

Cyber Resiliency with IBM Storage Sentinel and IBM Security

Marcelo Cristiano Mendes Ayaviri

Nezih Boyacioglu

Juan Carlos Jimenez Fuentes

Vasfi Gucer

Akash Kushwah

Michael Mirochnik

Earl Springer

Markus Standau

Christopher Vollmar



Storage

Data Resilience



IBM Redbooks

Cyber Resiliency with IBM Storage Sentinel and IBM Security

November 2024

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

First (November 2024)

This edition applies to IBM Storage Virtualize Version 8.6 and higher.

Contents

Figures	ix
Tables	xiii
Examples	xv
Notices	xvii
Trademarks	xviii
Preface	xix
Authors	xix
Now you can become a published author, too!	xxi
Comments welcome.	xxi
Stay connected to IBM Redbooks	xxi
Chapter 1. Introduction	1
1.1 Overview of data resilience	2
1.1.1 Cybersecurity versus cyber resiliency	3
1.2 Approaches to data resiliency	5
1.2.1 Considering the restoration of static and dynamic data	6
1.2.2 Time to Recover	7
1.2.3 Secondary workload cyber resiliency	9
1.2.4 IBM Storage Sentinel as part of IBM Storage Defender	10
1.2.5 Supported applications	11
1.2.6 Use cases for Storage Sentinel	15
1.2.7 IBM Storage Sentinel workflow	15
1.2.8 IBM Storage Sentinel components	16
Chapter 2. IBM Storage Sentinel: An end-to-end automated cyber resiliency solution. 19	
2.1 An end-to-end automated cyber resiliency solution	20
2.2 IBM Storage FlashSystem	20
2.2.1 IBM Storage Safeguarded Copy	21
2.2.2 IBM Storage Copy Data Management	22
2.2.3 IBM Storage Sentinel: Anomaly scan software	24
2.2.4 IBM Security QRadar	25
2.2.5 IBM Security Guardium	27
Chapter 3. Scanning engine and its technology	29
3.1 Storage Sentinel architecture	30
3.2 The advantage of anomaly scanning versus signature scanning	31
3.2.1 The scanning process	32
3.2.2 Scanning process for databases	33
3.2.3 Machine learning	33
3.2.4 Scanning encrypted data	34
3.2.5 How to recognize and handle alerts	35
3.2.6 After alert workflow	36
3.2.7 What to do when the scanning engine finds an issue	36
3.2.8 Dealing with false positives	37
3.3 Scanning engine planning considerations	37

3.3.1	Sizing considerations	37
3.3.2	Federation setup and scaling of scan workloads	37
3.3.3	Virtual versus physical servers	38
3.4	Administration	38
3.4.1	Monitoring the scanning engine	39
3.4.2	Backing up and restoring the scanning engine components	39
3.4.3	Adding new applications	40
3.4.4	Adding new scanning engines	40
Chapter 4.	Protecting SAP HANA databases	41
4.1	Protecting SAP-HANA: A comprehensive approach	42
4.1.1	Components	42
4.2	Protecting SAP-HANA: Architecture and step by step implementation	43
4.2.1	Validating the SAP-HANA server requirements for CDM	45
4.2.2	Configuring IBM FlashSystem credentials	47
4.2.3	Defining storage SLA policies	50
4.2.4	Creating SAP-HANA backup jobs	52
4.3	IBM Storage Sentinel	57
4.3.1	Architecture	57
4.3.2	Server requirements	60
4.3.3	Sentinel installation	62
4.3.4	Sentinel federation	62
4.3.5	Configuring the manager engine - Sentinel federation integration	63
4.3.6	Adding member engines to the Sentinel federation	64
4.4	Integrating IBM Storage CDM, IBM Storage FlashSystem and IBM Storage Sentinel	65
4.4.1	IBM FlashSystem best-practices for administrative users	70
4.5	IBM Security Guardium for SAP HANA	72
4.5.1	Data security concerns for SAP HANA environments	72
4.5.2	Data security controls for SAP HANA	73
4.5.3	Demo: IBM Guardium and QRadar Preventing Attackers with Safeguarded Copy and Copy Services Manager	74
Chapter 5.	Configuring IBM Storage Sentinel for VMware	77
5.1	VMware configuration on Copy Data Management	78
5.1.1	Registering VMware vCenter	78
5.1.2	Registering Storage Sentinel Security Scan Server	79
5.1.3	Configuring SLA policies	80
5.1.4	Safeguarded Snapshot of critical VMs	81
5.2	Scanning process	82
5.3	Monitoring	83
5.4	Restore and recovery	85
Chapter 6.	Protecting EPIC Cache and IRIS use case (on AIX)	87
6.1	Introduction	88
6.2	Supported configurations for IBM Storage Copy Data Management and IBM Storage Sentinel for Epic databases	88
6.3	Setting up a CDM and Storage Sentinel environment to scan Epic databases	88
6.4	Scanning process	99
6.5	Performing a restore of an Epic database backup	101
Chapter 7.	Secure and resilient AI	105
7.1	Exploring LLMs: Applications, risks, and security in AI Systems	106
7.1.1	Applications of Large Language Models	106
7.1.2	Grounding and context	106

