

Data Resiliency Designs: A Deep Dive into IBM Storage Safeguarded snapshots

Denis Olshanskiy

John Nycz

Jonathan Wilkie

Nezih Boyacioglu

Vasfi Gucer

Vineet Sharma



Storage



IBM Redbooks

**Data Resiliency Designs: A Deep Dive into IBM Storage
Safeguarded Snapshots**

September 2024

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

First Edition (September 2024)

This edition applies to IBM Storage Virtualize 8.7.

This document was created or updated on September 24, 2024.

© Copyright International Business Machines Corporation 2024. All rights reserved.

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

Contents

Notices	v
Trademarks	vi
Preface	vii
Authors	vii
Now you can become a published author, too!	ix
Comments welcome	ix
Stay connected to IBM Redbooks	ix
Chapter 1. Introduction	1
1.1 The importance of data resiliency	2
1.1.1 A three-tiered approach to data resiliency planning	2
1.1.2 Approaches to data resiliency	2
1.1.3 What is cyber resiliency?	5
1.2 Introducing Safeguarded snapshot	6
1.2.1 The business need for Safeguarded snapshot	6
1.3 Safeguarded snapshot as part of an overall cyber resiliency strategy	9
1.3.1 DR replication	9
1.3.2 Backup infrastructure architecture	9
1.3.3 Storage management hardening	10
1.3.4 Monitoring	10
1.3.5 Validation	11
1.3.6 Automation	11
Chapter 2. Safeguarded snapshot: Redefining data protection	13
2.1 Safeguarded snapshot core concept and components	14
2.2 Safeguarded snapshot core components	15
2.2.1 Safeguarded backup location	16
2.2.2 Volume groups	16
2.3 Safeguarded snapshots in action	18
2.3.1 Creating and configuring Internal Snapshot Policies: Overall flow and recommended practices	18
2.3.2 Creating and configuring Safeguarded snapshots: Overall flow and recommended practices	19
2.3.3 Near-instantaneous volume recovery and restore	21
Chapter 3. Designing with Safeguarded snapshot: Tailored solutions	27
3.1 Design scenarios	28
3.1.1 Protecting Oracle database with IBM Safeguarded snapshot	28
3.1.2 Building a Cyber Vault for critical MS SQL database	29
3.1.3 Enhanced resiliency for SAP HANA with IBM Storage Sentinel	31
3.1.4 Resilient data protection for a small business with VMware and FlashSystem 5300	32
3.2 Design considerations	33
3.2.1 Assessing your data protection needs	34
3.2.2 Designing your Safeguarded snapshot strategy	35
3.2.3 Integrating with security and management tools	35
3.2.4 Testing and validation	36

Chapter 4. Implementing Safeguarded snapshot.	37
4.1 Implementing a Safeguarded snapshot environment	38
4.2 Configuring the Safeguarded snapshot capacity	39
4.2.1 Configuring a Safeguarded snapshot location (optional)	39
4.3 Defining volume groups and Safeguarded policies.	42
4.3.1 Volume group creation	42
4.4 Creating a Safeguarded snapshot policy	47
4.4.1 Creating a Safeguarded snapshot policy	48
4.5 Managing snapshots	52
4.5.1 View a volume group snapshots	52
4.5.2 Adding a snapshot to a volume group	54
4.5.3 Restoring a volume group from a snapshot	56
4.5.4 Deleting a snapshot	60
4.5.5 Creating a clone or thin-clone from a snapshot	61
4.5.6 Refreshing a thin clone from a snapshot.	65
4.5.7 Converting a thin clone to a clone.	67
Chapter 5. Advanced considerations	69
5.1 Two person integrity	70
5.2 Performance considerations	74
5.3 Preserve parent.	75
5.4 Volume or snapshot deletion considerations	75
5.5 Integration with monitoring and automation tools	76
5.5.1 IBM Storage Defender	77
5.5.2 IBM Storage Sentinel	77
5.5.3 IBM Storage Insights.	77
Related publications	79
IBM Redbooks	79
Online resources	79
Help from IBM	79

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.

DS8000®	IBM Security®	Redbooks (logo)  ®
FlashCopy®	IBM Spectrum®	Storwize®
HyperSwap®	Insight®	Tivoli®
IBM®	QRadar®	
IBM FlashSystem®	Redbooks®	

The following terms are trademarks of other companies:

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

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

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

VMware, and the VMware logo are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

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

Preface

In today's digital landscape, ensuring the integrity and recoverability of critical data is paramount. This IBM® Redpaper delves into this essential concept, providing a comprehensive exploration of the IBM Safeguarded snapshot solution for data protection.

Safeguarded snapshot (called Safeguarded Copy prior to IBM Storage Virtualize 8.7) works by creating snapshots of your data volumes. These snapshots are then stored in a logically separate location, isolated from the original data. This isolation is key to its security; even if your main system is compromised, the attackers cannot get to the backups.

This book serves as a valuable resource for IT professionals seeking to safeguard their organization's information against various threats, including accidental deletions, ransomware attacks, and system malfunctions.

The target audience of this book is IT managers and directors, storage administrators, system administrators, storage solution architects and business continuity and disaster recovery (BCDR) professionals.

Authors

This paper was produced by a team of specialists from around the world.




Denis Olshanskiy is a seasoned Storage Specialist with a master's degree in Mechatronics, Robotics, and Automation Engineering from Budapest University of Technology and Economics. His expertise lies in storage area networks (SANs), data center management, and storage solutions. Notably, his skills extend beyond traditional storage to include proficiency in Arduino and Linux, demonstrating a well-rounded approach to technology.



John Nycz is an Advanced Subject Matter Expert for IBM Storage Virtualize and IBM FlashSystem®. With over 10 years of experience in systems management, networking hardware, and software, he has been with IBM for more than 20 years. John has served on numerous development, project management, and support teams. For the past seven years, he has been a Member of IBM's Storage Virtualization Support Team.



Jonathan Wilkie is an Advanced Subject Matter Expert/L3 support representative for IBM Spectrum Virtualize and IBM FlashSystem. He has more than 20 years of experience in IBM storage technical support. Over his career, he has provided technical support for Shark, DS4000, DS6000, and IBM DS8000® products. He has been supporting IBM Storage Virtualize-based products since 2010.



Nezhir Boyacioglu is an experienced SAN Storage Specialist with over 20 years of IT experience. He currently leverages his expertise at IBM Premier Business Partner, Istanbul Pazarlama, in Turkey. Nezhir's IBM storage journey began with Tivoli® Storage Manager (Spectrum Protect) and tape systems. For the past 10 years, his focus has shifted to the IBM Spectrum® Virtualize family (encompassing IBM SAN Volume Controller, and IBM FlashSystem) and Storage Area Networks (SANs). Nezhir's commitment to expertise is evident in his IBM certifications, including Enterprise Storage Technical Support, Flash Technical Solutions, Virtualized Storage, and Spectrum Storage software. He is a co-author of several IBM Redbooks® publications.



Vasfi Gucer leads projects for the IBM Redbooks team, leveraging his 20+ years of experience in systems management, networking, and software. A prolific writer and global IBM instructor, his focus has shifted to storage and cloud computing in the past eight years. Vasfi holds multiple certifications, including IBM Certified Senior IT Specialist, PMP, ITIL V2 Manager, and ITIL V3 Expert.



Vineet Sharma is a Master Certified Technical Specialist in Storage Systems and has expertise in several storage system technology areas in IBM Expert Labs, MEA. He works on planning, configuration, and implementation of the IBM Storage Virtualize family and IBM Enterprise Storage (DS8000®). He has delivered many complex solutions and complex data migrations for block storages. He actively works with IBM Business Partners and IBM clients for technology enablement and storage bootcamps. He has over 15 years of experience in IBM storage products such as IBM DS4000 series, IBM Storwize® Series, IBM SVC, IBM FlashSystem, IBM DS8000, IBM A9000 and IBM Software products like IBM Spectrum Control, IBM Copy Services Manager, and Storage Insight®.

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

Elias Luna, Andrew Greenfield, Yves Santos, Byron Grossnickle
IBM USA

Lucy Harris, Evelyn Perez, Chris Bulmer, Chris Canto, Daniel Dent, Bill Passingham, Nolan Rogers, David Seager, Russell Kinmond
IBM UK

Diana Laura Silva Gallardo
IBM Mexico

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 papers to be as helpful as possible. Send us your comments about this paper 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>



Introduction

In this chapter we introduce the importance of data resiliency, cyber resiliency and Safeguarded snapshot. We also address the limitations of other data recovery solutions and therefore the need for Safeguarded snapshot to protect data.

This chapter has the following sections:

- ▶ “The importance of data resiliency” on page 2
- ▶ “Introducing Safeguarded snapshot” on page 6
- ▶ “Safeguarded snapshot as part of an overall cyber resiliency strategy” on page 9

Safeguarded Copy is now Safeguarded snapshot: In IBM Storage Virtualize 8.7, the functionality previously called *Safeguarded Copy* is now referred to as *Safeguarded snapshot*.

This terminology change reflects a focus on the snapshot nature of the backups created by this feature. It highlights that these are point-in-time (PiT) copies of your data that are immutable and cannot be altered or deleted.

What about DS8000? The Safeguarded snapshot functionality discussed in this document applies to IBM Storage FlashSystem and IBM SAN Volume Controller.

See the IBM Redbooks *IBM Storage DS8000 Safeguarded Copy: Updated for DS8000 Release 9.3*, REDP-5506 for information on implementing this feature for DS8000.

1.1 The importance of data resiliency

In today's digital landscape, data resiliency is paramount for organizational success. It encompasses an organization's ability to effectively recover from data breaches and disruptions, ensuring business continuity. This involves the swift implementation of disaster recovery (DR) plans, the restoration of critical assets, and the ongoing protection of sensitive data.

Time is of the essence during data outages. Delays in recovery can significantly impact an organization's reputation, leading to customer distrust and potential financial losses. The severity of these consequences underscores the crucial role of data resiliency.

1.1.1 A three-tiered approach to data resiliency planning

A well-established approach to data resiliency planning involves considering disruptions of varying complexity:

- ▶ **Tier 1:** Simple recovery (for example, recovering a single, accidentally deleted file).
- ▶ **Tier 2:** Moderate recovery (for example, restoring a malfunctioning server).
- ▶ **Tier 3:** Highly complex recovery (for example, recovering from a ransomware attack and data theft, DR).

If an organization can handle all three types of disruption and still maintain normal business practices, that organization can be considered data resilient.

1.1.2 Approaches to data resiliency

Effective data recovery depends on having uncompromised backups available. These backups can come in various forms: PiT copies, array-based snapshots (like copies of your primary data, also called *primary workloads*), or data written to backup applications that store them in repositories like disks, tapes (including virtual tape libraries (VTLs)), or the cloud (also known as *secondary workloads*).

What is the difference between primary and secondary workloads?

Primary workloads refer to your critical data, the core information your system uses to function. Examples include:

- ▶ Operating system files.
- ▶ Database tables containing essential business data.
- ▶ User files and documents that are vital for daily operations.

They are typically backed up frequently (for example, hourly, daily) to minimize data loss in case of an incident.

Secondary workloads refer to data that is less critical but still valuable. Examples include:

- ▶ Logs and historical data.
- ▶ Development and testing environments.
- ▶ Backups of non-essential user files.

Backups of secondary workloads are considered lower priority than primary backups. They may be less frequent (for example, weekly, monthly) depending on the importance of the data. The focus here is on long-term retention and DR scenarios rather than immediate accessibility.

However, these recovery options rely on two critical assumptions that may not hold true in all scenarios, particularly ransomware attacks:

- ▶ **System availability:** A functioning system is needed to initiate the recovery process.
- ▶ **Backup integrity:** The backups themselves must be uncompromised, meaning the data they contain is not locked, corrupted, or encrypted.

Challenges of backup-based recovery

Both the primary and secondary copies that are available for use in recovery scenarios must be free of contamination to remove the risk of a repeated attack and reinfection.

Backup-based recovery has the following challenges:

- ▶ **Scalability:** Restoring massive datasets across multiple systems can be a time-consuming process.
- ▶ **Speed:** Retrieving data from tape-based systems, or VTLs, is significantly slower due to inherent mechanical limitations.
- ▶ **Immutable storage considerations:** Backups stored in immutable storage (unchangeable data repositories) require additional validation before recovery to ensure they have not been infected.

Figure 1-1 on page 4 shows the challenges of backup-based recovery.

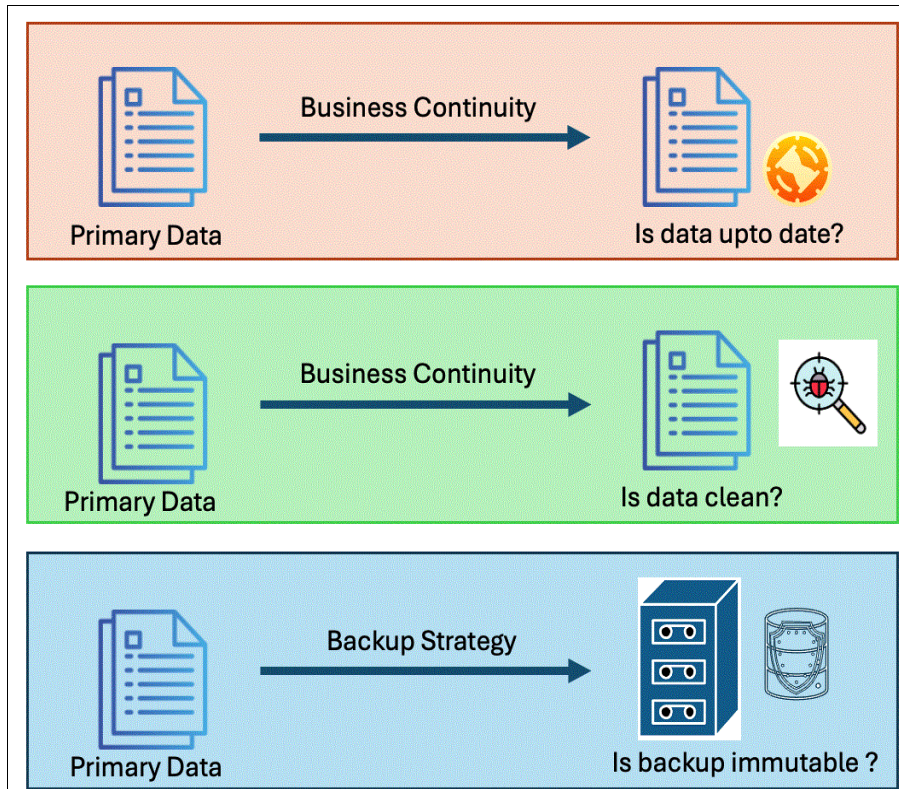


Figure 1-1 Challenges of backup-based recovery

Even when DR systems and restored data are available, there can be a lack of confidence in running business services at the secondary site. This stems from three main limitations:

- ▶ **Insufficient testing:** DR testing often falls short of real-world scenarios. It frequently focuses on isolated system or application failovers, not complex interactions between systems and interfaced components.
- ▶ **Non-representative testing:** Testing may involve gracefully stopping services at the primary site before starting them at the secondary. This does not reflect the chaos of actual disasters, potentially leading to data inconsistencies or service disruptions.
- ▶ **Limited transaction processing:** Testing on the secondary site might not include realistic workloads. Transactions may be run for a brief period, providing a false sense of security compared to sustained real-world operations.

These limitations can create a significant gap between recovery capabilities and actual business continuity.

Clients can use data from traditional DR testing such as the priority and sequencing of applications, and data and infrastructure dependencies, in support of the recovery of business-critical systems and applications and elements of Service Management. More mature traditional DR testing might consider the following aspects:

- ▶ Recovery from loss of data.
- ▶ Restarting systems.
- ▶ Recovering and restarting applications.
- ▶ Synchronous or asynchronous data mirroring.
- ▶ Recovery point objective that is greater than zero (RPO > 0).
- ▶ Restarting business services to an earlier PiT.
- ▶ Complexities associated with microservices.

- ▶ Distributed systems and the synchronizing of data across system boundaries.
- ▶ Interfaces with third parties.

1.1.3 What is cyber resiliency?

All businesses face the risk of cyberattacks, which often target critical applications. These attacks can encrypt, steal, or do both to data and applications. Traditionally, DR and business continuity efforts focused on mitigating software and hardware failures. Businesses achieved this by designing redundancy into systems and storage and utilizing technologies like backups and data replication to prevent data loss. Many businesses are not prepared for, or are unaware of, the extent of damage that a cyberattack can cause. They are also unaware of the costs of recovery from a cyberattack. Many of the businesses that do take steps to guard against cyber-attacks focus efforts on prevention and not how to recover quickly from an incident.

Cyber resilience is a business's ability to withstand and recover from disruptions, including cyberattacks and natural disasters, while maintaining normal operations. This translates to reduced financial losses and protects the organization's reputation. *A cyber-resilient company has a competitive advantage because of efficient and effective operations and can maintain or even grow business during a crisis if its competitors cannot.*

Cyber resiliency versus data resiliency: There is a subtle difference between cyber resiliency and data resiliency. Data resiliency focuses on protecting data from accidental loss or corruption. Cyber resiliency, on the other hand, is an organization's ability to prevent, withstand, and recover from cyberattacks, minimizing disruption to business operations.

In essence, data resiliency is a subset of cyber resilience. While data recovery is essential, cyber resilience takes a more holistic approach to ensure your organization can withstand and recover from various disruptions, including cyberattacks.

Cyber resiliency includes cybersecurity as a component. Cybersecurity is the methods or practices that an organization uses to protect its systems and critical information from digital attacks. It is also known as Information Technology (IT) security. Cybersecurity measures are designed to combat threats against applications and networked applications. These threats can come from both inside and outside of organizations.

Cyber resiliency begins with a *strategy or plan*. This strategy identifies the critical assets that matter most to the organization and its stakeholders. Assets include the data or information that must be protected and is critical to the function of the organization, and the systems and services that matter most. The strategy must also include identifying the vulnerabilities and risks and organization faces.

The next part of cyber resiliency is the *design*. Design work chooses the controls, procedures, and training that are appropriate to prevent harm to critical assets. However, the design must be practical. An impractical design that cannot or will not be implemented is not an effective one. The design work should also identify who has what authority to make decisions and act on them.

After the design is complete, the organization transitions to a *test operational state*. This tests where possible and closely monitors critical assets where it is not possible to test beforehand. The monitoring identifies when critical assets from the design phase are impacted by internal or external action. The design may be refined based on testing results.

After testing is complete the organization moves to an *operational state*. In this phase, the design has been deployed. There is still testing being done using controls to ensure that the operational state is effective and consistent.

From the operational state, an organization with a mature cyber resilient design will move to *evolution*. Environments are constantly changing with new threats and new technologies. Organizations will learn from incidents and how they recover from them. They will need to modify procedures, training, and even strategy as they learn.

1.2 Introducing Safeguarded snapshot

IBM Safeguarded snapshot, called *Safeguarded Copy* prior to IBM Storage Virtualize 8.7, integrated within the IBM Storage FlashSystem and IBM SAN Volume Controller family, empowers you to create cyber resilient, PiT copies of volumes. These copies are immutable, meaning they cannot be altered or deleted, even by user errors, malicious actions, or ransomware attacks.

Key features and benefits of IBM Safeguarded snapshots include:

- ▶ **Immutability:** Safeguarded snapshots are logically isolated, read-only¹ and cannot be altered or deleted, ensuring data integrity and protection from unauthorized changes.
- ▶ **Cyber resilience:** Designed to quickly recover from cyber-attacks from protected copies of the data².
- ▶ **Rapid recovery:** Enables rapid data recovery in the event of data loss or corruption.
- ▶ **Integration with automation tools:** IBM Safeguarded snapshots can be integrated with various automation tools, including IBM Copy Services Manager (CSM), IBM Storage Copy Data Management (SCDM), Ansible, external custom scripts, and an internal scheduler, providing the flexibility to tailor data protection processes and automate the creation of regular protected copies to meet specific needs.
- ▶ **Integration with security tools:** IBM Safeguarded snapshots can be integrated with various security tools, such as:
 - **SIEM (Security Information and Event Management) tools:** Correlates snapshot data with other security logs to detect anomalies and potential threats.
 - **SOAR (Security Orchestration, Automation, and Response) platforms:** Automates incident response actions based on snapshot data and security alerts.
 - **EDR (Endpoint Detection and Response) solutions:** Leverages snapshot data for forensic investigations and threat hunting.

1.2.1 The business need for Safeguarded snapshot

Safeguarded snapshot on Storage Virtualize supports the ability to create cyber-resilient PiT copies of volumes that cannot be changed or deleted through user errors, malicious actions, or ransomware attacks.

¹ It is read only by the storage subsystem itself. A host cannot read from this copy.

² Cyber resiliency focuses on your organization's ability to recover from a cyberattack and continue operations with minimal disruption. Cyber security focuses on preventing cyberattacks and minimizing their impact by building defenses to keep attackers out and safeguarding your systems and data.

The business needs of Safeguarded snapshot are as follows:

- ▶ “Ransomware and cyberattacks” on page 7
- ▶ “Regulatory requirements” on page 7
- ▶ “Data loss threats” on page 8

Ransomware and cyberattacks

Recent high-profile ransomware attacks have crippled critical infrastructure, disrupted global supply chains, and caused significant financial losses for organizations of all sizes. These devastating attacks highlight the far-reaching consequences of cybercrime, compromising sensitive data, disrupting essential services, and placing a growing burden on individuals, businesses, and society as a whole.

Impact of ransomware attacks on recovery time

Minimizing downtime after a ransomware attack requires meeting *Recovery Time Objectives (RTOs)* or staying within the *Maximum Tolerable Period of Disruption (MTPD)*. Several factors impact recovery time:

- ▶ **Attack analysis:** Understanding the attack scope is crucial for efficient recovery.
- ▶ **Data retrieval and restoration:** Affected data size and location heavily influence speed. Recovering large datasets from tapes is slower than flash storage.
- ▶ **Reinfection prevention:** Ensuring complete malware removal before recovery minimizes risk.
- ▶ **Data validation:** Verifying data integrity, especially if not pre-validated, can add time.
- ▶ **Transactional consistency:** Restoring data to a specific point might require replaying transactions for consistency, especially across systems.

Safeguarded snapshots can significantly improve ransomware recovery. These are immutable backups, meaning they cannot be altered by ransomware. This ensures readily available, validated data for fast restoration.

Benefits of Safeguarded snapshot:

- ▶ **Faster recoveries:** Reduce downtime with immediate access to uncorrupted backups.
- ▶ **Reduced risk of data loss:** Immutable backups are tamper-proof, protecting data from ransomware encryption.
- ▶ **Improved efficiency:** Pre-validated backups eliminate the need for additional data checks after an attack.

Regulatory requirements

The regulatory landscape for data protection and cybersecurity is rapidly evolving to address the growing sophistication and frequency of cyber threats. While established frameworks like [Health Insurance Portability and Accountability Act \(HIPAA\)](#) Health Insurance Portability and Accountability Act (HIPAA) and [General Data Protection Regulation \(GDPR\)](#) remain foundational, new regulations like [Digital Operational Resilience Act \(DORA\)](#) and [National Institute of Standards and Technology Cybersecurity Framework \(NIST CSF\) 2.0](#) raise the bar for operational resilience and risk management.

DORA: Fortifying the financial sector

EU's DORA regulation strengthens cyber resilience of financial institutions by mandating robust risk management, incident reporting, and third-party risk management. This ensures a more stable and secure financial ecosystem.

NIST CSF 2.0: A flexible framework for all

The NIST Cybersecurity Framework has long been a valuable resource for organizations seeking to improve their cybersecurity posture. The updated version, CSF 2.0, offers even greater flexibility and adaptability, making it suitable for organizations of all sizes and across industries. It emphasizes a risk-based approach, encouraging organizations to tailor their cybersecurity strategies to their specific needs and risk profiles. This includes identifying critical assets, assessing vulnerabilities, implementing protective measures, detecting threats, and responding effectively to incidents.

Sector-specific guidance: A layered approach

Beyond general frameworks, industries have specific cybersecurity guidelines. The Federal Financial Institutions Examination Council's (FFIEC) [Appendix J \(App-J\)](#) provides detailed recommendations for financial institutions, emphasizing the protection of backup data and the importance of air-gapped architectures. The [National Association of Insurance Commissioners \(NAIC\)](#) offers similar guidance for the insurance sector, while the [European Banking Authority \(EBA\)](#) outlines specific requirements for European banks.

Safeguarded snapshots: An immutable record for regulatory compliance

Taken together, these regulations and guidelines highlight the need for a multi-faceted approach to cybersecurity. This includes not only protecting data through encryption, immutability, and retention, but also ensuring the resiliency of systems and processes through robust backup measures, incident response planning, continuous monitoring, and a strong emphasis on risk management. By adhering to these evolving standards, organizations can better safeguard their sensitive information, maintain business continuity, and protect themselves from the ever-present threat of cyberattacks.

Safeguarded snapshots offer several advantages in the context of regulatory compliance:

- ▶ **Demonstrate data integrity:** Regulatory requirements often emphasize data immutability, ensuring data cannot be altered after creation. Safeguarded snapshots, by their unalterable nature, provide verifiable proof that your data remains unchanged, meeting compliance mandates.
- ▶ **Facilitate rapid recovery:** Regulations often dictate data recovery timelines in case of cyberattacks or system failures. Safeguarded snapshots enable organizations to restore data quickly and efficiently, minimizing downtime and ensuring business continuity.
- ▶ **Support incident response:** Regulatory frameworks often require comprehensive incident response plans. Safeguarded snapshots provide a pristine copy of data from a specific point in time, aiding forensic investigations and facilitating swift recovery after an incident.

Data loss threats

In today's data-driven landscape, ensuring data accessibility and integrity is fundamental for organizational success across all industries. While controlled outages for maintenance pose minimal disruption, unplanned data loss can have severe consequences. Financial strain from recovery efforts, reputational damage from breaches, and business disruptions due to missing critical data are just a few potential ramifications.

Safeguarded snapshots emerge as a critical tool in this context. Unlike traditional backups, these create immutable copies of data, providing an additional layer of protection against the ever-present threat of ransomware attacks that attempt to encrypt and hold data hostage. Safeguarded snapshots offer a significant advantage: their unalterable nature ensures organizations possess a reliable, uncorrupted copy of their data for swift restoration, minimizing downtime and safeguarding business continuity.

Safeguarded snapshots and workloads:

You can leverage Safeguarded snapshots in two ways depending on the workload:

- ▶ **Primary workloads:** Integrate them as part of your primary workload backup process. This creates an additional, ransomware-proof copy of your most important data alongside your regular backups.
- ▶ **Secondary workloads:** Use them alongside traditional secondary backups for DR. Because they are unalterable, Safeguarded snapshots offer a reliable, long-term archive that is not susceptible to tampering, making them valuable for less frequently accessed data as well.

