensitive systems. The recent attack on SolarWinds was detected through MFA, as publicly noted by FireEye. They detected multiple phones seeking to perform MFA authentication, which was not normal. They contacted the employee, found out they did not have a new phone, and found the initial discovery of the SolarWinds attack. However, MFA is not a guarantee that cyberattacks will be prevented. The IBM Z Cyber Vault solution provides protection in case an attacker breaks through IBM Z MFA and gains access to an IBM Z system.

For more information, see IBM Z Multi-Factor Authentication.

1.5.3 IBM Hyper Protect Data Controller

IBM Hyper Protect Data Controller is a security offering that is supported on the IBM z15 and IBM z14 systems. With this offering, you can create encrypted data objects, which add policy-based security to any compatible data object regardless of where it travels or who tries to access it.

A potential user must request access to the encrypted data object from IBM Hyper Protect Data Controller, which determines whether the user is allowed to access the data. If access is allowed, the level of access the user is entitled to is specified. The access level includes hidden data, hashed data, encrypted data, or clear text data. This protection is available for data that is created on all common server platforms and operating systems and consumed on these same platforms and public cloud providers. The IBM Z Cyber Vault solution focuses on data that is hosted on z/OS systems and is complemented by the protection that is offered off-platform by IBM Hyper Protect Data Controller.

For more information, see IBM Hyper Protect Data Controller.

1.6 IBM Z Cyber Vault service offerings

The IBM Garage for Systems and IBM Lab Services teams offer a set of service offerings to help you plan and deploy your IBM Z Cyber Vault capabilities (see Figure 1-5).

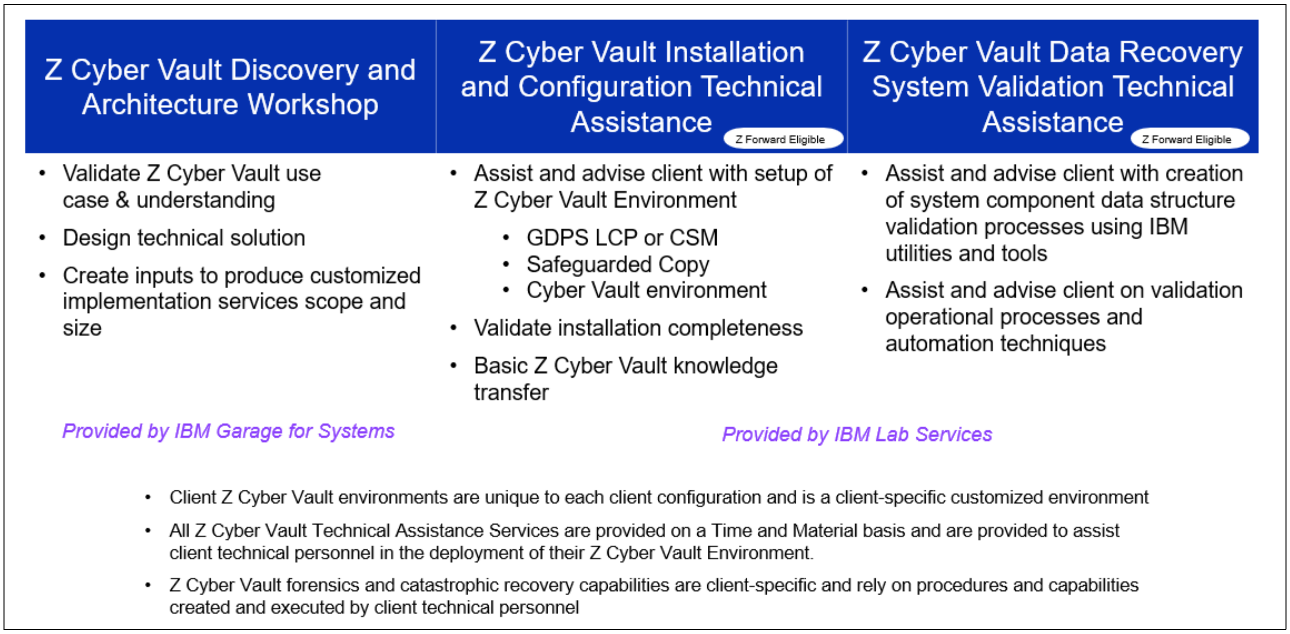

Z Cyber Vault Discovery and Architecture Workshop	Z Cyber Vault Installation and Configuration Technical Assistance (Z Forward Eligible)	Z Cyber Vault Data Recovery System Validation Technical Assistance (Z Forward Eligible)
• Validate Z Cyber Vault use case & understanding • Design technical solution • Create inputs to produce customized implementation services scope and size	• Assist and advise client with setup of Z Cyber Vault Environment • GDPS LCP or CSM • Safeguarded Copy • Cyber Vault environment • Validate installation completeness • Basic Z Cyber Vault knowledge transfer	• Assist and advise client with creation of system component data structure validation processes using IBM utilities and tools • Assist and advise client on validation operational processes and automation techniques
Provided by IBM Garage for Systems	*Provided by IBM Lab Services*	

- Client Z Cyber Vault environments are unique to each client configuration and is a client-specific customized environment
- All Z Cyber Vault Technical Assistance Services are provided on a Time and Material basis and are provided to assist client technical personnel in the deployment of their Z Cyber Vault Environment.
- Z Cyber Vault forensics and catastrophic recovery capabilities are client-specific and rely on procedures and capabilities created and executed by client technical personnel

Figure 1-5 Deployment service offerings for the IBM Z Cyber Vault solution

1.6.1 IBM Z Cyber Vault Discover and Architecture Workshop

The IBM Z Cyber Vault Discovery and Architecture workshop is a co-creation session with IBM Storage, GDPS, and IBM Z experts. The objective of the workshop is to understand your current environment and gain insights into cyber resiliency risks and challenges. During the workshop, the team defines the optimal IBM Z Cyber Vault architecture and creates a roadmap for implementing that solution.

During the workshop, the team discusses client objectives, solution options, and best practices. That information is used to design a solution that best meets your cyber resiliency objectives. At the completion of the workshop, the recommended architecture is documented along with an implementation roadmap and next steps.

Those next steps become part of the implementation plan and can include IBM Lab Services.

1.6.2 IBM Z Cyber Vault Installation and Configuration Technical Assistance

The IBM Z Cyber Vault Installation and Configuration Technical Assistance service provides IBM Storage and IBM Z implementation and consulting services. They include the installation and configuration of the IBM Z Cyber Vault solution components, validation of the environment, and the transfer of knowledge about the solution.

Specific services might include:

- Implement the Safeguarded Copy environment.
- Develop and implement a process for automatically creating Safeguarded Copy backups.
- Configure the IBM Z Cyber Vault environment.

- Verify the ability to perform an IPL of the system from Safeguarded Copy backups.
- Conduct a Cyber Incident Response recovery test Safeguarded Copy backups.

1.6.3 IBM Z Cyber Vault Data Recovery System Validation Technical Assistance

The IBM Z Cyber Vault Data Recovery System Validation Technical Assistance service provides advice and assistance with the design and creation of the system component data structure validation processes, also known as Type 2 validation, by using IBM utilities and tools. Because IBM Z Cyber Vault validation capabilities are custom-built for each specific client situation, this service relies on deep collaboration between the IBM Lab Services consultants and client systems programming and technical personnel. Advice and assistance on the design and creation of operational processes and automation techniques might also be part of this level-of-effort technical assistance.

Some areas of focus might include:

- How to confirm that the Safeguarded Copy backups can be used for recovery.
- Setup of IBM utilities to perform selected data structure validation for system components.
- Help with the selection and definition of use cases for IBM Z Cyber Vault Data Recovery System Validation.

The following items are true at the time of writing:

- IBM Z Cyber Vault environments are unique to each client configuration, and each one is a client-specific customized environment.
- All IBM Z Cyber Vault Technical Assistance Services are provided on a time and material basis and to help client technical personnel in the deployment of their IBM Z Cyber Vault environment.
- IBM Z Cyber Vault forensics and catastrophic recovery capabilities are client-specific and rely on procedures and capabilities that are created and run by client technical personnel.

For more information, see IBM Technology Technical Sales.

2

Planning and designing the IBM Z Cyber Vault environment

In this chapter, we describe a planning approach and the different components that are required to implement the IBM Z Cyber Vault solution. We describe planning and design considerations, and explain the prerequisites that are needed for the deployment of an IBM Z Cyber Vault environment.

The following topics are covered in this chapter:

- Planning approach for the IBM Z Cyber Vault environment
- IBM Z Cyber Vault reference architecture
- IBM Z Cyber Vault infrastructure design
- Prerequisites and considerations for implementing Safeguarded Copy
- Key operational considerations for Safeguarded Copy
- IBM Z platform considerations
- IBM z/OS considerations
- Application considerations
- Validation, forensic analysis, and recovery considerations

2.1 Planning approach for the IBM Z Cyber Vault environment

Before you implement an IBM Z Cyber Vault environment, you must consider your requirements:

- Which data do you need to protect?
- How do you want to isolate the IBM Z Cyber Vault environment from your production environment?
- How often do you want to create protected copies (backups / captures), and how long would you need to keep those copies?
- Which analysis and recovery capabilities exist today in your current production environment?
- How important are each of the IBM Z Cyber Vault capabilities to your organization?
- How can you achieve the correct balance between the depth of the IBM Z Cyber Vault capabilities, operational burden, and cost?

Implementing an IBM Z Cyber Vault environment requires separate 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 airgap systems or logical partitions (LPARs), if so designed, where you can run data validations, forensic analysis, or recovery that are based on protected copies of your data. Finally, to deliver on the IBM Z Cyber Vault capabilities, a set of IBM Z software tools are also required.

Software tools enable and define the level of automation, depth of analysis, and ease of use for the recovery of the solution, 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 minimizing risk and cost.

2.1.1 Implementation phases

When considering the deployment and usage of the IBM Z Cyber Vault solution, there is a logical progression for planning and implementation tasks that is a best practice. When implemented in a structured and managed manner, these tasks enable the IBM Z Cyber Vault solution to be built in a way that supports business goals and priorities. We describe the phases that guide you in determining the objectives of the IBM Z Cyber Vault implementation, the expected capabilities, and what needs to be considered to achieve these goals.

Collecting business and technical requirements

Before diving into 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. The reliance on digital technology to conduct business operations is expansive. Consider which IT services are the most critical among the ones that are delivered by the targeted IBM Z environments. Consider not only straightforward application services, but also what data services are provided by the IBM Z environment that can be relied upon by distributed processing or IT services that are not hosted on IBM Z.

For situations where the resources are constrained and a subset of services can be supported by the IBM Z Cyber Vault solution, consider developing a prioritization methodology to make the best choices for which IT assets will be supported in the IBM Z Cyber Vault environment. Many times, multiple dimensions of prioritization can be helpful, such as business criticality, revenue impact, or IT service outage sensitivity to BC.

Defining the IBM Z Cyber Vault configuration

After the appropriate IT services, both application and data, are identified and ranked based on business priority, the next step is to move toward defining and documenting the IBM Z Cyber Vault configuration and architecture. This task involves defining the specific IBM Z Cyber Vault configuration design to be used. It should include the number of LPARs within the IBM Z Cyber Vault and the role of LPARs in terms of support for Parallel Sysplex or IBM Geographically Dispersed Parallel Sysplex (GDPS) configurations.

Also, consider the customization of the IBM Z Cyber Vault I/O configuration for both mirrored I/O and any IBM Z Cyber Vault persistent resources definitions, such as storage for IBM Z Cyber Vault logging or event recording between Safeguarded Copy validation steps within the IBM Z Cyber Vault environment. Decisions about software configurations that will be required within the IBM Z Cyber Vault environment should be made now, and all the security and user access considerations should be evaluated now.

Because automation is key in guaranteeing a consistent, repeatable, and efficient validation of the IBM Z Cyber Vault environment, it should also be laid out now and checked during the Subsystem Data Validation stage of the IBM Z Cyber Vault operation.

Building and testing the IBM Z Cyber Vault environment

Building and testing the base IBM Z Cyber Vault environment involves the specific tasks that are associated with defining IBM Z Cyber Vault LPARs, network connectivity, and preparing IBM Z Cyber Vault for the first initial program load (IPL) and start. The customization of specific system and operational parameters and creation of any IBM Z Cyber Vault persistent resources are defined during this phase.

Establishing an IBM Z subsystem data structure validation design

Validating the structure of critical IBM Z subsystem resources such as catalogs, system data sets, and critical database structures provides the next level of resiliency checking and strengthening beyond a basic IPL operation. This step typically involves assembling a series of IBM Z utilities and tools to perform basic data structure checks in an efficient and expeditious way. In this area, each user determines the specific validation that they need, and build and customize their validation capabilities.

Designing application data validation

Beyond validating that a recovering system can be successfully started, ensuring that user application data is not compromised is a best practice. This activity depends on the format and nature of an application's data structures and formats. Referencing application data that might have been identified during the business and technical objectives and goals activities can be a way to guide and prioritize the focus in this area. Most installations have vast amounts of data, so prioritizing data to be validated will be essential. Considering BC and the criticality of various applications is a method to narrow the focus to the most important applications and data to be validated. Due to the custom nature of this data, customized programming by applications teams might be required to perform this type of validation.

Developing and testing validation capabilities and processes

Operational processes are important to use effectively the IBM Z Cyber Vault solution. Some of the areas to be addressed include defining the overall operational sequencing of daily IBM Z Cyber Vault procedures, automation strategies, and overall operational validation, and checking for successful validation analysis. The operation of the day-to-day IBM Z Cyber Vault process is important, and it can be performed by an individual or orchestrated across several operational actors. It can also be integrated into an installation's automation capabilities. Ensuring that each Safeguarded Copy is complete and valid for use is a key outcome of this activity.

Selecting 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 and more reliable, and help reduce downtime after a logical corruption incident. Standard utilities come with the operating system (z/OS) and IBM middleware products (such as Customer Information Control System (CICS), IBM Information Management System (IMS), Db2, and so on) for data validation and recovery purposes.

The tools and utilities that are described in this IBM Redbooks publication cover the most common scenarios. However, there might be other scenarios where you need to check different databases, control files, or other elements that you have not considered. The IBM Z Cyber Vault solution allows for the use of other tools and utilities that support your validation and recovery requirements.

For help with selecting IBM utilities and tools, see 1.6.3, “IBM Z Cyber Vault Data Recovery System Validation Technical Assistance” on page 17.

Testing the IBM Z Cyber Vault environment and operation

After the IBM Z Cyber Vault environment is defined and tested, it is time to activate it for operations. This task involves testing IBM Z Cyber Vault operational capabilities and procedures. Verify the basic IBM Z Cyber Vault environment with start and operate. Test the Type 1, 2, or 3 validation capabilities that were built for the IBM Z Cyber Vault environment (see 1.4.1, “Validation use cases” on page 11). Check to ensure that the operational and run procedures are documented and work properly. Make any adjustments that are necessary for smooth out future operation.

Production cutover and on-going monitoring

Production cutover and monitoring occur when the IBM Z Cyber Vault environment is started and perform as a part of the IT production environment. On-going monitoring is important to identifying and remediating any operational or technical adjustments that are needed for good operation. During this step, it is useful to consider periodic audit checks of the overall IBM Z Cyber Vault environment and operation. Ideally, this audit function should lie outside of the day-to-day IBM Z Cyber Vault operation team and be focused on inspecting and confirming a solid operation environment. Because a company's IT operations is always changing, the IBM Z Cyber Vault environment is affected by these changes. On-going monitoring should include the proper connections to Change Management to ensure that as environmental changes relate to the system, application, or data occur, the appropriate changes are also made to the IBM Z Cyber Vault environment.

2.1.2 IBM Z Cyber Vault environment operational roles

The following roles are critical to planning, deploying, and managing effectively an IBM Z Cyber Vault environment.

In addition, an essential attribute of any cybersecurity strategy is separation of duties between administrators. The concept of separation of duties suggests that more than one person is needed to complete a security-related task. This process helps avoid conflicts of interest and can better detect control failures that might lead to security breaches, information theft, and violations of corporate security controls and policies.

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 is important to effective planning and design. This role is responsible for overall design and close collaboration with operations to ensure efficient and effective day-to-day operations of IBM Z Cyber Vault.

Business continuity representative

This role represents the business requirements and works closely with the IBM Z Cyber Vault solution architect to ensure that BC, risk management, and recovery requirements are well understood. This role is also involved in helping to assess 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 generally available IBM Z utilities and tools. They also provide on-going technical support for the IBM Z Cyber Vault environment.

Application specialist and owners

This role understands various applications that take advantage of 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 specialist

This role performs day-to-day 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

These roles 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 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 administrator

Storage administrators play a key role in managing and using 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 administrator

Database administrators (DBAs) who must recover data, run forensics, and do other work in the IBM Z Cyber Vault environment will 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 will learn. This role is key for all IBM Z environments that use databases.

Network administrator

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 23.

2.2 IBM Z Cyber Vault reference architecture

Figure 2-1 depicts the reference architecture for the IBM Z Cyber Vault solution. It represents the infrastructure with the underlying components of the solution. This reference architecture can help locate the various components that are addressed in this chapter. The IBM Z Cyber Vault solution consists of two domains: the production environment and the IBM Z Cyber Vault environment. The IBM Z Cyber Vault environment has two subdomains: the recovery system environment and the storage environment. The storage environment (replication topology) spans the production environment.

For a list of the components that make up the solution, see 1.3.2, “IBM Z Cyber Vault solution technical summary” on page 8.

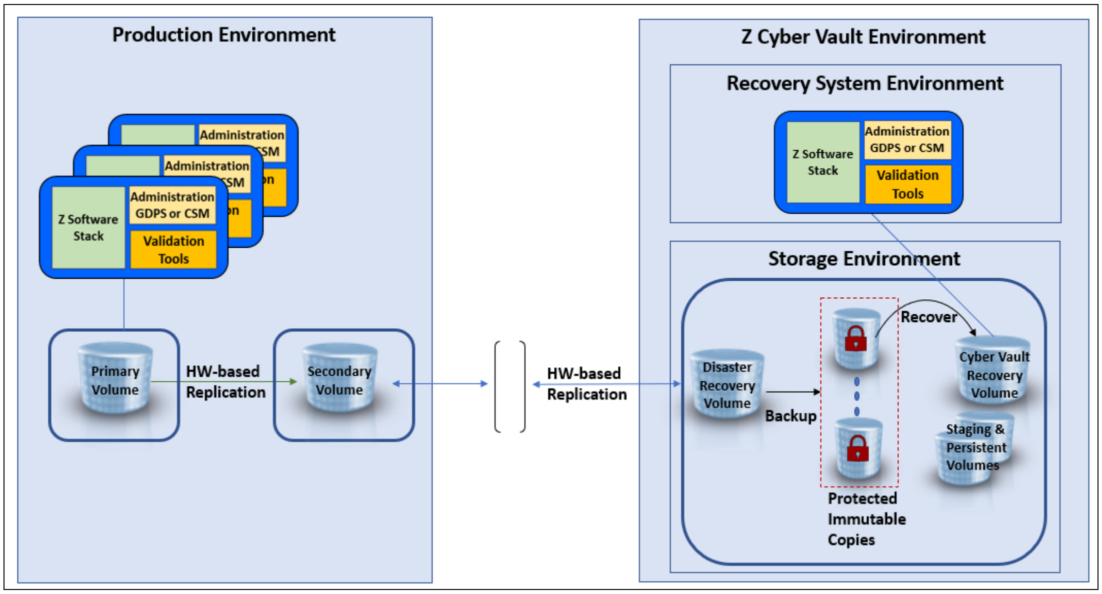

Figure 2-1 IBM Z Cyber Vault reference architecture (physically isolated)

Figure 2-1 on page 24 represents a physically isolated configuration that uses GDPS Metro Mirror (MM) (see “GDPS Logical Corruption Protection for the IBM Z Cyber Vault solution” on page 35). There are several infrastructure options that are supported by the IBM Z Cyber Vault solution, as described in 2.3, “IBM Z Cyber Vault infrastructure design” on page 34.

The storage environment and the recovery system environment are physically separated from the production systems (see 2.2.4, “Isolating the IBM Z Cyber Vault environment” on page 32). The distance between the production environment and IBM Z Cyber Vault environment can be hundreds of kilometers, like in a disaster recovery (DR) site.

The protected immutable copies (known as Safeguarded Copy backups) are on a separate storage environment and accessed only by the IBM Z Cyber Vault recovery system. From there, the Safeguarded Copy backups are restored on a separate pool of volumes, from which the recovery system performs an IPL, as part of the Type 1 validation (system data) process (see 1.4.1, “Validation use cases” on page 11).

In addition to the Safeguarded Copy backups, consider using *staging volumes*, and *persistent volumes*:

- Staging volumes are volumes that are used to temporarily store data sets, tables, or any other objects during validation or restoration to production.
- Persistent volumes are used for data that should not be lost between different validation runs, like historical data or Systems Management Facility (SMF) records, or data that is used to identify the last validated copy before corruption.

2.2.1 Production environment

Because the starting point for an IBM Z Cyber Vault recovery system is your existing z/OS production environment, it is where the implementation of the IBM Z Cyber Vault solution should begin.

The purpose is twofold: making your system more robust and resilient. When your system is robust, it means that it can endure cyberattacks without having to change, and keeping the expected service levels. Resiliency, in essence, is the ability to survive these cyberattacks despite severe impact. Here is where the IBM Z Cyber Vault recovery system plays a key role because it helps you be better prepared to recover faster, minimizing the impact of the attack.

We do not intend to cover how to make your production environment more robust because there are numerous publications and solutions that cover this aspect, but the implementation of some of the tools that are described here will improve the environment’s robustness while also increasing its resiliency.

Preparing the production environment for the IBM Z Cyber Vault solution involves adding some tools. These 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 recovery system, which keeps the software cost low (see “Software configuration and cost” on page 27). Other tools must be active in production because they must collect and create data and metadata that is used in the IBM Z Cyber Vault environment.

But the key point here is that there is no actual software installation or customization in the IBM Z Cyber Vault environment. All the preparations and configurations for the IBM Z Cyber Vault environment must be done in the production environment, which includes creating the special IPL procedures that select which products must be active or not in the IBM Z Cyber Vault environment.

The resulting software architecture is specific to each environment because it must consider existing tools and practices, and how the new tools will be integrated into the environment.

As a best practice, use the IBM Z Cyber Vault Discover and Architecture Workshop to make these decisions because it can help determine the cyber resiliency level and the cost of the IBM Z Cyber Vault solution. (See 1.6, "IBM Z Cyber Vault service offerings" on page 16.)

2.2.2 IBM Z Cyber Vault recovery system environment

The IBM Z Cyber Vault recovery system starts empty. Initially, it is just an LPAR that is configured in such a way that is isolated from the production environment.

When a Safeguarded Copy backup is selected to be used, it is first restored onto the IBM Z Cyber Vault recovery volumes, which are the volumes that are used to perform an IPL of the IBM Z Cyber Vault recovery system. The Safeguarded Copy is an exact replica of the production environment (see Figure 2-2). As a best practice, include all production system volumes in the Safeguarded Copy backups to ensure broader protection.

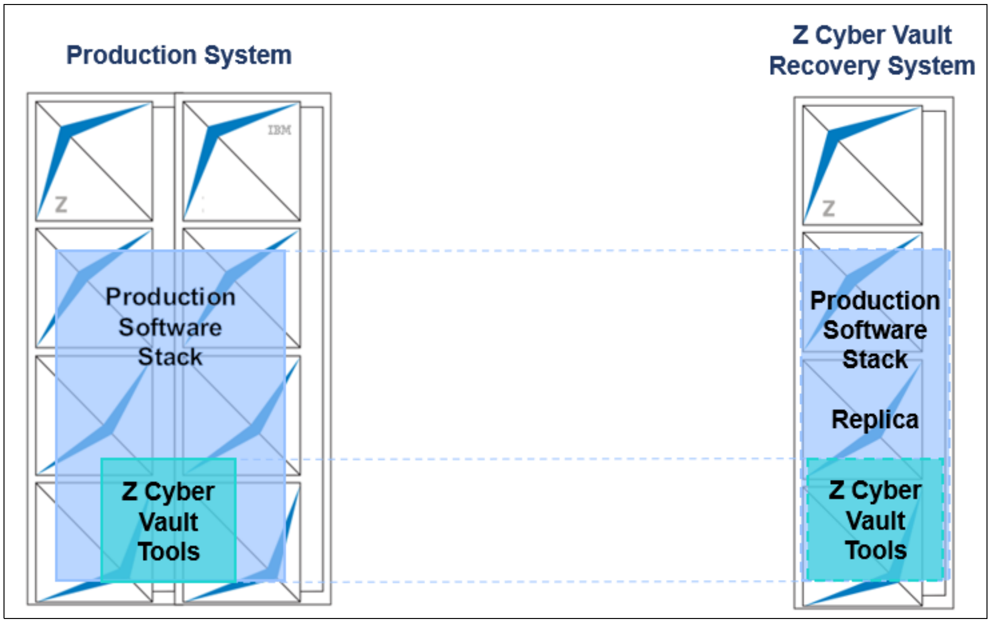
Production System
Z Cyber Vault Recovery System
Production Software Stack
Z Cyber Vault Tools
Production Software Stack
Replica
Z Cyber Vault Tools

Figure 2-2   IBM Z Cyber Vault recovery system loaded with a production software stack replica

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

After all the activities in the IBM Z Cyber Vault environment are complete, the recovery system (or LPARs) can be shut down, making them ready for the next Safeguarded Copy backups to be processed.

As a result of this mechanism, the IBM Z Cyber Vault environment remains isolated and Safeguarded Copy backups remain immutable. All the preparations and customizations that are required for the software stack that are used in the IBM Z Cyber Vault environment must be done in the actual production environment, which include preparing a special IPL procedure, setting up the required software, activating tools that will be collecting metadata and tracking activity, and getting all the automation and data validation tools ready.

Environments that are already running a GDPS environment or are using IBM Tivoli NetView for z/OS and IBM System Automation to automate their IPL procedures will benefit from these tools to automate the IBM Z Cyber Vault validation processes, which include the Type 1, Type 2, and Type 3 validation processes. For a description of the validation types, see 1.4, “Use cases for the IBM Z Cyber Vault solution” on page 11.

The IBM Z Cyber Vault recovery system is started from a point-in-time image of production, but it is started with special parameters that keep the network restricted to the IBM Z Cyber Vault environment and with limited user access.

The tools, utilities, and programs that are required to perform data structure and application data validation, forensic analysis, and surgical recovery or catastrophic recovery should be available in the IBM Z Cyber Vault recovery system. The tools for validation include IBM licensed software and other vendors’ software products.

Automation tools and procedures are available to run the data validation tasks as fast and error-free as possible.

These automation tasks and other operating procedures are defined with the user when implementing the solution because they vary depending on the tools that are selected for the IBM Z Cyber Vault environment, and existing recovery procedures and data management policies.

Designing and building the correct automation and data validation procedures is fundamental to the solution because these procedures determine how fast data corruption is detected, and how fast the system is validated (and able to move on to the next Safeguarded Copy backups). In 2.5.1, “Frequency of Safeguarded Copy backups” on page 47, we addressed the frequency of Safeguarded Copy backups, and in “Data validation frequency” on page 60, we provide extra considerations about the data validation frequency.

Software configuration and cost

Software in the IBM Z Cyber Vault recovery system has special pricing because there will be no business or production workload running in this system.

The full IBM software stack is licensed for limited use with the sole purpose of supporting all the capabilities that are available with the IBM Z Cyber Vault solution. Because this solution remains separate from your production, development, test, QA, and other environments, packaging it as an IBM Solution Edition can offer the best value because it does not affect existing contracts, and provides you with a fixed-term, full-capacity license of the entire software stack. With this limited use license for IBM software, you can run all the programs and procedures that are required to deliver the IBM Z Cyber Vault solution capabilities. Any production workload is excluded from this licensing.

The IBM Z Cyber Vault recovery system is different from your regular IBM Z operation. This system is started several times throughout the day. The objective is to run through the entire validation routine as quickly as possible so that the process can be repeated as often as needed. With full capacity licensing for the software that is running in the IBM Z Cyber Vault recovery system, there are no limitations to using all the available hardware capacity.

There might be software applications that you choose not to run in the IBM Z Cyber Vault environment, either because they cannot be supported in an isolated environment or you deem them as not critical to the restoration of your business. This decision always should be taken while understanding the benefits and risks of doing so, and any mitigation mechanisms to address the risks.

Balancing IT risk and the cost of adopting and operating the IBM Z Cyber Vault solution to protect your business from the consequences of cyberattacks is something that you must evaluate carefully. Data corruption events typically mean that the system, and your business, shut down at least until the source of the corruption and its scope are identified. How fast and thorough your validation, forensic analysis, and recovery capabilities are, and with them the cost of the mitigation solution, must be balanced with the risk of facing a cyberattack and not being able to restore your business.

2.2.3 IBM Z Cyber Vault storage environment

At the time of writing, IBM is the only storage vendor that supports Safeguarded Copy backups. The Safeguarded Copy function is integrated into the IBM DS8000 storage system and supported only on DS8880 and DS8900 storage system models with microcode Release 8.5 or later. Though the Safeguarded Copy is an IBM DS8000 feature that is best suited to supporting an IBM Z Cyber Vault project, it is possible to get started and provide IBM Z Cyber Vault capabilities with IBM FlashCopy® only.

Environments with non IBM storage (or older IBM Storage) can get started with the IBM Z Cyber Vault solution on storage platforms that support point-in-time copies.

Safeguarded Copy and FlashCopy

The DS8000 Copy Services (CS) functions (FlashCopy or Safeguarded Copy) are a fundamental part of IBM Z Cyber Vault. In this section, we provide a brief overview of the Safeguarded Copy function and its requirements. We compare Safeguarded Copy with DS8000 FlashCopy, which is another function that can be used in a logical corruption protection (LCP) context.

Using FlashCopy for logical corruption protection

FlashCopy can be used with IBM GDPS Logical Corruption Protection or IBM Copy Services Manager (CSM). With the introduction of Cascaded FlashCopy in DS8880 microcode Release 8.3, a concept was implemented that does not require direct access to the FlashCopy targets (see Figure 2-3 on page 29) to read the backup data. Instead of reading data directly from the FlashCopy targets, another cascaded FlashCopy can be established from one of the targets (F1, F2, or F3) to a recovery device R1 without losing the capability to restore one of the other FlashCopy copies if the first restored backup was already corrupted.

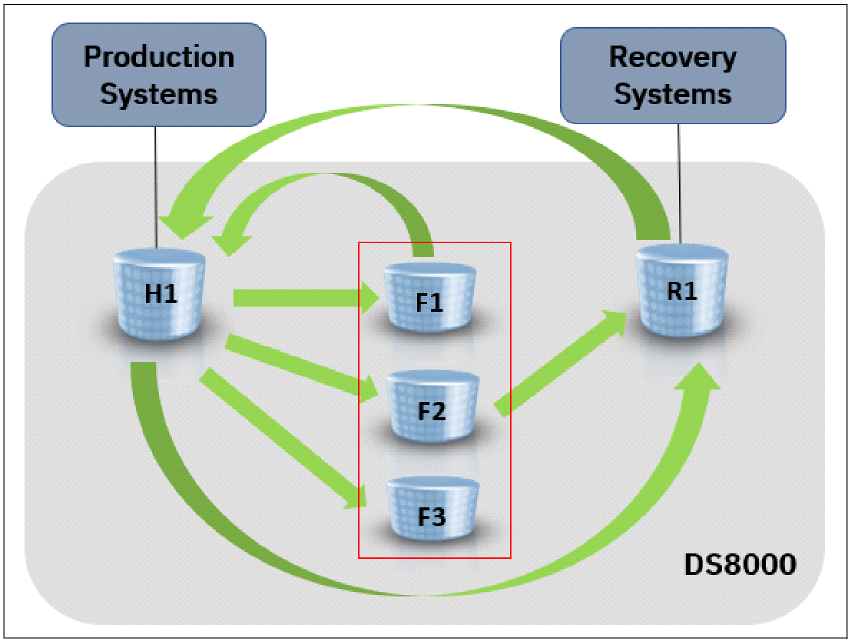

Figure 2-3 LCP with FlashCopy

GDPS Logical Corruption Protection Manager 4.1 and later or IBM CSM support using FlashCopy to establish an LCP solution, but because FlashCopy was originally designed with objectives other than LCP, it has some characteristics that do not fit ideally for this purpose:

- FlashCopy is limited to 12 relations per source volume for a maximum of 10 targets (point-in-time copies) in a GDPS Logical Corruption Protection concept.
- FlashCopy uses DS8000 device numbers, and therefore reduces the number of possible host volumes.
- FlashCopy targets are regular volumes. Therefore, they can be accessed by hosts and potentially be modified or corrupted, so extra efforts must be made to secure them.
- FlashCopy relations and targets can be deleted by the DS8000 administrator.
- Source volumes with multiple targets can suffer from a write performance impact because all targets are maintained individually.
- Space usage that uses thin-provisioned devices and the 21-cylinder allocation unit that is used for extent space efficient (ESE) devices is optimized for host and FlashCopy performance, but is not as efficient as a smaller (for example, track-level) allocation unit in Safeguarded Copy.

Safeguarded Copy does not have these limits. It is self-protected by design, allows more backups (up to 500), and is more space-efficient by using track level allocation, compared to the extent-level unit that is used by FlashCopy. Safeguarded Copy stores a changed track only once.

Performance-wise, you can consider the extra workload for Safeguarded Copy as one FlashCopy no-copy operation for all Safeguarded Copy backups (the number of backups does not matter). FlashCopy requires more DS8000 internal bandwidth if multiple FlashCopy relationships exist.

Using Safeguarded Copy for logical corruption protection

Safeguarded Copy provides immutable point-in-time copies to prevent data from being modified or deleted. The point-in-time copies (Safeguarded Copy backups) are not accessible by systems or applications because they do not have any logical control units (LCUs) or volume IDs. When backups exist for a volume, the volume (Safeguarded Copy Source volume) and its backups cannot be erased or deleted by using the DS8000 native interfaces (DS Command-line Interface (DS CLI) or DS GUI). To access a Safeguarded Copy backup, a recovery action to a so-called *recovery volume* is necessary so that the data can be accessed from a *recovery system* (see Figure 2-4).

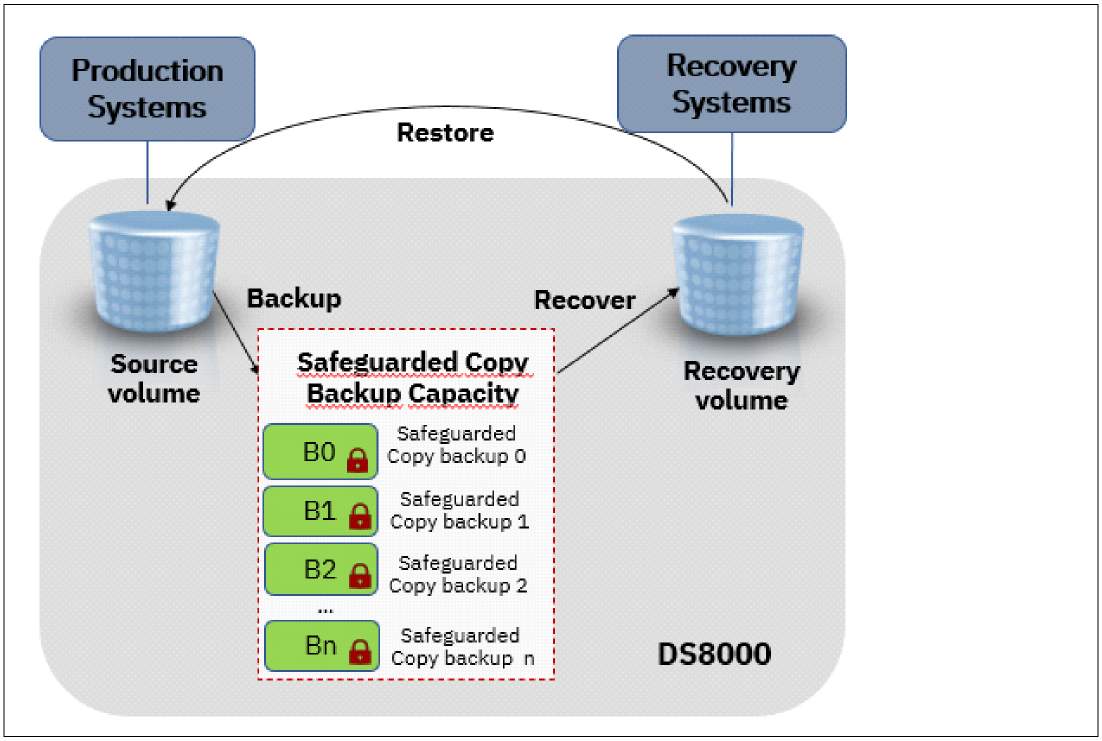
Production Systems
Recovery Systems
Restore
Backup
Recover
Source volume
Recovery volume
Safeguarded Copy Backup Capacity
B0
Safeguarded Copy backup 0
B1
Safeguarded Copy backup 1
B2
Safeguarded Copy backup 2
...
Bn
Safeguarded Copy backup n
DS8000

Figure 2-4 Safeguarded Copy overview diagram

With Safeguarded Copy, you can have up to 500 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 in a thin-provisioned Safeguarded Copy Backup Capacity copy. In addition, Safeguarded Copy backups do not require a DS8000 Volume ID or a host device address (UCB), and the function can be integrated with different DR or high availability (HA) configurations.

To manage, create, recover, and expire Safeguarded Copy backups, management software like IBM CSM or GDPS Logical Corruption Protection Manager is required.

Before you use Safeguarded Copy, you must specify in the IBM DS8000 the amount of space for each volume that you want to use for backups. The required capacity depends on the data change rate, the number of backups, and the time that you want to keep the backups (retention period). Therefore, you define a so-called Backup Capacity Multiplier per volume for Safeguarded Copy (see Figure 2-5 on page 31).

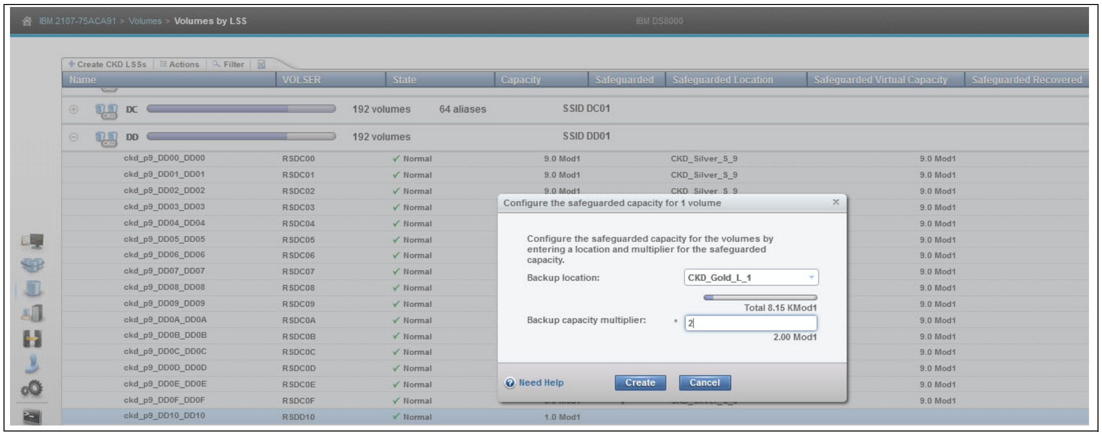

Name	VOLSER	State	Capacity	Safeguarded	Safeguarded Location	Safeguarded Virtual Capacity	Safeguarded Recovered
DC	192 volumes	64 aliases	SSID DC01				
DD	192 volumes		SSID DD01				
ckd_p9_DD00_DD00	RSDC00	✓ Normal	9.0 Mod1		CKD_Silver_S_9	9.0 Mod1	
ckd_p9_DD01_DD01	RSDC01	✓ Normal	9.0 Mod1		CKD_Silver_S_9	9.0 Mod1	
ckd_p9_DD02_DD02	RSDC02	✓ Normal	9.0 Mod1		CKD_Silver_S_9	9.0 Mod1	
ckd_p9_DD03_DD03	RSDC03	✓ Normal				9.0 Mod1	
ckd_p9_DD04_DD04	RSDC04	✓ Normal				9.0 Mod1	
ckd_p9_DD05_DD05	RSDC05	✓ Normal				9.0 Mod1	
ckd_p9_DD06_DD06	RSDC06	✓ Normal				9.0 Mod1	
ckd_p9_DD07_DD07	RSDC07	✓ Normal				9.0 Mod1	
ckd_p9_DD08_DD08	RSDC08	✓ Normal				9.0 Mod1	
ckd_p9_DD09_DD09	RSDC09	✓ Normal				9.0 Mod1	
ckd_p9_DD0A_DD0A	RSDC0A	✓ Normal				9.0 Mod1	
ckd_p9_DD0B_DD0B	RSDC0B	✓ Normal				9.0 Mod1	
ckd_p9_DD0C_DD0C	RSDC0C	✓ Normal				9.0 Mod1	
ckd_p9_DD0D_DD0D	RSDC0D	✓ Normal				9.0 Mod1	
ckd_p9_DD0E_DD0E	RSDC0E	✓ Normal				9.0 Mod1	
ckd_p9_DD0F_DD0F	RSDC0F	✓ Normal				9.0 Mod1	
ckd_p9_DD10_DD10	RSDD10	✓ Normal	1.0 Mod1				

Configure the safeguarded capacity for 1 volume

Configure the safeguarded capacity for the volumes by entering a location and multiplier for the safeguarded capacity.

Backup location: CKD_Gold_L_1

Total 8.15 KMod1

Backup capacity multiplier: * 2

2.00 Mod1

Figure 2-5 Creating a volume Backup Capacity Multiplier

To store the data changes, storage capacity in the DS8000 is required. Without a Safeguarded Copy backup, the capacity is pure virtual capacity; physical capacity is allocated as you create backups; and data that is overwritten in the original volume is saved in the backups. Backup data is saved with track granularity, which leads to better efficiency than with FlashCopy, which stores changes at the extent level.

After creating the Safeguarded Copy Backup Capacity for all volumes that should be part of your Safeguarded Copy backups, you can set up and manage Safeguarded Copy with your management software. The management software is used to initiate, recover, and expire backups.

Initiating Safeguarded Copy backups

When you initiate a Safeguarded Copy backup with your management software, a consistency group (CG) is created across all involved volumes and the DS8000 system. 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 copies data that was overwritten by host I/O from the Safeguarded Copy source volume to the CG log within the Safeguarded Copy Backup Capacity.

When you initiate the next backup, the DS8000 system closes the previous backup and creates a CG. Therefore, it does not have to maintain each backup individually. To restore a certain backup, the DS8000 system needs all the backups that are younger than the one that you selected to recover.

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

1. Reservation: In this step, the DS8000 sets up the required bitmap and prepares the metadata in the Safeguarded Copy Backup Capacity. It also makes sure that all changed data from the previous backup is stored. After all preparations are done, the actual CG formation can take place.
2. Check in: To create a CG, the DS8000 system must stop updates for all volumes within the CG for a short period. It does this task by changing to an Extended Long Busy (ELB) state. When the data in cache is consistent, the previous CG logs of all affected volumes are closed, and therefore are also consistent. From now on, the DS8000 system writes further backup data into the CG logs of the new backup.
3. Completion: The DS8000 lifts the ELB and write operations can continue.

The management software (either GDPS Logical Corruption Protection Manager or CSM) coordinates and performs these steps automatically and with minimal impact to the host operations, but consider the ELB time for your host write operations.

Recovering Safeguarded Copy backups

To access a Safeguarded Copy backup, a recovery action is required. 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 recovery 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 restore from a certain backup, the DS8000 system needs all backups that are younger than the one that will be recovered.

During the recovery action, the DS8000 system initiates a FlashCopy copy between the Safeguarded Copy source volume and the recovery volume, which makes both volumes identical. 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 younger backups rather than from the Safeguarded Copy source volume.

Now, you have read/write access to the recovery volumes. 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. If the same track shows up in more than one backup, you must use the *oldest* instance because it is closest to the backup that should be recovered.

Expiring Safeguarded Copy backups

With your management software, you can expire Safeguarded Copy backups manually or automatically after the retention period of those backups is over. Because the backups depend on each other, expiring a certain backup expires all older backups too.

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 copies per volume.
- A DS8000 storage pool runs out of physical space.
- The specified backup capacity (backup multiplier) for a particular volume is too small.

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

2.2.4 Isolating the IBM Z Cyber Vault environment

When implementing the IBM Z Cyber Vault solution, you must decide where and how you want to isolate, capture, and manage your protected data copies. Two different topologies are possible: *virtual isolation* and *physical isolation*, also referred as a *virtual airgap* model or *physical airgap* model.

With virtual isolation, the protected backup copies are created on one or more of the storage systems in your existing HADR topology. The storage systems are typically in the same physical SAN or IP network as the production environment, and perhaps the same group of administrators are managing both the HADR solution and the IBM Z Cyber Vault environment.

The protected backup copies are not addressable by the operating system, which makes these copies virtually isolated from the z/OS operating system, as shown in Figure 2-6.