7.2	Examples of risks associated with Large Language Models	106
7.2.1	Adversarial risks across AI models	106
7.2.2	Prompt injection attacks	107
7.3	Moderation and instruction tuning	108
7.3.1	Identity management	108
7.3.2	Identity Threat Detection and Response	108
7.3.3	Continuous monitoring and feedback loops	110
7.3.4	Conclusion	110
7.4	Understanding system instructions, fine-tuning, and vector databases	110
7.4.1	The nature of user interactions	111
7.4.2	System instructions: The foundation of interaction	111
7.4.3	The role of fine-tuning	111
7.4.4	Vector databases: Managing context	111
7.4.5	Retrieval-Augmented Generation	111
7.5	Effective risk management	112
7.6	Protecting sensitive information	112
7.6.1	IBM Guardium for AI - Discovering and protecting sensitive data	112
7.6.2	The interrelationship of data and models	115
7.6.3	Securing models: Addressing LLM injection	115
7.6.4	Conclusion	116
7.7	Securing Generative AI: Threat vectors, data protection, and advanced applications	117
7.7.1	Advanced applications: The role of agents in relevance to security	117
7.7.2	Understanding agents in generative AI	117
7.7.3	Threat vectors in Generative AI	117
7.7.4	Data protection mechanisms	118
7.8	Security measures for LLMs	119
7.9	Integrated security technologies for AI management	120
7.9.1	Toolkit for AI Governance	120
7.10	IBM Guardium AI Security: Manage data model security risk demo overview and key capabilities	122
	Related publications	123
	IBM Redbooks	123
	Online resources	123
	Help from IBM	123

Figures

1-1	Traditional recovery versus cyber recovery	2
1-2	Cyber security and cyber resilience	3
1-3	IBM Point of View: Levels of data resilience	6
1-4	Recovery from ransomware attack timeline	8
1-5	Segmentation of workloads	10
1-6	How Sentinel fits in IBM's data resiliency workflow and IBM Storage Defender	11
1-7	An example Oracle Backup job	13
1-8	An example of a Backup joblog containing a detected anomaly	13
1-9	An example of a Backup joblog with no detected anomaly	14
1-10	VM scanning job details	14
1-11	Storage Sentinel in Cyber Vault Blueprint	15
1-12	Storage Sentinel attack timeline	16
1-13	Five steps to cyber resilience	17
2-1	End-to-end automated cyber resiliency solution	20
2-2	IBM FlashSystem with Safeguarded Copy	21
2-3	IBM Storage Copy Data Management	23
2-4	IBM Storage Copy Data Management orchestrating and automating	24
2-5	Some of the validations performed by the anomaly scan software	25
2-6	IBM FlashSystem with Safeguarded Copy Triggered by IBM QRadar SEIM	26
3-1	IBM Storage Sentinel Scanning workflow	31
3-2	Threat detected message	35
3-3	Job log showing details on the detected corruption	35
3-4	IBM Storage Sentinel Scanning Engine dashboard	36
3-5	IBM Storage Sentinel configuration with an AIX proxy machine on the Power platform	38
3-6	IBM Storage Sentinel engine status	39
3-7	IBM Storage Sentinel service status	39
4-1	Common infrastructure to protect SAP-HANA	42
4-2	CDM snapshot process	44
4-3	SAP-HANA database and IBM CDM Server interaction	46
4-4	Create a storage administrative user through the GUI	47
4-5	Register storage subsystems and their credentials	48
4-6	Configure CDM to connect to the storage subsystem	49
4-7	Set the storage credentials	49
4-8	Storage credentials	50
4-9	IBM Storage Virtualize for Snapshot option	50
4-10	Add Safeguarded Copy option	51
4-11	Set the copy frequency	51
4-12	Options tab	52
4-13	SLA Policies tab	52
4-14	Select the Register option	53
4-15	Select SAP HANA as the server type	53
4-16	Register Application Server	54
4-17	Create the system credentials	55
4-18	Create the database credentials	55
4-19	Application Server Provider Browser	56
4-20	Choose Backup	56
4-21	Choose the database to be protected and the SLA policy	56
4-22	Enable Schedule	57

4-23	Next Runtime	57
4-24	Both SAP HANA and Sentinel servers are deployed physically	58
4-25	Architecture with physical and virtual SAP HANA servers	59
4-26	The third scenario with the iSCSI protocol for communication	60
4-27	Sentinel federation	63
4-28	Sign in panel	63
4-29	Setup License	63
4-30	Sentinel monitoring interface	64
4-31	Sign in panel	64
4-32	Deselect default index	65
4-33	Join Federation	65
4-34	Manager engine name is registered	65
4-35	Select Register	66
4-36	Register Security Scan Server	67
4-37	Security Scan API Credential	67
4-38	Create Credential	68
4-39	required indexes are provisioned on each engine	68
4-40	Index Manager section of the Sentinel administration console	68
4-41	Security Scan Server tab	69
4-42	Perform Security Scan every option	69
4-43	Security scans will run after each Safeguarded Snapshot	69
4-44	iSCSI connectivity between IBM FlashSystem, IBM Storage Sentinel servers, and host systems	70
4-45	SAN zoning and host connections for the ESXi hosts	70
4-46	Audit log showing IP address and hostname where the command originated	72
4-47	Data security proactive approach provided by IBM Security	74
4-48	Deployment options for using IBM Guardium and SAP-HANA	74
5-1	Registering VMware vCenter	78
5-2	Registering Sentinel Security Scan Server	79
5-3	Sentinel Indexes	79
5-4	Configuring SLA policy and Storage Sentinel Scan frequency	81
5-5	Defining a backup	82
5-6	Job details on CDM	82
5-7	VMware job history	83
5-8	IBM Storage Sentinel scan list	83
5-9	IBM Storage Sentinel Hosts scanned file statistics	84
5-10	Edit settings for creating alerts	84
5-11	IBM Storage Sentinel alerts	85
5-12	VMware restore options	85
6-1	Register the LDAP server	89
6-2	Import LDAP Group	90
6-3	Configure Sites	90
6-4	Register your storage components	90
6-5	Register your vCenter server(s)	91
6-6	Register your Epic DB application servers	92
6-7	Register your Epic DB application servers	92
6-8	Register your Storage Sentinel server(s)	93
6-9	Click on the job log hyperlink to open the job log page	94
6-10	Job log page	94
6-11	Define an SLA for your Epic DB data protection	95
6-12	Add Safeguarded Copy	95
6-13	Completing the SLA configuration	96
6-14	Register the AIX proxy FileSystem -1	97