1.3 Safeguarded snapshot as part of an overall cyber resiliency strategy

While this document focuses on the components and use of Storage Virtualize Safeguarded snapshots, it is important to place the solution in a larger context to illustrate how this capability might fit into an overall cyber resiliency strategy. Some other components to consider are:

- ▶ 1.3.1, “DR replication” on page 9
- ▶ 1.3.2, “Backup infrastructure architecture” on page 9
- ▶ 1.3.3, “Storage management hardening” on page 10
- ▶ 1.3.4, “Monitoring” on page 10
- ▶ 1.3.5, “Validation” on page 11
- ▶ 1.3.6, “Automation” on page 11

1.3.1 DR replication

One of the key components of a sound cyber resiliency solution is the notion of an air-gap that insulates the *safeguarded* or immutable copy from access or tampering. Storage Virtualize accomplishes this by making the immutable copies unmountable. Since the volumes cannot be mounted by a host system, significant protection against tampering is provided. For further isolation, if the environment already has replication in place, the choice could be made to take the Safeguarded snapshot at the DR or replication target site. This has the added benefits not only of physical and network isolation but also removes any additional load on the primary volumes incurred by maintaining the point in time copies.

1.3.2 Backup infrastructure architecture

Another consideration for a holistic data protection strategy is integrating Safeguarded snapshots with backup solutions like [Storage Protect](#) for long-term retention. Since Storage Protect is generally designed to have a much longer retention than what we would want to implement with online Safeguarded snapshot in Storage Virtualize (due to performance and most importantly, capacity considerations), it will be useful to balance the two complementary capabilities in a holistic data protection strategy.

Moreover, many environments utilize traditional FlashCopies with the “backup” setting (incremental FlashCopies with copy rate greater than zero) in order to decouple the heavy read I/O from the primary volume. It may be advantageous to create the Safeguarded snapshots against these *backup FlashCopies* instead of the primary volume.

This approach offers several potential benefits:

- ▶ **Reduced primary volume load:** Safeguarded snapshot creation against the primary volume can impose some I/O overhead. Using existing FlashCopies minimizes this impact.
- ▶ **Efficiency:** Since these FlashCopies are already capturing incremental changes, leveraging them for Safeguarded snapshots might improve efficiency.

1.3.3 Storage management hardening

Storage Virtualize offers automated management for deleting Safeguarded snapshots based on policies or retention periods. However, users with the Security Administrator role can still manually delete these snapshots. To mitigate this potential for unintended or unauthorized deletion, the [superuser account can be locked](#).³

However, essential maintenance tasks require the superuser account, such as running various service assistant commands (**satask** or **sainfo**) or accessing the service assistant tool GUI. Striking the right balance between security and functionality is essential.

Here are some approaches:

- ▶ **Implement two-person integrity (TPI):** TPI is a security control mechanism that requires two authorized individuals to be involved when completing a critical or sensitive task.⁴ For more information, see 5.1, “Two person integrity” on page 70.
- ▶ **Granular access controls (ideal):** Implement controls within the Security Administrator role, differentiating between snapshot deletion and superuser-only tasks.
- ▶ **Dedicated credentials:** If granular controls are not feasible, consider creating dedicated maintenance credentials with minimal superuser privileges.
- ▶ **Superuser access process:** Establish a documented process for granting temporary superuser access for authorized maintenance activities.

By implementing these recommendations, organizations can leverage Safeguarded snapshots for data protection while maintaining appropriate security controls.

1.3.4 Monitoring

A key component to any cyber resiliency solution is intrusion detection. While this is mainly implemented in the network or application layer, there are also tools associated with storage that can provide some early warning as well as direct integration into Storage Virtualize. One such tool is [Storage Insights](#), which has traditionally been used to provide performance and capacity reporting capabilities to storage environments, and in the case of Storage Insights, provide additional support enhancement benefits such as reducing the amount of time needed to create tickets and uploading logs to support.

³ IBM Storage Virtualize 8.6.0 and later.

⁴ TPI requires that you have two users with the Security Administrator role, so you cannot lock the superuser account if you are using TPI.

Given that these tools already have the capability to monitor the storage environment, they are perfectly positioned to detect sudden changes in storage consumption and decreased compressibility as would be indicative of an application level encryption-based ransomware attack.

Another IBM monitoring tool that is capable of correlating events that might be indicative of an intrusion is [IBM Security® QRadar® XDR](#). Both these tools, when incorporated into an overall cyber resiliency solution with the rapid recovery of Safeguarded snapshots, could provide a valuable alerting solution that might vastly reduce recovery time.

IBM Ransomware protection solutions: For detailed information on IBM Ransomware protection solutions, see [Ransomware protection solutions](#). Also, refer to IBM Redpaper *IBM FlashCore Module (FCM) Product Guide: Features the newly available FCM4 with AI-powered ransomware detection*, REDP-5725 for a detailed discussion on ransomware threat detection within IBM Storage Virtualize.

1.3.5 Validation

A complete cyber resilience solution necessitates the ability to validate the integrity and recoverability of backup copies. This ensures that backups can be utilized effectively in the event of a disaster or cyberattack.

There are various methods for backup validation, the most common being filesystem-level validation:

- ▶ **Filesystem validation:** For volumes formatted as file systems, the recovery volume can be mapped to a dedicated validation host. This host then verifies whether the filesystem can be mounted and accessed as expected.
- ▶ **Application-level validation:** In some cases, additional validation beyond the filesystem level may be necessary. This could involve testing backups with specific applications to ensure functionality.
- ▶ **Checkpoint files and alternative methods:** Other strategies like utilizing checkpoint files might also be employed for validation purposes.

Regardless of the chosen method, periodic validation of backups is essential for maintaining a robust cyber resilience posture. This practice, even beyond the context of cyberattacks, is a fundamental principle of good IT hygiene.

1.3.6 Automation

A complete cyber resilience solution extends beyond manual processes and embraces automation.

This begins with provisioning new application volumes. The automation framework should automatically configure these volumes, or their associated backups (FlashCopies or replication targets), to be placed in an appropriate volume group with a designated Safeguarded snapshot policy, specifying desired frequency and retention.

Furthermore, automation should encompass:

- ▶ **Periodic recovery and validation:** Regular automated recovery and validation of Safeguarded snapshots ensure their integrity and readiness for restoration.
- ▶ **Suspicious activity monitoring:** Automated monitoring for suspicious activity can trigger access lockdowns, mitigating potential threats.

- ▶ **DR automation:** If primary volumes are corrupted, automation can select the most recent valid Safeguarded snapshot and orchestrate its recovery to the original volumes.

By integrating automation throughout the life cycle, organizations can achieve a more cohesive and efficient cyber resilience strategy.



Safeguarded snapshot: Redefining data protection

This chapter introduces the IBM Storage Virtualize Safeguarded snapshot capability, a powerful tool for ensuring data resiliency and rapid recovery in your storage infrastructure. We break down the key elements that make up this function and how they work together to safeguard your data. We explore the entire process flow of creating Safeguarded snapshots, from automated scheduling with internal snapshot policies to crafting custom policies using the FlashSystem GUI. Finally, we delve into the recovery and restore capabilities, covering everything from full volume group restoration to granular, data-specific recovery options.

This chapter has the following sections:

- ▶ “Safeguarded snapshot core concept and components” on page 14
- ▶ “Safeguarded snapshot core components” on page 15
- ▶ “Safeguarded snapshots in action” on page 18

2.1 Safeguarded snapshot core concept and components

The Safeguarded snapshot capability leverages the volume group snapshot technology of FlashCopy®. Its core concept is to provide immutable, scheduled point-in-time (PiT) copies of your critical volumes or primary volumes. These snapshots cannot be accidentally or maliciously changed or deleted, offering resiliency from even from ransomware attacks. Also, Safeguarded snapshots optimize space utilization through thin-provisioning.

These safeguards are achieved through additional controls over PiT copies:

- ▶ **Separation of duties and access control:** This prevents unauthorized users from tampering with backup data. Only authorized roles with specific permissions on the FlashSystem or SVC can manage Safeguarded snapshots. See also [TPI](#), as discussed in 5.1, “Two person integrity” on page 70.
- ▶ **Safeguarded internal snapshot policies:** These are standard internal snapshot policies with an additional *Safeguarded* option enabled. They allow you to schedule snapshot creation and define retention periods.
- ▶ **Isolation of Safeguarded snapshots:** Starting from IBM Storage Virtualize 8.6, snapshot snapshots cannot be directly mapped to hosts. This ensures immutability as their content cannot be changed by users, applications, or hosts. Safeguarded snapshots can only be created for volumes or groups of volumes within a specific volume group with an *Internal Snapshot Policy*. The content of these Safeguarded snapshots resides in a dedicated *Safeguarded Backup Capacity* location. Accessing this data on hosts requires creating separate cloned volumes. Finally, only authorized users with access rights to the Safeguarded Backup Capacity can remove Safeguarded snapshots.

Scheduling and automation for Safeguarded snapshots are available in three ways:

- ▶ **Internal functionality:** Starting with IBM Storage Virtualize 8.5.2, you can schedule and automate snapshot creation directly on the system.
- ▶ **IBM Copy Services Manager:** This tool offers centralized automation capabilities and additional features for using Safeguarded snapshots even with systems running code versions earlier than 8.5.2.
- ▶ **IBM Storage Copy Data Management (SCDM):** SCDM can also be used for automated creation and management of Safeguarded snapshots.

Safeguarded snapshots offer powerful and near-instantaneous recovery and restore capabilities for entire volume groups or individual volumes within a group.

These capabilities provide several benefits:

- ▶ **Granular control and data validation:** You can easily validate the integrity of your backups, both routinely and before attempting a recovery, ensuring the chosen PiT copy is in good condition.
- ▶ **Forensic analysis:** Safeguarded snapshots allow forensic analysis of affected data or volumes saved in any PiT copy. This is invaluable for investigations and research purposes.
- ▶ **Flexible recovery options:** You can recover a subset of volumes for specific data extraction, such as recovering individual files or data images, to your production environment. Also, you have the option to perform a full environment restoration to a specific point in time.

Key takeaways: The following are the core strengths of the Safeguarded snapshot feature:

- ▶ **Immutability:** This ensures that your backups cannot be accidentally or maliciously altered. Once created, they become fixed points in time, providing a reliable source for recovery.
- ▶ **Access and activity control based on separation of duties:** This prevents unauthorized access and tampering with your backups. By following the principle of separation of duties, only authorized users with specific permissions can manage Safeguarded snapshots, minimizing the risk of human error or security breaches.
- ▶ **Instant recovery capabilities due to FlashCopy Snapshot functionality:** Safeguarded snapshots leverage FlashCopy, a technology that creates PiT copies of your data with minimal disruption. This allows for near-instantaneous recovery of entire volume groups, individual volumes, or even specific sets of data.

2.2 Safeguarded snapshot core components

There are a number of components and FlashCopy Volume Group Snapshot functionality in the heart of Safeguarded snapshot capability. These configuration components can be logically divided into standard capacity virtualization components and components that allow implementation of the specific features of Safeguarded snapshots, such as, automated scheduling for snapshot creation, consistent snapshots from several logically related volumes using volume groups, administrative access controls, restore and recovery capabilities.

The standard capacity virtualization components hierarchy is as follows:

- ▶ **Storage pool or MDisk group:** This is the highest level in the hierarchy. It acts as a pool of physical storage resources, typically combining storage from multiple physical disks (MDisks).
- ▶ **Production volume (Source Volume, volume/vdisk):** This is a logical unit of storage that is provisioned from the storage pool. It provides a structured and formatted space for storing data and is visible to the operating system and applications. Components that are mentioned before are used to define and implement standard capacity virtualization of storage space into volumes for our host applications.

To implement Safeguarded snapshots as PiT backup of existing production volumes, additional components are introduced and must be configured to provide necessary control, automation, and functionality.

Here are the control and functionality-related components, specific for snapshots overall and Safeguarded snapshots specifically:

- ▶ Safeguarded backup location/child pool with Safeguarded backup location capability
- ▶ Volume groups
- ▶ Internal Snapshot Policy
- ▶ Roles overall, and role groups for separation of duties
- ▶ Clone and thin-clone

Here is a breakdown of each component and its role in Safeguarded snapshot.

2.2.1 Safeguarded backup location

Starting with 8.5.2, Safeguarded snapshots no longer require manual creation of a separate Safeguarded child pool as the backup location.

The location is now automatically maintained by the system when using Safeguarded functionality and assigning a policy to the volume group.

However, the option to define a Safeguarded backup location as a child pool with Safeguarded capability still exists. This might seem to offer more manual control over sizing for Safeguarded snapshots, but it requires careful attention to capacity management. Mistakes can lead to outages.

Note: You cannot define child pool targets for Safeguarded snapshots using the internal scheduler. To configure child pool targets, use the command line or Copy Services Manager.

If you choose to create a dedicated Safeguarded child pool for your Safeguarded snapshots, keep these important points in mind:

- ▶ **Fixed capacity:** A Safeguarded child pool with a standard parent pool has a fixed capacity. This space cannot be used by your production volumes.
- ▶ **Space exhaustion:** If the child pool runs out of space, your Safeguarded snapshot volumes go offline. This will also cause a cascading outage of your production volumes. See also 5.3, “Preserve parent” on page 75.

Recommendation: In most scenarios, it is recommended to use Safeguarded snapshots without defining a dedicated child pool. With free space available in the parent pool and no child pool that is defined, the system automatically manages space allocation, ensuring that everything stays online.

2.2.2 Volume groups

In Storage Virtualize, consistency and policy control for all snapshots, including Safeguarded snapshots, are implemented through logical entities called volume groups. These groups ensure that all member volumes reach a consistent state at the time of a snapshot, similar to consistency groups. Safeguarded snapshots leverage the FlashCopy Volume Group Snapshot function. Starting from 8.5, all snapshots (including nonsafeguarded ones) cannot be directly mounted to hosts, making them immutable (isolated) and preventing accidental modifications.

Volume groups serve two key purposes:

- ▶ **Consistency:** They guarantee consistency within a group of volumes during snapshots.
- ▶ **Policy management:** They provide an interface for applying various policies, including Internal Snapshot Policies (as discussed in “Internal Snapshot Policies” on page 17), to groups of volumes.

Note: Policy-based replication is also controlled by a policy.

This simplifies scheduling and automation for snapshot creation, allowing you to define frequency and retention period.

Can you define number of retained copies? *There is no option to define number of retained copies.* This is determined by the snapshot frequency and retention period. However, snapshots might be retained even after their designated expiration date if certain dependencies exist. For example, thin clones relying on those snapshots might prevent automatic deletion.

Volume groups also streamline recovery and restoration operations for individual volumes or entire, logically related groups of volumes. Each volume group can only be assigned one Internal Snapshot Policy.

Recommendation: For complex scheduling or application-consistent copies, consider external tools like CDM (Copy Data Management) and CSM (Copy Services Manager). Internal scheduling is a good starting point, but external tools offer more flexibility for intricate scheduling needs. Also, when you need to synchronize snapshot creation with other processes or systems (for example, Ansible, Puppet, or Chef), an external scheduler can often handle these complex integrations more effectively.

Take a large financial institution with multiple data centers, varying data regulations, and different data sensitivities (for example, customer data versus market data). Here, external scheduling helps manage:

- ▶ **Differing retention policies:** Scheduling snapshots based on data type and regulations (for example, hourly for sensitive data, daily for less sensitive).
- ▶ **Data center load balancing:** Distributing snapshot creation across data centers to avoid overloading any single location.
- ▶ **Integration with monitoring tools:** Integration with monitoring systems to adjust snapshot schedules based on system performance and storage utilization.

Each volume group can only be assigned a single Internal Snapshot Policy. Starting with 8.5.1.0, snapshots are managed entirely at the volume group level. Safeguarded snapshots no longer require a separate, predefined storage pool (child pool) for implementation. This location is now automatically managed behind the scenes.

Internal Snapshot Policies

Internal Snapshot Policies control the creation, retention (how long snapshots are kept), and deletion of snapshot backup copies within the volume groups they're assigned to. The system provides five predefined policies that you can choose from, and these policies cannot be removed. However, you can create custom policies with your desired settings or use the predefined options as templates.

Note: Snapshot creation frequency is specified in hours, days, weeks, or months and retention is specified in days.

Separation of duties

Separation of duties is one of the cornerstones of security and data protection. Selecting the *safeguarded option* during policy assignment to the volume group enables additional protective measures for the snapshots created with this policy. Access to Safeguarded snapshots is strictly controlled. Only authorized users or groups with specific roles can perform operations on them. These roles determine the permitted actions, such as creation, deletion, or restoration.

Superuser

The superuser has the highest authority and can:

- ▶ Perform maintenance and configuration actions unrestricted.
- ▶ Remove Safeguarded snapshots and policies.

As it has such a high authority, it is possible to disable this role for extra security. While disabling the superuser role is a good security practice, it is important to have a recovery plan in case authorized personnel are accidentally locked out. Here are the potential recovery options:

- ▶ **Remote Support (IBM):** IBM support might be able to assist in regaining access remotely.
- ▶ **Technician Port:** As a last resort, physical access to the machine and the Technician Port could be used for recovery.

Security Administrator

The Security Administrator role holds extensive privileges within the system. They can:

- ▶ Manage users and security-related tasks.
- ▶ Remove Safeguarded snapshots.

Tip: You have the option to not create a specific user account with the Security Administrator role. It is not mandatory, but skipping a Security Administrator for Safeguarded snapshots poses a security risk and can lead to management difficulties. A better option is to use two person integrity, as discussed in 5.1, “Two person integrity” on page 70 for more information on TPI.

System Administrator

This role and any roles below this level are restricted for day-to-day operations and tasks. System Administrator can:

- ▶ Provision and configure Safeguarded snapshots.
- ▶ Create and assign Safeguarded snapshot policies and assign them to volume groups.

Note: A user with the System Administrator role cannot remove or change any existing safeguarded copy snapshots.

2.3 Safeguarded snapshots in action

In this section we discuss Safeguarded snapshots in more details and provide some best practices.

2.3.1 Creating and configuring Internal Snapshot Policies: Overall flow and recommended practices

Three predefined snapshot policies are available for use with volume groups. These policies offer preconfigured settings for snapshot creation, retention, and deletion. They cannot be removed or modified directly, ensuring a baseline level of consistency and preventing accidental configuration changes.

Note that policies can be created by almost all group of users mentioned in “Separation of duties” on page 17, but the ability to modify policies for existing snapshots, especially Safeguarded snapshots, is restricted to users with specific permissions. This is a crucial security measure to ensure the integrity and immutability of your backups.

Snapshot policies creation flow

Perform the following steps to create snapshot policies:

1. Access the list of available policies through the **Policies** menu on the right side of the screen. Within this menu, you find the Snapshot Policies list displaying all options.
2. Here, you can either:
 - Assign existing policies to existing volume groups.
 - Create a custom policy using the dedicated **Create Snapshot Policy** button.
3. When creating a custom policy, choose a descriptive and specific name. This is especially important for complex environments to aid future identification and management.
4. Define the snapshot creation schedule:
 - **Frequency:** How often snapshots of specific volume groups should be created (for example, daily, hourly).
 - **Timing:** When the snapshots should be taken (specific time of day).
 - **Retention:** How long the snapshots need to be kept on the system before deletion.

Recommended practices

By following the following recommendations, you can effectively create and manage snapshot policies to optimize data protection and storage utilization for your IBM FlashSystem and SVC environment:

- **Conservative approach:** Start conservative, assess the impact on the system and get more aggressive from there. The biggest pitfall is starting out too aggressive and facing disappointment.
- **Storage pool capacity:** Evaluate the current capacity of your storage pool. This includes estimating the maximum capacity of your production volumes and the anticipated number of data changes within those volumes. Frequent snapshots or those kept for long periods consume more storage space.
- **Snapshot impact on storage:** Remember that snapshots utilize storage from the source volume pool. Therefore, the frequency of snapshots and their retention period significantly impact your storage capacity planning and utilization.
- **Security:** Limit who can modify policies for existing snapshots, especially Safeguarded snapshots. This ensures the integrity and immutability of your backups.
- **Complexity:** If your solution architecture is complex, consider using even more descriptive policy names for better organization.
- **Compliance:** Some regulations might mandate specific controls over data backups. Setting appropriate retention periods in your policies can help ensure compliance.

2.3.2 Creating and configuring Safeguarded snapshots: Overall flow and recommended practices

In this section we discuss Safeguarded snapshots in more details and provide some best practices.

Safeguarded snapshot creation flow

Perform the following steps to create Safeguarded snapshots:

1. Identify the volumes that require backup (source volumes). These volumes must be assigned to the volume groups.
2. Volume groups can be created independently. Volumes can be assigned to a volume group either during its creation or afterward.
3. Once volumes are assigned to a volume group, you can perform the following operations:
 - **Create a new Internal Snapshot Policy:** Define custom settings for snapshot creation, retention, and deletion.
 - **Choose an existing Internal Snapshot Policy:** Select a preconfigured policy from available options.
4. Assign the selected policy to the volume group. When applying the policy, ensure the **Safeguarded** option is selected. This creates immutable, tamper-proof snapshots for enhanced protection according to the policy settings.

Refer to Chapter 4, “Implementing Safeguarded snapshot” on page 37 for step-by-step instructions on performing these actions.

Figure 2-1 shows the Safeguarded snapshot configuration elements. You have two options for the Safeguarded Backup Location (capacity): automatic allocation from the production pool when assigning a Safeguarded snapshot policy, or defining a dedicated child pool with safeguard capabilities.

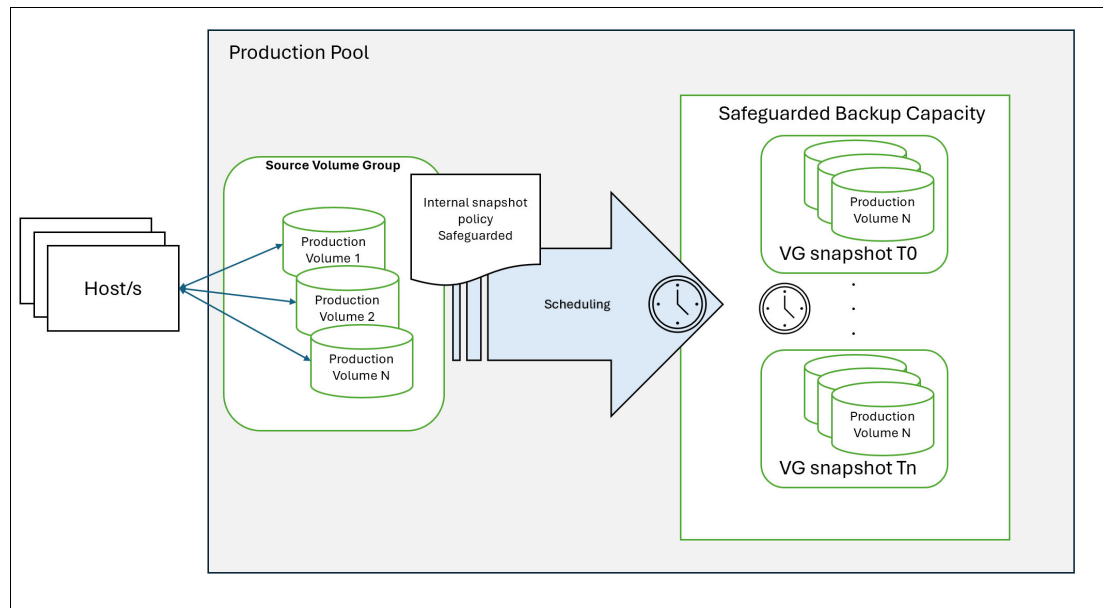


Figure 2-1 Safeguarded snapshot configuration elements

Things to consider and best practices

As mentioned in “Recommended practices” on page 19, Safeguarded snapshots utilize storage from the source pool for backups.

Here is an equation to estimate the required Safeguarded snapshot Backup capacity (SSB):

$$SSB = V \times DCR \times N \times R$$

Where:

- ▶ **SSB:** Safeguarded snapshot Backup capacity (TB)
- ▶ **V:** Total volume capacity (Volume Group) in TB
- ▶ **DCR:** Data Change Rate per copy (decimal)
- ▶ **N:** Number of snapshots per day
- ▶ **R:** Number of days the copies are retained

Determining Data Change Rate (DCR)

The most crucial factor in this equation is the Data Change Rate (DCR). You can obtain it through various methods:

- ▶ **System dashboards:** Check system dashboards for historical data change information.
- ▶ **Storage Insights:** Leverage Storage Insights for detailed data change rate analysis.
- ▶ **Snapshot measurement:** Create multiple snapshots at the same frequency and measure the data change between them.
- ▶ **Rule of thumb:** As a rough estimate, you can use 12-15% for the DCR x N portion of the equation. This approximates the daily data change rate.

Recovered volume capacity

It is generally recommended to allocate recovered volume capacity equivalent to the total volume capacity (volume group). This ensures sufficient space to restore entire volume groups in a disaster.

Performance considerations

The performance and latency of Safeguarded snapshots heavily depend on factors like:

- ▶ **Workload:** The type and intensity of activity on the production volume can impact performance. More frequent modifications require higher resource consumption for maintaining snapshot consistency.
- ▶ **FlashSystem model:** The capabilities of your specific FlashSystem model play a significant role in determining performance characteristics.

2.3.3 Near-instantaneous volume recovery and restore

Safeguarded snapshots offer a high level of protection by ensuring the immutability of your backups. That means:

- ▶ **Immutability:** Once created, Safeguarded snapshots cannot be accidentally or intentionally modified or deleted. This prevents unauthorized changes and protects your backup data integrity.
- ▶ **Isolation:** Since they cannot be directly mounted to hosts, Safeguarded snapshots are isolated from your production environment. This further safeguards your backups from potential corruption or accidental changes.

The system provides two main functionalities for recovering data from Safeguarded snapshots:

- ▶ Recovery copies
- ▶ Restoration

Recovery copies

Recovery copies with Safeguarded snapshots offer a versatile solution for various data access needs:

- ▶ **Testing and validation:** Quickly restore specific data points from a snapshot for testing or validating applications without impacting production environments.
- ▶ **Forensics investigations:** Easily access historical data from snapshots for forensic investigations or audits.

Data from recovery copies becomes available immediately by creating clones of the desired snapshot(s). These clones come in two options:

- ▶ **Thin clones:** Thin clones offer a space-efficient way to access data from snapshots for backup recovery. They are dependent on the source volume (the volume from the selected snapshot) for initial data. However, they provide several advantages:
 - **Immediate access:** You can map thin clones to hosts, allowing them to read data directly. This provides immediate access without waiting for a full copy to complete.
 - **Write redirection:** When a host writes data to the thin clone, those changes are saved only on the thin clone itself. The original source volume remains unmodified.
 - **Space efficiency:** Thin clones share data blocks with the source volume, minimizing storage consumption.
- ▶ **Full clones:** Full clones offer a different approach for backup recovery with Safeguarded snapshots:
 - **Independent copies:** Full clones create entirely new volumes based on the selected snapshot. These new volumes are independent of the source snapshot and source volume.
 - **Immediate access:** Similar to thin clones, data in full clones is available for use immediately. You can access the data while the copying process happens in the background.