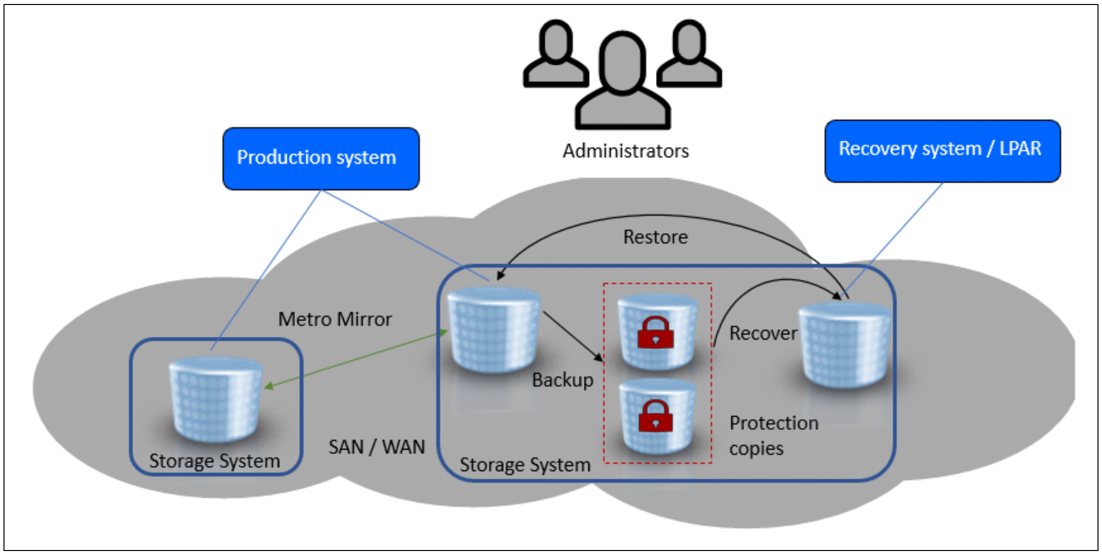

Figure 2-6 Virtual isolation

Using a virtual isolated environment for protected copies requires less resources compared to a physically isolated solution, but for scenarios like insider threats, virtual isolation provides less protection.

A physically isolated environment is when the protected copies exist on a storage system that is different from the one hosting the production data. Within a physically isolated environment, the data copies are not captured directly from the production data volumes. Instead, an extra replication (Global Copy (GC) / Global Mirror (GM)) leg is configured and cascaded from the production data volumes, which avoids the ELB times on the production volumes. These storage systems are typically not on the same physical SAN or IP network as the production environment. The systems might have restricted access, perhaps even with different administrators to provide separation of duties, as shown in Figure 2-7.

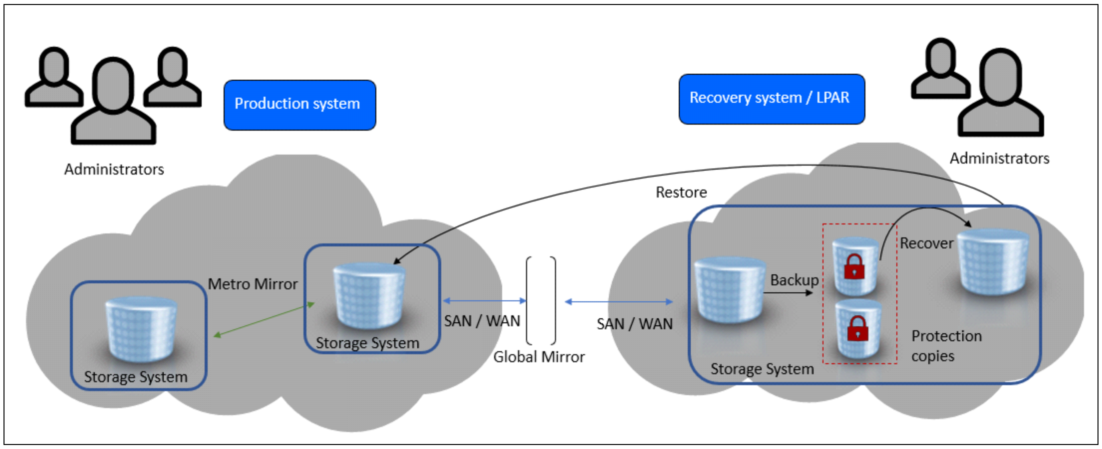

Figure 2-7 Physical isolation

A physically isolated environment requires more resources and investment, but provides better protection against various scenarios like insider threats. A virtual tape library that uses write once read many (WORM) tapes along with a large expiration date in the physical isolated environment to store more backups would improve the protection.

2.2.5 Network isolation for the recovery system

The recover system in the IBM Z Cyber Vault environment should run in an isolated network to ensure that the Safeguarded Copy backups cannot be accessed from the production environment or any other system. This isolation can be 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 can be used to define logical separation by configuring a separate virtual LAN (VLAN) and IP subnetwork with optional firewalls.

In addition, both the recovery system environment and the production system environment can be protected by using built-in z/OS Communications Server security features, such as:

- IP filtering blocks out all IP traffic that this system does not explicitly permit within its defined IP Filter Policy.
- Intrusion Detection Services (IDS) protect against attacks of various types on the system's services.
- Application Transparent Transport Layer Security (AT-TLS) provides SSL/TLS encryption services at the TCP transport layer to protect and secure “in-flight” sensitive application data. AT-TLS is transparent to the application.

Because a recovery system is started by using a replica of a production system, including its configuration, all the security features should be implemented in the production environment too.

2.3 IBM Z Cyber Vault infrastructure design

In this section, the administration and automation tools that are used for replication management are introduced, and various architectural options for the IBM Z Cyber Vault infrastructure.

2.3.1 Replication management

There are two IBM products that can be used to help manage and monitor the IBM Z Cyber Vault storage environment:

- IBM CSM

  CSM controls CS in storage environments. CS consists of features that are used by IBM storage systems, such as IBM DS8000 or IBM FlashSystem® to configure, manage, and monitor data replication. CSM runs on Windows, IBM AIX®, Linux, Linux on IBM Z, DS8000 Hardware Management Console (HMC), and z/OS operating systems. CSM is focused on CS and HyperSwap management, but it does not include any z/OS LPAR automation or LPAR monitoring functions.

  If you plan on using CSM to manage Safeguarded Copy in an IBM Z Cyber Vault environment, see *IBM DS8000 Safeguarded Copy (Updated for DS8000 R9.2)*, REDP-5506 for more information about the supported topologies.

- IBM GDPS

  GDPS provides copy services management (as CSM does) and also z/OS LPAR orchestration and automation. GDPS is an IBM Z centric family of solutions that support availability and DR goals. There are different solutions 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 either in the production or DR sites (or both), with appropriate management of Parallel Sysplex resources and orchestration of workflows to drive recovery actions when needed.

  A key part of any GDPS solution is to ensure that there is always a point of consistency for the production data that can be used to restore service from an alternative location. This point is often referred to as a *crash* or *power-fail consistent copy of data*. Different GDPS solutions employ different mechanisms that use the capabilities of IBM Z environment to achieve this point.

  GDPS provides automation capabilities and procedures that support the IBM Z Cyber Vault environment and greatly simplify operations for data validation (see Chapter 4, "Establishing a validation framework" on page 103).

Both replication management tools support implementation of virtually isolated or physically isolated environments. Depending on the selected isolation method, the implementation approach is different. For more information, see 2.2.4, "Isolating the IBM Z Cyber Vault environment" on page 32.

2.3.2 IBM GDPS capabilities and supported models

IBM GDPS is a set of different products that was developed to help achieve user-defined goals for continuous availability and DR. GDPS provides specific advantages when it is used with the IBM Z Cyber Vault solution, such as automation, protected copy management, Safeguarded Copy support, restore capabilities, and integration with existing GDPS environments.

GDPS Logical Corruption Protection for the IBM Z Cyber Vault solution

GDPS Logical Corruption Protection is a feature that is integrated with the GDPS family of products that is aimed at helping with recovery from cyberattacks, internal threats, and other forms of logical data corruption through administration and automation capabilities.

In the IBM Z Cyber Vault environment, the ability to test and validate point-in-time copies regularly and quickly make automation nearly mandatory, and GDPS Logical Corruption Protection can help satisfy this requirement.

GDPS Logical Corruption Protection can capture multiple, secure point-in-time copies of critical production data and restore the data back in to production, if necessary. GDPS Logical Corruption Protection also can recover a particular point-in-time copy from another set of volumes that can be used to start one or more isolated recovery systems, and to help analyze the scope of a distinct logical corruption event.

More security and protection capabilities are provided for the GDPS Logical Corruption Protection protection copies (Safeguarded Copy backups) than for copies that are taken with more traditional methods like FlashCopy because GDPS Logical Corruption Protection minimizes host access to these volumes by providing specific roles and rules (policies) for their management.

GDPS Logical Corruption Protection also provides users with key capabilities to minimize the risk of errors and reduced manual operations. For example, to simplify the administration of Safeguarded Copy and FlashCopy captures, GDPS Logical Corruption Protection has defined management profiles.

A GDPS Logical Corruption Protection management profile describes the management characteristics of the volume captures that are taken:

- The replication site (RS) where the captures are to be taken.
- The consistency group (CG) to be captured.
- 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 Safeguard Reservation Scan phase.
- The maximum permissible elapsed time for the Safeguard Check In phase.
- The maximum permissible time allowed for a Consistent Group Pause to complete in preparation for capture processing.

Another key capability of GDPS Logical Corruption Protection is the 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, such as:

- CAPTURE: This statement performs a consistent capture of the RS(n) volume set to an FC(n) or SGC(n) copy set.
- RELEASE: This statement performs a release of one or more expired captures from an FC(n) or SGC(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 recovery copy (RC) copy set from either an FC(n) or an SGC(n) copy set. When recovered, it is then possible to access the RC(n) copy set for data analysis, extraction, or test purposes.

GDPS Logical Corruption Protection models for the IBM Z Cyber Vault solution

The GDPS Logical Corruption Protection Manager can be deployed in several different ways to provide the required combination of isolation from production systems (either through a virtual or physical air gap) and recovery capabilities at the recovery site, in an IBM Z Cyber Vault solution. With GDPS Logical Corruption Protection Version 4.4, the following models are supported:

- GDPS Metro virtual airgap model
- GDPS Global virtual airgap model
- GDPS Metro physical airgap model
- GDPS Metro/Global Mirror 3-Site virtual airgap model
- GDPS MGM 4-Site physical airgap model

Common acronyms that are used in the GDPS Logical Corruption Protection models include:

- RS(n) represents the source volumes, where RS1 is the primary disk, and RS2 is the secondary disk. The RS(n) volume sets are the conventional RS volume sets that form the basis of the GDPS replication model and an IPL can be performed on them in a primary status.
- CS(n) is the copy set volumes (Safeguarded Copy or FlashCopy). Copy sets are a grouping of volumes that together can provide a consistent point in time copy of data.
- RC(n) is the RC set volumes that are used for recovery testing that are not Target Write Inhibited. Therefore, you can perform an IPL of a recovery system in support of forensic analysis or other recovery tests.
- A Kp controlling system runs in the production site. It monitors the resources and carries out reconfiguration and recovery actions for planned and unplanned outages.
- A Kr controlling system runs in the recovery site and manages the Safeguarded Copy backups.
- A Kg controlling system runs in the production site and sends information to the systems running in the recovery site.

GDPS Metro virtual airgap model

In the GDPS Metro virtual airgap model (see Figure 2-8), the internal GDPS Logical Corruption Protection Manager (MM) is operating in an MM solution. The internal GDPS Logical Corruption Protection Manager function is hosted and runs in the Kp controlling system. The copy sets (CS1, CS2, and CS3) that are managed by the GDPS Logical Corruption Protection Manager (MM) can be configured in A.RS1, A.RS2, or both.

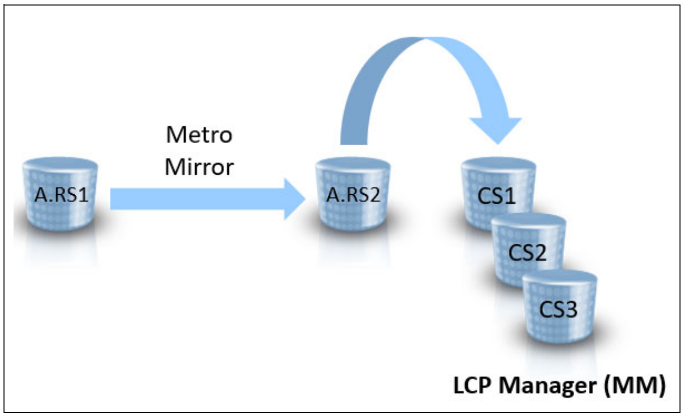

Figure 2-8 GDPS Metro virtual airgap model

GDPS Global virtual airgap model

In the GDPS Global virtual airgap model (see Figure 2-9), the GDPS Logical Corruption Protection Manager (GM) operates in a GM environment. The internal GDPS Logical Corruption Protection Manager function is hosted and runs in the Kr controlling system. The GDPS Kg controlling system is used to set up, monitor, and manage the GM, and is also responsible for managing the CGPAUSE function that the GDPS Logical Corruption Protection Manager uses to create consistent copy sets as part of the CAPTURE script command.

The copy sets that are managed by the GDPS Logical Corruption Protection Manager (GM) can be configured only at the recovery site (labeled as B.RS1).

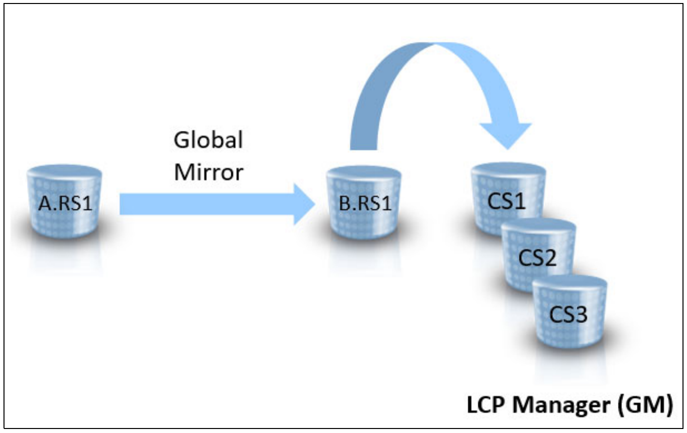

Figure 2-9 GDPS Global GM virtual airgap model

GDPS Metro physical airgap model

In the GDPS Metro physical airgap model (see Figure 2-10), the GDPS Logical Corruption Protection Manager (GM) is external and not tied to the GDPS Metro solution to provide a physical airgap. The external GDPS Logical Corruption Protection Manager function is hosted and runs on the GDPS Logical Corruption Protection Manager (GM) in the GDPS Kr controlling system.

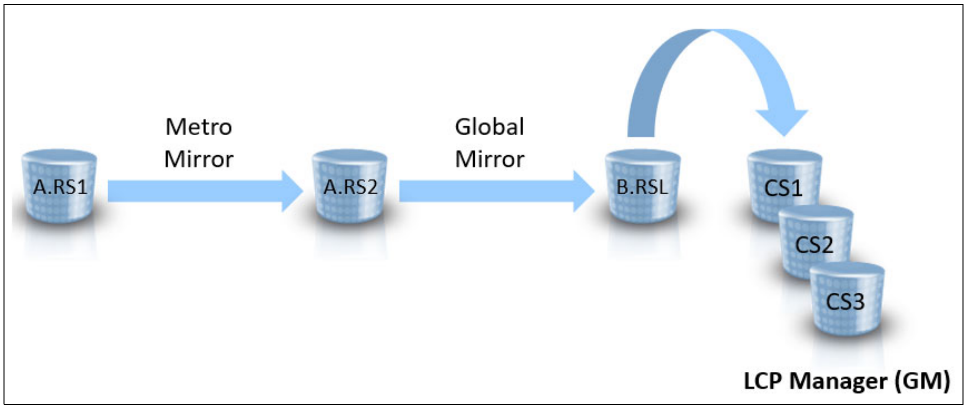

Figure 2-10 GDPS Metro physical airgap model

The copy sets in this configuration rely on Safeguarded Copy backups and are captured only from the GM secondary volumes in the recovery site, which are labeled as “B.RSL”.

The GDPS Kg controlling system is used to set up, monitor, and manage the GM, and it is responsible for managing the CGPAUSE function that the GDPS Logical Corruption Protection Manager uses to create consistent copy sets as part of the CAPTURE script command.

GDPS Metro/Global Mirror 3-Site virtual airgap model

In the GDPS Metro/Global Mirror (MGM) 3-Site virtual airgap model (see Figure 2-11), there are two internal GDPS Logical Corruption Protection Managers (GDPS Logical Corruption Protection Manager (MM) and GDPS Logical Corruption Protection Manager (GM)), which both operate in an MGM 3-Site solution. The internal GDPS Logical Corruption Protection Manager (MM) function is hosted and runs in the GDPS Metro Kp controlling system, and the internal GDPS Logical Corruption Protection Manager (GM) function is hosted and runs in the Kr controlling system.

With this model, you can have a captured copy either on the MM volume or on the GM secondary volume (or on both).

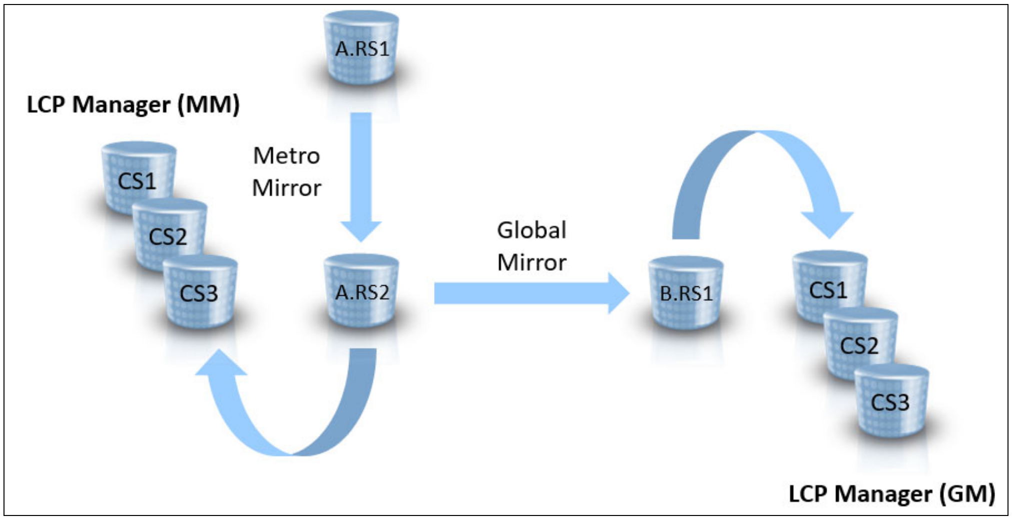

Figure 2-11   GDPS MGM 3-Site virtual airgap model

GDPS MGM 4-Site physical airgap model

In the GDPS MGM 4-Site physical airgap model (Figure 2-12), captured copies are isolated physically from both the production site and the DR site.

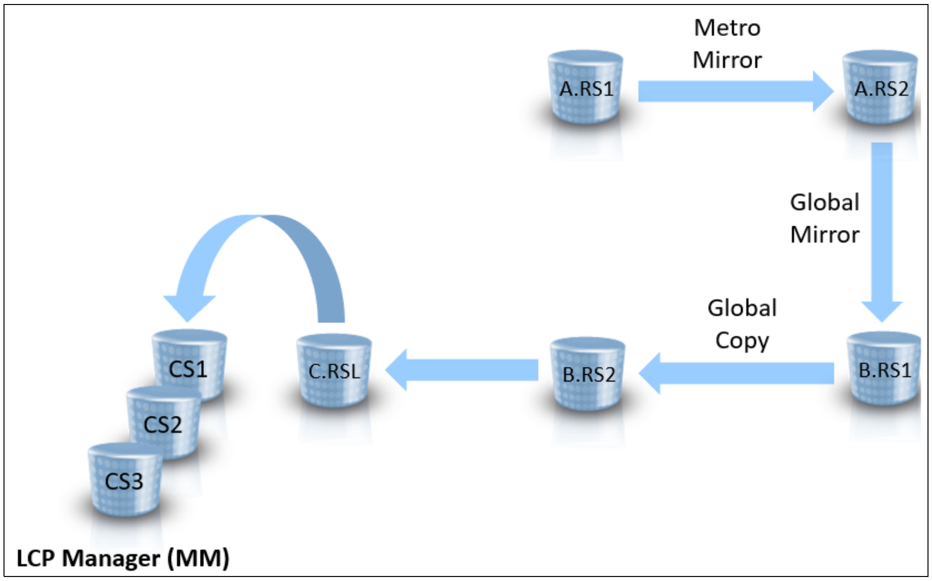

Figure 2-12 GDPS MGM 4-Site physical airgap model

The LCP environment in a GDPS MGM 4-site solution is managed by an instance of GDPS Metro that runs on a separate controlling system from the GDPS MGM controlling system.

Capturing an external LCP protection copy requires task coordination between the GDPS Logical Corruption Protection Manager controlling system and the controlling systems that make up the GDPS MGM 4-site solution. The key coordination task is to create a consistent data point on the B.RS2 volume set so that it can be replicated to the C.RSL volume set and then captured to a corresponding CS(n) copy set. Because GC does not provide data consistency, GDPS has automated a sequence of actions to allow a consistent point-in-time copy to be captured on the physically separated LCP volume environment.

When a DR exercise is in process (test or real event), access to point-in-time copies (captured copies) are always possible.

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

2.4 Prerequisites and considerations for implementing Safeguarded Copy

Before implementing Safeguarded Copy, you must fulfill some prerequisites, and you should consider some aspects that are required for a successful Safeguarded Copy implementation.

2.4.1 Safeguarded Copy prerequisites

The Safeguarded Copy function is integrated into the IBM DS8000 microcode, and it is supported on DS888x storage system models with microcode Release 8.5 or later and on all DS8900F storage system models.

To start using Safeguarded Copy, you must have a 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 V6.2.3.1 or later or GDPS Logical Corruption Protection Manager V4.2 SP2 or later is required. You cannot use the DS8000 interfaces DS GUI / DS CLI or a z/OS interface to initiate, recover, or expire a Safeguarded Copy backup.

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

- The changed data that is stored in Safeguarded Copy Backup Capacity over the retention period
- The small Safeguarded Copy overhead for each backup
- Recovery volumes
- Safeguarded Copy Source volumes (for physical isolation)

2.4.2 Safeguarded Copy considerations

Consider the following points 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 less than 10 minutes, IBM requires that you submit a request price quotation (RPQ) for approval and support.
- The maximum Safeguarded backup capacity for a volume is 16 TiB.
- A Safeguarded Copy source cannot be a FlashCopy target.
- The source and recovery volumes must be managed by the same DS8000 internal server. Therefore, they must both be either in an even or odd logical subsystem (LSS).
- 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 DS8000 Safeguarded Copy (Updated for DS8000 R9.2)*, REDP-5506.

2.4.3 Sizing extra storage for Safeguarded Copy

In this section, we describe the capacity and performance sizing for Safeguarded Copy. We indicate which information is required for the Safeguarded Copy capacity sizing. In addition, we describe performance considerations when Safeguarded Copy is used.

Safeguarded Copy sizing overview

It is crucial to do an accurate Safeguarded Copy capacity sizing. It is 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 steps 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 will be used in different use cases.
- Size the Safeguarded Copy recovery volume and source volume 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 an DS8000 system depend mainly on cache size of the system, as shown in Table 2-1.

Table 2-1 Capacity limits of an IBM DS8000

System memory	Maximum number of physical extents	Maximum number of volume extents	Maximum physical size small extent (Fixed Block (FB) or Count Key Data (CKD)	Maximum virtual size small extents (FB/CKD)
<= 512 GB	32 million	64 million	512 TB (FB)	1024 TB (FB)
			551 TiB (CKD)	913 TB (CKD)
> 512 GB?	128 million	256 million	2024 TB (FB)	3968 - 4096 TB (FB)
			2206 TB (CKD)	3538 - 3652 TB (CKD)

It is necessary to estimate the physical and virtual capacity of the following components in the DS8000 storage system, as shown in Figure 2-13 on page 43:

- Safeguarded Copy Backup Capacity.
- Recovery volume.
- Safeguarded Copy source volume if a physical isolation approach will be implemented.

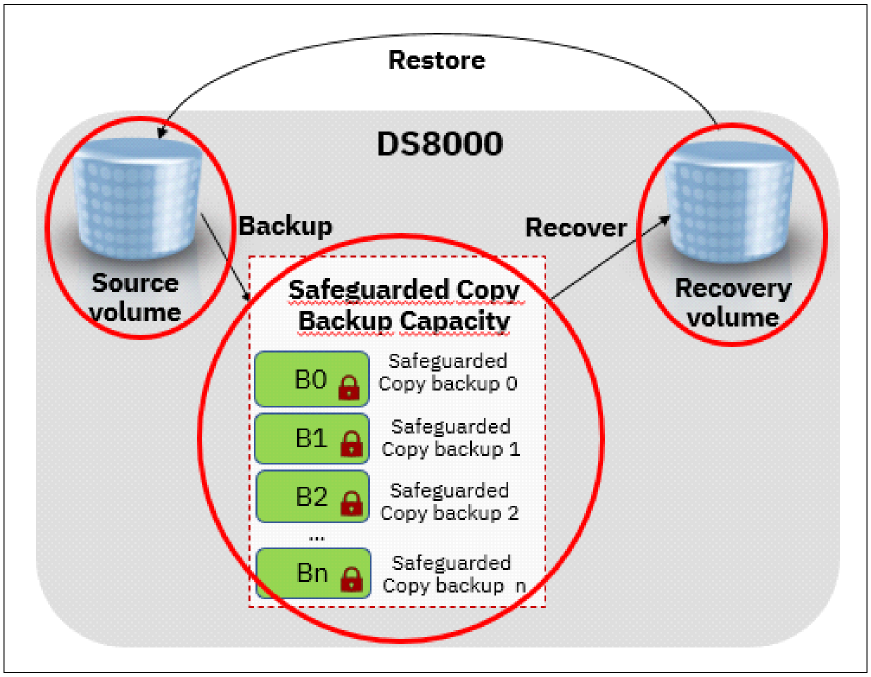
Restore
DS8000
Source volume
Backup
Recover
Recovery volume
Safeguarded Copy Backup Capacity
B0
Safeguarded Copy backup 0
B1
Safeguarded Copy backup 1
B2
Safeguarded Copy backup 2
...
Bn
Safeguarded Copy backup n

Figure 2-13 Components for Safeguarded Copy that must be sized

Physical capacity estimation is required to determine how much capacity is required to implement Safeguarded Copy (for example, 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 following backup management policies:

- Frequency of backups to be taken
- Retention period for the backups

For example, if you were creating a backup every 6 hours and retaining it for 48 hours, you would need to know the total data change rate over a 48-hour retention period.

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 change while the relationship exists.

Note: It is a best practice to add ~ 20% of the physical source volume capacity for the recovery volumes so that you can use the recovery volumes for data validations during normal operations.

The sizing for the Safeguarded Copy capacity (physical and virtual) and recovery volumes can be done by using the methods and modeling in “Safeguarded Copy capacity sizing” on page 44 or “Safeguarded Copy performance sizing” on page 44.

The methods determine the data change or destage rate in tracks. This absolute number is then 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 interval for as many intervals as the length of the retention period. You must do this task for each involved DS8000 system to calculate the required physical capacity for the Safeguarded Copy backups.

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

Note: Use the simple approach for determine the Backup Capacity Multiplier per volume only if a low amount of Safeguarded Copy is required and the source volume capacity is small. Only then will you not reach a DS8000 capacity limit.

Safeguarded Copy capacity sizing

Different methods are available to do a capacity sizing for Safeguarded Copy. 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 functionality*.

- Analyzing performance data, such as Resource Measurement Facility (RMF) data in z/OS or IBM Storage Insights.

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

  For more information, see *IBM DS8000 Safeguarded Copy (Updated for DS8000 R9.2)*, REDP-5506.

Safeguarded Copy performance sizing

In addition to Safeguarded Copy capacity sizing, you should do a performance sizing of the new or upgraded DS8000 systems by contacting your IBM representative.

For performance considerations, consider the extra workload for Safeguarded Copy in your DS8000 system as one FlashCopy no-copy operation that constantly runs for all involved volumes.

Here are some performance considerations:

- Hybrid DS8880 performance considerations

  There are two kinds of metadata that are used in this case:

  - Global metadata is used for the overall system.
  - Volume metadata is used for individual volumes.

  If your DS8000 is a hybrid environment, the global metadata is stored on flash drives. As a best practice, at least 10% of the volume capacity should be available as flash storage to make sure that all volume metadata is also on the flash drives.

- Overall DS8000 system performance sizing

  The DS8000 system must be able to handle peak write workload of the production volumes and the Safeguarded Copy backup workload. The peak workload is basically two (for an MM solution) or three times (for a GM solution) the write workload of the Safeguarded Copy source volumes. Make sure that the utilization of DS8000 resources such as arrays and device interfaces is at or below 30% based on the production workload alone.

  If you plan to do regular data validation in parallel with creating Safeguarded Copy, you must use three (for an MM solution) to four times (for a GM solution) the write workload of the source volumes for the performance modeling.

Sizing is a crucial part of a Safeguarded Copy implementation, so for the Safeguarded Copy sizing, we recommend involving IBM to support you. Depending on your region, contact either the IBM Advanced Technical Group or the IBM EMEA Storage Competence Center.

2.4.4 IBM Z Cyber Vault storage sizing

The IBM Z Cyber Vault implementation uses the amount of storage that is required for the Safeguarded Copy Backup Capacity and the recovery volume (RC1) to implement Safeguarded Copy, as described in 2.2.3, “IBM Z Cyber Vault storage environment” on page 28. Some extra storage capacity is required for implementing the IBM Z Cyber Vault environment.

For example, extra storage capacity is needed for forensic data analysis and for surgical recovery. We call these extra volumes *staging volumes*. During a surgical recovery action, you copy the data that you identified for recovery from the recovery volume RC1 to the staging volumes. In a surgical recovery or forensic analysis, the restoration to production happens mainly from these staging volumes by either bringing the volumes online in your production environment (if the staging volumes are in an MM secondary (virtual isolated)) or by using the DS8000 GC function to copy the data to another set of staging volumes in your production environment. In addition, you require some persistent volumes in your IBM Z Cyber Vault environment to store reports and other historical data sets.

Some clients might also use a second set of recovery volumes (RC2) to compare the current data of the Safeguarded Copy source volumes (RS2) with an older Safeguarded Copy backup (or GDPS Capture).

A better example of an IBM Z Cyber Vault storage architecture and sizing is shown in Figure 2-14 (for more information, see 3.2, "Description of our IBM Z Cyber Vault environment" on page 71).

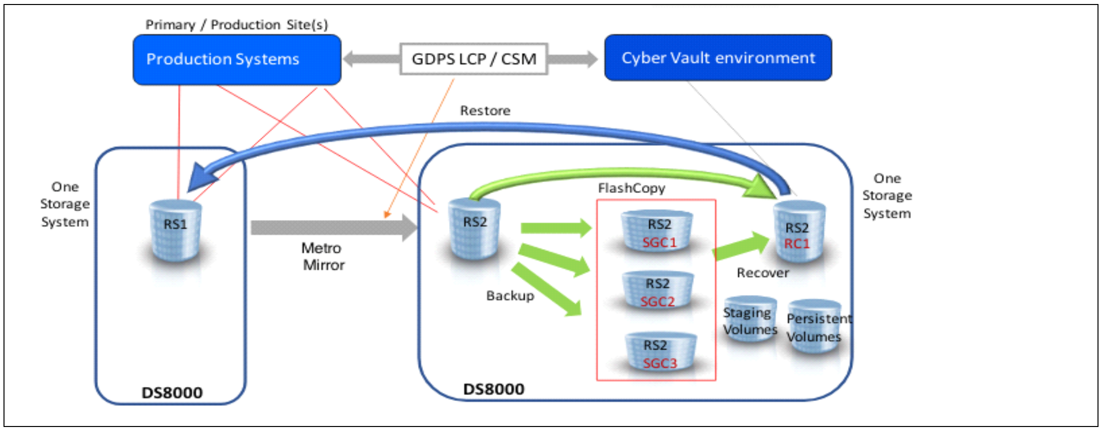

Figure 2-14 GDPS Metro with virtual airgap LCP

We used GDPS Logical Corruption Protection in an MM setup. The MM primary DS8000 system is RS1, and the secondary DS8000 system is RS2. The Safeguarded Copy backups are created from the RS2 set of volumes.

In this environment, we used two sets of recovery volumes that are called RC1 and RC2. The set RC1 is a FlashCopy copy from the source volume RS2, and RC2 is the Safeguarded Copy backups (recovery volumes) that can be used for data or structure validation.

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

Generally, usage of the RC1, RC2, and staging volumes is for a short duration, so they are allocated as ESE (thin-provisioned). The required capacity depends on the usage during a forensic analysis, data validation, or surgical recovery. Usually, the required capacity is far less than the source volume RS1 capacity, and might 20 - 30% per set.

Note: It is a best practice to add ~ 20% of the physical source volume capacity for the recovery volumes so that you can use the recovery volumes for data validation during normal operations. The sizing for the Safeguarded Capacity (physical and virtual) and recovery volumes can be done by using one of the methods that are described in 2.4.1, "Safeguarded Copy prerequisites" on page 41.

You must add this storage capacity to the Safeguarded Copy Backup Capacity and recovery volume sizing, as described in 2.2.3, "IBM Z Cyber Vault storage environment" on page 28.

Note: You cannot recover two Safeguarded Copy backups in parallel from the same Safeguarded Copy source volume to RC1 or RC2.

2.5 Key operational considerations for Safeguarded Copy

During the implementation of an IBM Z Cyber Vault environment, you must consider how often you create Safeguarded Copy backups and how long you keep these backups. These considerations might depend on regulatory or business requirements.

In addition, you must consider how often you must validate data in your IBM Z Cyber Vault environment, and how long that validation takes.

A higher backup (capture) frequency and a high regular frequency of data validation means that you are losing less data in the case of logical corruption, and faster detection of logical corruption during data validation is possible.

However, the Safeguarded Copy backup frequency, the backup retention period of your backups, and the data change rate are the key factors that influence how much capacity you need to store the backups. In addition, you need capacity for the recovery volumes on which you are doing the data validation.

2.5.1 Frequency of Safeguarded Copy backups

When you have a high frequency of Safeguarded Copy backups and captures, you reduce the recovery point objective (RPO). A low RPO is wanted, but it might require more capacity to store the changed tracks and data in your DS8000 system. It might be that some business requirements expect a specific frequency, so you must consider the backup frequency based on your requirements for your environment.

The Safeguarded Copy backups require consistent data. When backups are taken in an MM environment, you must “freeze” all write I/O to the volumes being backed up. A higher backup frequency would result in a freeze frequently impacting production. However, if the Safeguarded Copy backups are taken on a GM DR DS8000 system, or on isolated third or fourth site, such a freeze does not impact production, which enables more frequent backups. Our experience shows that a common frequency is a backup frequency of 4 - 6 hours. The DS8000 supports a frequency of every 10 minutes.

Figure 2-15 shows the backup frequencies that a client implemented with an indication how often they are taken.

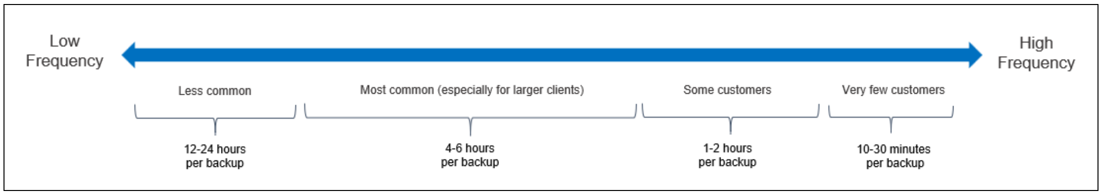

Figure 2-15 Backup frequencies

The backup frequency depends on the size of the environment, the data change rate, the required retention period, and the data validation frequency. However, compared to the change rate and the retention period, the backup frequency does not increase the storage requirements in the same way, which leads to a good value (increased protection) compared to the overall extra cost for a higher backup frequency. Even if you decide to not validate every copy of data (because the depth of validation might be more important for a client than the frequency), it is a best practice to keep more backups to restore to in case of an event.

2.5.2 Retention period of Safeguarded Copy backups

Beside the backup frequency, you must decide how long that you want to keep the Safeguarded Copy backups or captures. The longer that your retention period is, the more capacity that is required to store the changes in the DS8000 system. You must consider different aspects before you define the retention period for your environment, for example:

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

Today, the most common retention period is 2 - 5 days. Some clients want longer retention periods, and for these clients it might make sense to copy regularly validated data from the DS8000 system to WORM tapes to increase the retention period, which also would reduce the amount of capacity that is required in the DS8000 system.

Figure 2-16 shows the requested retention period, including an indication of how often they are implemented.

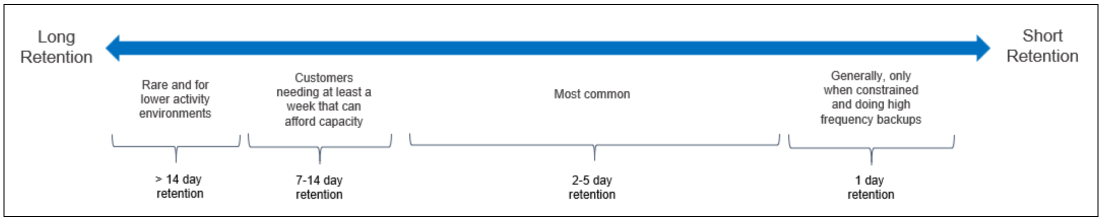
Long Retention
Short Retention
Rare and for lower activity environments
Customers needing at least a week that can afford capacity
Most common
Generally, only when constrained and doing high frequency backups
> 14 day retention
7-14 day retention
2-5 day retention
1 day retention

Figure 2-16 Safeguarded Copy backup retention period

You do not need to select one of the shown retention periods if you want to define the retention period based on your requirements. Make sure you do a Safeguarded Copy capacity sizing, as described in 2.4.3, “Sizing extra storage for Safeguarded Copy” on page 42 that reflects your requirements in terms of retention period and backup frequency.

2.6 IBM Z platform considerations

The IBM Z Cyber Vault environment is intended to work autonomously. It begins with the activation of the LPARs to send the results of data validation tests.

From the hardware point of view, a dedicated LPAR environment should be planned with network isolation.

The Base Control Program internal interface (BCPii) is used to automate the activation and the load of IBM Z Cyber Vault LPARs. You may also use IBM GDPS. The solution provides many features, such as the automating the management of the mainframe environment either through a GUI or by using scripts.

The system should be ready without any manual action. You can complete an IPL by defining an auto-reply policy for write-to-operator-with-replys (WTORs), setting up the COMMNDxx, and using an automation tool such as IBM Z System Automation.

Data structure validation may use REXX, z/OS utilities, and middleware-specific tools. The result of the validation is sent to the specific personnel by using SMTP.

If you plan on using the GDPS Logical Corruption Protection feature, it must be enabled with GDPS. With that in place, it is possible to create Safeguarded Copy backups, which are the foundation for IBM Z Cyber Vault. Extra procedures, tools, and utilities must be in place to support IBM Z Cyber Vault data validation. Some examples of IBM middleware tools and utilities that can be used in IBM Z Cyber Vault and production environments are described in 2.8, “Application considerations” on page 54.

2.6.1 IBM Z System Recovery Boost and IBM Z Cyber Vault

IBM Z System Recovery Boost was introduced with the IBM z15. This tool can deliver higher processor capacity for an IPL, for a limited time after an IPL, and during the shutdown of the system (LPARs). The increased capacity can be provided by one or more of the following items:

- In an LPAR on a subcapacity machine by using the full speed for the general-purpose processors of the boosted logical partition (Speed Boost)
- Dispatching some work to z Integrated Information Processors (zIIPs), even if the work is not eligible for zIIP or zIIP Boost

The *Startup Boost* function is activated by default during IPL and lasts for 1 hour. There is no specific configuration that is needed for running the IBM Z Cyber Vault environment, and the best System Recovery Boost option depends on the current system environment.

For more information about System Recovery Boost, see IBM Z System Recovery Boost content solution and *Introducing IBM Z System Recovery Boost*, REDP-5563*.*

2.6.2 Sizing the environment

Sizing the IBM Z Cyber Vault environment mostly depends on the planned work for it. The IBM Z Cyber Vault environment will not drive any production workload. However, resources are required to run all the data validation, forensic analysis, and recovery and backup procedures.

One LPAR is required to run the IBM Z Cyber Vault. An isolated coupling facility (CF) also can be defined to simplify the setup of the IBM Z Cyber Vault environment if the production system is using a CF. The number of CPs and zIIPs depend on the time that it takes and the procedures that are required to validate each Safeguarded Copy in the IBM Z Cyber Vault environment.

As a best practice, use roughly 10% of the production MIPS in IBM Z Cyber Vault, but no more than two CPs. This best practice depends on many different factors, for example:

- The number of sysplexes to be checked. Because IBM Z Cyber Vault checks data per sysplex, IPLs are needed on a per sysplex base. if a client runs three separated sysplex environments, three IPLs are needed to check the separated environments. These IPLs can be done one after the other, although a client might decide to check the sysplex environments in parallel. If so, the required resources increase.
- Depth of validation versus frequency of validation. There is no general rule that can be applied here. For example, some clients might have a data validation procedure that is available for their most critical databases and that this program must run as often as possible for the environment, which means frequency is more important to them. Other clients might say that they want to have a deep analysis of their environment, which might take some hours to complete and cannot check every Safeguarded Copy.

- The amount of resources that is needed depends on the number of data sharing members in the sysplex. The more systems (and application instances) that participate in a sysplex, the more Db2 data sharing group members (for example) must be restarted during the validation process to reach transactional consistency.

In addition, the following hardware components should be considered when sizing the environment:

- Extra IBM FICON® channels might be needed to access the new storage controller (physical isolation) or the recovery volumes in a virtually isolated environment. The minimum number is two FICON channels, but it is a best practice to use four FICON channels per storage controller from the IBM Z Cyber Vault environment if it is not used to run production in a DR case.
- Extra FICON channels also might be necessary after separated tape pools for the IBM Z Cyber Vault environment are used. Because this item is highly dependent on the client's current configuration and many options exist, there is no best practice.
- Extra OSA-Express features are required if a physically isolated network for IBM Z Cyber Vault is implemented. However, networks can be virtualized at the OSA-Express port level. As a best practice, use isolated features to avoid any compromises. For fault tolerance, at least two OSA-Express features for the IBM Z Cyber Vault environment are recommended.
- CF LPARs in the IBM Z Cyber Vault are required after a sysplex runs in production. Because there is no production workload, the performance of the CF is not as important as it is in production. Shared integrated catalog facility (ICF) LPARs are acceptable in the IBM Z Cyber Vault environment. Because fault tolerance also is not part of the IBM Z Cyber Vault ICF setup, you can use internal CFs.
- Extra memory on the machine that the IBM Z Cyber Vault LPAR and ICF is running on is required. There is no best practice because the size of the z/OS image and the ICF depends on the amount of memory that is used in production. For example, for ICFs, all the structures that are placed in the CF in the production environment must be placed in the IBM Z Cyber Vault environment, and the memory for doing so must be available to the IBM Z Cyber Vault ICF.

2.7 IBM z/OS considerations

z/OS and the subsystems that it supports rely on data structures that contain programs, configurations, and information, and they provide the system services that activate and maintain the z/OS operating environment.

These data structures can be the object of data corruption, so as key components of the main system services, they are where you start the Type 2 validation process. (Type 1 validation (IPL of the IBM Z Cyber Vault recovery system) is already complete, as described in 4.2, “Type 1 validation” on page 111.)

The Type 2 validation process can consist of, but is not limited to, these data structures:

- Catalogs
- Basic catalog structure - Virtual Storage Access Method Volume Data Set (BCS-VVDS)
- VSAM key sequenced data set (KSDS) files

Sample scripts and Job Control Language (JCL) code to run the validations by running utilities are provided with z/OS, such as Integrated Data Cluster Access Method Services (IDCAMS). You can find them in 4.3, “Type 2 validation” on page 114.

Also, IBM Tivoli Advanced Catalog Management for z/OS provides vital data protection that enhances system and data availability and helps you recover quickly from unexpected events. It has features that enable you to perform operations that are either difficult or impossible with IDCAMS alone:

- Back up, recover, and forward recover ICF catalogs and VVDS quickly, safely, and comprehensively.
- Clean hundreds of thousands of catalog entries in a few minutes and save hours at the recovery site by using the versatile **CATSCRUB** command.
- Audit your basic catalog structure (BCS), VVDS, and volume table of contents (VTOC) structures with fix commands that are automatically generated for correcting all identified error conditions.
- Back up and recover VSAM data sets and repair corrupted VSAM data sets.
- Create data set reports from catalog and VTOC information with ease.
- Simulate major commands, such as forward recovery and **MERGECAT**, to document procedures and verify results before running.
- Use the interactive system productivity facility (ISPF) interface to help create commands and batch jobs.
- Merge, split, or clone catalogs at many times the speed of IDCAMS REPRO MERGECAT and relax knowing that you have full restart or backout capabilities.

IBM Tivoli Advanced Catalog Management for z/OS also provides features for the correlation of entries in the tape management catalog with ICF catalogs that can help identify and resolve:

- Tape data that is not cataloged in tape management catalogs
- Tape data that is cataloged on tapes in scratch status
- Tape data that is cataloged in catalogs that are not connected to the master catalog
- Data in ICF catalogs that are not found in the tape management catalog
- Missing catalog records for tape data by building new entries

After you are done with these data structures, focus on DFSMShsm because it is the program that automatically performs space management and availability management in a storage device hierarchy.