6-15 Register the AIX proxy FileSystem -2	97
6-16 Defining a backup job	98
6-17 Select your SLA	98
6-18 Monitor your data protection	99
6-19 No threats detected	99
6-20 Threats detected	100
6-21 IBM Storage Sentinel scan list	100
6-22 IBM Storage Sentinel alerts	101
6-23 Select Restore	101
6-24 Instant Database Restore	102
6-25 Select the database to be restored	102
6-26 Click the Copy icon	102
6-27 Select a specific version	102
6-28 Create a Job	103
7-1 Adversarial risks across AI models	107
7-2 Model inference	107
7-3 IBM Verify Identity Management Suite	108
7-4 Identity Threat Detection with Identity Fabric	109
7-5 Verify Identity Protection Platform	109
7-6 Verify Identity Protection	110
7-7 Securing AI pipeline - Attacker versus security perspective	112
7-8 IBM Guardium AI Security overview	113
7-9 RAG customer example securing data with Guardium Data Protection	113
7-10 IBM Guardium AI-specific use cases	114
7-11 Guardium for AI Data Monitoring and Protection	114
7-12 Guardium Data Protection for vector capable databases	118
7-13 Extend exiting security across the underlying AI infrastructure	119
7-14 IBM Guardium AI Security	120
7-15 IBM watsonx.governance	121
7-16 IBM watsonx.governance and IBM Guardium AI Security: A unified approach	121
7-17 Integration use cases IBM watsonx.governance with IBM Guardium AI Security	122

Tables

4-1 Ports used for Sentinel 61

Examples

4-1	hdbuserstore list	45
4-2	Installing the HANA client	45
4-3	Create the cdmgroup group in a Linux server.	46
4-4	Create the cdmgroup group in a Linux server.	46
4-5	Configuring sudoers for the cdmagent user	46
4-6	Creating a user to interact with CDM	46
4-7	mkuser	47
4-8	systemctl status firewalld	60
4-9	Configuring firewall exceptions	61
4-10	Check if Mail Transfer Agent (MTA) is running.	61
4-11	/etc/hosts/	62
4-12	mkusergrp command	70

Notices

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

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

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

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

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

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

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

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

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

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

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

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


COPYRIGHT LICENSE:

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

Trademarks

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

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

AIX®	IBM FlashSystem®	Redbooks (logo)  ®
DS8000®	IBM Security®	Resilient®
FlashCopy®	IBM Spectrum®	Tivoli®
Guardium®	Passport Advantage®	X-Force®
IBM®	QRadar®	
IBM Cloud®	Redbooks®	

The following terms are trademarks of other companies:

ITIL is a Registered Trade Mark of AXELOS Limited.

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

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

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

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

VMware, VMware vSphere, 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 data-driven world, safeguarding sensitive information and ensuring uninterrupted business operations is a top priority. This IBM® Redbooks® explores how the integrated solution of IBM Storage Sentinel, IBM Storage FlashSystem, IBM Storage Copy Data Management, IBM Security® Guardium®, and IBM Security QRadar® empowers organizations to protect their valuable assets, minimize downtime, and maintain operational resilience. By leveraging these products, organizations can effectively detect, prevent, and recover from cyberattacks, ensuring business continuity and minimizing downtime.

This book empowers security, AI, and storage professionals to build a secure and resilient IT infrastructure. You will gain expert guidance on configuring IBM Storage Sentinel for critical applications like VMware, EPIC, and SAP HANA. By leveraging Sentinel's advanced features, you will optimize performance, enhance security, and ensure business continuity.

Additionally, the book explores how IBM Guardium strengthens AI environments by protecting sensitive data, monitoring user access, and detecting potential threats, safeguarding the integrity and confidentiality of AI models and data.

Authors

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

Marcelo Cristiano Mendes Ayaviri is a Senior Storage Consultant in IBM Expert Labs with years of experience, specializing in storage and backup solutions for South America and Mexico. His expertise spans a wide range of IBM technologies, including IBM Storage Protect and Software-Defined Storage. He has a proven track record of supporting customers and partners, implementing complex solutions, and delivering insightful training sessions. As a frequent speaker at Tech-U from 2014 to 2020, he has shared his knowledge with the broader storage community, particularly focusing on data protection against cyber threats.

