Redbooks
ibm.com/redbooks

Getting Started with IBM Z Cyber Vault

Bill White
Dino Amarini
Diego Bessone
Tom Bish
Nathan Brice
Richard Cairns
Giovanni Cerquone
Nick Clayton
Michael Frankenberg
Nathan Gurley
Maryellen Kliethermes
David Mateo
Kevin Miner
Nadim Shehab
Karen Smolar
Paolo Vitali
Joseph Welsh II

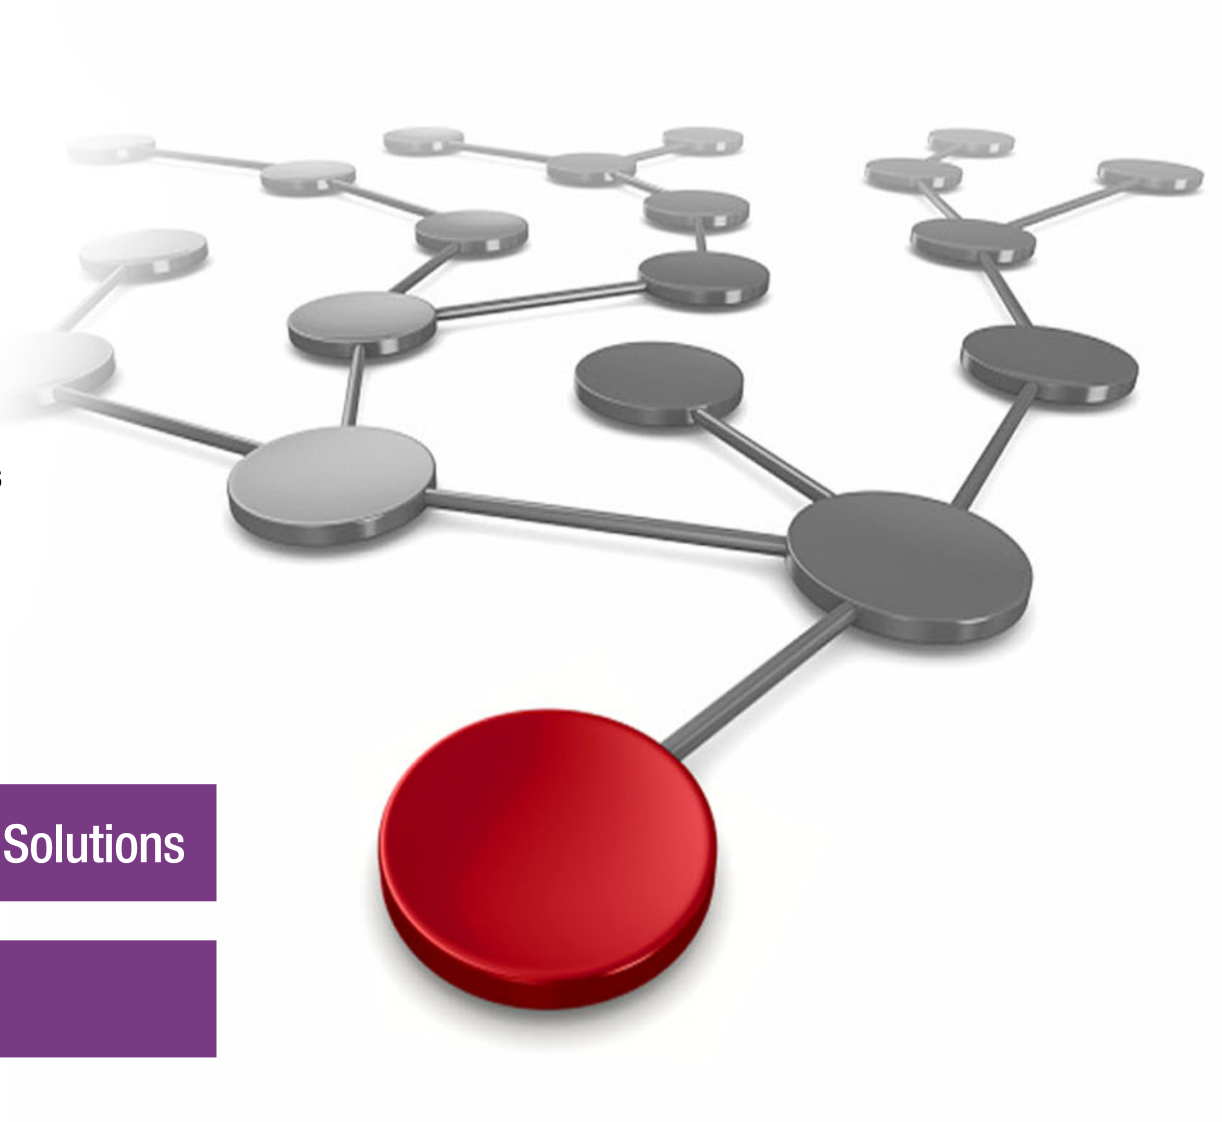

Infrastructure Solutions

IBM Z

IBM Redbooks

IBM Redbooks

Getting Started with IBM Z Cyber Vault

December 2024

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

Second Edition (December 2024)

This edition applies to the IBM Z Cyber Vault solution.

This document was updated on June 4, 2025.

Contents

Notices

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

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

Trademarks

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

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

AIX®
CICS®
Db2®
DS8000®
FICON®
FlashCopy®
GDPS®
HyperSwap®
IBM®
IBM Security®
IBM Z®
Parallel Sysplex®
RACF®
Redbooks®
Redbooks (logo) ®
S/390®
Tivoli®
VTAM®
z/Architecture®
z/OS®

The following terms are trademarks of other companies:

ITIL is a Registered Trade Mark of AXELOS Limited.

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

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

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

Preface

With cyberattacks on the rise, cyber resiliency is becoming more important than ever. Cyber resiliency provides the required capability to prevent significant impact in the event of an attack. A cybersecurity strategy might minimize the risk of attacks getting through to systems, applications, and data, but a cyber resiliency strategy is needed to recover 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. It 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 examples for the deployment of the IBM Z Cyber Vault solution are also included, and a suggested framework with advice for conducting basic data validation, analysis, and recovery.

This publication is intended for IT 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 the IBM Redbooks team, Poughkeepsie Center.

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

Dino Amarini

Diego Bessone is a Global Software Sales Director for IBM Z. He got immersed in the mainframe world in 1987 when he started working at the data center of Aerolíneas Argentinas, the flag carrier of Argentina. After 10 years he continued leveraging his experience working as a technical consultant for the IT departments of multinational companies, finance, utility, and government agencies. Diego joined IBM in 1998 in Argentina, leading the sales strategy for IBM mainframe systems management portfolio, then moving to the USA in 2001 to lead the Tivoli® S/390® Latin America team, adding to his role business operations management, and then leading IBM Z middleware software sales at a worldwide level since 2009. In July 2022 Diego was promoted to an Executive role as IBM Z Sales Director, managing global sales of IBM Z high-end platforms, and since 2024 leading the IBM Z Software sales strategy for pricing, as well as strategic initiatives including IBM Z Cyber Vault, IBM Z Cyber Security, and future IBM Z programs. He has written extensively on data management, systems management, mission critical systems, and resiliency.

Tom Bish has been with IBM for 34 years and is a IBM Senior Technical Staff Member, Principal Architect, and Master Inventor. Tom was in IBM storage product development and architecture for 25 years, and then joined IBM Global Technology Services to define their software define storage strategy. Tom is currently within IBM Advanced Technology Group (ATG) within technical sales to help clients and sellers best utilize IBM storage technologies.

Nathan Brice is the GDPS® and IBM Z Cyber Vault Product Manager. He is based in the UK and has worked for IBM in many roles for over 25 years. Over the last 10 years, he has held Product Manager roles in GDPS, CICS®, and - on assignment in North Carolina - in the IBM AIOps team. During this time he has presented at conferences around the world and written many articles on IBM products and technology.

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

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

Nick Clayton currently works in IBM Infrastructure development as the Solution Architect for the DS8000® storage system focusing on solution integration and product strategy. His specific interests include storage performance, replication technology and business resilience and he is a member of the GDPS design team. He also works with clients on their storage strategy and advising them on their deployment of IBM storage technology. Nick is an author and co-author of many patents, articles, whitepapers and IBM Redbooks publications on IBM storage technology. He is a regular presenter at storage conferences. Nick graduated from Trinity College Cambridge in 1994 with a degree in Mathematics and has had previous roles both within and outside IBM in the areas of parallel sysplex, performance, and enterprise storage.

Michael Frankenberg is a Certified IT Specialist in Germany with more than 25 years of experience in high-end storage. He currently works in IBM Technical Sales Support, EMEA and his area of expertise includes performance analysis, establishing high availability and disaster recovery solutions and the implementation of IBM storage systems. Michael supports the introduction of new IBM storage products and provides advice to clients, business partners, and IBM Technical Sales personnel. He holds a degree in Electrical Engineering / Information Technology from University of Applied Sciences Bochum, Germany.

Nathan Gurley is a GDPS developer in the United States. He has 5 years of experience in the disaster recovery and corruption protection field. He holds a degree in Computer Science from Miami University. His areas of expertise include GDPS - Logical Corruption Protection and security on z/OS. He has presented at conferences in the U.S. and Europe about role-based and dual control security within GDPS and enhancing the capabilities of the IBM Z Cyber Vault solution.

Maryellen Kliethermes is a Senior Project Manager in the United States with over 30 years of experience in Information Technology. She has held roles in the Federal Government and Utility industry as a developer, Manager of IT Software Development, and Manager of the IT Project Management Office. She holds a Bachelors of Science in Computer Information Systems, a Masters in International Business from St. Louis University, and holds certifications in PMP, Agile and ITIL. Most recently, she has been Lead Project Manager at IBM for deploying the IBM Z Cyber Vault solution and IBM Z pervasive encryption solutions at clients worldwide.

David Mateo

Kevin Miner

Nadim Shehab is a Senior GDPS developer in United States and has 20 years of experience in the disaster recovery and corruption protection fields. His areas of expertise include development of GDPS - Logical Corruption Protection and Copy Services. Nadim has presented at conferences around the world regarding Logical Corruption Protection and Copy Services. He holds a Bachelors of Science in Computer Engineering from The University of Arizona.

Karen Smolar is a Principal Solution Architect in the United States. She has 36 years of experience in the information technology field. She holds a degree in Computer Science from Clarkson University. Her areas of expertise include helping clients achieve success by designing infrastructure and application solutions that meet the demands of their business. She has written extensively on all aspects of resiliency and business continuity planning.

Paolo Vitali

Joseph Welsh II is an IBM Senior Management IT Consultant and certified IT Network Specialist in the USA. He joined IBM in 1988 with a Bachelors degree in Computer Science from Transylvania University. Joe has held the roles of software developer, designer, and tester for IBM z/OS Communications Server (VTAM®, TCP/IP). Since 1998, he has performed IT consulting services engagements focused on SNA, Advanced Peer-to-Peer Networking (APPN) and high-performance routing (HPR), Enterprise Extender, VTAM, TCP/IP, and Networking security for the Communications Server for (IBM AIX®, Linux, Microsoft Windows, and z/OS) at Fortune 500 companies around the world. This includes providing SNA, APPN and HPR, TCP/IP, and IP Security education and training, developing network designs, and migrations, strategy and product direction, problem determination, implementation, and installation and migration assistance. In the most recent years, Joe has been providing IBM Z Networking Security consulting services and deploying IBM Z Cyber Vault solutions to clients worldwide.

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

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

Dave Clitherow
IBM GDPS Product Manager (Retired)

Dave Petersen
IBM Distinguished Engineer

Vijay Radhakrishnan
IBM GDPS Solution Architect

Thanks to the authors of the first edition, “Getting Started with IBM Z Cyber Vault”, published in July 2022:

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

Now you can become a published author, too!

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

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

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:

  https://www.linkedin.com/groups/2130806

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

  https://www.redbooks.ibm.com/subscribe

- Stay current on recent Redbooks publications with RSS Feeds:

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

1

Business resiliency: Proactive analysis and expedited recovery

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

Victims of data corruption typically suffer in the following ways:

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

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

The following topics are covered in this chapter:

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

1.1 Common data corruption threats and remediation

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

Malware 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.

Ransomware and malware attacks are typically launched by external entities who have successfully attacked a business by using phishing, identity theft, or other scams. The most common attacks start with an e-mail with infected attachments, smartphone message, or multimedia. Another common attack vector is infected websites exploiting unpatched browser vulnerabilities, which, although less common than email attacks, is still a successful way to gain entry to IT environments.

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

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

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

- Computer systems and infrastructure itself, such as operating systems, clustering technology, storage systems, and disaster recovery (DR) replication
- Application middleware and runtimes, database servers, and core data file systems
- Application and business data that is stored in databases and file systems

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

Another example for unintentional data corruption is when bad logic is introduced into an application code update, sometimes exacerbated by subsequent patches to fix this problem that actually end up leading to further data loss. Alternatively simple human error (either in commands or processes) or inadequate testing and backup practices, can also be factors in this type of data errors.

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

Another increasingly common attack vector involves accessing systems with stolen credentials. Sophisticated, targeted “Spear phishing” attacks on individuals are becoming far

[1] Jointly produced by the Ponemon Institute and IBM Security®.

more common. The *IBM X-Force Threat Intelligence Index report* finds that there has been a 71% year-over-year increase in cyberattacks that used stolen or compromised credentials.

As the IBM Z platform is a critical part of the data infrastructure for many organizations, particularly in industries like finance, healthcare, and insurance, it is now more than ever the focus of malicious attacks and intents to steal or corrupt data. With just a simple online search is very easy to find publicly available exploits that could surface unwanted exposures and take advantage of insufficient security practices. While IBM Z may be the most securable platform, data never rests or stays in one place. IBM Z is still a server, and no server is immune to attacks; they are all at risk. While it can be argued that the likelihood of a breach occurring on IBM Z is lower than other types of servers, it can similarly be argued that the impact of a breach on an IBM Z platform is much higher than other types of servers (due to the critical role IBM Z platforms play). As a result, the risk is still a real and significant concern.

1.1.1 Main factors that determine favorable outcomes

Detailed planning and preparation must be given to evaluating the impact on business reputation, implications of regulatory compliance, and the real 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 *Cost of a Data Breach Report* also found the following actions reduced the financial and brand impacts of a data breach while also reducing the time to detect and respond, and recover from the breach:

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

These recommendations build on what are already very well-established guidelines:

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

A perfect cybersecurity solution does not exist. However, a well thought-out, architected and tested cybersecurity strategy can minimize the risk of attacks. Perhaps most importantly, a comprehensive 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 should always includes the following elements:

- **Identification**: focuses on the need of an enterprise to understand their most important assets and resources. This includes 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**: 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**: focuses on measures that alert an organization to cyberattacks, which might include anomalies and events, continuous security monitoring, and early detection processes.
- **Incident response**: ensures an appropriate response to cyberattacks and other cybersecurity events. Categories include response planning, communications, analysis, mitigation, and improvements.
- **Recovery**: covers the 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.

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

There are also several frameworks available that can be utilized to improve readiness while building and enhancing your cyber resiliency strategy.

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*. Another is from the Digital Operational Resilience Act (DORA), called the *ICT risk management framework*, which should be evaluated by those doing business in the European Union.

Complementary *International Organization for Standardization (ISO)* documents ISO 31000 and ISO 27005 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.

1.2 What is cyber resiliency?

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

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

1.2.1 Traditional resiliency vs. cyber resiliency

A key difference between designing and deploying a traditional resiliency solution versus a cyber resiliency solution is that a traditional resiliency solution must protect against 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.

For example, a traditional resiliency solution aims to provide protection against power failure or other physical issues. Examples of these include a weather event such as hurricane that causes localized flooding and takes a production data center offline. In all cases, however, the primary goal of traditional resiliency solutions is to manage multiple copies of data as consistent to the production data as possible. That may be via a synchronized write to a second copy of data with 100% consistency or via an asynchronous write where the secondary copy of data is several seconds behind the production copy.

In the event of a cyberattack, any corruption, encryption or erasure of data in production will get rapidly replicated out to other copies of data and corrupt those additional copies of data. So, for cyber resiliency, what is required is protection again logical data corruption rather than physical issues. This is why a cyber resiliency solution takes a fundamentally different approach to a traditional resiliency solution and must compliment your resiliency solution rather than be a replacement.