DFSMShsm tracks all migration and backup activity by recording information in the control data sets, the journal data set, and the DFSMShsm log. These data structures also need to be validated. To do so, use two software tools: IBM Tivoli Advanced Audit for DFSMShsm and IBM Tivoli Advanced Reporting and Management for DFSMShsm, which work together to help facilitate a healthy DFSMShsm environment.

- Tivoli Advanced Audit for DFSMShsm delivers a set of auditing commands that help support the health of the control data sets. With this tool, storage administrators can maintain healthy metadata environments and prevent temporary or permanent loss of data access. It provides fast and accurate audits and diagnostics, and automatic corrective actions to resolve error conditions:
  - Nondisruptive audits for the metadata structures in your environment: catalogs, DASD VTOC records, DFSMShsm, and tape management systems.
  - Online selection and submission of audits, diagnostic aids, and corrective actions.
  - Online viewing of audit results.
  - Complete control over which metadata structures are audited.

  - Complete control over which diagnostic aid or corrective action is implemented for your unique environment.
  - Flexibility to edit JCL and corrective action commands to apply fix control statements to only the records that you want to fix.
  - Automatic application of fixes to reported errors.
  - Capability to turn on or off individual error correction actions.
  - Automatic halting of audits if a high threshold of errors is reported.
  - Ability to apply automatic fixes to specific high-level qualifiers (HLQs).
  - Support for both DFSMShsm standard TTOC and extended TTOC.
- Tivoli Advanced Reporting & Management for DFSMShsm is a tool that helps gather critical information from the DFSMS hierarchical storage manager environment. The tool provides detailed information about the records from the control data sets, health indicators for your DFSMShsm environment, information about your DFSMShsm host environment, and information about your DFSMShsm activity records. It also provides functions that help you manage your DFSMShsm environment. You can issue commands (such as **`HMIGRATE`**, **`HRECOVER`**, **`HBACKDS`**, and **`HRECALL`**), test user access to commands, reorganize your control data sets, and run reports that provide detailed information about the backup, migration, and offline control data sets that are used by DFSMShsm to store status information about the data sets under its control.
  - Report on all DFSMShsm activity.
  - Zoom in on problem areas.
  - Provide daily and historical information.
  - Provide Health Reports.
  - Select Error Corrections.
  - Process Corrective Commands.

The two DFSMShsm tools help you ensure access to DFSMShsm managed data:

- Data can be recalled from migration.
- Data can be restored from a DFSMShsm backup.
- Migration and backup tapes can be recycled.
- Corrupted DFSMShsm tape records can be rebuilt.
- DFSMShsm and the tape management system are synchronized.
- Control data sets are healthy and free of errors.
- Data can be recovered at the DR site.

2.7.1 Non-database managed files

Database managers track database activity (logs) and provide tools to recover to a consistency point. IBM Z Batch Resiliency provides log and recover capabilities for non-database managed data, such as libraries, flat files, and VSAM data sets.

This tool is fundamental in understanding how useful Safeguarded Copy will be when you must recover non-database managed data because it tracks all data set activity in your system and knows which files where open at the time of the Safeguarded Copy. In fact, this tool is the basis for the IBM Z Cyber Vault Health Check report that IBM Z Batch Resiliency provides and that must be run for every Safeguarded Copy. This report also helps to fine-tune the scheduling of the Safeguarded Copy copies.

IBM Z Batch Resiliency delivers high-value resiliency management of non-database managed data and applications, leveraging detailed analytic reporting to provide insights to reduce manual approaches that are required to manage data outside of database control. IBM Z Batch Resiliency reduces dependency on domain expertise and time-consuming, error-prone analysis that is needed to determine the impact of data corruption incidents. Like subsystem database tools for non-database managed applications and data, it improves resiliency and reduces enterprise business risk by offering immediate insight into application data inter-dependencies and vulnerabilities.

Key features of IBM Z Batch Resiliency include the following items:

- Capability to recover non-database-managed data and applications at a data set, job, or step level, which enables recovery from data corruption, operational errors, or data impacting events
- Journal-like capability to identify affected workloads and data sets that might be impacted by an event or a data set restoration
- Visibility into running jobs and open data sets that might be at risk
- Point-in-time recovery capability for non-database managed systems, such as batch workloads
- Near real-time comprehensive inventory of data sets, including detailed attributes and status, plus a complete repository of backups that are created by using any methodology (full-volume and logical) to help enable faster, panel-driven recovery

IBM Z Batch Resiliency also provides identification of all data sets in Safeguarded Copy, all data sets that were open at the time that a Safeguarded Copy copy was taken, and the capability for surgical recovery of any data sets from a Safeguarded Copy copy:

- Understanding of application data interdependencies and recovery points
- Reporting of resiliency readiness and audit gaps

2.7.2 IBM z/OS security

There are several external security managers (ESMs) that can be used to secure the z/OS environment, including IBM RACF (z/OS Security Server), and CA ACF2 or CA Top Secret. These ESMs must be checked for integrity, changes, and corruption. IBM Security zSecure Audit provides all the required capabilities and is available in three versions, one for each of the ESMs.

IBM Security zSecure Audit for RACF provides RACF and z/OS monitoring, SMF reporting, z/OS integrity checking, change tracking, and library change detection. This product can be used in the production environment to extend and enrich the existing IBM Z security by enforcing and enhancing security policies in a repeatable, sustainable, and automated fashion.

The automated reports can also help in the IBM Z Cyber Vault environment to quickly locate vulnerabilities and issues through audit trails and customized reports. Forensic analysis will benefit from the usage of data analytics to detect system changes.

IBM Security zSecure Audit for RACF can obtain and save point-in-time snapshots of the system settings and extra security-related data. This information is stored in a data set containing resource information that is gathered by the IBM Security zSecure Collect job, which is called CKFREEZE, and the supported data includes live settings based on information that is directly read from the control blocks of the current live system.

The full-size CKFREEZE data set contains information from all system and user catalogs, backup, migration and tape catalogs, all VVDSs, and all VTOCs. It also contains the directory information from the APF, linklist, lpalib, parmlib, and proclib data sets. This type of CKFREEZE data set also has information about files in UNIX HFS or zFS data sets. Program and transaction information for CICS and IMS systems is collected, and Db2 subsystem information like tables, packages, and other information.

Besides the auditing capabilities of CKFREEZE, the *library analysis* function is what should be used when performing data structure validations (Type 2). With library analysis, CKFREEZE has checksum information for the specified data sets. Because calculating checksum information is a time-consuming process, you create this type of CKFREEZE data set only when you want to do library analysis. So, you can create a CKFREEZE data set with checksums in production, and then each time you are validating a Safeguarded Copy backup in the IBM Z Cyber Vault environment, you can run CKFREEZE again to obtain new checksums, which can be compared against those from the production environment.

2.8 Application considerations

The IBM Z environment is driven and controlled by the z/OS operating system. However, applications are using different subsystems that were made for specific purposes. As such, they require distinct validation and recovery procedures, which can be based on utilities that are provided already with the corresponding subsystems, or leverage extra software tools and utilities that greatly improve on these processes.

Application subsystems enable transaction and data management, providing data and application integrity. Their software controls the creation, organization, and modification of data, and its access. 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. So, it is important to run Type 2 data validations against all these subsystems to make sure that their inner workings are not compromised.

In this section, we provide an overview of the tools and utilities that must be considered for data structure validation and recovery for the Db2, IMS, and CICS subsystems.

2.8.1 Db2 subsystem

Backup and recovery are one of the most complicated areas of database management. Having the correct resources to do a recovery is critical. Without them, you risk the loss of key data.

Database backup and recovery tasks vary from recovering from a dropped object to rebuilding after a major disaster. Recoveries that are done manually can be error-prone, time-consuming, and resource-intensive.

The backup and recovery solutions that are provided by IBM Db2 tools can help you reduce the business risks that are associated with data loss. In an IBM Z Cyber Vault configuration, these tools are the main ones that will help implement capabilities across the existing production environment and the IBM Z Cyber Vault recovery system environment:

- Db2 Administration Tool

  With its in-depth catalog navigation features, the Db2 Administration Tool minimizes the time that is required to review the Db2 catalog when performing forensic analysis. It interprets catalog information, and displays authorization information, SQL statements, and any other object in the catalog. It provides a convenient audit trail that can be used to determine the status of objects that were changed and allows you to recover changes and restore database objects to their previous state.

  You can use the Db2 Administration Tool to identify and fix problems with your databases. With its ability to navigate the catalog and use Db2 commands on objects, this tool can help you discover, analyze, and fix database problems.

  TheDb2 Administration Tool can be used with one or multiple copies of the Db2 system catalog.

- Db2 Log Analysis Tool

  Db2 Log Analysis Tool helps you to ensure HA and complete control over data integrity. It allows you to monitor data changes by automatically building reports of changes that are made to database tables.

  Some of the reporting capabilities that are provided include:

  - View data changes by dates, users, tables, and other criteria.
  - Create summary and detail reports that show:
    - Original state of the data.
    - Current state.
    - When data values changed and who made the changes.
    - Change activity and see a history of changes by time, by user, and by application.
  - Use the robust reporting capabilities, including reports on rollbacks, the ability to identify changes that are related to referential integrity, and expanded column data filtering.
  - View current-and post-row image differences to let you easily identify changes that occurred after the DML report was run.
  - Report on all log records that are associated with units of recovery when the commit (or abort) point for the UR is found.

  You can use these reports to institute tighter controls over the data to ensure that it can no longer be compromised.

  The auditing features also help with the analysis in case of data corruption by monitoring and auditing table activity:

  - Update, insert, delete, and much more.
  - Determine who changed the data.
  - Determine the sequence of the changes.

  If you must perform surgical recovery, the Db2 Log Analysis Tool can help to:

  - Generate SQL to undo or redo changes that are recorded in the log.
  - Support dropped object recoveries.

  - Report on and recover data for dropped objects by using both old and new Db2 identifiers.
  - Write data to an SQL file or to a LOAD-format file.
  - After DDL is re-created, restore the data in the regenerated table back to its state before the table was dropped.

  This tool automatically selecting the best way to run reports, determine table space names, and even dynamically allocating data sets that were previously hardcoded in the JCL.

- Db2 Recovery Expert

  Db2 Recovery Expert is a self-managing backup and recovery solution that protects mission-critical data. Db2 Recovery Expert can help you avoid accidental data loss or corruption by providing the fastest, least costly method of recovery when time is critical.

  The Db2 for z/OS environment can experience different kinds of failures, such as application errors, Db2 subsystem failures, Internal Resource Lock Manager (IRLM) failures, disk failures, z/OS failures, power failures, and site failures. When these failures occur, appropriate recovery procedures must be run, which is where Db2 Recovery Expert can help address your day-to-day data management problems.

2.8.2 IMS subsystem

IBM IMS Tools is a respected, industry-standard family of software products that support the operation and maintenance of IMS databases that are used in many of the world’s leading businesses.

Many IMS database environments rely on IMS Tools features to provide enhanced performance; improved efficiency, backup, and recovery processes; security; smart analysis and reports; and a multitude of regularly required database maintenance operations.

The IMS tools portfolio is composed of many different capabilities to support the IBM Z Cyber Vault recovery system environment:

- IMS Data Structure Validation

  IMS validation starts with detecting and reporting physical or logical direct pointer problems and detecting changes in database characteristics, such as size and number of segments. IMS High Performance Pointer Checker does this task for full function databases, or IMS Fast Path Solution Pack for fast path databases. These tools analyze a corrupted database as part of the repair process, reducing the diagnostic and repair time that is spent by programmers or analysts.

  To verify the assets that are needed for recovery and make sure that they are available, the IMS Recovery Solution Pack is recommended for all IMS databases. Among the tools that are included in this solution pack are IMS Database Recovery Facility, which integrates with other IMS tools to allow you to create image copies, rebuild indexes, and validate recovered databases. IMS Recovery Expert has integrated and intelligent recovery and DR managers that analyze recovery assets and establish optimal recovery procedures to minimize recovery time and meet RPOs.

- IMS Forensic Analysis

  After data corruption is detected, to understand where it is coming from and to what extent it has spread throughout your database, it is necessary to collect data in the production environment. IMS creates log data automatically, including all database updates. But if you are using IMS connect, you need IMS Connect Extensions to collect data for applications connecting to IMS through TCP/IP.

You can streamline your problem analysis process with IMS Performance Analyzer. It can create a transaction index from the IMS logs to quickly locate your problem transactions. This index contains the execution and performance metrics for all your transactions. You can use this index as the source input for subsequent IMS Performance Analyzer forms-based reporting or as input to IMS Problem Investigator to locate a problem transaction in the IMS log, and in-depth analysis of IMS transactions. The Connect Event Trace can include IMS Connect and IMS log records in a single report.

- IMS Recovery

  When surgical recovery is required, you must repair specific segments of the database so that few data points are impacted. IMS Database Repair Facility is a powerful tool to repair VSAM and OSAM organized IMS databases that contain pointer or data errors.

  This tool can be run both interactively and in batch mode for VSAM and OSAM data sets. In batch mode for VSAM data sets, it supports both IMS data sets and non-IMS data sets. The interactive capabilities of IMS Database Repair Facility help you make the repair in a short time, which means that the affected database must be taken offline only for a brief period. In this way, IMS Database Repair Facility enhances the integrity and availability of IMS databases.

  Using IMS Database Repair Facility, you can change the bad pointers and data to the values that you think they should have.

  IMS Database Repair Facility also provides the Indirect List Key (ILK) Repair utility. Use the ILK Repair utility together with the HD Pointer Checker utility (of IMS HP Pointer Checker) to repair HALDB databases that corrupted HALDB partition reorganization numbers, duplicate ILKs, or potentially duplicate ILKs.

  To complete a database recovery, IMS Database Recovery Facility can perform the following recovery tasks:

  - Simultaneously recover multiple full-function databases, HALDB partitions, and DEDB areas in a single pass of the change accumulation and log data sets.
  - Perform timestamp recovery by using timestamps that are not restricted to allocation boundaries. The ability to recover to any point in time allows you to recover IMS databases to the same point in time as other databases in your environment, such as Db2.
  - Recover database data sets and areas without needing to process IMS log data with a change accumulation utility.

  After you recover up to a recovery point, you can use the IMS Queue Control Facility to replay valid transactions beyond this recovery point, skipping invalid transactions based on the forensic analysis. It allows you to query, browse, unload, and load IMS messages from or to the IMS message queues.

2.8.3 CICS subsystem

A CICS application program can access data that is stored in direct-access data sets and the following VSAM file types: entry sequenced data sets (ESDSs), KSDSs, and relative record data sets (RRDS). Data structure validation for these types of data sets can be done with the utilities and tools that are described in 4.1.3, “Constructing data structure validation jobs” on page 106 and 4.1.4, “IBM Z Cyber Vault validation framework” on page 106.

CICS relies on a series of VSAM files to run, which are files holding temporary storage, definition of transactions, programs, and so on. If these files are not intact, the CICS system does not start and errors are logged.

To cope with corrupted VSAM files and CICS definitions, CICS VSAM Recovery (CICS VR) provides surgical recovery capabilities.

CICS VR is used to recover lost or damaged VSAM data sets. It determines which CICS logs and VSAM backups are needed and constructs the recovery jobs. In a CICS TS or batch environment, CICS VR forward recovery performs the same set of tasks:

- Restores the VSAM sphere from a logical backup, if available.
- Forward recovers all updates that are made by CICS and batch applications since the backup was taken.

When recovering VSAM transactional data from damage that is caused by logical corruption, the ISPF panels are used to build a forward recovery job where you can specify forward recovery criteria to exclude any updates that were made but are not required to your VSAM data by CICS transactions.

The *multiple undo logs* function enables you to use any number of undo logs instead of a single undo log. You can select a particular undo log to use that is based on user ID, job-name prefix, and VSAM sphere-name HLQ. The function also enables you to segment use of function, for example, to separate test logs from production ones.

In the case of a batch job that updates VSAM data but encounters a failure, it might leave VSAM data in an inconsistent state. So in this case, you use CICS VR batch backout to remove updates that are made to VSAM data by a batch job step that failed.

2.8.4 Application dependency

One of the biggest challenges of a logical corruption event is to identify the dependencies resulting from it. Imagine that a database was corrupted by accident, by malicious activity, or another event. To recover from such a situation, the following questions must be answered:

- Which applications are using this database?
- After this database is recovered to a status an hour ago (meaning data is lost and must be recovered from logs), what processes must be reviewed and maybe rerun to correct the corrupted data?
- Is the data that is used in the affected database used only on IBM Z or is it used outside the platform too?
- If the database is recovered to an earlier point in time, which other databases or sequential files must be recovered to the same point in time to make the application available in a consistent way?

There is no simple answer to these questions. However, IBM Z Batch Resiliency can help because it maintains an *application table* to identify which data sets and databases belong to or are shared among applications. Also, personnel that are responsible for the applications should always be involved to ensure they are recovered to the required consistent point-in-time.

2.9 Validation, forensic analysis, and recovery considerations

In 1.3, “IBM Z Cyber Vault solution” on page 6, we describe the cyber resiliency capabilities of the solution that are enabled by using a combination of hardware, software, storage, and services. Those capabilities can be used 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 RPO and the recovery time objective (RTO). The implementation of the cyber resiliency capabilities and the data validation frequency can help meet RPO and RTO targets in the event of a cyberattack.

As a reference, the IBM Z Cyber Vault solution supports a repeatable data validation process, and if necessary enables forensic analysis and recovery. Data validation should be done at intervals that align with your recovery strategy.

If no corruption is detected, the environment can be saved to offline tape media for added security. If data corruption is found, you 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 2-17 shows the validation, forensic analysis, and recovery phases that span the IBM Z Cyber Vault and production environments.

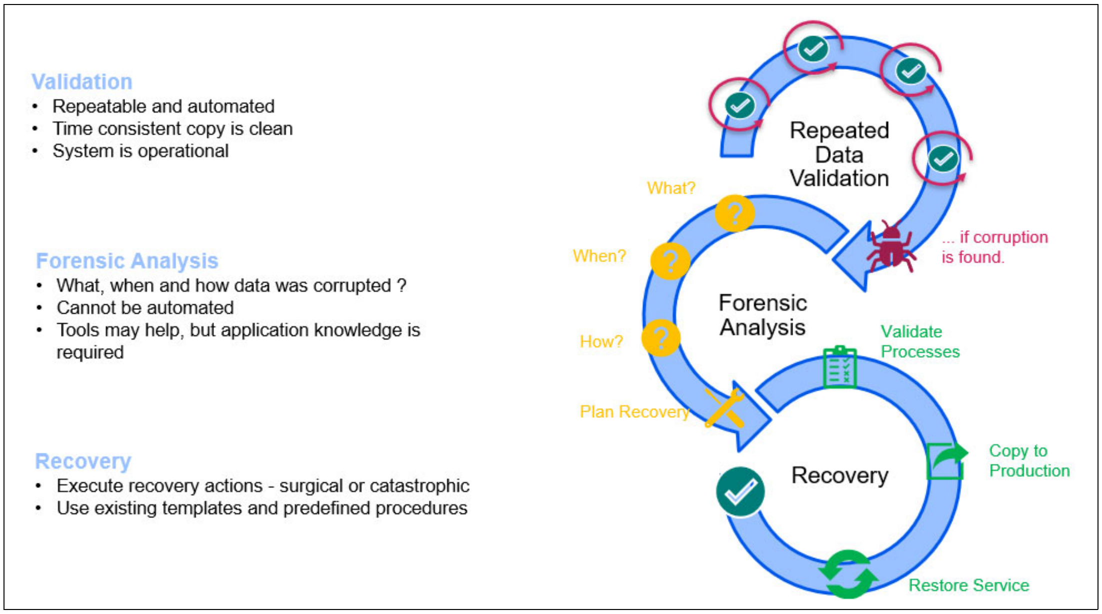

Validation

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

Forensic Analysis

- What, when and how data was corrupted ?
- Cannot be automated
- Tools may help, but application knowledge is required

Recovery

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

Figure 2-17 IBM Z Cyber Vault validation, forensic analysis, and recovery processes

2.9.1 Validation

When considering or implementing cyber resiliency 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 will allow your business to get back to normal in the shortest possible time while losing little to no data.

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

The next step is to take the created point-in-time image of the production environment and validate it to ensure that there is no data corruption. Data corruption can affect the system's software structure or the application data.

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 described as Type 1 and Type 2 validations in 1.4.1, “Validation use cases” on page 11.

Application data validation has to do with each user's own application data, and because the lifecycle of this data is managed by business applications, it is only the user that 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 validation is defined as Type 3 in 1.4.1, “Validation use cases” on page 11.

Figure 2-18 shows an outline of the validation process.

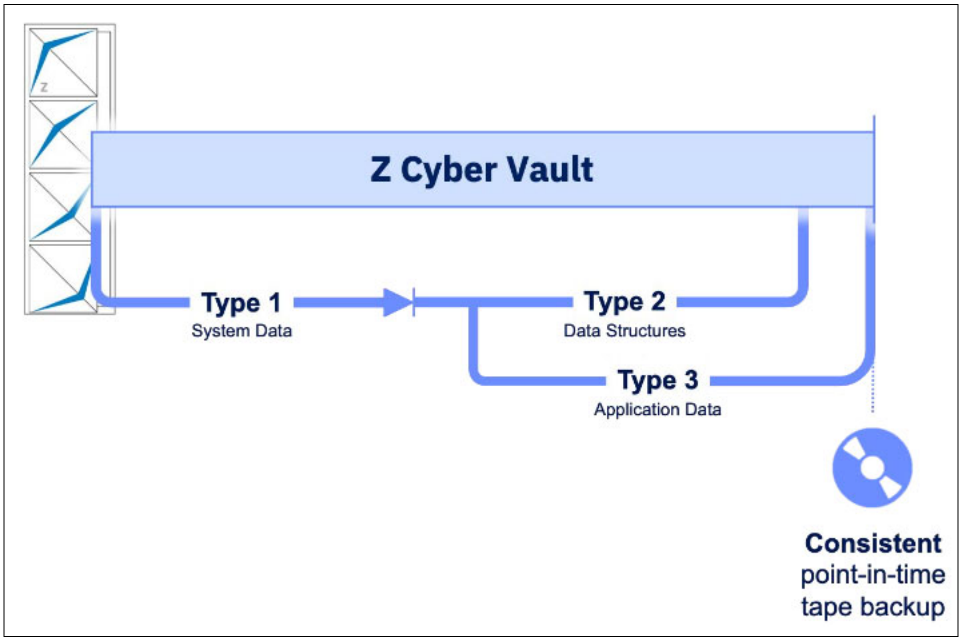

Figure 2-18 IBM Z Cyber Vault validation process

After successfully completing all the data validation procedures, you can choose to take a tape backup copy of this consistent, validated environment as a second layer of protection.

If data corruption was encountered, you must understand how it happened and to what extent through forensic analysis (see 2.9.2, “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. It includes data structure validation and application data validation. By running data validation procedures, you can detect logical corruptions in your environment.

As a best practice, run this function frequently, which means that the best case is to perform the data validation for every Safeguarded Copy backup. In this way, you can detect a logical corruption earlier and reduce the impact of the corruption.

To accomplish this task, appropriate tools and automation are required because the time that it takes to run a full validation of your complete environment ultimately determines how frequently you may 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 2-19).

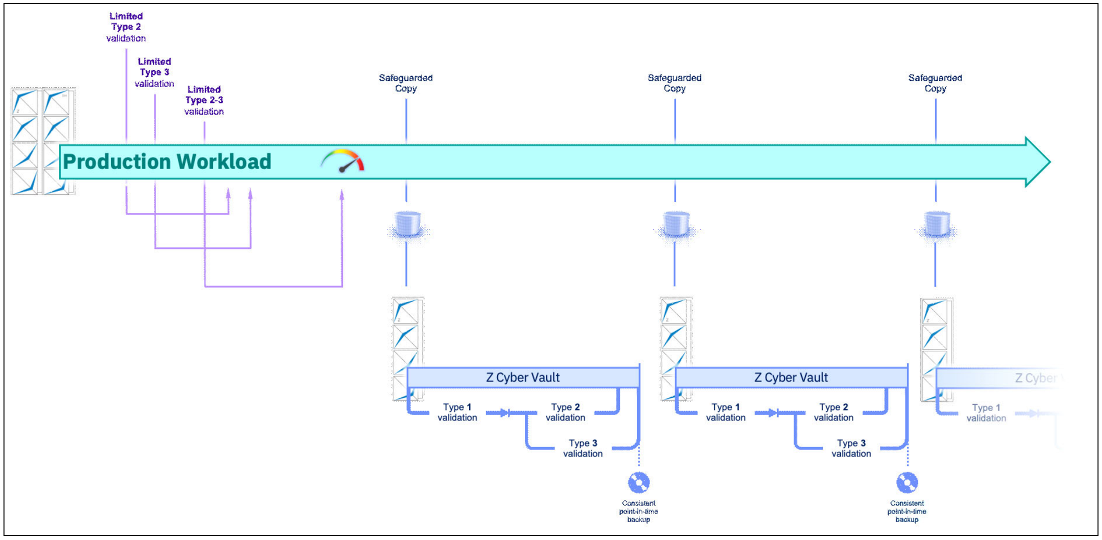

Figure 2-19 Data validation frequency

Some of the validation processes can be run in the production environment to detect data corruption as early as possible. However, these validation processes always are limited in scope and depth to avoid affecting system and applications performance and response times, which can increase the cost of running your business-critical applications.

Data validation tools

In our environment, we 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, we developed 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 allow you to validate your data effectively and efficiently and be better prepared in case of 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 were not designed to be used in a cyberattack scenario where a fast response and accurate actions are required. Therefore, IBM and several other vendors developed software tools that deliver features and functions to address these requirements in a more efficient and comprehensive way.

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, easier, 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 stress conditions. Therefore, it makes sense to think beyond the *normal* when selecting extra software.

IBM Z Cyber Vault Discover and Architecture Workshop helps you determine how existing tools are leveraged in the IBM Z Cyber Vault, and what are the gaps that must be addressed through the implementation of extra software. (For more information, see 1.6, “IBM Z Cyber Vault service offerings” on page 16.)

2.9.2 Forensic analysis

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® and operational understanding, and potentially specific application and database knowledge.

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, say at 10:15 AM in the morning.
- Your quick analysis shows that the cyberattack already has 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 using a validated Safeguarded Copy backup and perform a restart.

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 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. At this point, 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 will be running several tools to their full potential to investigate where the data corruption started and how. Many of these tools are described throughout this publication, but also can include other IBM or non IBM software that you are using in your production environment. Log analysis tools are some of the key utilities that you should consider for everything that creates a log in z/OS including (but not limited to) 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 procedure. All software tools that are required for this purpose should be made available in the IBM Z Cyber Vault environment.

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 Logical Corruption Protection support.

In a Db2 scenario (see Figure 2-20), 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.

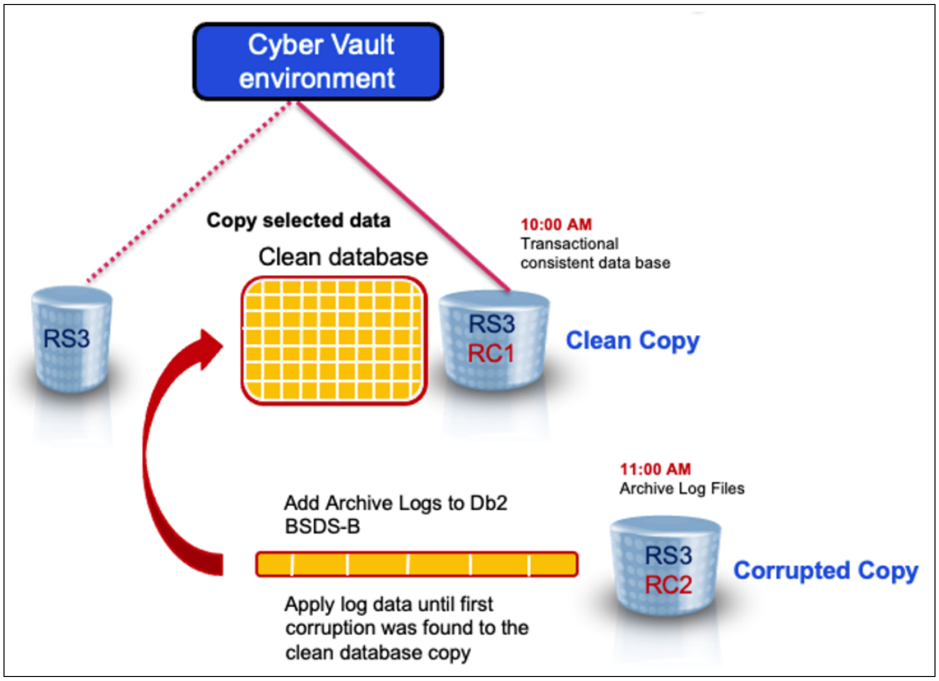

Figure 2-20 Forensic analysis in a Db2 scenario

You can either establish a FlashCopy relationship directly from the RS3 set of volumes to the recovery set of volumes (RC1), or you 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 will be ESE volumes (thin-provisioned). The data will be 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.

If the RS3 will be changing at the rate of data change in the production environment and the change rate in the recovery environment is low, little space for the RC1 set of volumes is needed. However, if the change rate in your environment is high, then you will copy more data to RC1.

In general, you will do sizing for a catastrophic recovery, which should suffice for forensic analysis and surgical recovery (see 2.4.4, “IBM Z Cyber Vault storage sizing” on page 45).

2.9.3 Recovery

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 different ways to restore data back to production.

With 3-site and 4-site configurations, you can have variations that are not covered in this section. For more information about system recovery offerings and options for the IBM Z platform, contact the IBM GDPS team at mailto:gdps@us.ibm.com.

Surgical recovery

If you must recover only portions of data, you can do a selective restore, which is also known as a *surgical recovery*.

This capability can be provided at a z/OS system level (see “System-level surgical recovery” on page 64) and for each of the transaction and data management subsystems in use. We selected Db2 to describe in “Db2 subsystem surgical recovery” on page 65 how surgical recovery would work for it. 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 a virtually isolated environment, system-level surgical recovery can be done by copying the data from the recovery volumes to a set of staging volumes (see Figure 2-21) and making the staging volumes available to the production system LPARs by using standard storage-level CS. Use standard z/OS tools or application methods to copy the data that you need from the (production) staging volumes to the production volumes.

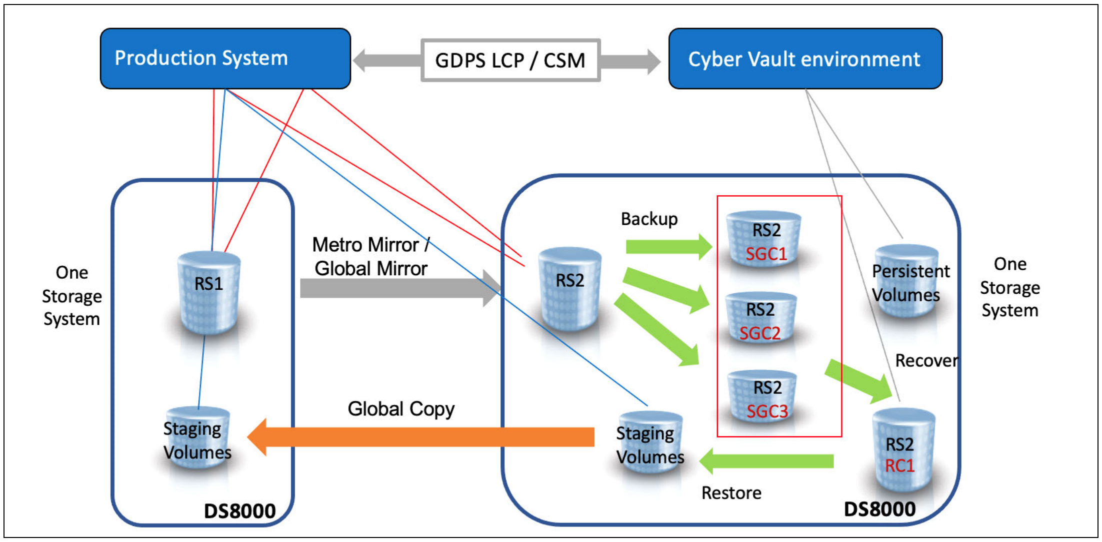

Figure 2-21 Using staging volumes for surgical recovery and partial restore to production

In a physically isolated environment, copy the data from the recovery volume (RC1) to a set of staging volumes in the IBM Z Cyber Vault environment, and copy the established GC relationships between a set of staging volumes in the IBM Z Cyber Vault environment and a set of staging volumes in the production sysplex. When 100% of the data is copied, you can remove the GC; bring the staging volumes online in your production environment; and copy the data to the production volumes by using standard operating system methods.

Db2 subsystem surgical recovery

We now describe typical surgical recovery scenarios for Db2 that can be used as an examples 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:

- Scenario 1: Your database system log files are available in production, as well as your image copies and any incremental copies. These files and copies also are 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.

So, let us see what can be done in each of these situations:

- Scenario 1: 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: 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, it should be easy to determine which is the last valid database image copy, which would have been 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 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: 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, 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, then here are the steps that you should follow:

a. Recover the most recent Safeguarded Copy backups that contain a copy of the clean database into the RC1 volume.

b. 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.

c. Copy the logs and BSDSs from the RC2 to the RC1 volumes.

d. Reassemble the Db2 zParms with DEFER ALL.

e. 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.

f. Determine the LRSN corresponding to the recovery point chosen above.

g. Use the DSNJ003 utility to update the ENDLRSN on all Db2 member BSDSs.

h. 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 for z/OS, there is no need to use two recovery sets of volumes. Only the latest Safeguarded Copy backups are required: After you have identified 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.

Catastrophic recovery

If you must do a catastrophic recovery because the data corruption is extensive, a full restore of a Safeguarded Copy backup (GDPS capture) is required.

Because the amount of data to move back with GC from RC1 to production volumes (RS1 or RS2) is large, this action can take time.

Note: With recent enhancements in IBM DS89xx storage and microcode Release 9.2, it is possible to copy data back to production incrementally. This capability drastically reduces the amount of time that is needed to copy data back to the production system. However, you must do a fresh Safeguarded Copy Recover and use the CSM V6.3.0 or later Restore option, which requires that the production volumes that are being restored are not online.

For that full restore, you must establish an infrastructure that allows you to restore the data with GC. This infrastructure will require DS8000 resources and adequate bandwidth in the SAN infrastructure between the recovery volumes (RC1) and the production volumes (RS1).

The preferred way for catastrophic recovery is to use GC to replicate the data from the recovery volumes (RC1) to the production volumes (RS1), as shown by the orange arrow (GC) in Figure 2-22 on page 67.

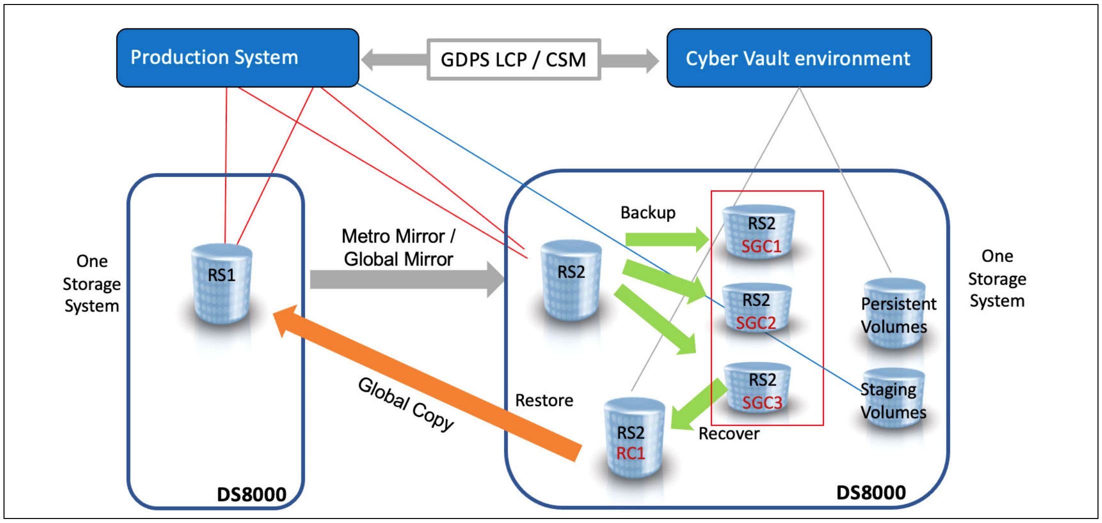

Figure 2-22 Restore to production

You can use different interfaces like the DS CLI, CSM, GDPS Logical Corruption Protection, or z/OS commands to establish the GC for the restore process.

> **Note:** For DS89xx with microcode Release 9.2 or later and using CSM V6.3.0 or later, it is possible to use GC to restore in increments to the production system.
>
> With earlier DS8000 storage systems, an incremental resynchronization from RC1 with GC is not possible. So, a GC restore process might take a while to complete.

Depending on your requirements and topology, you might decide to perform an IPL from a validated Safeguarded Copy backup by using your RC1 volumes to run production either in your IBM Z Cyber Vault environment, or directly from your production environment or DR LPARs. While you run production by using the RC1 set of volumes, you can start restoring the RC1 data back to a production set of volumes in parallel.

With this procedure, you can reduce the RTO dramatically, and run your production system immediately while restoring data to production is done in the background.

To see how this method is used in our example environment, see Chapter 3, “Deploying the IBM Z Cyber Vault environment” on page 69.

> **Note:** During IPL and running production from RC1, every I/O operation will go through the Safeguarded Copy metadata, impacting the response time. It might be possible to change the existing Recover command to Full Copy instead of NOCOPY. A full copy of the Safeguarded Copy data to RC1 would not cause such a slowdown. Consider this option when you decide to restart immediately on RC1 to reduce the RTO, or wait for the background copy to complete and get better performance.

Regardless of whether you run production in the IBM Z Cyber Vault environment or not, 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 Logical Corruption Protection 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 Chapter 4, “Establishing a validation framework” on page 103) 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 (MM) to GC.
4. Establish the GC relationship between your recovery volumes (RC1) to the production volumes (RS1) with the interface of your choice (DS CLI, GDPS, or CSM).
5. You should stop the Safeguarded Copy backup (GDPS capture) until the data is copied and verified in production. The new Safeguarded Copy backup might expire older backups. You do not want the older Safeguarded Copy backups to expire in case you need them.
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 now must 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 will go 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. Reestablish your Safeguarded Copy environment if you stopped it.

The process above describes the general steps to restore to production.

> **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. The current DS8000 microcode does not support FlashCopy from RC1, which is why GC is the best practice to do a restore to production.

3

Deploying the IBM Z Cyber Vault environment

This chapter describes the deployment process for the IBM Z Cyber Vault environment in a multi-site configuration. We set up our environment to emulate an IBM Geographically Dispersed Parallel Sysplex (GDPS) Metro solution (single-leg configuration with GDPS Logical Corruption Protection Manager). This solution is the most commonly one that is used in an IBM Z multi-site configuration. GDPS also provides an automation layer with GDPS Logical Corruption Protection Manager to enable the capture of multiple, secure, and point-in-time copies of production data, which can later be used for validation, repair, or replacement of compromised production data.

Other configurations are also supported, as described in 2.3.2, “IBM GDPS capabilities and supported models” on page 35.

The following topics are covered in this chapter:

- Preparing to deploy the IBM Z Cyber Vault environment
- Description of our IBM Z Cyber Vault environment
- Setting up IBM Z Cyber Vault storage environment
- Setting up Safeguarded Copy in IBM Z Cyber Vault storage
- Preparing the IBM Z Cyber Vault environment for validation

3.1 Preparing to deploy the IBM Z Cyber Vault environment

Before you deploy the IBM Z Cyber Vault environment, you should become familiar with the various capabilities and supported configurations that are described in Chapter 2, “Planning and designing the IBM Z Cyber Vault environment” on page 19.

The following items should also be considered:

- The management method for generating Safeguarded Copy backups of your data by using either GDPS Logical Corruption Protection or IBM Copy Services Manager (CSM).
- The topology of your IBM Z Cyber Vault environment and replication technique. For example, a 2-site, 3-site, or 4-site configuration.
- Whether you will use virtual or physical airgap storage for your Safeguarded Copy backups of data.
- The sizing of extra storage that is required for your Safeguarded Copy backups based on the location, frequency, retention period of your copies, and the data change rate.
- The number of IBM Z Cyber Vault logical partitions (LPARs) that are required for performing validation of your production environments:
  - You might have chosen to start all critical subsystems on one IBM Z Cyber Vault LPAR for validation (preferred).
  - You might need to have extra IBM Z Cyber Vault LPARs.

  If your catalogs are not synchronized between all of your production LPARs in your environment, then you need multiple IBM Z Cyber Vault LPARs.
- Whether your IBM Z Cyber Vault environment will use an existing IBM Z platform (if it is an already installed disaster recovery (DR) environment) or your IBM Z Cyber Vault environment will be located elsewhere.
- The sizing of your IBM Z Cyber Vault LPARs and the needed IBM Z Cyber Vault CFs if you are verifying a sysplex.

Note: Your IBM Z Cyber Vault LPARs should contain dedicated resources when sharing an IBM Z platform with an existing environment so that you can continuously perform automated initial program loads (IPLs) and validation of your IBM Z Cyber Vault environment and Safeguarded Copy backups in a consistent manner.

A phased approach is a best practice so that you get, as a priority, the benefit of logical corruption protection (LCP) that is provided by an airgap Safeguarded Copy solution.

The phased approach should consist of discrete milestones, such as:

- Establishment of an airgap Safeguarded Copy backup environment and Safeguarded Copy recovery environment
- Ability to perform Type 1 validation, that is, the basic IPL (manually) to confirm a working IBM Z Cyber Vault environment in a network that is isolated system from your Safeguarded Copy recovery volumes without affecting the production environment
- Ability to perform Type 1 validation (automating IPLs in IBM Z Cyber Vault environment from Safeguarded Copy recovery volumes) in the allotted time
- Ability to perform Type 2 validation (running of a set of critical subsystem data structure validation procedures) in the allotted time

- Ability to perform Type 3 validation (running a set of data validation procedures) in the allotted time
- Ability to report on validation results and confirm receipt of those results

You can phase in the more automated capabilities of the IPL process for the IBM Z Cyber Vault environment and perform critical data structure validation and application data content validation scripts in steps.

> **Note:** Depending on the benchmarked duration of your baseline run of data structure validation (such as catalogs) and the duration of your baseline run of data content validation (jobs that are built by your applications), the actual IBM Z Cyber Vault validation cycles that are performed might be less, equal, or higher to the frequency of your Safeguarded Copy operation.
>
> For this reason, you must first perform baseline execution of the validation jobs that you plan to run to determine the reasonable validation activities that you can perform in the time allotted.
>
> Consider running an IBM Z Cyber Vault validation cycle at least once a day.

3.2 Description of our IBM Z Cyber Vault environment

In our environment, we use the GDPS Logical Corruption Protection feature. Several multi-site topologies are supported by the GDPS Logical Corruption Protection feature with virtual or physical airgap configurations. The basic difference for determining where GDPS Logical Corruption Protection Manager is implemented depends on whether you have a virtual or physical airgap configuration.

In a virtual airgap configuration, the protection copies (in our case, Safeguarded Copy backups, also known as captures in GDPS) are within the same storage system as one of the existing production or DR copies in the GDPS environment. GDPS Logical Corruption Protection Manager is implemented in one of the existing controlling systems of the GDPS environment. Extra controlling systems are not required for GDPS Logical Corruption Protection Manager.

For more information about supported configurations, see 2.3.2, “IBM GDPS capabilities and supported models” on page 35.

We set up our environment as a GDPS Metro single-leg configuration with a virtual airgap GDPS Logical Corruption Protection Manager. Single-leg configurations consist of only two copies of the production data: a primary copy and a secondary copy. The primary and secondary copies of data are called replication sites (RSs). The copy in Site 1 is known as RS1, and the copy in Site 2 is known as RS2.

Figure 3-1 shows the logical flow within the environment of a GDPS Metro single-leg configuration.