Nezih 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. Nezih'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 Storage Virtualize family and Storage Area Networks. Nezih's commitment to expertise is evident in his IBM certifications, including Enterprise Storage Technical Support, Flash Technical Solutions, Virtualized Storage, and Storage Virtualize Storage software. He is a co-author of several IBM Redbooks publications.

Juan Carlos Jimenez Fuentes is the World-Wide Data Resiliency Product Manager based in Dallas, Texas. He is focused on defining roadmap, initiatives, and strategy within the various data resiliency software products that he manages alongside his team. Juan Carlos brings an end-to-end view to cyber resilience leveraging his expertise in both storage and security. Juan Carlos developed the IBM Cyber Resiliency Assessment Tool which has been helping numerous enterprises identify and close gaps in their IT environments. He holds a Management Information Systems Degree from the University of Arizona.

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.

Akash Kushwah is a Lead Application Developer in Storage Copy Data Management and the IBM Storage Sentinel product with IBM Systems, ISDL Lab Pune, India. He has worked extensively on products that offer data protection, including Backup and Restore, Disaster Recovery, Business Continuity, and Cyber Resiliency. Recently He has developed the cyber resilience feature for the IBM Storage Sentinel product, which involves scanning AIX-based workloads for various applications like Oracle, EPIC, and SAP HANA, as well as VMware VM Sentinel scanning. He has been a key developer who worked on onboarding the support for different Storage vendors in the SCDM product.

Michael Mirochnik is a seasoned Principal Security Architect at IBM Security, with over 25 years of expertise in cybersecurity. His extensive career spans several key areas, including Quantum Computing, Artificial Intelligence, Cloud Security, Zero Trust, and Ransomware Detection & Response. Additionally, Michael is highly proficient in conducting Penetration Testing, Red Team Testing, and Risk Management. He has been instrumental in helping organizations design and implement strategic security initiatives, with a focus on Data Security, Zero Trust, Cyber Resiliency, AI, Identity, policy development, and remediation efforts. His practical experience also extends to intrusion protection and performing comprehensive gap analysis to identify and mitigate security risks. Michael's career is defined by his commitment to assisting enterprises in protecting their digital environments against emerging threats, ensuring a proactive and robust defense.

Earl Springer is a Senior Partner Technical Specialist supporting IBM Storage hardware and software technologies for IBM Ecosystem. Working with IBM Value Add Distributors and IBM Business Partners across North America, Earl helps co-create opportunity-driven solutions, develops and delivers enablement, and helps IBM Business Partners understand the value and differentiators of IBM storage technologies. Originally working with IBM Canada, Earl now works with IBM US in Tampa, FL. He has an extensive technical background in software development, compute technologies including IBM Power, networking, virtualization, cloud, and storage technologies. Earl is a designated Subject Matter Expert for IBM Storage Virtualize, helped develop certification and education content for several hardware and software storage certification exams, and helped develop and deliver education for various IBM technical events, including IBM TechU.

Markus Standau works for IBM Germany. He has more than 20 years of experience in the storage field in roles such as services, technical sales, and worldwide product management. He currently works in the storage sales acceleration team as the offering leader for Storage Virtualize, FlashSystem, and the Storage Control family in DACH. In his current role Markus organizes various business partner and customer events in DACH, such as IBM Storage Strategy Days, the SVC/FlashSystem user group and more. He holds a degree in Computer Science from Baden-Wuerttemberg Cooperative State University. Markus is the co-author of several IBM Redbooks on IBM Spectrum® Control and its predecessor, Tivoli Storage Productivity Center (TPC).

Christopher Vollmar is an Principal, World Wide Storage Data Resiliency Architect. Christopher is an IBM Certified IT Specialist (Level 3 Thought Leader) and Storage Architect. He is focused on helping customers design solutions to support Operational and Cyber Resiliency on primary and backup data to complement their Cyber Security practices. He is an author of several IBM Redbooks, an Enterprise Design Thinking Co-Creator, and a frequent speaker at events like IBM THINK, and TechXchange.

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

Pepe Lam
IBM USA

Shashank Shingornikar
IBM India

Jamie Roszel
IBM RTP

Now you can become a published author, too!

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

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

Comments welcome

Your comments are important to us!

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

- ▶ Use the online **Contact us** review Redbooks form found at:

ibm.com/redbooks

- ▶ Send your comments in an email to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, IBM Redbooks
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- ▶ Find us on LinkedIn:

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

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

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

- ▶ Stay current on recent Redbooks publications with RSS Feeds:

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



Introduction

This chapter introduces the concepts of cyber resiliency and cybersecurity, highlighting their key differences. It also provides an overview of IBM Storage Sentinel, including supported applications, use cases, and the workflow.

This chapter has the following sections:

- ▶ “Overview of data resilience” on page 2
- ▶ “Approaches to data resiliency” on page 5

1.1 Overview of data resilience

Just a few decades ago, data resilience was a much simpler matter. If a company lost or damaged an important file or folder, they'd simply load up the previous day's backup tape, make a copy of the missing data, and continue from there. Those days are long gone. Today, the volume of data and diverse range of workloads have made backup and restore operations much more complex. Regardless of their size, industry, or location, every organization must have an active security perimeter to keep out the bad actors, plus effective recovery mechanisms to get back up and running quickly when an attack gets through. Although the IT world may be a dangerous place these days, careful planning and execution of the appropriate data security and data resilience processes can enable organizations to gracefully recover from otherwise dire situations. This IBM Redbooks publication provides guidance on one of IBM's solutions dedicated to these use cases, enabling customers to recover rapidly, at scale.