Recovery copies allow you to restore:

- ▶ **Individual volumes:** Recover specific volumes from a snapshot for targeted data retrieval.
- ▶ **Entire volume group:** Restore the entire volume group to a specific point in time represented by the snapshot, ideal for disaster recovery scenarios.

Considerations for recovery

Figure 2-2 on page 23 shows the Safeguarded snapshot recovery process.

As mentioned in “Recommended practices” on page 19, it is important to consider capacity for the clones of the snapshot volumes before recovery.

In case of recovery production volumes can still be used by production hosts, while separated clones of snapshots can be used for other hosts activities. These clones are fully separated and any data changes on them do not change anything on the snapshots or production volumes.

Once the recovery is complete, verify the integrity and consistency of the recovered data. This ensures a successful restoration process.

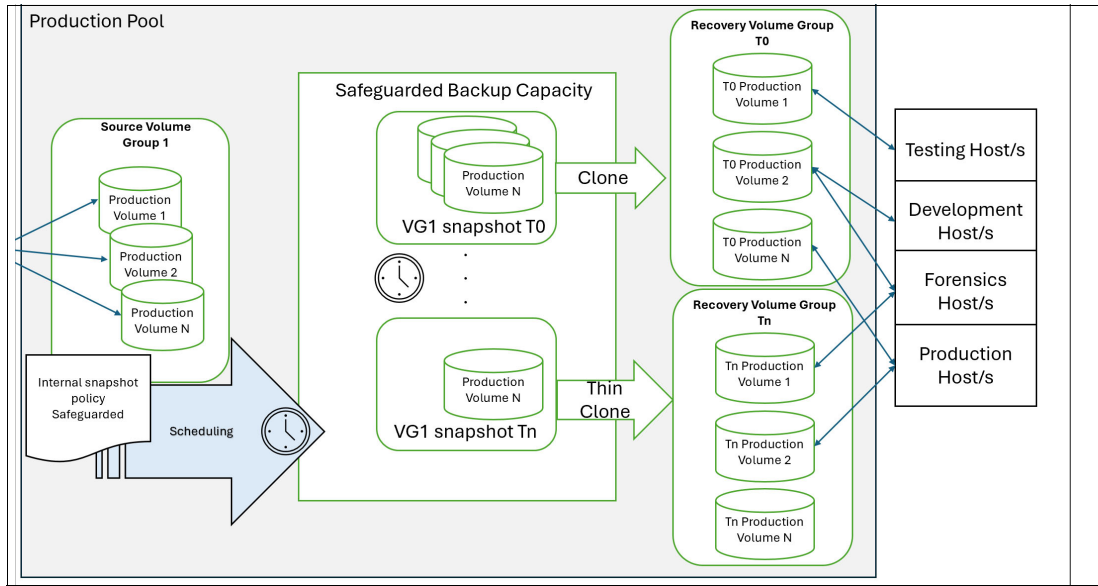


Figure 2-2 Safeguarded snapshot recovery process

Recovery flow

The following details the typical workflow for recovering data from Safeguarded snapshots:

1. In the Storage Virtualize GUI, navigate to the **Volume Groups** section. Choose the volume group that contains the desired snapshot.
2. Within the chosen volume group, locate the **Local Snapshots** tab. Browse through the available snapshots and select the one representing the point in time from which you want to recover your data.
3. Locate the **action menu that is associated** with the chosen snapshot. It is represented by three vertical dots on the right side of the snapshot entry. Select the **Clone** option (You cannot directly restore from Safeguarded snapshots. They are used to create clones for recovery).
4. Choose the clone type: **Thin Clone** or **Full Clone**.
5. Follow the on-screen instructions and prompts within the interface to complete the cloning process. The system creates a new volume group for the clones. You also have the option to select specific volumes within the snapshot for recovery.
6. Once the recovery is complete, you may need to present one or more recovered volumes to the hosts and mounting the volumes within the host environment.

For recovery, data is written to the clone volumes, which can later be used by the host for data consistency check before backup restoration, or for further granular data restoration (like specific file restoration) or forensic investigation.

Restoration

Restoration is used to restore the data back to the production volumes directly. It is possible to restore data to all volumes in the whole source (production) volume group or to choose a specific volume to restore from the snapshot.

Figure 2-3 shows Safeguarded snapshot restoration.

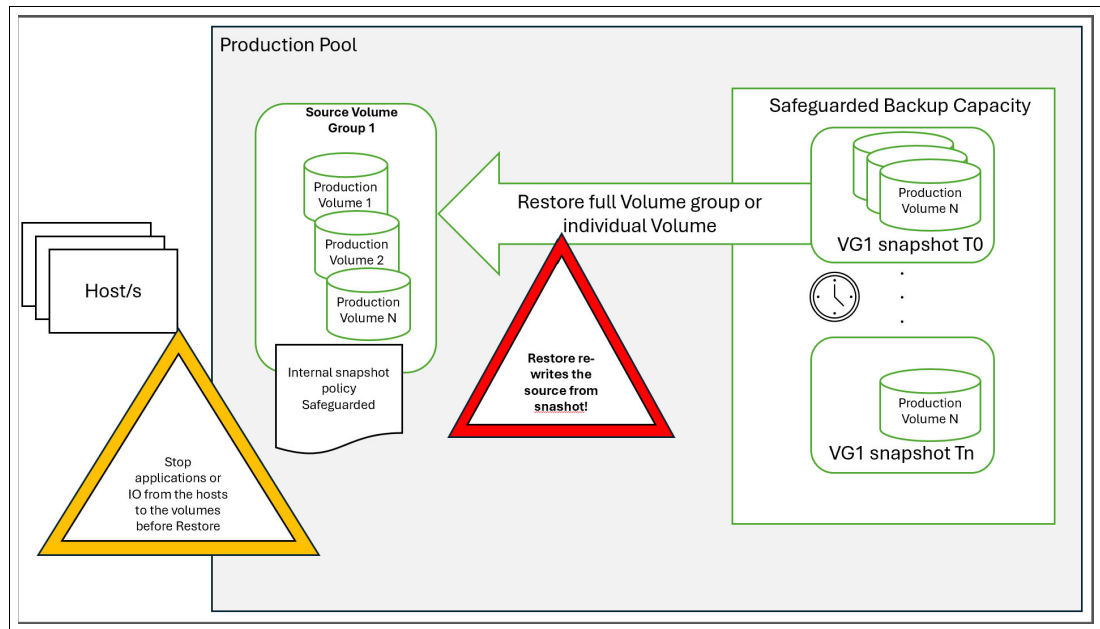


Figure 2-3 Safeguarded snapshot restoration

Considerations for restoration

Restoration is a critical process for recovering data in case of failures. However, it is important to be aware of these key considerations:

- ▶ **Disruptive activity:** Restoration rewrites data on the production volumes, potentially impacting ongoing operations. It's recommended to schedule restorations during maintenance windows or minimize downtime whenever possible.
- ▶ **Snapshot validation:** Ensure that you select a valid and recent snapshot for restoration to avoid restoring corrupted or outdated data.
- ▶ **I/O quiescing:** Before restoring, it is crucial to stop I/O operations (for example, unmount the volume) on the production volumes involved. Sometimes production volumes must be unmounted. This prevents data inconsistencies during the process.
- ▶ **Post-restore actions:** Depending on your operating system and applications, you might need to re-present the restored volume on the host level (re-mounting the volumes or rescanning the resources). This makes the restored data accessible to your applications.

Restoration flow

Before restoring a volume group or volume from a Safeguarded snapshot, it is crucial to ensure minimal disruption to ongoing operations:

- ▶ **Stop applications:** Suspend or terminate any applications currently accessing the volumes that are slated for restoration. This prevents data inconsistency during the process.
- ▶ **Quiesce I/O:** Stop all I/O operations directed at the volumes to be restored. This might involve unmounting the volumes depending on your specific operating system and application environment.

Important: Restore action overwrites data on the production volume/s or volume group from the snapshot.

1. In the Storage Virtualize GUI, navigate to the **Volume Groups** section. Choose the volume group that contain the desired snapshots. This will usually be the volume group where your production data resides.
2. Within the chosen volume group, locate the **Local Snapshots** tab, and choose the snapshot to recover from.
3. Locate the **action menu** associated with the chosen snapshot. It is represented by three dots on the right side of the snapshot entry. Select the **Restore** option.
4. The system presents a restore settings menu. Here, you can choose between:
 - **Restoring the entire volume group:** This recovers all volumes within the selected volume group to the state captured in the chosen snapshot.
 - **Restoring specific volumes:** This allows you to recover only a subset of volumes from the volume group based on your specific needs.
5. Follow the on-screen instructions and prompts within the interface to complete the restore process.



Designing with Safeguarded snapshot: Tailored solutions

This chapter describes various approaches to using IBM Safeguarded snapshot, providing tailored solutions for different environments and requirements. Whether you're managing a small business infrastructure like a point-of-sale system in a retail store, a large enterprise data center with mission-critical databases, or a hybrid cloud environment with virtual machines spread across on-premises and cloud resources, Safeguarded snapshot offers robust data protection and recovery capabilities.

Safeguarded snapshot is a powerful tool in your data protection suite, but its true potential is unleashed when it is precisely tailored to your organization's unique needs and environment. A one-size-fits-all approach simply will not suffice in today's complex and ever-evolving threat landscape. By understanding your specific recovery objectives, data sensitivity levels, storage infrastructure, and budgetary constraints, you can craft a Safeguarded snapshot solution that delivers the data protection you need, minimizes downtime, and ensures business continuity all while maximizing your return on investment and achieving the best possible value for your organization.

This chapter includes the following topics:

- ▶ “Design scenarios” on page 28
- ▶ “Design considerations” on page 33

3.1 Design scenarios

Real-world examples will illustrate how Safeguarded snapshots adapt to address diverse challenges, such as safeguarding mission-critical applications and defending against ransomware attacks. We will delve into specific scenarios, outlining the ideal configurations and best practices for each.

Protecting mission-critical applications

Mission-critical applications, such as financial systems or healthcare databases, are the heart of many organizations. They hold sensitive data, support essential operations, and directly impact revenue generation and customer service. Any disruption to these systems can lead to significant financial losses, reputational damage, and even legal consequences.

Therefore, ensuring the continuous availability and integrity of these applications is paramount. Traditional backup strategies, with their infrequent backups and lengthy recovery times, often fall short in meeting the stringent requirements of mission-critical environments.

Frequent, low-impact copies are crucial because they minimize the amount of data lost in case of a failure or outage. The shorter the interval between backups, the less data needs to be recreated or recovered, reducing the Recovery Point Objective (RPO) and ensuring that the most up-to-date information is available.

Moreover, low-impact copies are essential to avoid performance degradation on the primary system during the backup process. For critical applications, any disruption to normal operations can have severe consequences.

Rapid recovery is equally important as it minimizes downtime and ensures that the application is back online and operational as quickly as possible. This is where the Recovery Time Objective (RTO) comes into play – the shorter the RTO, the faster the application can be restored, minimizing the impact on business operations.

Frequent copies of critical applications

Customers can utilize Safeguarded snapshots to achieve far more frequent backups of critical applications compared to traditional backup systems. This is especially beneficial when dealing with massive datasets. Imagine a critical server holding 20 terabytes of data. Transferring such volumes over a network for hourly backups with traditional methods can violate business continuity requirements due to excessive transfer times. Even with incremental or differential backups, restoring from them involves applying multiple steps, significantly increasing recovery time. Safeguarded snapshots address these challenges by offering quicker and more efficient data protection, ensuring critical applications remain secure and readily available.

3.1.1 Protecting Oracle database with IBM Safeguarded snapshot

A large insurance company relies on a massive 57 TiB Oracle database running on an IBM Power Server to manage critical transactions and customer data. Full backups of this database traditionally took 16 hours, causing significant downtime and potentially millions of dollars in lost revenue if a critical system failure occurred during that window. Furthermore, restoring the database using their existing infrastructure could take up to 24 hours, hindering real-time access to data crucial for the insurance industry. In addition to the challenges of backup time, the company also sought a way to quickly create consistent database copies for testing and analytics purposes without impacting production performance.

The solution

To maximize data protection and storage efficiency, IBM Safeguarded snapshots were scheduled to run every 2 hours in conjunction with their existing IBM FlashSystem 7300 storage array and IBM Storage Copy Data Management (SCDM). SCDM orchestrated the snapshot process, ensuring data consistency and enabling fast, instant accessing of the database copies to the same or different servers within the data center. This solution dramatically reduced their RTO from 24 hours to just minutes and their RPO to two hours, ensuring rapid restoration of the database to a recent point in time in the event of an outage or data corruption.

Capacity calculation

- ▶ Database size: 57 TiB
- ▶ Daily change rate: 4%
- ▶ Change per snapshot (every 2 hours): 4% / 12 snapshots per day = 0.33%
- ▶ Estimated capacity consumption per snapshot: 0.33% of 57 TiB = 0.1881 TiB (approximately 192 GiB)

Total estimated capacity consumption for 3 days: 192 GiB/snapshot * 12 snapshots/day * 3 days = 6.912 TiB (approximately 7 TiB).

By scheduling Safeguarded snapshots every 2 hours and retaining them for 3 days, the customer ensured a high level of data protection with minimal data loss potential. This approach efficiently managed storage capacity consumption. This strategy, coupled with the significant reductions in RTO and RPO, allowed them to recover from any point in time within the 3-day retention period. This translates to greater flexibility, resilience, and business continuity in the face of data loss or corruption.

Outcome

By using IBM Safeguarded snapshot and SCDM, the company achieved dramatic improvements in data protection, recovery and operational efficiency. The solution's ability to create near-instantaneous, consistent copies of their massive Oracle database and quickly mount them for various purposes proved to be a game changer, enabling them to meet stringent service level agreements.

For consistent, protected snapshot integration with Oracle, see the IBM Redbooks [Cyber Resiliency with IBM Storage Sentinel and IBM Storage Safeguarded Copy, SG24-8541](#).

3.1.2 Building a Cyber Vault for critical MS SQL database

Recognizing the growing threat of ransomware attacks, a mid-sized financial institution, such as a credit union, heavily reliant on a 12 TiB Microsoft SQL Server database for its core operations, sought to fortify its data protection. To address this critical need, they implemented IBM Cyber Vault. Cyber Vault's air-gapped isolation capabilities create an unbreachable barrier for their critical data, ensuring it remains inaccessible even in the event of a sophisticated ransomware attack. Additionally, they integrated Cyber Vault with their existing QRadar SIEM for real-time threat detection and automated responses. This comprehensive solution empowers the financial institution with a robust defense against cyberattacks, ensuring their critical data remains secure and readily available.

Solution architecture

The solution implemented in this company is described below:

Isolated ESXi host

The institution established a dedicated, isolated ESXi host within their VMware environment. This host was physically and network-segmented from the production environment, providing a secure staging area for the protected data.

Scheduled and on-demand Safeguarded snapshots

IBM Storage Copy Data Management (SCDM) was configured to take Safeguarded snapshots of the 12 TiB MS SQL database every 3 hours. Additionally, it was integrated with QRadar to trigger unscheduled snapshots whenever QRadar detected suspicious activity in the network, providing an extra layer of protection.

Automated mounting, validation, and deep scan

Each morning at 8:00 AM, before the start of the business day, SCDM automatically mounted the latest snapshot to the isolated ESXi host. Upon mounting, a multi-layered validation process was initiated:

- ▶ **DBCC CHECKDB:** This script meticulously verified the logical and physical integrity of the database.
- ▶ **Deep Scan with EDR:** An Endpoint Detection and Response (EDR) solution was integrated to perform a deep scan of the mounted database, leveraging behavioral analysis and threat intelligence to detect any signs of malicious activity or hidden malware that may have evaded traditional antivirus scans.

Benefits

The benefits of the solution are as follows:

Enhanced ransomware protection

The air-gapped nature of IBM Cyber Vault, coupled with the immutable snapshots, made the financial institution's data virtually immune to ransomware attacks. Even if production data was compromised, the isolated copies remained untouched and readily available for recovery.

Rapid recovery and minimal downtime

In the event of a cyberattack or data corruption, the financial institution could quickly revert to a known good copy of the database within minutes, minimizing downtime and ensuring business continuity.

Data integrity assurance

The automated DBCC CHECKDB script and deep scan with EDR provided a comprehensive validation process, ensuring that the mounted database copy was free of errors, inconsistencies, and potential threats.

Proactive threat detection and response

The integration of QRadar and SCDM enabled the institution to detect and respond to advanced threats in real time. QRadar's ability to trigger on-demand snapshots ensured that a clean copy of the database was immediately available for analysis and recovery, minimizing the potential impact of an attack.

Outcome

By implementing IBM Cyber Vault and integrating it with a robust EDR solution and QRadar SIEM, the financial institution achieved a proactive, multi-layered defense strategy that significantly reduced their risk of data loss due to ransomware or other malicious activities. The solution's automated processes, comprehensive validation, and intelligent threat response capabilities provided a high level of confidence in the security and integrity of their critical data.

3.1.3 Enhanced resiliency for SAP HANA with IBM Storage Sentinel

A leading automotive supplier depends on an 11 TiB SAP HANA database hosted on Intel servers. They encountered significant challenges in maintaining the continuous availability and integrity of this essential data. Their current backup and recovery processes were time-consuming, exposing them to risks of data loss and prolonged downtime during system failures or cyber-attacks. Moreover, they lacked a proactive approach to detecting subtle data corruption or sophisticated ransomware attacks. Any disruption to the SAP HANA environment threatened severe consequences, including costly production delays, missed deadlines, and irreparable damage to their reputation with key customers.

Solution architecture

The solution implemented in this company is as follows:

Hourly Safeguarded snapshots

SCDM was configured to orchestrate hourly Safeguarded snapshots of the 11 TiB SAP HANA database and store these copies on the IBM FlashSystem 9500 for two days. These immutable snapshots ensured that a recent, recoverable copy of the database was always available, minimizing the potential for data loss in the event of an unexpected outage.

IBM Storage Sentinel integration:

IBM Storage Sentinel, an AI-powered anomaly detection solution, integrates with Storage Copy Data Management (SCDM) to continuously monitor the integrity of data within Safeguarded snapshots stored on the IBM FlashSystem 9500. Sentinel performs deep scans every three hours, ensuring the data is regularly checked for ransomware encryption and other anomalies. This scan schedule is strategically designed to balance the need for frequent monitoring with minimal resource consumption.

Benefits

The benefits of the solution are as follows:

Early detection of intermittent encryption attacks

IBM Storage Sentinel's unique ability to detect intermittent encryption attacks proved critical for the customer. These stealthy attacks attempt to encrypt data in small batches spread over time, often with long gaps in between, to evade traditional security measures, potentially leading to catastrophic data loss or disruption. Sentinel's AI-driven algorithms, however, identified the subtle patterns of such attacks and alerted the IT team promptly. This early warning, compared to traditional tools that might miss these intermittent attempts, allowed the customer to take immediate action and potentially avoid a major security incident.

Proactive identification of database corruption

Sentinel's deep understanding of the SAP HANA database structure enabled it to detect even small instances of data corruption, such as checksum inconsistencies or metadata errors.

This allowed the customer to proactively address these issues, preventing them from escalating into major problems and potentially causing service outages or data inconsistencies that could impact critical business operations.

Enhanced ransomware resilience

The combination of Safeguarded snapshots on the FlashSystem 9500 and Storage Sentinel provided a robust defense against ransomware attacks. Even if production data was compromised, clean and verified copies were readily available for rapid recovery, minimizing downtime and ensuring business continuity, potentially leading to reduced financial losses.

Continuous monitoring and peace of mind

The automated, continuous monitoring of critical SAP HANA data on the IBM FlashSystem 9500 by Storage Sentinel gave the customer peace of mind, knowing their data was safe and production processes protected from disruption. This allowed them to focus on their core business activities with confidence.

Outcome

By leveraging the capabilities of their IBM FlashSystem 9500 and implementing this comprehensive solution, the automotive supplier achieved significant improvements in the resiliency and security of their SAP HANA environment. This proactive approach to data protection resulted in reduced risk of downtime, improved operational efficiency, and a stronger foundation for their business operations and customer relationships. Additionally, they gained the ability to detect and respond to a wide range of threats, ensuring the integrity and availability of their critical data, ultimately minimizing the risk of costly production delays and damage to their reputation.

For more information on how to implement Storage Sentinel and SAP HANA integration refer to the IBM Redbooks [Cyber Resiliency with IBM Storage Sentinel and IBM Storage Safeguarded Copy, SG24-8541](#).

3.1.4 Resilient data protection for a small business with VMware and FlashSystem 5300

A small business with limited IT resources relies on a VMware cluster with multiple ESXi hosts and an IBM FlashSystem 5300 storage array for their primary data storage. However, their current backup solution, a slow near line SAS (NL-SAS) based NAS device, creates a critical bottleneck. Full backups can take several hours to complete, significantly impacting their daily operations and potentially leading to data loss or corruption during extended downtime. Recognizing these risks, the business urgently seeks a more efficient and reliable backup solution. However, budget constraints limit their ability to invest in a dedicated data protection product. This situation compels them to explore alternative functionalities within their existing infrastructure to address their backup challenges.

The solution

To address their backup challenges while staying within budget constraints, the company chose to implement IBM Safeguarded snapshots, a data protection solution utilizing the built-in scheduler of their existing IBM FlashSystem 5300 storage array. Safeguarded snapshots offer faster restores and immutability (unchangeable copies), making them ideal for protecting critical data from accidental deletion or ransomware attacks. The company strategically selected two of their ten data stores to host their business-critical virtual machines (VMs) supporting core operations, such as customer relationship management systems or financial databases.

They configured the internal scheduler to create daily Safeguarded snapshots of these data stores, ensuring they have readily available backups for a one-week recovery window.

Recovery

To ensure they could recover their critical VMs in the event of data loss or corruption, the company was provided with a detailed recovery procedure:

1. **Identifying the snapshot:** The administrator identifies the relevant Safeguarded snapshot to restore from, based on the desired recovery point.
2. **Creating a thin clone:** The administrator creates a thin clone of the selected snapshot, which is a space-efficient copy of the data that only consumes storage space as changes are made.
3. **Mapping the thin clone:** The thin clone is then mapped to the appropriate ESXi hosts in the VMware cluster.
4. **Re-scanning and re-signature:** The ESXi hosts are instructed to re-scan their storage adapters to detect the newly mapped thin clone volumes. Any VMFS data stores on the thin clones are re-signed to ensure compatibility.
5. **Mounting and accessing VMs:** The administrator mounts the thin clone volumes to the ESXi hosts, making the VMs and their data accessible for recovery or further analysis.

Benefits

The benefits of the solution are as follows:

Cost-effective data protection

The use of Safeguarded snapshots with the internal scheduler provided a low-cost solution for data protection, eliminating the need for additional software licenses.

Resilience and rapid recovery

The daily snapshots and seven-day retention ensured that the company could recover their critical VMs and data within a reasonable time frame, minimizing downtime and data loss.

Operational simplicity

While manual intervention was required for recovery, the process was straightforward and well-documented, allowing the company to restore their data without specialized expertise.

Outcome

By implementing Safeguarded snapshots on their IBM FlashSystem 5300, the small business achieved a cost-effective and resilient solution tailored to their specific needs. Reduced backup storage costs compared to their previous NL-SAS NAS solution was a significant benefit. The ability to create and recover from daily snapshots provided peace of mind and ensured business continuity in the face of unexpected disruptions, such as power outages or hardware failures. While the recovery process from Safeguarded snapshots required some manual intervention compared to fully automated solutions, the clear instructions and low cost made it a viable option for the budget-conscious organization.

3.2 Design considerations

While 3.1, “Design scenarios” on page 28 illustrates the adaptive use cases of IBM Safeguarded snapshot, this section delves deeper into the critical design considerations that ensure successful implementation and optimal outcomes.

Whether you are safeguarding a massive Oracle database, a business-critical SAP HANA environment, or seeking to enhance ransomware resilience, these considerations will guide you through key decision points like snapshot frequency, data retention policies, integration with existing backup solutions, and considerations for specific applications.

3.2.1 Assessing your data protection needs

When assessing your data protection needs, consider the following:

Identifying critical data

To ensure successful disaster recovery (DR), the first and most important step is to clearly define which data sets are essential for your business operations, ensuring you have the information needed to continue functioning in the event of an unexpected incident. This process should consider factors such as:

- ▶ The impact of data loss on your business continuity and reputation.
- ▶ Regulatory requirements, such as the Digital Operational Resilience Act (DORA) for financial institutions, which mandates the identification and protection of critical systems and data.
- ▶ Recovery time objectives (RTOs): The maximum acceptable time to restore data after a disruption.
- ▶ Recovery Point Objectives (RPOs): The maximum acceptable amount of data loss in the event of a disruption.

Understanding data change rates

Analyze the rate at which your data on the storage volume changes to determine the optimal frequency for taking Safeguarded snapshots. This can be estimated by reviewing the volume's write activity. The most accurate estimation of the change rate is to create a test snapshot and keep it for a period like a week to capture sudden increases in write activity, such as monthly batch jobs or database loads. Then, divide the size of the snapshot by the number of days it was kept to get the average daily change rate.

The above calculation provides a general idea of space consumption under normal conditions. It is a fairly simple calculation of $(A \div B) \times C$, where:

The above calculation provides a general idea of space consumption under normal conditions. It is a fairly straightforward calculation of $(A \div B) \times C$, where:

- ▶ A = number of FlashCopies per day
- ▶ B = average daily change rate
- ▶ C = retention in number of days

Evaluating storage requirements

Calculate the additional storage capacity needed to accommodate Safeguarded snapshots based on your data size and change rates. Consider both short-term and long-term retention needs. However, a more accurate picture of capacity planning also accounts for the impact of an actual ransomware attack by factoring it in:

Loss of compression: The most popular form of ransomware attack encrypts the data with a key for which the victim must pay to unlock their data. Since encrypted data (done at the Operating System through a nefarious application instead of at the storage level) defeats compression, a comprehensive capacity plan accounts for this increase in capacity.

Recovery volume space: During Safeguarded snapshot recovery, customers need to double their protected application capacity. This is because both the original production volume and the recovery volume (where the snapshot is restored) need to exist simultaneously, ensuring uninterrupted operations and providing a safety net for testing and verification. This doubling of capacity requirement must be considered when designing safeguarded snapshots.

3.2.2 Designing your Safeguarded snapshot strategy

Consider the following when designing your Safeguarded snapshot strategy:

Snapshot frequency

Determine how often you need to create Safeguarded snapshots to meet your RPO requirements. Hourly, daily, or weekly snapshots may be appropriate depending on the criticality of the data and its rate of change.

Retention period

Decide how long you need to retain each snapshot. This will depend on your regulatory requirements, data recovery needs, and available storage capacity.

Snapshot validation

Consider implementing a validation process to ensure the integrity and consistency of each snapshot. This can include running database checks or using tools such as IBM Storage Sentinel to detect anomalies.