Another key differentiator is how a decision to execute a recovery procedure is taken. In an infrastructure related incident, the system operations team has clear instructions and procedures in place to revert to a specific point in time or move the production site to an alternate location depending on the issue flagged by their monitoring systems. In fact, most of these procedures are usually fully automated since they are addressing known situations that have to do with systems or infrastructure malfunctions. However, when a data corruption event happens, since the nature and extent of the incident is not known, the operations team needs to work with the applications and line of business teams to determine which is the best approach to recovering corrupted data. This requires additional analysis and resources that need to be run in a clean-room environment, with enough and readily available capacity and tooling to help throughout the analysis and recovery processes.

1.2.2 Cyber security verses cyber resiliency

It is important to understand the difference between cyber security and cyber resiliency. The best way to think about this is that cyber security is trying to prevent a cyberattack from happening, while cyber resiliency is about minimizing the impact a cyberattack or data breach could have on service delivery.

In today's world, a hardened perimeter network with physical security is no longer an option. Having all employees access a company network only from a physical terminal in a secure office building is not viable. Working from home and mobile access is critical. To overcome this, many businesses and organizations are adopting a zero trust approach with three core principles for cyber security:

- **Continuous monitoring and validation**

  Zero trust makes all network assets inaccessible by default. Users, devices and workloads must pass continuous, contextual authentication and validation to access any resources, and they must pass these checks every time they request a connection.

  Dynamic access control policies determine whether to approve requests based on data points such as a user's privileges, physical location, device health status, threat intelligence and unusual behavior. Connections are continuously monitored and must be periodically reauthenticated to continue the session.

- **The principle of least privilege**

  In a zero trust environment, users and devices have least-privilege access to resources. This means they receive the minimum level of permission required to complete a task or fulfill their role. Those permissions are revoked when the session is over.

  Managing permissions in this way limits the ability of threat actors to gain access to other areas of the network.

- **Assume breach**

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

However, even with the best cyber security protection in place it is impossible to prevent 100% of all potential attacks. A successful cyber resiliency strategy considers how the impact of a successful cyberattack or data breach can be minimized.

For example, after a ransomware attack, a business would likely wish to have the capability to restart their production environment from a recent, known good state to get back online as quickly as possible. This capability can help avoid paying a ransom to get data unlocked – even assuming the decryption capability works as expected. In some cases, even after paying to receive the decryption key, a company may not get fully up and running for a longer period of time after the attack.

In the subsequent sections we look closely at capabilities of an ideal cyber resiliency strategy and IBM's solution, called IBM Z Cyber Vault.

1.3 Capabilities for an effective cyber resiliency strategy

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

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

Full cyber resiliency requires intrusion detection, monitoring for unusual behavior by individuals, programs, and systems, and reporting and dashboards to alert teams 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. 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 cyber security and cyber resiliency.

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

1.3.1 Characteristics of an ideal cyber resiliency solution

An ideal cyber resiliency solution must provide the ability to not only protect against the unique challenges of a cyberattack, but also provide all the necessary capacity, tooling and resources required to address it.

The most critical characteristic is the ability to take regular, immutable copies of the production environment, that cannot be corrupted or erased. These copies should be stored in a secured location, separated from production, with access governed separately from the production environment. By taking regular point in time copies, as opposed to continuous replication, it will always be possible to go back to a ‘good’ copy of data that was captured before the corruption event.

The second key characteristic is that the system processes should be fully automated and not require any manual intervention. The goal is to take backups as regularly and often as possible, to reduce the window of time between backups. Only when this process is fully automated will it be possible to bring the interval between backups down to hours instead of days.

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

1.3.2 Capture, analysis, recovery and restore characteristics

In the event of a data breach, this capability will enable the IT staff to recover a specific point-in-time backup into the vault environment to enable the analysis required to assess the incident. This will consist of formulating optimal recovery strategies and options; and determine the scope of recovery, files, databases, or even entire systems.

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

This capability will provide the applications and line-of-business teams all the tools required to make the best decision to restore their business to the best possible state after a serious cyber event.

1.3.3 Data validation characteristics

On a continuous basis, automation should be in place to recover, start, and validate these backups in advance, in the vault environment. The benefits to proactive, automated validations are:

- Provide confidence that if a specific backup is needed, it has already been validated as “good”, and is usable. In the event of any data breach, there will be tremendous pressure to restore production as quickly as possible. If the backup candidate has already been validated, then the IT staff is able to focus on the recovery strategy without the need for further tests to validate that this specific backup is going to be a suitable candidate for the restore into production.
- Secondly, by automating and regularly validating backups, any failure of validation may provide an early warning that production has been corrupted. This corruption may not have been detected in production yet, but a validation failure could be a trigger for further investigation and could potentially reduce the time taken to identify certain types of data breaches.

1.3.4 Backup characteristics

It is important to understand that there is a tradeoff between the capture interval and the retention duration for point-in-time backups. Ideally, the capture rate should be as low as possible, with the retention duration as long as possible. However, moving these numbers in different directions results in exponential growth in storage requirements. To address this tradeoff, different types of storage may be used across the solution. Ideally, disk should be used for the low interval initial captures. This capture rate should be in the low single digit hours range for example, every hour to four hours. Then each backup is retained on disk for between seven and 28 days to ensure that if needed, these backups can be rapidly restored.

However, it is entirely possible that the system was compromised further back in time, and it may be necessary to go back to a backup that is several months old. This can be achieved by archiving older backups on a slower access, but cheaper storage, such as tape. By combining disk, virtual and physical tapes devices, it is possible to get the best of both worlds with rapid access for the most commonly needed data, while allowing a long duration retention for infrequently needed data.

1.4 IBM Z Cyber Vault solution and key capabilities

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

Anticipating and responding to an event is increasingly complex and the speed and precision of response is increasingly critical. Figure 1-1 on page 9 show the areas of concern when dealing with an event.

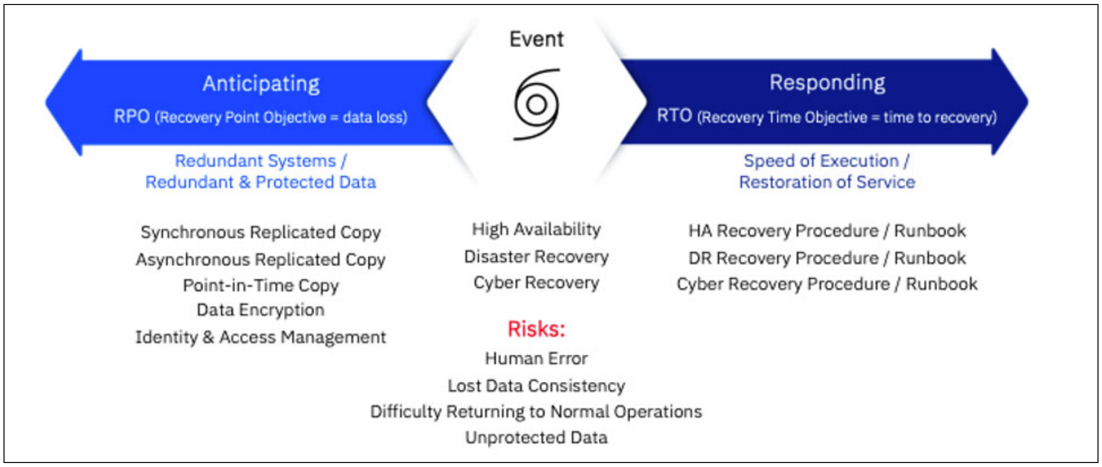
Event
Anticipating
RPO (Recovery Point Objective = data loss)
Responding
RTO (Recovery Time Objective = time to recovery)
Redundant Systems / Redundant & Protected Data
Speed of Execution / Restoration of Service
Synchronous Replicated Copy
Asynchronous Replicated Copy
Point-in-Time Copy
Data Encryption
Identity & Access Management
High Availability
Disaster Recovery
Cyber Recovery
HA Recovery Procedure / Runbook
DR Recovery Procedure / Runbook
Cyber Recovery Procedure / Runbook
Risks:
Human Error
Lost Data Consistency
Difficulty Returning to Normal Operations
Unprotected Data

Figure 1-1 Anticipating and responding to an event

IBM learned a great deal collaborating with businesses and organizations around the world, across every industry, and leveraging best practices for systems and data center resiliency. Through that experience, IBM has identified three areas or domains that any cyber resiliency solution should entail. Only when addressing these three domains with the right technology, will business services not only be better protected, but also capable of being restored to a point in time that makes it possible to restart business services quicker and with the right data.

These three domains can also be considered as entry-points to a full implementation of the IBM Z Cyber Vault solution. You can decide where to start, while putting together a deployment roadmap that will ultimately implement all the provided capabilities.

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

- **IBM Z Cyber Vault Storage**

  IBM Storage DS8000 is the latest innovation in enterprise-class storage for IBM Z. It is designed to ensure the availability of critical business workloads and lowering business risk of outages, while safeguarding sensitive data, meeting regulatory compliance requirements with Advanced Security Features, like enhanced encryption, access controls, and data protection measures.

  As the cornerstone of the IBM Z Cyber Vault solution, it prevents your data from being modified or deleted due to user errors, malicious destruction, malware or ransomware attacks with hundreds of immutable copies per volume, which can be used as a trusted source for surgical or full recovery of a production environment.

The Safeguarded Copy function supports the ability to create cyber resilient point-in-time copies of volumes that cannot be changed or deleted through user errors, malicious actions, or ransomware attacks. The DS8000 system integrates with IBM Copy Services Manager to provide automated backup copies and data recovery.

- **IBM Z Cyber Vault Automation**

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

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

  IBM Technology Expert Labs has developed special assets leveraging GDPS and LCP, that drive data validation in the IBM Z Cyber Vault environment.

- **IBM Z Cyber Vault Environment**

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

  This isolated (also known as air-gapped) environment can be physically or virtually separated from other production, development, and test environments, implementing network and security safeguards that will maintain the nature of the clean room.

  The IBM Z Cyber Vault Environment needs to be sized according to the amount of data that will be managed, and must include all the following to execute:

  – IBM Z Cyber Vault Environment Licensing (5770-ZCV) - the full IBM Z Software Production Stack licensed to be executed in the CV environment

  – IBM Z Software Tools that are expected to be used for validation, analysis and/or recovery. These tools will also be installed and available in the production environment as they typically need to collect data that will be used afterwards in the vault

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

1.4.1 IBM Z full stack resiliency

Designed for continuous availability and rapid disaster recovery, IBM Z provides industry-leading resiliency to protect the business from downtime (see Figure 1-2 on page 11). Implementing an IBM Z Cyber Vault solution will help optimize availability, keep systems running, detect problems in advance and recover critical data.

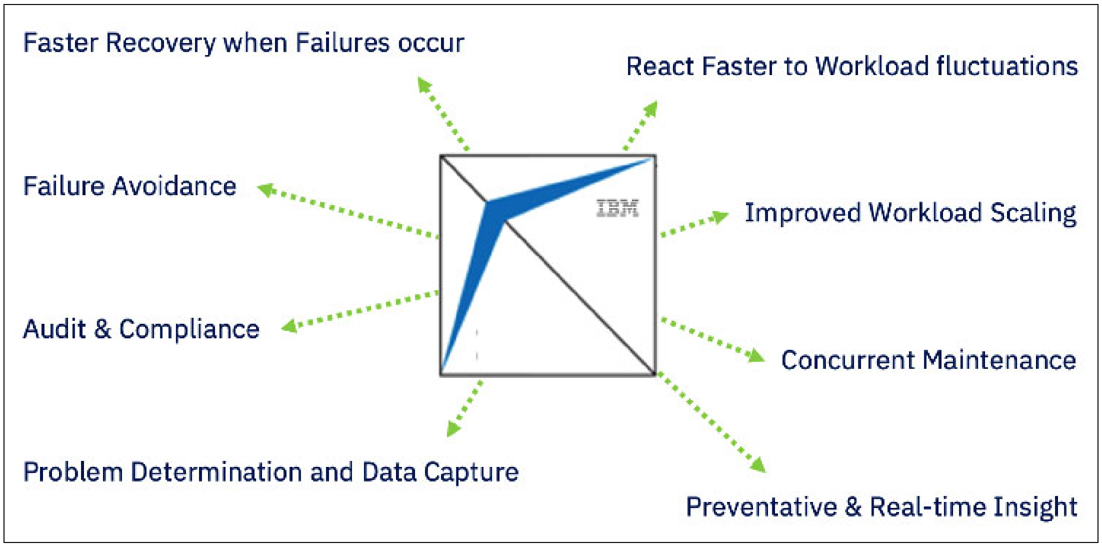
Faster Recovery when Failures occur
React Faster to Workload fluctuations
Failure Avoidance
IBM
Improved Workload Scaling
Audit & Compliance
Concurrent Maintenance
Problem Determination and Data Capture
Preventative & Real-time Insight

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

1.4.2 IBM Z Cyber Vault capabilities

By implementing the IBM Z Cyber Vault solution, five key capabilities (see Figure 1-3) that will allow your business to be ready and prepared to work through any type of data corruption event are delivered.

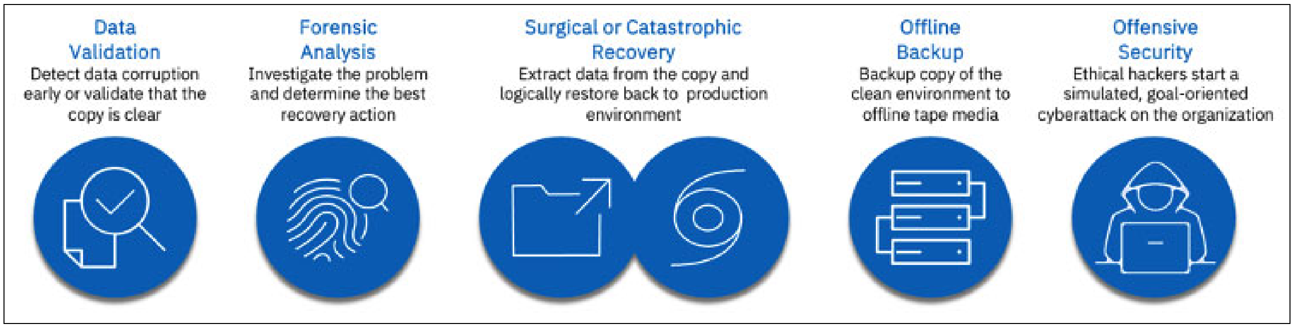
Data Validation
Detect data corruption early or validate that the copy is clear
Forensic Analysis
Investigate the problem and determine the best recovery action
Surgical or Catastrophic Recovery
Extract data from the copy and logically restore back to production environment
Offline Backup
Backup copy of the clean environment to offline tape media
Offensive Security
Ethical hackers start a simulated, goal-oriented cyberattack on the organization

Figure 1-3 IBM Z Cyber Vault - key capabilities

Data validation

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

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

- **Type 1**: Infrastructure validation, which consists of recovering a point-in-time copy of the production environment data into the IBM Z Cyber Vault Environment and starting the system and its subsystems from this recovered copy. If the system starts correctly and resolves all inconsistencies when restarting the middleware (transactions in flight and so on), a working replica of production, which means that the infrastructure to run the business applications is in good shape.

- **Type 2**: Data structure validation, where data structures are actively examine to identify inconsistencies or errors, often requiring specialized software to analyze data integrity. Some of the data structures and software solutions that help with validation are:
  - Db2, using Db2 Utilities, CHECK DATA/INDEX and Log analysis.
  - IMS, using IMS Utilities such as Pointer Checker.
  - z/OS Catalog, using IDCAMS and system management tools from IBM or other software vendors
  - VSAM, using Indexcheck and Datacheck
  - DFSMShsm Control Data Sets and DFSMSrmm catalog, using z/OS and other IBM tools
  - RACF (IRRUT200) database, using IBM zSecure Audit
  - ISV software
- **Type 3:** Data content validation, where application data is analyzed to see if the actual business data makes sense. This can be accomplished by special validation programs and procedures that need to be developed by the application development teams, as they are the ones that understand the business data, and can determine if it is valid or not. The IBM Z Cyber Vault solution provides the framework to automate the execution and verification of these validation processes.