Figure 1-1 compares traditional recovery with cyber recovery.

Paradigm shift: Need for Cyber Recovery
 Cyber Recovery is fundamentally different from traditional recovery

Category	Traditional Recovery	Cyber Recovery	
Nature of impact	Random e.g. natural disasters	Targeted engineered for maximum impact	} Need Early Detection
Scope of impact	Local / Regional	Global can affect any connected systems	
Backup repository affected	Not typical	Possible	
Recovery point	Known	Unknown need most recent uninfected copy	} Need Safe Recovery
Mitigation objective	RPO/RTO	RPO/RTO + Safe Recovery	
Duration of impact	Hours to Days	Days to Weeks	
Relative probability of occurrence	Low	High	

Figure 1-1 Traditional recovery versus cyber recovery

All businesses can be subjected to cyberattacks. These attacks often target applications that are critical to a business. Data or applications can be encrypted, stolen or both. Historically, disaster recovery and business continuity efforts focused on environmental, software and hardware failures. Businesses design redundancy into systems and storage, while also using technologies such as backups and data replication to try to prevent the loss of data.

Many businesses are not prepared for, or are unaware of, the extent of the damage that a cyberattack, such as ransomware, 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 cyberattacks focus efforts only on prevention and not on how to recover quickly from an incident.

Data resilience is a measure of how well the data within applications and other infrastructure of a business can withstand events such as cyberattacks, human error, and natural disasters and still deliver business operations at a normal level.

Cyber resiliency is critical for business continuity. It helps reduce financial losses, downtime, and may reduce the damage to the business's reputation.

Important: A data-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.

1.1.1 Cybersecurity versus cyber resiliency

A Cybersecurity framework (CSF), such as that defined by the National Institute of Standards and Technology (NIST) is heavily focused on security and not as much on infrastructure, which is responsible for operational recovery. Recovery in the event of a malware incident such as a ransomware attack can be challenging, as key service data is increasingly fragmented across different data stores, both within and external to the organization. This complexity makes it even more difficult to identify the ransomware attack variant, identify the entry point of a corruption or encryption event, and then eradicate the risk of reinfection. Multiple data stores can also increase the complexity of recovering and testing data to ensure parity and synchronicity across data stores, which is required to ensure that the service is recovered in a safe and correct sequence. One of the biggest challenges with cyber threats is identifying which copy or recovery point to restore from. Without proper knowledge of the incident's onset, blast radius, and affected copies it becomes almost impossible to recover quickly, the need for a scanning solution becomes apparent.

Organizations that are able to focus on both cybersecurity and data resilience improve their ability to recover from corruption or encryption-based events and ransomware style attacks. The two principals working in concert provide not only the capability of mitigating the attempts to disrupt the business by bad actors, but also provide for the ability to quickly recover the data that supports the organization. Data resilience can be seen as the way for an organization to recover, test, restore environments and in the face of a ransomware, data corruption or encryption event.

Figure 1-2 shows the differences between cybersecurity and cyber resiliency.

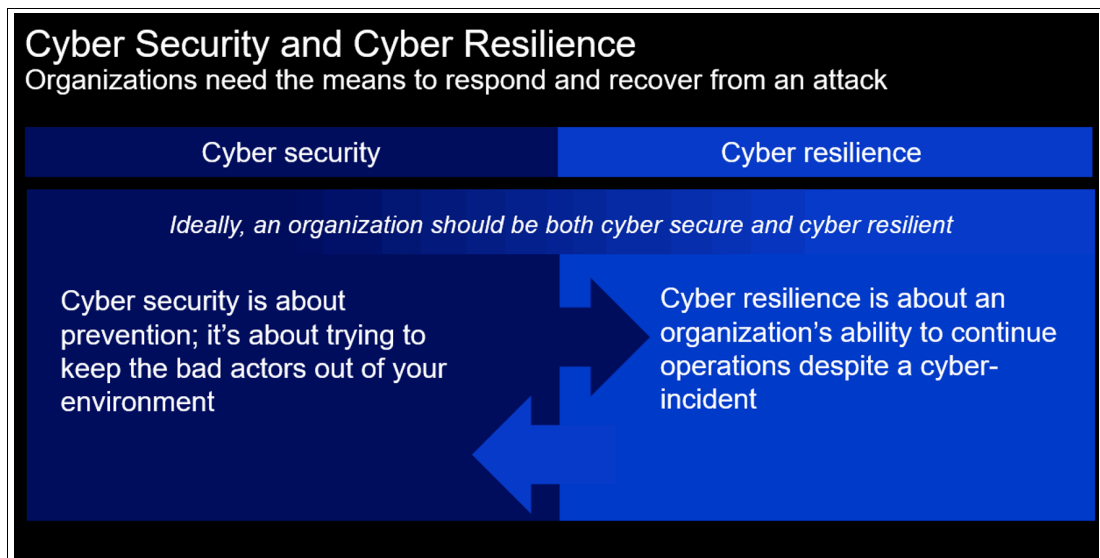


Figure 1-2 Cyber security and cyber resilience

Cybersecurity

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.