3.2.3 Integrating with security and management tools

You can integrate IBM Safeguarded snapshot with the following with security and management tools:

IBM Storage Copy Data Management (SCDM)

SCDM provides a centralized platform for orchestrating and automating the creation, management, and mounting of Safeguarded snapshots.

Copy Services Manager (CSM)

CSM acts as a central interface for managing and automating Safeguarded snapshot operations. You can use CSM to:

- ▶ Define Safeguarded snapshots.
- ▶ Schedule automated backups.
- ▶ Monitor and manage Safeguarded snapshots.
- ▶ Recover data.

IBM Storage Sentinel

This AI-powered tool can monitor your Safeguarded snapshots for anomalies, such as data corruption or ransomware encryption, providing early warning of potential threats.

IBM Security QRadar SIEM

Integrate Safeguarded snapshot with your Security Information and Event Management (SIEM) system to enable automated responses to security events, such as triggering an unscheduled snapshot when suspicious activity is detected.

3.2.4 Testing and validation

The following considerations apply for testing and validating your Safeguarded snapshot solution:

Regular testing

Conduct regular tests of your Safeguarded snapshot solution to ensure it performs as expected in the event of a data loss or cyberattack. Test the recovery process, validate the integrity of recovered data, and measure recovery times.

Monitoring and optimization

Monitor your snapshot environment for performance and capacity utilization. Adjust snapshot frequencies or retention periods as needed to optimize resource usage.

By carefully considering these design aspects, you can tailor a Safeguarded snapshot solution that meets your organization's specific data protection requirements, ensuring the resilience, security, and availability of your critical data assets.



Implementing Safeguarded snapshot

The Storage Virtualize Safeguarded snapshot functionality offers a robust solution for protecting your critical data from various threats and ensuring compliance with regulations.

This chapter covers the Safeguarded snapshot feature and has the following sections:

- ▶ “Implementing a Safeguarded snapshot environment” on page 38
- ▶ “Configuring the Safeguarded snapshot capacity” on page 39
- ▶ “Defining volume groups and Safeguarded policies” on page 42
- ▶ “Creating a Safeguarded snapshot policy” on page 47
- ▶ “Managing snapshots” on page 52

4.1 Implementing a Safeguarded snapshot environment

Before you configure a Safeguarded snapshot environment, complete the planning phase, which includes the following tasks:

- ▶ Sizing the Safeguarded snapshot Backup capacity.
- ▶ Verifying the prerequisites.
- ▶ Deciding which topology to implement.
- ▶ Defining the backup frequency.
- ▶ Specifying the retention period.

For more information about planning Safeguarded snapshot environment, see Chapter 2, “Safeguarded snapshot: Redefining data protection” on page 13.

Figure 4-1 shows the configuration that we use as an example to set up a Safeguarded snapshot environment.

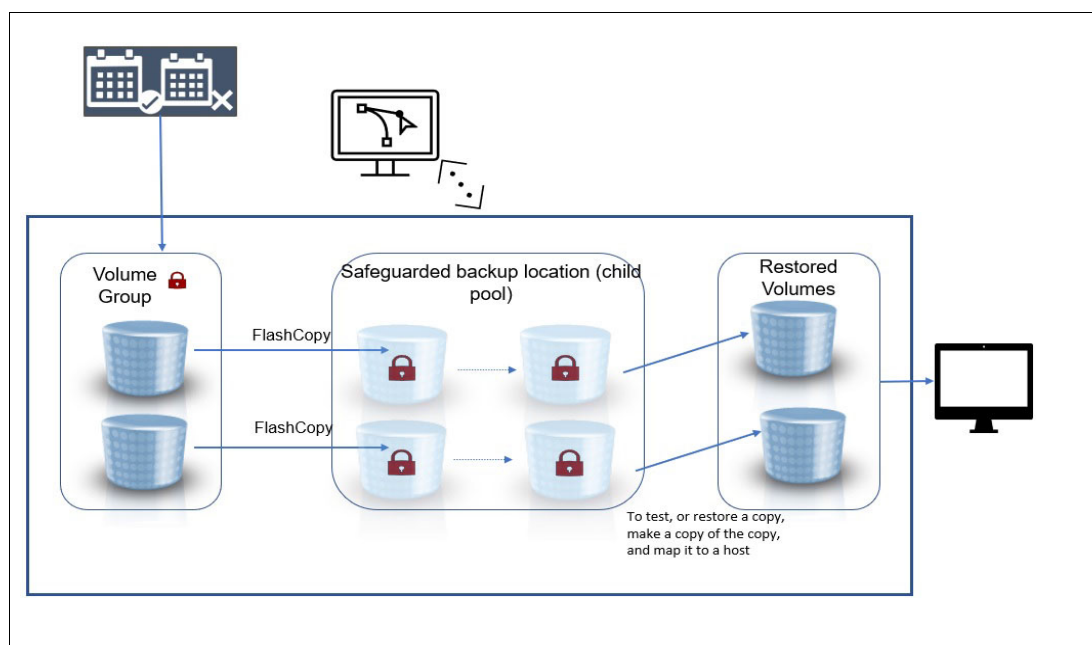


Figure 4-1 Safeguarded snapshot environment

The Safeguarded snapshot environment requires a set of Safeguarded snapshot source volumes and an equal number of recovery volumes. The recovery volumes are necessary to recover data from a Safeguarded snapshot backup. For more information about this requirement, see Chapter 2, “Safeguarded snapshot: Redefining data protection” on page 13.

Note: The following terms are used in the examples in this document:

- ▶ *Source volume* refers to the volume where the Safeguarded snapshot relationship is defined.
- ▶ *Production volume* refers to the volume that is active to the host.

4.2 Configuring the Safeguarded snapshot capacity

Before IBM Storage Virtualize 8.6, Safeguarded snapshots could be stored in specific locations within the storage pool. However, since 8.6, they leverage the functionality of volume group snapshots. As a result, Safeguarded snapshots inherit the same storage pool requirements as regular volume group snapshots - they can be placed in either a parent or child pool within the storage system.

The properties and restriction of Safeguarded pools are as follows:

- ▶ By default Safeguarded snapshots are in the same pool as the parent volume.
- ▶ Safeguarded snapshots can be added to a child pool of the parent volume's pool.
- ▶ If the parent volume is in a child pool itself, Safeguarded snapshots can be added to a child pool that has the same parent pool as the parent volumes pool.
- ▶ A Safeguarded snapshot location cannot be deleted if:
 - The Safeguarded snapshot location contains all Safeguarded snapshots (except by the Security Admin).
 - The Safeguarded snapshot location is associated to a Safeguarded source. This restriction also applies to the Security Admin.
- ▶ Safeguarded snapshot quotas can be increased, but reduction might require Security Administrator privileges.

Note: See this [video](#) on how to size for Safeguarded snapshot.

4.2.1 Configuring a Safeguarded snapshot location (optional)

To configure a Safeguarded snapshot location, you must create a dedicated safeguarded child pool within the parent pool where your source volumes reside.

Although Safeguarded snapshots no longer require child pools like they did with legacy FlashCopy mappings, creating a separate child pool can still isolate snapshots, protect them from parent pool issues, and optimize production workload performance.

A Safeguarded snapshot backup location can be created by using both the CLI and the GUI. Here are the steps for the GUI method:

1. In the management GUI, select **Pools** → **Pools**. Right-click a parent-pool and select **Create Child Pool**. See Figure 4-2 on page 40.

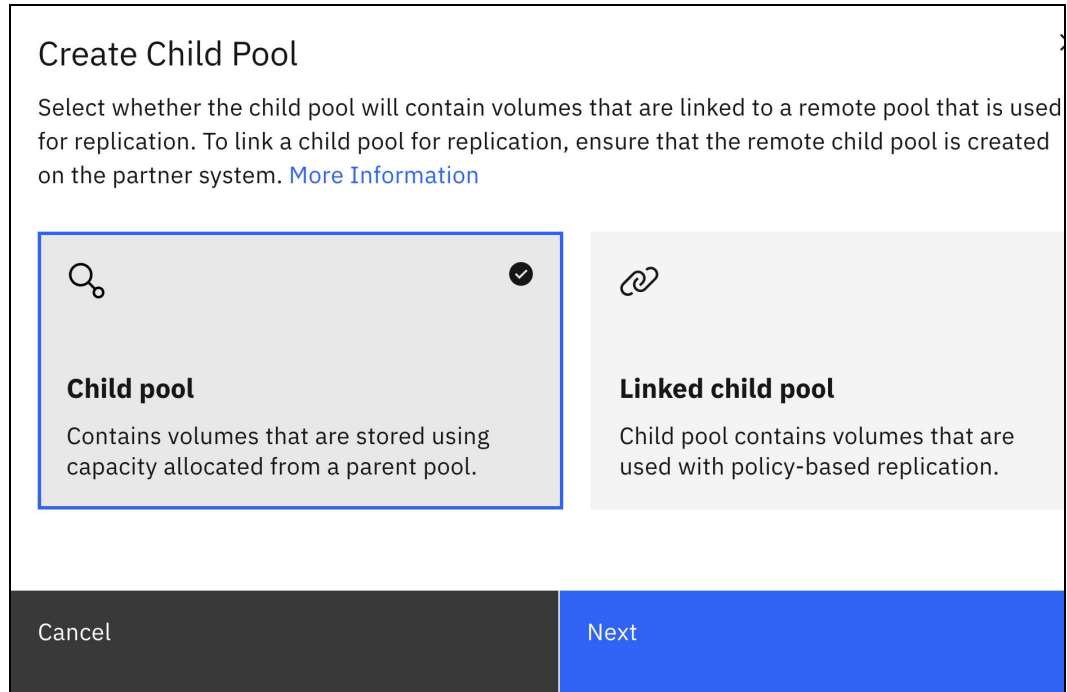


Figure 4-2 Create Child Pool

2. On the Create Child Pool page, select the **Child pool** option since the capacity will be allocated from the parent pool. Click **Next**. See Figure 4-3 on page 41.

Create Child Pool ✕

Create a child pool or a Safeguarded child pool for a parent storage pool. Use safeguarded child pools to store backups for a volume that is protected by the Safeguarded Copy function.

Parent Pool

Pool0 ▼

Available pool capacity : 99.01 TiB

Child pool name and capacity

Child Pool (optional)

Pool0_SG_0

Child Pool Capacity

10

-

|

+

TiB

▼

Extent-Rounded Capacity: 10.00 TiB ⓘ

Safeguard ⓘ

Encryption ⓘ

Provisioning Policy ⓘ

Provisioning Policy (optional)

Select... ▼

Cancel

Create

Figure 4-3 Create Child Pool details

3. Enter a name for the child pool (for example, Pool0_SG_0). Specify the amount of storage space that you want to allocate to this child pool. This space is dedicated to storing Safeguarded snapshots. Enable the **Safeguard** check-box to indicate that the child pool is used as the Safeguarded snapshot location. Click **Create**.
4. Child pools that are Safeguarded are marked with a shield icon on the Pools page. See Figure 4-4 on page 42.

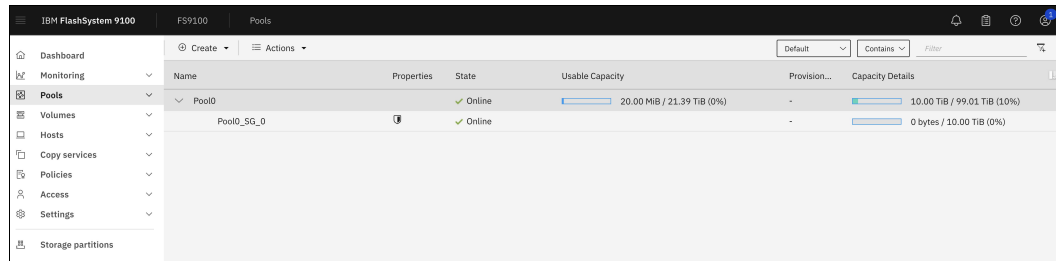


Figure 4-4 Safeguarded child pool

You can create a safeguarded child pool using the command-line interface (CLI). Here is the command for creating a safeguarded pool through the CLI:

```
mkdiskgrp -parentdiskgrp SVPC Pool1 -size 100 -unit gb -safeguarded
```

Explanation of parameters:

- ▶ **-safeguarded**: This parameter marks the child pool as safeguarded.
- ▶ **-parentdiskgrp <parent_pool_name>**: This parameter specifies the name of the parent pool where the safeguarded child pool resides.
- ▶ **<child_pool_name>**: Replace this with the desired name for your safeguarded child pool.
- ▶ **-size <size_in_GB>**: Indicate the size that you want to allocate for the child pool in gigabytes (GB).

4.3 Defining volume groups and Safeguarded policies

Volume groups are the way Safeguarded snapshot manages a group of related volumes. A *volume group* is a set of related volumes that can be managed and configured collectively.

Not all volumes in a Safeguarded volume group must belong to the same parent pool. However, all volumes in the Safeguarded volume group must have a Safeguarded snapshot location. This condition must also be met when a volume is added to a Safeguarded volume group. Otherwise, adding the volume fails.

Additional restrictions to volumes can be included in a Safeguarded volume group. If these restrictions are not met, a volume group might be prevented from becoming Safeguarded or a volume might be prevented from being added to a Safeguarded volume group.

A volume can be designated as *Safeguarded source volume* only if its volume group is associated with a Safeguarded policy. A Safeguarded source is automatically associated with a Safeguarded snapshot location. For a mirrored Safeguarded source, each volume copy is associated with a Safeguarded snapshot location.

4.3.1 Volume group creation

Before you create a volume group, determine of which source volumes you want to create Safeguarded snapshots. A volume group becomes *Safeguarded* when it is associated with a Safeguarded policy.

A volume group can be considered Safeguarded, but not have any volumes in it or any Safeguarded snapshots that are created yet.

Creating a volume group using the GUI

To create a volume group using the GUI, complete the following steps:

1. In the management GUI, select **Volumes** → **Volumes Groups**.
2. Click **Create Volume Group**. See Figure 4-5.




Figure 4-5 Create Volume Group

3. The first time a volume group is defined the Create Volume Group page enables the entry a name for the volume group and whether existing volumes will be selected and click **Next** as shown in Figure 4-6.

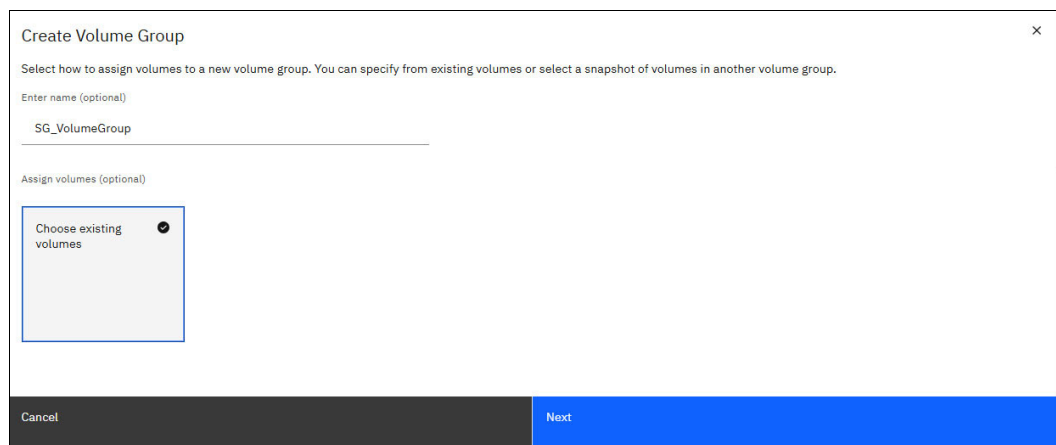


Figure 4-6 Create Volume Group - Choose existing volumes

4. From the list of volumes, select one of the volumes that you want in the volume group and click **Create Volume Group**. See Figure 4-7 on page 44.

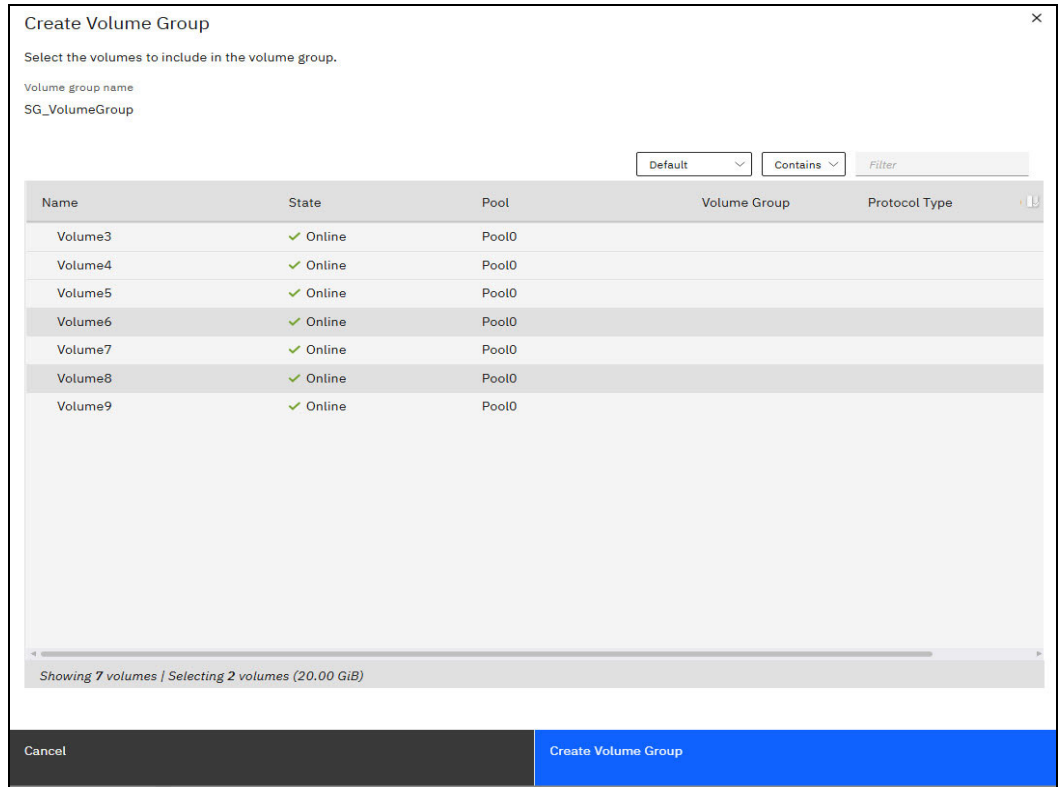


Figure 4-7 Create Volume Group - Volume selection

A task to handle the creation of the volume group is initiated. The progress of the task is displayed dynamically. See Figure 4-8 on page 45.

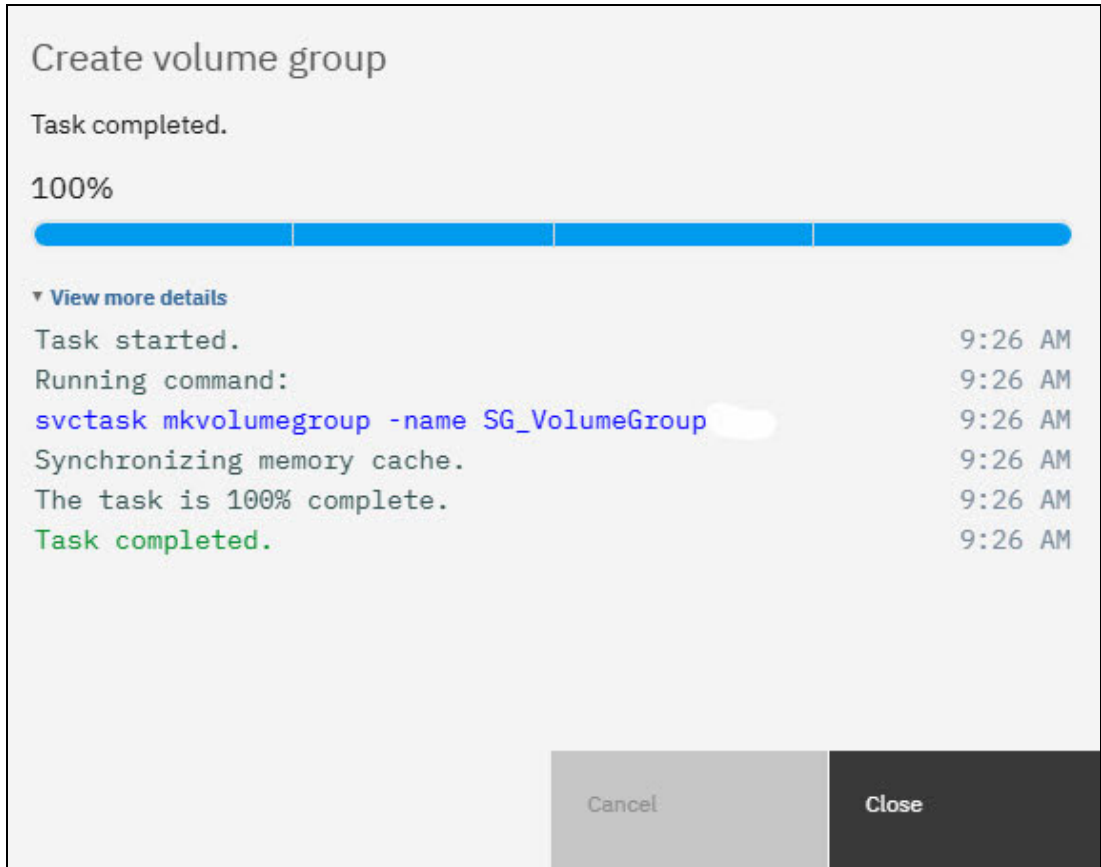


Figure 4-8 Create a volume group

When the task is closed control returns to the main Volume Groups view. See Figure 4-9.

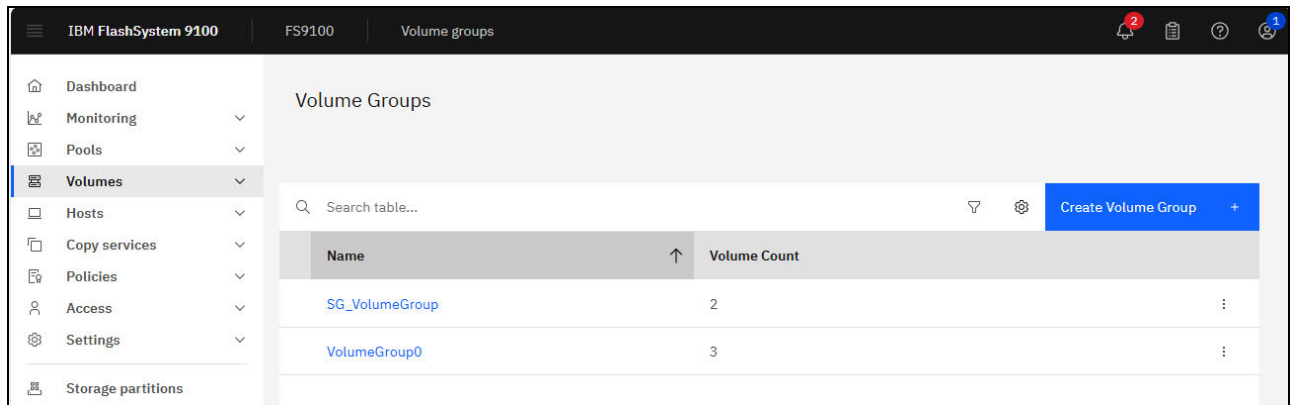


Figure 4-9 Volume Groups

- After the volume group is created numerous actions can be taken within the context of the group. In this example, two source volumes were added to the volume group, **Volume6** and **Volume8**. Clicking the volume group name displays the attributes of the group (Figure 4-10 on page 46).

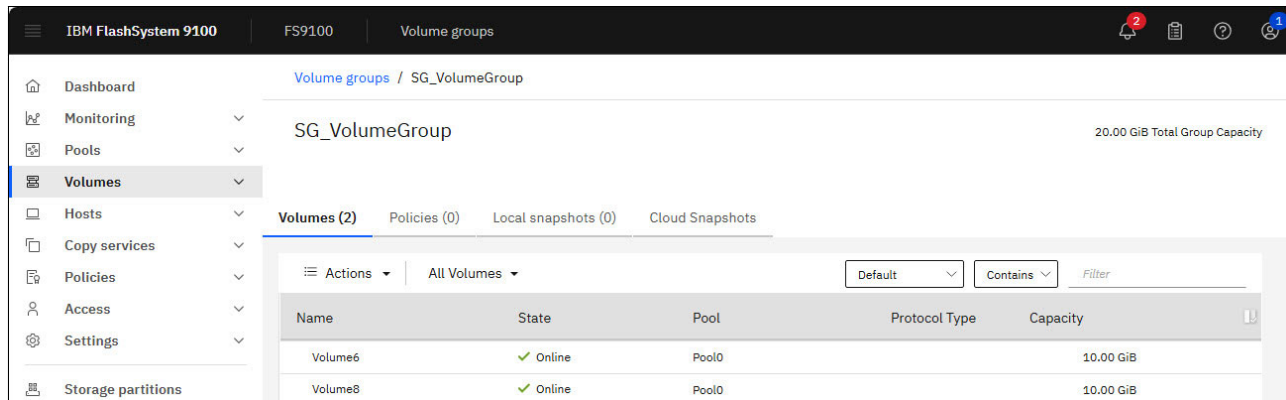


Figure 4-10 Volume Group attributes

It is important to understand that the group created is not Safeguarded at this point. The assignment of a Safeguarded policy is required.

Creating a volume group using the CLI

To create and configure a new volume group and assign volumes to that group, complete these steps:

1. To create the new volume group, enter the following command:

```
mkvolumegroup -name volumegroup_name
```

where:

- volumegroup_name specifies a volume group name

If a Safeguarded policy has been defined already the volume group can be designed safeguard by modifying the command syntax:

```
mkvolumegroup -name volumegroup_name -safeguardedpolicy safeguarded_policy_id | safeguarded_policy_name
```

where either the ID or name of the safeguard policy is specified.