Forensic analysis

Forensic analysis (after a data corruption event) involves using specialized software to examine a corrupted storage device, database or file, in order to put together a plan to recover as much usable data as possible while meticulously documenting the process to preserve evidence, even if the data is partially damaged or overwritten, with the goal of identifying the cause of corruption and potentially recovering valuable information.

Surgical or catastrophic recovery

Once the nature and extent of the data corruption event is understood, and a plan to recover lost data is determined, it is time to start the recovery process. In most cases it should be a surgical recovery, which would allow the user to restore into production just the affected data, usually by recovering data from one or more clean backups, and re-running some processes to synchronize data across applications to achieve an application consistency point. Transactions forward recovery is also part of this process, provided all the required logs and data is available and free of data corruption.

In a worst-case scenario, a catastrophic recovery would be required, which involves taking the most recent, already validated and clean Safeguarded copy, and completely restoring it into the production environment.

Offline backup

Offline backups add a second layer of protection. Once a recovered Safeguarded backup is validated and determined to be clean, it can be stored in a completely isolated, offline environment, typically physically separated from any other environment to protect it from any type of attack or cyber threats, ensuring a clean and uncompromised data recovery point in case of a major security breach. Magnetic tapes serve this function perfectly, by storing the validated environment in an offline, physically isolated manner, typically stored in a secure offsite facility, creating an “air gap” that protects the data from cyber threats by making it inaccessible to online systems, essentially acting as a last line of defense against cyberattacks by keeping backups completely disconnected from the network.

Offensive security

Offensive security (or OffSec) uses adversarial tactics, the same tactics that bad actors use in real-world at-tacks to strengthen network security rather than compromise it. Offensive security is conducted typically by ethical hackers, cybersecurity professionals who use hacking skills to detect and fix not only IT system flaws, but security risks and vulnerabilities in the way users respond to attacks.

Offensive security measures that can help strengthen insider threat programs include phishing simulations and red teaming, in which a team of ethical hackers start a simulated, goal-oriented cyberattack on the organization.

The IBM Z Cyber Vault environment not only provides an exact replica of the production environment to perform these exercises, but also because it is completely isolated there is no risk of harming any other part of the system.

For more information about using the IBM Z Cyber Vault capabilities, refer to Chapter 3, “IBM Z Cyber Vault capabilities” on page 49.

2

Planning and designing the IBM Z Cyber Vault solution

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 solution.

The following topics are covered:

2.1 Planning approach for the IBM Z Cyber Vault solution

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

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

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

Terminology: Throughout this document, we refer to LPARs, sysplex, and Parallel Sysplex. The environment(s) you are trying to protect and the related recovery environment(s) could be any of these. Since it is more common to have sysplexes (or Parallel Sysplexes), much of this document uses those terms. If that does not apply to your environment, you can substitute LPAR(s) in most cases.

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

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

Once you understand these, you will be ready to define your IBM Z Cyber Vault solution architecture. IBM can help you get through this process. Reach out to your IBM team and ask about the “IBM Z Cyber Vault Discovery and Architecture Workshop”. The workshop will help you with the initial planning, define the architecture, and provide an implementation roadmap.

2.1.1 Defining the architecture

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

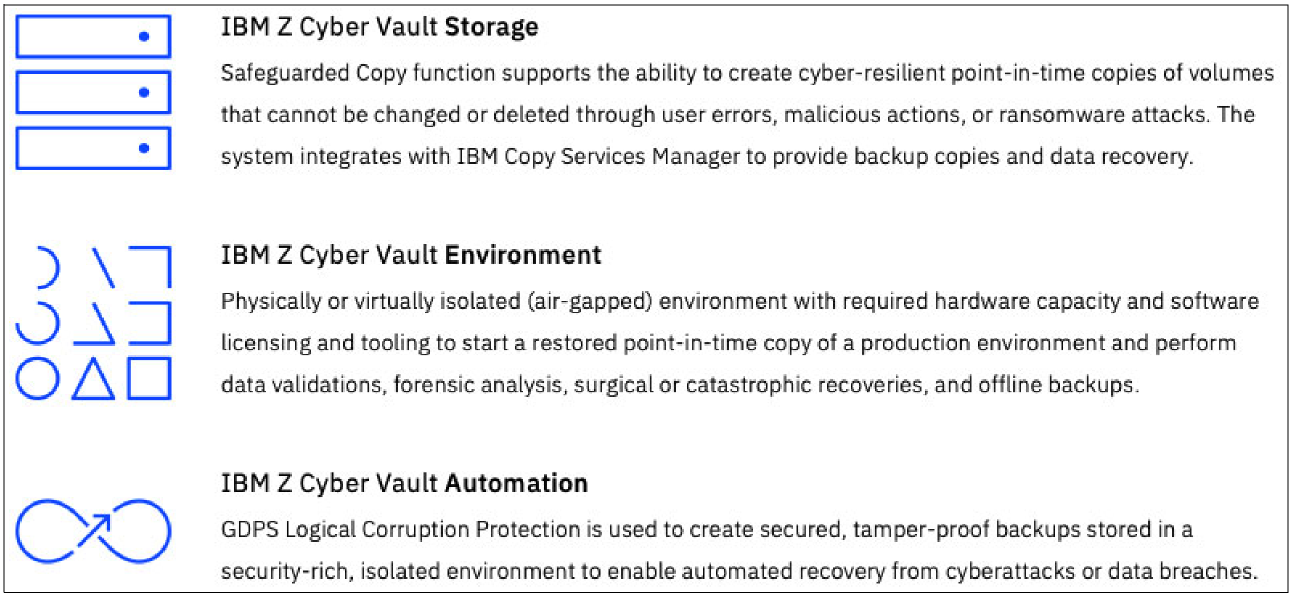

IBM Z Cyber Vault **Storage**

Safeguarded Copy function supports the ability to create cyber-resilient point-in-time copies of volumes that cannot be changed or deleted through user errors, malicious actions, or ransomware attacks. The system integrates with IBM Copy Services Manager to provide backup copies and data recovery.

IBM Z Cyber Vault **Environment**

Physically or virtually isolated (air-gapped) environment with required hardware capacity and software licensing and tooling to start a restored point-in-time copy of a production environment and perform data validations, forensic analysis, surgical or catastrophic recoveries, and offline backups.

IBM Z Cyber Vault **Automation**

GDPS Logical Corruption Protection is used to create secured, tamper-proof backups stored in a security-rich, isolated environment to enable automated recovery from cyberattacks or data breaches.

Figure 2-1 IBM Z Cyber Vault solution framework

The details in the rest of this chapter will provide the information and guidance needed to make these architectural decisions:

- Where will your IBM Z Cyber Vault solution be located? Will it be one of your existing data centers or a new location?
- Will you implement one of the virtual isolation topologies and leverage existing storage devices or opt for physical isolation on dedicated storage devices?
- Will you use GDPS LCP Manager or CSM for storage management? Will you use GDPS LCP Manager for automation?
- Will you use GDPS LCP Manager for validation and recovery automation or will you develop your own automation for these?
- Will your IBM Z Cyber Vault Environment be on IBM Z hardware that is also used for other workload or will it be dedicated IBM Z Cyber Vault processing?
- Will you have access to any virtual tape data from your IBM Z Cyber Vault Environment? If so, how will that data be accessed and protected?

2.1.2 Implementation considerations

Once your architecture has been defined, you will know where your IBM Z Cyber Vault solution will be, the source of the backups, whether the storage topology will be virtually or physically isolated, and whether the IBM Z Environment will be on dedicated hardware or share existing hardware. Then it will be time to talk about what it will take to build, test, and operate it.

Considerations for IBM Z Cyber Vault Storage implementation

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

Considerations for IBM Z Cyber Vault Environment implementation

This task involves defining the specific configuration for the IBM Z Cyber Vault Environment that will be used for validation, forensic analysis, and recovery planning. Let us look briefly at some of the topics to be considered:

- Network isolation: your IBM Z Cyber Vault Environment will need to be network isolated. You do not want your Cyber Vault processing to impact your production workload. You will also need to limit access to your IBM Z Cyber Vault Environment. Your network team can help determine the best way to achieve that isolation based on your current network architecture.
- Security: you will want to limit user access to your IBM Z Cyber Vault Environment. Your Safeguarded backup capacity is immutable, but you want to ensure that the GDPS LCP or CSM policies you have are secured properly so that bad actors cannot change those polices. You will also want to consider the implications of user access to the IBM Z Cyber Vault Environment.
- I/O connectivity: your IBM Z Cyber Vault Environment will need access to the recovery volumes where the Safeguarded backups are recovered and the persistent volumes that will contain scripts and reporting data.
- Data validation approach: you will need to make decisions about things like how often you plan to validate environments, whether you want to leverage a FlashCopy® or a recovered Safeguarded Copy, what specific types of data corruption you would like to detect, and how you want to report the results. These decisions will be reflected in the automation scripts.
- z/OS configuration: best practice is to duplicate the primary environments that you are protecting minus the GDPS K-Systems if you have GDPS. That means you should have the same general configuration from an LPAR and IPL process perspective. For example, if you have an 8-way sysplex that you are protecting, you will define an 8-way sysplex in your IBM Z Cyber Vault Environment. If you are protecting individual LPARs, you will define those LPARs in your IBM Z Cyber Vault Environment. Resources like processor capacity, memory, and coupling facilities will likely be significantly less. How much less will depend on factors like the size of the environment and complexity of the validations.
- GDPS: if you are leveraging GDPS and GDPS LCP, you will need to consider how it needs to be configured based on the topology you have chosen and the security requirements you have.
- Access to tape: if tape mounts are needed from your IBM Z Cyber Vault Environment for validation or forensic analysis, access from this environment will be needed to provide access to the tape storage environment.

Considerations for IBM Z Cyber Vault Automation implementation

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

If you do not use GDPS and GDPS LCP, you will still want to define processes and build automation scripts that can be used in your environment. Copy Services Manager (CSM) will provide the Safeguarded Copy management. You will need a solution for automating the validation processing. This needs to include a completely automated IPL process.

2.1.3 Testing considerations

There are several types of testing you will need to do for your IBM Z Cyber Vault solution. Each type has a different objective, for example:

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

2.1.4 Operational considerations

Ideally, your IBM Z Cyber Vault solution will not require much manual intervention. GDPS LCP will automate the storage management policies and initiate the daily validations. To minimize the resources needed to manage your IBM Z Cyber Vault solution, you should implement monitoring and alerts for it.

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

- Monitoring detects infrastructure issues (like storage filling up or LPARs not starting properly)
- Validation processes fail and alerts are created
- Response time or performance problems are identified
- Corruption is suspected or identified
- Security breach has been detected

Your environment is network isolated and you will want to limit user access to it. It is important to understand the roles needed on a daily basis and provide a mechanism for providing additional access only when needed. Daily operations may require access for a limited number of storage administrators, system programmers, network administrators, and monitoring/operations. There are two operational objectives. First, ensure that the Safeguarded Copy captures are working and there is no risk of running out of space. Second, ensure that validations are running as designed and if alerts are generated, they are investigated and addressed quickly.

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

In the case of corruption or some other problem, additional teams or team member may need access to the environment. For example, database administrators and application support teams may need access to perform forensic analysis and define recovery processes. You will need a process to grant temporary access.

2.2 IBM Z Cyber Vault solution reference architecture

In this section we introduce the infrastructure and the underlying components that make up the IBM Z Cyber Vault solution. We also explore several high-level architectural decisions that need to be made when designing and planning for the implementation of the IBM Z Cyber Vault solution.

Figure 2-2 on page 21 depicts the reference architecture for the IBM Z Cyber Vault solution. This reference architecture will help you become familiar with the terminology and locate the various components addressed in this chapter. The IBM Z Cyber Vault solution consists of three domains: IBM Z Cyber Vault Storage, IBM Z Cyber Vault Environment, and IBM Z Cyber Vault Automation, as presented in 2.2, “IBM Z Cyber Vault solution reference architecture” on page 20.

For simplicity, we show only the primary and Cyber Vault source volumes in the data capture process. In most configurations there will be additional volumes created by replication tools to support business continuity requirements. It is not uncommon to have up to five copies of production data, with any of those being the source for the Cyber Vault volumes.

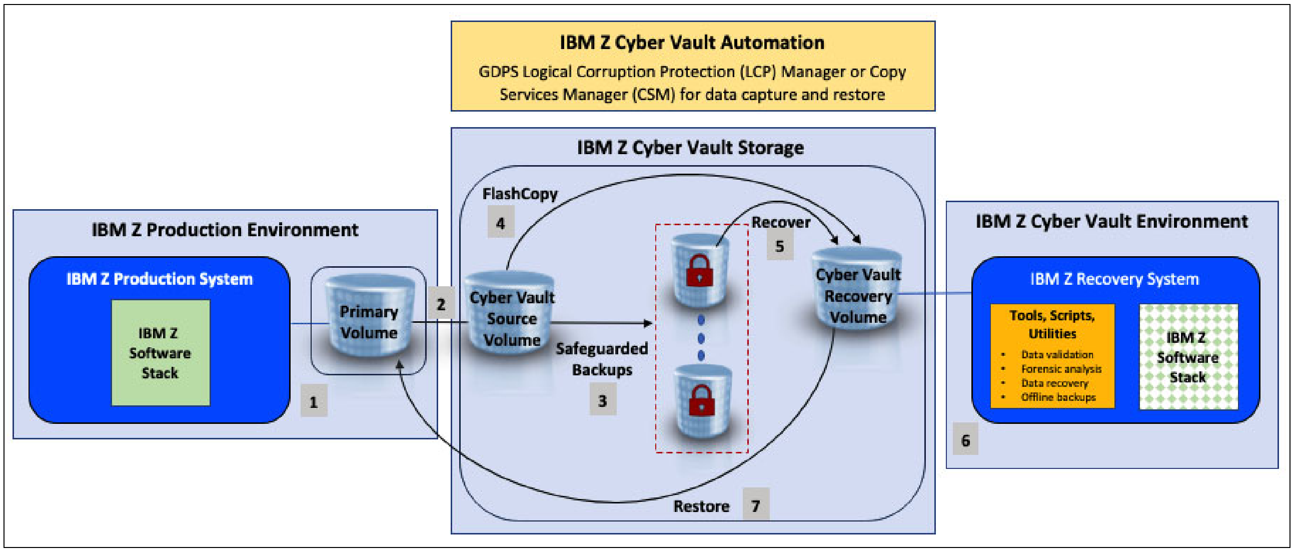
IBM Z Cyber Vault Automation
GDPS Logical Corruption Protection (LCP) Manager or Copy Services Manager (CSM) for data capture and restore
IBM Z Cyber Vault Storage
FlashCopy
4
Recover
5
IBM Z Production Environment
IBM Z Production System
IBM Z Software Stack
Primary Volume
1
2
Cyber Vault Source Volume
Safeguarded Backups
3
Cyber Vault Recovery Volume
IBM Z Cyber Vault Environment
IBM Z Recovery System
Tools, Scripts, Utilities
Data validation
Forensic analysis
Data recovery
Offline backups
IBM Z Software Stack
6
Restore
7

Figure 2-2 IBM Z Cyber Vault reference architecture

At its core, the IBM Z Cyber Vault solution is established with isolated, immutable copies of the production data, taken at multiple points in time, with recovery and restore functions. In addition, this solution helps identify data corruption through a validation process. The IBM Z Cyber Vault solution capabilities are discussed in detail in Chapter 3, “IBM Z Cyber Vault capabilities” on page 49.

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

1. Workloads running in an IBM Z Production System read and write data on the Primary Volume.
2. That data is replicated as part of your business continuity plan and one of those copies becomes the Cyber Vault Source Volume.
3. The procedure to capture production data from the Cyber Vault Source Volume is invoked by GDPS LCP Manager running in either in the production environment or the IBM Z Cyber Vault Environment. CSM is an alternative to GDPS LCP Manager for the creation and management of the Safeguarded backups. Those copies will ultimately be recovered to the Cyber Vault Recovery Volume when needed for data validation, analysis, recovery, or offline backup.
4. FlashCopy is initiated with GDPS LCP Manager or CSM to create the Cyber Vault Recovery Volume for data validation of the source data of any point in time in the IBM Z Recovery System. This step can be replaced or substituted with the next step (5) to validate captured data. FlashCopy is typically faster than a recover action against the Safeguarded backup so it may be preferred for validation.
5. Safeguarded backup is recovered to the Cyber Vault Recovery Volume to carry out data validation, analysis, recovery, or offline backup.
6. Data validation, analysis, recovery, or offline backup is performed in the IBM Z Cyber Vault Environment using various tools, scripts, and utilities. The IBM Z Cyber Vault Environment is a network isolated, secured system that prevents contamination of the validated copy and impacts to the production environment.
7. Corrupted data in the IBM Z Production System can be restored from the IBM Z Cyber Vault solution. If the event was targeted, only the data that was affected can be restored. If the event was widespread, the entire production environment can be restored.