Cybersecurity includes IT intrusion detection and prevention, data loss or theft prevention, ransomware protection and protections for services that are running in the cloud. As organizations go mobile and employees work from mobile devices, mobile security is also a growing concern and must be included in a robust cybersecurity policy. Other concerns include minimizing potential damage from internal threat actors who already have privileged access to critical IT infrastructure.

Cybersecurity attacks often result in data theft, ransomware attacks or both. This data can include company secrets such as source code, or valuable customer and employee personal data.

Cyber resiliency

Where cybersecurity is focused solely on intrusion detection, response, and data loss prevention, cyber resiliency is a measure of an organization's systems to survive a disaster or cyberattack and still allow the organization to function. Cyber resilience includes cybersecurity as a component. Cyber resilient organizations will be more resilient than an organization that is not cyber resilient, but cyber resiliency is a measure of how well the organization functions during a cyber event and how quickly it can recover after an event (or a disaster event) occurs.

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. This is sometimes referred to as identifying the *Minimum Viable Company (MVC)*. The strategy must also include identifying the vulnerabilities and risks an 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 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 progresses to a test operational state. Where it is possible, resiliency is tested. Also, organizations must closely monitor 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 can be refined based on testing results.

After testing is complete, the organization moves to an operational state. In this phase, the design is deployed. Testing continues and uses controls to ensure that the operational state is effective and consistent

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

IBM Storage Defender and its IBM Storage Sentinel component can enhance an organization's cyber resilience strategy and processes.

1.2 Approaches to data resiliency

The ability to recover to a prior point in time relies on the availability of copies and backups - either point-in-time, array-based snapshots (such as copies made of Primary Workloads) or written to backup applications and their repositories such as disk, tape, virtual tape library (VTL) or cloud (sometimes called Secondary Workloads). These recovery options require the availability of a system to use for recovery, and require the data being restored to not be compromised in any way. The assumptions of system availability and uncorrupted data are often false when faced with recovering from a ransomware level event where the system or data is locked, corrupted, or encrypted.

Both the Primary and Secondary (backup) copies that are available for use in recovery scenarios must be free of contamination to remove the risk of a repeated attack and reinfection. Also, where the backup is written to immutable storage, the backup must be validated and verified as being free of infection before supporting recovery. In those instances, immutability can be a double-edged sword, making sure the immutable copies are clean is paramount. Restoration from backups is traditionally limited to a single system, single files, or relatively small volumes of data. Restoration from traditional backups is not designed to restore mass volumes of data to multiple systems in a short space of time.

Recovery from tape-based systems (or Virtual Tape Libraries) is limited by read time (sequential), retrieval times, and network bandwidth, restricting the ability to recover at scale and speed.

Even when the Disaster Recovery systems, and restored data, are available, confidence in running business services by using the secondary site is typically low, either because it is not tested sufficiently, or because of differences between the primary and secondary configurations, or their integration to other interfaced components. Highly available environments typically allow malware to spread between sites through the replication itself. Testing is often not representative of realistic failure scenarios and is typically based around site switching off single systems or single applications, which are quiesced and stopped gracefully first at the active site before being initialized at the recovery site, randomized and surprise testing are common best practices here.

Clients can use data from traditional Disaster Recovery (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 point-in-time
- ▶ Complexities associated with microservices
- ▶ Distributed systems and the synchronizing of data across system boundaries
- ▶ Interfaces with third parties
- ▶ Severing replication to quarantine malware
- ▶ Mass recovery scenarios
- ▶ Surprise recovery testing

Existing business impact assessments (BIA) that include potential loss scenarios, prioritization, and service dependencies can be used to help determine appropriate recovery strategies for business services.

IBM has several IBM funded assessments like the [Secure and Resilient Assessment](#) to help enterprises better understand their cyber resiliency posture. Reach out to your IBM Storage Sales representative or IBM Business Partner for more information.

1.2.1 Considering the restoration of static and dynamic data

Some data is transactional and volatile whereas other data is relatively static. However, both types of data are important and support the organization in different ways. For example, in the energy and utilities sector, trading applications generate highly active transactional data but seismic records are static. However, both data types are regulated and might easily constitute the Minimum Viable Company of the organization.

Both of those data types can be protected but in different ways. Figure 1-3 highlights that not all data needs as much protection because of factors, such as the frequency that the data is accessed. An organization can protect its data in various ways to support both its importance and its activity.