2. If you are assigning a new volume to the volume group, enter the following command:

```
mkvolume -pool <pool_name_or_id> -volumegroup <volumegroup_name_or_id> -size <disk_size>
```

where:

- <pool_name_or_id> is the name or identifier of the parent pool
- <volumegroup_name_or_id> is the name or identifier of the volume group
- <disk_size> indicates the capacity that is provisioned for the volume from the parent pool

3. If you are assigning existing volumes to the volume group, enter the following command:

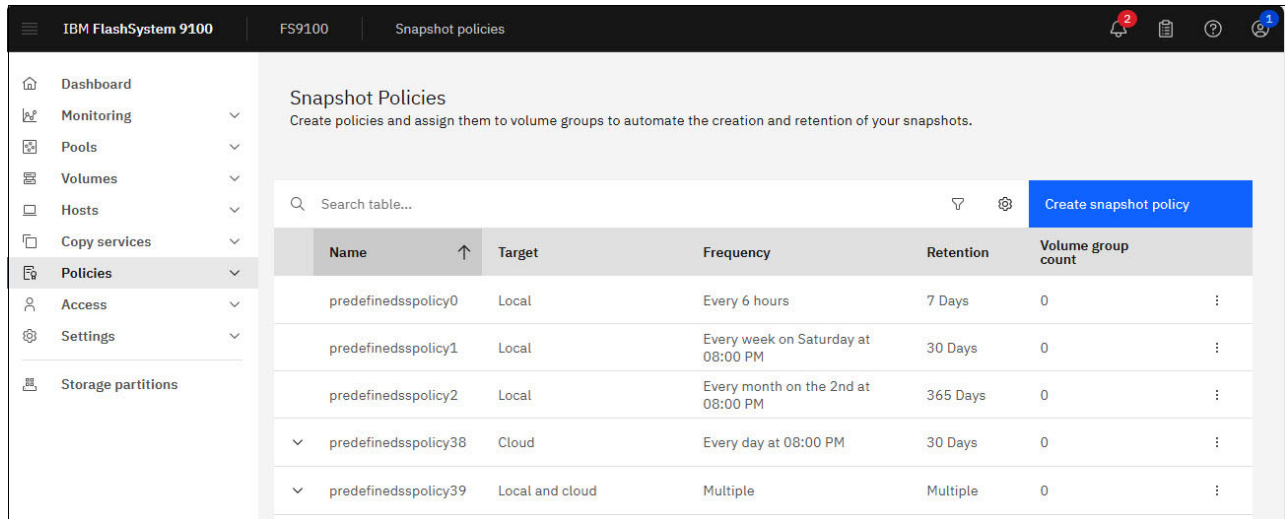
```
chvdisk -volumegroup <volumegroup_name_or_id> <name_or_id> ,
```

where:

- <volumegroup_name_or_id> is the name or identifier of the volume group
- <name_or_id> of the volume

4.4 Creating a Safeguarded snapshot policy

A Safeguarded snapshot policy controls the creation, retention, and expiration of Safeguarded snapshot copies of source volumes. The management GUI supports the display of predefined and user-defined Safeguarded snapshot policies. The policies are accessed through the navigation bar options **Policies** → **Snapshot policies** (Figure 4-11).



Name	Target	Frequency	Retention	Volume group count
predefinedsspolicy0	Local	Every 6 hours	7 Days	0
predefinedsspolicy1	Local	Every week on Saturday at 08:00 PM	30 Days	0
predefinedsspolicy2	Local	Every month on the 2nd at 08:00 PM	365 Days	0
predefinedsspolicy38	Cloud	Every day at 08:00 PM	30 Days	0
predefinedsspolicy39	Local and cloud	Multiple	Multiple	0

Figure 4-11 Snapshot policies view

By default, five policies for local snapshots are defined:

- predefinedsspolicy0

Select this policy for the most frequent copies and retention. When this policy is assigned to a volume group snapshots for the volume within the group are created at six hour intervals, and each snapshot is retained for seven days.

- predefinedsspolicy1

Select this policy for less frequent copies and medium retention. This policy creates snapshots once a week at 8 PM and retains them for 30 days.

- predefinedsspolicy2

Select this policy for less frequent copies and longer retention. This policy creates snapshots once a month on the second day of the month at 8 PM and retains each snapshot for 365 days.

- Predefinedsspolicy38 (new in 8.7)

Select this policy to create cloud snapshots only. For this policy, cloud snapshots are created daily and retained for a month. Use this policy for storing snapshots of the volumes to cloud storage.

- Predefinedsspolicy39 (new in 8.7)

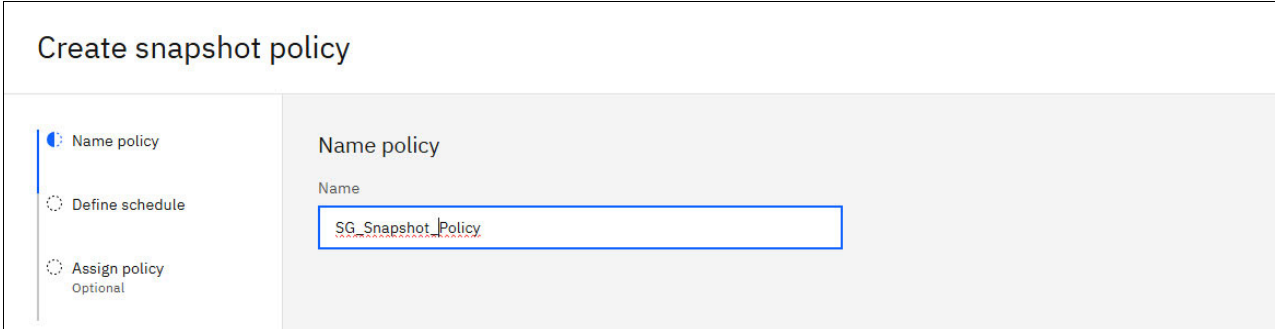
Select this hybrid policy to create schedules for local snapshots and cloud snapshots. Select this policy for creating both local snapshots and cloud snapshots. For this policy, local snapshots are created every six hours and retained for a week. It also creates cloud snapshots daily and are retained for a month. Use this policy for volume data that requires the highest recovery point objective (RPO) and for storing snapshots of the volumes to cloud storage.

Important: It is important to remember that a snapshot policy is a set of rules that controls the creation and expiration of snapshots. Safeguarded status is assigned during snapshot creation, independent of any defined policies. Users can create Safeguarded snapshots using the `addsnapshot` command or the GUI, even if a policy is not defined.

4.4.1 Creating a Safeguarded snapshot policy

New snapshot policies can be defined to suit the needs of your requirements.

1. From the main snapshot policies view (Figure 4-11 on page 47), click **Create snapshot policy** (Figure 4-12)



The screenshot displays the 'Create snapshot policy' web interface. On the left, a vertical sidebar contains three steps: 'Name policy' (indicated by a blue circle), 'Define schedule' (indicated by a grey circle), and 'Assign policy' (Optional, indicated by a grey circle). The main content area is titled 'Name policy' and features a 'Name' label above a text input field. The input field contains the text 'SG_Snapshot_Policy' and is highlighted with a blue border.

Figure 4-12 Create a snapshot policy

2. Enter the name for the policy (SG_Snapshot_Policy) and click **Next**. The frequency at which the snapshot will be created must be specified. They can be created:
 - ▶ Hourly
 - ▶ Daily
 - ▶ Weekly
 - ▶ Monthly

Once the generation period is selected the UI context changes prompting for specific attributes of the schedule. How the UI context changes depends on the period selected. For this example snapshots will be generated every four hours, initially starting at 1:00 AM and be retained for three days (Figure 4-13 on page 49).

Create snapshot policy

- ⊙ Name policy
- ⊙ Define schedule**
- Assign policy
Optional

Define schedule

Specify the frequency, time, day of the week, day of the month and retention period for snapshots.

Frequency

Hourly interval

Retention interval (days)
 - | +

Base time
 AM

The snapshot is taken at the specified hourly interval from the base time.

Summary

Name SG_Snapshot_Policy

Snapshots schedule
 Every 4 hours, initially starting at 01:00 AM, and retained for 3 days

Figure 4-13 Create a snapshot policy schedule

3. Clicking **Create Policy** initiates a task to create the new policy. Feedback on task progress is displayed (Figure 4-14).

Create Snapshot Policy

Task completed.

100%

▼ View more details

```

task started.
Running command:
svctask mksnapshotpolicy -backupinterval 4 -
backupstarttime 2406220500 -backupunit hour -name
SG_Snapshot_Policy -retentiondays 3
The snapshot policy (ID 3) was successfully created.
Synchronizing memory cache.
The task is 100% complete.
Task completed.
  
```

Cancel

Close

Figure 4-14 Create snapshot policy task

- After the **Close** button is clicked the policy can be assigned to one or more volume groups through the assignment that the safeguarded attributes are added (Figure 4-15). In addition to the volume group selection, enabling the Safeguarded checkbox before clicking **Assign policy** is critical.

Create snapshot policy

- Name policy
- Define schedule
- Assign policy** (Optional)

Assign policy (optional)
Assign this snapshot policy to one or more volume groups.

Policy name
SG_Snapshot_Policy

Local schedule
Every 4 hours, retained for 3 days

Select the volume groups to associate with this policy

Search

Volume group	Volume count	Snapshot policy
<input checked="" type="checkbox"/> SG_VolumeGroup	2	No
<input type="checkbox"/> VolumeGroup0	3	No

2 items

You can specify an alternate start date/time below.

Choose start date (optional): mm/dd/yyyy

Choose a time (optional): hh:mm AM

Safeguarded

Total volumes assigned (max: 454): 2

Buttons: Cancel, Skip, Back, **Assign policy**

Figure 4-15 Safeguarded policy assignment

- The **Assign Snapshot policy to volume groups** task (Figure 4-16 on page 51) provides feedback on the completion of the assignment.

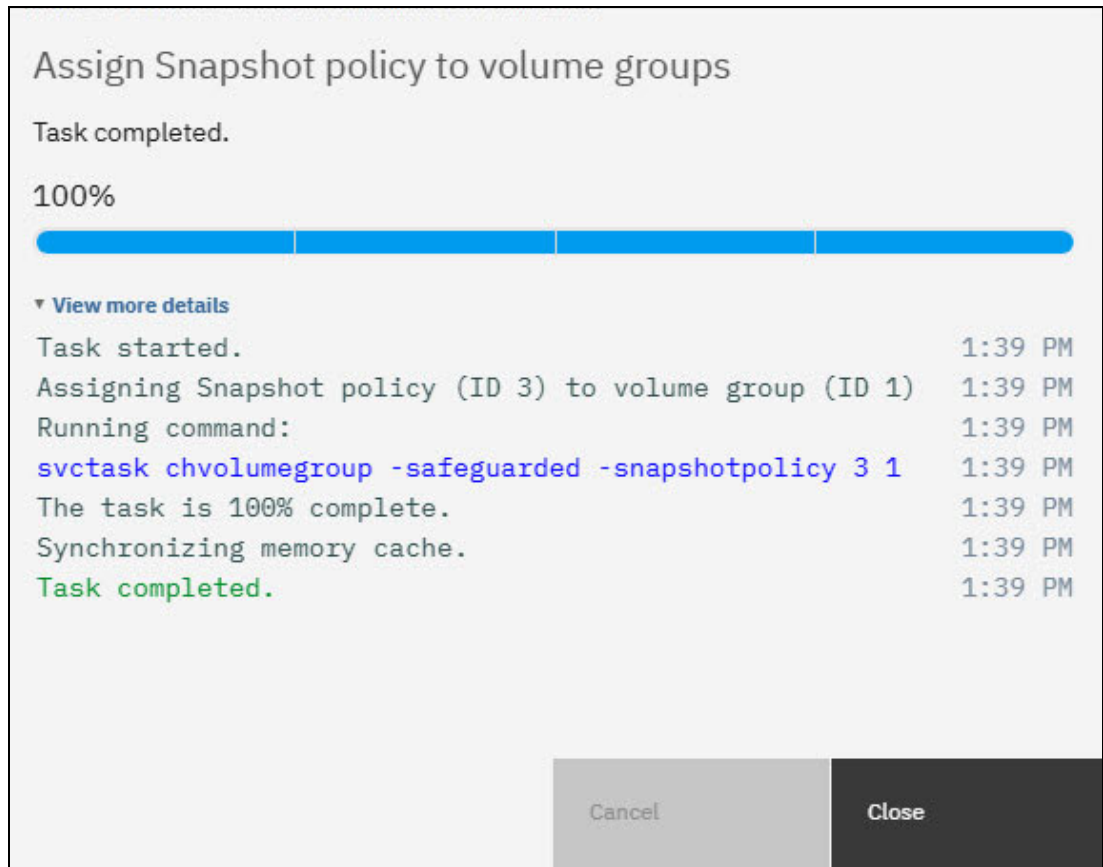


Figure 4-16 Safeguarded snapshot policy assignment task

If a policy intended for safeguarding volume groups is created and assigned to volume groups without enabling the safeguarded attribute the policy needs to be unassigned to the volume groups, a new policy created and assigned.

- Unassigning a policy is done using the **Volumes** → **Volume Groups** view. For each volume group, select the **Policies** tab (Figure 4-17 on page 52). A policy can only be deleted when it is assigned to no volume groups.

Note: Only a security administrator or higher can unassign a Safeguarded snapshot policy.

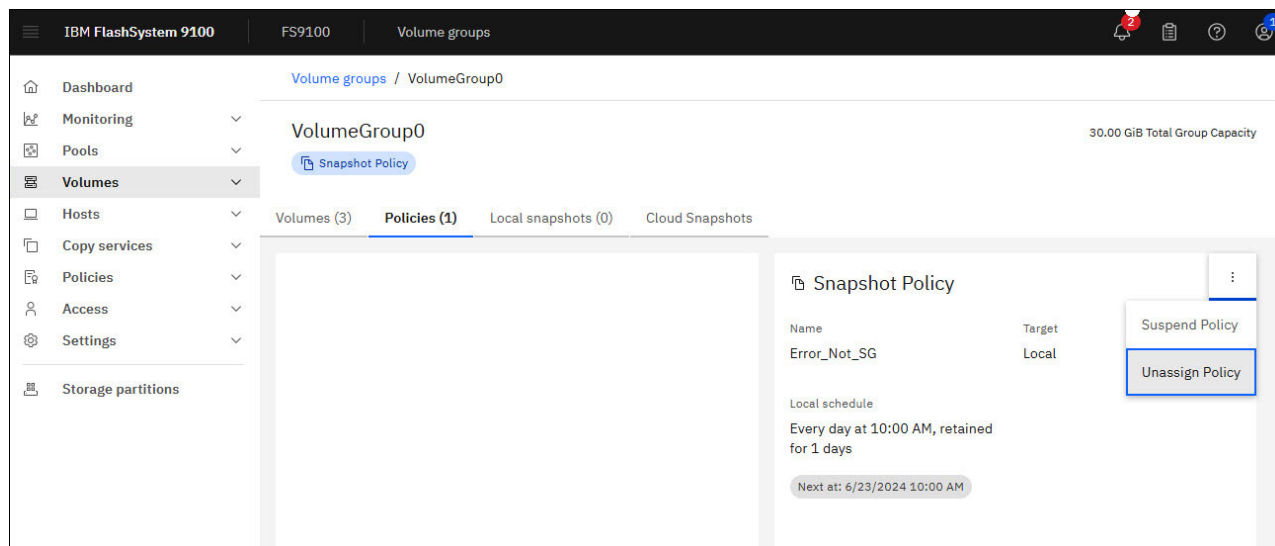


Figure 4-17 Unassign policy

IBM Storage Virtualize automates snapshot creation once a Safeguarded policy is assigned. This eliminates the need for manual intervention, requiring only standard monitoring procedures.

Tip: Snapshot policy is only needed if automatic creation of snapshots is needed (the internal scheduler). Customers using custom scripts or external schedulers do not need to define a snapshot policy.

4.5 Managing snapshots

At some point after snapshot creation has been initiated an administrator needs to manage the available snapshots. The actions include:

- ▶ View snapshots for a volume group.
- ▶ Add a snapshot to a volume group.
- ▶ Restore a volume group from a snapshot.
- ▶ Delete a snapshot.
- ▶ Convert a thin-clone to a clone.
- ▶ Refresh a thin-clone from a snapshot.

Each of these actions will be described in the following sections.

4.5.1 View a volume group snapshots

Viewing the available snapshots in the management UI can be done the **Volumes** → **Volume Groups** view (Figure 4-18 on page 53). This view will display the following information about all volume groups:

- ▶ Volume group name.
- ▶ Number of volumes within the volume group.
- ▶ Name of a snapshot policy, if one has been assigned.
- ▶ Number of snapshots that have been generated.
- ▶ Whether the volume group is safeguarded.

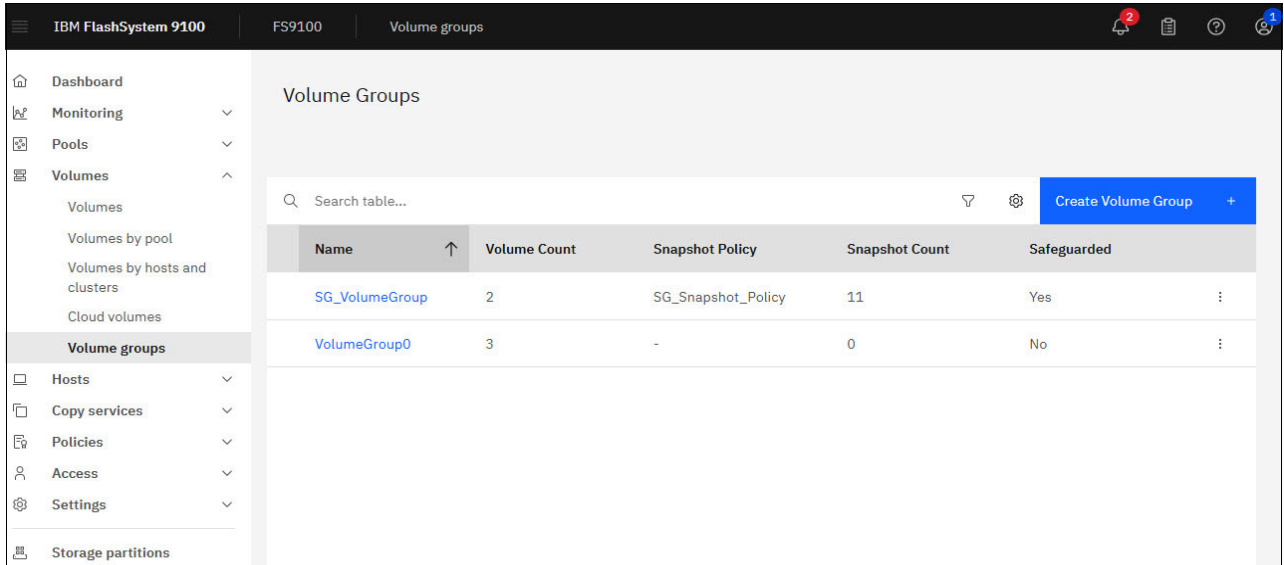


Figure 4-18 Volume Groups View

Clicking the name of a volume group, in this case SG_VolumeGroup, will display detailed information (Figure 4-19). All safeguarded volume groups are readily identified by the blue **Safeguarded snapshot Policy** icon below the group name.

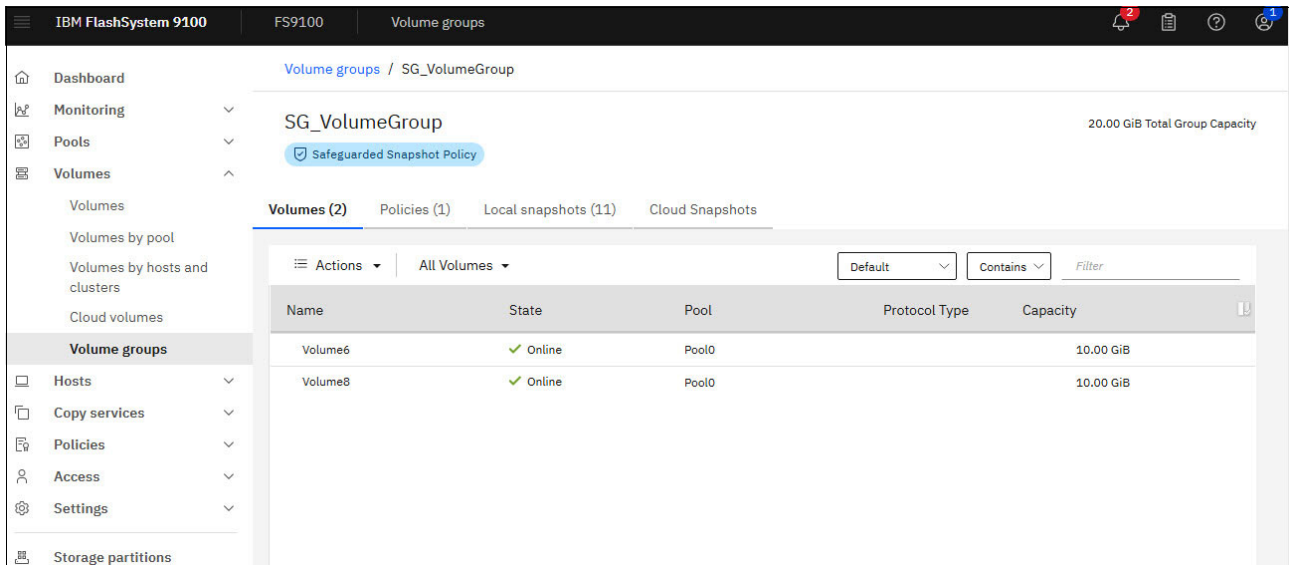


Figure 4-19 Volume group details

Information for a volume group is subdivided into four views:

- ▶ Volumes
- ▶ Policies
- ▶ Local snapshots
- ▶ Cloud snapshots

The view being displayed is denoted by a blue line under the view name. The available snapshots are listed when the **Local snapshots** view has focus (Figure 4-20 on page 54).

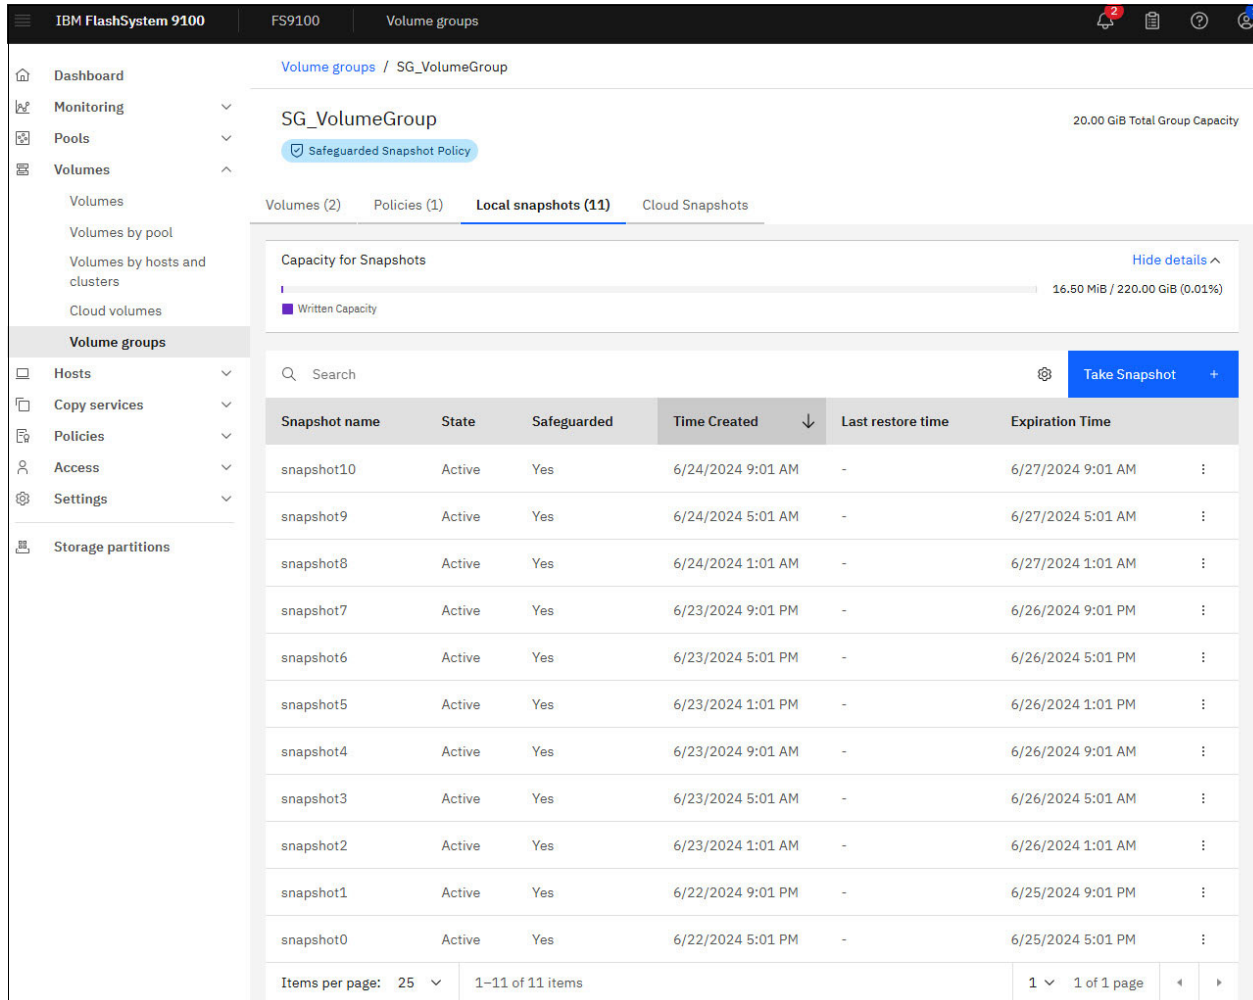


Figure 4-20 Volume group snapshots

Within this context one can take a snapshot or by clicking the ellipses icon at the end of a snapshot listing:

- ▶ Create a thin clone
- ▶ Create a clone
- ▶ Restore a snapshot
- ▶ Delete a snapshot

These actions are described in the following sections.

Viewing snapshot through the CLI can be done using the `lsvolumegroupsnapshot` command. To view Safeguarded snapshots the `-filter value safeguarded=yes` parameter can be included in the command.

4.5.2 Adding a snapshot to a volume group

The proceeding sections in this chapter have been focused on automatic snapshot generation using the internal scheduler. Periodically, depending on the requirements of the solution environment, manual snapshot generation may be required.

In the management UI one can add a snapshot to a given volume group by clicking the **Take Snapshot** button (Figure 4-20 on page 54). A task to create the snapshot will be launched and a pop-up to communicate the progress of the task will be displayed (Figure 4-21).