Note: FlashCopy and Safeguarded Copy require the IBM Z Cyber Vault Source and IBM Cyber Vault Recovery volumes to be in the same storage system.

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

- IBM DS8000 Storage with Safeguarded Copy – to enable the creation of point-in-time copies of production volumes that cannot be changed or deleted (immutable copies).
- IBM Software for data validation, forensic analysis, and surgical recovery.
- IBM Geographically Dispersed Parallel Sysplex (GDPS) Logical Corruption Protection (LCP) Manager or IBM Copy Services Manager (CSM) – to create secured, tamper-proof backups stored in an isolated environment.
- IBM Z (isolated LPARs) – to restore point-in-time copies of the production environment and perform data validations, forensic analysis, surgical or catastrophic recoveries, and offline backups when needed.

2.2.1 Production environment

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

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

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

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

Regular validation of the production data in your IBM Z Cyber Vault Environment will give a greater degree of confidence in your Safeguarded backups. Those validations are not done in real-time. You should also consider having real-time detection capability in the production environment. It can provide immediate or earlier warning of some form of a cyberattack or data corruption taking place.

2.2.2 IBM Z Cyber Vault Storage

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

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

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

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

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

2.2.3 IBM Z Cyber Vault Environment

The IBM Z Cyber Vault Environment is where the IBM Z Cyber Vault solution capabilities described in Chapter 3, “IBM Z Cyber Vault capabilities” on page 49, will be executed. The capabilities include data validation, forensic analysis, recovery, and potentially any offline backups that are necessary to satisfy your data retention requirements, as well as offensive security procedures.

The IBM Z Cyber Vault Environment consists of one or more recovery systems (LPARs) that start out empty and are isolated from the production environment. A recovery system is started from a point-in-time image of a production system with special parameters that keep the network restricted to the IBM Z Cyber Vault Environment with limited user access.

When Safeguarded backups are selected, they are first restored onto the IBM Z Cyber Vault Recovery volumes, which are then used to perform an IPL of the IBM Z Cyber Vault Environment. The Safeguarded backups are an exact replica of the production environment. All production system volumes are included in the Safeguarded backups to ensure broader protection.

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

After all 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 backups to be processed. As a result of this mechanism, the IBM Z Cyber Vault Environment remains isolated and Safeguarded backups remain immutable.

The recovery system in the IBM Z Cyber Vault Environment should run in an isolated network to ensure that the Safeguarded 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 as well.

For more information see 2.4, “IBM Z Cyber Vault Environment considerations” on page 42.

2.2.4 IBM Z Cyber Vault Automation

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

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

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

GDPS LCP Manager also provides automation capabilities and procedures that support the IBM Z Cyber Vault Environment and greatly simplify operations for data validation (see 2.4, “IBM Z Cyber Vault Environment considerations” on page 42 for details).

2.2.5 Architectural decisions

In this section we discuss several high-level architectural considerations that will influence the design and implementation of the IBM Z Cyber Vault solution.

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

Figure 2-3 on page 25 represents a physically isolated topology. Hardware-based replication is used to copy the production data into IBM Z Cyber Vault Storage environment via a Secondary volume. An asynchronous replication method called Global Mirror or Global Copy is used for the replication into the IBM Z Cyber Vault Storage Source volume. The distance between the production environment and IBM Z Cyber Vault Storage can potentially be hundreds or thousands of kilometers, or it can be adjacent to either the primary or secondary volumes.

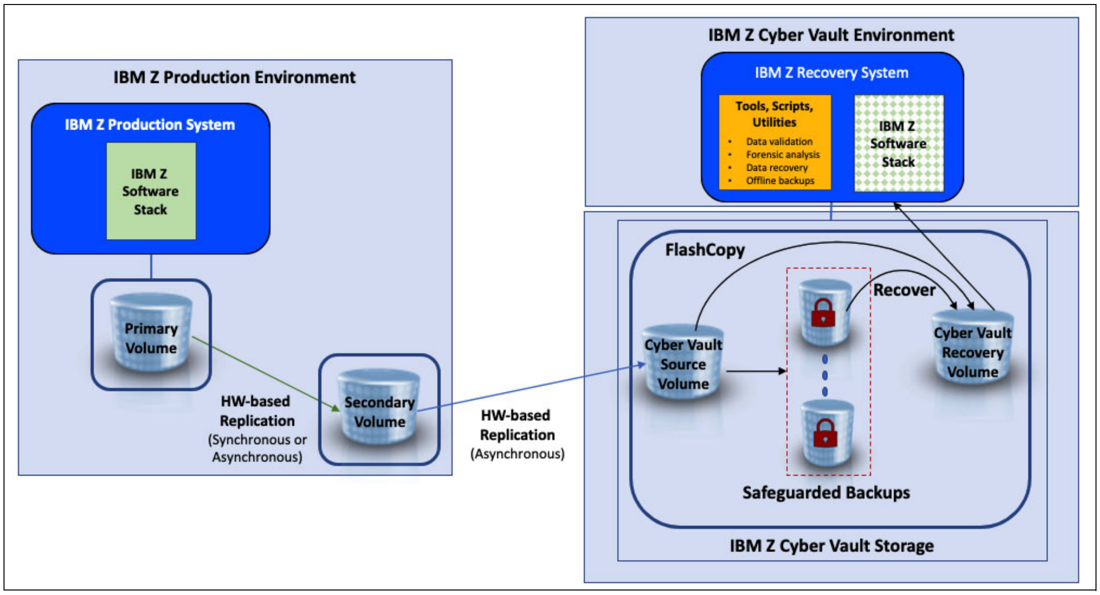
IBM Z Production Environment
IBM Z Production System
IBM Z Software Stack
Primary Volume
HW-based Replication (Synchronous or Asynchronous)
Secondary Volume
HW-based Replication (Asynchronous)
IBM Z Cyber Vault Environment
IBM Z Recovery System
Tools, Scripts, Utilities
Data validation
Forensic analysis
Data recovery
Offline backups
IBM Z Software Stack
FlashCopy
Recover
Cyber Vault Source Volume
Cyber Vault Recovery Volume
Safeguarded Backups
IBM Z Cyber Vault Storage

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

The immutable copies (Safeguarded backups) are in an isolated storage environment and inaccessible from all systems. Using IBM Z Cyber Vault Automation, the Safeguarded backups are recovered from a combination of the IBM Z Cyber Vault source volume and the immutable backups onto an isolated pool of Cyber Vault recovery volumes. The recovery systems are then IPLed as needed to support the relevant IBM Z Cyber Vault solution capabilities described in Chapter 3, “IBM Z Cyber Vault capabilities” on page 49.

We can contrast the physical isolation depiction in Figure 2-3 with a virtual isolation topology as shown in Figure 2-4 on page 26.

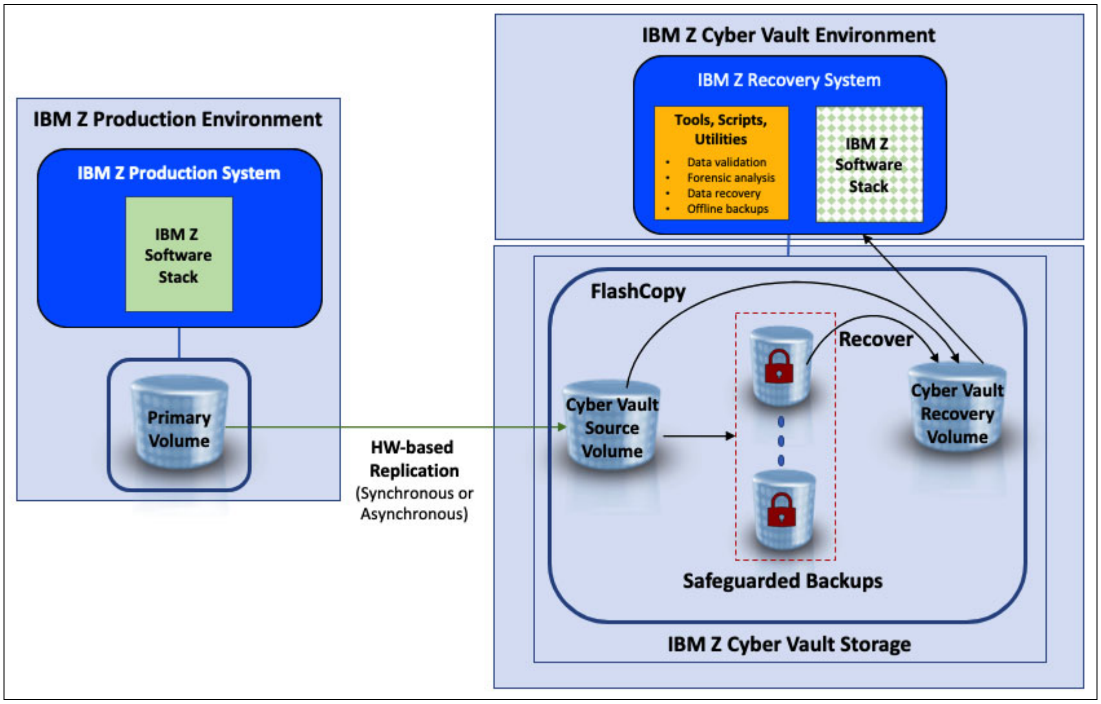

Figure 2-4 IBM Z Cyber Vault solution reference architecture (virtually isolated)

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

One characteristic to consider is the mechanism for capturing a consistent, point-in-time copy of the production data. When capturing production data, the I/O traffic to the primary volume must be paused across all volumes in the consistency group.[1] With physical isolation, this results in the replication to the IBM Z Cyber Vault Storage environment being paused in a consistent state, which does not directly impact the production I/O traffic. However, with virtual isolation, as it is the production data is copied directly from the Primary volume where synchronous replication (Metro Mirror) is in use, I/O traffic to all volumes must be paused while the point-in-time consistent backup is taken. Since the impact is like a consistent FlashCopy for all devices in production, you might want to use that to determine the impact.

A second consideration is where the capacity in the IBM Z Cyber Vault Environment for the validation process should be provisioned. That could be in the same site as your production environment or a remote site. Remember that the Recovery System will need connectivity to the Cyber Vault Recovery volumes in the IBM Z Cyber Vault Storage Environment.

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

Another characteristic to be considered is the separation from a management environment perspective. With virtual isolation, using GDPS LCP Manager, the control point for operations

[1] A consistency group can manage the independent relationships for all volumes that are related to an application, ensuring consistent, synchronized data.

such as taking the backup, recovering a backup or releasing a backup will be using the same control point (or GDPS controlling system) as used for the management of day-to-day production data replication.

GDPS LCP Manager supports both virtually isolation or physically isolation of the IBM Z Cyber Vault Storage environment. Depending on the selected isolation method, the implementation approach is different.

With physical isolation, a separate control point is used, allowing for another degree of separation from the production environment and the ability to have a different security database with rules specifically tailored just for the IBM Z Cyber Vault Environment. This consideration borders on the interface between cyber resilience which is where the IBM Z Cyber Vault solution fits and cyber security as described in 1.2.2, “Cyber security verses cyber resiliency” on page 5.

In addition to the IBM Z Cyber Vault Storage environment where you need to determine whether the physical or virtual isolation best meets your requirements, there are considerations for your IBM Z Cyber Vault Environment. You will also need to decide if you have dedicated IBM Z hardware for your IBM Z Cyber Vault Environment. Both are technically viable. However, this decision will have implications on your current configuration and software licensing. You will need to discuss options with your IBM team and any vendors that supply your software products.

2.3 IBM Z Cyber Vault Storage considerations

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

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

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

2.3.1 DS8000 Safeguarded Copy concepts

Safeguarded Copy provides immutable point-in-time copies that are not accessible by systems or applications because they do not have logical control units (LCUs) or volume IDs associated with them. Safeguarded Copy Source volumes and their backups cannot be erased or deleted using the DS8000 native interfaces (DS Command-line Interface (DS CLI) or DS GUI). To access a Safeguarded backup, a recovery action to a IBM Z Cyber Vault recovery volume is necessary so that the data can be accessed from a IBM Z Cyber Vault recovery system (see Figure 2-5 on page 28).

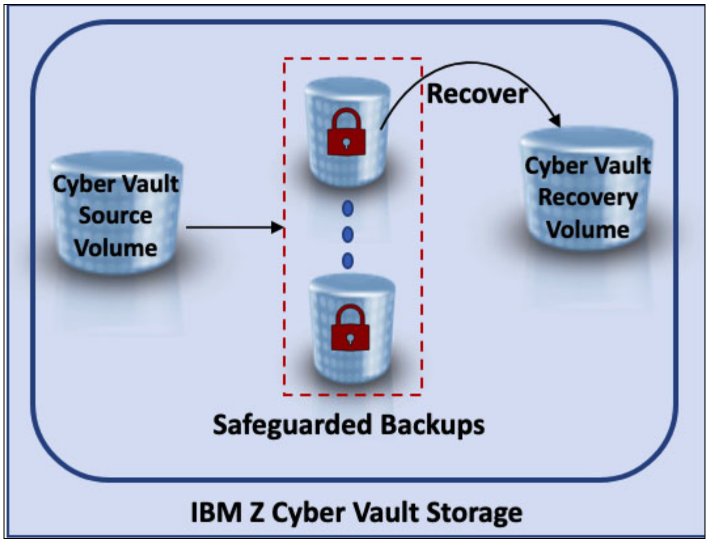

Figure 2-5 IBM Z Cyber Vault Storage

With Safeguarded Copy, you can have up to 500 crash-consistent point-in-time copies per volume. To optimize capacity usage, the Safeguarded backups depend on each other, and changed tracks are stored only once per backup in the Safeguarded backup capacity. The Safeguarded backup capacity is similar to an undo log.[2] The Safeguarded Copy capability can be integrated with different disaster recovery (DR) or high availability (HA) configurations.

Either GDPS Logical Corruption Protection (LCP) Manager or IBM Copy Services Manager (CSM) is required to manage Safeguarded backups. For more details about management software options, see 2.2.3, “IBM Z Cyber Vault Environment” on page 23.

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

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

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

Once Safeguarded backup capacity is defined, physical storage capacity in the DS8000 is required to store the data changes. Without a Safeguarded backup, the backup capacity is pure virtual capacity; physical capacity is allocated as you use your desired management software (GDPS LCP Manager or CSM) to set up and manage Safeguarded Copy. The

[2] When a change is made to data, an undo record that describes the change is logged, so that a transaction can be rolled back if needed, essentially reversing the changes.

management software is used to capture, recover, restore, and expire Safeguarded backups and contains the policy to control them.

Create and capture Safeguarded backups

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

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

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 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 consistent backup, the DS8000 system must stop updates for all volumes within the session for a short period. It does this task by setting an Extended Long Busy (ELB) state and then performing the final tasks to create the backup. When the check in is completed for all volumes and all volumes in ELB, the backup is consistent.
3. **Check out or Completion**: The DS8000 lifts the ELB and write operations can continue. From now on, the DS8000 system writes further backup data into the CG logs of the new backup

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

Recovering Safeguarded backups

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

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

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

Note: The recover action is providing access to only a selected backup on recovery volumes which allows you to do data validation and forensic analysis. The data on the recovery volumes can be then be used to restore that data to the production volumes.

Expiring Safeguarded backups

With your management software, you can expire Safeguarded 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.

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

- The system reaches 500 Safeguarded Copy copies per volume.
- A DS8000 storage pool runs out of physical space.
- The specified Safeguarded 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.

Note: Above mentioned default behavior for roll-off backups can be modified in DS8900F system with Rel. 9.3.2 and above. For more information see, IBM DS8000 Safeguarded Copy.