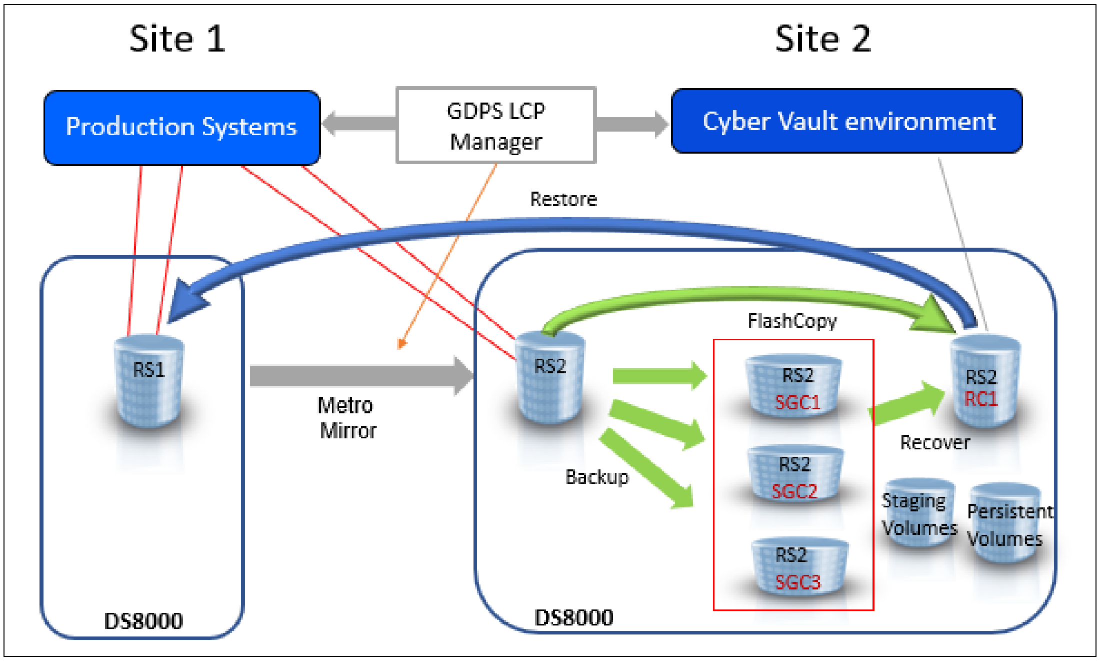

Figure 3-1 GDPS Metro solution with virtual airgap GDPS Logical Corruption Protection Manager

The details of our environment are as follows:

- All the LPARs are built on an IBM z15 (model 8561-729) system.
- The applications are running in a multi-site data sharing sysplex in two systems (MVS1 and MVS2) with two internal coupling facilities. In the same sysplex, we have the GDPS Metro controlling system (GVC2), which is also our virtual airgap GDPS Logical Corruption Protection Manager.
- Each application system LPAR has 10 LCP and two IBM Z Integrated Information Processors (zIIPs) with a real memory of 64 GB. The two application systems are running on z/OS V2.3, with Db2 V11 (DBW1) and Db2 V12 data sharing (DBX1/X2).
- The workload is simulated through an application program package that is developed by Sopra Banking that is called Sopra Banking Account Management. Account Management is a component of the Sopra Banking Platform, which is a complete and modular banking solution.
- In the environment, a Db2 batch workload is used. It is composed of several batch chain-called scenarios. The three most important scenarios of the workload for the IBM Z Cyber Vault test are Financial Operations (B3), Accounts Accruals (B2), and Accounts Valuation (B1).
- These scenarios are run in sequence and are composed of many jobs. One part of the processing is done by a single job, which reads and sorts the input sequential data set. The other part of the processing is done in parallel by 20 jobs, each job working on its own set of database partitions. A massive system and Db2 activity is triggered by these scenarios, allowing the environment to emulate a banking production environment.
- In the recovery region, there are two spare LPARs ready to be used to perform an IPL of the recovered systems, and two extra internal coupling facilities. We use 10 logical CPs plus two zIIPs and 64 GB of real memory.

The hardware configuration is illustrated in Figure 3-2.

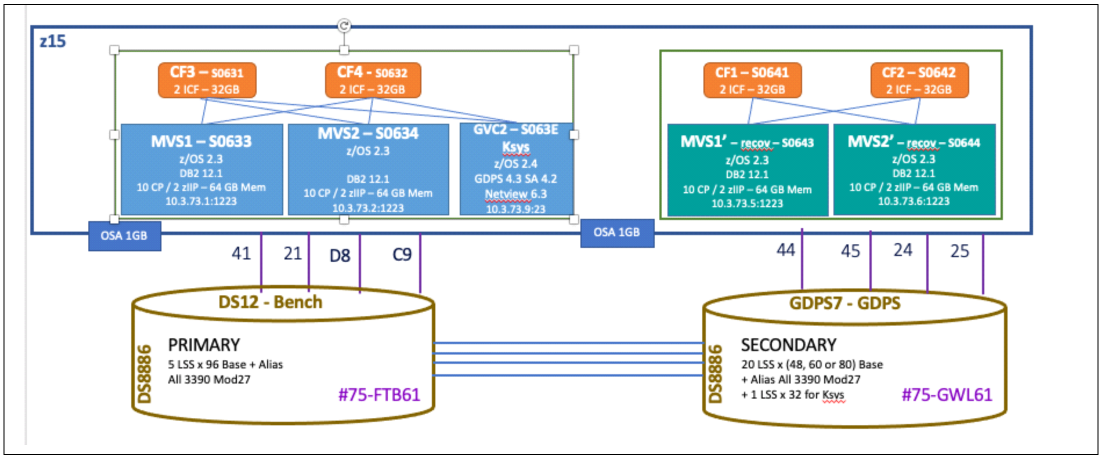

Figure 3-2 Hardware and sysplex configuration

The storage systems are DS8886 systems that are connected by four IBM FICON channels to the z15 LPARs, and with four Fibre Channel PPRC links for the GDPS Metro configuration.

The GDPS Metro primary DS8886 systems have five logical subsystems (LSSs) with 96 base addresses. All volumes are on 3390 Mod-27 systems.

The volume layout of our environment is shown in Figure 3-3.

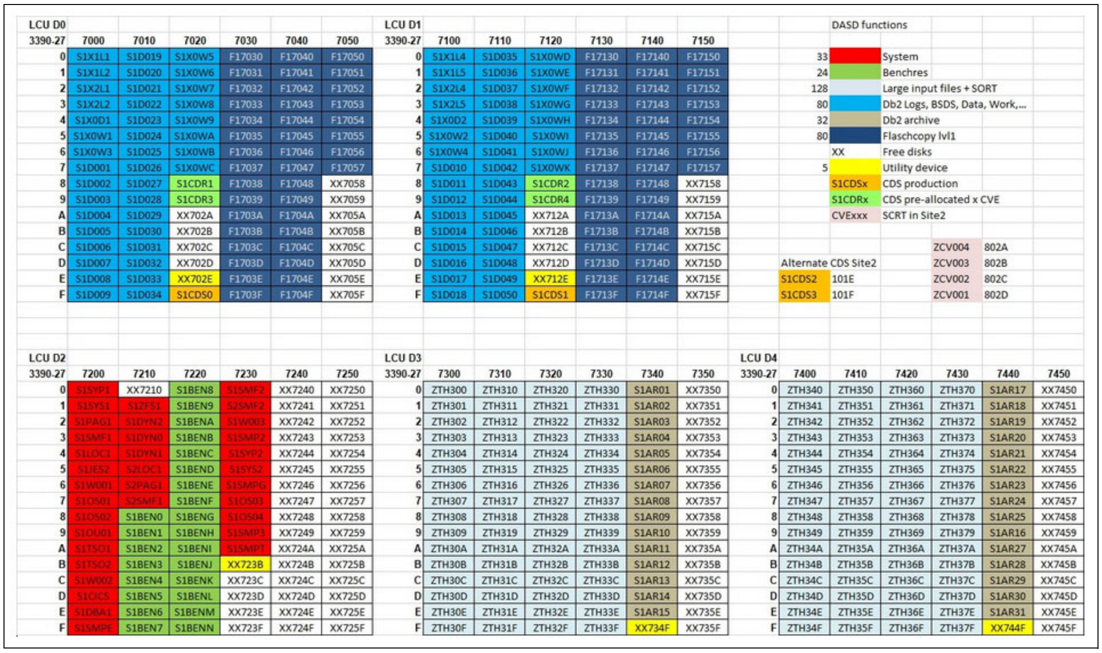

LCU D0						
3390-27	7000	7010	7020	7030	7040	7050
0	S1X1L1	S1D019	S1X0W5	F17030	F17040	F17050
1	S1X1L2	S1D020	S1X0W6	F17031	F17041	F17051
2	S1X2L1	S1D021	S1X0W7	F17032	F17042	F17052
3	S1X2L2	S1D022	S1X0W8	F17033	F17043	F17053
4	S1X0D1	S1D023	S1X0W9	F17034	F17044	F17054
5	S1X0W1	S1D024	S1X0WA	F17035	F17045	F17055
6	S1X0W3	S1D025	S1X0WB	F17036	F17046	F17056
7	S1D001	S1D026	S1X0WC	F17037	F17047	F17057
8	S1D002	S1D027	S1CDR1	F17038	F17048	XX7058
9	S1D003	S1D028	S1CDR3	F17039	F17049	XX7059
A	S1D004	S1D029	XX702A	F1703A	F1704A	XX705A
B	S1D005	S1D030	XX702B	F1703B	F1704B	XX705B
C	S1D006	S1D031	XX702C	F1703C	F1704C	XX705C
D	S1D007	S1D032	XX702D	F1703D	F1704D	XX705D
E	S1D008	S1D033	XX702E	F1703E	F1704E	XX705E
F	S1D009	S1D034	S1CDS0	F1703F	F1704F	XX705F

LCU D1						
3390-27	7100	7110	7120	7130	7140	7150
0	S1X1L4	S1D035	S1X0WD	F17130	F17140	F17150
1	S1X1L5	S1D036	S1X0WE	F17131	F17141	F17151
2	S1X2L4	S1D037	S1X0WF	F17132	F17142	F17152
3	S1X2L5	S1D038	S1X0WG	F17133	F17143	F17153
4	S1X0D2	S1D039	S1X0WH	F17134	F17144	F17154
5	S1X0W2	S1D040	S1X0WI	F17135	F17145	F17155
6	S1X0W4	S1D041	S1X0WJ	F17136	F17146	F17156
7	S1D010	S1D042	S1X0WK	F17137	F17147	F17157
8	S1D011	S1D043	S1CDR2	F17138	F17148	XX7158
9	S1D012	S1D044	S1CDR4	F17139	F17149	XX7159
A	S1D013	S1D045	XX712A	F1713A	F1714A	XX715A
B	S1D014	S1D046	XX712B	F1713B	F1714B	XX715B
C	S1D015	S1D047	XX712C	F1713C	F1714C	XX715C
D	S1D016	S1D048	XX712D	F1713D	F1714D	XX715D
E	S1D017	S1D049	XX712E	F1713E	F1714E	XX715E
F	S1D018	S1D050	S1CDS1	F1713F	F1714F	XX715F

DASD functions		
33	🟥	System
24	🟩	Benchres
128	⬜	Large input files + SORT
80	🟦	Db2 Logs, BSDS, Data, Work,…
32	🟫	Db2 archive
80	🟦	Flaschcopy lvl1
	XX	Free disks
5	🟨	Utility device
	S1CDSx	CDS production
	S1CDRx	CDS pre-allocated x CVE
	CVExxx	SCRT in Site2

		ZCV004	802A
Alternate CDS Site2		ZCV003	802B
S1CDS2	101E	ZCV002	802C
S1CDS3	101F	ZCV001	802D

LCU D2						
3390-27	7200	7210	7220	7230	7240	7250
0	S1SYP1	XX7210	S1BEN8	S1SMF2	XX7240	XX7250
1	S1SYS1	S1ZFS1	S1BEN9	S2SMF2	XX7241	XX7251
2	S1PAG1	S1DYN2	S1BENA	S1W003	XX7242	XX7252
3	S1SMF1	S1DYN0	S1BENB	S1SMP2	XX7243	XX7253
4	S1LOC1	S1DYN1	S1BENC	S1SYP2	XX7244	XX7254
5	S1JES2	S2LOC1	S1BEND	S1SYS2	XX7245	XX7255
6	S1W001	S2PAG1	S1BENE	S1SMPG	XX7246	XX7256
7	S1OS01	S2SMF1	S1BENF	S1OS03	XX7247	XX7257
8	S1OS02	S1BEN0	S1BENG	S1OS04	XX7248	XX7258
9	S1OU01	S1BEN1	S1BENH	S1SMP3	XX7249	XX7259
A	S1TSO1	S1BEN2	S1BENI	S1SMPT	XX724A	XX725A
B	S1TSO2	S1BEN3	S1BENJ	XX723B	XX724B	XX725B
C	S1W002	S1BEN4	S1BENK	XX723C	XX724C	XX725C
D	S1CICS	S1BEN5	S1BENL	XX723D	XX724D	XX725D
E	S1DBA1	S1BEN6	S1BENM	XX723E	XX724E	XX725E
F	S1SMPE	S1BEN7	S1BENN	XX723F	XX724F	XX725F

LCU D3						
3390-27	7300	7310	7320	7330	7340	7350
0	ZTH300	ZTH310	ZTH320	ZTH330	S1AR01	XX7350
1	ZTH301	ZTH311	ZTH321	ZTH331	S1AR02	XX7351
2	ZTH302	ZTH312	ZTH322	ZTH332	S1AR03	XX7352
3	ZTH303	ZTH313	ZTH323	ZTH333	S1AR04	XX7353
4	ZTH304	ZTH314	ZTH324	ZTH334	S1AR05	XX7354
5	ZTH305	ZTH315	ZTH325	ZTH335	S1AR06	XX7355
6	ZTH306	ZTH316	ZTH326	ZTH336	S1AR07	XX7356
7	ZTH307	ZTH317	ZTH327	ZTH337	S1AR08	XX7357
8	ZTH308	ZTH318	ZTH328	ZTH338	S1AR09	XX7358
9	ZTH309	ZTH319	ZTH329	ZTH339	S1AR10	XX7359
A	ZTH30A	ZTH31A	ZTH32A	ZTH33A	S1AR11	XX735A
B	ZTH30B	ZTH31B	ZTH32B	ZTH33B	S1AR12	XX735B
C	ZTH30C	ZTH31C	ZTH32C	ZTH33C	S1AR13	XX735C
D	ZTH30D	ZTH31D	ZTH32D	ZTH33D	S1AR14	XX735D
E	ZTH30E	ZTH31E	ZTH32E	ZTH33E	S1AR15	XX735E
F	ZTH30F	ZTH31F	ZTH32F	ZTH33F	XX734F	XX735F

LCU D4						
3390-27	7400	7410	7420	7430	7440	7450
0	ZTH340	ZTH350	ZTH360	ZTH370	S1AR17	XX7450
1	ZTH341	ZTH351	ZTH361	ZTH371	S1AR18	XX7451
2	ZTH342	ZTH352	ZTH362	ZTH372	S1AR19	XX7452
3	ZTH343	ZTH353	ZTH363	ZTH373	S1AR20	XX7453
4	ZTH344	ZTH354	ZTH364	ZTH374	S1AR21	XX7454
5	ZTH345	ZTH355	ZTH365	ZTH375	S1AR22	XX7455
6	ZTH346	ZTH356	ZTH366	ZTH376	S1AR23	XX7456
7	ZTH347	ZTH357	ZTH367	ZTH377	S1AR24	XX7457
8	ZTH348	ZTH358	ZTH368	ZTH378	S1AR25	XX7458
9	ZTH349	ZTH359	ZTH369	ZTH379	S1AR16	XX7459
A	ZTH34A	ZTH35A	ZTH36A	ZTH37A	S1AR27	XX745A
B	ZTH34B	ZTH35B	ZTH36B	ZTH37B	S1AR28	XX745B
C	ZTH34C	ZTH35C	ZTH36C	ZTH37C	S1AR29	XX745C
D	ZTH34D	ZTH35D	ZTH36D	ZTH37D	S1AR30	XX745D
E	ZTH34E	ZTH35E	ZTH36E	ZTH37E	S1AR31	XX745E
F	ZTH34F	ZTH35F	ZTH36F	ZTH37F	XX744F	XX745F

Figure 3-3 Sysplex volume layout

Figure 3-4 shows the details of our sysplex application volumes, which are the GDPS Metro primary volumes.

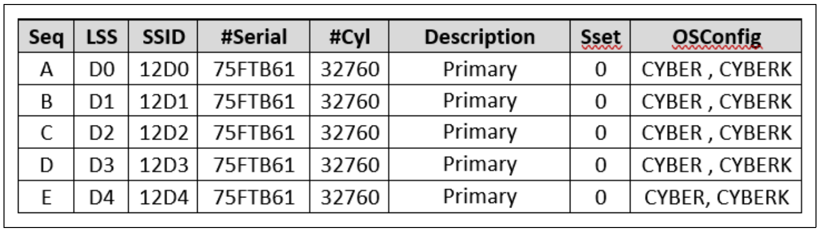

Seq	LSS	SSID	#Serial	#Cyl	Description	Sset	OSConfig
A	D0	12D0	75FTB61	32760	Primary	0	CYBER , CYBERK
B	D1	12D1	75FTB61	32760	Primary	0	CYBER , CYBERK
C	D2	12D2	75FTB61	32760	Primary	0	CYBER , CYBERK
D	D3	12D3	75FTB61	32760	Primary	0	CYBER , CYBERK
E	D4	12D4	75FTB61	32760	Primary	0	CYBER, CYBERK

Figure 3-4 Application sysplex volume layout: Metro primary

In Figure 3-4:

- CYBER is the OSConfig that is used by the application systems.
- CYBERK is the OSConfig for the GDPS controlling system.

All other details are used when preparing the GDPS Logical Corruption Protection configuration.

3.3 Setting up IBM Z Cyber Vault storage environment

This section describes the steps that are needed to set up the storage environment for IBM Z Cyber Vault. The steps include defining the Safeguarded Copy storage environment, and staging volumes that can be used if a surgical or catastrophic recovery is required. Also, persistent volumes contain information that is produced by the IBM Z Cyber Vault validation process. For a description of the volume types, see 2.2, “IBM Z Cyber Vault reference architecture” on page 24.

The input/output definition file (IODF) definitions and considerations for the z/OS couple data set (CDS) definitions and IBM Workload Manager for z/OS considerations are also described.

Our IBM Z Cyber Vault environment is based on a GDPS Metro solution with virtual airgap LCP to manage the Safeguarded Copy backups by using the Metro secondary, and the creation of a recover copy (RCx) set of volumes to be used for the validation process.

3.3.1 Metro secondary and Safeguarded Copy

Our GDPS Metro secondary volumes are a few hundred meters distant from the primary volumes. The environment is HyperSwap enabled for high availability (HA). All the GDPS Metro secondary volumes are defined in multiple subchannel set 1 (MSS1).

Figure 3-5 contains the details of the Metro secondary DS8886 system.

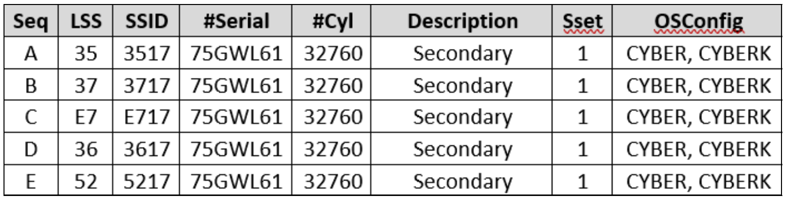

Seq	LSS	SSID	#Serial	#Cyl	Description	Sset	OSConfig
A	35	3517	75GWL61	32760	Secondary	1	CYBER, CYBERK
B	37	3717	75GWL61	32760	Secondary	1	CYBER, CYBERK
C	E7	E717	75GWL61	32760	Secondary	1	CYBER, CYBERK
D	36	3617	75GWL61	32760	Secondary	1	CYBER, CYBERK
E	52	5217	75GWL61	32760	Secondary	1	CYBER, CYBERK

Figure 3-5 Metro secondary volume layout

Each Metro secondary volume has a definition for Safeguarded Copy backup capacity in the DS8886 system, as described in 2.4, “Prerequisites and considerations for implementing Safeguarded Copy” on page 41. The physical and virtual capacity depends on your retention period, capture, or backup frequency, and the data change rate.

For each volume, we must define the virtual capacity multiplier.

We decided to take a Safeguarded Copy backup (capture) every 4 hours with a retention period of four days. Therefore, our virtual capacity is set to a maximum value of 24 for all the volumes, which is the number of backups during our retention period. We can use this simple approach for our small environment. Depending on your environment, you might need to do a proper sizing, as described in 3.2, “Description of our IBM Z Cyber Vault environment” on page 71. The Safeguarded Copy virtual capacity must be set up in the DS8000 system, either by using the GUI or the `dscli` commands.

> **Tips:** A best practice is to define all the volumes and Safeguarded Copy space in the same storage extent pools of the DS8000 system.
>
> The virtual multiplier can be dynamically increased, but not decreased.

3.3.2 Recovery copy

In our IBM Z Cyber Vault environment, we define two recovery sets of volumes, RC1 and RC2, for every Metro secondary volume. We defined them as extent space efficient (ESE) volumes and used the attributes that are shown in Figure 3-6.

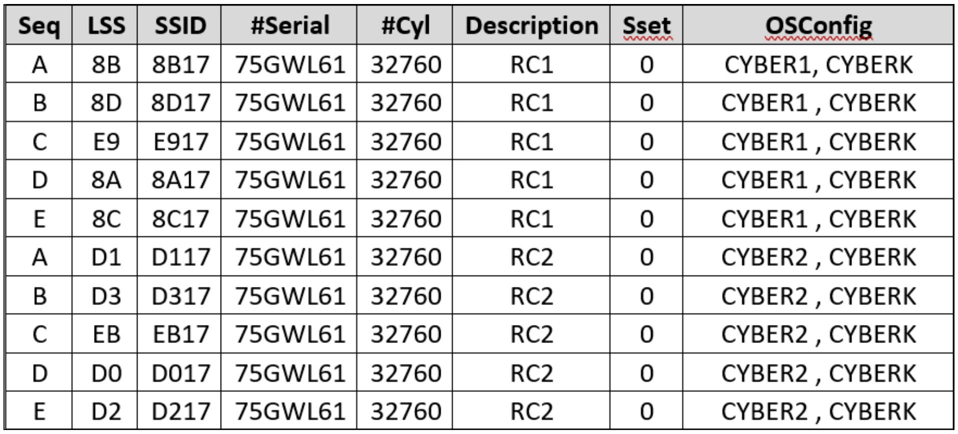

Seq	LSS	SSID	#Serial	#Cyl	Description	Sset	OSConfig
A	8B	8B17	75GWL61	32760	RC1	0	CYBER1, CYBERK
B	8D	8D17	75GWL61	32760	RC1	0	CYBER1 , CYBERK
C	E9	E917	75GWL61	32760	RC1	0	CYBER1 , CYBERK
D	8A	8A17	75GWL61	32760	RC1	0	CYBER1 , CYBERK
E	8C	8C17	75GWL61	32760	RC1	0	CYBER1 , CYBERK
A	D1	D117	75GWL61	32760	RC2	0	CYBER2 , CYBERK
B	D3	D317	75GWL61	32760	RC2	0	CYBER2 , CYBERK
C	EB	EB17	75GWL61	32760	RC2	0	CYBER2 , CYBERK
D	D0	D017	75GWL61	32760	RC2	0	CYBER2 , CYBERK
E	D2	D217	75GWL61	32760	RC2	0	CYBER2 , CYBERK

Figure 3-6 RC1 and RC2 set of volumes that are defined in the Site 2 DS8000 system

The volumes are defined as follows:

- RC1 volumes are defined in a new CYBER1 OSConfig.
- RC2 volumes are defined in a new CYBER2 OSConfig.

These OSConfigs are used to perform an IPL of the recovered systems in the IBM Z Cyber Vault environment. They are both also defined in the OSConfig of GDPS Logical Corruption Protection Manager.

Tip: All the volumes for RC1 and RC2 in the OSConfig for the recovered systems must be defined in MSS0 because of the need to perform an IPL of the recovered systems.

If a full copy on RC1 is required (depending on the planned usage and based on your capacity sizing (see 2.4, “Prerequisites and considerations for implementing Safeguarded Copy” on page 41)), the space must be the same as that on the RS2 volume that is provisioned on the DS8000 system to be able to convert the relationship from NOCOPY to COPY.

3.3.3 Staging and persistent volumes

In the case of a surgical recovery, a set of staging volumes can be used in the site that is hosting the LCP environment.

A volume that contains data sets, tables, or any other object to be restored in production can be copied to the staging volumes from RC1 or RC2. At the time of writing, you cannot take a FlashCopy copy from Safeguarded Copy recovered volumes, so in such cases you must perform a Global Copy (GC) to the staging volumes. You can use a FlashCopy command if RC1 or RC2 was established by using an **FCESTABLISH** script command.

In a physically isolated LCP environment and a Global Mirror (GM) virtual isolated environment where the primary volumes are asynchronously mirrored to the secondary volumes, the staging volumes must be copied back to the production site by using GC.

In our GDPS Metro Virtual Airgap environment, the staging volumes in Site 2, which are defined in MSS0, can be added to both the production OSConfig and the OSConfig that are used by the recovered systems. This way, we can vary the production sysplex staging volumes online after changing the labels.

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

To save historical files from a validation run, we defined more *persistent volumes* in Site2, again in MSS0, that can also be defined to the production OSConfig as OFFLINE at IPL. *Persistent* means that these volumes must not be lost between different persistent volumes in Site 2 in MSS0. These volumes also can be defined to the production OSrent validation runs. They can be Systems Management Facility (SMF) records and historical information that should be used to identify the last validated copy before a corruption.

In asynchronous mirroring solutions, these persistent volumes might need to be copied back to the production site by using GC.

In Figure 3-7 on page 77, note the following items:

- The first RC1 *LSS* is defined as 8B in the DS8886 75GWL61 storage system (as shown in Figure 3-2 on page 73).
- The first 42 volumes are defined as RC1, range 00 - 29.
- The *cca* range 2A - 2D is defined as persistent volumes for historical files for any validation run (label ZCV001-004).
- The *cca* range 2E - 2F is defined as spare space in the first *LSS* containing RC1, instead of generating more LSSs.
- The *cca* range 30 - 5F (8B30 - 8B5F) is defined as staging volumes.

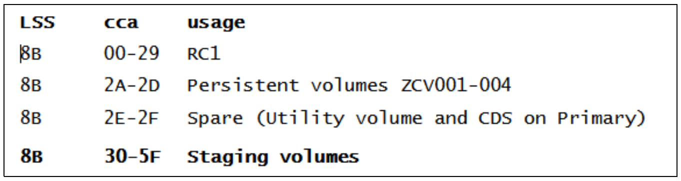

LSS	cca	usage
8B	00-29	RC1
8B	2A-2D	Persistent volumes ZCV001-004
8B	2E-2F	Spare (Utility volume and CDS on Primary)
8B	**30-5F**	**Staging volumes**

Figure 3-7 Staging volumes (8B30-8B5F)

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

On one of the persistent volumes is 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/Job Control Language (JCL), and it is where validation output is written. This user catalog is also connected to the production sysplex.

Connecting to the production system allows us to see output that is produced during the validation process in the production environment.

3.3.4 Recovery in case of full data corruption

For a full recovery in an asynchronous solution, the validated RC1 volumes should be copied back to the RS1 production volumes by using GC. The GC function restores the good consistent data preceding the corruption. In such a configuration, it is convenient to restart the application systems in the DR region from the RC1 volumes because the full copy might take longer than the restart.

> **Note:** When using CSM V6.3.0 or later and a DS89xx storage system with Microcode Release 9.2 or later, it is possible to send data incrementally by using GC.
>
> CSM V6.3.0 has a function that is called Restore in the Safeguarded Copy session that will copy data back to the production volumes. First, you must suspend the original replication from the production system to the Safeguarded Copy location, and then use the Restore option to do a GC back to the production system incrementally.

In our GDPS Metro environment, we have a quick option to restart on RC1, that is, reverting the definition of RC1 with RS2 in the GEOPARM, with an OSConfig that is prepared for the production systems with the former RC1 and the CDS volumes defined ONLINE at IPL, and the RS1 OFFLINE at IPL, so that we can issue a GDPSCHG ZDISK STOR(RS2) and perform another IPL of our production systems directly from the former RC1, now RS2.

Next, we reestablish the GDPS Metro Mirror (MM) from RS2 to RS1, HyperSwap to RS1, and reestablish the normal configuration. For this task, we define an extra CYBERR OSConfig.

For more information about this type of recovery, see the following resources:

- *IBM DS8000 Safeguarded Copy (Updated for DS8000 R9.2)*, REDP-5506
- *IBM GDPS Family: An Introduction to Concepts and Capabilities*, SG24-6374

> **Note:** The GDPS and LCP documentation is available only to users with a specific GDPS license. For more information about GDPS system recovery offerings and options for the IBM Z platform, contact the IBM GDPS team at mailto:gdps@us.ibm.com.

3.3.5 IODF definitions for the IBM Z Cyber Vault environment

The OSConfig definitions that are defined in our IODF for the IBM Z Cyber Vault environment are summarized in Figure 3-8.

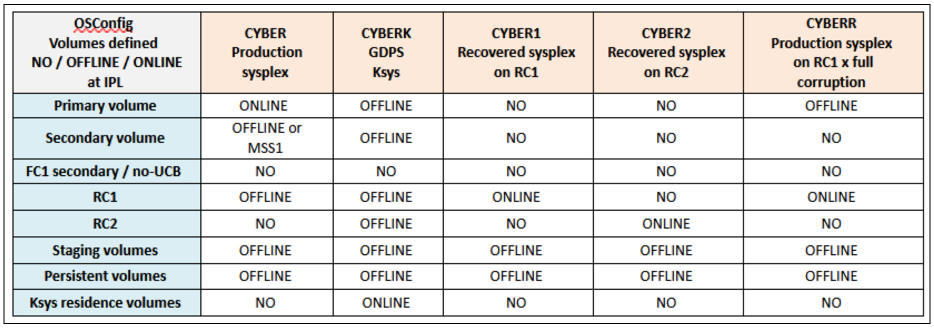

OSConfig Volumes defined NO / OFFLINE / ONLINE at IPL	CYBER Production sysplex	CYBERK GDPS Ksys	CYBER1 Recovered sysplex on RC1	CYBER2 Recovered sysplex on RC2	CYBERR Production sysplex on RC1 x full corruption
Primary volume	ONLINE	OFFLINE	NO	NO	OFFLINE
Secondary volume	OFFLINE or MSS1	OFFLINE	NO	NO	NO
FC1 secondary / no-UCB	NO	NO	NO	NO	NO
RC1	OFFLINE	OFFLINE	ONLINE	NO	ONLINE
RC2	NO	OFFLINE	NO	ONLINE	NO
Staging volumes	OFFLINE	OFFLINE	OFFLINE	OFFLINE	OFFLINE
Persistent volumes	OFFLINE	OFFLINE	OFFLINE	OFFLINE	OFFLINE
Ksys residence volumes	NO	ONLINE	NO	NO	NO

Figure 3-8 IODF OSConfig

Tip: Define all volumes OFFLINE at IPL to the GDPS controlling system. This action also can be used in an emergency to change any wrong parameters at IPL of RC1/RC2, or any rescue function, by varying the needed volumes online.

3.3.6 Couple data sets for the IBM Z Cyber Vault environment

Special considerations must be accounted for regarding the CDS in the IBM Z Cyber Vault environment.

In a HyperSwap setup, the production CDSs volumes are not mirrored, but the LOGR CDSs are. Therefore, the production CDSs volumes are not in the GDPS GEOPARM, and the Safeguarded Copy and RC1 or RC2 copy set do not contain those volumes.

To restart the sysplex from the RC1 or RC2 set of volumes, we need more CDSs. The two options for restarting the sysplex are as follows:

- Pre-allocate CDSs on dedicated volumes that are defined in MSS0 and ONLINE at IPL to the CYBER1 or CYBER2 OSConfig in the recovery region, except for the LOGR.
- Pre-allocate on mirrored volumes in Site 1 all the CDSs, except for the LOGR. The LOGR will be a GDPS MM to Site 2 and included in the Safeguarded Copy captures and the FlashCopy to RC1. For more information, see Figure 3-1 on page 72.

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

Tip: Because the Cross System Coupling Facility (XCF) CDS is newly allocated and never activated, it does not contain the name of the last CFRM policy that was used. Therefore, the CFRM policy name must be specified in the COUPLExx member.

These pre-allocated CDSs must be primed with the current policies that are used in production. To accomplish this task, specify the CDS name in the IXCMIAPU utility JCL for the ARM, CFRM, and SFM policies. Do not specify the CDS for the Workload Manager (WLM) policy. For more information, see 3.3.7, “Considerations for pre-allocated WLM CDSs” on page 79.

Important: Remember to update your day-to-day CDS management procedure to include the proper steps to refresh changes to the production policies. This update also applies to the pre-allocated CDSs.

3.3.7 Considerations for pre-allocated WLM CDSs

The WLM policy cannot be primed in the pre-allocated CDSs by using a batch job (JCL), but it can be exported from the production WLM CDS to a partitioned data set, and then imported from that partitioned data set to a pre-allocated CDS for the first system IPL in the IBM Z Cyber Vault environment.

The WLM utility to import the policy is IWMARIDU, and a sample batch job (JCL) can be found in `SYS1.SAMPLIB(IWMINSTL)`. The import batch job can be automatically started at IPL.

A default WLM policy is activated if you use pre-allocated CDSs without importing any policy.

3.3.8 Considerations for pre-allocated XCF CDSs

From an XCF perspective, pre-allocated CDSs are sufficient. However, there are other sysplex users that store information in the XCF CDS. These users expect this content to persist.

For example, Virtual Storage Access Method (VSAM) record-level sharing (RLS) stores the name of the sharing control data sets (SHCDS) and IBM MQ registers information regarding queue-sharing groups in the XCF CDS.

Such information cannot be re-created in the pre-allocated 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 pre-allocated XCF CDS, when the Short Message Service Virtual Storage Access Method (SMSVSAM) address space is started, the first production system issues a prompt because it does not find SHCDS information in the XCF CDS.

When SMSVSAM is initialized in a sysplex environment with pre-allocated XCF CDSs, the information in the SHCDSs is lost if you use the **`V SMS,SHCDS(dsn),NEW`** command. This command destroys lost lock data and other important information in the SHCDS.

To remedy this situation, a new option is added to the SMS SHCDS command that is called **`V SMS,SHCDS(dsn),OLD`**. This command allows an existing SHCDS to be made ACTIVE without formatting its contents.

At this point, you must run commands to define the SHCDS to be stored in the XCF CDS, for example:

- **`V SMS,SHCDS(PRIMARY.Z1SHC1),OLD`**
- **`V SMS,SHCDS(SECONDRY.Z1SHC2), OLD`**
- **`V SMS,SHCDS(SPARE.Z1SHCS),NEWSPARE`**

The interface between SMSVSAM and the Customer Information Control System (CICS) allows CICS to understand and resolve these lost locks in most cases. The only remaining action to be taken is to start the CICS regions, which communicate with SMSVSAM and work to resolve outstanding transactions based on the CICS logging.

For more information about the commands, see the following resources:

- The *DFSMS Storage Administration* manual
- The *MVS System Commands* manual
- APAR OA58064

Before starting a workload that requires IBM MQ services be started, use the IBM MQ -supplied CSQ5PQSG utility to reregister the IBM MQ sharing-groups in the XCF CDS.

For more information about using this utility, see the IBM MQ publications for your release of IBM MQ. For IBM MQ V8, the queue-sharing group utility (CSQ5PQSG) publications are available at IBM Documentation.

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

3.4 Setting up Safeguarded Copy in IBM Z Cyber Vault storage

In this section, we provide the necessary information to guide you in the setup of the Safeguarded Copy environment for the GDPS Metro virtual airgap solution.

For more information about the GDPS Logical Corruption Protection setup, see *GDPS Logical Corruption Protection Manager V4R4 Planning, Implementation and Operations Guide*, ZG24-4673.

> **Note**: The GDPS and LCP documentation is available only to users with a specific GDPS license.

The management profiles appear when you first enter the GDPS LPS panel from the main GDPS panel by using the L option, as shown in Figure 3-9.

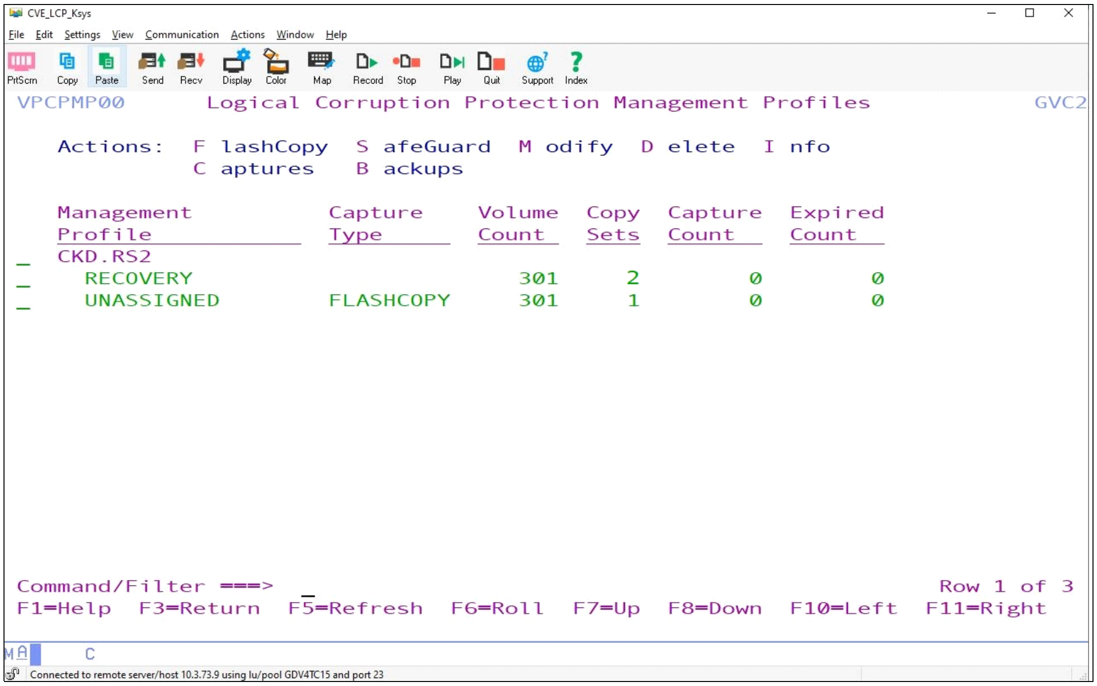
CVE_LCP_Ksys
File Edit Settings View Communication Actions Window Help
PrtScrn Copy Paste Send Recv Display Color Map Record Stop Play Quit Support Index
MA
C

```
 VPCPMP00       Logical Corruption Protection Management Profiles           GVC2

    Actions:  F lashCopy  S afeGuard  M odify  D elete  I nfo
              C aptures   B ackups

    Management          Capture    Volume  Copy  Capture  Expired
    Profile             Type       Count   Sets  Count    Count
 _  CKD.RS2
 _    RECOVERY                        301     2        0        0
 _    UNASSIGNED        FLASHCOPY     301     1        0        0











 Command/Filter ===>  _                                              Row 1 of 3
 F1=Help  F3=Return  F5=Refresh  F6=Roll  F7=Up  F8=Down  F10=Left  F11=Right
```

Connected to remote server/host 10.3.73.9 using lu/pool GDV4TC15 and port 23

Figure 3-9 Logical Corruption Protection panel at first entry

The following details are shown in Figure 3-9 on page 80:

- The LCP environment is defined on our RS2 DS8000, which is the Metro secondary one.
- Two RECOVERY sets are defined with 301 volumes, which is the full Metro configuration.
- One FlashCopy set is defined. However, the volumes are UNASSIGNED and not used in our scenario.
- You cannot see a Safeguarded Copy set because you do not have a Safeguarded Copy management profile that is defined yet.
- Captures (also known as Safeguarded Copy backups) are not yet taken.

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

1. Type `S` on the `CKD.RS2` line to define a Safeguarded Copy management profile with the following definitions:
   - `Management Profile`: SGC1GOLD
   - `Retention Period`: 4 days
   - `Minimum Interval`: 1 hour

   You do not change the default `Reservation Time` and `Check In Time` default values, as shown in Figure 3-10.

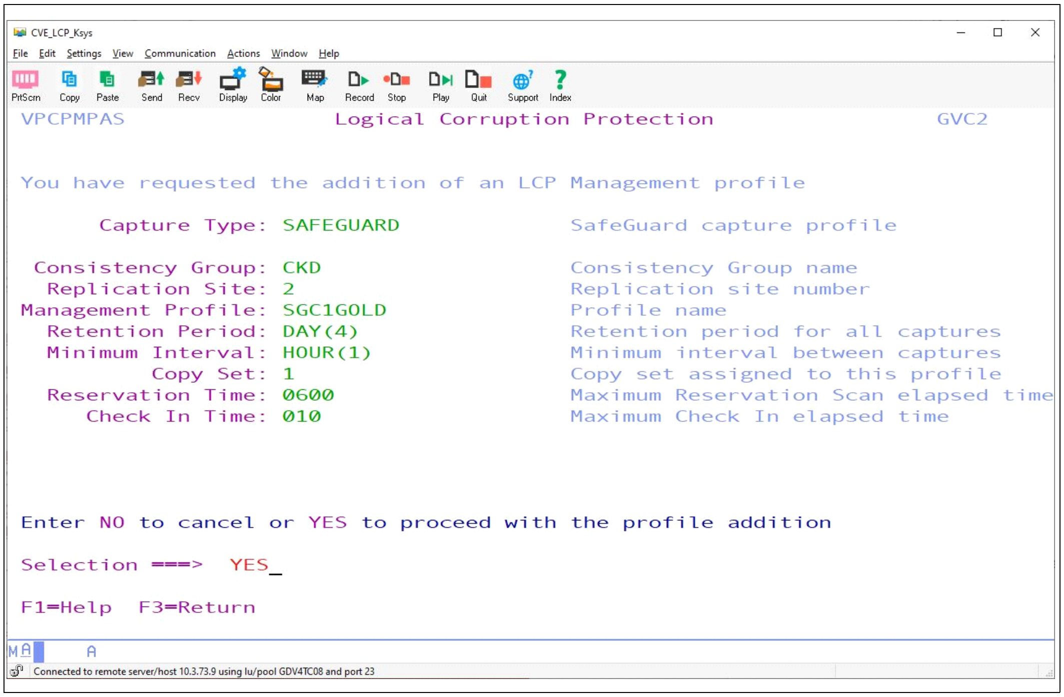
CVE_LCP_Ksys
File Edit Settings View Communication Actions Window Help
PrtScrn Copy Paste Send Recv Display Color Map Record Stop Play Quit Support Index

```
VPCPMPAS                  Logical Corruption Protection                  GVC2


You have requested the addition of an LCP Management profile

      Capture Type:  SAFEGUARD                 SafeGuard capture profile

 Consistency Group:  CKD                       Consistency Group name
  Replication Site:  2                         Replication site number
Management Profile:  SGC1GOLD                  Profile name
  Retention Period:  DAY(4)                    Retention period for all captures
  Minimum Interval:  HOUR(1)                   Minimum interval between captures
          Copy Set:  1                         Copy set assigned to this profile
  Reservation Time:  0600                      Maximum Reservation Scan elapsed time
     Check In Time:  010                       Maximum Check In elapsed time



Enter NO to cancel or YES to proceed with the profile addition

Selection ===>  YES_

F1=Help  F3=Return
```

MA A
Connected to remote server/host 10.3.73.9 using lu/pool GDV4TC08 and port 23

Figure 3-10 Safeguarded Copy management profile definition

In the Logical Corruption Protection Management Profiles panel, you find the SGC1GOLD management profile with 301 volumes and captures not yet taken (see Figure 3-11).

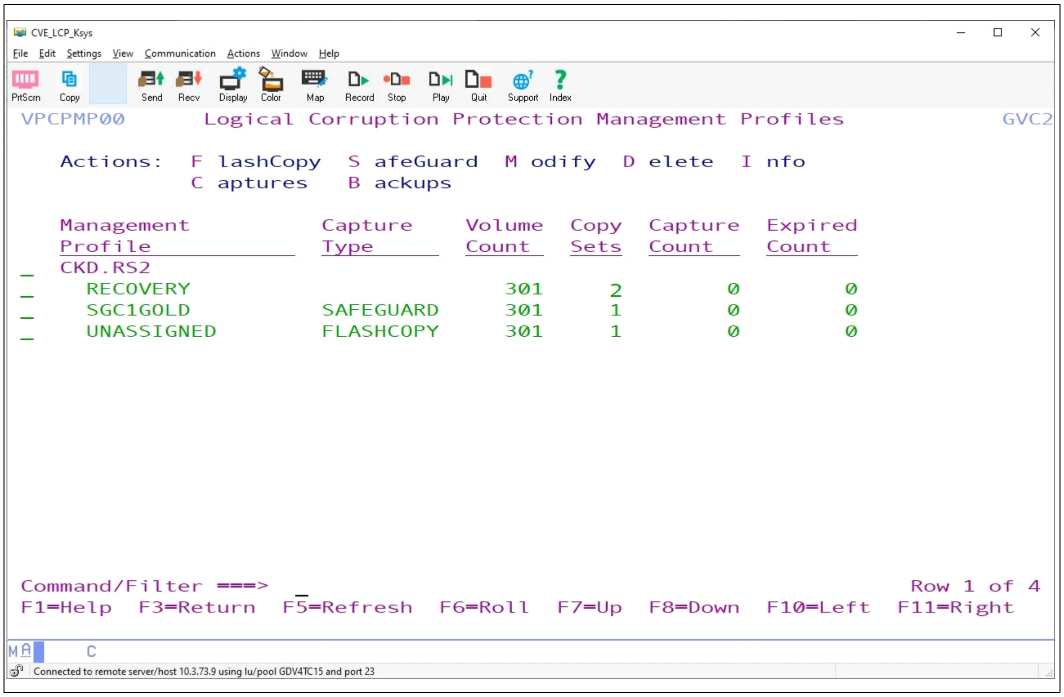
CVE_LCP_Ksys
File Edit Settings View Communication Actions Window Help
PrtScrn Copy Send Recv Display Color Map Record Stop Play Quit Support Index

```
VPCPMP00       Logical Corruption Protection Management Profiles             GVC2

    Actions:  F lashCopy  S afeGuard  M odify  D elete  I nfo
              C aptures   B ackups

    Management                Capture     Volume  Copy  Capture  Expired
    Profile                   Type        Count   Sets  Count    Count
    CKD.RS2
 _    RECOVERY                              301     2       0        0
 _    SGC1GOLD                SAFEGUARD     301     1       0        0
 _    UNASSIGNED              FLASHCOPY     301     1       0        0

Command/Filter ===> _                                                Row 1 of 4
F1=Help  F3=Return  F5=Refresh  F6=Roll  F7=Up  F8=Down  F10=Left  F11=Right
```

MA   C
Connected to remote server/host 10.3.73.9 using lu/pool GDV4TC15 and port 23

Figure 3-11 LCP Management Profiles panel with the new SGC1GOLD profile

2. Take the first Safeguarded Copy capture. To do this task, you first must define a GDPS script, like the script that is shown in Figure 3-12.