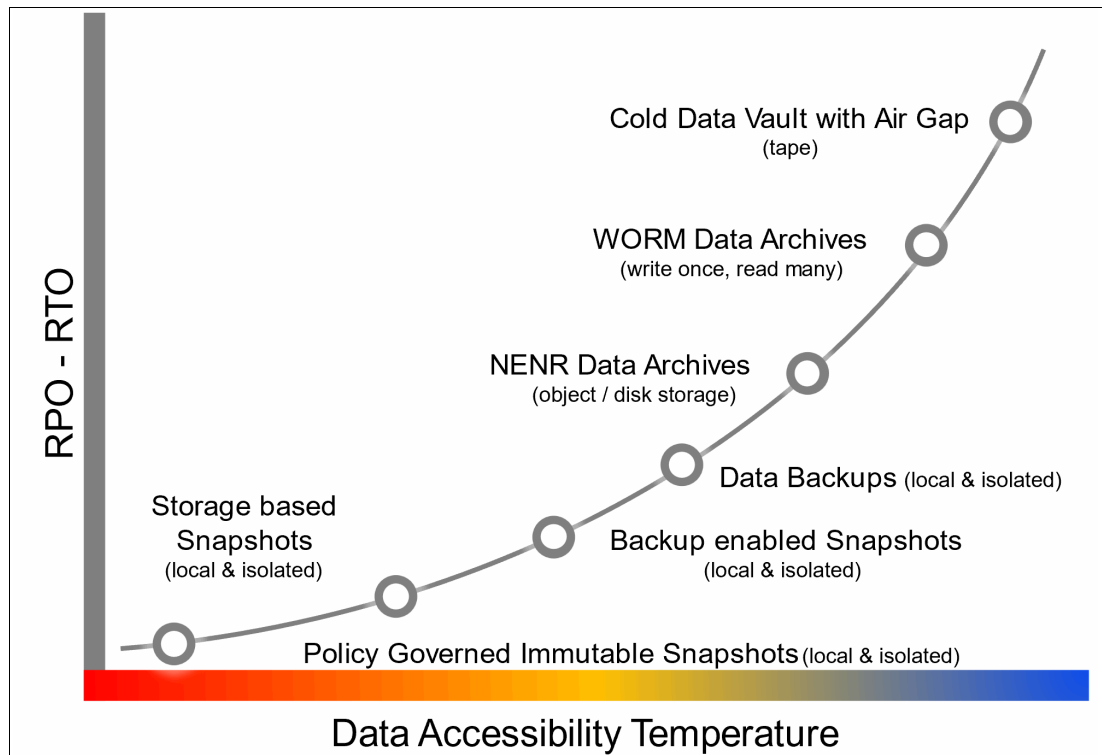


Figure 1-3 IBM Point of View: Levels of data resilience

Highly transactional workloads, such as those in critical applications, require minimal recovery time objectives (RTOs) and recovery point objectives (RPOs). To ensure rapid accessibility to cyber-resilient copies, technologies like policy-based immutable snapshots are essential. In contrast, static data, such as seismic data or long-term patient records, may have less stringent recovery requirements due to their nature and size. While still needing protection to comply with regulatory retention policies, these datasets can be efficiently managed using cold storage solutions, including tape-based media.

The most critical and highly transactional workloads generally need the lowest recovery time objective (RTO) and recovery point objective (RPO) and need the quickest accessibility to the cyber resilient copies by using something like policy-based immutable snapshots.

Whereas, static data like seismic data or long-term records, that might have regulatory retention requirements, such as hospital patient data, still need to be protected, but given their nature and size might need a different level of protection because they are not changing. This might be achieved by using something such as cold data storage, which might be using tape-based media.

1.2.2 Time to Recover

The ability of the organization to restore business services in a timely manner is key to success and in remaining within Impact Tolerance or in meeting the *Maximum Tolerable Period of Disruption (MTPD)*. This section identifies key considerations, which contribute to the elapsed duration of outage and recovery time.

The identification of the extent of the attack, including point of infection, nature of the attack, and extent of the damage is critical. All these factors affect the time that is needed to recover. Understanding these elements helps to determine the scope of recovery and what specific recovery actions are required. It helps to determine what services are impacted and what services can continue as they are.

The number and nature of infected systems directly affects the time needed to recover. Retrieval of data and restoration of data at a large scale can be limited by network bandwidth and the I/O capability of the source data repository, for example, the backup system. Whether recovery involves repairing a single file, data set or the complete restoration of a data volume, complete system, or server farm influences the time that is spent on recovery. The bandwidth available from the source repository, its media, and the network impacts how quickly data can be transported. Recovering large volumes of data from magnetic tape takes more time than recovering a volume from flash or disk media.

Isolated Recovery Environments (IREs) also known as clean rooms or isolated sandboxes help prevent reintroduction of contaminated data to the system after a ransomware attack. Before you restore data and restart services on the affected systems, verify that all traces of the ransomware are removed and that the restored / mounted copies are clean. Recovery of services on an alternate, isolated system can be done concurrently during recovery of the primary system or proactively in order to validate data before creating those copies, especially when those copies are to become immutable.

As traditional HADR solutions are not content aware, additional checks, validation, and technical solutions are required to restore service. Ideally, these validations and checks are made before any need for recovery so that the integrity and usability of the stored data that is used for recovery is already known. However, if the state of the stored data is not known, then validation of data might be required during recovery, which can increase the recovery time.

Where data is restored to an earlier point-in-time (RPO>0), the client might need to roll forward or replay transactions that occurred between the last know good backup of data and the time of the attack. The roll forward ensures that the system returns to a transactionally consistent state. Before you restart business services, verify the restored data is valid and free of malware. Reconciliation of data across multiple, interfaced systems and 3rd parties is also a key consideration and can add to the recovery time. Reconciliation of data might be required on applications that are hosted on the same platform (intra-system), on disparate systems (inter-system), and on 3rd-party systems and across the network in which the recovered system operates.

Technology solutions, such as the IBM Cyber Vault blueprint for the protection of mainframe and open systems data, provide an ability to regularly backup data to immutable storage on the primary storage systems at production and DR sites. The solutions provide validation of data in an air-gapped system.

Validated copies of data and an environment to restore, inspect, and enact recovery, significantly reduces the time to execute recovery and greatly improves the chance of a successful recovery then otherwise would be available. See Figure 1-4.