Safeguarded backup restore- to-production

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

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

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

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

The process is managed and controlled with the management software.

These are the steps that must be done during this restore-to-production process:

1. Recover selected backup to Recovery volumes with your management software.
2. Validate data on Recovery volumes.
3. Perform forensic analysis to check if a catastrophic recovery is required.
4. If not already done, stop production / applications.
5. Suspend replication relationships.
6. Make sure no replication relationship exists on Recovery volumes.
7. Start restore-to-production with your management software.
8. After all data are copied via incremental copy process you might validate data on production system.
9. Reestablish suspended replication relationship.
10. Restart production environment.

For more details about the restore-to-productions process review your management software documentation and IBM DS8000 Safeguarded Copy.

2.3.2 Safeguarded Copy prerequisites

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

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

For managing Safeguarded Copy, either a fully licensed CSM V6.2.3.1 or later or GDPS LCP Manager V4.2 SP2 or later is required. You cannot use the DS8000 interfaces DS GUI / DS CLI or a z/OS interface to capture, recover, expire, or restore a Safeguarded backup.

Make sure you consider the following when planning for a Safeguarded Copy implementation:

- Safeguarded Copy operates at the volume level.
- The DS8000 system maintains a maximum number of 500 backups per volume.
- If you intend to use a backup frequency of less than 10 minutes, IBM requires that you submit a Request Price Quotation (RPQ) for approval and support.
- The maximum Safeguarded backup capacity for a volume is 14.6 TiB for CKD volumes.
- A Safeguarded Copy source can be a FlashCopy target if DS8900F microcode Release 9.3 or later is installed.
- The source volume and Recovery volume must be managed by the same DS8000 internal server. Therefore, they must both be either in an even or odd logical subsystem (LSS).
- Only a single Safeguarded backup for a volume can be recovered at the same time.
- During a Safeguarded Copy recovery action (with no-copy option), a cascaded FlashCopy from the Recovery volume to another volume is not possible.
- DS8000 Dynamic Volume Expansion (DVE) is not supported for Safeguarded Copy source volumes in the same way as it is not supported for other replication types.
- 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.

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 backup capacity over the retention period.
- The small Safeguarded Copy overhead for each backup
- Recovery volumes
- Safeguarded Copy Source volumes (for physical isolation)
- Provisioning of Global Mirror FC Journal volumes if GM is used in a physical isolated solution.

2.3.3 Safeguarded Copy capacity sizing considerations

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

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

In this section, we describe the capacity sizing for Safeguarded Copy.

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 and source volumes physical and virtual capacity.
- Size the Safeguarded 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, The limits for DS8000 systems with below 1TB cache is only ~25% of the maximum capacity limit. For more information about capacity limits refer to the DS8000 product documentation for your DS8000 model.

It is necessary to estimate the physical and virtual capacity of the following components in the DS8000 storage system:

- Safeguarded backup capacity
- Recovery volume
- Safeguarded source volume if a physical isolation approach will be implemented

Physical capacity estimation is necessary to determine how much capacity is required to store all changed data within the Safeguarded 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 backup capacity and the Safeguarded virtual capacity depend on the data change rate and following backup management policies:

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

Frequency of Safeguarded backups

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

There are also implications based on the source of your backups. When backups are taken in a Metro Mirror environment, you must “freeze” all write I/O to the volumes being backed up. A higher backup frequency would result in a freeze frequently impacting production. However, if the Safeguarded 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.

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

Retention period of Safeguarded backups

Beside the backup frequency, you must decide how long you want to keep Safeguarded backups. The longer your retention period is, the more capacity is required to store the changes in the DS8000 system. The following considerations will help you define the best retention period to meet your needs:

- Do you have regulatory or business requirements that define how long you must keep the backups?
- 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 7 - 14 days. Copying validated data from the DS8000 system to LWORM tapes is a good option for longer term retention on lower cost media. This in turn would also reduce the amount of capacity that is required in the DS8000 system.

You should define the retention period based on your requirements. Make sure you do a Safeguarded backup capacity sizing (as described in the subsequent section), to reflect your retention period and backup frequency requirements.

Safeguarded backup capacity sizing

As an example, if you are creating a backup every 6 hours and retaining it for 7 days, you would need to know the data change rate of 42 backup intervals over a 7-days retention period to estimate the Safeguarded backup capacity.

In addition to the Safeguarded backup capacity, the Safeguarded recovery volumes must be sized as well. 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 source volumes change while the relationship exists.

> **Note:** Best practice is to add ~20% of the physical source volume capacity for the recovery volumes so that they can be used for data validations during normal operations.

The sizing for the Safeguarded backup capacity (physical and virtual) and recovery volumes can be done by using the methods described below. 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 backup interval for as many intervals as the length of the retention period. You must do this task for each DS8000 system to calculate the required physical capacity for the Safeguarded backups.

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

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

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

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

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

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

  IBM Storage pre-sales and IBM Business partner can use IBM Storage Modeller to do a sizing based on RMF data.

Estimating the virtual capacity and assigning a correct Backup Capacity Multiplier is very important. It is not uncommon to run out of virtual capacity for volumes or reach the DS8000 virtual capacity limit. To avoid this, a preferred practice is to assign the same Backup Capacity Multiplier to all volumes that belong to same z/OS DFSMS Storage Group after you have estimated the virtual capacity per volume. You should assign the highest Backup Capacity Multiplier estimate for a DFSMS Storage Group. To simplify storage management, you can reduce the number of different Backup Capacity Multiplier to three to four per environment. Of course, that approach will increase required virtual capacity, but it avoids losing a Safeguarded backup.

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

For more information, see IBM DS8000 Safeguarded Copy.

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

Additional storage for IBM Z Cyber Vault solution

Additional storage capacity is required for the IBM Z Cyber Vault Environment beyond the backup capacity and recovery volume capacity. For example, storage volumes are needed for surgical recovery. We call these extra volumes staging volumes.

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

You will also require some persistent volumes in your IBM Z Cyber Vault Environment to store reports and other historical data sets created by the data validation process.

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

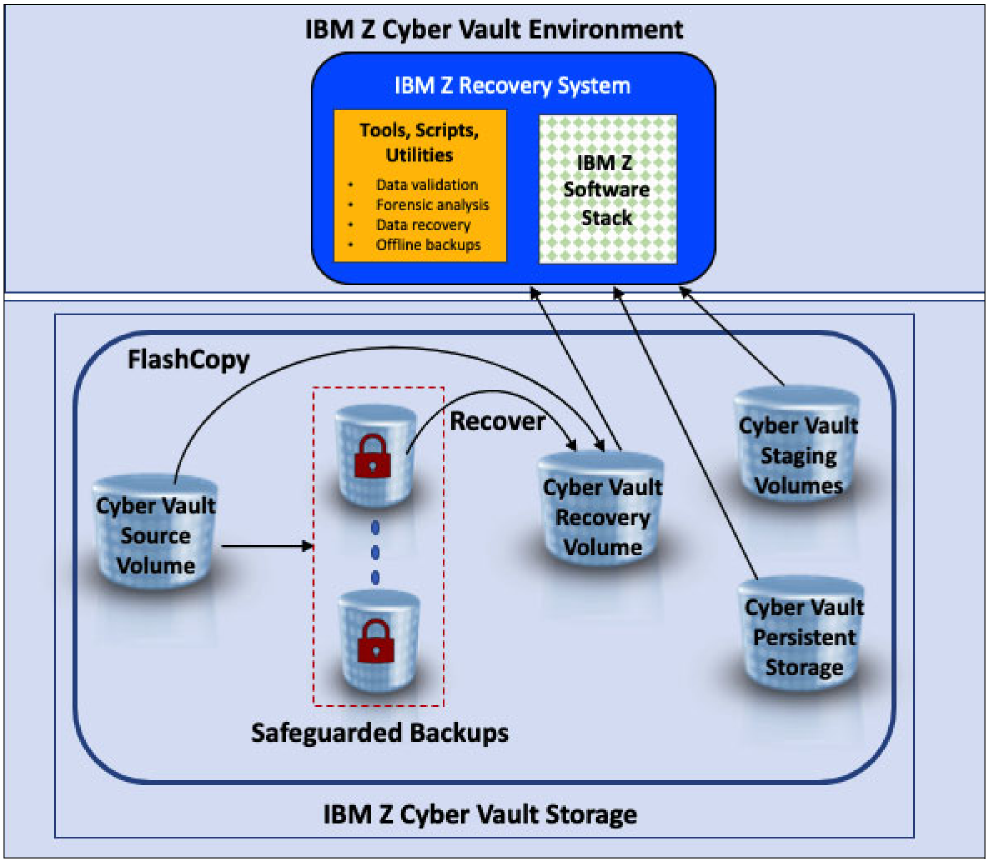
IBM Z Cyber Vault Environment
IBM Z Recovery System
Tools, Scripts, Utilities
Data validation
Forensic analysis
Data recovery
Offline backups
IBM Z Software Stack
FlashCopy
Recover
Cyber Vault Source Volume
Cyber Vault Recovery Volume
Cyber Vault Staging Volumes
Cyber Vault Persistent Storage
Safeguarded Backups
IBM Z Cyber Vault Storage

Figure 2-6 IBM Z Cyber Vault Storage architecture

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

In addition, a smaller set of volumes are used, called staging volumes for surgical recovery and a few persistent volumes to store 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.

You can also use a second set of recovery volumes so that regular tasks like data validation at the time as recovering a Safeguarded backup for surgical recovery or forensic analysis can be performed.

2.3.4 DS8000 performance considerations for Safeguarded backup

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

There will be a number of different workloads running on the storage systems that are being used by your IBM Z Cyber Vault Environment.

- The production write workload updating the Safeguarded Copy source volumes
- The saving of data into the Safeguarded backups
- FlashCopy activity for the recovery copy being used for data validation
- FlashCopy activity for the Global Mirror journals (if you are using Global Mirror for replication into IBM Z Cyber Vault Storage)
- The data validation workload which is typically a read intensive workload on the recovery volume

The Safeguarded Copy and FlashCopy activity can each be considered worst case to be an additional backend read and write I/O for every write to the Safeguarded Copy source volume. Most of the time it will be significantly less than this, but for sizing purposes you may want to consider the worst-case scenario.

Another important aspect that must be considered when running a Safeguarded backups on a Global Mirror secondary volume, is the use of FlashCopy with Global Mirror. By using Remote Pair FlashCopy (RPFC) in the Metro Mirror environment, this function can help keep replication relationships in full duplex even though a FlashCopy operation is done on a Metro Mirror primary volume.

However, if RPFC is used in a Multi-Target Metro Mirror or in a Metro-Global-Mirror environment data must be copied thru the PPRC-link to either the second Metro Mirror leg or to the Global Mirror leg. That means a burst write activity with a hundred percent data change rate for all involve volumes or datasets. During this time the second Metro Mirror leg will be in copy pending and a Global Mirror leg potentially cannot create a Global Mirror consistency group.

If you use Remote Pair FlashCopy for many volumes, you must coordinate RPFC and Safeguarded capturing. Otherwise, you might not be able to create a Safeguarded backup as the Global Mirror RPO is too high to successfully pause Global Mirror.

The best approach would be do all Remote Pair FlashCopy at the same time and create a Safeguarded backup after a second Metro Mirror leg is in full duplex or after Global Mirror has created a new consistency group.

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

2.3.5 DS8000 Safeguarded backup operational considerations

For Safeguarded Copy environments, you will have additional operational considerations. In this section we discuss, operational changes like impact of capturing a Safeguarded backup, adding and removing volumes, changing Safeguarded Copy Volume Backup Multiplier when adding volumes to an Storage Management Subsystem (SMS) Storage Group. We also discuss monitoring and alerting and DS8000 security considerations.

Impact of capturing Safeguarded backups

A DS8000 Safeguarded backup is a protected crash consistent point-in-time copy. To create consistency the DS8000 pauses all updates to volumes within the Consistency Group (CG) for a very short period by presenting an Extended Long Busy (ELB) state. If multiple DS8000 systems are involved in a Safeguarded Copy CG the management software GDPS LCP Manager or CSM makes sure that the backup is crash consistent across all volumes. The fewer number of volumes, the smaller the impact will be.

In a Metro Mirror environment where either Safeguarded Copy is running on Metro Mirror primary or Metro Mirror secondary, the z/OS LPARS will notice this short pause. This could be as low a 100-200ms for an environment with only a few thousand volumes up to perhaps 1-2 seconds for environments with 10,000 or more volumes. For CSM managed environment make sure you have implemented the z/OS IOS enhancement APAR OA59561 and establish a secure IP connection to the z/OS HyperSwap® Address space to reduce freeze impact, for more information refer to IBM DS8000 Safeguarded Copy. The desire to avoid an ELB in production is one reason clients implement a physically isolated solution with a Global Mirror relationship to a separated system that is running Safeguarded Copy.

In a Global Mirror environment where Safeguarded Copy is running on a Global Mirror secondary, it is required to pause the Global Mirror with consistency before capturing or creating a Safeguarded backup. That will create a consistent Safeguarded backup from the Global Mirror secondary. This GM pause will increase GM RPO (Recovery Point Object). The GM pause and creating a Safeguarded backup can be automated by the management software GDPS LCP Manager.

Monitoring and alerting considerations

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

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

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

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

- IEC817I - provides warning for virtual capacity out-of-space situation.
- IEA499E – warning related to physical capacity of the storage pools.

GDPS LCP Manager also provides message that allow monitoring and alerting.

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

Changing Volume backup multipliers

During the lifetime of a Safeguarded Copy environment there might be circumstances that require changes of the Volume Backup Multipliers. For example, a new volume must be added to the SMS Storage Group, or the workload has changed. Requirements may change over time and a higher frequency or retention period may be required.

The DS8000 provides a dynamic expansion of the Safeguarded Copy Backup Capacity. You can do that by increasing the volume backup multiplier in the DS8000 via the DS Command Line Interface (DSCLI) or via the DSGUI. But an expansion might be delayed until the retention period is over or backups that are blocking an expansion are manually deleted. Unfortunately, you cannot predict if a backup will block expansion or not. If expansion is not possible the volume will stay in “expanding state” for a longer period and both the CSM Safeguarded Copy session and GDPS LCP Manager will indicate which backups are blocking the expansion. At that point, you can decide if you manually delete backups or just wait until the retention period of these backups is over.

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

If your requirements have changed, and you must increase the backup frequency and/or the retention period, IBM recommends doing again a Safeguarded Copy capacity sizing with the ESESizer / Safeguarded Copiesizer tool, compare the results with your current configuration and adjust the volume backup multipliers accordingly.

DS8000 security aspects related to Safeguarded Copy

Safeguarded Copy provides immutable point-in-time copies to protect your data in the event of a logical corruption event. To ensure that those copies are protected, you should consider increasing security in a Safeguarded Copy environment. You must make sure that a single user cannot modify the policies that control the Safeguarded backup creation and expiration and damage the DS8000 configuration.

We recommend that you follow these best practices:

- Implement a separation of duty policy so that a single user has not access to both the management software GDPS LCP or CSM and the DS8000 system.
- Restrict DS8000 user access and limit user rights in the management software. There are different functions available to restrict access in the DS8000 and in the management software.
- Use multi-factor authentication in the DS8000 to improve security for the login process and implement custom user roles which allows a very granular definition of user rights. See

DS8000 documentation for more information:
https://www.ibm.com/docs/en/ds8900/9.4.0?topic=security

Both GDPS LCP Manager and CSM provide role-based security and dual-control functions for Safeguarded Copy. With role-based security you can restrict the user rights in the management software by assigning a specific role. The dual control function provides additional protection for certain tasks like expire Safeguarded backups or changing GDPS LCP management profiles / CSM properties. After enabling dual control, a second user with the same privileges must approve the task before it gets executed. For more information about GDPS LCP Manager dual control refer to the GDPS LCP Manager documentation. For CSM dual control function, see:
https://www.ibm.com/docs/en/csm/6.3.11?topic=security-dual-control

You might also consider an integration of your Safeguarded Copy and IBM Z Cyber Vault Environment into a SIEM (Security information and event management) solution to detect unauthorized actions in this environment. You can forward audit logs and messages from the DS8000 and the management software into a SIEM solution. Depending on log entries or messages, certain actions such as sending a notification or take another Safeguarded Copy Backup could be initiated.