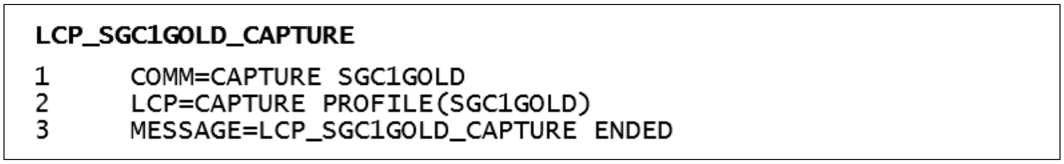

```
LCP_SGC1GOLD_CAPTURE

1      COMM=CAPTURE SGC1GOLD
2      LCP=CAPTURE PROFILE(SGC1GOLD)
3      MESSAGE=LCP_SGC1GOLD_CAPTURE ENDED
```

Figure 3-12 Sample script to take a Safeguarded Copy capture

> **Note**: The `MESSAGE` statement in each script is needed only if you intend to use the automation and check process, as described in 3.4.1, “Automating Safeguarded Copy captures” on page 85.

Figure 3-13 on page 83 shows the SDF trace entries that are generated by the GDPS script that captures a Safeguarded Copy backup.

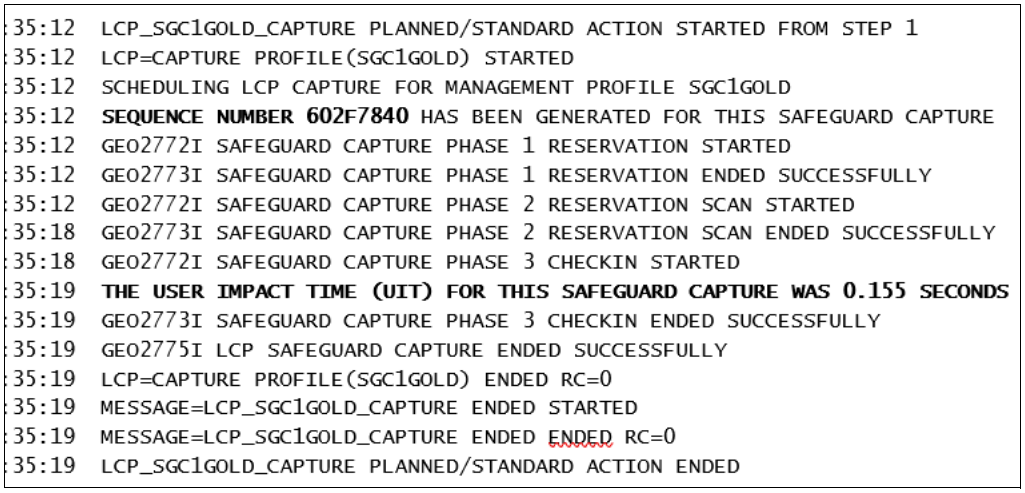

```
:35:12  LCP_SGC1GOLD_CAPTURE PLANNED/STANDARD ACTION STARTED FROM STEP 1
:35:12  LCP=CAPTURE PROFILE(SGC1GOLD) STARTED
:35:12  SCHEDULING LCP CAPTURE FOR MANAGEMENT PROFILE SGC1GOLD
:35:12  SEQUENCE NUMBER 602F7840 HAS BEEN GENERATED FOR THIS SAFEGUARD CAPTURE
:35:12  GEO2772I SAFEGUARD CAPTURE PHASE 1 RESERVATION STARTED
:35:12  GEO2773I SAFEGUARD CAPTURE PHASE 1 RESERVATION ENDED SUCCESSFULLY
:35:12  GEO2772I SAFEGUARD CAPTURE PHASE 2 RESERVATION SCAN STARTED
:35:18  GEO2773I SAFEGUARD CAPTURE PHASE 2 RESERVATION SCAN ENDED SUCCESSFULLY
:35:18  GEO2772I SAFEGUARD CAPTURE PHASE 3 CHECKIN STARTED
:35:19  THE USER IMPACT TIME (UIT) FOR THIS SAFEGUARD CAPTURE WAS 0.155 SECONDS
:35:19  GEO2773I SAFEGUARD CAPTURE PHASE 3 CHECKIN ENDED SUCCESSFULLY
:35:19  GEO2775I LCP SAFEGUARD CAPTURE ENDED SUCCESSFULLY
:35:19  LCP=CAPTURE PROFILE(SGC1GOLD) ENDED RC=0
:35:19  MESSAGE=LCP_SGC1GOLD_CAPTURE ENDED STARTED
:35:19  MESSAGE=LCP_SGC1GOLD_CAPTURE ENDED ENDED RC=0
:35:19  LCP_SGC1GOLD_CAPTURE PLANNED/STANDARD ACTION ENDED
```

Figure 3-13 Sample SDF trace entries when capturing a Safeguarded Copy

> **Note**: The user impact time (UIT) in the SDF trace entries is the time that your applications were frozen to take consistent data, which is 0.155 seconds, as shown in Figure 3-13.

`SEQUENCE NUMBER 602F7840` is the hexadecimal representation of the timestamp for the capture. The timestamp is found in the Logical Corruption Protection SafeGuard Captures panel.

Figure 3-14 shows the GDPS Logical Corruption Protection Safeguarded Captures panel with the first backup copy captured.

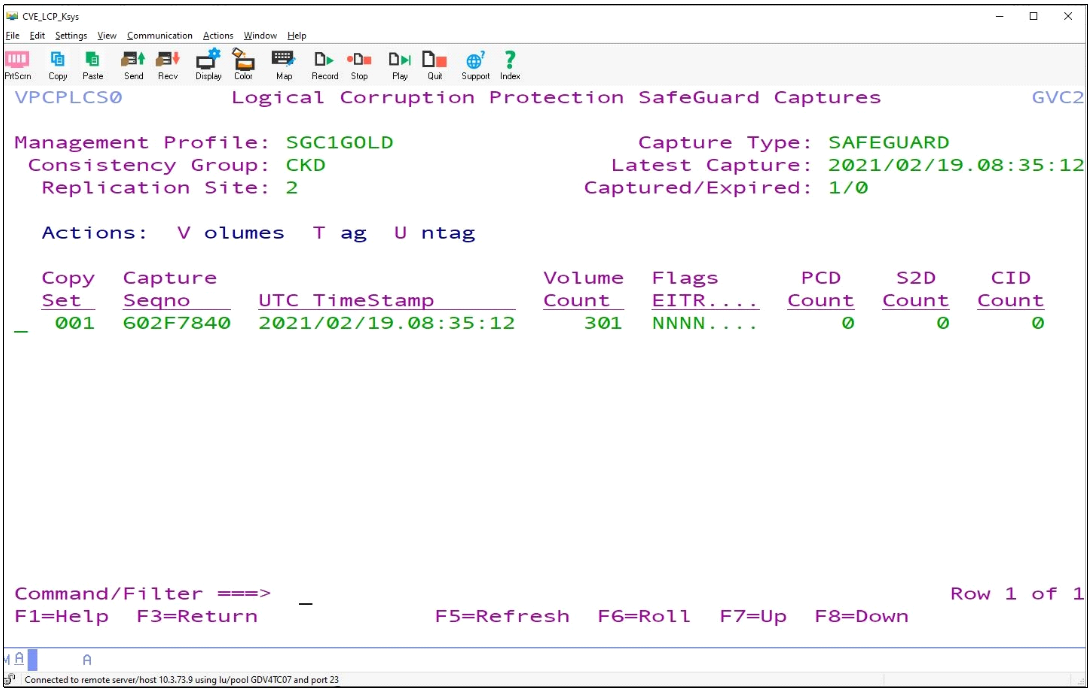

CVE_LCP_Ksys
File Edit Settings View Communication Actions Window Help
PrtScrn Copy Paste Send Recv Display Color Map Record Stop Play Quit Support Index

```
VPCPLCS0         Logical Corruption Protection SafeGuard Captures          GVC2

Management Profile: SGC1GOLD                Capture Type: SAFEGUARD
 Consistency Group: CKD                   Latest Capture: 2021/02/19.08:35:12
  Replication Site: 2                   Captured/Expired: 1/0

  Actions:  V olumes  T ag  U ntag

  Copy  Capture                        Volume  Flags      PCD    S2D    CID
  Set   Seqno     UTC TimeStamp        Count   EITR....   Count  Count  Count
_  001  602F7840  2021/02/19.08:35:12    301   NNNN....      0      0      0








Command/Filter ===>  _                                                Row 1 of 1
F1=Help  F3=Return            F5=Refresh  F6=Roll  F7=Up  F8=Down
```

Connected to remote server/host 10.3.73.9 using lu/pool GDV4TC07 and port 23

Figure 3-14 Logical Corruption Protection Safeguard Captures

In our IBM Z Cyber Vault use case, we plan to automatically take one Safeguarded Copy backup (capture) every 4 hours. Concurrently with one of the Safeguarded Copy backups (captures), we also take a direct FlashCopy from RS2 to RC1. The validation process runs on this direct RC1 FlashCopy to make sure that data corruption does not exist up to this point in time.

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

- Only a single Safeguarded Recovery relationship is supported. With this single relationship, you can do a Safeguarded Recovery, if needed, without interrupting the current data-validation process.
- The read I/O on the recovered relationship goes through the Safeguarded Copy metadata for every read, which negatively impacts the overall DS8000 workload perspective and the speed of the validation process. Also, if you are taking further backups after the recovery, the read performance gets worse over time because you continuously have more metadata to scan for the reads.

To take a direct FlashCopy from RS2 to RC1, use a script like the script that is shown in Figure 3-15.

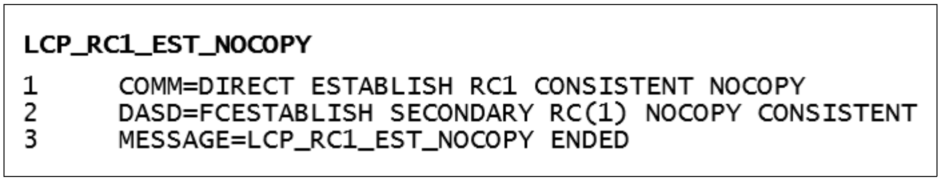

```
LCP_RC1_EST_NOCOPY
1      COMM=DIRECT ESTABLISH RC1 CONSISTENT NOCOPY
2      DASD=FCESTABLISH SECONDARY RC(1) NOCOPY CONSISTENT
3      MESSAGE=LCP_RC1_EST_NOCOPY ENDED
```

Figure 3-15 Sample script to take a direct FlashCopy from RS2 to RC1

When you run this script, you receive the following trace entries, as shown in Figure 3-16.

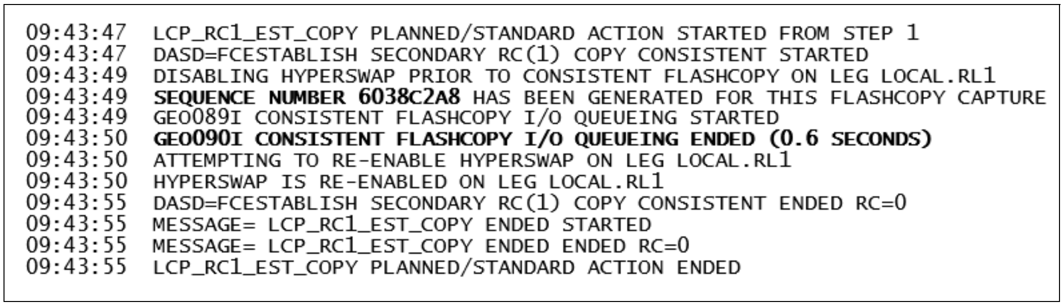

```
09:43:47  LCP_RC1_EST_COPY PLANNED/STANDARD ACTION STARTED FROM STEP 1
09:43:47  DASD=FCESTABLISH SECONDARY RC(1) COPY CONSISTENT STARTED
09:43:49  DISABLING HYPERSWAP PRIOR TO CONSISTENT FLASHCOPY ON LEG LOCAL.RL1
09:43:49  SEQUENCE NUMBER 6038C2A8 HAS BEEN GENERATED FOR THIS FLASHCOPY CAPTURE
09:43:49  GEO089I CONSISTENT FLASHCOPY I/O QUEUEING STARTED
09:43:50  GEO090I CONSISTENT FLASHCOPY I/O QUEUEING ENDED (0.6 SECONDS)
09:43:50  ATTEMPTING TO RE-ENABLE HYPERSWAP ON LEG LOCAL.RL1
09:43:50  HYPERSWAP IS RE-ENABLED ON LEG LOCAL.RL1
09:43:55  DASD=FCESTABLISH SECONDARY RC(1) COPY CONSISTENT ENDED RC=0
09:43:55  MESSAGE= LCP_RC1_EST_COPY ENDED STARTED
09:43:55  MESSAGE= LCP_RC1_EST_COPY ENDED ENDED RC=0
09:43:55  LCP_RC1_EST_COPY PLANNED/STANDARD ACTION ENDED
```

Figure 3-16 Sample SDF trace entries when taking a direct FlashCopy from RS2 to RC1

Tip: The UIT for a direct FlashCopy is found in the GEO090I message. For this FlashCopy, it is 0.6 seconds.

In Figure 3-16, the `SEQUENCE NUMBER 6038C2A8` is the hexadecimal representation of the timestamp for the direct FlashCopy. The timestamp is displayed in the Logical Corruption Protection Recovery Captures panel, as shown in Figure 3-17 on page 85.

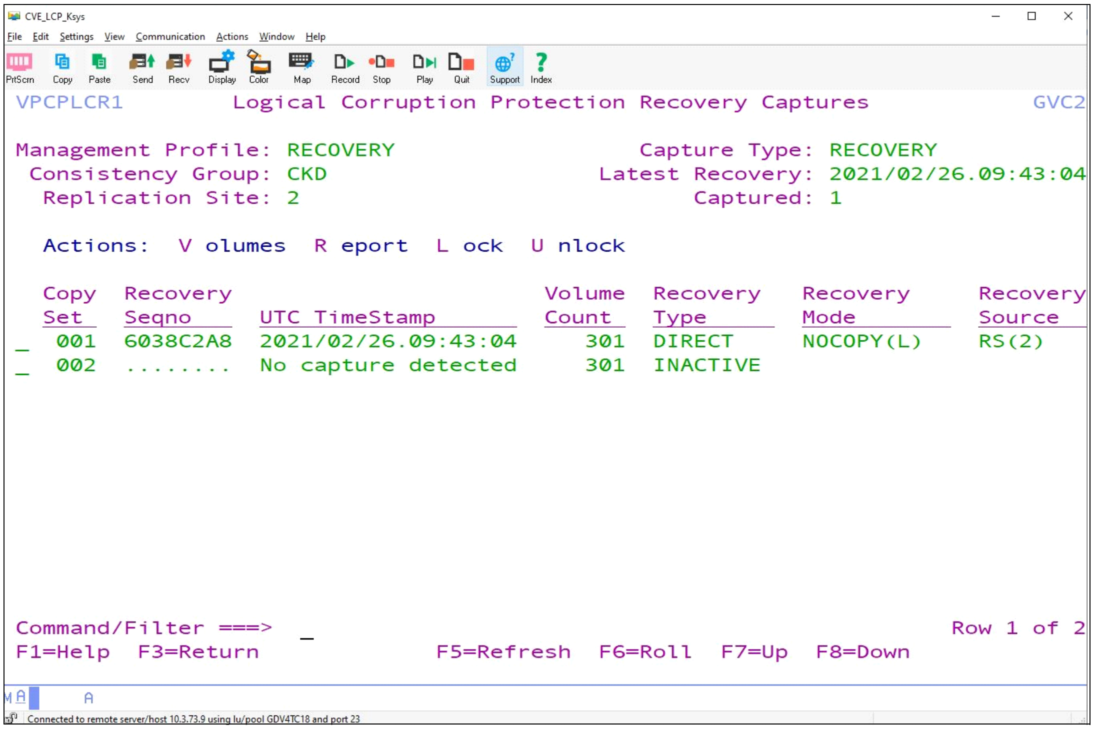
CVE_LCP_Ksys
File Edit Settings View Communication Actions Window Help
PrtScrn Copy Paste Send Recv Display Color Map Record Stop Play Quit Support Index
MA A

```
VPCPLCR1        Logical Corruption Protection Recovery Captures            GVC2

Management Profile: RECOVERY                  Capture Type: RECOVERY
 Consistency Group: CKD                    Latest Recovery: 2021/02/26.09:43:04
  Replication Site: 2                             Captured: 1

  Actions:  V olumes  R eport  L ock  U nlock

  Copy  Recovery                       Volume  Recovery   Recovery     Recovery
  Set   Seqno     UTC TimeStamp        Count   Type       Mode         Source
_  001  6038C2A8  2021/02/26.09:43:04     301  DIRECT     NOCOPY(L)    RS(2)
_  002  ........  No capture detected     301  INACTIVE








Command/Filter ===>  _                                               Row 1 of 2
F1=Help  F3=Return             F5=Refresh  F6=Roll  F7=Up  F8=Down
```

Connected to remote server/host 10.3.73.9 using lu/pool GDV4TC18 and port 23

Figure 3-17 Logical Corruption Protection Recovery Captures taken directly from RS2

The RC1 set of volumes is immediately available to perform an IPL of the recovered systems in the IBM Z Cyber Vault environment and start the validation process.

3.4.1 Automating Safeguarded Copy captures

Taking a new Safeguarded Copy capture is done by running a GDPS script on the GDPS K-System, which is a manual action. In an IBM Z Cyber Vault environment, you must automate the captures by using GDPS Logical Corruption Protection Manager to the time of day that is most appropriate for the business operations.

To do this task, you need a way to start a GDPS script by using a batch JCL, which can be done by using the GDPS RESTful API.

To use the RESTful API, you need a way to transfer data to and from servers by using HTTP or HTTPS. The samples in this section use cURL.

cURL is an open source tool that can transfer data to and from servers by using one of the supported protocols (HTTP, HTTPS, FTP, FTPS, TFTP, DICT, TELNET, LDAP, or FILE). This tool can work without user interaction.

The cURL functions include the following ones:

- Proxy support
- User authentication
- FTP upload
- HTTP post
- SSL connections
- Cookies
- File transfer
- Resume

To proceed, you first must install Gzip and cURL on the z/OS of the GDPS Logical Corruption Protection Manager. For more information, see Open Source Languages and Tools for z/OS.

When cURL is installed, you can submit the JCL, as shown in Figure 3-18, to start any GDPS script.

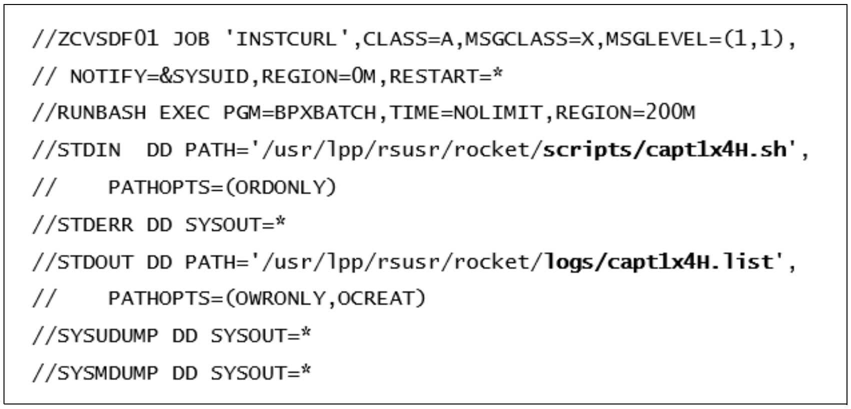

```
//ZCVSDF01 JOB 'INSTCURL',CLASS=A,MSGCLASS=X,MSGLEVEL=(1,1),
// NOTIFY=&SYSUID,REGION=0M,RESTART=*
//RUNBASH EXEC PGM=BPXBATCH,TIME=NOLIMIT,REGION=200M
//STDIN   DD PATH='/usr/lpp/rsusr/rocket/scripts/captlx4H.sh',
//     PATHOPTS=(ORDONLY)
//STDERR DD SYSOUT=*
//STDOUT DD PATH='/usr/lpp/rsusr/rocket/logs/captlx4H.list',
//     PATHOPTS=(OWRONLY,OCREAT)
//SYSUDUMP DD SYSOUT=*
//SYSMDUMP DD SYSOUT=*
```

Figure 3-18   RESTful API JCL to start a GDPS script

In Figure 3-18:

- cURL is installed in the `/usr/lpp/rsusr/rocket/` directory.
- We created a directory for all the scripts in `/scripts`.
- **`captlx4H.sh`** is the shell that we want to run, but it changes depending on the script that we want to run. The output is stored in `/logs`.

Figure 3-19 shows an example of RESTful API shell.

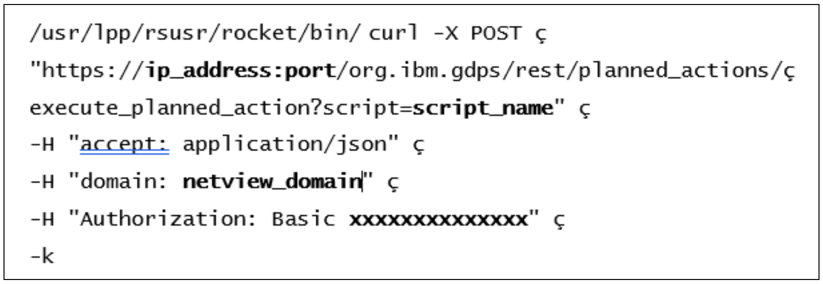

```
/usr/lpp/rsusr/rocket/bin/ curl -X POST ç
"https://ip_address:port/org.ibm.gdps/rest/planned_actions/ç
execute_planned_action?script=script_name" ç
-H "accept: application/json" ç
-H "domain: netview_domain" ç
-H "Authorization: Basic xxxxxxxxxxxxxx" ç
-k
```

Figure 3-19   RESTful API shell to start a GDPS script

In Figure 3-19, you must replace each parameter in **bold** with the values that match your environment:

- `ip_address`: IP address of the system running Liberty.
- `port`: Port number to access the Liberty server.
- `script_name`: Name of the GDPS script to run.
- `netview_domain`: The Tivoli NetView for z/OS domain name of your GDPS master K-System.
- `xxxxxxxxx`: User ID:password in the HTTP basic authentication form (user ID:password) encoded in Base64, for example, `oper5:mypass` in Base64 becomes `b3BlcjU6bXlwYXNz`. For more information, see RFC 7617.

Consider the following points:

- The specified user must be:
  - A TSO user with the OMVS segment and a Home directory.
  - Defined in Tivoli NetView for z/OS, but not in System Automation. The user must be dedicated to this service and must not be logged on at the time that the JCL is started.
- In the script name, the \ character must precede any $ to tell **bash** not to interpret the script name as a variable.

Tip: The continuation character that shows-up as ç in Figure 3-19 depends on your 3270-emulation charset. It must be hexadecimal x’E0’.

When a script is run, you can retrieve the execution SDF traces entries that result from using another RESTful API service, as shown in Figure 3-20.

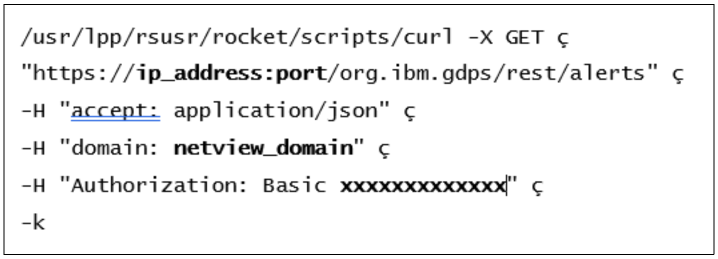

```
/usr/lpp/rsusr/rocket/scripts/curl -X GET ç
"https://ip_address:port/org.ibm.gdps/rest/alerts" ç
-H "accept: application/json" ç
-H "domain: netview_domain" ç
-H "Authorization: Basic xxxxxxxxxxxxx" ç
-k
```

Figure 3-20   RESTful API shell to retrieve SDF trace entries of a GDPS script run

The trace entries that are retrieved are in the format that is shown in Figure 3-21:

- In blue: The trace entries that are related to the capture process
- In bold: The most relevant trace entries

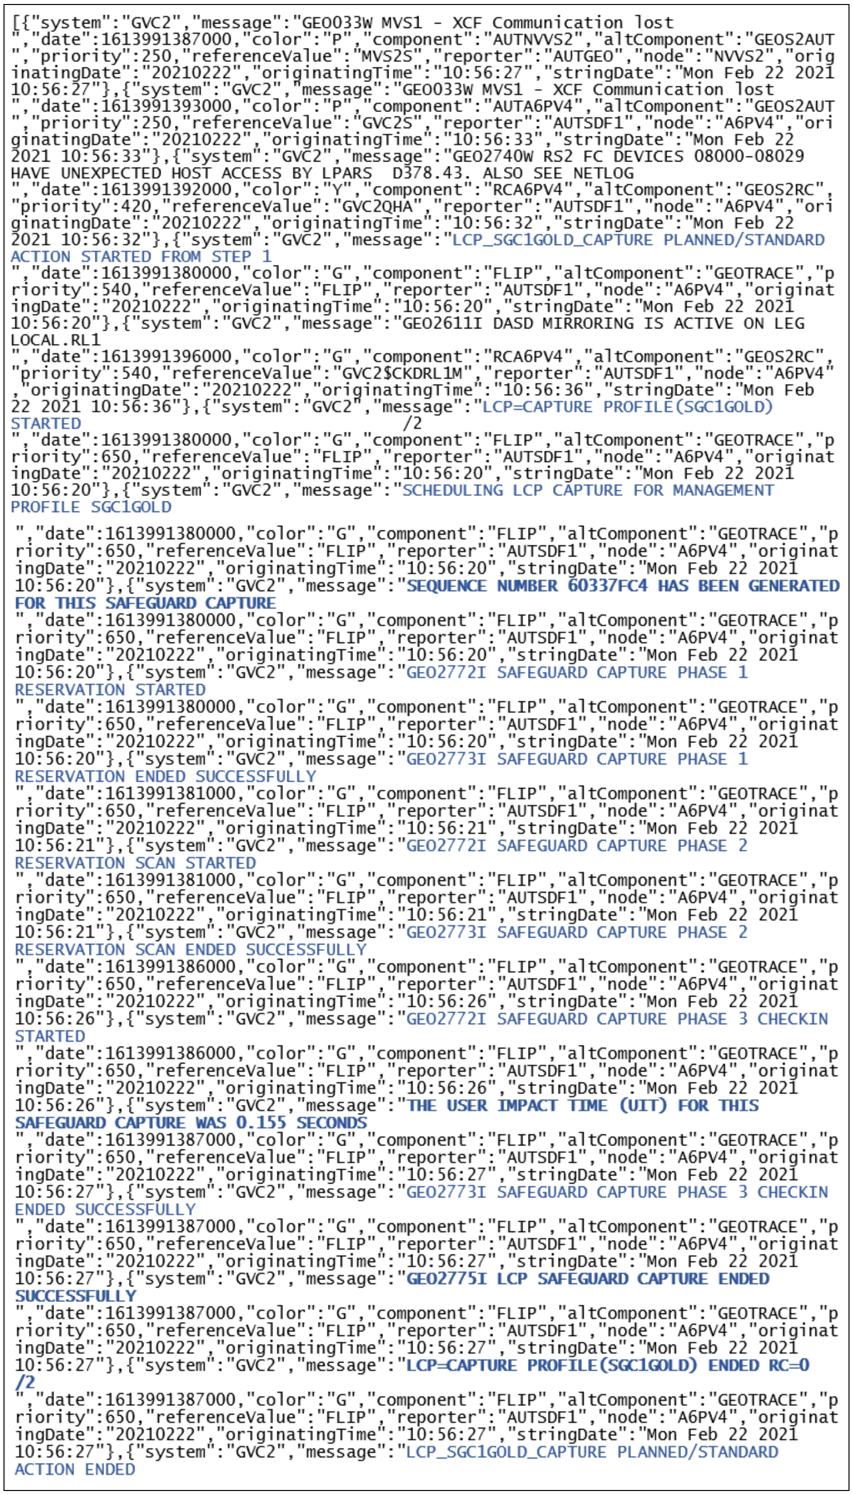

```
[{"system":"GVC2","message":"GEO033W MVS1 - XCF Communication lost
","date":1613991387000,"color":"P","component":"AUTNVVS2","altComponent":"GEOS2AUT
","priority":250,"referenceValue":"MVS2S","reporter":"AUTGEO","node":"NVVS2","orig
inatingDate":"20210222","originatingTime":"10:56:27","stringDate":"Mon Feb 22 2021
10:56:27"},{"system":"GVC2","message":"GEO033W MVS1 - XCF Communication lost
","date":1613991393000,"color":"P","component":"AUTA6PV4","altComponent":"GEOS2AUT
","priority":250,"referenceValue":"GVC2S","reporter":"AUTSDF1","node":"A6PV4","ori
ginatingDate":"20210222","originatingTime":"10:56:33","stringDate":"Mon Feb 22
2021 10:56:33"},{"system":"GVC2","message":"GEO2740W RS2 FC DEVICES 08000-08029
HAVE UNEXPECTED HOST ACCESS BY LPARS  D378.43. ALSO SEE NETLOG
","date":1613991392000,"color":"Y","component":"RCA6PV4","altComponent":"GEOS2RC",
"priority":420,"referenceValue":"GVC2QHA","reporter":"AUTSDF1","node":"A6PV4","ori
ginatingDate":"20210222","originatingTime":"10:56:32","stringDate":"Mon Feb 22
2021 10:56:32"},{"system":"GVC2","message":"LCP_SGC1GOLD_CAPTURE PLANNED/STANDARD
ACTION STARTED FROM STEP 1
","date":1613991380000,"color":"G","component":"FLIP","altComponent":"GEOTRACE","p
riority":540,"referenceValue":"FLIP","reporter":"AUTSDF1","node":"A6PV4","originat
ingDate":"20210222","originatingTime":"10:56:20","stringDate":"Mon Feb 22 2021
10:56:20"},{"system":"GVC2","message":"GEO2611I DASD MIRRORING IS ACTIVE ON LEG
LOCAL.RL1
","date":1613991396000,"color":"G","component":"RCA6PV4","altComponent":"GEOS2RC",
"priority":540,"referenceValue":"GVC2$CKDRL1M","reporter":"AUTSDF1","node":"A6PV4"
,"originatingDate":"20210222","originatingTime":"10:56:36","stringDate":"Mon Feb
22 2021 10:56:36"},{"system":"GVC2","message":"LCP=CAPTURE PROFILE(SGC1GOLD)
STARTED                                      /2
","date":1613991380000,"color":"G","component":"FLIP","altComponent":"GEOTRACE","p
riority":650,"referenceValue":"FLIP","reporter":"AUTSDF1","node":"A6PV4","originat
ingDate":"20210222","originatingTime":"10:56:20","stringDate":"Mon Feb 22 2021
10:56:20"},{"system":"GVC2","message":"SCHEDULING LCP CAPTURE FOR MANAGEMENT
PROFILE SGC1GOLD

","date":1613991380000,"color":"G","component":"FLIP","altComponent":"GEOTRACE","p
riority":650,"referenceValue":"FLIP","reporter":"AUTSDF1","node":"A6PV4","originat
ingDate":"20210222","originatingTime":"10:56:20","stringDate":"Mon Feb 22 2021
10:56:20"},{"system":"GVC2","message":"SEQUENCE NUMBER 60337FC4 HAS BEEN GENERATED
FOR THIS SAFEGUARD CAPTURE
","date":1613991380000,"color":"G","component":"FLIP","altComponent":"GEOTRACE","p
riority":650,"referenceValue":"FLIP","reporter":"AUTSDF1","node":"A6PV4","originat
ingDate":"20210222","originatingTime":"10:56:20","stringDate":"Mon Feb 22 2021
10:56:20"},{"system":"GVC2","message":"GEO2772I SAFEGUARD CAPTURE PHASE 1
RESERVATION STARTED
","date":1613991380000,"color":"G","component":"FLIP","altComponent":"GEOTRACE","p
riority":650,"referenceValue":"FLIP","reporter":"AUTSDF1","node":"A6PV4","originat
ingDate":"20210222","originatingTime":"10:56:20","stringDate":"Mon Feb 22 2021
10:56:20"},{"system":"GVC2","message":"GEO2773I SAFEGUARD CAPTURE PHASE 1
RESERVATION ENDED SUCCESSFULLY
","date":1613991381000,"color":"G","component":"FLIP","altComponent":"GEOTRACE","p
riority":650,"referenceValue":"FLIP","reporter":"AUTSDF1","node":"A6PV4","originat
ingDate":"20210222","originatingTime":"10:56:21","stringDate":"Mon Feb 22 2021
10:56:21"},{"system":"GVC2","message":"GEO2772I SAFEGUARD CAPTURE PHASE 2
RESERVATION SCAN STARTED
","date":1613991381000,"color":"G","component":"FLIP","altComponent":"GEOTRACE","p
riority":650,"referenceValue":"FLIP","reporter":"AUTSDF1","node":"A6PV4","originat
ingDate":"20210222","originatingTime":"10:56:21","stringDate":"Mon Feb 22 2021
10:56:21"},{"system":"GVC2","message":"GEO2773I SAFEGUARD CAPTURE PHASE 2
RESERVATION SCAN ENDED SUCCESSFULLY
","date":1613991386000,"color":"G","component":"FLIP","altComponent":"GEOTRACE","p
riority":650,"referenceValue":"FLIP","reporter":"AUTSDF1","node":"A6PV4","originat
ingDate":"20210222","originatingTime":"10:56:26","stringDate":"Mon Feb 22 2021
10:56:26"},{"system":"GVC2","message":"GEO2772I SAFEGUARD CAPTURE PHASE 3 CHECKIN
STARTED
","date":1613991386000,"color":"G","component":"FLIP","altComponent":"GEOTRACE","p
riority":650,"referenceValue":"FLIP","reporter":"AUTSDF1","node":"A6PV4","originat
ingDate":"20210222","originatingTime":"10:56:26","stringDate":"Mon Feb 22 2021
10:56:26"},{"system":"GVC2","message":"THE USER IMPACT TIME (UIT) FOR THIS
SAFEGUARD CAPTURE WAS 0.155 SECONDS
","date":1613991387000,"color":"G","component":"FLIP","altComponent":"GEOTRACE","p
riority":650,"referenceValue":"FLIP","reporter":"AUTSDF1","node":"A6PV4","originat
ingDate":"20210222","originatingTime":"10:56:27","stringDate":"Mon Feb 22 2021
10:56:27"},{"system":"GVC2","message":"GEO2773I SAFEGUARD CAPTURE PHASE 3 CHECKIN
ENDED SUCCESSFULLY
","date":1613991387000,"color":"G","component":"FLIP","altComponent":"GEOTRACE","p
riority":650,"referenceValue":"FLIP","reporter":"AUTSDF1","node":"A6PV4","originat
ingDate":"20210222","originatingTime":"10:56:27","stringDate":"Mon Feb 22 2021
10:56:27"},{"system":"GVC2","message":"GEO2775I LCP SAFEGUARD CAPTURE ENDED
SUCCESSFULLY
","date":1613991387000,"color":"G","component":"FLIP","altComponent":"GEOTRACE","p
riority":650,"referenceValue":"FLIP","reporter":"AUTSDF1","node":"A6PV4","originat
ingDate":"20210222","originatingTime":"10:56:27","stringDate":"Mon Feb 22 2021
10:56:27"},{"system":"GVC2","message":"LCP=CAPTURE PROFILE(SGC1GOLD) ENDED RC=0
/2
","date":1613991387000,"color":"G","component":"FLIP","altComponent":"GEOTRACE","p
riority":650,"referenceValue":"FLIP","reporter":"AUTSDF1","node":"A6PV4","originat
ingDate":"20210222","originatingTime":"10:56:27","stringDate":"Mon Feb 22 2021
10:56:27"},{"system":"GVC2","message":"LCP_SGC1GOLD_CAPTURE PLANNED/STANDARD
ACTION ENDED
```

Figure 3-21 SDF trace entries that are retrieved by the RESTful API service after capturing a Safeguarded Copy

The SDF trace entries that are retrieved can be parsed by a program that can display them, evaluate the execution return code, and eventually register the capture token.

Several JSON parsers exist in z/OS. For more information about the z/OS client web enablement toolkit, see z/OS JSON parser.

In our IBM Z Cyber Vault environment, we decided to take a Safeguarded Copy capture every 4 hours and one direct FlashCopy from RS2 to RC1 every day, and run the validation process on the RC1 set of volumes.

> **Important:** When you recycle GDPS Tivoli NetView for z/OS, you must also recycle the GDPS GUI server address space. Otherwise, a **bash** shell returns RC=0, which results in an operation not being performed.

3.4.2 Using digital certificates to authenticate to the GDPS RESTful API

A more secure way to authenticate the RESTful API request is to use digital certificates instead of a BASE64-encoded user and password.

To use digital certificates to authenticate to the GDPS RESTful API, complete the following steps:

1. Create a user SSL certificate (see Figure 3-22).

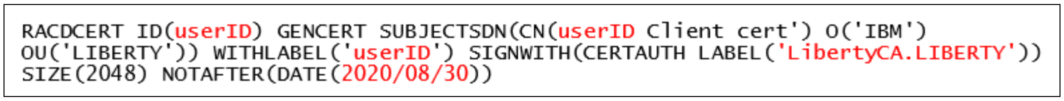

```
RACDCERT ID(userID) GENCERT SUBJECTSDN(CN(userID Client cert') O('IBM')
OU('LIBERTY')) WITHLABEL('userID') SIGNWITH(CERTAUTH LABEL('LibertyCA.LIBERTY'))
SIZE(2048) NOTAFTER(DATE(2020/08/30))
```

Figure 3-22 Commands to create a user SSL certificate

Customize the values in red, where:

- userID is the user that logs in to the RESTful API.
- LibertyCA.LIBERTY is the certificate authority (CA) certificate that is used to sign the GDPS Liberty server certificate.

> **Note:** Make sure that the date that is specified is within the range of the CA certificate's validity.

2. Export the client certificate (RACDCERT) to a data set.

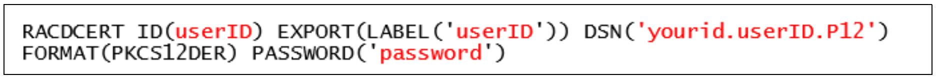

```
RACDCERT ID(userID) EXPORT(LABEL('userID')) DSN('yourid.userID.P12')
FORMAT(PKCS12DER) PASSWORD('password')
```

Figure 3-23 Exporting the client certificate

As shown in Figure 3-23:

- Set the userID to your user ID, which is used as a HLQ for the data set that contains the exported certificate.
- Set the password, which is used to encrypt the client certificate. This password is needed when you import the client certificate into the workstation.

> **Note**: Do not send this password in the same email as the certificate itself. The password is case-sensitive.

3. Export the CA certificate to a data set (see Figure 3-24).

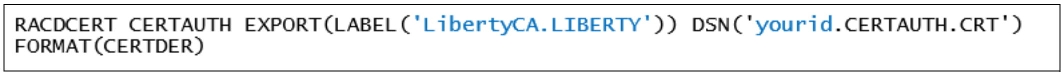

```
RACDCERT CERTAUTH EXPORT(LABEL('LibertyCA.LIBERTY')) DSN('yourid.CERTAUTH.CRT')
FORMAT(CERTDER)
```

Figure 3-24 Exporting the CA certificate

4. Download and import the client and CA certificates into the GDPS Logical Corruption Protection Manager instance that will use the API:
   - FTP or copy the client certificate (`yourid. userID.P12`) as a binary file to your system.
   - FTP or copy the CA certificate (`yourid.CERTAUTH.CRT`) as a binary file to your system.
   - Import the certificates into your system.

You can now use your P12 certificate in your RESTful API request to authenticate to the GDPS Liberty server (see Figure 3-25).

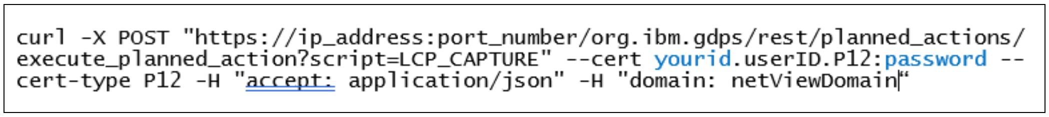

```
curl -X POST "https://ip_address:port_number/org.ibm.gdps/rest/planned_actions/
execute_planned_action?script=LCP_CAPTURE" --cert yourid.userID.P12:password --
cert-type P12 -H "accept: application/json" -H "domain: netViewDomain“
```

Figure 3-25 RESTful API cURL request that uses digital certificates instead of a BASE64-encoded user:password

To verify the RC and results of the SCRIPT execution and eventually send WTO and SDF trace entries in case of failure, we prepared a set of REXX and customization files that are provided *as is*.

3.4.3 Automating the process

To automate this process without checks, create the following procedures in `SYS1.PROCLIB`:

- ZCVRELBJ for the **`RELEASE`** command
- ZCVCAPBJ for the CAPTURE operations
- ZCVACTBJ to take a direct FlashCopy on RC1
- ZCVVALBJ to start the validation process (see Figure 3-26)

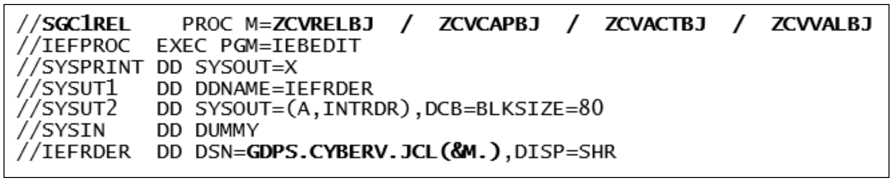

```
//SGC1REL    PROC M=ZCVRELBJ  /  ZCVCAPBJ  /  ZCVACTBJ  /  ZCVVALBJ
//IEFPROC   EXEC PGM=IEBEDIT
//SYSPRINT DD SYSOUT=X
//SYSUT1    DD DDNAME=IEFRDER
//SYSUT2    DD SYSOUT=(A,INTRDR),DCB=BLKSIZE=80
//SYSIN     DD DUMMY
//IEFRDER   DD DSN=GDPS.CYBERV.JCL(&M.),DISP=SHR
```

Figure 3-26 SGC1REL, SGC1CAP, ZCVACTV, and ZCVVALD procs in SYS1.PROCLIB

The members ZCVRELBJ, ZCVCAPBJ, ZCVACTBJ, and ZCVVALBJ in the IEFRDER library are in RESTful API JCL. They point to the bash scripts that start the following specific scripts:

- **`LCP_SGC1GOLD_RELE`** to release all expired Safeguarded Copy versions
- **`LCP_SGC1GOLD_CAP`** to capture a new Safeguarded Copy
- **`ZCV_RESET`** to activate the IBM Z Cyber Vault environment
- **`ZCV_VALIDATION`** to take the FlashCopy RS2 to RC1 and start the validation process

For example, the **START SGC1CAP** command from the console or syslog starts the GDPS script to capture a new Safeguarded Copy.

You can automate the **START** command in one of the following ways:

- By setting Job Entry Subsystem 2 (JES2) Automatic Commands
- BY setting System Automation TIMERS
- By adding them to a job scheduler like Tivoli Workload Scheduler

Using a scheduler implies starting a tracker on the GDPS Logical Corruption Protection Manager. We describe the two preferred methods of using a schedule: JES2 or TIMERS.

Setting up a release/capture process by using JES2 Automatic Commands

The commands that are shown in Figure 3-27 can be run and added to your JES2PARM.

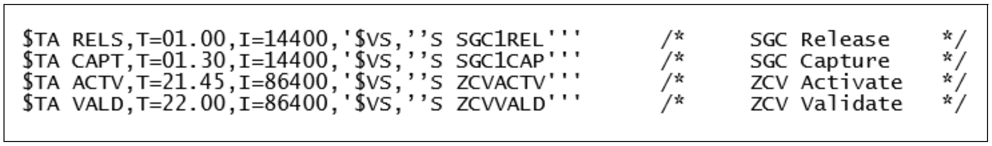

```
$TA RELS,T=01.00,I=14400,'$VS,''S SGC1REL'''      /*     SGC Release    */
$TA CAPT,T=01.30,I=14400,'$VS,''S SGC1CAP'''      /*     SGC Capture    */
$TA ACTV,T=21.45,I=86400,'$VS,''S ZCVACTV'''      /*     ZCV Activate   */
$TA VALD,T=22.00,I=86400,'$VS,''S ZCVVALD'''      /*     ZCV Validate   */
```

Figure 3-27 Setting JES2 Automatic Commands

You can display the active JES2 Automatic Commands by issuing the **$TA,ALL** command, as shown in Figure 3-28.

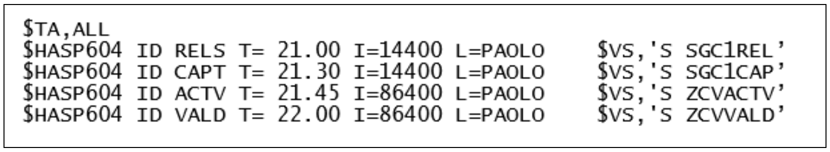

```
$TA,ALL
$HASP604 ID RELS T= 21.00 I=14400 L=PAOLO    $VS,'S SGC1REL'
$HASP604 ID CAPT T= 21.30 I=14400 L=PAOLO    $VS,'S SGC1CAP'
$HASP604 ID ACTV T= 21.45 I=86400 L=PAOLO    $VS,'S ZCVACTV'
$HASP604 ID VALD T= 22.00 I=86400 L=PAOLO    $VS,'S ZCVVALD'
```

Figure 3-28 Displaying JES2 Automatic Commands

You can cancel the active JES2 Automatic Commands either singularly or all together by issuing the **$CA** command, as shown in Figure 3-29.