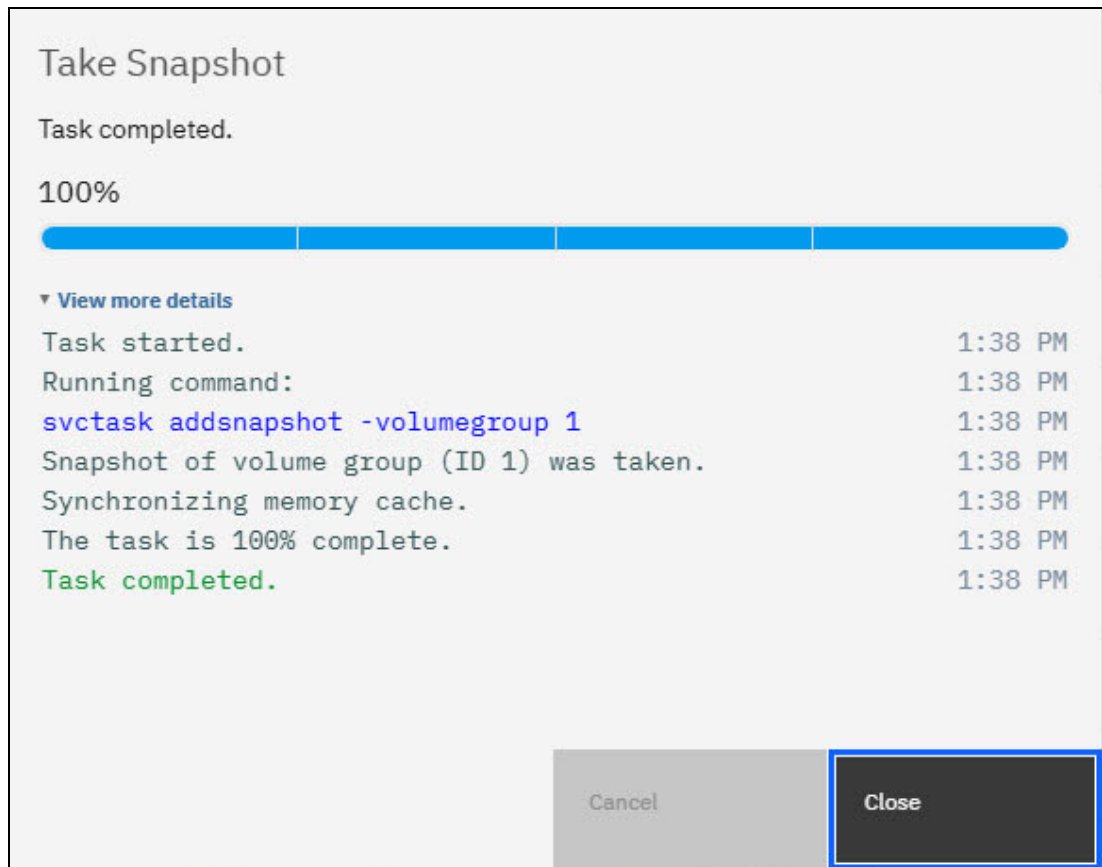


Figure 4-21 Add snapshot

It is important to note that the type of snapshot that is created is not safeguarded. This fact is reflected in the volume group's local snapshot view. There is no expiration time and the Safeguarded attribute is set to **No**.

The `addsnapshot` command has a diverse set of parameters. As an example, let us assume a manually generated Safeguarded snapshot of `SG_VolumeGroup` is needed at the current time and it is needed for a longer period of time, say 5 days. Invoking the command `addsnapshot -volume group 1 -safeguarded -retentiondays 5` will capture the volumes within the volume group at the current point in time. Shortly thereafter the new snapshot is displayed within the local snapshot view (Figure 4-22 on page 56).

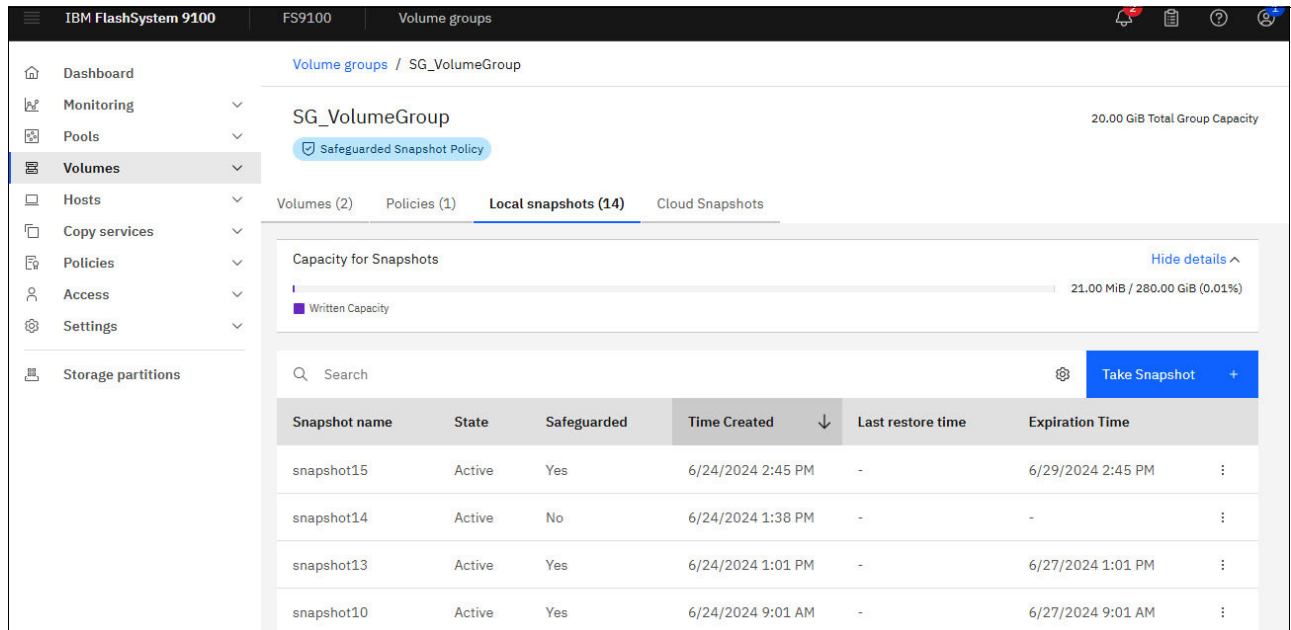


Figure 4-22 Safeguarded snapshot added

For a Safeguarded snapshot added to a volume group the Time Created and Expiration Time attributes reflect the time of command invocation.

4.5.3 Restoring a volume group from a snapshot

Moving onto the next snapshot management topic, the content of a volume group may become invalid and require restoration from a prior snapshot. In a best case scenario one can image a testing environment that leverages Safeguarded snapshots and periodically restores data based on test needs. In a worst case a ransomware event may necessitate the restoration of one or more volume groups or parts of groups.

Whether through the IBM Storage Virtualize management UI or CLI a user with sufficient authority can quickly restore the necessary volumes. The first step is to identify the snapshot that will be the restoration source. All snapshots can be accessed in a volume group's local snapshot view (Figure 4-23 on page 57).

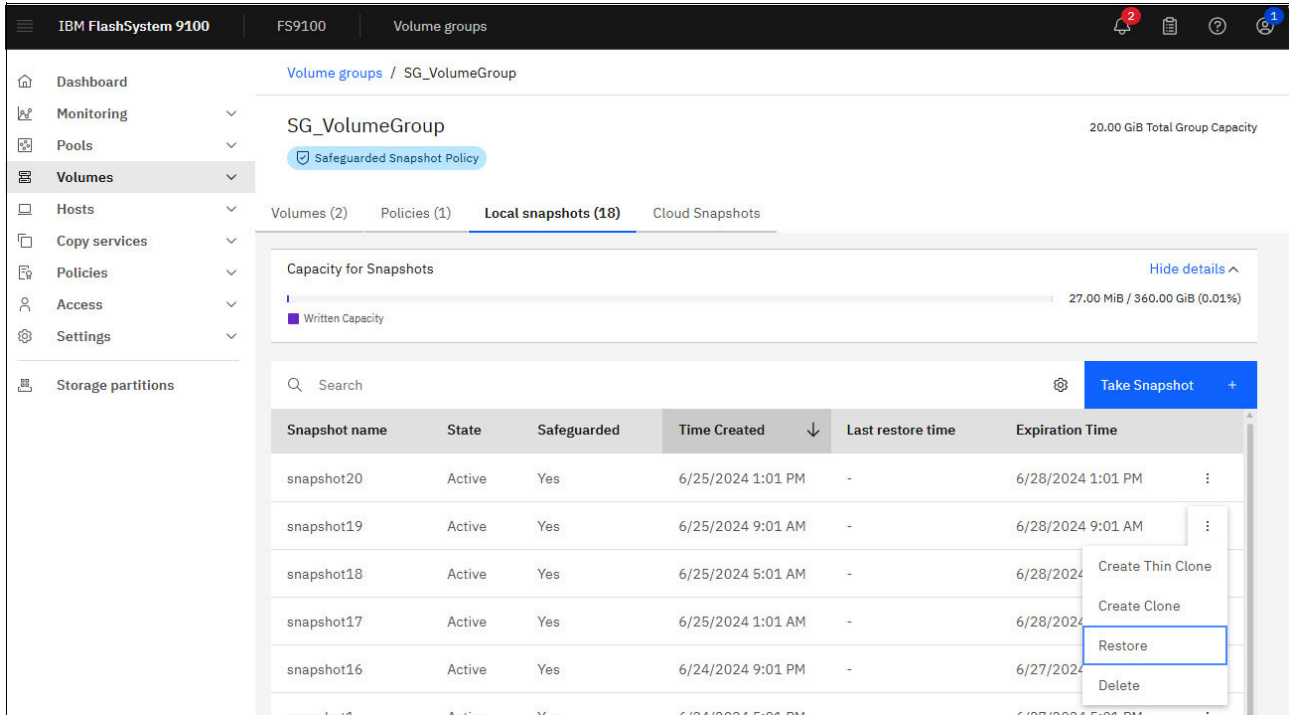


Figure 4-23 Restore snapshot

Selecting the **Restore** option will initiate a task to restore production volumes using the snapshot's content. As a precaution the dialog displayed will include the snapshot name and time of its generation (Figure 4-24). It is the user's responsibility to ensure the required snapshot is chosen. Once a restore action is complete it cannot be undone, the production data will be overwritten.

Though volume groups are treated as one logical unit it does not need to be restored in its entirety.

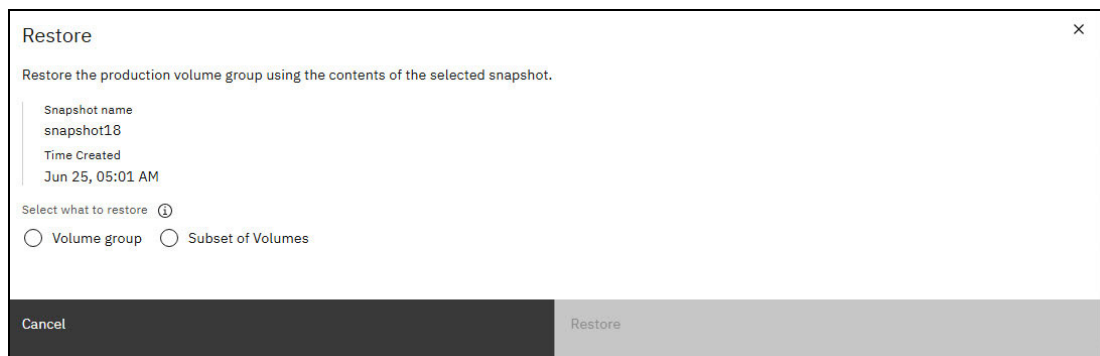


Figure 4-24 Restore snapshot selection

The context of the dialog will change based on whether an entire volume group (Figure 4-25 on page 58) or a subset of the volumes within a group (Figure 4-25 on page 58) need to be restored.

The next step when restoring an entire snapshot is to confirm the name of the volume group. This is a precautionary action. The dialog emphasizes that data will be overwritten, restates the source snapshot and that the action cannot be reversed.

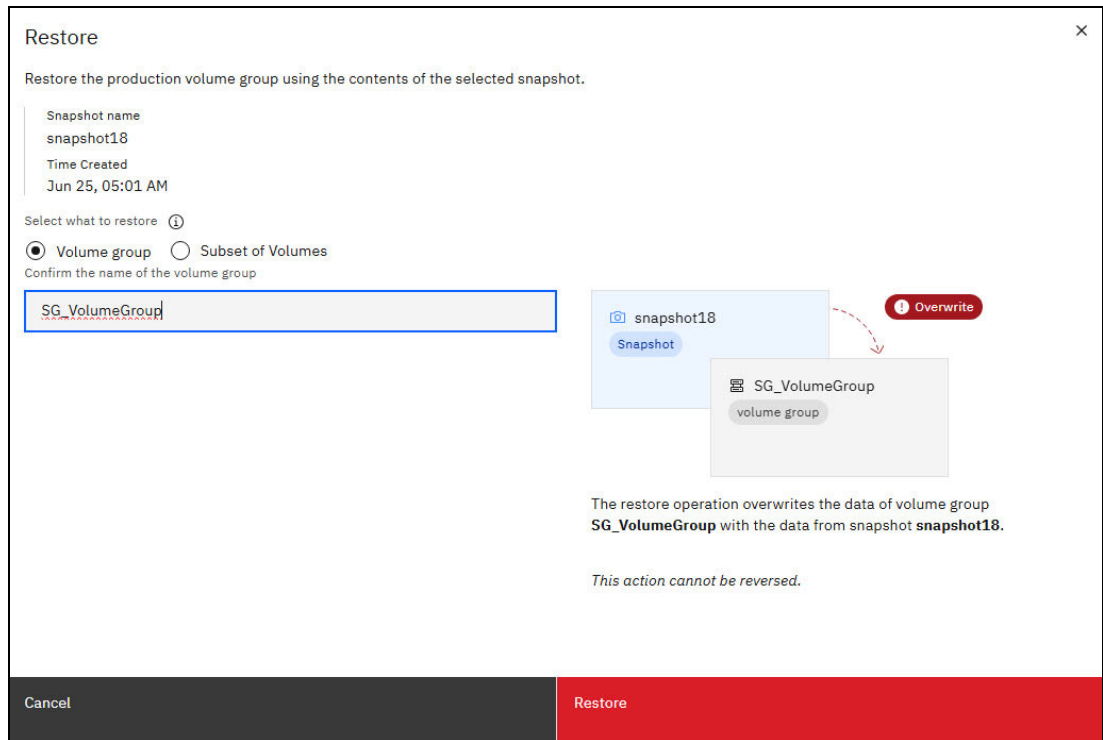


Figure 4-25 Restore entire snapshot to volume group

When restoring a subset of volume the content of the snapshot is displayed. The volumes to restore are selected using the check-box. See Figure 4-26 on page 59.

Attention: It is recommended to exercise caution when considering the restoration of a subset of volumes. Due to the potential for unknown interdependencies between volumes within a volume group, restoring only a select portion could compromise data integrity.

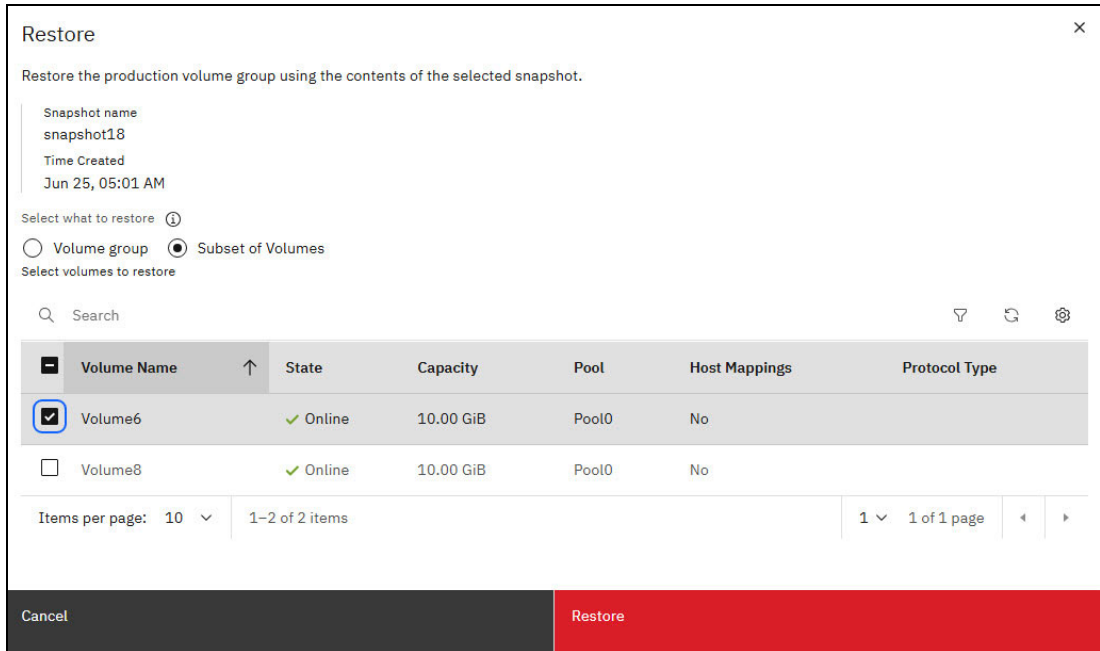


Figure 4-26 Restore subset of snapshot volumes to volume group

After the necessary volumes are selected confirm of the volume group name is required. The source snapshot will be displayed, the number of volumes being restored and the fact that the action cannot be reversed. See Figure 4-27.

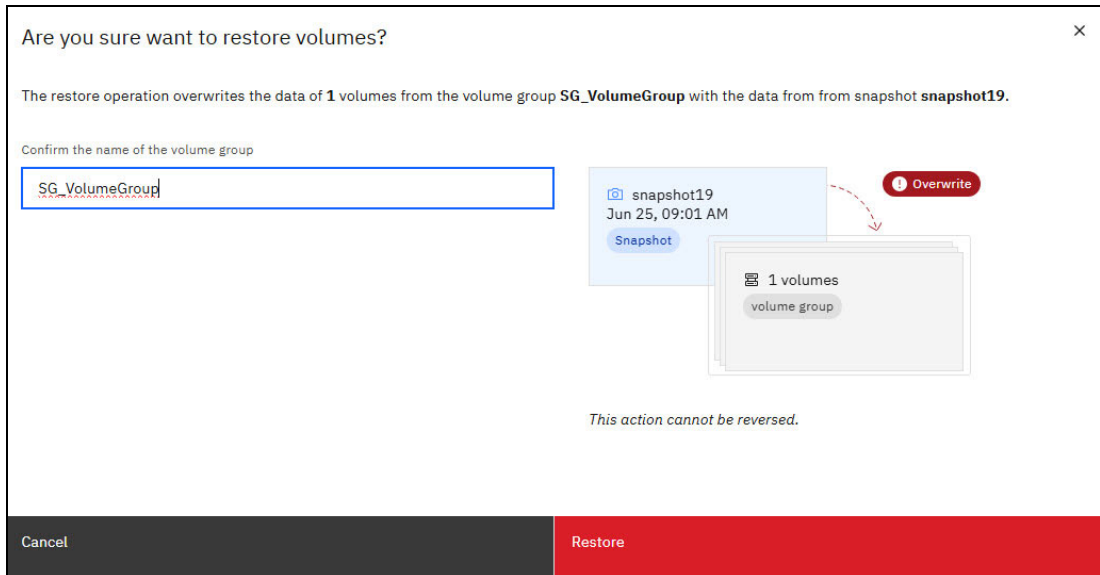


Figure 4-27 Restore subset of snapshot volumes

Regardless of where an entire snapshot or a subset of it is restored the restoration time is noted in the **Last restore time** column within the volume groups snapshot view (Figure 4-28 on page 60).

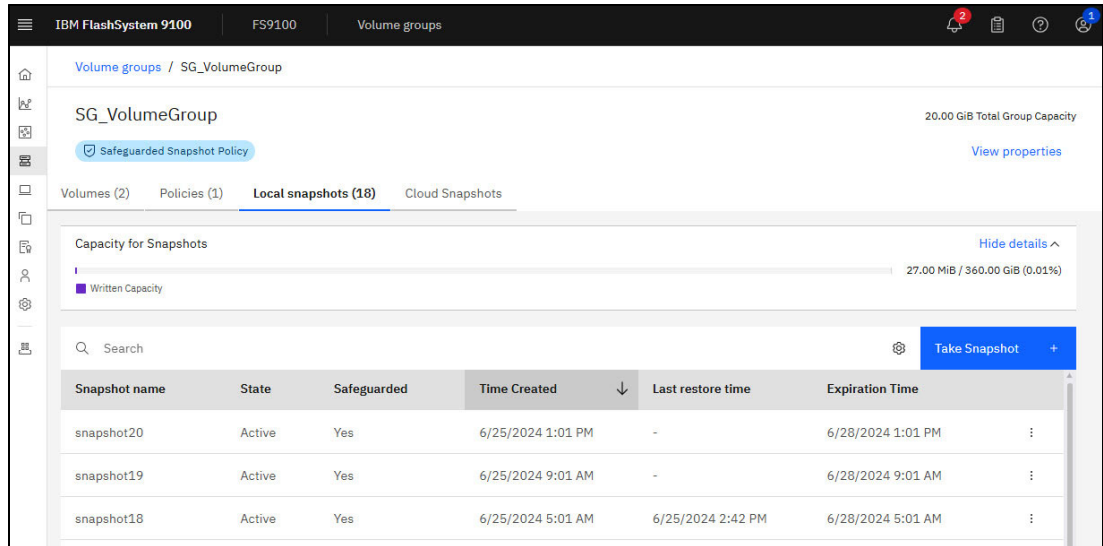


Figure 4-28 Snapshot last restore time

As with all management UI actions, these restores could have been executed using the CLI. The invocation for the restoring the entire volume group from the snapshot was:

```
restorefromsnapshot -snapshotid 19
```

and when a subset of volumes from a snapshot are restored the `-volumes` parameter is added:

```
restorefromsnapshot -snapshotid 18 -volumes 8
```

For more information, see [restorefromsnapshot](#).

4.5.4 Deleting a snapshot

When the end of a snapshot's retention period has been reached it will be automatically deleted. From time to time there will be a need to delete snapshots that have not reached the end of their retention period. It is important to remember that only a security administrator has the authority to delete a Safeguarded snapshot.

Note: If snapshot is safeguarded the delete option will not show up unless you are a security administrator or higher.

In the management UI when a non-security administrator views Safeguarded snapshots the only actions that they can initiate are **Create Thin Clone**, **Create Clone** and **Restore** (Figure 4-29 on page 61). A security administrator will see a fourth option, **Delete**. Once initiated the standard task pop-up will be displayed to report command progress. Control returns to the volume group's local snapshots view once the pop-up is closed.

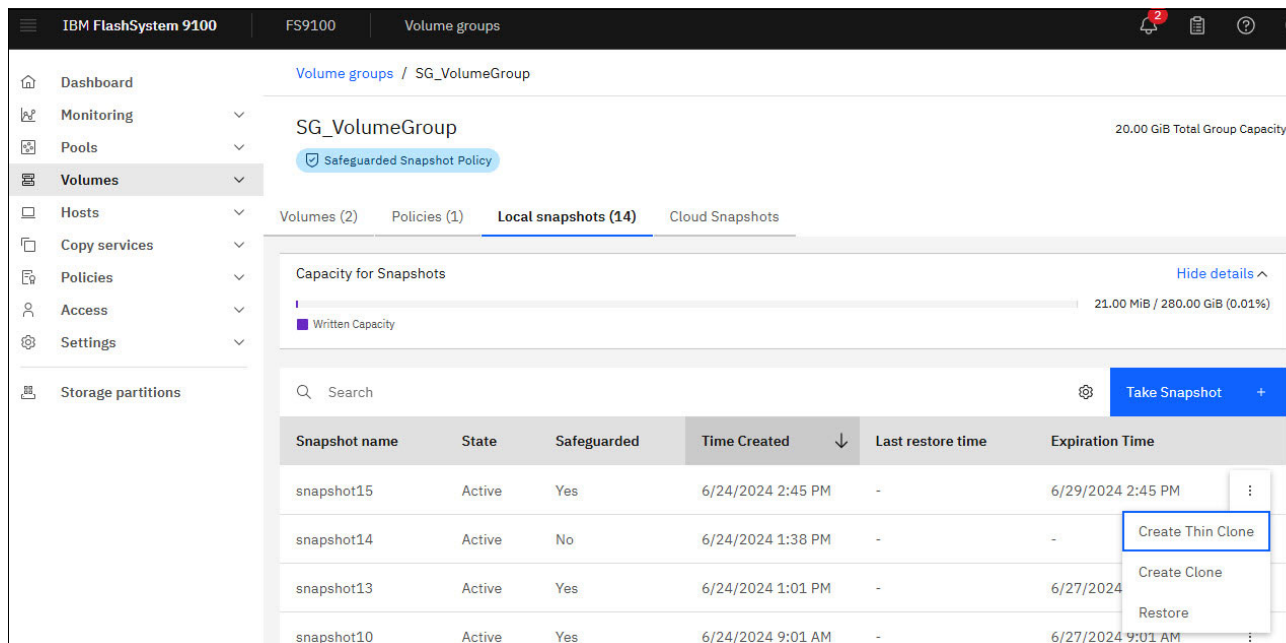


Figure 4-29 Non-security Administrator Safeguarded snapshot actions

If a non-security administrator were to execute the `rmsnapshot` command for a Safeguarded snapshot they would be told the snapshot was not deleted. Specifically, the message “*CMMVC8851E The command failed because the snapshot is Safeguarded and your requested action cannot be completed under your current user role.*” would be displayed.

Tip: You could run into a snapshot with a status of *dependent deleting*. This occurs when a snapshot has reached its expiration date, or has been deleted but has contingencies against it such as a thin clone.

4.5.5 Creating a clone or thin-clone from a snapshot

As indicated earlier Safeguarded snapshots prevent the accidental deletion or modification of critical data, ensuring data integrity. However, at some point, modification of the data may be necessary. A volume group cloning function has been enabled for this purpose.

There are two types of clones within IBM Storage Virtualize:

- ▶ Thin-Clone

A thin-clone is populated by a specified snapshot at the time of its creation. It will always be dependent on its source but can be modified by a host. Because the thin-clone is always dependent on the source it will normally be created in the same storage pool as the source, but this is not a hard requirement. Similarly, a thin-clone will normally be created in the same IO group as the source, though again it can be created in a different IO group.

- ▶ Clone

A clone, like a thin-clone, is populated by a specified snapshot at the time of its creation. Unlike a thin-clone it will become independent of its source after all the required data from the source has been copied by a background process. The clone can be modified by a host while the background process is in progress.

Once the background process is complete the clone will become stand-alone. A clone will typically be created in a separate storage pool to the source but this is not a requirement.

Note: All clones will *not* be safeguarded given they are modifiable by design.

Once the source snapshot is identified creating a clone is initiated using the snapshot actions (Figure 4-30).