2.3.6 TS7700 / VTS considerations for IBM Z Cyber Vault solution

Most z/OS environments include the use of the IBM TS7700 or virtual tape solutions from other vendors. The information described in this section will focus on the TS7700 and its capabilities, but similar parallels might be available with other vendor solutions.

Virtual tape is a widely used technology to provide lower cost media for the typical backup and restore use cases, but is also key in supporting an “operational tape” function for primary production operations.

Examples of operational tape use might include:

- DFSMShsm migrated data
- OAM object data
- Batch input and output datasets
- Reporting data

Some of this data is no longer stored on the primary storage system and therefore is the “master data” and should be protected beyond normal HA/DR best practices by using additional protections against compromised environments as part of the overall IBM Z Cyber Vault solution.

There are several security and data gapping related capabilities that should be considered for improving the security of an IBM TS7700 environment across all facets of the solution. Some of these include:

- Expire Hold
- Logical WORM (LWORM) and LWORM Retention
- Secure Data Transfer
- Data at rest encryption
- Event and rsyslog logging
- Dual admin authorization
- Physical Tape and Copy Export
- Cloud tiering and Cloud Export
- Selective Device Access Controls (SDAC)

All of these capabilities are covered in more detail in:
https://www.redbooks.ibm.com/redpieces/abstracts/sg248464.html

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

- Providing backups of disk volumes and application data. This is the well-known, well used capability of using tools like DFSMSdss, ImageCopy, and so on to create a point in time dump of critical data to tape volumes.

  For IBM Z Cyber Vault, this includes backing up the disk volumes associated with a point in time recovery volume that can be retained for a longer duration than the Safeguarded backup restore depth held by the DS8K.

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

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

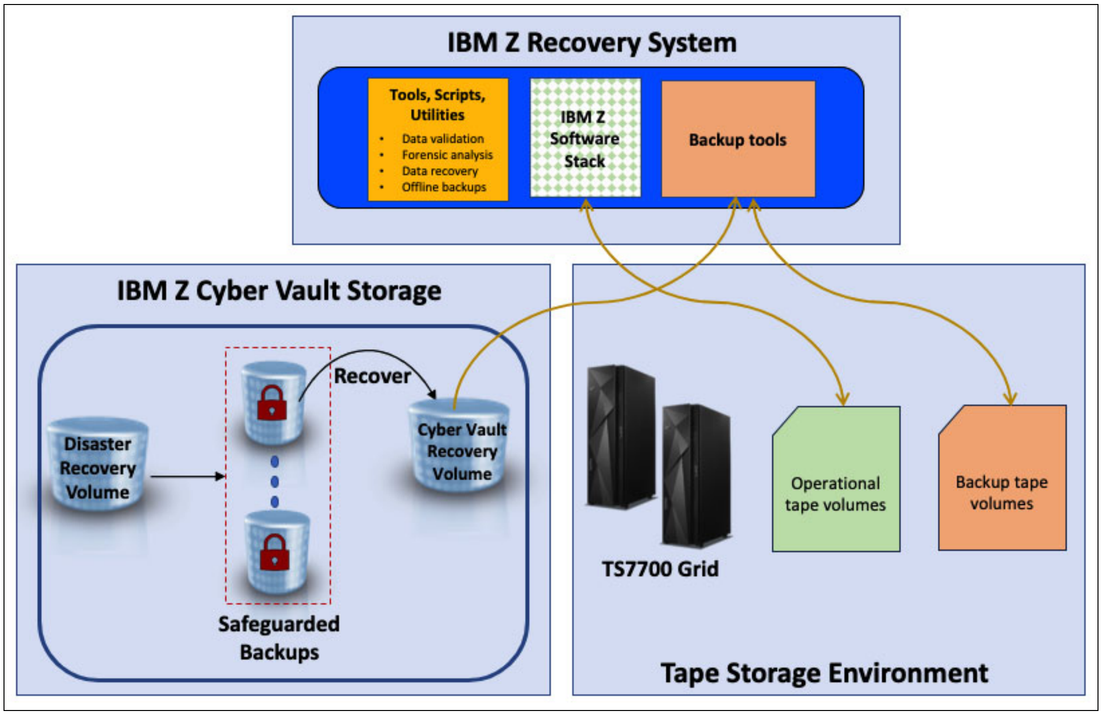

Figure 2-7 Tape storage for the IBM Z Cyber Vault solution

Protecting data with Logical WORM

While the deletion of tape data is different than disk, such that tape volumes are moved to a scratch category to make them eligible for future use and overwrite vs immediate impact, a savvy attacker could know how to accelerate this process. These protections should be considered for both backup and restore as well as operational tape volumes such as DFSMShsm.

In order to protect data stored within the TS7700 grid the combination of Logical WORM (LWORM), LWORM Retention, and the Expire Hold functionality should be used.

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

Note: the LWORM retention duration is defined to this tape volume when the volume is created and will not change for this volume even if this retention period is modified for future volumes. For example, if the LWORM retention is set to 3 months when the tape volume VOL001 is written from BOT, this 3-month value will remain with this tape volume regardless if the volume is copy exported, cloud exported, or if the retention period is changed for this data class. Any new retention period would be applied to any future tape volumes written from BOT.

Expire Hold will prevent any reuse of volumes that are returned to scratch for the hold period. This places a defined grace period from the time a tape volume has been returned to scratch until the tape environment looks to reuse that tape volume again to overwrite it with new data.

2.3.7 TS7700 / VTS best practice for IBM Z Cyber Vault solution

Leveraging tape to create backups of data is a well-established practice with known tools and practices. The additional thoughts to consider are aimed at improving the protection and gapping of these dumps away from a compromised environment. The following are the primary considerations to augmenting this capability:

- LWORM and LWORM retention – this is an option for preventing a compromised environment from deleting or overwriting data on tape volumes. In a virtual tape environment, there are no additional requirements other than defining the retention policies and data class constructs to instruct the TS7700 and tape management software for defining this level of protection. However, these tape volumes could still be cataloged, visible and accessible (read) to a compromised environment. LWORM will protect the data from alteration, but they are still accessible.
- SDAC controls – this would partition the access to the tape volumes for the dumps to a specific device range that only a dedicated backup host LPAR would have access to. This should be considered on top of LWORM to isolate the actual tape catalog and management system to a very specific and controlled LPAR.
- Physically isolated the TS7700 cluster or grid – provides the ultimate separation of tape dumps to a cluster or grid for the sole purpose of holding the dumps. This leverages a different set of hardware which could have different physical controls, management access, and administrator sets.

Access operational tape from an IBM Z Cyber Vault Environment

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

Since the IPL of your IBM Z Cyber Vault Environment LPAR might be from a Safeguarded backup instance that is days old, it will have a different tape management catalog view of the tape categories, tape data, and so on, compared to the “current” TS7700 database and tape management system in production. Protections of impacts from the IPL of a different system need to be put in place to ensure that the “older” system does not push changes to the environment. These can be done at the TS7700 hardware level or at the tape management level (RMM, CA1, and so forth) depending on your operational needs.

This is very similar to considerations that need to be put in place for a normal DR test, but with the addition that the tape management system view being used can be from “back in time”.

Ensuring that the “older” system being IPLed from days ago will still find the tape data on the tape volumes it has cataloged, which may have been scratched by the current system. Expire Hold should be used to ensure tape volumes moved to scratch by the current production view will not be overwritten until the oldest possible Safeguarded Copy recovery depth that might be IPLed to a cyber vault environment.

LWORM and LWORM retention should be considered for the “master data” instances on tape to further protect accidental or malicious attempts at overwriting data via direct FICON® CCW tape commands or bypassing the tape management software.

There are multiple tradeoffs between the options above and what might make the most sense for a given environment, business security or compliance requirements, or objectives that would find the optimal configuration for a given environment. These should be explored with a storage subject matter experts to help weigh these options.

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

For more details about see: https://www.redbooks.ibm.com/abstracts/sg248529.html

2.4 IBM Z Cyber Vault Environment considerations

In this section we explore considerations for the design and implementation of your IBM Z Cyber Vault Environment.

2.4.1 z/OS recovery system configuration

The purpose of the IBM Z Cyber Vault Environment is to provide a clean-room system where data corruption activities and actions can be safely performed. These include data validation processes, forensic analysis and recovery activities, and security strengthening by running additional offline backup procedures, or running offensive security exercises. Planning for the z/OS in the IBM Z Cyber Vault Environment should consider the following:

- As a best practice, at an LPAR definition level, the recovery system should mimic the production environment being protected in number of z/OS LPARs, minus the GDPS Kg/Kr systems, as well as cryptographic hardware. Doing so reduces complexity in managing the IPL and validation processes in the recovery environment. You will only need a single Coupling Facility (CF) LPAR.
- The size of the LPARs for the recovered systems and coupling facilities (CFs) depends on the workload and amount of data that you intend to validate and process. When sizing the recovery environment, keep in mind that the workload profile of the IBM Z Cyber Vault Environment will likely be considerably different than that of the production counterparts in terms of I/O activity and CPU capacity when running potentially many thousands of concurrent structure and data validation jobs. On the other hand, production workload will never run in an IBM Z Cyber Vault Environment, so sizing considerations are not related to actual transaction processing in production.
- Persistent volumes are recommended for the storage of the validation scripts and associated artifacts, and for data that should not be lost between different validation runs,

such as historical data or Systems Management Facility (SMF) records. Consider defining multiple volumes for this purpose with a unique esoteric to avoid having to hardcode the VOLSERs when creating output files during validation runs, and to I/O contention during validation cycles.

- Ensure operational procedures are reviewed and updated on a regular basis to accommodate changes in both the production and Cyber Vault environments. Processes such as master key rotation, or LPAR configuration changes will need to account for corresponding updates necessary in the Cyber Vault environment.
- The software stack in the IBM Z Cyber Vault Environment will be an exact replica of the production environment. Therefore, any tools, utilities, procedures to manage, analyze, validate, and restore system components need to be already identified and installed in production. This includes preparing a special IPL procedure, setting up the required software, activating tools that will be collecting metadata, tracking activity, and getting all the automation and data validation tools ready.
- Specialized configuration items required to properly direct the IPLs of the recovery systems and subsequent processing should be present in the production environment, so they are available as necessary in the IBM Z Cyber Vault Environment. This includes IBM Z Cyber Vault-specific PARMLIB members, PROCs, automation routines, and XCF policy definitions for CFRM (Coupling Facility Resource Manager) and WLM (Workload Manager).

2.4.2 Network isolation

The IBM Z Cyber Vault Environment should be constructed in a logical bubble with no intersection with the existing production environment. The IBM Z Cyber Vault Environment should reside in an isolated network to ensure that the Safeguarded 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 IBM Z Cyber Vault Environment and the production 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.

From the router network perspective, the IBM Z Cyber Vault Environment should be fenced off from the existing production environments with extremely limited access points. The IBM Z Cyber Vault Environment should only be accessible through a very small number of predetermined IP addresses, ports, and protocols as well as protected by firewalls and routers that requires VPN access and encrypted flows to communicate with the IBM Z Cyber Vault Environment (z/OS LPARs). The z/OS LPARs in the IBM Z Cyber Vault Environment should only be accessible through the network using Network Address Translated (NAT) IP addresses since the IBM Z Cyber Vault environment (z/OS LPARs) are activated using the exact same configurations as production and duplication of IP addresses on the network must be avoided.

2.4.3 z/OS security

As a best practice, the IBM Z Cyber Vault Environment should mimic the production environment including the external security managers (ESMs) that are used to secure the z/OS environment: IBM RACF, CA ACF2 or CA Top Secret. The security policies defined in production should be replicated in the IBM Z Cyber Vault Environment without modifications.

To perform validation within the IBM Z Cyber Vault Environment, additional security access may need to be permitted. For example, to execute data validation scripts and utilities within IBM Z Cyber Vault Environment:

- RACF started task user ID may require “System Special” access in the existing production environment so that the access is propagated into IBM Z Cyber Vault Environment when activated using production volumes.
- PARMLIB methodology activation is recommended in the startup of RACF on IBM Z Cyber Vault Environment.
- Started Task user ID activating/owning validation scripts and utilities must be defined in RACF as “Trusted” and will be restricted to only operations capability within the IBM Z Cyber Vault Environment.

2.4.4 Database and middleware consideration

Applications use different subsystems and are developed for specific purposes. 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.

Data validation

Application subsystems enable transaction and data management, providing data and application integrity. Their software controls the creation, organization, and modification of data, and 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. As a result, it is important to perform data validations against all these subsystems to make sure that their inner workings are not compromised. In the IBM Z Cyber Vault Environment, it is recommended all applications and databases are activated by system automation during the IPL process to enable a comprehensive validation of the database and middleware in conjunction with the infrastructure.

For an IBM Z Cyber Vault Environment configuration, the recommended database and middleware capabilities include, but are not limited to the following:

- Validate database structures
- Monitor data changes through database log reporting
- Identify and fix problems with your databases
- Undo or roll-back changes to data and database structures utilizing log report

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

- Db2 Utilities (CHECK DATA/INDEX, Log analysis)
- IMS Utilities (Pointer checker)
- Catalog tools (Tivoli, IDCAMS, ISV products)
- VSAM Indexcheck, Datacheck

- DFSMShsm, DFSMSrmm tools
- RACF (IRRUT200), zSecure-Audit
- ISV software
- Custom-built programs

An example of CICS validation might include issuing CICS Commands such as CEMT against CICS Regions. The CICS Installation Verification Program “Catalog Manager” could also be run if it is available within the CICS Regions to validate.

An example of Db2 database validation might include performing DSN1COPY CHECKs against System and Applications tables.

An example of IMS validation might include providing Pointer Checker jobs to the IBM Z Cyber Vault Environment to validate the IMS Databases.

An example of MQ for z/OS validation might include running the Installation Verification Program ‘CSQ4IVP1’, the Dead Letter Verification Program ‘CSQUDLH’, or issuing commands to an MQ Queue via the ‘CSQUTIL’ program.

Backup and recovery

Backup and recovery are the most complicated areas of database management. Having the correct resources to perform a recovery is critical. Without them, 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 design must ensure:

- Process is Repeatable and Automated
- Time Consistent Copy is clean
- IPLed system is operational

The recovery design must account for surgical and catastrophic recovery. Surgical recovery would apply if only a small portion of the production data is corrupted and if consistency between current production data and the restored parts can be re-established. A catastrophic recovery would apply in the case of massive corruption to all or most of the data in the environment.

2.4.5 Offensive security environment

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

Here are some techniques used in offensive security:

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

The IBM Z Cyber Vault Environment provides a ideal setup to perform any of these activities, as it is not only an exact replica of the actual production system, but it is also isolated so that any action performed in this system will have no consequences outside of it. As a result of these exercises weaknesses and vulnerabilities might be identified, which can then be remediated in the actual production system.

Benefits of regularly executing offensive security practices are:

- Reduced Risk of Cyber Attacks
- Enhanced Security Posture
- Improved Incident Response
- Increased Security Awareness
- Cost Savings and Efficiency
- Competitive Advantage
- Regulatory Compliance and Risk Management

2.5 IBM Z Cyber Vault Automation considerations

There are essentially two levels of automation we will discuss here. Fully automated tasks are those that can be scheduled or triggered based on policy with no manual interaction. Although it would be ideal to fully automate all processes, that is not always possible. Partial automation includes scripts that can be manually triggered based on decisions made during incident management or daily operations. For example, the process of creating images on the recovery volumes and then activating images and initiating validation is fully automated. If you need to initiate a recover action and IPL for purposes of forensic analysis, that would require someone to manually initiate that action but the process itself is scripted.

In this section we will explore the automation considerations when designing or planning for the implementation of an IBM Z Cyber Vault solution. We will cover automation for Safeguarded Copy and related storage processing, recovery system operation, and data validation.

2.5.1 Safeguarded Copy and related storage processing

Managing, creating, recovering, and expiring Safeguarded Copy backups requires management software like IBM CSM or GDPS Logical Corruption Protection Manager. They coordinate and perform these steps automatically and with minimal impact to the host operations. They are described in 2.3, “IBM Z Cyber Vault Storage considerations” on page 27.