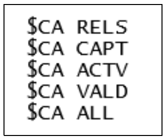

```
$CA RELS
$CA CAPT
$CA ACTV
$CA VALD
$CA ALL
```

Figure 3-29 Canceling JES2 Automatic Commands $CA RELS

Setting up a capture/release process by using SA TIMERS

TIMERS can be set in System Automation by adding them to the CNMSTYLE, as shown in Figure 3-30.

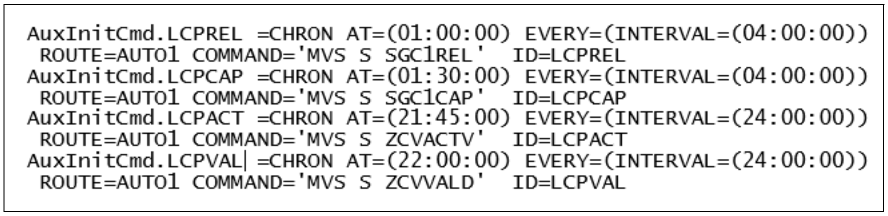

```
AuxInitCmd.LCPREL =CHRON AT=(01:00:00) EVERY=(INTERVAL=(04:00:00))
 ROUTE=AUTO1 COMMAND='MVS S SGC1REL'  ID=LCPREL
AuxInitCmd.LCPCAP =CHRON AT=(01:30:00) EVERY=(INTERVAL=(04:00:00))
 ROUTE=AUTO1 COMMAND='MVS S SGC1CAP'  ID=LCPCAP
AuxInitCmd.LCPACT =CHRON AT=(21:45:00) EVERY=(INTERVAL=(24:00:00))
 ROUTE=AUTO1 COMMAND='MVS S ZCVACTV'  ID=LCPACT
AuxInitCmd.LCPVAL| =CHRON AT=(22:00:00) EVERY=(INTERVAL=(24:00:00))
 ROUTE=AUTO1 COMMAND='MVS S ZCVVALD'  ID=LCPVAL
```

Figure 3-30 Setting TIMERS in CNMSTYLE

The timers can be displayed, modified, or deleted in Tivoli NetView for z/OS by using the **TIMERS** command, as shown in Figure 3-31.

:

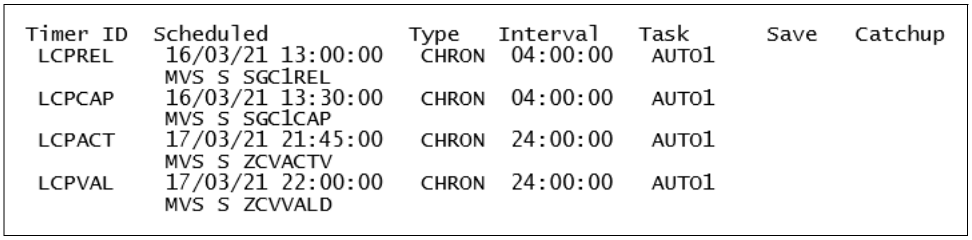

```
Timer ID  Scheduled           Type   Interval   Task      Save   Catchup
 LCPREL     16/03/21 13:00:00   CHRON   04:00:00   AUTO1
            MVS S SGC1REL
 LCPCAP     16/03/21 13:30:00   CHRON   04:00:00   AUTO1
            MVS S SGC1CAP
 LCPACT     17/03/21 21:45:00   CHRON   24:00:00   AUTO1
            MVS S ZCVACTV
 LCPVAL     17/03/21 22:00:00   CHRON   24:00:00   AUTO1
            MVS S ZCVVALD
```

Figure 3-31 Viewing TIMERS in Tivoli NetView for z/OS

The timers can be deleted by using the **PURGE** command, as shown in Figure 3-32.

.

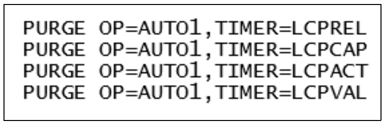

```
PURGE OP=AUTO1,TIMER=LCPREL
PURGE OP=AUTO1,TIMER=LCPCAP
PURGE OP=AUTO1,TIMER=LCPACT
PURGE OP=AUTO1,TIMER=LCPVAL
```

Figure 3-32 Deleting TIMERS by using the PURGE command

For more information about GDPS system recovery offerings for the IBM Z platform and specific REXX scripts to perform more validation checks, contact the IBM GDPS team at `mailto:gdps@us.ibm.com`.

3.5 Preparing the IBM Z Cyber Vault environment for validation

This section describes how to set up the environment and check whether the **IPL LOAD** command to the IBM Z Hardware Management Console (HMC) has started and successfully processed.

To prepare for Type 1 validation, activate the necessary hardware resources and perform the IPL of the systems that are needed for the Type 2 and 3 validation, as described in Chapter 4, “Establishing a validation framework” on page 103.

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

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

One requirement for an IBM Z Cyber Vault solution is creating LPARs for regular testing of the ability to perform an IPL from a recovered copy of your production system. To accomplish this testing, an isolated environment is built to perform the IBM Z Cyber Vault IPL and validation.

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

The new CFs and system LPARs for the IBM Z Cyber Vault environment must be defined in the *Geoplex Domains* of the GDPS Metro, which results in a Site Table, as shown in Figure 3-33 on page 93.

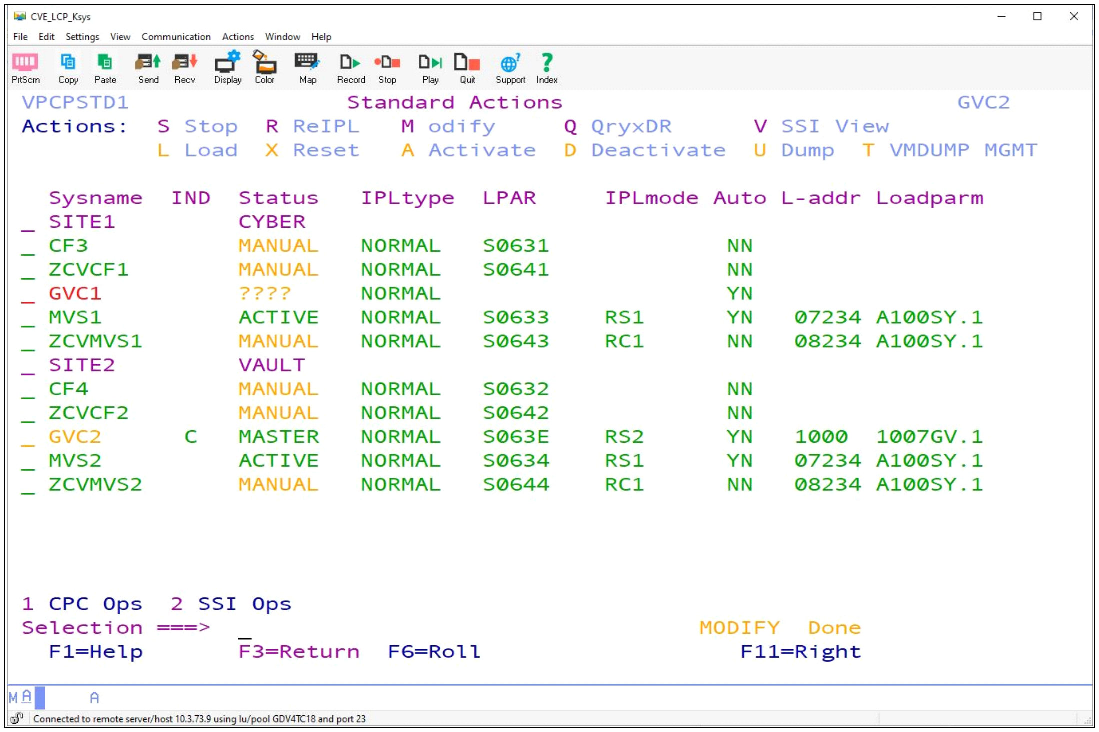

```
VPCPSTD1                Standard Actions                             GVC2
Actions:  S Stop  R ReIPL   M odify     Q QryxDR      V SSI View
          L Load  X Reset   A Activate  D Deactivate  U Dump  T VMDUMP MGMT

  Sysname  IND  Status   IPLtype  LPAR     IPLmode Auto L-addr Loadparm
_ SITE1         CYBER
_ CF3           MANUAL   NORMAL   S0631             NN
_ ZCVCF1        MANUAL   NORMAL   S0641             NN
_ GVC1          ????     NORMAL                     YN
_ MVS1          ACTIVE   NORMAL   S0633    RS1      YN   07234 A100SY.1
_ ZCVMVS1       MANUAL   NORMAL   S0643    RC1      NN   08234 A100SY.1
_ SITE2         VAULT
_ CF4           MANUAL   NORMAL   S0632             NN
_ ZCVCF2        MANUAL   NORMAL   S0642             NN
_ GVC2      C   MASTER   NORMAL   S063E    RS2      YN   1000  1007GV.1
_ MVS2          ACTIVE   NORMAL   S0634    RS1      YN   07234 A100SY.1
_ ZCVMVS2       MANUAL   NORMAL   S0644    RC1      NN   08234 A100SY.1



1 CPC Ops  2 SSI Ops
Selection ===>  _                                 MODIFY  Done
  F1=Help       F3=Return  F6=Roll                   F11=Right
```

```
MA      A
Connected to remote server/host 10.3.73.9 using lu/pool GDV4TC18 and port 23
```

Figure 3-33   GDPS Site Table

In the Site Table that is shown in Figure 3-33, note the following items:

- CF3 and CF4 are the CFs that are used by the production application sysplex.
- ZCVCF1 and ZCVCF2 are the CFs that are used in the IBM Z Cyber Vault recovered sysplex.
- MVS1 and MVS2 are the production application systems.
- ZCVMVS1 and ZCVMVS2 are systems names that are defined to GDPS to send Base Control Program internal interface (BCPii) commands to the LPAR to *activate* them, and perform IPLs of the recovered systems in the IBM Z Cyber Vault environment. The recovered systems must be defined with the `NN` automation parameter, like a foreign system.
- GVC2 is the Master GDPS controlling system.
- GVC1 is a dummy Alternate GDPS controlling system.

  In GDPS Metro, as a best practice, use two controlling systems for the best HA. In the context of LCP, you need only one controlling system, so for the sake of simplicity, we defined only one.

The two recovered systems, MVS1 and MVS2, can have an IPL from either RC1 or RC2, with two different OSConfigs that contain the two sets of volumes, which are defined in MSS0 online at IPL.

In the Site Table for the recovered systems, you can prepare both IPL parameters, as shown in Figure 3-34.

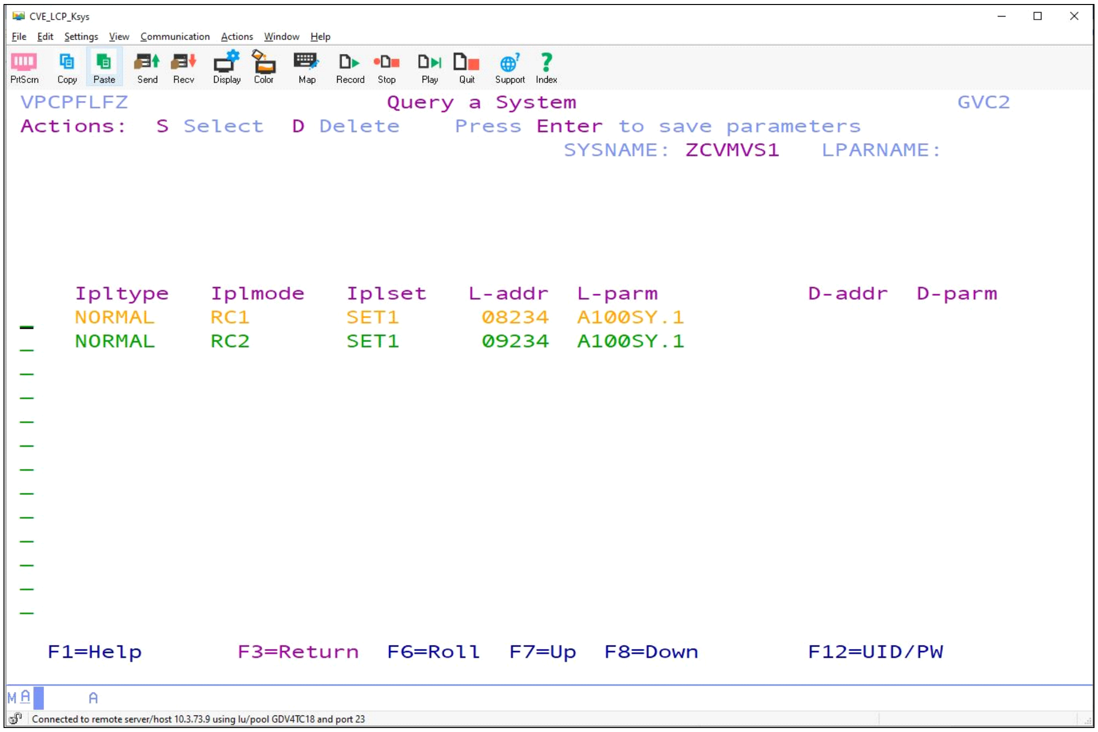
CVE_LCP_Ksys
File Edit Settings View Communication Actions Window Help
PrtScrn Copy Paste Send Recv Display Color Map Record Stop Play Quit Support Index
M A A
Connected to remote server/host 10.3.73.9 using lu/pool GDV4TC18 and port 23

```
VPCPFLFZ                    Query a System                           GVC2
Actions:  S Select  D Delete    Press Enter to save parameters
                                        SYSNAME: ZCVMVS1   LPARNAME:



    Ipltype    Iplmode    Iplset    L-addr   L-parm            D-addr  D-parm
    NORMAL     RC1        SET1       08234   A100SY.1
_   NORMAL     RC2        SET1       09234   A100SY.1
_
_
_
_
_
_
_
_
_
_
_

  F1=Help       F3=Return  F6=Roll   F7=Up  F8=Down         F12=UID/PW
```

Figure 3-34 GDPS Site Table IPL parameters for the systems to be started in the IBM Z Cyber Vault environment

IBM Z Cyber Vault MVS1 is started initially. IBM Z Cyber Vault MVS2 if started only if needed.

The two systems are supposed to be reset at the end of the previous validation process. Otherwise, they must be reset by adding a **sysplex reset** command at the beginning of the **Z Cyber Vault_reset** script.

The **IBM Z Cyber Vault_RESET** and **IBM Z Cyber Vault_VALIDATION** scripts, as shown in Example 3-1, are used to perform an IPL of the recovery system off RC1 and start a validation process. The scripts reset the actual recovery system, refresh the RC1 set of volumes, and perform an IPL of systems off RC1 to start the validation process.

Example 3-1 Scripts to start the validation process

```
Script name ZCV_RESET
1     COMM=RESET CYBER VAULT ENVIRONMENT
2     SYSPLEX=ACTIVATE SYSTEM(ZCVCF1,ZCVCF2)
3     SYSPLEX=ACTIVATE SYSTEM(ZCVMVS1,ZCVMVS2)
4     MESSAGE=ZCV_RESET ENDED

Script name ZCV_VALIDATION
1    COMM=START CYBER VALIDATION
2    DASD=FCWITHDRAW SECONDARY RC(1)
3    DASD=FCESTABLISH SECONDARY RC(1) NOCOPY CONSISTENT
4    SYSPLEX=LOAD SYSTEM(ZCVMVS1)
5    MONITOR=QHADISABLE FCSEC LEG(RL1)
6    MESSAGE=ZCV_VALIDATION ENDED
```

Note: When the systems are started off RC1, in GDPS the following SDF alert is displayed:

```
GEO274OW RS2 FC DEVICES 08000-08029 HAVE UNEXPECTED HOST ACCESS BY LPARS
```

You can avoid the warning by using the **monitor qhadisable** command. When the validation process is complete and the system resets, you can enable Query Host Access (QHA) by using the **monitor qhaenable** command.

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

Automating the validation process once per day can be achieved by adding the two last timers that are shown in Figure 3-30 on page 91.

At this point, the initiation of the IPL of the IBM Z Cyber Vault environment is complete.

3.5.2 Network isolation considerations for the IBM Z Cyber Vault environment

The recovered sysplex in the IBM Z Cyber Vault environment for validation must run in an isolated network from the application sysplex and be accessible by using different IP addresses.

In our environment, different Open Systems Adapter (OSA) cards and virtual LANs (VLANs) are used for the application sysplex and the IBM Z Cyber Vault environment.

3.5.3 Security access considerations for the IBM Z Cyber Vault environment

The IBM Z Cyber Vault environment should be secured to limit access to the environment and to perform all its tasks with automation.

From the hardware point of view, the IBM Z Cyber Vault LPARs must be configured in the HMC to allow the automation of the IPL.

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

From a z/OS point of view, the RACF database is like the production database. All users have the same rights. Access to the IBM Z Cyber Vault environment is limited to users who have access to the isolated VLAN setup for this environment.

The user ID for the data validation jobs must be set up to run in the production environment so that it is ready when an IPL of it is performed on the IBM Z Cyber Vault environment. In our configuration, a user ID is set up to start the data validation jobs that have the required access.

Also, if a specific user was created for running the data validation jobs, make sure that all user rights are configured. This user might need access to the `CSFKEY` resource profile, Db2 SYSADM, or other specific rights.

3.5.4 IODF requirements for the IBM Z Cyber Vault environment

To maintain the information that is gathered from the three phases of validation, you must define a minimal number of volumes that persist in the IBM Z Cyber Vault environment across IPLs. Define a unique esoteric for these IBM Z Cyber Vault persistent volumes to avoid having to hardcode the VOLSERs to create the output files.

In our environment, a pool of four 3390 Mod.27 volumes (IBM Z Cyber Vault 001 - IBM Z Cyber Vault 004) in Site 2 is defined as ZCV001-ZCV004 in MSS0 for both the application sysplex and IBM Z Cyber Vault OSconfigs. Therefore, the volumes can be varied online and the information made available in production.

In an asynchronous solution, these persistent volumes must be GC copied back to the production site.

3.5.5 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 will have an IPL performed on it has the same system definitions as the main recovered LPAR, some changes must be performed on the LOADPARM.

Based on the OSConfig configuration on the IODF, another OSGroup ID is used to allow the Cyber Vault LPARs to undergo an IPL with the RC1 or RC2 volumes.

The **IEASYMxx** parameter should be changed so that changes can be made in the system symbol definitions, such as a CYBERID symbol.

Example 3-2 shows the LOADPARM of one LPAR of the production system and the IBM Z Cyber Vault environment.

Example 3-2 LOADPARM Member example of one production system and one IBM Z Cyber Vault LPAR

```
****************************************************************** CYBER VAULT
Redbooks - PROD SYSPLEX LPAR 1 - MVS1
*****************************************************************
LPARNAME S0633
IODF     **          CYBER
NUCLEUS  1
NUCLST   06
INITSQA  0512K 0002M
SYSCAT   S1LOC1113CSYS1.MCAT.VS1CAT1
IEASYM   00
SYSPLEX  PORTINS1
PARMLIB  SYSM.PARMLIB.PORTINS1                        S1LOC1
PARMLIB  SYSM.PARMLIB.MVS1                            S1LOC1
PARMLIB  SYSM.PARMLIB                                 S1LOC1
PARMLIB  SYS1.IBM.PARMLIB                             ******
****************************************************************** CYBER VAULT
Redbooks - CYBER VAULT ENV. LPAR 1 - MVS1
*****************************************************************
LPARNAME S0643
IODF     **          CYBER1
NUCLEUS  1
NUCLST   06
```

```
INITSQA  0512K 0002M
SYSCAT   S1LOC1113CSYS1.MCAT.VS1CAT1
IEASYM   01
SYSPLEX  PORTINS1
PARMLIB  SYSM.PARMLIB.PORTINS1                        S1LOC1
PARMLIB  SYSM.PARMLIB.MVS1                            S1LOC1
PARMLIB  SYSM.PARMLIB                                 S1LOC1
PARMLIB  SYS1.IBM.PARMLIB                             ******
```

3.5.6 PARMLIB changes

In the PARMLIB, as a best practice, use system symbols for simplifying the deployment of the IBM Z Cyber Vault environment.

In our environment, as shown on the LOADPARM, two **IEASYMxx** are used: one for the production system and one for the IBM Z Cyber Vault environment.

The &CYBERID symbol is created to match our residency environment configuration. It is defined as follows:

- 0 for production system (in IEASYM00)
- R for Cyber Vault environment (in IEASYM01)

Two differences exist between the LPAR name and SYSPARM IEASYMxx members. The IBM Z Cyber Vault environment uses SYSPARM=01, which allows two changes to the PARMLIB members that are used during IPL:

- Some lines might need to be added to the COMMNDxx member to start the validation when the system is ready. In our environment, most of the tasks that are required for the IBM Z Cyber Vault environment are started by using the COMMND01 member, as shown in Example 3-3.

Example 3-3 IBM Z Cyber Vault environment tasks

```
COM='T MPF=CL'
COM='S LLA,LLA=00,SUB=MSTR'
COM='S VLF,NN=00,SUB=MSTR'
COM='S RRS,SUB=MSTR'
COM='SET SLIP=00'
COM='SET DAE=00'
COM='CD SET,SDUMP=(ALLPSA,NUC,SQA,LSQA,RGN,LPA,TRT,SWA,CSA,SUM),Q=YES'
COM='CD SET,SDUMP,MAXSPACE=3000M'
COM='DD ADD,VOL=&VOLDUMP.'
COM='DD NAME=SYSM.DUMP.&SYSNAME..D&YYMMDD..T&LHHMMSS..S&SEQ'
COM='DD ALLOC=ACTIVE'
COM='S IRRDPTAB'
COM='S JES2,PARM=(WARM,NOREQ)'
COM='S NET,,,(LIST=00)'
COM='S RMF'
COM='SET PROG=EX'
COM='SET PROG=LP'
COM='SET PROG=WS'
COM='SET PROG=CO'
COM='S SDSFSRV.SDSF'
COM='SETLOGRC LOGSTREAM'
COM='S AUTONET'
```

```
COM='S AUTOLIN'
COM='S AUTOSTC'
COM='S APPC,SUB=MSTR'
COM='S ASCH,SUB=MSTR'
COM='S BPO1,SUB=MSTR'
COM='S ICEOPT,ICEPRM=01'
COM='S AUTOSTCV'
```

The final line in COMMND01 runs a procedure to start the Db2 regions, where DBX1 and DBX2 are in a data-sharing complex. Therefore, the second member, DBX2, is started with LIGHT(YES) so that it releases any held locks and then go back down. In the S4 step, the CV#MON IBM Z Cyber Vault validation-monitor task is started, as shown in Example 3-4. This task controls and monitors the validation process. For more information, see 4.1.7, "Creating an IBM Z Cyber Vault validation monitoring task" on page 109.

Example 3-4 STARTSTC Proc

```
//STARTSTC PROC
//SW1      EXEC PGM=STCWAIT,PARM=60
//STEPLIB DD DISP=SHR,DSN=SYSP.LINKLIB.AUTOSTR
//S1       EXEC PGM=COMMAND,PARM='-DBX1 START DB2'
//STEPLIB DD DISP=SHR,DSN=SYSP.LINKLIB.AUTOSTR
//SW2      EXEC PGM=STCWAIT,PARM=80
//STEPLIB DD DISP=SHR,DSN=SYSP.LINKLIB.AUTOSTR
//S2       EXEC PGM=COMMAND,PARM='-DBW1 START DB2'
//STEPLIB DD DISP=SHR,DSN=SYSP.LINKLIB.AUTOSTR
//SW3      EXEC PGM=STCWAIT,PARM=80
//STEPLIB DD DISP=SHR,DSN=SYSP.LINKLIB.AUTOSTR
//S3       EXEC PGM=COMMAND,PARM='-DBX2 START DB2 LIGHT(YES)'
//STEPLIB DD DISP=SHR,DSN=SYSP.LINKLIB.AUTOSTR
//SW4      EXEC PGM=STCWAIT,PARM=80
//STEPLIB DD DISP=SHR,DSN=SYSP.LINKLIB.AUTOSTR
//S4       EXEC PGM=COMMAND,PARM='S CV#MON'
//STEPLIB DD DISP=SHR,DSN=SYSP.LINKLIB.AUTOSTR
```

- The second PARMLIB change is the AUTO member. It should be changed to allow other modifications regarding the automation of the IPL. In our environment, the only addition concerns the use of CDSs. The warning displays that some CDSs already might be active in another sysplex. However, the IBM Z Cyber Vault environment does not work on production volumes, so these CDSs are not used. The following line is added to the AUTORxx member:

  ```
  Msgid(IXC247D)   Delay(05S) Reply(U)
  ```

JES2PARM considerations

The data-validation process can generate several job entries in the Job Entry Subsystem (JES), which can create a JES2 RESOURCE SHORTAGE issue. To avoid this type of issue, use specific configuration options in the JES2PARM, as shown in Example 3-5.

Example 3-5 JES2PARM configuration options

```
OUTCLASS(O) BLNKTRNC=YES,
         OUTDISP=(PURGE,KEEP),
         OUTPUT=PRINT,
         TRKCELL=YES
```

If MSGCLASS=0 is used in the data validation jobs, the output of jobs is automatically purged upon completion. With an abnormal completion, the output is kept in the joblog.

The COUPLExx member, which is used by both the production system and the IBM Z Cyber Vault environment, includes the new &CYBERID symbol. Except for the LOGR CDS, every CDS is the COUPLExx is defined with the previously created system symbol, as shown in Example 3-6.

Example 3-6 COUPLExx example with &CYBERID system symbol

```
COUPLE SYSPLEX(&SYSPLEX.)
        PCOUPLE(SYSM.XCF.&SYSPLEX..CDS&CYBERID.2)
        ACOUPLE(SYSM.XCF.&SYSPLEX..CDS&CYBERID.3)
        CFRMPOL(POL2)
        . . . /* Other XCF Definitions */
DATA TYPE(CFRM)
        PCOUPLE(SYSM.CFRM.&SYSPLEX..CDS&CYBERID.3)
        ACOUPLE(SYSM.CFRM.&SYSPLEX..CDS&CYBERID.2)
DATA TYPE(WLM)
        PCOUPLE(SYSM.WLM.&SYSPLEX..CDS&CYBERID.2)
        ACOUPLE(SYSM.WLM.&SYSPLEX..CDS&CYBERID.3)
DATA TYPE(SFM)
        PCOUPLE(SYSM.SFM.&SYSPLEX..CDS&CYBERID.3)
        ACOUPLE(SYSM.SFM.&SYSPLEX..CDS&CYBERID.2)
DATA TYPE(ARM)
        PCOUPLE(SYSM.ARM.&SYSPLEX..CDS&CYBERID.3)
        ACOUPLE(SYSM.ARM.&SYSPLEX..CDS&CYBERID.2)
DATA TYPE(LOGR)
       PCOUPLE(SYSM.LOGR.&SYSPLEX..CDS02)
       ACOUPLE(SYSM.LOGR.&SYSPLEX..CDS03)
DATA TYPE(BPXMCDS)
       PCOUPLE(SYSM.BPX.&SYSPLEX..CDS&CYBERID.2)
       ACOUPLE(SYSM.BPX.&SYSPLEX..CDS&CYBERID.3)
```

3.5.7 CFRM changes

The CFRM CDS, allocated with the &CYBERID system symbol for recovery, must be defined with the same production policy, and the PREFLIST of the structures must contain the CF in the IBM Z Cyber Vault environment. The hardware connectivity determines where the STRs are allocated: The CF of the IBM Z Cyber Vault environment must not be connected to the production LPARs and vice versa.

In our environment, the definition is done with four CFs: two for the production sysplex and two for the IBM Z Cyber Vault environment. All four CFs are defined in CFRM policy and in the correct order in the PREFLIST (the ordered list of CF, where preferably STRs is allocated), as shown in Example 3-7.

Example 3-7 IXCMIAPU SYSIN for the IBM Z Cyber Vault CFRM CDS

```
/*  DEFINE POLICY TYPE CFRM          */
      DATA TYPE(CFRM) REPORT(YES)
      DSN(SYSM.CFRM.PORTINS1.CDSR3)   /* Define to Dataset */
      VOLSER(S1CDR3)
        DEFINE POLICY NAME(POL2) REPLACE(YES)
       CF NAME(CF1)   /*  NAME OF CYBER VAULT CF          */
               TYPE(008561)
```

```
             MFG(IBM)
             PLANT(02)
             SEQUENCE(0000000xxxxx)
             PARTITION(41)
             CPCID(00)
             DUMPSPACE(20)

   CF NAME(CF2) /*  NAME OF CYBER VAULT CF        */
             TYPE(008561)
             MFG(IBM)
             PLANT(02)
             SEQUENCE(0000000xxxxx)
             PARTITION(42)
             CPCID(00)
             DUMPSPACE(20)

  CF NAME(CF3)  /* NAME OF PROD CF - No Path Available */
             TYPE(008561)
             MFG(IBM)
             PLANT(02)
             SEQUENCE(0000000xxxxx)
             PARTITION(31)
             CPCID(00)
             DUMPSPACE(20)

     CF NAME(CF4)  /*  NAME OF PROD CF - No Path Available */
             TYPE(008561)
             MFG(IBM)
             PLANT(02)
             SEQUENCE(0000000xxxxx)
             PARTITION(32)
             CPCID(00)
             DUMPSPACE(20)

/*  DEFINE STRUCTURE SIGNALING        */
        STRUCTURE NAME(IXC1_GRS)
             SIZE(20000)
             PREFLIST(CF1,CF2,CF3,CF4)  /* PREFLIST Example */
```

3.5.8 TCP/IP configuration changes

You can manage the TCP/IP and IBM VTAM® configuration in the production system to perform an IPL of the IBM Z Cyber Vault environment in parallel.

As an example, the system symbol `&CYBERID` is used in the TCP/IP procedure to select either the production system profile or the IBM Z Cyber Vault environment profile. With this method, you can have the two different profiles available and configurable. The same symbol is used with IPNODES to configure the resolver.

For VTAM, the IBM Z Cyber Vault `PORTNAME` environment and devices must be defined. In our environment, we define the IBM Z Cyber Vault information in the same technology readiness level (TRL) as the production system. Depending on which OSA is physically seen by the LPAR that is undergoing IPL, the correct devices will be accessible.

3.5.9 IBM Z Multi-Factor Authentication consideration

If multi-factor authentication (MFA) is used, a copy of the setup must be implemented in the IBM Z Cyber Vault environment.

3.5.10 Pervasive encryption considerations

If you implement IBM Z pervasive encryption in your environment, you must ensure that the following tasks are complete:

- Your master keys are loaded into the crypto-card domains for your IBM Z Cyber Vault LPAR.
- The rotation of your master keys in your production environment account for your IBM Z Cyber Vault environment.

3.5.11 Determining tasks to start on IBM Z Cyber Vault

For licensing purposes, in your IBM Z Cyber Vault environment, you might not want to start certain tasks that are not critical to data validation, forensic analysis, or surgical recovery. You might run certain tasks that you want to start on a single IBM Z Cyber Vault LPAR, or on different LPARs in your sysplex. Identifying required tasks (and any modifications to normal startup procedures) is important here.

In our environment, we started the required tasks by using our COMMND01 member, as shown in Example 3-3 on page 97.

3.5.12 Considerations for Db2 data sharing

The purpose of data sharing is that tables and all other objects in Db2 should be accessible if just *one* member of the data-sharing complex is started. Consider a sysplex with two application systems. If one of the systems is taken down for maintenance, Db2 should still be fully operational from the other members. This situation also applies to the IBM Z Cyber Vault environment. Although one member can perform a restart on behalf of the group, you should restart all the nonquiesced and nonstarting members, it is best to start all members of the group for maximum data availability.

For our environment, we have a Db2 data sharing complex that is composed of two members. We started one of the members with the option of `LIGHT(YES)` so that locks that are held at the point of the data capture for the Safeguarded Copy are released.

3.5.13 Automating your IBM Z Cyber Vault IPL procedures

Managing target system hardware usually means full control (from ACTIVATE and IPL to DEACTIVATE). Both IBM GDPS and IBM System Automation Processor Operations (SA ProcOps) both provide this capability.

SA ProcOps is the right choice if you do not have GDPS, but you want to fully manage and control all target system hardware across multiple CPCs and physical sysplexes.

If you are not running GDPS or SA ProcOps, check with your Automation Product vendor about capabilities to schedule and initiate an unattended automated IPL of a system.

At this point, we have a virtual airgap-protected environment with Safeguarded Copy backups of our production data that is taken at four-hour durations with a retention period of four days. Also, the RC1 volumes are ready to be used for the validation process that is described in Chapter 4, “Establishing a validation framework” on page 103.

4

Establishing a validation framework

Although data validation is a key element of the IBM Z Cyber Vault solution, data validation requirements, design, and capabilities are unique to every environment and need investment to build, customize, and maintain.

This chapter provides a suggested validation framework and process flow with advice for conducting basic data validation. There are samples that can be tailored to individual business needs, priorities, and IT configurations.

The three basic types of data validation that can be regularly performed in an IBM Z Cyber Vault environment include:

- Type 1 validation: Check that initial program load (IPL) is possible.
- Type 2 validation: Check the data structural integrity of key middleware or databases.
- Type 3 validation: Check that the application data that is stored in data sets or databases is valid.

For more information about the validation types, see 1.4.1, “Validation use cases” on page 11.

In the case of Type 1 validation, much of the automation that is needed to set up and perform an IPL on the IBM Z Cyber Vault environment is managed and tracked by the IBM Geographically Dispersed Parallel Sysplex (GDPS) K-sys. However, GDPS can only initiate the IPL; it cannot manage and track activities after the IPL because the IBM Z Cyber Vault recovery system is in another network-isolated sysplex. Therefore, suggestions are offered for ways that validation can be tracked, monitored, and reported on in an automated fashion within your IBM Z Cyber Vault environment.

The ways in which you define jobs or processes to be run for validation can be flexible to mechanize. You can use different vendor z/OS Automation and z/OS Batch Scheduling products. We created a generic process for validation independent of the specific vendor automation or batch-scheduling product. For the three types of validations, you can use non-vendor-dependent methods to perform validation and monitor the results.

This chapter also describes a method for ensuring the collection of feedback from the validation process and ways to automate it.

The following topics are covered in this chapter:

- Creating a validation framework
- Type 1 validation
- Type 2 validation
- Type 3 validation

4.1 Creating a validation framework

The basic assumptions that are used to design our framework for performing data validation are listed in this section. The framework is described in terms of naming conventions for libraries, jobs, and feedback data sets. In addition, we describe the creation of a monitoring task that controls the sequencing of validation jobs and detects the completion of those jobs so that the collection process of feedback from a validation run can be started.

> **Note:** IBM Z Cyber Vault validation is *not* a software product, but rather a series of jobs and processes that must be individually constructed for your environment. For this IBM Redbooks publication, sample code is built to accomplish the automation of this process.

4.1.1 Data validation assumptions

Consider the following IBM Z Cyber Vault data validation assumptions:

- 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.
- 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 Virtual Storage Access Method. (VSAM) files that are added or deleted) in the construction of data-validation jobs.
- Must identify and document errors that are encountered for IBM Z Cyber Vault validation jobs.
- Must allow for exclusion of certain conditions during data-validation checking. However, the excluded conditions must be documented.
- Must avoid any write-to-operator-with-reply (WTOR) operations from being generated by a validation job (such as a reference to a data set that is cataloged to an offline volume).
- Must have a mechanism for creating IBM Z Cyber Vault data structure validation output (summary and detail of errors) to send home. Also, you must have an auditable trail of all validation jobs that run.
- Must have a mechanism such that data-content validation jobs fit into the IBM Z Cyber Vault validation framework.
- Must have the necessary security requirements documented for an automated process to perform validation (that is, for access to the files to validate).

4.1.2 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 a data management product or OEM product.
- Data content validation: Refers to validation of the content of the 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 to construct these jobs.
- Product: Refers to the specific product in use, which might be:
  - Provided as a part of the base z/OS operating system that is provided by IBM.
  - Available as “open source” for z/OS (no charge, or optional fee-based service plans available) that is 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 the *data content*. Many products use published metadata schema 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 regarding the metadata that is used or maintained by that product. Sometimes, the terms “products” and “subsystems” are used interchangeably, such as Db2, IBM Information Management System (IMS), IBM Customer Information Control System (CICS). There might be several products that are installed at an installation: Some products are related specifically to management and enhanced functions of data management subsystems, and some are related to enhanced functions of general 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 used by user-supplied jobs and processes to 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 or data set, or a unique Db2 table space.
- Validation build library: Refers to a partitioned data set extended (PDSE) library that contains the processes that are needed to determine 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.
  - Basic job card to use for your environment.
  - Input files that are needed to define objects that you want to validate.
  - Job Control Language (JCL) processes to invoke for the following reasons:
    - Submitting a simple validation.
    - Submitting a more complex validation process that requires REXX (or other language) code to compose JCL in a runtime library and then submit the JCL.
  - JCL to invoke for monitoring more complex validation that requires a specific sequencing of validation for a product (for data-structure validation) or application (for data-content validation).
  - REXX code members invoked by JCL in the validation build library.

- Validation runtime library: Refers to a PDSE library that is newly built with each validation run. It contains JCL members that are built and submitted during the validation run.

4.1.3 Constructing data structure validation jobs

In our environment, several jobs or processes and the REXX code were constructed to enable the following items:

- Auto-build of validation jobs
- Sequencing of validation jobs when required
- Monitoring of the jobs to ensure completion
- Collecting validation information for reporting and history archival of information

Validation processes are built for the following six types of data structure validation activities:

- Catalog Integrated Data Cluster Access Method Services (IDCAMS) Diagnose and Examine activities
- Basic catalog structure - Virtual Storage Access Method Volume Data Set (BCS-VVDS) Examine activities
- VSAM key sequenced data set files (KSDS) Examine activities
- RACF database data structure validation
- Db2 database data structure validation for a Db2 V11 stand-alone subsystem
- Db2 database data structure validation for a Db2 V12 data sharing subsystem

4.1.4 IBM Z Cyber Vault validation framework

For our environment, 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 might be provided from a pre-determined input list, or alternatively, needed to be self-discovered at the time the validation is run.

Here are the three variations for building validation processes:

- Simple validation: Validation for a single object, based on the concept of running a single job. The job writes the output to a pre-determined error file name, and if the return code is successful, it 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 simple validation is the data-structure validation job that was built for the RACF database (see 4.1.6, “IBM Z Cyber Vault validation data set naming conventions” on page 108).
- Validation jobs based on the input of objects: In this type of validation, an input file of objects to be validated is used. A build process is constructed for individual jobs for each object. You should incorporate checks for the existence of the object because a validation job for a nonexistent object results in a JCL failure. Note the cases where a validation job for a particular object is not built (by checking whether the object is cataloged to an offline volume). An example of this type of validation is the data-structure validation process that we built for Catalog IDCAMS Diagnose and Examine activities. For objects that pass the check, concatenate all jobs to a single validation `runlib` member to be submitted at the conclusion of the build process.
- Validation jobs based on self-discovery of objects: In this type of validation, we built in self-discovery of objects to validate, constructed build-process individual jobs for each object, and incorporated checks for the existence of the object based on the object type.

For example, for the BCS-VVDS data-structure validation, complete the following steps:

a. Each catalog is interrogated for the existence of data set entries of 'SYS1.VVDS.*'.

b. Each catalog is checked to determine whether that entry is known to the master catalog:

   - If the entry is known, check whether the volume is online.
   - If the entry is not known to the master catalog, or a volume is detected offline, then an error message is produced.

c. Determine whether the object passes the initial checks:

   - If it does not pass the initial checks, then a validation job is not built (because it would result in a JCL failure). Note the cases where self-discovery of an object indicates an anomaly in the environment, which might indicate the necessity for clean-up activities.
   - If it passes the initial checks, concatenate all jobs to a validation runlib member that will be submitted at the conclusion of the build process.

If you are performing multiple types of validations for a product but they must be performed in a specific sequence, you must perform the following extra tasks:

- Define separate, unique job prefixes and validation runlib members for each sequence of jobs to run.
- Construct a validation monitoring job to ensure that each sequence is completed before triggering the next sequence.

4.1.5 IBM Z Cyber Vault job naming conventions

Validation jobs are assigned a unique three-character prefix. The first two characters are unique to IBM Z Cyber Vault jobs. Other started tasks in your IBM Z Cyber Vault environment cannot begin with these first two characters because all jobs beginning with these first two characters are for exclusive use by IBM Z Cyber Vault validation activities. This concept is important when it comes to monitoring the completion of the IBM Z Cyber Vault validation cycle because the main IBM Z Cyber Vault monitoring task CV#MON periodically examines the System Display and Search Facility (SDSF) DA panel for any running jobs with the 2-character prefix. In our environment, we chose "CV" to be the two character job prefix.

The third character denotes the type of validation being performed, which is documented in a library member in our build jobs PDSE, as shown in Example 4-1.

Example 4-1 Legend for job prefix for all validation jobs

```
*****************************************************************
* Legend for Job prefix for all validation jobs
*****************************************************************
Prefix   Description
-------- --------------------------------------------------------
CV@      Build jobs for validation functions
CV#      Monitoring jobs
CVA      Cyber Vault Information Collection
CVB      BCS-VVDS validation jobs
CVC      Catalog diagnose and examine jobs
CVD      VSAM KSDS examine jobs
CVE      RACF structure validation job
CVS      Db2 DBW1 system DSN1COPY jobs
CVT      Db2 DBW1 application DSN1COPY jobs
```

```
CVU      Db2 DBW1 system CHECK INDEX jobs
CVV      Db2 DBW1 application CHECK INDEX jobs
CVW      Db2 DBX1 system DSN1COPY jobs
CVX      Db2 DBX1 application DSN1COPY jobs
CVY      Db2 DBX1 system CHECK INDEX jobs
CVZ      Db2 DBX1 application CHECK INDEX jobs
*****************************************************************
```

In the first three job prefixes, the third character should be the value that is shown in Example 4-1 on page 107.

For the remaining job prefixes, the prefix is passed by the build JCL for a validation category. So, there can be up to 99,999 individual objects that are validated for that validation category. For Db2 validations, you need four validation job prefixes per Db2 SSID validated. In our case, we performed IBM Z Cyber Vault validations on two Db2 subsystems, so we set aside a total of eight job prefixes to account for the different categories of Db2 validation.

4.1.6 IBM Z Cyber Vault validation data set naming conventions

Important data set naming conventions are used for certain types of data sets that are required or used in the validation process. These data sets are written to the persistent volumes that were defined to the IBM Z Cyber Vault environment.

Figure 4-1 on page 109 shows the data-set naming conventions that we built, which are based on the process that is set up in our environment.