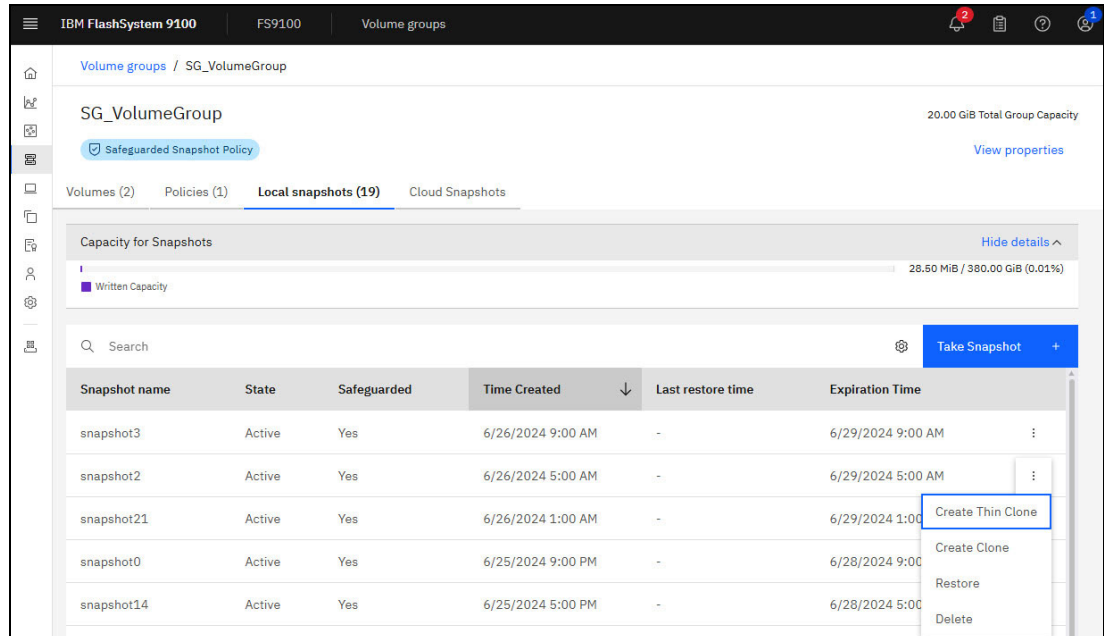


Figure 4-30 Initiate clone

Depending on whether the **Create Thin Clone** or **Create Clone** option is selected different choices will be presented.

When a clone is being created the **Create Volume Group** dialog (Figure 4-31 on page 63) will:

- ▶ Allow the user to define a name for the new group.
- ▶ Confirm the source volume group and snapshot.
- ▶ Enable the user to switch clone type, if required.
- ▶ Select a target pool.

If a name is not specified the system appends an index number to the source volume group. For example, if the source volume group name is *TestVG* the name for the cloned group would be *TestVG-0*. By default type Clone is enabled (denoted by check mark icon). One can switch to Thin-clone, if required. Finally, the target pool is specified using the drop-down menu. When the required information has been specified the task to create the clone is initiated by clicking **Create Volume Group**.

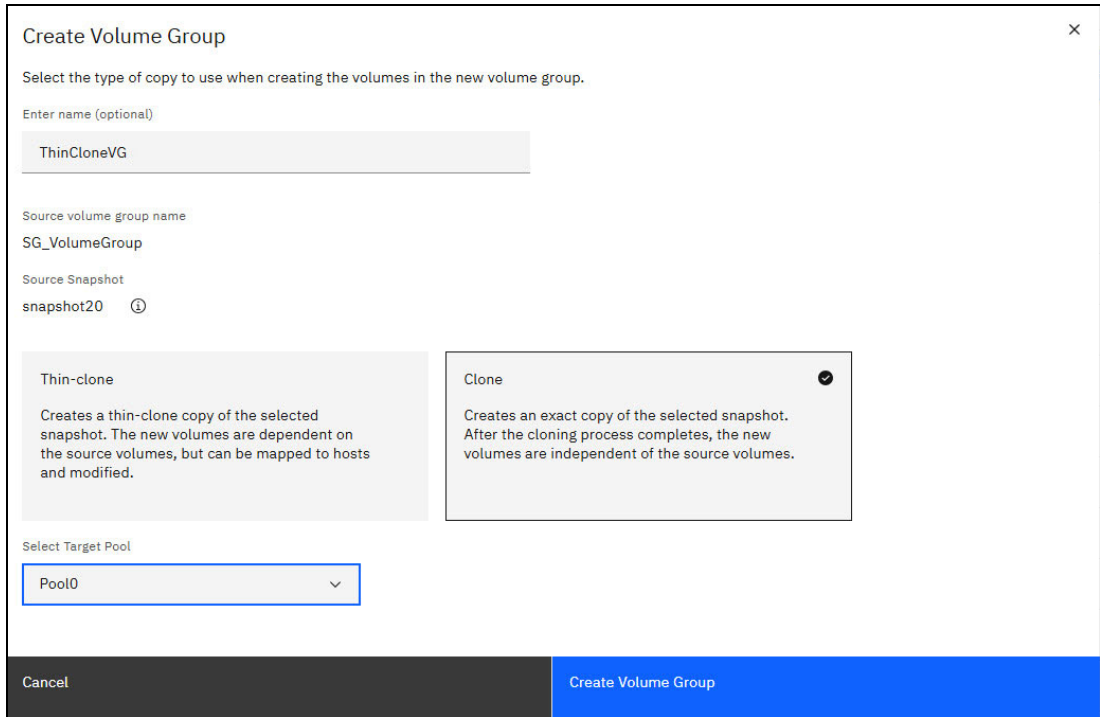


Figure 4-31 Create clone volume group

Progress on the creation of the clone is communicated using the task dialog (Figure 4-32).

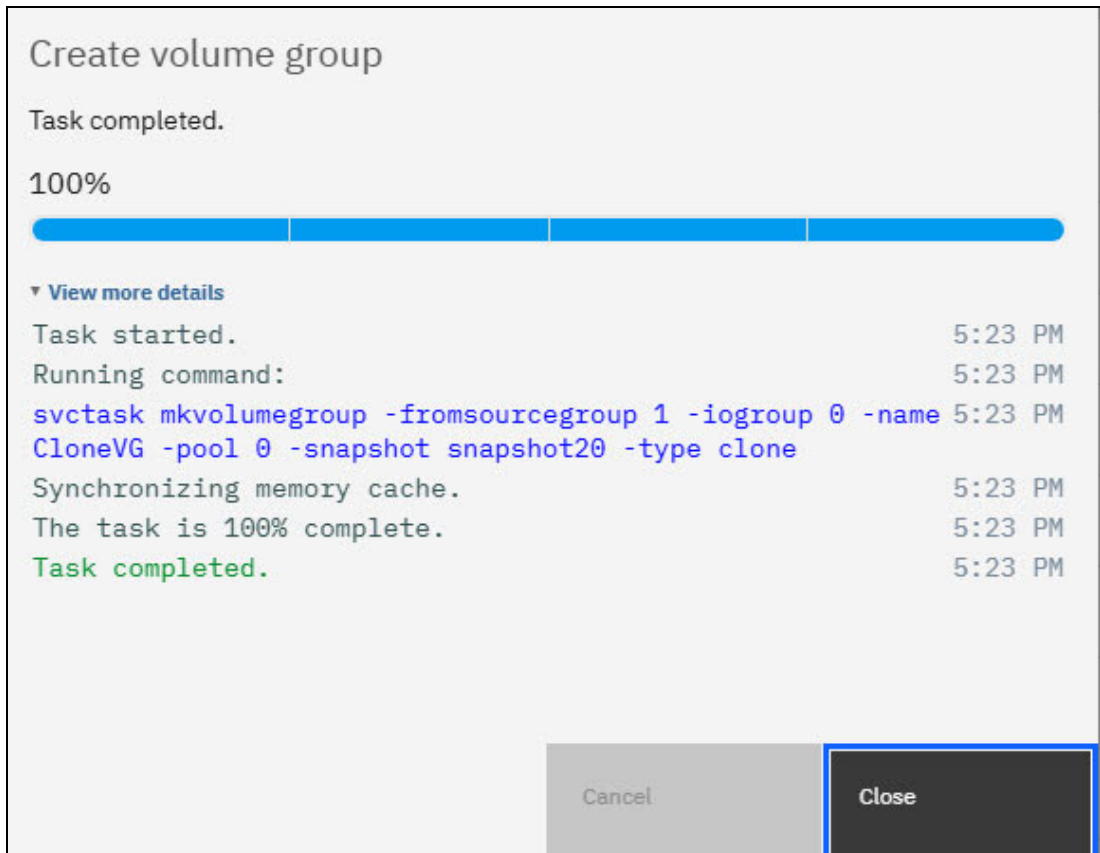


Figure 4-32 Create clone task

When the dialog is closed control returns to the source volume groups snapshots view.

The process flow when creating a thin clone is different than the clone process.

- ▶ Confirm the source volume group and snapshot.
- ▶ Allow the user to clone the entire volume group or subset of it.
- ▶ Specify a name for the new group.
- ▶ Select a target pool.

See the Create a thin clone dialog in Figure 4-33.

The screenshot shows a dialog box titled "Create a thin clone" with a close button (X) in the top right corner. Below the title is a descriptive sentence: "Creates a thin-clone copy of the selected snapshot. The new volumes are dependent on the source volumes, but can be mapped to hosts and modified." The dialog contains several input fields and controls: "Source volume group name" with the value "SG_VolumeGroup"; "Source snapshot" with the value "snapshot20"; "Time snapshot created" with the value "6/25/2024 1:01 PM"; "Select what to copy" with two radio buttons, "Volume groups" (selected) and "Subset of volumes"; "Enter name (optional)" with the text "ThinCloneVGx"; and "Select Target Pool" with a dropdown menu showing "Pool0". At the bottom, there are two buttons: "Cancel" on the left and "Create Volume Group" on the right.

Figure 4-33 Create thin clone

When creating a thin clone from a selection of volumes within the volume group, the dialog (Figure 4-34 on page 65) allows selecting specific volumes. If the entire volume group snapshot is the basis for the thin clone, **Create Volume Group** is enabled. Clicking this button initiates the task that creates the new volume group.

Note: Regardless of the volume selection during the thin clone creation, the resulting new volume group is accessible and manageable from the main volume groups view.

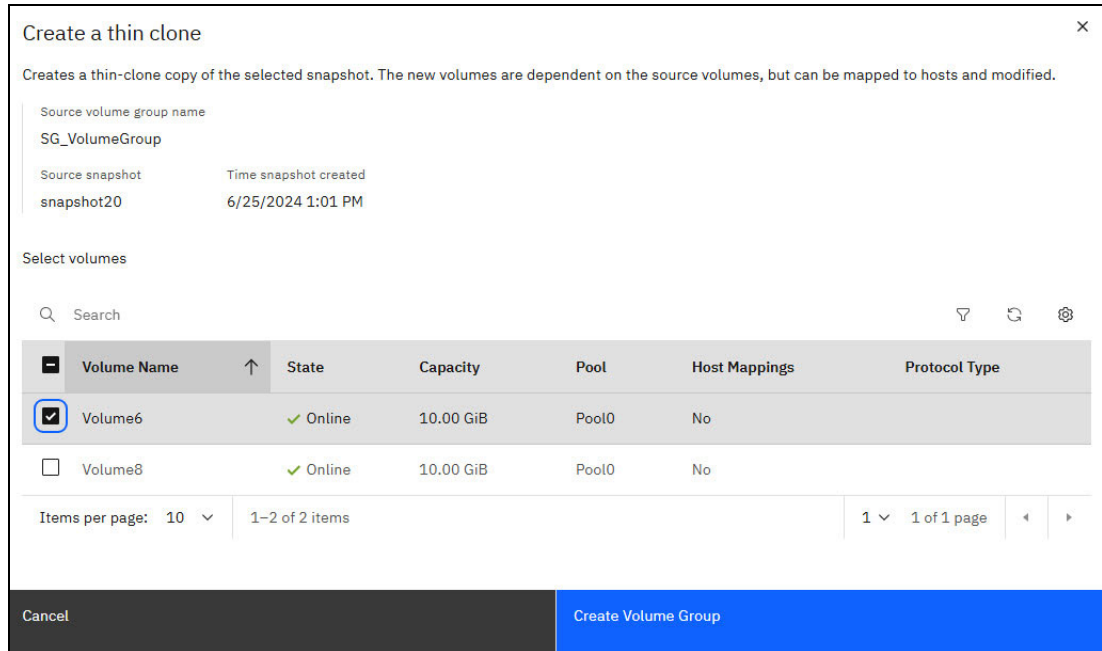


Figure 4-34 Thin clone snapshot volume selection

4.5.6 Refreshing a thin clone from a snapshot

After a thin clone has been created it is possible to refresh the content of the volume group from a snapshot.

This process is initiated from the volume group itself within the main volume groups view. (Figure 4-35).

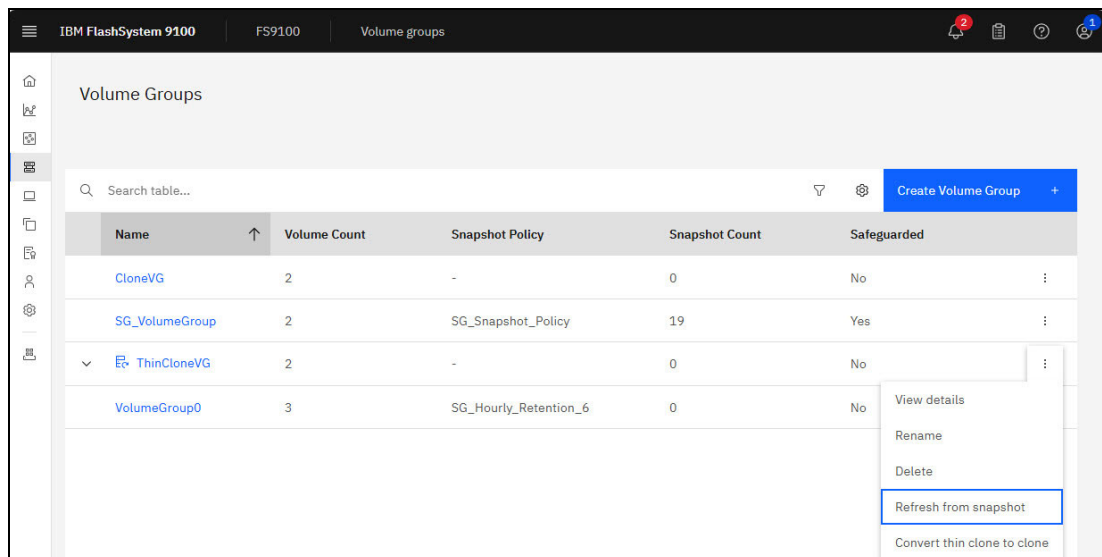


Figure 4-35 Refresh from snapshot

The dialog displayed (Figure 4-36 on page 66) will:

- ▶ Confirm the source volume group and snapshot used to create the thin clone.
- ▶ Display the list of snapshot available in the source volume group.

It is important to note that any of the listed snapshots can be used for the refresh. The one selected is dependent on the user's requirements.

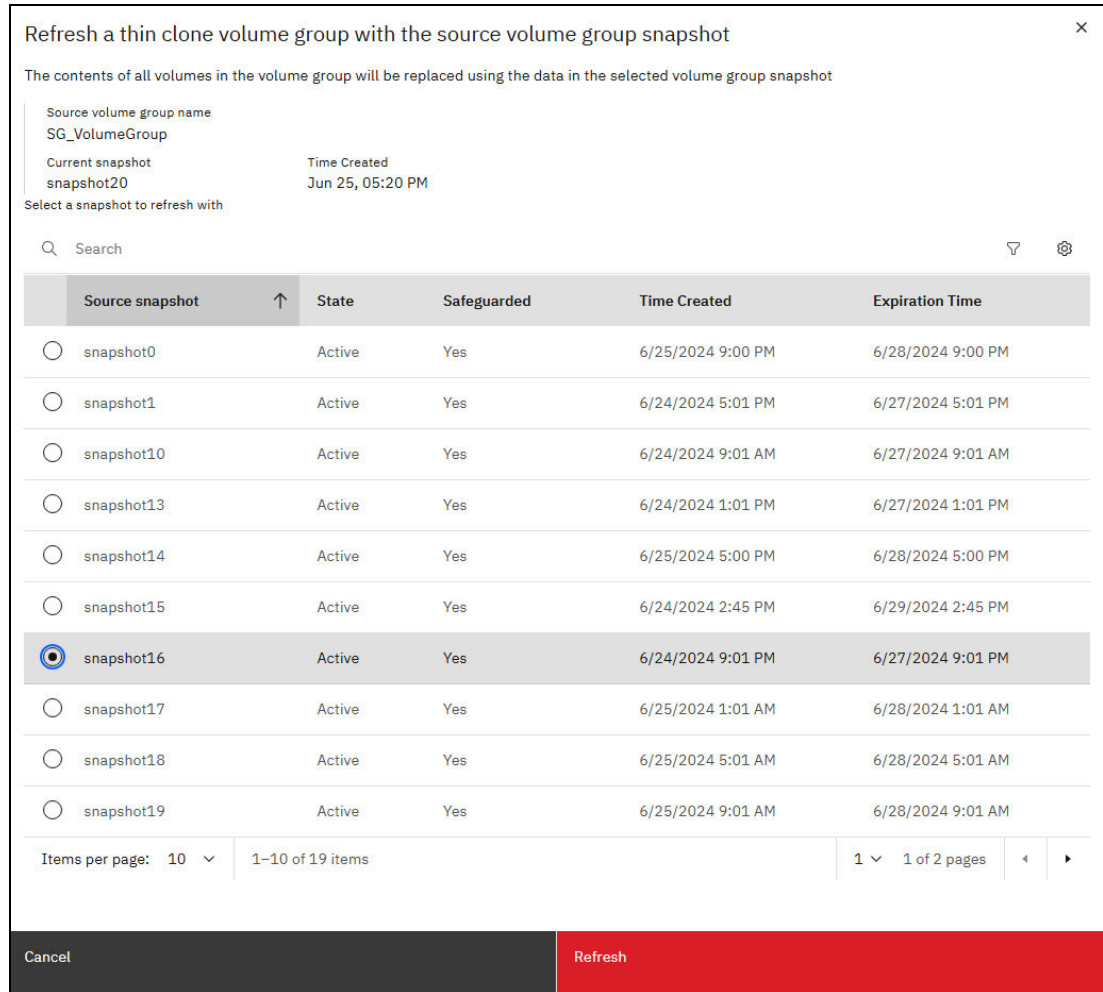


Figure 4-36 Thin clone refresh select

The refresh task is initiated once the user has provided confirmation, by entering the name of the thin cloned volume group and clicking the **Refresh** button (Figure 4-37 on page 67).

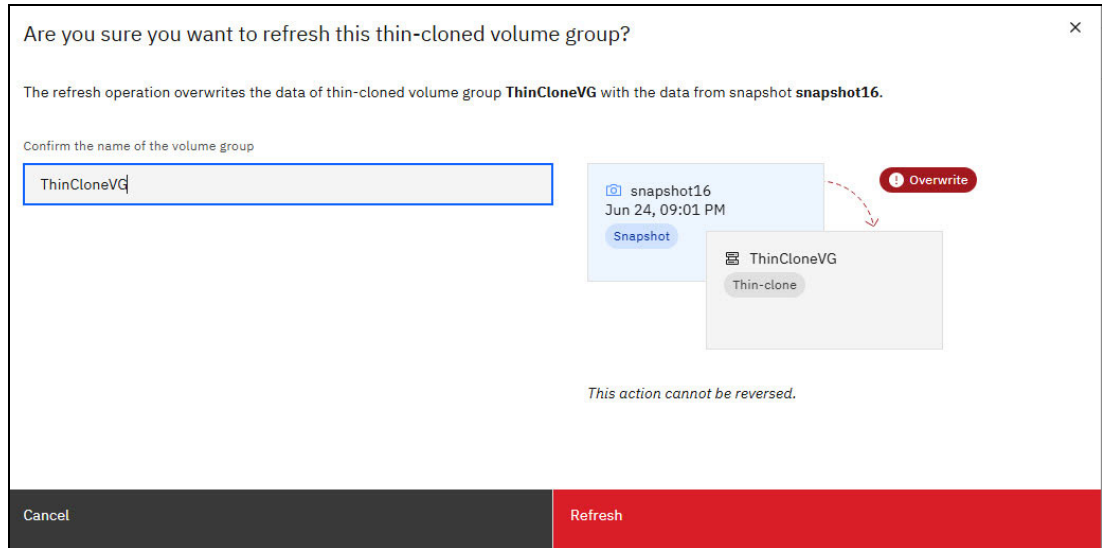


Figure 4-37 Confirm thin clone refresh

A task progress dialog is displayed. At this point the volumes with the refreshed volume group reflect the content of the selected snapshot.

4.5.7 Converting a thin clone to a clone

As noted in 4.5.5, “Creating a clone or thin-clone from a snapshot” on page 61, thin clone volume groups are dependent on the source volume group. In certain scenarios, the solution might require breaking the dependency on the source volume. This can be achieved by converting the thin clone into a fully independent clone.

To convert a thin clone select the **Convert thin clone to clone** option for the given volume group within the main volume groups view (Figure 4-35 on page 65).

The conversion process is simple.

Important: Converting a thin clone to a full clone severs its dependency on the source volume. This action cannot be undone and the management UI will not prompt for confirmation. Use this option only when complete independence is absolutely necessary.

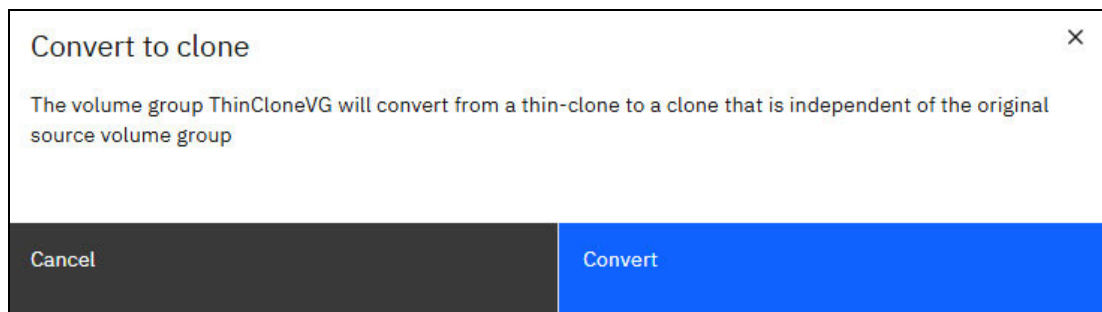


Figure 4-38 Convert to clone

The task is launched once the **Convert** button is clicked. See Figure 4-39 on page 68.

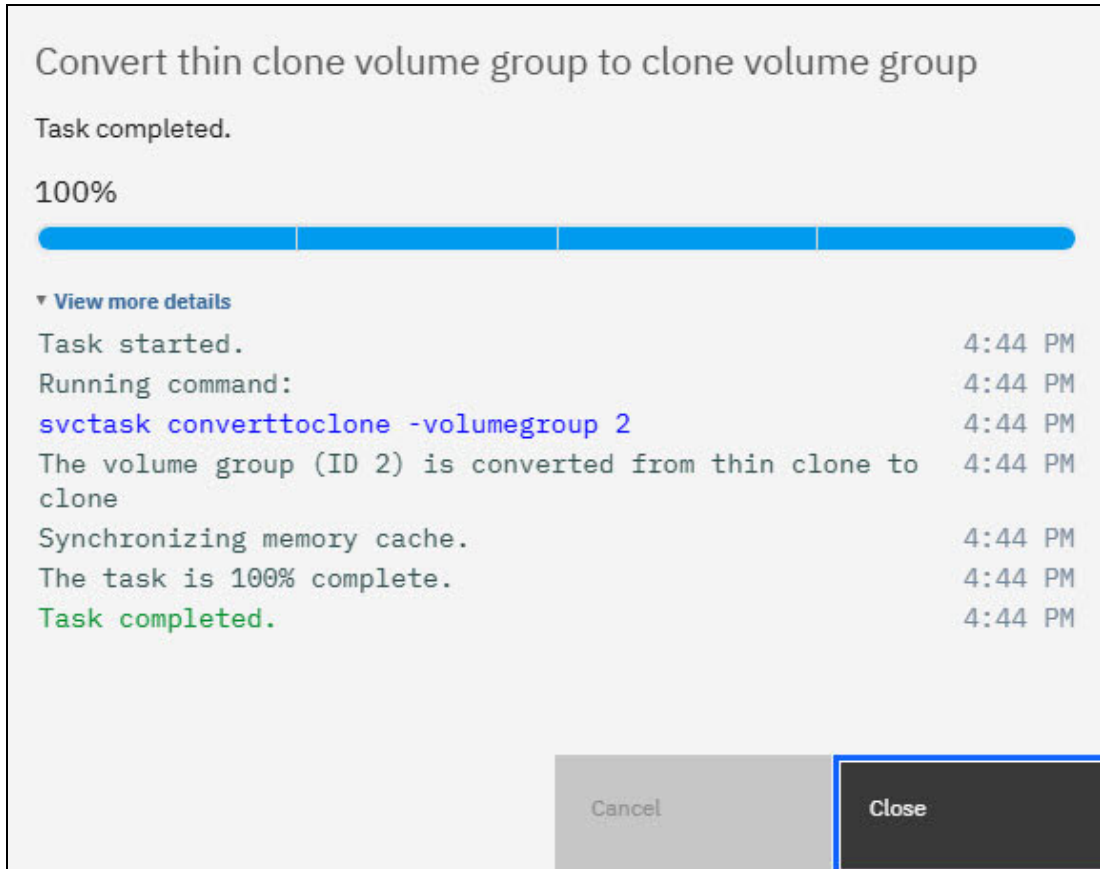


Figure 4-39 Convert to clone task

Once the task dialog is closed, control returns to the main volume group view. The converted volume group is now fully independent of the source volume group, and the thin clone visual cues are no longer present. See Figure 4-40.

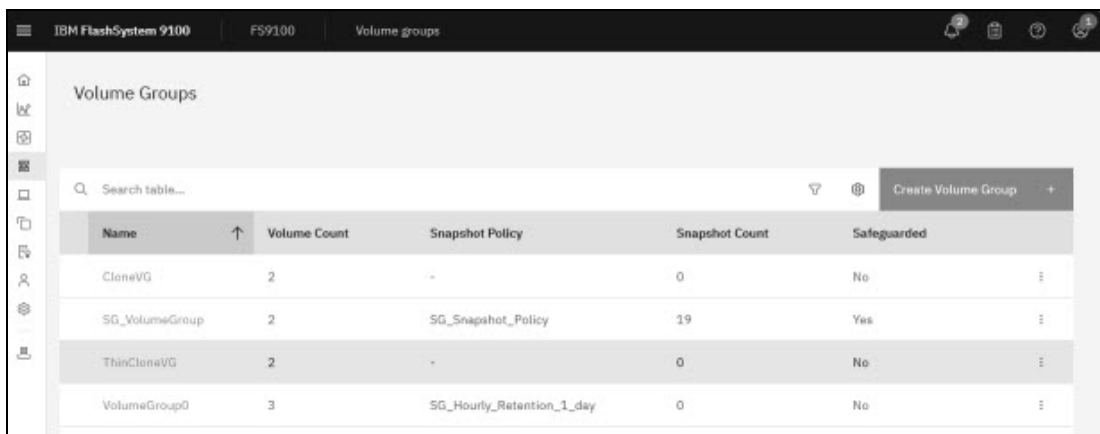


Figure 4-40 Convert to clone complete



Advanced considerations

This chapter discusses advanced considerations to be aware of when using Safeguarded snapshots. This includes using two person integrity (TPI) to protect against insider threats, performance considerations, and considerations when deleting volumes with active snapshots or snapshots with active thin clones. It also discusses IBM Storage Defender and IBM Storage Sentinel, which are two automation tools that utilize Safeguarded snapshots for data recovery.