GDPS LCP Manager minimizes the risk of errors by reducing manual operations. To simplify the administration Safeguarded Copy captures, it has defined management profiles. A 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 LCP Manager 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 Safeguarded Copy(n) copy set.
- **RELEASE**: This statement performs a release of one or more expired captures from an FC(n) or Safeguarded Copy(n) copy set.
- **RESTORE**: This statement performs a restore from an FC(n) copy set to the RS volume set. When they are restored, it is then possible to access the RS(n) volume set for data analysis, extraction, or test purposes. The RESTORE operation is not supported for Safeguarded Copy backup sets.
- **RECOVER**: This statement performs a recovery to the recovery copy (RC) copy set from either an FC(n) or a Safeguarded Copy(n) copy set. When recovered, it is then possible to access the RC(n) copy set for data analysis, extraction, or test purposes.

2.5.2 Recovery system operation

The recovery system is where all validations, forensic analysis, and recovery planning is done. The first step is creating the images on the recovery volumes. Those are commands discussed in the previous paragraph. Once that is complete, the systems must be activated and IPLed. GDPS LCP Manager scripts can initiate these actions. For CSM environments, that automation will need to be created. For that reason, this section will focus on the capabilities included in GDPS LCP Manager. Understanding the GDPS capabilities will provide the background necessary to implement manual processes or custom automation where required.

Once the recovery volumes are created, you are ready to IPL the systems and run your validations. GDPS, using the Base Control Program internal interface (BCPii), can be used to automate the activation and load (IPL) of the IBM Z Cyber Vault Environment LPARs. Your GDPS script will include the activation of each LPAR in your sysplex or protected environment. The IPL process will leverage your existing IPL automation. It should include all your system and subsystem components. It is important that you start all of your database and middleware subsystems so that they go through their normal restart and recovery processing. If you do not have this level of IPL automation in place today, you will need to enhance your production IPL processes.

The IPL process for the cyber vault systems will need to differ slightly from their production versions and will need to accommodate changes in the I/O configuration used, as well as changes in PARMLIB members to facilitate the validation process. These changes are driven by the LOADPARM used to IPL the cyber vault systems and the introduction of a system symbol (in IEASYMxx) used to differentiate between IPLs of the systems in the production and cyber vault LPARs. The IBM Z Cyber Vault-specific members will need to be present in the production environment so that they are copied to the recovery volumes used to IPL the IBM Z Cyber Vault Environment.

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

2.5.3 Data validation

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

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

A report of the results of the validation process is sent to specific email addresses using SMTP or forwarded to a monitoring solution.

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

More information about data validation and the IBM Z Cyber Vault capabilities, see Chapter 3, “IBM Z Cyber Vault capabilities” on page 49.

3

IBM Z Cyber Vault capabilities

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

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

The following topics are covered in this chapter:

3.1 Key design considerations

Before making use of the IBM Z Cyber Vault supported capabilities it is important to consider the following:

- Which data is to be protected by the IBM Z Cyber Vault solution? This will help determine the storage and system capacity, as well as the location for protecting and validating the environment. The implementation of IBM Z Cyber Vault solution can be done in stages across those systems based on requirements and priorities of the application environments.
- Are there security or other business requirements that will prescribe the location for the IBM Z Cyber Vault environment and degrees of isolation for data and processing? These will drive the topology and locations of your storage and system environments for data validation.
- What is the desired capture frequency and retention of the data? The frequency of the Safeguarded backups, the length of time they are retained, and the measured change rate of the application storage environment will be used to determine the storage capacity required for the Safeguarded backups.
- What is your current architecture and replication environment? Using the current environment as a starting point, combined with the above decisions for fulfilling protection and recovery requirements, will help provide the next steps for planning and acquisitions needed to implement the blueprint.

3.1.1 Applying 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.

3.1.2 IBM Z Cyber Vault environment design

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

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

- 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 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, analysis and/or recovery of your production environments and data.

- To be compliant with security requirements consider using a *jump server*[1] or deploy additional IBM Z Cyber Vault capabilities, such as offensive security exercises or offline tape backups.
- Whether your IBM Z Cyber Vault environment will use an existing IBM Z system (such as a current disaster recovery (DR) environment) or your IBM Z Cyber Vault environment will run stand-alone in its own IBM Z system and possibly in a different location.

As part of the IBM Z Cyber Vault environment implementation,consider scheduling an IBM Z Cyber Vault workshop (your IBM representative or business partner can set this up for you), where subject matter experts can assess your production environment (storage, base z/OS, RACF, VSAM, CICS, Db2, Network, IMS and other subsystems) to determine the best possible environment setup. In that context, IBM can also help create the parameter files for the IBM Z Cyber Vault data validation asset as a service offering.

3.1.3 Developing and testing validation processes

Operational processes are important to effectively utilize the IBM Z Cyber Vault solution. Some of the areas to be addressed include defining the operational sequencing of daily IBM Z Cyber Vault processes and procedures. This includes automation strategies, scheduling regular daily operational validation, and checking validation status. The operation of the day-to-day IBM Z Cyber Vault processes are important, and can be fully automated (preferred) or performed by an individual or orchestrated across several operational actors. Safeguarded backups are taken at regular intervals and validation is performed one or more times a day. The capture frequency and validation frequency are independent. It is common to capture Safeguarded backups more frequently than performing validations. Validations can be performed on recovered Safeguarded backups or on FlashCopy of the source. You should strive to perform as many validations as possible during a 24-hour period. That number will depend on several factors, such as:

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

3.1.4 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, consistent, 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 CICS,IMS, Db2, and so on) that can be used for data validation and recovery purposes.

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

[1] A jump server is an intermediary device responsible for moving data through firewalls using a secure channel.

3.1.5 Testing the IBM Z Cyber Vault environment

After the IBM Z Cyber Vault environment is defined and verified, 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.

3.1.6 Production cutover and on-going monitoring

Production cutover and monitoring occur when the IBM Z Cyber Vault environment is started and performed 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 IT operations are always changing, the IBM Z Cyber Vault environment can be affected by these changes. On-going monitoring should include the proper connections to change management to ensure changes relate to the systems, applications, or data occur.

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

3.2 IBM Z Cyber Vault environment operational roles

Operational roles are critical to planning, deploying, and effectively managing an IBM Z Cyber Vault environment. 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 Business Continuity (BC) requirements and collaboration with application and line of business (LOB) teams are important to effective planning and design. This person is responsible for overall design and close collaboration with operations to ensure efficient and effective day-to-day operations of IBM Z Cyber Vault.

Business continuity representative

The person in this role represents the business requirements and works closely with the IBM Z Cyber Vault solution architect to ensure that business continuity risk management, and recovery requirements are well understood. They are 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 specialists and owners

Understand 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 specialists Perform 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 the DS8K, the overall storage, flash, mirroring, etc. They work closely with the IBM Z systems programming team and the IBM Z security administrators to set up and operate this environment.

Automation Administrator

Automation administrators (GDPS/CSM) play a key role in managing the scripts to run the validation 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, forward recovery, 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.

3.3 Environment setup

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

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.

3.3.1 Networking best practices

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

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

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

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

3.3.2 Required software and tools

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

IBM offers best of breed automation solutions that are policy based and provide the foundation for fully automating a Cyber Vault environment.

Required software should be installed in the production environment prior to starting the IBM Z Cyber Vault implementation. Because the IBM Z Cyber Vault environment is an exact replica of the production environment, the following software should included:

- z/OS and JES (SDSF REXX and REXX Runtime)
- IBM GDPS and Logical Corruption Protection Manager or IBM Copy Services Manager
- IBM Z NetView and IBM Z System Automation

3.3.3 Utilities for data validation by subsystem

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

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

Data Type	Validation Program
Type 2 (Data Structure)	
RACF Database	IRRUT200 (INDEX MAP)
ICF Catalog (BCS Only)	IDCAMS DIAGNOSE IDCAMS EXAMINE
ICF Catalog (BCS + VVDS)	IDCAMS DIAGNOSE[a]
VSAM-KSDS (Inc. AIX, VRRDS)	IDCAMS EXAMINE[b]
PDS	IEBCOPY UNLOAD
PDSe	IEBPDSE
CICS/TS	IDCAMS
MQ	CSQ4IVP1 CSQUTIL CSQUDLQH
Db2	DSN1COPY DSNUTILB[c]
IMS	FABPMAIN
Endevor[d]	NDVRC1(C1BM5000)
ADABAS[d]	ADAREP
Type 3 (Data Content)	
Application Data[e]	Client application dependent

a. DIAGNOSE with compare of each owned VVDS

b. Optional VERIFY also performed4 Optional INDEX CHECK for system and/or application databases

c. Data type validation under development
d. Requires separately licensed software (like IBM IMS High Performance Pointer Checker)
e. All Type 3 validation processes are client-developed and provided

3.4 Validation, forensic analysis, and recovery processes

In 1.4.2, “IBM Z Cyber Vault capabilities” on page 11 the capabilities of the IBM Z Cyber Vault solution are described. The 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 recovery point objective RPO and the recovery time objective (RTO). The implementation of the IBM Z Cyber Vault capabilities and the data validation frequency can help meet RPO and RTO targets in the event of a cyberattack.

Figure 3-1 depicts the IBM Z Cyber Vault capabilities as 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 and longer-term retention. 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.

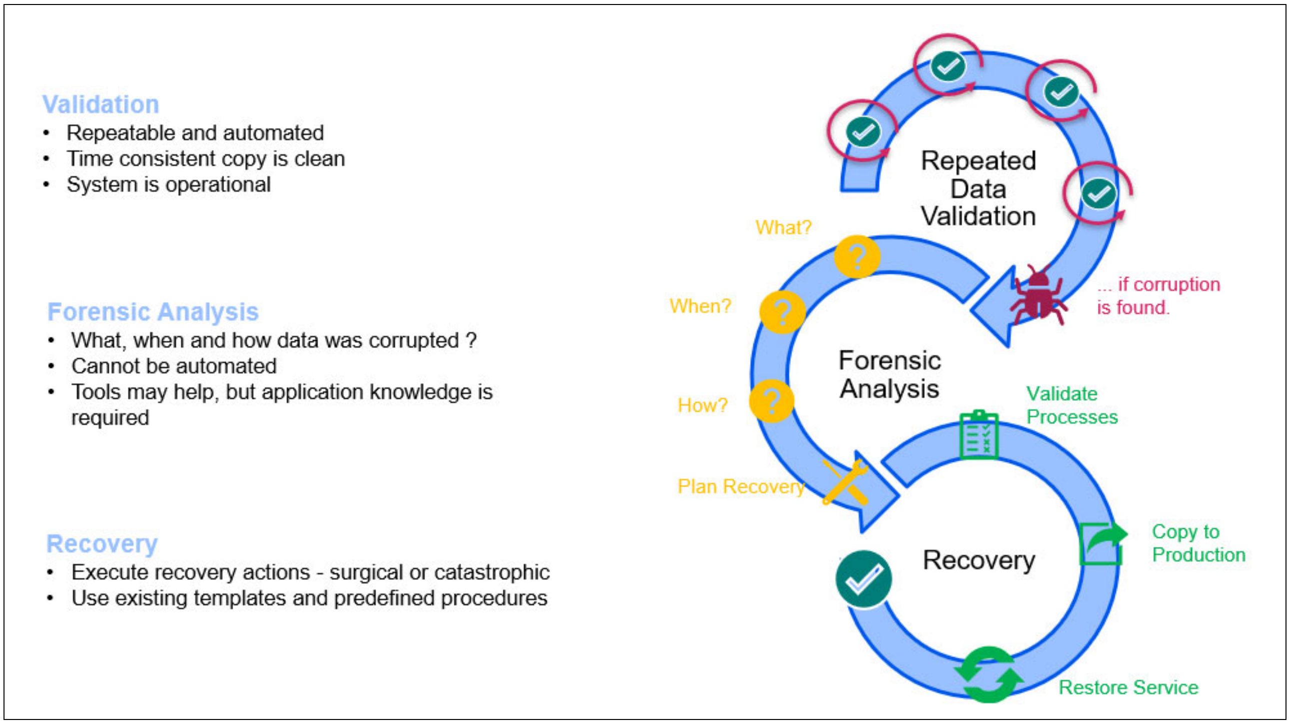

Validation
- Repeatable and automated
- Time consistent copy is clean
- System is operational

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 3-1 IBM Z Cyber Vault capabilities

3.4.1 Validation

When considering or implementing the IBM Z Cyber Vault capabilities, and specifically an IBM Z Cyber Vault environment, the software stack plays a fundamental role because it determines the strength of your system when facing cyberattacks and the resiliency level that 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 air gapped 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. The options include FlashCopy or Safeguarded backup for validation. Data corruption can affect the system's software structure or the application data.

The data validation process consists of three different types that are described in “Data validation” on page 11. These include:

- *Infrastructure validation* is the process of IPL’ing from a Safeguarded backup and validating no corruption has taken place that would prevent a successful system IPL. This process is defined as Type-1 validation.
- *Data structure validation* is the process of checking the z/OS system to search for potential structural corruption in any of the control files, configuration libraries, catalogs, repositories, file systems, database systems, or any other component that makes your production z/OS environment run. This process is defined as Type 2 validations.
- *Data content 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 process is defined as Type 3 validation.

Figure 3-2 on page 58 outlines the validation process.

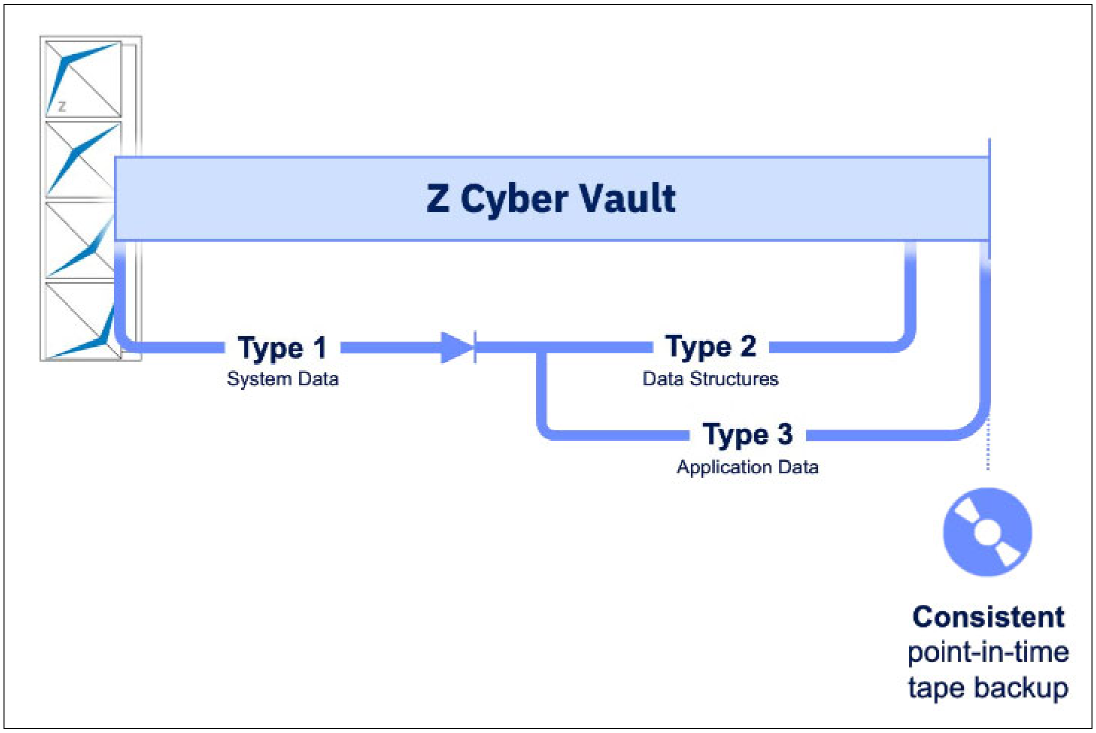

Figure 3-2 IBM Z Cyber Vault validation process

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 Forensic analysis). 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 you should be able to perform validation on as many Safeguarded backups as possible within the confines of your environment. In this way, you can detect a logical corruption earlier and reduce the impact of the corruption.

To accomplish this task, 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 3-3 on page 59).