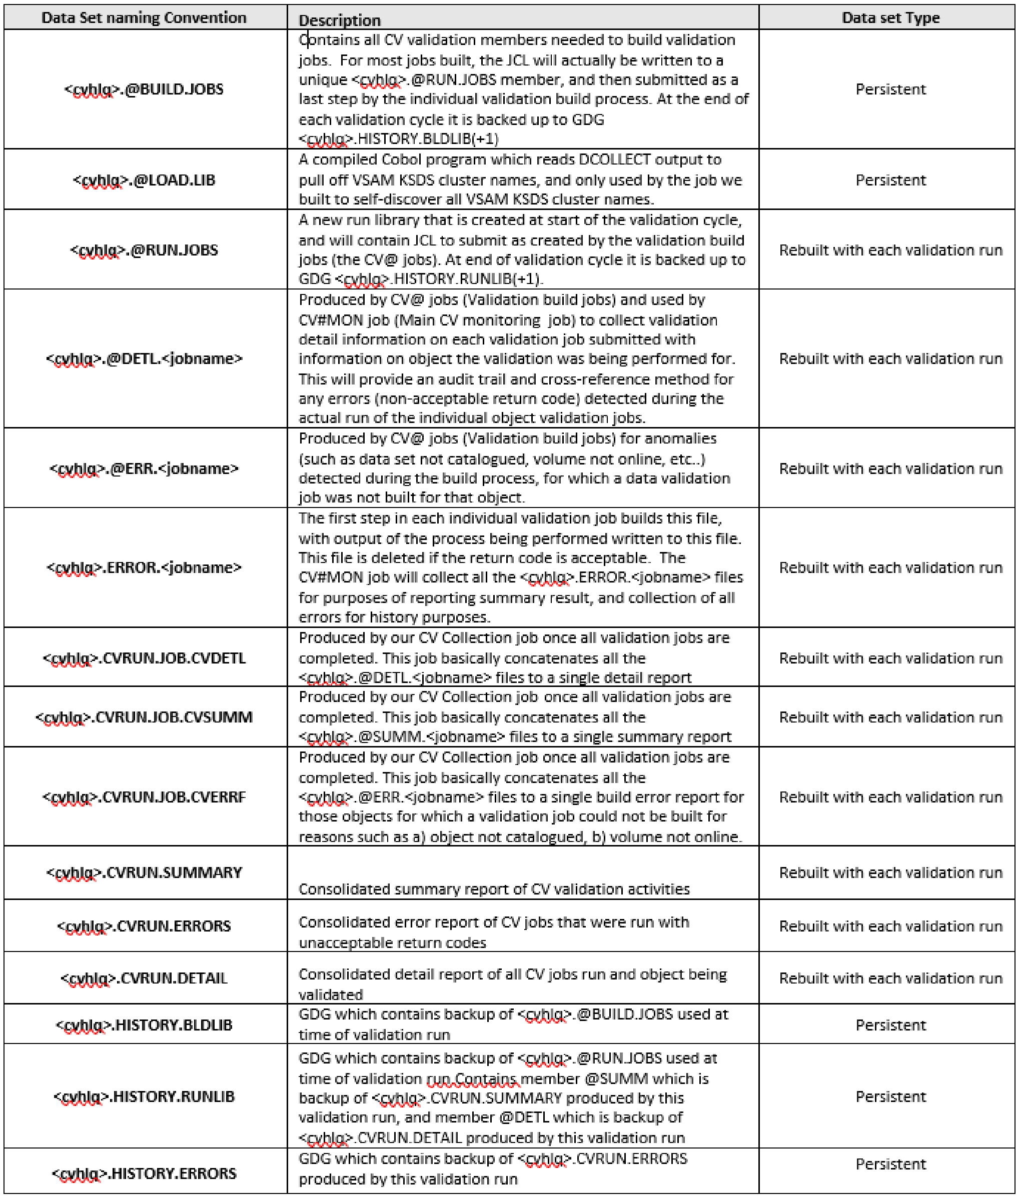

Data Set naming Convention	Description	Data set Type
<cvhlq>.@BUILD.JOBS	Contains all CV validation members needed to build validation jobs. For most jobs built, the JCL will actually be written to a unique <cvhlq>.@RUN.JOBS member, and then submitted as a last step by the individual validation build process. At the end of each validation cycle it is backed up to GDG <cvhlq>.HISTORY.BLDLIB(+1)	Persistent
<cvhlq>.@LOAD.LIB	A compiled Cobol program which reads DCOLLECT output to pull off VSAM KSDS cluster names, and only used by the job we built to self-discover all VSAM KSDS cluster names.	Persistent
<cvhlq>.@RUN.JOBS	A new run library that is created at start of the validation cycle, and will contain JCL to submit as created by the validation build jobs (the CV@ jobs). At end of validation cycle it is backed up to GDG <cvhlq>.HISTORY.RUNLIB(+1).	Rebuilt with each validation run
<cvhlq>.@DETL.<jobname>	Produced by CV@ jobs (Validation build jobs) and used by CV#MON job (Main CV monitoring job) to collect validation detail information on each validation job submitted with information on object the validation was being performed for. This will provide an audit trail and cross-reference method for any errors (non-acceptable return code) detected during the actual run of the individual object validation jobs.	Rebuilt with each validation run
<cvhlq>.@ERR.<jobname>	Produced by CV@ jobs (Validation build jobs) for anomalies (such as data set not catalogued, volume not online, etc..) detected during the build process, for which a data validation job was not built for that object.	Rebuilt with each validation run
<cvhlq>.ERROR.<jobname>	The first step in each individual validation job builds this file, with output of the process being performed written to this file. This file is deleted if the return code is acceptable. The CV#MON job will collect all the <cvhlq>.ERROR.<jobname> files for purposes of reporting summary result, and collection of all errors for history purposes.	Rebuilt with each validation run
<cvhlq>.CVRUN.JOB.CVDETL	Produced by our CV Collection job once all validation jobs are completed. This job basically concatenates all the <cvhlq>.@DETL.<jobname> files to a single detail report	Rebuilt with each validation run
<cvhlq>.CVRUN.JOB.CVSUMM	Produced by our CV Collection job once all validation jobs are completed. This job basically concatenates all the <cvhlq>.@SUMM.<jobname> files to a single summary report	Rebuilt with each validation run
<cvhlq>.CVRUN.JOB.CVERRF	Produced by our CV Collection job once all validation jobs are completed. This job basically concatenates all the <cvhlq>.@ERR.<jobname> files to a single build error report for those objects for which a validation job could not be built for reasons such as a) object not catalogued, b) volume not online.	Rebuilt with each validation run
<cvhlq>.CVRUN.SUMMARY	Consolidated summary report of CV validation activities	Rebuilt with each validation run
<cvhlq>.CVRUN.ERRORS	Consolidated error report of CV jobs that were run with unacceptable return codes	Rebuilt with each validation run
<cvhlq>.CVRUN.DETAIL	Consolidated detail report of all CV jobs run and object being validated	Rebuilt with each validation run
<cvhlq>.HISTORY.BLDLIB	GDG which contains backup of <cvhlq>.@BUILD.JOBS used at time of validation run	Persistent
<cvhlq>.HISTORY.RUNLIB	GDG which contains backup of <cvhlq>.@RUN.JOBS used at time of validation run.Contains member @SUMM which is backup of <cvhlq>.CVRUN.SUMMARY produced by this validation run, and member @DETL which is backup of <cvhlq>.CVRUN.DETAIL produced by this validation run	Persistent
<cvhlq>.HISTORY.ERRORS	GDG which contains backup of <cvhlq>.CVRUN.ERRORS produced by this validation run	Persistent

Figure 4-1   Data set descriptions

4.1.7 Creating an IBM Z Cyber Vault validation monitoring task

The key to creating an automated and unattended validation process is using 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. In our data structure validation, for some types of validation we used discovery steps to retrieve a list of objects and build the JCL by using the job-naming convention that is described in 4.1.5, “IBM Z Cyber Vault job naming conventions” on page 107. Then, you can batch-submit the constructed jobs.

Figure 4-2 shows the simple sequence of steps that are performed by this monitoring task. It is possible to build in dependencies such that Type 3 data-content validation can be submitted on completion of a subset of Type 2 data-structure validation jobs.

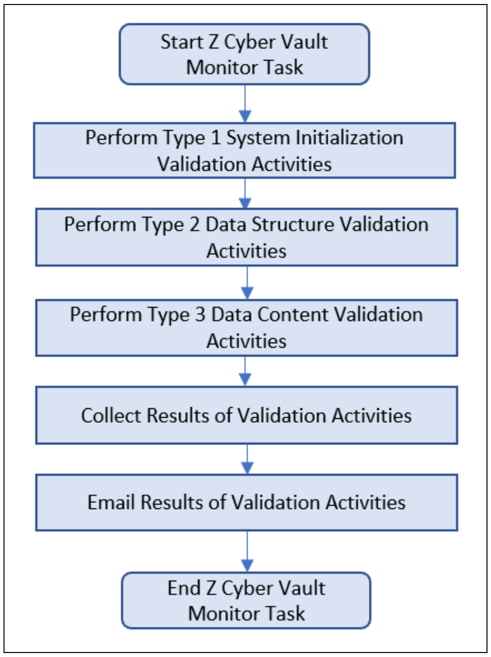
Start Z Cyber Vault Monitor Task
Perform Type 1 System Initialization Validation Activities
Perform Type 2 Data Structure Validation Activities
Perform Type 3 Data Content Validation Activities
Collect Results of Validation Activities
Email Results of Validation Activities
End Z Cyber Vault Monitor Task

Figure 4-2 Process flow for the SCV Monitor task

This process is written by using System REXX and the SDSF REXX API, which became available with IBM z/OS SDSF V1.9. By using the SDSF REXX API, you can query the output of tasks or jobs and display currently running jobs to determine when a specific type or phase of validation is complete. You can use a combination of the batch scheduling product and automation product to perform the sequencing of a validation job submittal. Then, you trigger the collection of validation results when all submitted validation jobs are complete.

Validation results are collected for jobs that are submitted in the following ways:

- By a batch scheduler or automation
- Internally (that is, the JCL for each object to be validated is constructed and submitted when the validation job is run)

Tip: You can find samples of REXX coding that show how to monitor jobs and search for messages within jobs or started tasks in *Implementing REXX Support in SDSF*, SG24-7419.

If you are not using IBM z/OS SDSF, ask your vendor which interfaces are available for invoking displays of jobs, tasks, and job output.

The processes that are invoked for each type of validation are described in the following sections.

4.2 Type 1 validation

This section describes how we monitor and collect feedback for Type 1 validation, which is the completion of a successful IPL in the IBM Z Cyber Vault environment.

Although much of the automation that is needed to set up and perform an IPL of the IBM Z Cyber Vault environment is managed and tracked by the GDPS K-sys, GDPS can initiate only the IPL. GDPS cannot manage and track activities after the IPL because the IBM Z Cyber Vault environment is in another network-isolated sysplex.

4.2.1 Collecting feedback from a Type 1 validation

To obtain visibility into the IPL process, we create a started task that is composed of a series of steps to control and monitor the validation process. This task is started by using the `COMMND01` member in our IBM Z Cyber Vault environment. A PROC, which starts the Db2 subsystems that are used in our environment, is invoked and starts the IBM Z Cyber Vault validation monitor (`CV#MON`).

As a first step, `CV#MON` accesses SDSF by using REXX code and performs the following tasks:

- Checks critical tasks to ensure that they are running.
- (Optional) Checks for a specific message to indicate that the tasks are in the correct state.

Example 4-2 shows the tasks that are checked in our environment. An input file is created for those tasks that must be active. Optionally, the message that we want to look for can be included in the input file.

Example 4-2 Tasks in our environment

```
*JOBNAME MESSAGE
GRS
SMS
ALLOCAS
TCPIP
RACF
TSO
JES2
CSF      CSFM001I ICSF INITIALIZATION COMPLETE
SDSF
DBW1MSTR DSN9022I  -DBW1 DSNYASCP 'START DB2' NORMAL COMPLETION
DBX1MSTR DSN9022I  -DBX1 DSNYASCP 'START DB2' NORMAL COMPLETION
```

The basic process flow of the REXX code that we created uses the process that is described in *Implementing REXX Support in SDSF*, SG24-7419 with some modifications to suit our requirements. Instead of using SDSF to interrogate the STATUS queue, specific jobs in the DISPLAY ACTIVE queue are interrogated. Optionally, you can use SDSF to look for a specific message for a job.

Figure 4-3 shows the process flow for system initialization validation.

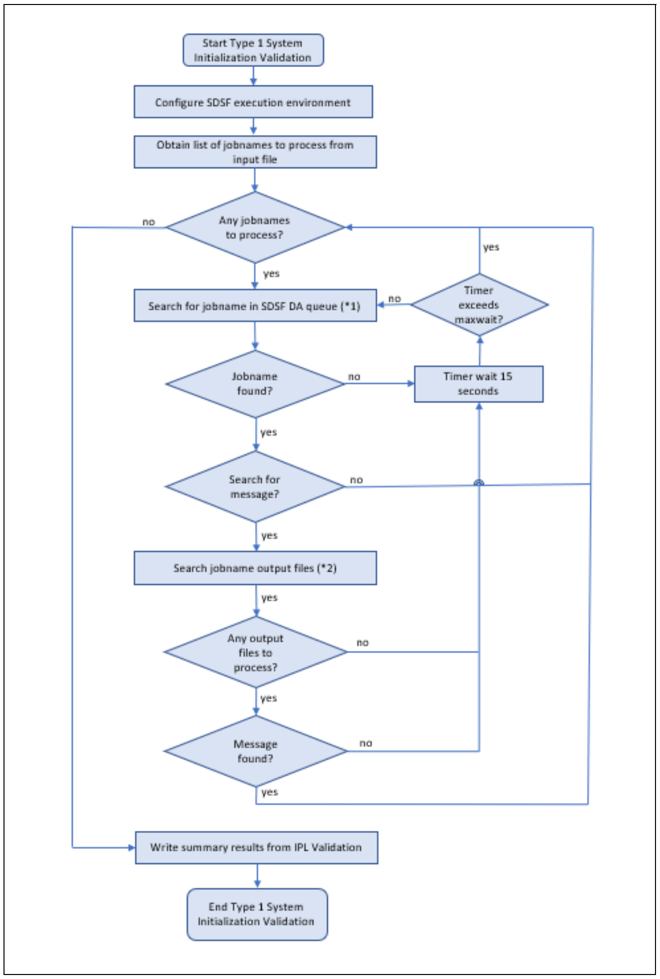
Start Type 1 System Initialization Validation
Configure SDSF execution environment
Obtain list of jobnames to process from input file
Any jobnames to process?
no
yes
Search for jobname in SDSF DA queue (*1)
Timer exceeds maxwait?
no
yes
Jobname found?
no
Timer wait 15 seconds
yes
Search for message?
no
yes
Search jobname output files (*2)
yes
Any output files to process?
no
yes
Message found?
no
yes
Write summary results from IPL Validation
End Type 1 System Initialization Validation

Figure 4-3 Process flow for system initialization validation

The following steps correspond to the numbers that are denoted in Figure 4-3:

1. Search for the job name in SDSF DA Queue.

   Example 4-3 shows the snippet of code that is used to accomplish the search.

   Example 4-3 Commands that are used to search for the job name in SDSF DA Queue

   ```
   call activate_SDSF_REXX_support
   call set_SDSF_special_variables
   call exec_sdsf "0 ISFEXEC" command
   ```

   Example 4-4 shows the details of the `activate SDSF REXX support` routine.

   Example 4-4 activate_SDSF_REXX_support routine

   ```
   /* In order to use REXX with SDSF is mandatory to add a host command  */
   /* environment prior to any other SDSF host environment commands      */
   activate_SDSF_REXX_support:
    rc_isf = isfcalls("ON")
    select
     when rc_isf = 00 then return
     when rc_isf = 01 then msg_isf = "Query failed, environment not added"
   ```

```
  when rc_isf = 02 then msg_isf = "Add failed"
  when rc_isf = 03 then msg_isf = "Delete failed"
  otherwise do
    msg_isf = "Unrecognized Return Code from isfCALLS(ON): "rc_isf
  end
 end
 if rc_isf <> 00 then do
    say "Error adding SDSF host command environment." msg_isf
    retcode = rc_isf * 10
    exit retcode
 end
 return
```

Example 4-5 shows the details of the `Set_SDSF_special_variables` routine.

Example 4-5 Set_SDSF_special_variables routine

```
set_SDSF_special_variables:
 isfprefix = jobname                       /* Only syslog jobs
*/
 isfowner  = "*"                               /* Owner does not care
*/
 isfcols   = "JNAME TOKEN JOBID"   /* Only retrieve certain columns   */
 command   = "DA"                            /* SDSF panel Display Active
*/
 return
```

Example 4-6 shows the details of the `exec SDSF` routine.

Example 4-6 exec_SDSF routine

```
/*-------------------------------------------------------------------*/
/* Subroutine to execute an SDSF REXX command testing its return code */
/*-------------------------------------------------------------------*/
exec_sdsf:
 parse arg maxrc command
 address SDSF command "(VERBOSE ALTERNATE)"
 if (maxrc = "*") then
    return rc
 if (rc > maxrc | rc < 0) then do
    call SDSF_msg_rtn
    exit rc
 end
 return 0

/*-------------------------------------------------------------------*/
/* Subroutine to list SDSF error messages                            */
/*-------------------------------------------------------------------*/
SDSF_msg_rtn:
     /***********************************************/
     /* The isfmsg variable contains a short message */
     /***********************************************/
if isfmsg<>"" then
  Say "isfmsg is:" isfmsg

     /***************************************************/
     /* The isfmsg2 stem contains additional descriptive */
     /* error messages                                   */
     /****************************************************/
do ix=1 to isfmsg2.0
  Say "isfmsg2."ix "is:" isfmsg2.ix
end
return
```

2. Search for the job name output files. Example 4-7 shows a snippet of the code that is used to accomplish this search.

Example 4-7 Commands that are used to search for the job name output file

```
njob=1
call exec_sdsf "0 ISFACT DA TOKEN('"TOKEN.njob"') PARM(NP SA)"
Say "Number of data sets allocated:" isfdsname.0
if jobup='y' then leave
do loopdd=1 to isfddname.0
  if debug > 0 then
      Say "Now reading" isfdsname.loopdd
  /* Loop reading one line each time */
  eof = 'NO'
  do while(eof = 'NO')
    "EXECIO 1 DISKR" isfddname.loopdd "(STEM line."
    if (rc = 2) then
       eof = 'YES'
    else do
      if wordpos(msgparm,line.1)>0 then do
        Say 'Verified' jobname 'is up' date() time()
        jobup='y'
        eof='YES'
        /* Indicate that message was found */
        Say  " ** Verified' jobname 'is up at"' date() time()
        Say " Message trapped' msgparm
      end
    end
  end
end
```

4.3 Type 2 validation

When you are able to perform an IPL for your IBM Z Cyber Vault environment in a controlled manner, you can construct your validation jobs or scripts for your critical subsystems.

An advantage of this critical subsystem validation is that your critical subsystems are validated at the same point-in-time. This type of validation is typically not possible in a running production system because the checks might run at different frequencies and times in production, usually during periods of low activity for that subsystem.

An important factor in reducing the time to perform validations is to achieve as much parallelism as possible when you run the checks. This step is important because it gives you a baseline of the time that is needed to perform these checks. If your IBM Z Cyber Vault logical partition (LPAR) is running on a z15, you can speed up the validation process by using IBM Z System Boost, as described in 2.6.1, “IBM Z System Recovery Boost and IBM Z Cyber Vault” on page 49.

Figure 4-4 shows the process for building a job for Type 2 structure validation for an object validation that was constructed in our environment. As listed on the process chart, code snippets and JCL are explained.

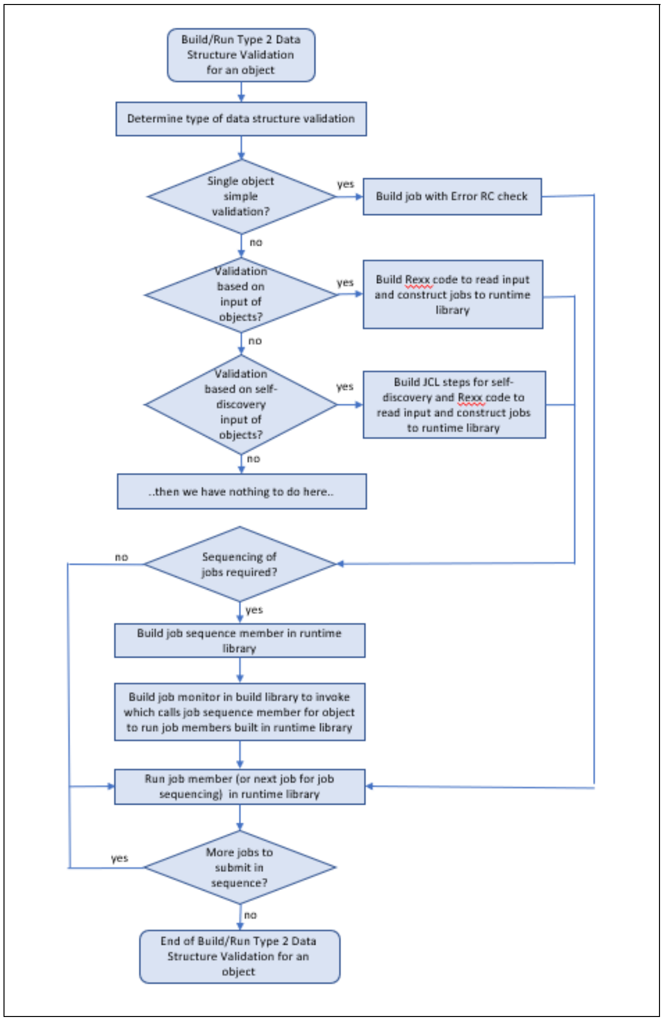
Build/Run Type 2 Data Structure Validation for an object
Determine type of data structure validation
Single object simple validation?
yes
Build job with Error RC check
no
Validation based on input of objects?
yes
Build Rexx code to read input and construct jobs to runtime library
no
Validation based on self-discovery input of objects?
yes
Build JCL steps for self-discovery and Rexx code to read input and construct jobs to runtime library
no
..then we have nothing to do here..
no
Sequencing of jobs required?
yes
Build job sequence member in runtime library
Build job monitor in build library to invoke which calls job sequence member for object to run job members built in runtime library
Run job member (or next job for job sequencing) in runtime library
yes
More jobs to submit in sequence?
no
End of Build/Run Type 2 Data Structure Validation for an object

Figure 4-4 Process flow for build job: Example 1

4.3.1 Description of the construction of a simple validation job

In this section, we describe a simple validation process for a single object based on the concept of running a single job. The example of simple validation is the data-structure validation job that was built for the RACF database. We use the job naming convention that is described in 4.1.6, “IBM Z Cyber Vault validation data set naming conventions” on page 108. the legend for job prefixes for all validation jobs in Example 4-1 on page 107, and the data set naming convention in Figure 4-1 on page 109.

The results are collected and recorded in the following two files:

- `CYBERV.@SUMM.CVERACF1`: This file indicates that a data structure validation job was submitted for the RACF database.
- `CYBERV.ERROR.CVERACF1`: This file contains the message output for the RACF utility that is run for data structure validation (`IRRUT200`). If the return code is good, this file is deleted as a final step. Otherwise, the error file is retained and noted when the results of all validation jobs are collected.

Example 4-8 shows a JCL sample of a simple validation job.

Example 4-8 JCL sample of a simple validation job

```
//CVERACF1 JOB 0001,BUILDJOBS,CLASS=A,MSGCLASS=X,//
MSGLEVEL=(1,1)//**************************************************************//
SET  CVHLQ=CYBERV                    <= CV HLQ//*         SET
CVUNIT='3390,VOL=SER=ZCV001'   <= CV VOLSER//       SET  CVUNIT=3390
<= CV VOLSER//         SET  RACF=LABPLEX.RACF.PRIMARY.DATABASE  RACF DB//
SET  JOBN=CVERACF1                   <= THIS JOBNAME//**** THE FOLLOWING ARE
BUILT BASED ON THE ABOVE SET STATEMENTS ****//          SET
ERRF=&CVHLQ..ERROR.&JOBN         <= ERROR RPT//          SET
SUMM=&CVHLQ..@SUMM.RACFCHK       <= SUMMARY RPT//**** THE FOLLOWING USED TO CLEAR
OUT PREVIOUS CV REPORT RUNS **//          SET
ODEL='DISP=(MOD,DELETE),UNIT=3390,SPACE=(TRK,0)'//********************************
******************************//* Change SYSRACF DD to current RACF database//*
-To display current RACF DB, issue TSO RVARY
LIST//**************************************************************//CLEANUP EXEC
PGM=IEFBR14//ERRF   DD &ODEL,DSN=&ERRF//SUMM   DD
&ODEL,DSN=&SUMM//**************************************************************//*
Write summary line to indicate RACF DB check
done//**************************************************************//SUMMRPT EXEC
PGM=IEBGENER//SYSPRINT DD SYSOUT=*//SYSUT1   DD *RACFCHK  ** RACF database
structure validation was
submitted**************************************************************/*
//SYSUT2   DD DISP=(NEW,CATLG),//         DSN=&SUMM,//
SPACE=(TRK,(1,1),RLSE),DSORG=PS,RECFM=FB,//
LRECL=80,BLKSIZE=0,UNIT=&CVUNIT//SYSIN    DD
DUMMY//**************************************************************//* Change
SYSRACF DD to current RACF database//* -To display current RACF DB, issue TSO
RVARY
LIST//**************************************************************//RACFCHK
EXEC PGM=IRRUT200//SYSRACF DD DISP=SHR,DSN=&RACF//SYSPRINT DD DISP=(NEW,CATLG),//
DSN=&ERRF,//         SPACE=(TRK,(15,15),RLSE),DSORG=PS,RECFM=FB,//
LRECL=132,BLKSIZE=0,UNIT=&CVUNIT//SYSUT1  DD UNIT=&CVUNIT,SPACE=(CYL,(10)),//
DCB=(LRECL=4096,RECFM=F)//SYSUT2   DD SYSOUT=*//SYSPRINT DD SYSOUT=*//SYSIN    DD
*INDEXMAPEND/*//**************************************************************//*
Delete error file dataset if RC is
good//**************************************************************//GOOD EXEC
PGM=IEFBR14,COND=(0,GT)//ERRF   DD &ODEL,DSN=&ERRF
```

4.3.2 Detailed description of the Db2 data structure validation jobs

This section describes the code and JCL that were built for one of the most complex data structure validation processes, that is, for our Db2 data structure validation.

In our environment, Db2 data structure validation jobs are set up for the following two Db2 subsystems:

- A Db2 stand-alone region running at Version 11 (Db2 SSID of DBW1)
- A Db2 data-sharing region running at Version 12 (Db2 SSIDs of DBX1 running on system MVS1, and DBX2 running on system MVS2)

A key component of the Db2 data structure validation for a selected Db2 subsystem is that layers of validation can be performed. As a best practice, run them in the following sequence:

1. Db2 system catalog and directory (DSNDB01 and DSNDB06 databases) VSAM file DSN1COPY CHECK jobs.
2. Db2 INDEX CHECK for DSNDB01 (Db2 directory table) SYSUTILX. This table must be processed by itself with no other CHECK INDEX processes running.
3. Db2 system catalog and directory CHECK INDEX for table spaces.
4. Db2 application VSAM file DSN1COPY CHECK jobs.
5. Db2 application CHECK INDEX for table spaces (optional and dependent on the length of time that is needed to complete the validation).

Figure 4-5 shows a process flow for the build job for the Db2 subsystem object validation that was constructed in our environment. As listed on the process chart, code snippets and JCL are explained.

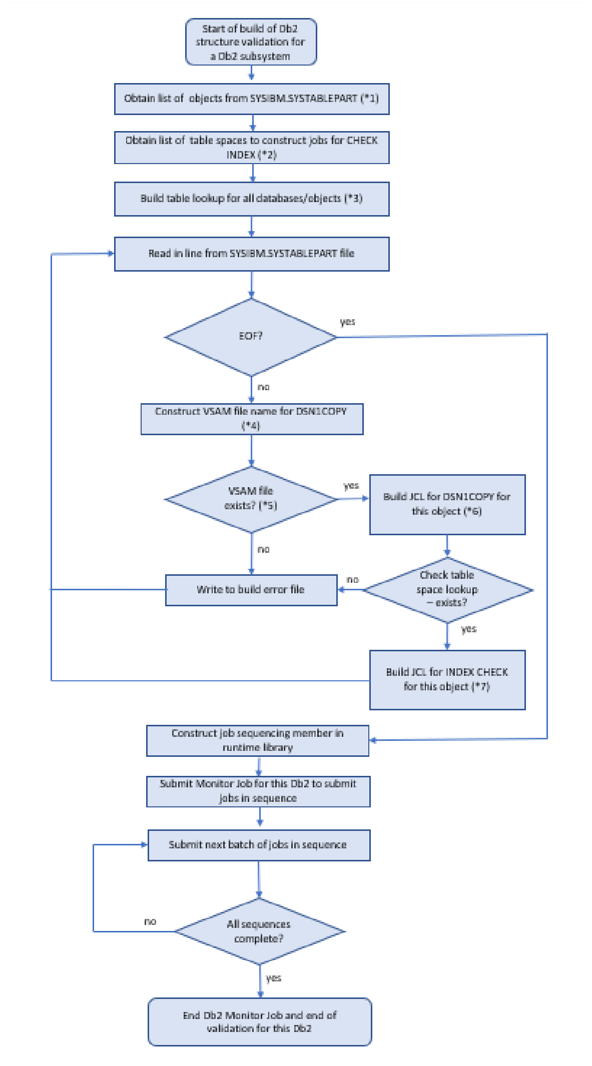
Start of build of Db2 structure validation for a Db2 subsystem
Obtain list of objects from SYSIBM.SYSTABLEPART (*1)
Obtain list of table spaces to construct jobs for CHECK INDEX (*2)
Build table lookup for all databases/objects (*3)
Read in line from SYSIBM.SYSTABLEPART file
EOF?
yes
no
Construct VSAM file name for DSN1COPY (*4)
VSAM file exists? (*5)
yes
Build JCL for DSN1COPY for this object (*6)
no
Write to build error file
no
Check table space lookup – exists?
yes
Build JCL for INDEX CHECK for this object (*7)
Construct job sequencing member in runtime library
Submit Monitor Job for this Db2 to submit jobs in sequence
Submit next batch of jobs in sequence
no
All sequences complete?
yes
End Db2 Monitor Job and end of validation for this Db2

Figure 4-5 Process flow for build job: Example 2

The following steps correspond to the numbers that are denoted in Figure 4-5 on page 117:

1. Obtain a list of objects from SYSIM.SYSTABLEPART, as shown in Example 4-9. In this case, we exclude the Db2 workfile databases (database DSNDB07 and databases beginning with WRK* in our environment).

Example 4-9 Obtaining a list of objects from SYSIM.SYSTABLEPART

```
//****************************************************
//* DB2DSPLY: RUN SQL QUERY ON SYSIBM.SYSTABLEPART
//****************************************************
//QUERY  EXEC PGM=IKJEFT01,COND=(4,LT),DYNAMNBR=20
//STEPLIB DD DISP=SHR,DSN=DSNB10.SDSNLOAD
//SYSTSPRT DD SYSOUT=*
//SYSPRINT DD  DISP=(NEW,CATLG),DSN=&DDSN,
//            UNIT=&CVUNIT,SPACE=(TRK,(30,30)),LRECL=133,RECFM=FB,
//            DSORG=PS,BLKSIZE=0
//SYSTERM  DD SYSOUT=*
//CEEDUMP  DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SYSOUT   DD SYSOUT=*
//****************************************************
//* WILL NEED TO SPECIFY SYSTEM AND LIB ON SYSTSIN
//****************************************************
//SYSTSIN  DD *
 DSN SYSTEM(DBW1)
 RUN   PROGRAM(DSNTEP2) PLAN(DSNTEP11) -
       LIB('DSNDBWO.RUNLIB.LOAD') PARMS('/ALIGN(LHS)')
 END
//SYSIN    DD *
SELECT DBNAME,TSNAME,PARTITION,VCATNAME,IPREFIX
FROM SYSIBM.SYSTABLEPART
WHERE DBNAME NOT IN ('DSNDB07')
AND DBNAME NOT LIKE 'WRK%';
/*
```

2. In this case, we created REXX code to invoke Db2 commands and then invoked our table space lookup routine, as shown in Example 4-10.

Example 4-10 Obtaining a list of table spaces

```
READ_OBJECTS:
Makebuf
Queue "-DIS DATABASE(*) LIMIT(*)"  /*DB2 COMMAND */
Queue "END"
Address(TSO)
Y = OUTTRAP(db2cmd.,"*","NOCONCAT")
"DSN SYSTEM("||db2system||")"      /*talk to the Db2 SSID */
Y = OUTTRAP("OFF")
Dropbuf
ADDRESS(ISPEXEC)
IX.=0;TS.=0;LS.=0
Do I = 1 to db2cmd.0
  var1='';var2='';var3='';var4='';var5='';var6='';
  parse var db2cmd.I var1 var2 var3 var4 var5 var6
  if wordpos(db2cmd.I,'NORMAL COMPLETION')>1 then NOP
  else Call BUILD_OBJECT_TYPE_LOOKUP
```

```
End
Return
```

3. Build a table space lookup by using REXX code, as shown in Example 4-11.

Example 4-11 Building a table space lookup

```
BUILD_OBJECT_TYPE_LOOKUP:
var1=strip(var1)
var2=strip(var2)
var5=strip(var5)
if var3 = 'DATABASE' then do
  database=var5
  if (database='DSNDB07') |,
     (substr(database,1,3) = 'WRK') then skip_db = 1
  else skip_db = 0
  if skip_db=1 then say '**Skipping database' database
End
SELECT
  When skip_db=1 then NOP
  When var2 = 'TS' then do
    if TS.database.var1 <> 1 then TS.database.var1 = 1
  End
  When var2 = 'LS' then do
    if LS.database.var1 <> 1 then LS.database.var1 = 1
  End
  When var2 = 'IX' then do
    if IX.database.var1 <> 1 then IX.database.var1 = 1
  End
OTHERWISE NOP
END
Return
```

4. Construct a VSAM file name for the DSN1COPY job, as shown Example 4-12.

Example 4-12 Constructing a VSAM file name

```
READ_IN_SYSTABLEPART:
Makebuf
in.0=0
"execio * diskr infile ( finis stem in."
/* in is the stem (array) that holds the input data */
"execio "o" diskr infile (finis "
Say 'in.0 =' in.0
Dropbuf
/* Nows the fun part - parsing thru output to get fields we want */

Do  ST=2 to in.0
  var1='';var2='';var3='';var4='';var5='';var6=''
  parse var in.ST var1 '_|' var2 '|' var3 '|' var4 '|' var5 '|',
        var6 '|'
  /* var1 will always be numeric for a line that contains the
     information we need */
  if datatype(var1,'N')=1 then do
    dbname=strip(var2)
    tsname=strip(var3)
    partition=strip(var4)
    vcatname=strip(var5)
    iprefix=strip(var6)
    SELECT
      When partition = 0 then partition = 'A001'
      When partition > 0000 and partition < 1000 then
       partition='A'||right(partition,3,'0')
      When partition > 0999 and partition < 2000 then
       partition='B'||right(partition,3,'0')
      when partition > 1999 and partition < 3000 then
       partition='C'||right(partition,3,'0')
      when partition > 2999 and partition < 4000 then
       partition='D'||right(partition,3,'0')
      when partition > 3999 and partition < 4097 then
       partition='E'||right(partition,3,'0')
    END
    instance=iprefix||'0001'

    /* Check table to see type: TS, LS, or IX */
   Call CHECK_TYPE

   if wordpos(dbname,'DSNDB01 DSNDB06') > 0 then
     vcatname=db2hlq
   /* Construct dsname to feed into DSN1COPY CHECK job */
   dsname=vcatname||'.DSNDBD.'||dbname||'.'||tsname||'.'
   dsname=dsname||instance||'.'||partition
   dsname=strip(dsname)
 Say 'SQLout=' type vcatname dbname tsname instance partition
 Say '*dsname=' dsname

   /* Check dsname that it is accessible */
   Call CHECK_DSNAME

   If ds_ok='y' then do
     Call WRITE_DSN1COPY
     If (type='TS') then Call WRITE_INDEX_CHECK
   End
 End
End
Return
```

5. Check whether the VSAM file name that was constructed exists, as shown in Example 4-13.

Example 4-13 Checking the VSM file name

```
CHECK_DSNAME:
Makebuf
ds_ok='y'
x=outtrap(line.)
"listc ent('"||dsname||"')"
retcode = rc
x=outtrap('off')
if retcode<>0 then do
 ds_ok='n'
 err_cnt=err_cnt+1
err.1='DB2CHK   **' db2system 'Error LISTC chk for' dsname
```

```
 "execio "1" diskw errf (stem err."
end
Dropbuf
Return
```

6. Build a JCL for the DSN1COPY job. Example 4-14 shows a sample of a JCL that is built for one object.

Example 4-14 Sample JCL for DSN1COPY

```
//* Delete error file if it exists
//CLEANUP EXEC PGM=IEFBR14
//DD1    DD DISP=(MOD,DELETE),UNIT=3390,SPACE=(TRK,0),
//          DSN=CYBERV.ERROR.CVW00001
//DSN1COPY EXEC PGM=DSN1COPY,PARM='CHECK'
//STEPLIB  DD  DISP=SHR,DSN=DSNC10.DBX0.SDSNEXIT
//         DD  DISP=SHR,DSN=DSNC10.SDSNLOAD
//SYSPRINT DD DISP=(NEW,CATLG),
//         DSN=CYBERV.ERROR.CVW00001,
//         SPACE=(TRK,(15,15),RLSE),DSORG=PS,RECFM=FB,
//         LRECL=132,BLKSIZE=0,UNIT=3390
//SYSUT1 DD DISP=SHR,DSN=DSNDBX0.DSNDBD.ACDB.SYSTSTAB.I0001.A001,DISP=SHR
//SYSUT2   DD DUMMY
//* Delete error file to dataset if RC is good
//GOOD EXEC PGM=IEFBR14,COND=(4,LE)
//DD1    DD DISP=(MOD,DELETE),UNIT=3390,SPACE=(TRK,0),
//          DSN=CYBERV.ERROR.CVW00001
```

7. Build a JCL for the INDEX CHECK job. Example 4-15 shows a sample of a JCL that is built for one table space.

Example 4-15 Sample JCL for INDEX CHECK

```
//* Delete error file if it exists
//CLEANUP EXEC PGM=IEFBR14
//DD1    DD DISP=(MOD,DELETE),UNIT=3390,SPACE=(TRK,0),
//          DSN=CYBERV.ERROR.CVX00001
//DB2CHK   EXEC PGM=DSNUTILB,REGION=0M,PARM=(DBX1)
//STEPLIB  DD  DISP=SHR,DSN=DSNC10.DBX0.SDSNEXIT
//         DD  DISP=SHR,DSN=DSNC10.SDSNLOAD
//SYSPRINT DD DISP=(NEW,CATLG),
//         DSN=CYBERV.ERROR.CVX00001,
//         SPACE=(TRK,(15,15),RLSE),DSORG=PS,RECFM=FB,
//         LRECL=122,BLKSIZE=0,UNIT=3390
//* Below include member has additional SORTWKxx DD statements
//INCLSORT INCLUDE MEMBER=@SORTWRK
//SYSOUT   DD SYSOUT=*
//UTPRINT  DD SYSOUT=*
//SYSIN    DD *
 CHECK INDEX(ALL) TABLESPACE ACDB.SYSTSTAB
 SORTNUM 24
/*
//* Delete error file to dataset if RC is good
//GOOD EXEC PGM=IEFBR14,COND=(5,LT)
//DD1    DD DISP=(MOD,DELETE),UNIT=3390,SPACE=(TRK,0),
//          DSN=CYBERV.ERROR.CVX00001
```

4.3.3 Automating data structure validation

The `CV#MON` job is the first one that is run in the IBM Z Cyber Vault validation cycle. It can run as a started task or a batch job. Here are the key inputs to the `CV#MON` job:

- `<cvhlq>.@BUILD.JOBS` - members @JOBIN1-4. These processes are performed in sequence and correspond to the following items:

  – `@JOBIN1`: Type 2 - Data Structure Validation Jobs. The job prefix of `CV@` denotes a build job for the validations. Example 4-16 shows a sample of our `@JOBIN1` member. Each build job submits the JCL member that built in to the `<cvhlq>.@RUN.JOBS` library for that type of data-structure validation. Some build jobs might create several members in `<cvhlq>.@RUN.JOBS`, such as Db2 structure validation jobs, where a monitoring job that is specific to that Db2 subsystem is started with a job name of `CV#<ssid>`, where `ssid` is the Db2 subsystem ID. The build jobs for validation for Db2 feed in to the specific validation sequence as required and ensure that each sequence is completed before the next sequence of jobs is submitted.

  Example 4-16 JOBIN1

  ```
  CYBERV.@BUILD.JOBS(CV@CATCK)
  CYBERV.@BUILD.JOBS(CV@VVDCK)
  CYBERV.@BUILD.JOBS(CV@VSMCA)
  CYBERV.@BUILD.JOBS(CV@RACF)
  CYBERV.@BUILD.JOBS(CV@DBW1)
  CYBERV.@BUILD.JOBS(CV@DBX1)
  ```

  – `@JOBIN2`: Type 3 - Data Content Validation Jobs (for specific application validation jobs that are to be inserted into the IBM Z Cyber Vault validation cycle). `@JOBIN2` might point to one or more data content validation jobs. In our sample, specific data content validation jobs are not built. Therefore, a dummy job (which basically does nothing) is submitted to satisfy the Type 3 requirement, as shown in Example 4-17.

  Example 4-17 JOBIN2

  ```
  CYBERV.@BUILD.JOBS(CV@DUMMY)
  ```

  > **Note:** For @`JOBIN1` and `@JOBIN2`, you can specify `CYBERV.@BUILD.JOBS(CV@DUMMY)` as input (where `CYBERV` is your `<cvhlq>`) if you want to bypass validation testing for that level. If you specify a `CV@DUMMY` job for both levels, then you are not doing any validations beyond a simple IPL.

  – `@JOBIN3`: When this job is triggered it signals that all data validation jobs (Type 2 and Type 3) are complete and starts the collection of all summary, detail, and error reporting for the validation cycle run. Generally, `@JOBIN3` has a single input of the generic data collection job, as shown in Example 4-18.

  Example 4-18 JOBIN3

  ```
  CYBERV.@BUILD.JOBS(CV@COLL1)
  ```

  – `@JOBIN4`: When this job is run, it collects the IBM Z Cyber Vault summary, detail, and error collection into the respective generation data group (GDG) history files. Generally, `@JOBIN4` has a single input of the generic data history collection job, as shown in Example 4-19 on page 123.

Example 4-19 JOBIN4

```
CYBERV.@BUILD.JOBS(CV@COLL2)
```

JCL for CV#MON

Example 4-20 shows the JCL that is submitted for the `CV#MON` IBM Z Cyber Vault Monitoring task. It is the first job that is submitted in the validation cycle, and it triggers the downstream jobs in the sequence that is defined in the `@JOBINx` members. The `CV#MON` job stays in execution mode until it detects that there are no more validation jobs to be run.

Example 4-20 JCL for CV#MON

```
//CV#MON   JOB 0001,CVMONITOR,CLASS=A,MSGCLASS=X,
//        MSGLEVEL=(1,1)
//****************************************************************
//*  THIS JOB IS THE FIRST JOB SUBMITTED FOR VALIDATION.
//*  IT WILL STAY IN THE SYSTEM TO MONITOR JOBS IN SDSF DA QUEUE.
//*  ONCE THERE ARE NO CV* JOBS RUNNING THEN WE ARE COMPLETE WITH
//*  THE VALIDATION CYCLE.
//****************************************************************
//* EDIT @JOBIX  MEMBER TO SPECIFY JOB INDEX (3 CHARACTER PREFIX
//*      DENOTING TYPE OF VALIDATION JOB)
//* EDIT @JOBIN<N> MEMBER (WHERE <N> IS SEQUENCE OF RUNNING OF
//*      VALIDATION JOBS. THIS JOB WILL MONITOR FOR COMPLETION OF ALL
//*      JOB IN CURRENT SEQUENCE BEFORE SUBMITTING JOBS IN NEXT
//*      SEQUENCE.
//****************************************************************
//          SET  CVHLQ=CYBERV                     <= CV HLQ
//          SET  CVVOL=ZCV001                     <= CV VOLSER
//****************************************************************
//**** EDIT THESE MEMBERS TO SPECIFY WHAT JOBS TO RUN IN SEQUENCE
//****************************************************************
//          SET  IDSN1=&CVHLQ..@BUILD.JOBS(@JOBIN1) <= 1ST SERIES JOBS
//          SET  IDSN2=&CVHLQ..@BUILD.JOBS(@JOBIN2) <= 2ND SERIES JOBS
//          SET  IDSN3=&CVHLQ..@BUILD.JOBS(@JOBIN3) <= 3RD SERIES JOBS
//          SET  IDSN4=&CVHLQ..@BUILD.JOBS(@JOBIN4) <= 4TH SERIES JOBS
//****************************************************************
//**** THE FOLLOWING ARE STANDARD FILE NAMES USED, DO NOT CHANGE
//**** THEM AS SOME AUTO-BUILD OF JCL VIA REXX CODE WILL DEPEND ON
//**** THIS NAMING CONVENTION.
//****************************************************************
//          SET  RUNLIB=&CVHLQ..@RUN.JOBS         <= RUN LIBRARY
//          SET  SUMM=&CVHLQ..@SUMM.@CVEMON       <= CV SUMMARY
//          SET  DETL=&CVHLQ..@DETL.@CVEMON       <= CV DETAIL
//          SET  ERRF=&CVHLQ..@ERR.@CVEMON        <= CV ERRORS
//          SET  CVSUMM=&CVHLQ..CVRUN.SUMMARY     <= CV BUILD JOB SUMM
//          SET  CVERRF=&CVHLQ..CVRUN.ERRORS      <= CV BUILD JOB SUMM
//          SET  CVDETL=&CVHLQ..CVRUN.DETAIL      <= CV RUN DETAIL
//          SET  BSUMM=&CVHLQ..CVRUN.JOB.CVSUMM    = BUILD SUMMARY
//          SET  BERRF=&CVHLQ..CVRUN.JOB.CVERRF    = BUILD ERRORS
//          SET  BERRD=&CVHLQ..CVRUN.JOB.CVERRD    = ERROR FILE DSN
//          SET  BDETL=&CVHLQ..CVRUN.JOB.CVDETL    = BUILD DETAIL
//          SET  XLIB=&CVHLQ..@BUILD.JOBS          = EXEC LIB
//**** THE FOLLOWING USED TO CLEAR OUT PREVIOUS CV REPORT RUNS **
//          SET  ODEL='DISP=(MOD,DELETE),UNIT=SYSDA,SPACE=(TRK,0)'
```

```
//********************************************************
//CLEANUP  EXEC  PGM=IEFBR14
//DETL     DD  &ODEL,DSN=&DETL
//SUMM     DD  &ODEL,DSN=&SUMM
//ERRF     DD  &ODEL,DSN=&ERRF
//CVDETL   DD  &ODEL,DSN=&CVDETL
//CVERRF   DD  &ODEL,DSN=&CVERRF
//CVERRD   DD  &ODEL,DSN=&CVERRD
//CVSUMM   DD  &ODEL,DSN=&CVSUMM
//BSUMM    DD  &ODEL,DSN=&BSUMM
//BERRF    DD  &ODEL,DSN=&BERRF
//BERRD    DD  &ODEL,DSN=&BERRD
//BDETL    DD  &ODEL,DSN=&BDETL
//RUNLIB   DD  &ODEL,DSN=&RUNLIB
//********************************************************
//* DELETE ANY FILES BUILT FROM A PREVIOUS VALIDATION RUN
//********************************************************
//DELERR   EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//SYSIN    DD *
 DELETE 'CYBERV.ERROR.**' MASK
 DELETE 'CYBERV.CVRUN.**' MASK
 DELETE 'CYBERV.@DETL.**' MASK
 DELETE 'CYBERV.@ERR.**'         MASK
 DELETE 'CYBERV.@SUMM.**' MASK
 SET MAXCC=0
/*
//********************************************************
//* BUILD NEW RUN LIBRARY
//********************************************************
//BUILDRUN EXEC PGM=IEFBR14
//DD1      DD DISP=(NEW,CATLG),
//           DSN=&RUNLIB,
//           UNIT=3390,SPACE=(TRK,(60,60,10)),
//           RECFM=FB,LRECL=80,DSORG=PO,BLKSIZE=0,
//           DSNTYPE=LIBRARY
//********************************************************
//* START SUBMITAL AND MONITORING OF CVE JOBS
//********************************************************
//CVEMON1  EXEC  PGM=IKJEFT01,DYNAMNBR=20
//SYSEXEC  DD DISP=SHR,DSN=&XLIB
//SYSOUT   DD SYSOUT=*
//SYSTSPRT DD SYSOUT=*
//INFILE1  DD DISP=SHR,DSN=&IDSN1
//INFILE2  DD DISP=SHR,DSN=&IDSN2
//INFILE3  DD DISP=SHR,DSN=&IDSN3
//INFILE4  DD DISP=SHR,DSN=&IDSN4
//DETL     DD DISP=(NEW,CATLG),DSN=&DETL,
//           UNIT=3390,SPACE=(TRK,(90,90)),
//           RECFM=FB,LRECL=80,DSORG=PS,BLKSIZE=0,FREE=CLOSE
//SUMM     DD DISP=(NEW,CATLG),DSN=&SUMM,
//           UNIT=3390,SPACE=(TRK,(90,90)),
//           RECFM=FB,LRECL=80,DSORG=PS,BLKSIZE=0,FREE=CLOSE
//ERRF     DD DISP=(NEW,CATLG),DSN=&ERRF,
//           UNIT=3390,SPACE=(TRK,(90,90)),
//          RECFM=FB,LRECL=80,DSORG=PS,BLKSIZE=0,FREE=CLOSE
//SYSTSIN  DD *
 %XCVEMON
/*
//********************************************************
//* EMAIL THE SUMMARY REPORT
//* (jcl to email the report in &CVHLQ..CVRUN.SUMMARY goes here)
//********************************************************
//********************************************************
//* EMAIL THE ERRORS REPORT
//*  (jcl to email the report in &CVHLQ..HISTORY.ERRORS(0) goes here)
//********************************************************
//********************************************************
//* EMAIL THE DETAIL REPORT WHICH LISTS ALL JOBS AND OBJECTS VALIDATED
//*  (jcl to email the report in &CVHLQ..CVRUN.DETAIL goes here)
//********************************************************
```

REXX for XCVEMON

The code that is shown in Example 4-21 is run by CV#MON, and it contains timestamps for all validations that are performed. It also monitors for completion of all validation jobs.