This chapter has the following sections:

- ▶ “Two person integrity” on page 70
- ▶ “Performance considerations” on page 74
- ▶ “Preserve parent” on page 75
- ▶ “Volume or snapshot deletion considerations” on page 75
- ▶ “Integration with monitoring and automation tools” on page 76

5.1 Two person integrity

Two person integrity (TPI) can help protect against insider threats. An insider threat refers to security risks that come from within an organization, often involving individuals who have legitimate access to systems and data. These individuals may misuse their access for personal gain, causing harm to the company. It is crucial to recognize the potential for malicious intent from within your organization and take the necessary precautions to mitigate these risks.

Examples of insider threats can include:

- ▶ An employee with Security Administrator privileges sharing their credentials with a malicious third party, allowing them to gain unauthorized access to sensitive company data.
- ▶ A disgruntled employee who intentionally sabotages company systems or data to harm the organization.
- ▶ A contractor or temporary employee who retains access to company systems after their employment has ended, posing a risk of continued unauthorized access.

Insider threats can have severe consequences for organizations, including financial loss, damage to reputation, legal liabilities, and regulatory penalties. Therefore, it is essential for computer system administrators to implement robust access control measures.

Once TPI is enabled, a role elevation request and approval process is required to perform certain sensitive tasks:

- ▶ The restricted security administrator can issue a role elevation request on its own behalf to complete certain tasks in the system.
- ▶ Another restricted security administrator or a security administrator must approve the role elevation request.

For example, this role elevation request and approval process is required to remove a Safeguarded snapshot. The restricted security administrators or security administrators can approve or deny role elevation requests, cancel role elevation requests, or revoke a role elevation request that was approved.

When you enable TPI, the users that belong to user groups of security administrator role are assigned the restricted security administrator role and the superuser id is locked. However, their user groups retain their security administrator role.

After TPI is enabled, a role elevation request and approval process is required to perform certain sensitive tasks:

- ▶ The restricted security administrator can issue a role elevation request on its own behalf to complete certain tasks in the system.
- ▶ Another restricted security administrator or a security administrator must approve the role elevation request. For example, this role elevation request and approval process is required to remove a Safeguarded snapshot.
- ▶ The restricted security administrators or security administrators can approve or deny role elevation requests, cancel role elevation requests, or revoke a role elevation request that was approved.

Recommendation: As a best practice the approver should be monitoring the audit log of the system and be ready to revoke privileges if necessary.

The following actions are available for a restricted security administrator that has an approved role elevation request:

- ▶ Create, change, or remove security administrator user groups.
- ▶ Change the non-security administrator user group attribute on an existing local user to a security administrator user group.
- ▶ Modify attributes on existing local users that are members of the security administrator user groups.
- ▶ Change the role of existing non-security administrator user groups to the security administrator role.
- ▶ Change the security administrator role of an existing user group to a non-security administrator role.
- ▶ Remove and change Safeguarded backups and Safeguarded backup locations.
- ▶ Delete Safeguarded snapshots.
- ▶ Use a provisioning policy to define a set of rules that are applied when volumes are created within a storage pool or child pool.
- ▶ Change the single sign-on credentials that are used for the system.
- ▶ Remove the Safeguarded snapshot policy association from a volume group.

Enabling TPI

Before enabling TPI, the following prerequisites must be met:

- ▶ Two users with the security administrator role (excluding superuser).
- ▶ The two users can be local, remote, or a combination of both.
- ▶ If you use remote users, a remote user group of security administrator role must be defined on the system and the remote authentication service must be enabled.

To enable TPI, complete the following steps:

1. In the management GUI, select **Settings** → **Security** → **User Access** → **Two person integrity**.
2. Select **Enabled**.
3. Click **Save**. See Figure 5-1 on page 72.

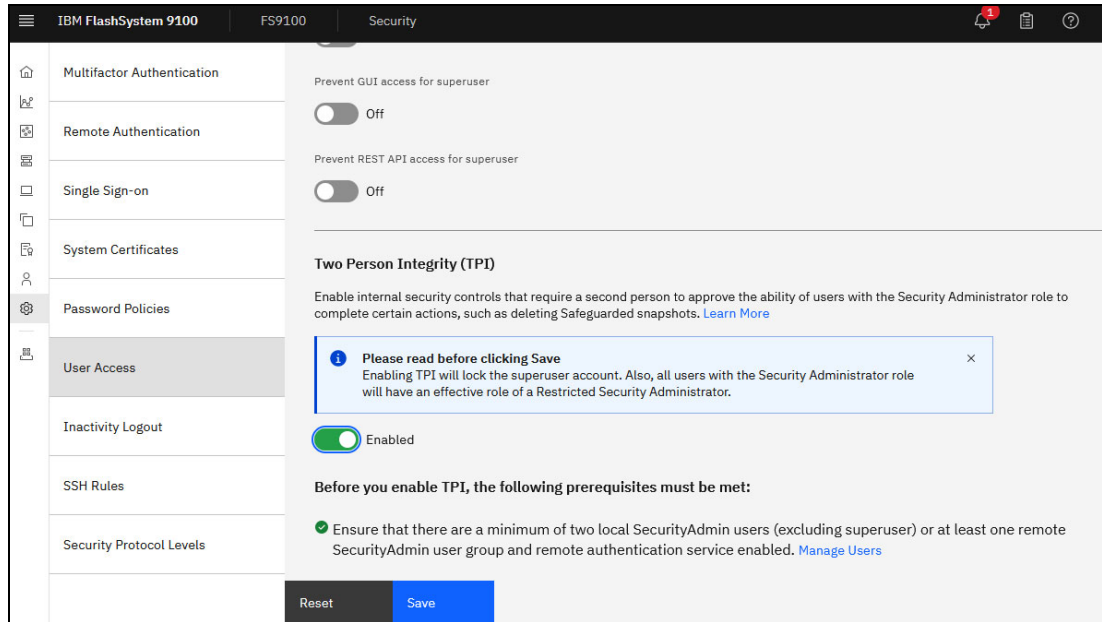


Figure 5-1 Enable TPI

4. If logged in as a superuser there will be a warning “The login credentials are no longer valid. Re-enter your credentials on the login panel” and you will be required to sign back in.
5. After you sign in, the page displays that the current role is updated to restricted security administrator, and *Manage Role Elevation Requests* and *Request Elevated role* are displayed in the user menu list see Figure 5-2.

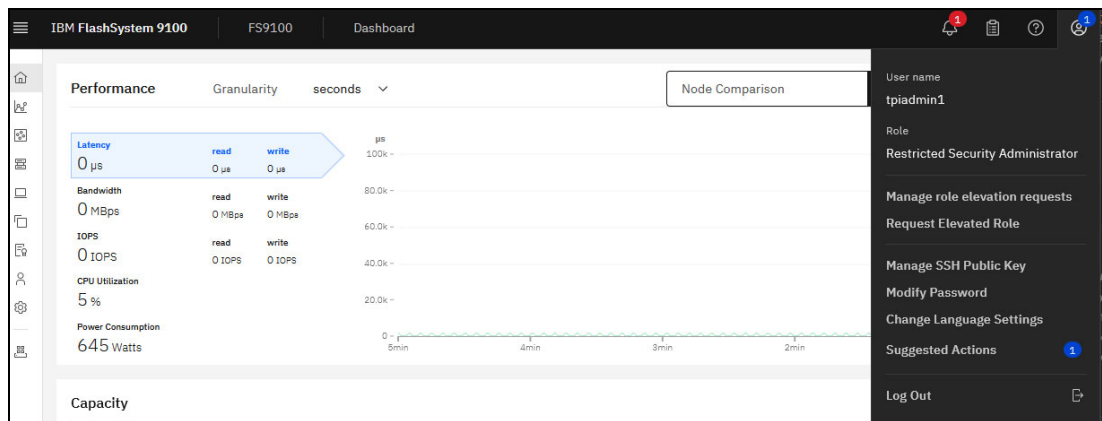


Figure 5-2 TPI User Menu

6. To disable TPI, click **Disabled**. After TPI is enabled, a user with an approved TPI request can disable TPI.

Note: After TPI is enabled, a user with an approved TPI request can disable TPI. When TPI is disabled the superuser id will be unlocked.

Managing TPI requests

Selecting **Manage Role Elevation Requests** will open a panel where requests can be made, approved, and revoked. In the example shown in Figure 5-3 on page 73, tpiadmin1 has requested a role elevation for one hour, which is now pending approval.

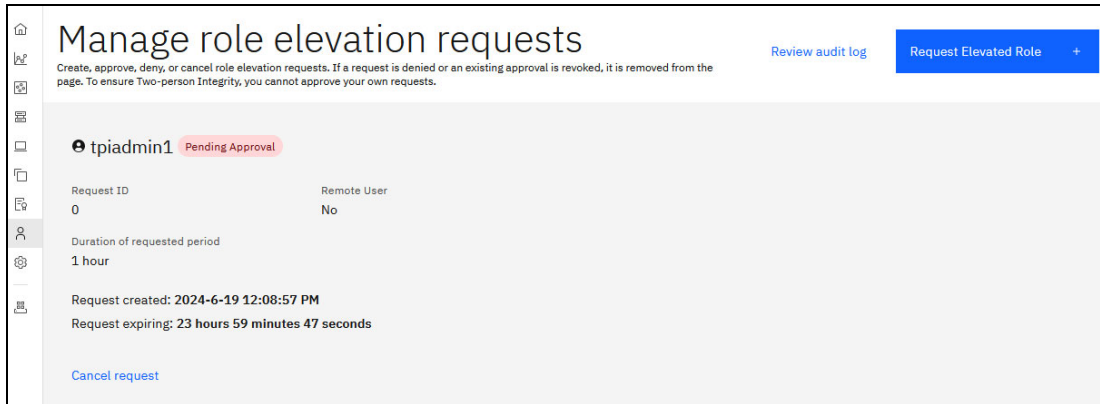


Figure 5-3 Role elevation request pending

A second restricted security administrator can then select *Manage Role Elevation Requests* and approve or deny the request. See Figure 5-4.

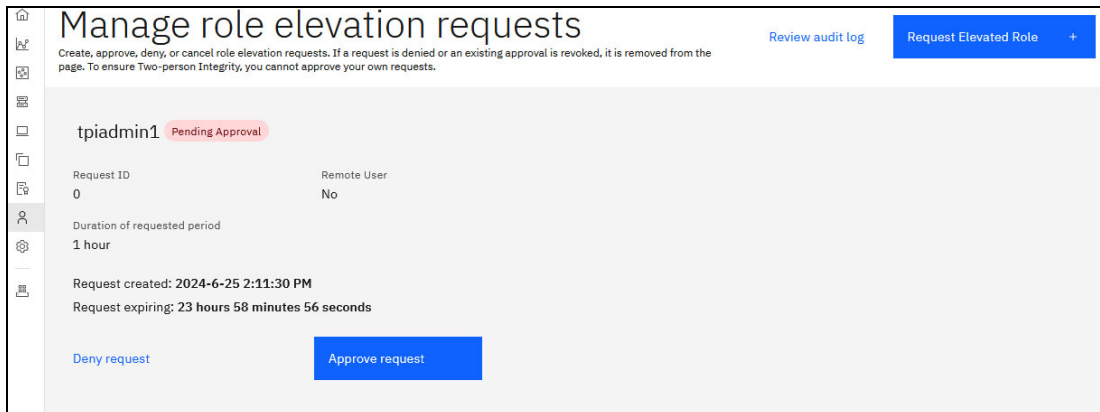


Figure 5-4 TPI approve request

Once the request is approved the second restricted security administrator can revoke the elevated role at any time. See Figure 5-5 on page 74.

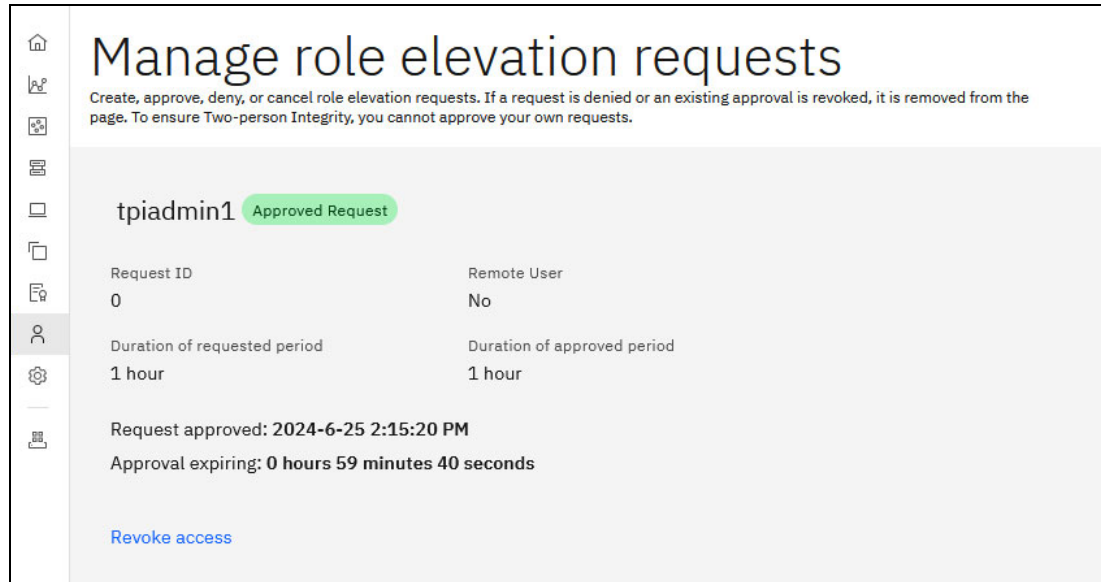


Figure 5-5 TPI Request Approved

Demonstration video: Although created in 8.6.0, the following video is still applicable to 8.7.0: [Configuring Two Person Integrity with IBM Storage Virtualize](#).

5.2 Performance considerations

There are two critical points in which performance may be affected due to snapshots: during the initial creation of the snapshots (trigger), and later as the source volume is written and the snapshot copy is maintained in the background.

When a trigger Snapshot operation starts, the write cache is flushed and a checkpoint is made of the source volumes. No data is copied at the time the trigger operation occurs. Instead, the checkpoint creates a bitmap that indicates that no part of the source volume was copied. Each bit in the bitmap represents one region of the source volumes called a grain. At the creation of the snapshot (previously called a FlashCopy mapping), the table is filled with zeros, which indicates that no grain is copied yet.

In 8.7.0 there is no performance impact at the very start of the snapshot despite the creation of this bitmap because of checkpointing technology allowing IO to continue. Code levels below 8.7.0 can see a short increase of write response time on the source volumes during this time. This is in the order of 10s of milliseconds when a snapshot is triggered. Most applications are not sensitive to this increase.

A second potential performance impact is dependent on the number of source volumes and the write workload on those volumes throughout the life of the snapshot. The bitmap updates and maintenance of the snapshots can cause I/O write amplification due to the grain size tracked by the snapshots. The write cache and extra internal bandwidth in the controller helps reduce this impact for certain workloads, but heavy write workloads can see an increased response time during write peaks. The lower-end IBM FlashSystem storage systems are more susceptible to this potential impact.

5.3 Preserve parent

As discussed in Chapter 2, “Safeguarded snapshot: Redefining data protection” on page 13, Safeguarded snapshots are point-in-time captures of the source (parent) volume, relying on data in the source volume for any grains that have not changed since the snapshot was initiated. By default, the Storage Virtualize system places the source volume offline when an event occurs that prevents it from maintaining the snapshot, such as when the pool containing the snapshots reaches its capacity. This action ensures that the system can copy the original grain to the snapshot volume before committing the write to the source, and the snapshots contain all valid data from the time the snapshot was initiated. However, this setting can be altered so that the parent volume remains online, and the snapshot is deleted instead. To modify this setting, utilize the `chsystem` and `lssystem` commands. See Example 5-1.

Example 5-1 lssystem command

```
IBM_FlashSystem:FS5200_6H2:superuser>lssystem | grep preserve
snapshot_preserve_parent no
IBM_FlashSystem:FS5200_6H2:superuser>chsystem -snapshotpreserveparent yes
IBM_FlashSystem:FS5200_6H2:superuser>lssystem | grep preserve
snapshot_preserve_parent yes
```

Attention: The downside to using preserve parent is that an attacker could fill up the system. When the pool fills up, Safeguarded snapshots are deleted to keep the parent online. This would leave no Safeguarded snapshots to restore from if the data were then corrupted.

5.4 Volume or snapshot deletion considerations

Snapshots are a point in time copy of the source volumes and in order to access the data on a snapshot a thin-clone or clone must be created from a snapshot to map to a host. This means that a snapshot is always dependent on the source volume, thin-clones are always dependent on the snapshot they were created from, and clones are dependent on the snapshot they were created from until the background copy has completed for the clone.

When a volume is deleted that has dependent snapshots, the volume transitions to deleting state. When a volume is in the deleting state it is no longer visible in the GUI and does not respond to host IO but capacity used is not released until the dependent snapshots have been deleted either because their retention time has expired or they have been manually deleted.

When a snapshot is deleted that has dependent thin-clones or clones where the background copy is complete, the snapshot will show dependent deleting. A snapshot will remain in dependent deleting state until the dependent thin-clones are deleted and clones have completed their background copy. Snapshot10 is Dependent deleting, as shown in Figure 5-6 on page 76.

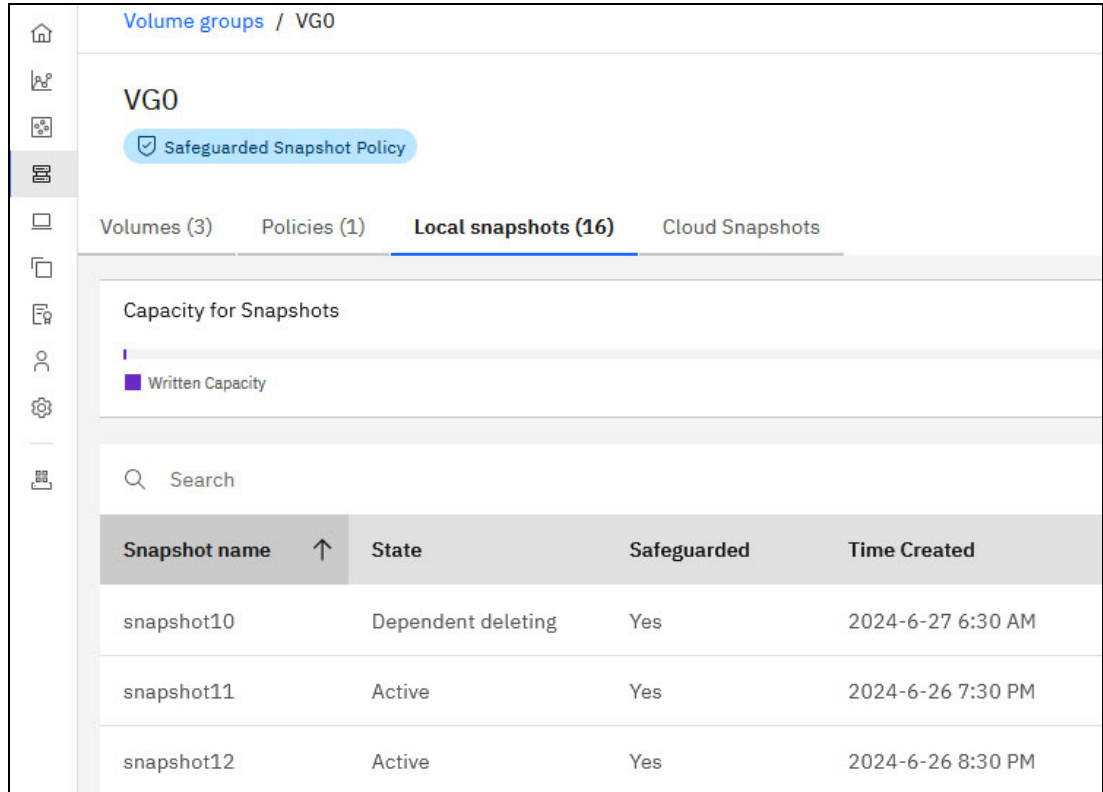


Figure 5-6 Dependent deleting

Viewing the thin clone volume group details will show the source volume group name and the snapshot it is dependent on. See Figure 5-7.

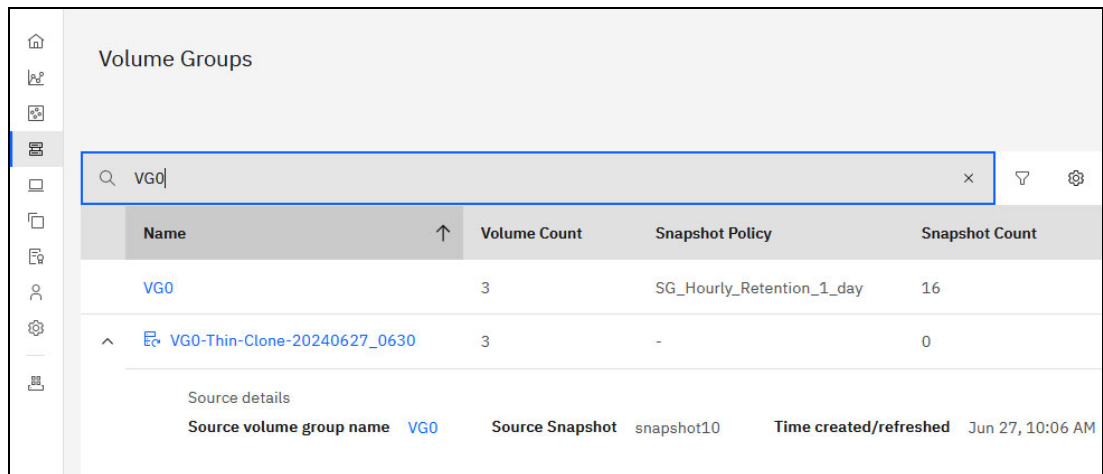


Figure 5-7 Thin clone details

5.5 Integration with monitoring and automation tools

Safeguarded snapshot can integrate with several integration tools, such as IBM Storage Defender and IBM Storage Sentinel. These are discussed in this section.

5.5.1 IBM Storage Defender

IBM Storage Defender is a comprehensive solution for end-to-end data resiliency across primary and secondary workloads. It provides multiple layers of protection to safeguard data against threats like ransomware, human error, disasters, and hardware failures. Key features of IBM Storage Defender include:

- ▶ **Threat detection:** Uses AI-driven intelligence and anomaly detection to identify potential threats early.
- ▶ **Data protection:** Offers immutable storage, hardware snapshots, backups, and air gap capabilities.
- ▶ **Recovery orchestration:** Helps identify the safest recovery points and orchestrates recovery at scale.
- ▶ **Unified management:** Provides a dashboard and management console for visibility across storage systems.
- ▶ **Integration:** Supports various enterprise workloads, databases, and cloud environments.
- ▶ **Compliance support:** Assists in meeting complex compliance requirements like GDPR, CCPA, PCI, and HIPAA.

IBM Storage Defender aims to streamline data recovery, enhance cyber resilience, and reduce downtime in case of data breaches or system failures. For more information see [IBM Storage Defender](#).

5.5.2 IBM Storage Sentinel

IBM Storage Sentinel is a cyber resiliency solution designed to help organizations enhance ransomware detection and incident recovery. Key features include:

- ▶ Creation of immutable application-specific primary storage snapshots using IBM Safeguarded snapshot technology.
- ▶ Frequent scanning of data copies to check for evidence of malware or ransomware damage.
- ▶ Use of anomaly detection and machine learning to identify potential threats.
- ▶ Generation of forensic reports to help quickly diagnose and identify attack sources.
- ▶ Orchestration of recovery from verified and validated backup copies.
- ▶ Intelligent isolation of infected backups to accelerate recovery time.

IBM Storage Sentinel is specifically configured for enterprise workloads like SAP HANA, Oracle and Epic Healthcare Systems. It is not meant to replace existing real-time security applications but serves as a last line of defense against data corruption during an attack. The solution aims to help organizations quickly identify ongoing attacks and recover clean data copies, reducing downtime and mitigating the impact of cyber threats. For more information see *Cyber Resiliency with IBM Storage Sentinel and IBM Safeguarded Copy*, [SG24-8441](#).

5.5.3 IBM Storage Insights

IBM Storage Insights has introduced new views for volume groups and snapshots, particularly focusing on Safeguarded snapshots. These enhancements provide deeper insights into storage management and protection.

- ▶ **Volume Group View:**
 - Detailed information about the volume group, including its relationship to Safeguarded snapshots.
 - Metrics related to capacity utilization, performance, and replication status (if applicable).

- Visualization of the volume group's hierarchy and its constituent volumes.
- Potential insights into the volume group's role in data protection strategies.
- ▶ **Snapshot View:**
 - Specific details about Safeguarded snapshots, including creation time, retention policy, and protection status.
 - Comparison of Safeguarded snapshots to standard snapshots.
 - Metrics related to snapshot size, growth rate, and impact on performance.
 - Visualization of snapshot relationships within the volume group and their role in data recovery plans.

For more information, see [IBM Documentation - What's new in IBM Storage Insights](#).

Ransomware threat detection alerts and marking volume snapshots as compromised

In addition to these new views, IBM Storage Insights Pro now alerts users about potential ransomware attacks across IBM Storage Virtualize systems, marking infected snapshots based on detection results. These alerts can be addressed directly or integrated into other security tools for further analysis.

For more information, see [IBM Documentation - What's new in IBM Storage Insights](#).

You can also refer to the IBM Redpaper *IBM FlashCore Module (FCM) Product Guide: Features the newly available FCM4 with AI-powered ransomware detection*, [REDP-5725](#).

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this paper.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- ▶ *Cyber Resiliency with IBM Storage Sentinel and IBM Storage Safeguarded Copy*, SG24-8541
- ▶ *IBM FlashCore Module (FCM) Product Guide: Features the newly available FCM4 with AI-powered ransomware detection*, REDP-5725
- ▶ *IBM FlashSystem Safeguarded Copy Implementation Guide*, REDP-5654
- ▶ *Unleash the Power of Flash: Getting Started with IBM Storage Virtualize Version 8.7 on IBM Storage FlashSystem and IBM SAN Volume Controller*, SG24-8561

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Online resources

These websites are also relevant as further information sources:

- ▶ [IBM Documentation - IBM Security QRadar XDR](#)
- ▶ [IBM Documentation -Safeguarded snapshot](#)
- ▶ [IBM Documentation - Storage Insights](#)
- ▶ [IBM Documentation - Storage Protect](#)

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



REDP-5737-00

ISBN 0738461717

Printed in U.S.A.

Get connected