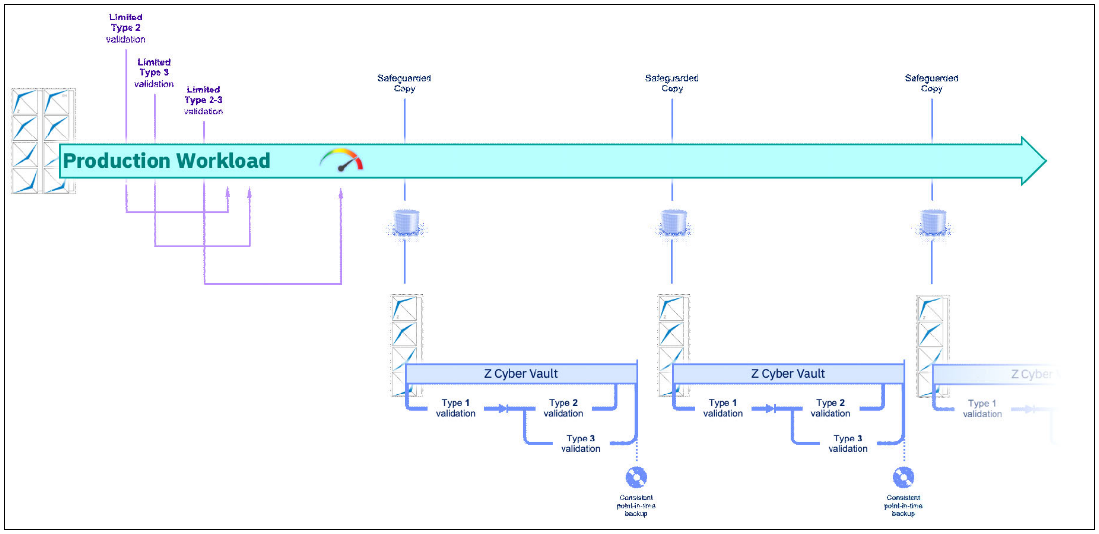

Figure 3-3 Data validation frequency

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

It is the combination of the correct tools and best practices that 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.

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

One key aspect of data corruption recovery is that in a data corruption scenario the decision to restore data is not driven by the operations team, nor can it be established in advance. A careful analysis of the applications involved, and the data lost, needs to be performed in conjunction with the lines of business. This is exactly what the forensic analysis is about. This process will determine what data can be recovered, what data makes sense to recover, and what data – corrupted or not – will be kept.

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

- You experience a malicious activity, at 10:15 AM in the morning.
- Your quick analysis shows that the cyberattack 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 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 this analysis, we need to consider the real impact of the cyber event. We could have discovered data corruption in some files, databases and/or applications, and we are already taking action based on this analysis to recover what’s needed. But how do we know that we have actually discovered all the corrupted data, or even worse, that we have effectively removed all the causes of data corruption? This requires an analysis that goes further than the immediate recovery of known data corruption and should continue even after the recovery actions.

During the forensic analysis process, you investigate problems and check which recovery actions must be carried out. To successfully do so, you must have a system that resembles the production environment at the time of the cyberattack. 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. You may have other IBM or non-IBM software that you are using in your production environment that will also help. Log analysis tools are some of the key utilities that you should consider for everything that creates a log in z/OS including (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 plan and procedure. All software tools that are required for this purpose should already be in the production environment since the IBM Z Cyber Vault environment is a replica of production, so you will only have available the tools already installed in the production environment. Also, it is necessary for some of these tools to be monitoring in real time what is happening in production, continuously collecting data that can then be used in the forensic analysis, and/or recovery of data and applications.

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

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

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 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.

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

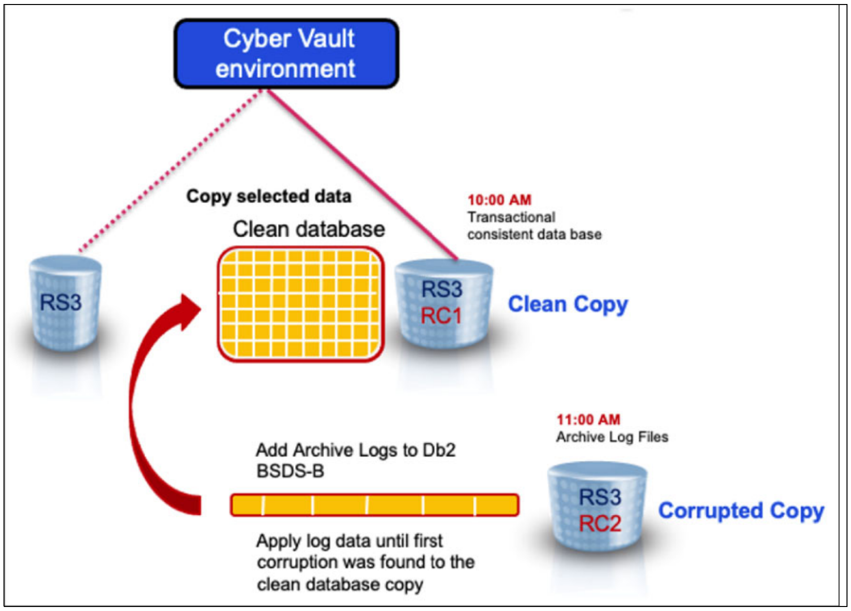
Cyber Vault environment
Copy selected data
Clean database
10:00 AM
Transactional consistent data base
RS3
RS3
RC1
Clean Copy
11:00 AM
Archive Log Files
Add Archive Logs to Db2 BSDS-B
RS3
RC2
Corrupted Copy
Apply log data until first corruption was found to the clean database copy

Figure 3-4 Forensic analysis in a Db2 scenario

3.4.3 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 ways to restore data back to production.

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 type of recovery can involve recovering data from different backups and media, and also running processes to re-create lost data.

This capability can be provided at a z/OS system level (see System-level surgical recovery) and for each of the transaction and data management subsystems in use. We selected Db2 to describe in Db2 subsystem surgical recovery 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 Cyber Vault environment that is located in the same datacenter as the primary production data, system-level surgical recovery can be done by copying the data from the recovery volumes to a set of staging volumes (see Figure 3-5) and making the staging volumes available to the production system LPARs 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.

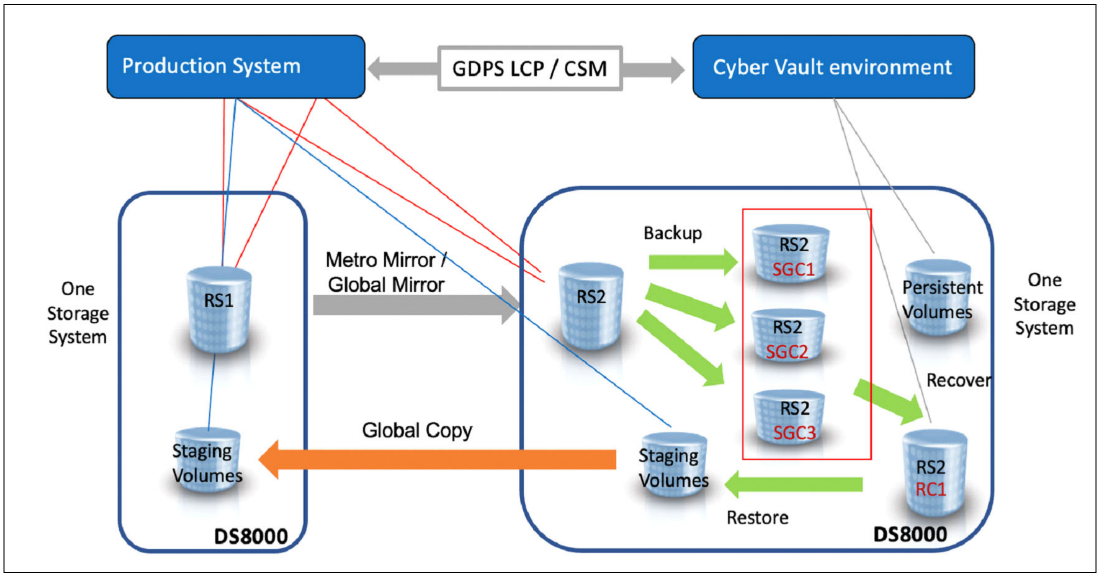

Figure 3-5 Using staging volumes for surgical recovery and partial restore to production

In an IBM Z Cyber Vault environment that is located in a remote datacenter, 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 will now describe typical surgical recovery scenarios for Db2 that can be used as an example and reference for what must be considered when designing an IBM Z Cyber Vault environment.

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

- Scenario 1: Your database system log files 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 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 backup, but no image copy to recover from is available. In this case, you must apply the available log data to the last clean version of the corrupted database to minimize the loss of data.

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

1. Recover the most recent Safeguarded backups that contain a copy of the clean database into the RC1 volume.
2. Recover the Safeguarded backups that contain the database log files into the RC2 volumes. Both sets of volumes, RC1 and RC2, must be online to the IBM Z Cyber Vault recovery system.
3. Copy the logs and BSDSs from the RC2 to the RC1 volumes.
4. Reassemble the Db2 zParms with DEFER ALL.

5. Apply the database log records to the database, up to the point right before the data corruption happened (Db2 LOGONLY recovery). This recovery point already was identified through the forensic analysis.
6. Determine the LRSN corresponding to the recovery point chosen above.
7. Use the DSNJ003 utility to update the ENDLRSN on all Db2 member BSDSs.
8. Run Db2 object LOGONLY Recovery (Cat/Dir and user objects).

If you use tools such as Db2 Log Analysis Tool for z/OS and IBM Db2 Recovery Expert for z/OS, there is no need to use two recovery sets of volumes. Only the latest Safeguarded 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.

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

3.4.4 Catastrophic recovery

If you must do a catastrophic recovery because the data corruption is extensive, a full restore of a Safeguarded backup is required. We understand this would rarely be the case as it would have a huge impact on the business if the recovery point is too far from the current time.

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

Since 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 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).

In this case the global copy is chosen to replicate data from the recovery volumes (RC1) to the production volumes (RS1) and is represented by the orange arrow (global copy) in Figure 3-6 on page 65.

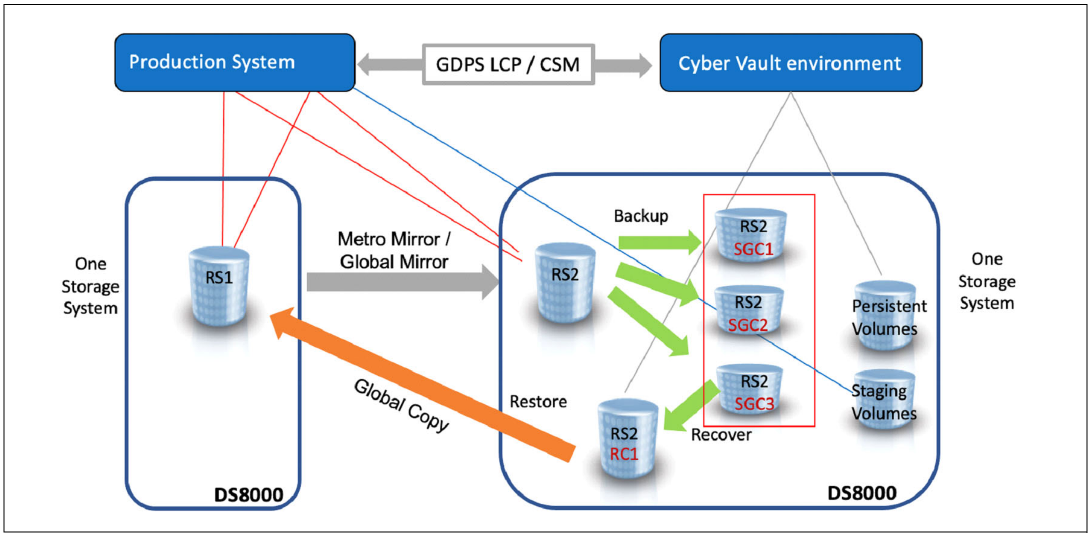

Figure 3-6 Restore to production

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

With earlier DS8000 storage systems, an incremental resynchronization from RC1 with global copy is not possible; therefore, a global copy 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 backup by using your RC1 volumes to run production 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.

Note: During IPL and running production from RC1, every I/O operation goes 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.

The following general steps are required to restore a valid Safeguarded backup to your production environment:

1. Select a validated Safeguarded 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 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 (metro mirror) to global copy.
4. Establish the global copy 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 backup until the data is copied and verified in production. The new Safeguarded backup might expire older backups. You do not want the older Safeguarded 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 global copy relationship. You cannot perform an IPL of the production system if the RS1 volumes are the global copy 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.

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

3.5 Extend the air gap

Tapes are durable and can last for 20-30 years, making them ideal for archiving data that requires long-term preservation. Additionally, they offer an air gap between the network and the data, safeguarding it from ransomware and malware. By offloading a validated and consistent point-in-time copy of the production environment onto tape and storing it in a secure, off-site physical vault, you create an impenetrable layer of protection, rendering the data virtually invulnerable to any cyber-attack.

3.6 Enhance security posture

To enhance an organization’s security posture, adopt a comprehensive approach that includes continuous vulnerability assessments, employee training, and robust incident response plans. An air-gapped environment, mimicking a real production system, serves as an ideal sandbox for conducting these assessments, learning from the results, and determining the necessary actions to safeguard and recover applications.

Offensive security, or “OffSec,” encompasses a range of proactive security strategies that employ tactics similar to those employed by malicious actors in real-world attacks. Common offensive security methods include red teaming, penetration testing, and vulnerability assessment. These operations are typically carried out by ethical hackers, cybersecurity professionals who leverage their hacking skills to identify and rectify IT system vulnerabilities.

Offensive security complements defensive security by enabling security teams to uncover and respond to unknown attack vectors that might otherwise go unnoticed. Furthermore, it offers a more proactive approach compared to defensive security, as offensive security measures proactively identify and address flaws before attackers can exploit them.

In essence, offensive security provides valuable insights that enhance the effectiveness of defensive security measures. It also reduces the burden on security teams. Consequently, offensive security has become an industry standard in certain highly regulated sectors due to its benefits.

4

Deploying the IBM Z Cyber Vault solution

This chapter is under construction.

5

The art of the possible with the IBM Z Cyber Vault solution

This chapter is under construction.

To determine the spine width of a book, you divide the paper PPI into the number of pages in the book. An example is a 250 page book using Plainfield opaque 50# smooth which has a PPI of 526. Divided 250 by 526 which equals a spine width of .4752". In this case, you would use the .5" spine. Now select the Spine width for the book and hide the others: **Special>Conditional Text>Show/Hide>SpineSize(-->Hide:)>Set** . Move the changed Conditional text settings to all files in your book by opening the book file with the spine.fm still open and **File>Import>Formats** the Conditional Text Settings (ONLY!) to the book files.

Redbooks

Getting Started with
IBM Z Cyber Vault

SG24-8511-01

ISBN

(1.5" spine)
1.5"<-> 1.998"
789 <->1051 pages

Redbooks

Getting Started with
IBM Z Cyber Vault

SG24-8511-01

ISBN

(1.0" spine)
0.875"<->1.498"
460 <-> 788 pages

Redbooks

Getting Started with IBM Z Cyber Vault

SG24-8511-01

ISBN

(0.5" spine)
0.475"<->0.873"
250 <-> 459 pages

Redbooks

Getting Started with IBM Z Cyber Vault

(0.2"spine)
0.17"<->0.473"
90<->249 pages

(0.1"spine)
0.1"<->0.169"
53<->89 pages

To determine the spine width of a book, you divide the paper PPI into the number of pages in the book. An example is a 250 page book using Plainfield opaque 50# smooth which has a PPI of 526. Divided 250 by 526 which equals a spine width of .4752". In this case, you would use the .5” spine. Now select the Spine width for the book and hide the others: **Special>Conditional Text>Show/Hide>SpineSize(-->Hide:)>Set** . Move the changed Conditional text settings to all files in your book by opening the book file with the spine.fm still open and **File>Import>Formats** the Conditional Text Settings (ONLY!) to the book files.

Redbooks

Getting Started with IBM Z Cyber Vault

SG24-8511-01

ISBN

(2.5” spine)
2.5”<->nnn.n”
1315<-> nnnn pages

Redbooks

Getting Started with IBM Z Cyber Vault

SG24-8511-01

ISBN

(2.0” spine)
2.0” <-> 2.498”
1052 <-> 1314 pages

SG24-8511-01

ISBN

Printed in U.S.A.