Example 4-21 REXX for XCVEMON

```
/* REXX - XCVEMON
   This REXX submits jobs @JOBIN1-4 members in sequence. It
   monitors all CV* jobs for completion for each @JOBIN<n> before
   moving on to the next @JOBIN<n> series
   */ trace 0
ww.1= '***************************************************************'
ww.2= '************ Cyber Vault Validation Summary Report ************'
ww.3= '*                                                             *'
ww.4= '* This report provides timestamps of validation activities.   *'
ww.5= '*                                                             *'
ww.6= '* It also notes validations that were not performed for noted *'
ww.7= '* objects due to anomalies seen in your environment. This may *'
ww.8= '* be due to transient activity occurring at the time of data  *'
ww.9= '* capture, or may indicate some clean up needed in your       *'
ww.10='* environment for the objects noted.                          *'
ww.11='***************************************************************'
ww.12='CV#MON   ** Started Cyber Vault validation at' date() time()
w=12
Call WRITE_SUMM
o=0
/* There must always be 4 @jobin<x> members in the build library
   where generally these members will contain jobs to submit and
   monitor in the following sequence:
   @JOBIN1 - Data Structure Validation jobs
   @JOBIN2 - Application Validation jobs
   @JOBIN3 - End of submittal of all validation jobs, create reports
   @JOBIN4 - Create history PDS GDGs for each validation run
*/
i=1
Do While i <= 4
  in.0=0
  infile='infile'||i
  Say 'Reading in @JOBIN'||i 'to determine what to execute'
  /* Read in the list of jobs to submit and submit them */
  "execio * diskr " infile " ( finis stem in."
  "execio "o" diskr " infile " (finis "
  Say 'Number of lines in' infile 'is' in.0
  /* Infile3 and Infile4 are for collections of output so must
     free the files to do so */
  if i=3 then do
    "execio "o" diskw summ (finis"
    "execio "o" diskw detl (finis"
    "execio "o" diskw errf (finis"
    "Free fi(summ)"
    "Free fi(detl)"
    "Free fi(errf)"
  end
  s=1
  do while s <= in.0
    say infile 'contents' in.s
    job=strip(in.s)
    "submit '"||job||"'"
    ww.1='CV#MON   ** Submitted' job 'at' date() time()
    w=1
    call WRITE_SUMM
    s=s+1
  end

  /* Now start monitoring jobs */
  call MONITOR_JOBS
  i=i+1
end
exit

MONITOR_JOBS:
jobsdone='n'
Do while jobsdone='n'
  rc=isfcalls('ON')
        /* Set the jobname prefix and owner */
  isfprefix="CV*"
  isfowner="*"
      /* Access the ST panel.  A TOKEN variable is  */
      /* created for each row which is subsequently */
      /* needed to perform actions                  */
  Address SDSF "ISFEXEC DA"
  lrc=rc
  if lrc<>0 then do
    call MSGRTN  /* List any error messages */
ww.1='*~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~*'
ww.2='CV#MON   ** Error encountered at' date() time()
ww.3='CV#MON   ** SDSF monitoring of jobs had RC=' lrc
ww.4='CV#MON   ** Exiting Cyber Vault validation with RC=99'
ww.5='*~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~*'
    Say '** Exiting MONITOR_JOBS with RC of' lrc 'encountered'
    w=5
    Call WRITE_SUMM
  end
```

```
  numrows=isfrows
  Say 'Number of lines captured in SDSF DA panel is'  numrows
  job = word(isfcols,1)
  /* There will always be this job in the system so check */
  /* if 2 or more jobs in DA queue                        */
  if numrows >1 then do
    do ix=2 to numrows    /* Loop for all rows returned */
      Say 'Job running is' JNAME.ix
      ix=ix+1
    end
  end
  else do
ww.1='CV#MON   ** All jobs complete for' infile 'at' date() time()
ww.2='***************************************************************'
    Say '** All jobs complete for infile' infile 'at' date() time()
    w=2
    Call WRITE_SUMM
    jobsdone='y'
  end
  if jobsdone='n' then Call SLEEP_TIME
end
rc=isfcalls('OFF')
Return

/* Subroutine to list error messages */
MSGRTN: procedure expose isfmsg isfmsg2.
/* The isfmsg variable contains a short message */
if isfmsg<>"" then
  Say "isfmsg is:" isfmsg
/* The isfmsg2 stem contains additional descriptive */
/* error messages                                   */
do ix=1 to isfmsg2.0
  Say "isfmsg2."ix "is:" isfmsg2.ix
end
Return

SLEEP_TIME:
CALL SYSCALLS 'ON'
ADDRESS SYSCALL
"SLEEP" 10
CALL SYSCALLS 'OFF'
Return

WRITE_SUMM:
if i < 3 then
"execio "w" diskw summ (stem ww."
Return
```

4.3.4 Collecting feedback from data structure validation

The collection of feedback from an IBM Z Cyber Vault validation run is performed by the CV#MON job when it is determined that all IBM Z Cyber Vault validation jobs have completed. The CV#MON job triggers the CV@COLL1 job to collect information and produce a summary report.

A sample of an IBM Z Cyber Vault summary email report is shown in Example 4-22.

Example 4-22 IBM Z Cyber Vault validation

```
***************************************************************
*  Cyber Vault Validation - Jobs with exception return codes  *
***************************************************************
*  ** Note: Reported below are validations performed that did *
*           not meet the criteria for passing a validation    *
*           based on coded criteria (return code checks)      *
***************************************************************
CVECOLL  ** Started Exception collection at 22 Mar 2021 03:13:26
            Any exception output will be listed in data sets
            noted below - a Legend is provided to indicate
            the type of validation being performed
***************************************************************
*****************************************************************
* Legend for Job prefix for all validation jobs
*****************************************************************
Prefix   Description
-------- --------------------------------------------------------
CV@      Build jobs for validation functions
CV#      Monitoring jobs
CVA      Cyber Vault Information Collection
CVB      BCS-VVDS validation jobs
CVC      Catalog diagnose and examine jobs
CVD      VSAM KSDS examine jobs
CVE      RACF structure validation job
CVS      DBW1 system DSN1COPY jobs
CVT      DBW1 application DSN1COPY jobs
CVU      DBW1 system CHECK INDEX jobs
CVV      DBW1 application CHECK INDEX jobs
CVW      DBX1 system DSN1COPY jobs
CVX      DBX1 application DSN1COPY jobs
CVY      DBX1 system CHECK INDEX jobs
CVZ      DBX1 application CHECK INDEX jobs
*****************************************************************
Check error file CYBERV.ERROR.CVD00456
Check error file CYBERV.ERROR.CVD00478
Check error file CYBERV.ERROR.CVT00693
Check error file CYBERV.ERROR.CVT00694
Check error file CYBERV.ERROR.CVT00695
Check error file CYBERV.ERROR.CVT00696
Check error file CYBERV.ERROR.CVT00697
Check error file CYBERV.ERROR.CVT00698
Check error file CYBERV.ERROR.CVT00699
Check error file CYBERV.ERROR.CVT00700
Check error file CYBERV.ERROR.CVT00701
Check error file CYBERV.ERROR.CVT00702
Check error file CYBERV.ERROR.CVT00703
Check error file CYBERV.ERROR.CVT00704
Check error file CYBERV.ERROR.CVT00705
Check error file CYBERV.ERROR.CVX03246
CVECOLL  **************************************************
CVECOLL  ** Count of jobs run with exception RCs= 16
CVECOLL  **************************************************
```

```
***************************************************************
************ Cyber Vault Validation Summary Report ************
*                                                             *
* This report provides timestamps of validation activities.   *
*                                                             *
* It also notes validations that were not performed for noted *
* objects due to anomalies seen in your environment. This may *
* be due to transient activity occurring at the time of data  *
* capture, or may indicate some clean up needed in your       *
* environment for the objects noted.                          *
***************************************************************
CV#MON   ** Started Cyber Vault validation at 22 Mar 2021 03:10:43
CV#MON   ** Submitted CYBERV.@BUILD.JOBS(CV@CATCK) at 22 Mar 2021 03:10:43
CV#MON   ** Submitted CYBERV.@BUILD.JOBS(CV@VVDCK) at 22 Mar 2021 03:10:43
CV#MON   ** Submitted CYBERV.@BUILD.JOBS(CV@VSMCA) at 22 Mar 2021 03:10:43
CV#MON   ** Submitted CYBERV.@BUILD.JOBS(CV@RACF) at 22 Mar 2021 03:10:43
CV#MON   ** Submitted CYBERV.@BUILD.JOBS(CV@DBW1) at 22 Mar 2021 03:10:43
CV#MON   ** Submitted CYBERV.@BUILD.JOBS(CV@DBX1) at 22 Mar 2021 03:10:43
CV#MON   ** All jobs complete for infile1 at 22 Mar 2021 03:13:26
***************************************************************
CV#MON   ** Submitted CYBERV.@BUILD.JOBS(CV@DUMMY) at 22 Mar 2021 03:13:26
CV#MON   ** All jobs complete for infile2 at 22 Mar 2021 03:13:26
***************************************************************
CATCHK   ** Started build for catalog jobs at 22 Mar 2021 03:10:43
CATCHK   ** No errors encountered during build of catalog jobs
CATCHK   ** Total number of Catalog verify jobs = 14
CATCHK   ** Submitting jobs created by CATCHK at 22 Mar 2021 03:10:44
***************************************************************
*~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~*
CATVVD   ** Error(s) encountered during build of BCS-VVDS jobs
CATVVD   ** Number of errors = 29
CATVVD   ** Need to examine errors reported for this job
*~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~*
CATVVD   ** Total number of BCS-VVDS verify jobs = 201
CATVVD   ** Submitting jobs created by CATVVD at 22 Mar 2021 03:10:45
***************************************************************
DB2CHK   ** DBW1 started build jobs at 22 Mar 2021 03:10:44
DB2CHK   ** DBW1 building CHECK INDEX for system DBs
DB2CHK   ** DBW1 skipping CHECK INDEX for application DBs
*~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~*
DB2CHK   ** DBW1 Errors encountered during build of jobs
DB2CHK   ** DBW1 Number of errors => 20
DB2CHK   ** DBW1 Check the errors reported in this job
*~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~*
DB2CHK   ** DBW1 Total system DSN1COPY jobs = 116
DB2CHK   ** DBW1 -System DSN1COPY for LOBs = 24
DB2CHK   ** DBW1 Total system INDEX CHECK jobs = 92
DB2CHK   ** DBW1 Total application DSN1COPY jobs = 717
DB2CHK   ** DBW1 -Application DSN1COPY for LOBs = 80
DB2CHK   ** DBW1 Total applic INDEX CHECK jobs = 0
DB2CHK   ** DBW1 ** Launching monitor at 22 Mar 2021 03:10:51
***************************************************************
DB2CHK   ** DBX1 started build jobs at 22 Mar 2021 03:10:44
DB2CHK   ** DBX1 building CHECK INDEX for system DBs
DB2CHK   ** DBX1 skipping CHECK INDEX for application DBs
```

```
*~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~*
DB2CHK   ** DBX1 Errors encountered during build of jobs
DB2CHK   ** DBX1 Number of errors => 26
DB2CHK   ** DBX1 Check the errors reported in this job
*~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~*
DB2CHK   ** DBX1 Total system DSN1COPY jobs = 130
DB2CHK   ** DBX1 -System DSN1COPY for LOBs = 31
DB2CHK   ** DBX1 Total system INDEX CHECK jobs = 99
DB2CHK   ** DBX1 Total application DSN1COPY jobs = 3252
DB2CHK   ** DBX1 -Application DSN1COPY for LOBs = 79
DB2CHK   ** DBX1 Total applic INDEX CHECK jobs = 0
DB2CHK   ** DBX1 ** Launching monitor at 22 Mar 2021 03:11:04
***************************************************************
RACFCHK  ** RACF database structure validation was submitted
***************************************************************
VSAMCHK  ** Started build for VSAMCHK jobs at 22 Mar 2021 03:10:56
VSAMCHK  ** Will VERIFY be added to jobs ==> yes
*~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~*
VSAMCHK  ** Error(s) encountered during build of VSAMCHK jobs
VSAMCHK  ** Number of errors = 2
VSAMCHK  ** Need to examine error file for this job
*~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~*
VSAMCHK  ** Total number of VSAM KSDS verify jobs = 492
VSAMCHK  ** Submitting jobs created by VSAMCHK at 22 Mar 2021 03:10:58
***************************************************************
CATVVD   ** Error for SYS1.VVDS.VA2LOC1 in catalog IODF.CAT0
CATVVD       -Error reported was DATASET NOT FOUND
CATVVD   ** Error for SYS1.VVDS.VCICSCF in catalog IODF.CAT0
CATVVD       -Error reported was DATASET NOT FOUND

CATVVD   ** Error for SYS1.VVDS.VSHZFS0 in catalog SYS1.MCAT.VS1CAT1
CATVVD       -Error reported was ERROR PROCESSING REQUESTED DATASET
CATVVD   ** Error for SYS1.VVDS.VS1OS05 in catalog SYS1.MCAT.VS1CAT1
CATVVD       -Error reported was ERROR PROCESSING REQUESTED DATASET

DB2CHK   ** DBW1 Error LISTC chk for DSNDBW0.DSNDBD.ADBDCHG.ADBSCHGM.I0001.A001
DB2CHK   ** DBW1 Error LISTC chk for DSNDBW0.DSNDBD.ADBDCHG.ADBSPF1.I0001.A001
DB2CHK   ** DBW1 Error in lookup for DSN8D11X XPRO0000
DB2CHK   ** DBW1 Error in lookup for DSN8D11X XCUS0000

VSAMCHK  ** Error for CATALOG.USER.COMIOD LISTC check, RC= 4
VSAMCHK  ** Error for SYS1.UCAT.VSHSMPE LISTC check, RC= 4
```

In Example 4-22 on page 128, 16 validation jobs had return code values greater than wanted. The individual output for the job is in the listed file. Additionally, all errors are collected in the `<cvhlq>.HISTORY.ERRORS(+1)` history file. When the validation object is not explicitly listed in the error output, `<cvhlq>.HISTORY.ERRORS(+1)` can be cross-reference with the detail history file `<cvhlq>.HISTORY.RUNLIB(+1)` member `@DETL`. The `runlib` history file also contains JCL members (by validation category) for all validation jobs that are submitted, and a `@SUMM` member, which is the summary report that is shown in Example 4-22 on page 128.

A final history file, `<cvhlq>.HISTORY.BLDLIB(+1)`, serves as a backup of the `<cvhlq>.@BUILD.JOBS` library.

At the bottom of the summary report validation section, the objects that were bypassed for validation checking are noted.

Within each validation category, a total count of jobs that are submitted is given for that category. In this sample, a total of 5,224 validation jobs were run in approximately 3 minutes. In this sample, `CHECK INDEX` for the Db2 application tables was not run.

Example 4-23 shows excerpts of the summary report from which you can derive the total number validation jobs that were run.

Example 4-23   Sample of summary report

```
CATCHK   ** Total number of Catalog verify jobs = 14
CATVVD   ** Total number of BCS-VVDS verify jobs = 201
DB2CHK   ** DBW1 Total system DSN1COPY jobs = 116
DB2CHK   ** DBW1 Total system INDEX CHECK jobs = 92
DB2CHK   ** DBW1 Total application DSN1COPY jobs = 717
DB2CHK   ** DBW1 Total applic INDEX CHECK jobs = 0
DB2CHK   ** DBX1 Total system DSN1COPY jobs = 130
DB2CHK   ** DBX1 -System DSN1COPY for LOBs = 31
DB2CHK   ** DBX1 Total system INDEX CHECK jobs = 99
DB2CHK   ** DBX1 Total application DSN1COPY jobs = 3252
DB2CHK   ** DBX1 -Application DSN1COPY for LOBs = 79
DB2CHK   ** DBX1 Total applic INDEX CHECK jobs = 0
RACFCHK  ** RACF database structure validation was submitted
VSAMCHK  ** Total number of VSAM KSDS verify jobs = 492
```

> **Note:** The IBM Z Cyber Vault environment is a closed network environment. As a best practice, you can communicate the results of a validation cycle through email, so set up a collective email address that is actively monitored for the receipt of IBM Z Cyber Vault validation results emails.

The JCL and REXX code to accomplish the collection of validation information is shown in Example 4-24.

JCL for CV@COLL1

The code that is shown in Example 4-24 is run after all validation jobs are complete. The code collects all the error files that are written if validation for an individual object does not receive a good return code.

Example 4-24   JCL for CV@COLL1

```
//CV@COLL1 JOB 0001,CVCOLL1,CLASS=A,MSGCLASS=X,
//        MSGLEVEL=(1,1)
//****************************************************************
//          SET  CVHLQ=CYBERV                     <= CV HLQ
//          SET  CVVOL=ZCV001                     <= CV VOLSER
//****************************************************************
//**** THE FOLLOWING ARE STANDARD FILE NAMES USED, DO NOT CHANGE
//**** THEM AS SOME AUTO-BUILD OF JCL VIA REXX CODE WILL DEPEND ON
//**** THIS NAMING CONVENTION.
//****************************************************************
//          SET  BSUMM=&CVHLQ..CVRUN.JOB.CVSUMM    = BUILD SUMMARY
//          SET  BERRF=&CVHLQ..CVRUN.JOB.CVERRF    = BUILD ERRORS
//          SET  BERRD=&CVHLQ..CVRUN.JOB.CVERRD    = ERROR DSN FILES
//          SET  BDETL=&CVHLQ..CVRUN.JOB.CVDETL    = BUILD DETAIL
//          SET  LEGEND=&CVHLQ..@BUILD.JOBS(@JOBIX) = LEGEND FOR JOBS
//          SET  JDSN=&CVHLQ..@BUILD.JOBS(#JCK)     = JOBCARD TEMPLATE
//          SET  XLIB=&CVHLQ..@BUILD.JOBS          = EXEC LIB
//********************************************************
//* BUILD NEW HISTORY PDS GDG FOR ERROR COLLECTION
//********************************************************
//HISTERR  EXEC PGM=IEFBR14
//DD1      DD DISP=(NEW,CATLG),
//           DSN=&CVHLQ..HISTORY.ERRORS(+1),
//           UNIT=3390,SPACE=(TRK,(60,60)),
//           RECFM=FB,LRECL=132,DSORG=PS,BLKSIZE=0,
//           DSNTYPE=LIBRARY
//********************************************************
//* BUILD NEW HISTORY PDS GDG FOR CYBERVAULT RUNLIB
//********************************************************
//HISTRUNL EXEC PGM=IEFBR14
//DD1      DD DISP=(NEW,CATLG),
//           DSN=&CVHLQ..HISTORY.RUNLIB(+1),
//           UNIT=3390,SPACE=(TRK,(60,60,10)),
//           RECFM=FB,LRECL=80,DSORG=PO,BLKSIZE=0,
//           DSNTYPE=LIBRARY
//********************************************************
//* BUILD NEW HISTORY PDS GDG FOR CYBERVAULT BUILD LIBRARY
//********************************************************
//HISTBLD  EXEC PGM=IEFBR14
//DD1      DD DISP=(NEW,CATLG),
//           DSN=&CVHLQ..HISTORY.BLDLIB(+1),
//           UNIT=3390,SPACE=(TRK,(60,60,10)),
//           RECFM=FB,LRECL=80,DSORG=PO,BLKSIZE=0,
//           DSNTYPE=LIBRARY
//********************************************************
//* NOW COLLECT THE OUTPUT TO CREATE REPORTS
//********************************************************
//CVEMON2  EXEC  PGM=IKJEFT01,DYNAMNBR=20
//SYSEXEC  DD DISP=SHR,DSN=&XLIB
//SYSOUT   DD SYSOUT=*
//SYSTSPRT DD SYSOUT=*
//LEGEND   DD DISP=SHR,DSN=&LEGEND
//JOBCARD  DD DISP=SHR,DSN=&JDSN
//OUTFILE  DD DISP=(NEW,CATLG),DSN=&BSUMM,
//           UNIT=3390,SPACE=(TRK,(30,30)),
//           RECFM=FB,LRECL=80,DSORG=PS,BLKSIZE=0
//ERRFILE  DD DISP=(NEW,CATLG),DSN=&BERRF,
//           UNIT=3390,SPACE=(TRK,(30,30)),
//           RECFM=FB,LRECL=80,DSORG=PS,BLKSIZE=0
//ERRDSN   DD DISP=(NEW,CATLG),DSN=&BERRD,
//           UNIT=3390,SPACE=(TRK,(30,30)),
//           RECFM=FB,LRECL=80,DSORG=PS,BLKSIZE=0
//*=======================================================
//* WILL NEED TO CHANGE INFILE PREFIX OF DATASETS TO MATCH
//* WHAT YOU HAVE SPECIFIED FOR CVHLQ
//*=======================================================
//INFILE   DD *
CYBERV.@SUMM        SUMMARY_OF_VALIDATION_JOB_BUILD_PROCESS
CYBERV.@ERR         SUMMARY_OF_VALIDATION_JOB_BUILD_ERRORS
```

```
/*
//*=======================================================
//* SYSTSIN PARMS PASSED:
//* 1ST ARGUMENT - UNIQUE 3 CHARACTER HLQ FOR CV INFORMATION JOBS
//*                 (SHOULD MATCH WHAT YOU HAVE IN @JOBIX MEMBER)
//* 2ND ARGUMENT - THE HLQ YOU ARE USING FOR CYBER VAULT REPORTING
//*                 (SHOULD MATCH WHAT YOU HAVE SPECIFIED FOR
//*                  SET STATEMENT FOR CVHLQ IN THIS JOB)
//*=======================================================
//SYSTSIN  DD *
 %XCVECOLL CVA CYBERV
/*
//************************************************
//* SUBMIT THE SUMMARY COLLECTION JOB
//************************************************
//SUBMIT    EXEC  PGM=IEBGENER
//SYSPRINT  DD    SYSOUT=*
//SYSIN     DD    DUMMY
//SYSUT2    DD    SYSOUT=(A,INTRDR)
//SYSUT1    DD    DISP=SHR,DSN=&BSUMM
//********************************************************
//* COLLECT THE DETAIL INFORMATION OF ALL VALIDATION JOBS
//********************************************************
//CVEMON3  EXEC  PGM=IKJEFT01,DYNAMNBR=20
//SYSEXEC  DD DISP=SHR,DSN=&XLIB
//SYSOUT   DD SYSOUT=*
//SYSTSPRT DD SYSOUT=*
//JOBCARD  DD DISP=SHR,DSN=&JDSN
//OUTFILE  DD DISP=(NEW,CATLG),DSN=&BDETL,
//           UNIT=3390,SPACE=(TRK,(30,30)),
//           RECFM=FB,LRECL=80,DSORG=PS,BLKSIZE=0
//*=======================================================
//* WILL NEED TO CHANGE INFILE PREFIX OF DATASETS TO MATCH
//* WHAT YOU HAVE SPECIFIED FOR CVHLQ
//*=======================================================
//INFILE   DD *
CYBERV.@DETL        DETAIL_ON_VALIDATION_JOBS
/*
//*=======================================================
//* SYSTSIN PARMS PASSED:
//* 1ST ARGUMENT - UNIQUE 3 CHARACTER HLQ FOR CV INFORMATION JOBS
//*                 (SHOULD MATCH WHAT YOU HAVE IN @JOBIX MEMBER)
//* 2ND ARGUMENT - THE HLQ YOU ARE USING FOR CYBER VAULT REPORTING
//*                 (SHOULD MATCH WHAT YOU HAVE SPECIFIED FOR
//*                  SET STATEMENT FOR CVHLQ IN THIS JOB)
//*=======================================================
//SYSTSIN  DD *
 %XCVEDETL CVA CYBERV
/*
//************************************************
//* SUBMIT THE DETAIL INFORMATION COLLECTION JOB
//************************************************
//SUBMIT    EXEC  PGM=IEBGENER
//SYSPRINT  DD    SYSOUT=*
//SYSIN     DD    DUMMY
```

```
//SYSUT2    DD    SYSOUT=(A,INTRDR)
//SYSUT1    DD    DISP=SHR,DSN=&BDETL
```

REXX code for XCVECOLL

The code that is shown in Example 4-25 captures the details on all the jobs that are run and the object that was being validated by each job.

Example 4-25 REXX code for XCVECOLL

```
/* REXX - XCVECOLL
   This REXX collects files written by validation process
   and then submits jcl to write to summary report
   */ trace 0

/* Argument is passed on prefixes of any error output written
   from individual validation jobs */
Parse arg prefix cvhlq
job_errors=cvhlq||'.ERROR'

Say '**Starting CVECOLL at' date() time()
w=1;o=0;n=1;err_cnt=0

/* start of jobcard creation */
jobname=prefix||'#SUMM'
Makebuf
"execio * diskr jobcard ( finis stem jci."
jj=1
Do while jj <= jci.0
 if substr(jci.jj,1,3)='//*' then NOP
 else if jj=1 then queue '//'||jobname jci.jj
 else queue jci.jj
 jj=jj+1
End
"execio " queued() "diskw outfile"
Dropbuf
/* end of jobcard creation */

/* Create JCL for job to concatenate all validation information */
card.1 ='//          SET  CVHLQ='cvhlq
card.2 ='//          SET  BERRD=&CVHLQ..CVRUN.JOB.CVERRD = ERROR DSN'
card.3 ='//          SET  XLIB=&CVHLQ..@BUILD.JOBS      = EXEC LIB'
card.4 ='//CLEANUP EXEC PGM=IEFBR14'
card.5 ='//DD1    DD DISP=(MOD,DELETE),UNIT=SYSDA,SPACE=(TRK,0),'
card.6 ='//          DSN='||cvhlq||'.CVRUN.SUMMARY'
card.7 ='//**************************************************'
card.8 ='//SUMMITUP EXEC PGM=IEBGENER'
card.9 ='//SYSPRINT DD SYSOUT=*'
card.10='//SYSIN    DD DUMMY'
card.11='//SYSUT2   DD DISP=(MOD,CATLG),'
card.12='//          DSN='||cvhlq||'.CVRUN.SUMMARY,'
card.13='//         SPACE=(TRK,(30,30),RLSE),DSORG=PS,RECFM=FB,'
card.14='//         LRECL=80,BLKSIZE=0,UNIT=SYSDA '
card.15='//SYSUT1   DD DISP=SHR,DSN='||cvhlq||'.CVRUN.JOB.CVERRF'
"execio "15" diskw outfile (stem card."
ww.1='***************************************************************'
ww.2='*  Cyber Vault Validation - Jobs with exception return codes  *'
ww.3='***************************************************************'
ww.4='*  ** Note: Reported below are validations performed that did *'
ww.5='*           not meet the criteria for passing a validation    *'
ww.6='*           based on coded criteria (return code checks)      *'
ww.7='***************************************************************'
ww.8 ='CVECOLL  ** Started Exception collection at' date() time()
ww.9 ='            Any exception output will be listed in data sets'
ww.10='            noted below - a Legend is provided to indicate'
ww.11='            the type of validation being performed'
ww.12='***************************************************************'

"execio "12" diskw errfile (stem ww."
/* Read in Legend of job prefixes for validation jobs */
"execio * diskr legend (stem ll."
"execio "ll.0" diskw errfile (stem ll."

errdetl.0=0
Call CHECK_ERROR
if err_cnt = 0 then do
err.1='CVECOLL  ** No Exception files found for any validation job'
err.2='CVECOLL  ** All validation jobs that were run had acceptable'
err.3='CVECOLL  ** return codes.'
  "execio "3" diskw errfile (stem err."
End
else do
err.1='CVECOLL  ***************************************************'
err.2='CVECOLL  ** Count of jobs run with exception RCs=' err_cnt
err.3='CVECOLL  ***************************************************'
  "execio "3" diskw errfile (stem err."
end
/* Read in the list of reports to read in and summarize */
"execio * diskr infile ( finis stem in."
"execio "o" diskr infile (finis "
s=1
do while s <= in.0
  parse var in.s dslevel comment
  Say 'Reading in files' dslevel
  Call CHECK_FILES
  s=s+1
End

/* At end of summary report copy the error contents into the
   error history pds */
If err_cnt >0 then do
  i=1
  do while i <=err_cnt
    dsname=errdetl.i
    Call CREATE_ERROR_REPORT
    i=i+1
  end
  card.1='//SYSTSIN DD *'
  card.2=' %XCV@ERRH' err_cnt
  card.3='/*'
  "execio "3" diskw outfile (stem card."
end
/* We are done at this point with validation activities */
exit

CHECK_FILES:
oo=0;out.0=0
x=outtrap("out.")
address tso "listc level('"||dslevel||"')"
x=outtrap("off")
if rc <> 0 then say "**listc lvl("||dslevel||") rc=" rc
else
do oo=1 to out.0
  var1='';var2='';var3=''
  Parse var out.oo var1 var2 var3
  if var1='NONVSAM' then do
    Say '**Concatenating this file in SYSUT1' var3
    dsname=strip(var3)
    Call CAPTURE_SUMM
  End
End
Return

CHECK_ERROR:
x=outtrap("out.")
address tso "listc level('"||job_errors||"')"
x=outtrap("off")
if rc <> 0 then say "**listc lvl("||job_errors||") rc=" rc
else
do oo=1 to out.0
  Parse var out.oo var1 var2 var3
  if var1='NONVSAM' then do
    Say '**Will communicate there are errors found' var3
    dsname=strip(var3)
    Call CAPTURE_ERRORS
  End
End
Return

CAPTURE_SUMM:
card.1 ='//         DD DISP=SHR,DSN='||dsname
"execio "1" diskw outfile (stem card."
Return

CAPTURE_ERRORS:
err_cnt=err_cnt+1
card.1='Check error file' dsname
  "execio "1" diskw errfile (stem card."
card.1=dsname
  "execio "1" diskw errdsn (stem card."
errdetl.err_cnt=dsname
Return

CREATE_ERROR_REPORT:
if i=1 then do
```

```
card.1 ='//ERRCOLL  EXEC  PGM=IKJEFT01,DYNAMNBR=20'
card.2 ='//SYSEXEC  DD DISP=SHR,DSN=&XLIB'
card.3 ='//SYSOUT   DD SYSOUT=*'
card.4 ='//SYSTSPRT DD SYSOUT=*'
card.5 ='//ERRDSN   DD DISP=SHR,DSN=&BERRD'
card.6 ='//OUTFILE  DD DISP=(MOD,KEEP),'
card.7 ='//           DSN=&CVHLQ..HISTORY.ERRORS(0)'
"execio "7" diskw outfile (stem card."
end
card.1='//ERR'||i '  DD DISP=SHR,DSN='||dsname
"execio "1" diskw outfile (stem card."
Return
```

REXX code for XCVEDETL

The code that is shown in Example 4-26 captures details on all jobs that are run and the object that was being validated by each job.

Example 4-26 REXX code for XCVEDETL

```
/* REXX - XCVEDETL
   This REXX collects detail files written by validation process
   and then submits jcl to write to detail job report
   */ trace 0
Parse arg prefix cvhlq
Say '**Starting CVEDETL at' date() time()
w=1;o=0;n=1

/* start of jobcard creation */
jobname=prefix||'#DETL'
Makebuf
"execio * diskr jobcard ( finis stem jci."
jj=1
Do while jj <= jci.0
 if substr(jci.jj,1,3)='//*' then NOP
 else if jj=1 then queue '//'||jobname jci.jj
 else queue jci.jj
 jj=jj+1
End
"execio " queued() "diskw outfile"
Dropbuf
/* end of jobcard creation */

/* Create JCL for job to concatenate all validation information */
card.1 ='//CLEANUP EXEC PGM=IEFBR14'
card.2 ='//DD1    DD DISP=(MOD,DELETE),UNIT=SYSDA,SPACE=(TRK,0),'
card.3 ='//          DSN='||cvhlq||'.CVRUN.DETAILS'
card.4 ='//**************************************************'
card.5 ='//SUMMITUP EXEC PGM=IEBGENER'
card.6 ='//SYSPRINT DD SYSOUT=*'
card.7 ='//SYSIN    DD DUMMY'
card.8 ='//SYSUT2   DD DISP=(MOD,CATLG),'
card.9 ='//          DSN='||cvhlq||'.CVRUN.DETAIL,'
card.10='//         SPACE=(TRK,(30,30),RLSE),DSORG=PS,RECFM=FB,'
card.11='//         LRECL=80,BLKSIZE=0,UNIT=SYSDA '
card.12='//SYSUT1   DD DISP=SHR,DSN='||cvhlq||'.@BUILD.JOBS(@JOBIX)'
```

```
"execio "12" diskw outfile (stem card."

/* Read in the list of reports to read in and summarize */
"execio * diskr infile ( finis stem in."
"execio "o" diskr infile (finis "
s=1
do while s <= in.0
  parse var in.s dslevel comment
  Say 'Reading in files' dslevel
  Call CHECK_FILES
  s=s+1
End
exit

CHECK_FILES:
oo=0;out.0=0
x=outtrap("out.")
address tso "listc level('"||dslevel||"')"
x=outtrap("off")
if rc <> 0 then say "**listc lvl("||dslevel||") rc=" rc
else
do oo=1 to out.0
  var1='';var2='';var3=''
  Parse var out.oo var1 var2 var3
  if var1='NONVSAM' then do
    Say '**Concatenating this file in SYSUT1' var3
    dsname=strip(var3)
    Call CAPTURE_SUMM
  End
End
Return

CAPTURE_SUMM:
card.1 ='//         DD DISP=SHR,DSN='||dsname
"execio "1" diskw outfile (stem card."
Return
```

JCL for CV@COLL2

The final step in the validation cycle is to collect information for the GDG history files that contain JCL runs, copies of the created reports, and a copy of `<cvhlq>.@BUILD.JOBS` for backup purposes. The code is shown in Example 4-27.

Example 4-27 JCL for CV@COLL2

```
//CV@COLL2 JOB 0001,CVCOLL2,CLASS=A,MSGCLASS=X,
//        MSGLEVEL=(1,1)
//****************************************************************
//          SET  CVHLQ=CYBERV                     <= CV HLQ
//          SET  CVVOL=ZCV001                     <= CV VOLSER
//****************************************************************
//**** THE FOLLOWING ARE STANDARD FILE NAMES USED, DO NOT CHANGE
//**** THEM AS SOME AUTO-BUILD OF JCL VIA REXX CODE WILL DEPEND ON
//**** THIS NAMING CONVENTION.
//****************************************************************
//          SET  RUNLIB=&CVHLQ..@RUN.JOBS         <= RUN LIBRARY
//          SET  BLDLIB=&CVHLQ..@BUILD.JOBS       <= BUILD LIBRARY
//          SET  CVSUMM=&CVHLQ..CVRUN.SUMMARY     <= CV BUILD JOB SUMM
//          SET  CVERRF=&CVHLQ..CVRUN.ERRORS      <= CV BUILD JOB SUMM
//          SET  CVDETL=&CVHLQ..CVRUN.DETAIL      <= CV RUN DETAIL
//*******************************************************
//* COPY IN BLDLIB MEMBERS TO THE BLDLIB HISTORY FILE GDG
//*******************************************************
//RUNHIST  EXEC PGM=IEBCOPY
//SYSPRINT DD SYSOUT=*
//DDIN     DD DISP=SHR,DSN=&BLDLIB
//DDOUT    DD DISP=(MOD,KEEP),
//           DSN=&CVHLQ..HISTORY.BLDLIB(0)
//SYSUT3  DD UNIT=3390,SPACE=(TRK,(30,30))
//SYSIN   DD *
  COPYGRP  OUTDD=DDOUT,INDD=DDIN
/*
//*******************************************************
//* COPY IN RUNLIB MEMBERS TO THE RUNLIB HISTORY FILE GDG
//*******************************************************
//RUNHIST  EXEC PGM=IEBCOPY
//SYSPRINT DD SYSOUT=*
//DDIN     DD DISP=SHR,DSN=&RUNLIB
//DDOUT    DD DISP=(MOD,KEEP),
//           DSN=&CVHLQ..HISTORY.RUNLIB(0)
//SYSUT3  DD UNIT=3390,SPACE=(TRK,(30,30))
//SYSIN   DD *
  COPYGRP  OUTDD=DDOUT,INDD=DDIN
/*
//********************************************************
//* COPY IN THE SUMMARY REPORT TO THE RUNLIB HISTORY GDG
//********************************************************
//SUMMHIST EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSUT1   DD DISP=SHR,DSN=&CVSUMM
//SYSUT2   DD DISP=(MOD,KEEP),DSN=&CVHLQ..HISTORY.RUNLIB(0)
//SYSIN    DD *
 GENERATE MAXNAME=1
 MEMBER NAME=(@SUMM)
/*
//********************************************************
//* COPY IN THE DETAIL REPORT TO THE RUNLIB HISTORY GDG
//********************************************************
//DETLHIST EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSUT1   DD DISP=SHR,DSN=&CVDETL
//SYSUT2   DD DISP=(MOD,KEEP),DSN=&CVHLQ..HISTORY.RUNLIB(0)
//SYSIN    DD *
 GENERATE MAXNAME=1
 MEMBER NAME=(@DETL)
/*
```

4.4 Type 3 validation

Data content validation can be input to the validation process. Data content validation is specific to your application (data that is stored in data sets, databases, or other System of Record environments). For that reason, this section contains only basic concepts to consider when you construct your validation jobs.

4.4.1 Constructing data content validation jobs

Using the established framework, the jobs should follow these conventions:

- Contain a unique 3-character prefix of *XXy*, where:
  - *XX* is the standard 2-character prefix that is used by all IBM Z Cyber Vault validation jobs. In CV#MON, CV is the standard 2-character prefix.
  - *y* is a unique character that distinguishes validation jobs.
- Write the applicable output to a file that follows the output file's naming convention. In Example 4-28, the naming convention is CYBERV.ERROR.<*jobname*>.
- Pass a return code that indicates whether the check was successful:
  - If successful: The last step should delete the CYBERV.ERROR.<*jobname*>.
  - If not successful: In the final phase of validation that collects the CYBERV.ERROR.<*jobname*> files, it is noted on the validation summary report that the contents of the error file should be checked.
- Contain a step that writes a line to a summary data set that indicates the job that was submitted and a description so that it will be noted on the validation summary report. In Example 4-28, the summary data sets have the naming convention of the summary data set, which is CYBERV.@SUMM.<*jobname*>.

Example 4-28 Writing a line to a summary data set

```
//**************************************************************
//* Write summary line to indicate RACF DB check done
//**************************************************************
//SUMMRPT EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSUT1   DD *
CVERACF  ** RACF database structure validation was submitted
***************************************************************
/*
//SYSUT2   DD DISP=(NEW,CATLG),
//         DSN=CYBERV.@SUMM.CVERACF
//         SPACE=(TRK,(1,1),RLSE),DSORG=PS,RECFM=FB,
//         LRECL=80,BLKSIZE=0,UNIT=&CVUNIT
//SYSIN    DD DUMMY
```

4.4.2 Automating data content validation

The jobs that are submitted can be in the build library member, which contains a list of Type 3 data-content validation jobs to be submitted by the IBM Z Cyber Vault monitoring job (CV#MON) after the Type 2 data-structure validation jobs are complete.

`@JOBIN2` contains the Type 3: Data Content Validation Jobs (for specific application validation jobs that will be inserted into the IBM Z Cyber Vault validation cycle). `@JOBIN2` might point to one or multiple data-content validation jobs. In our sample, which is shown in Example 4-28 on page 140, specific data content validation jobs are not built. Therefore, a dummy job (that basically does nothing) is submitted to satisfy the Type 3 requirement.

4.4.3 Collecting feedback from data content validation

According to the naming convention that is used for output files, the feedback files are automatically noted on the validation summary report:

- `<cvhlq>.@SUMM.<`*`jobname`*`>`

  Contains one line with the job name and a description of validation.

- `<cvhlq>.ERROR.<`*`jobname`*`>`

  Contains multiple lines with the output of validation or a job that did not result in a satisfactory return code.

Abbreviations and acronyms

BC	business continuity
BCPii	Base Control Program internal interface
BCS	basic catalog structure
CA	certificate authority
CDS	couple data set
CF	coupling facility
CG	consistency group
CICS	Customer Information Control System
CICS VR	CICS VSAM Recovery
CKD	Count Key Data
CS	Copy Services
CSM	Copy Services Manager
DBA	database administrator
DR	disaster recovery
DS CLI	DS Command-line Interface
DVE	Dynamic Volume Expansion
ESDS	entry sequenced data set
ELB	Extended Long Busy
ESE	extent space efficient
ESM	external security manager
FB	Fixed Block
GC	Global Copy
GDG	generation data group
GDPS	Geographically Dispersed Parallel Sysplex
GDPS MzGM	GDPS Metro Global - XRC
GM	Global Mirror
HA	high availability
HADR	high availability and disaster recovery
HLQ	high-level qualifier
HMC	Hardware Management Console
IBM	International Business Machines Corporation
ICF	integrated catalog facility
IDCAMS	Integrated Data Cluster Access Method Services
IDS	Intrusion Detection Services
ILK	indirect list key
IMS	Information Management System
IODF	input/output definition file
IPL	initial program load
IRLM	Internal Resource Lock Manager
ISO	International Organization for Standardization
ISPF	interactive system productivity facility
JES2	Job Entry Subsystem 2
JES	Job Entry Subsystem
JCL	Job Control Language
KSDS	key sequenced data set
LCM	License Compliance Manager
LCP	logical corruption protection
LCU	logical control unit
LOB	line of business
LPAR	logical partition
LSS	logical subsystem
MFA	Multi-Factor Authentication
MGM	Metro/Global Mirror
MM	Metro Mirror
MSS	multiple subchannel set
MTMM	Multi-Target Metro Mirror
OSA	Open Systems Adapter
PDSE	partitioned data set extended
PoC	proof of concept
QHA	Query Host Access
RC	recovery copy
RLS	record-level sharing
RMF	Resource Measurement Facility
RPO	recovery point objective
RPQ	request price quotation
RRDS	relative record data set
RS	replication site
RTO	recovery time objective
SDSF	System Display and Search Facility
SHCDS	sharing control data sets
SLA	service level agreement
SMF	Systems Management Facility
SMTP	Simple Mail Transfer Protocol
SOAR	security orchestration, automation and response

TRL	technology readiness level
UIT	user impact time
VLAN	virtual LAN
VSAM	Virtual Storage Access Method
VTOC	volume table of contents
VVDS	VSAM volume data set
WLM	Workload Manager
WORM	write once read many
WTOR	write-to-operator-with-reply
zCX	z/OS Container Extensions
zIIP	z Integrated Information Processors

Redbooks

Getting Started with IBM Z Cyber Vault

(0.2”spine)
0.17”<->0.473”
90<->249 pages

SG24-8511-00

ISBN 0738460117

Printed in U.S.A.