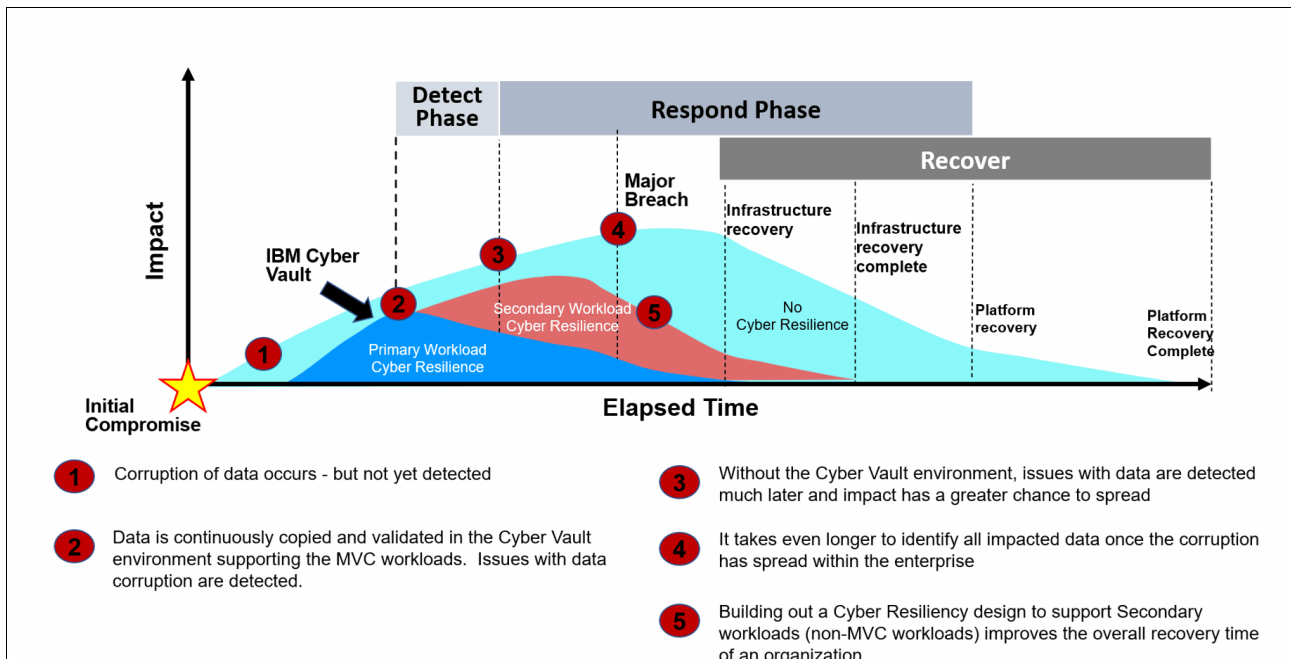


Figure 1-4 Recovery from ransomware attack timeline

You can find a full report on cyberattacks and their effects on companies at the [IBM Security X-Force® Threat Intelligence Index](#).

IBM has also released a study on the impact of data breaches at [Cost of a Data Breach Report](#).

Ransomware: Ransomware is an online attack that is perpetrated by cyber criminals or nation state-sponsored groups who demand a monetary ransom to release their hold on encrypted or stolen data.

A ransomware infection can be costly and disruptive if the only solution to return to normal business operations is to pay the cyber criminals' ransom. Statistics show, that only 50% of ransomware victims get back access to their data, even if the ransom was paid. One alarming trend is that cyber criminals now install malware and leave it dormant for 100 days or more before springing the trap. At that point, the malicious code has infected not only the target's production data systems and snapshots, but all of their backup copies, even if they use a *30 – 60 – 90 backup policy*. The victims have little choice but to pay.

Ransomware attacks can use several methods to infect a device or network. Some of the most prominent malware infection methods include:

- ▶ Phishing emails and other social engineering attacks: Phishing emails manipulate users into downloading and running a malicious attachment (that contains the ransomware that is disguised as a harmless looking .pdf, Microsoft Word document, or another file), or into visiting a malicious website that passes the ransomware through the user's web browser.
- ▶ Operating system and software vulnerabilities: Cyber criminals often exploit existing vulnerabilities to inject malicious code into a device or network.

IBM Storage Sentinel is a solution that is designed to help organizations detect ransomware and recover from cybersecurity incidents. It automatically creates immutable copies of data, then will use application-aware artificial intelligence and machine learning to detect possible corruption. It can generate forensic reports to help diagnose problems and find the source of an attack. Because it can isolate infected backups, you can identify which backup copies are verified and which are the most recent ones. This accelerates your time to recovery. IBM Storage Sentinel is available as a component of IBM Storage Defender.

1.2.3 Secondary workload cyber resiliency

For a more rapid recovery, an organization can prioritize its *Minimum Viable Company (MVC)* assets from the typically large collections of systems and data in an enterprise. Depending on the spread of the corruption or encryption, the organization might need to address existing applications. Including a data resiliency strategy for existing applications can reduce recovery time for the remaining workloads. Part of a secondary workload strategy can include the following steps:

- ▶ Further prioritization of workloads for testing, recovery, and validation procedures.
- ▶ Regular backup copies being moved to alternate immutable data platforms.
- ▶ Moving copies from primary systems to secondary environments such as off the array.
- ▶ Including a random sampling of workload for full recovery and validation.

By separating MVC recovery from secondary workload recovery, both can be recovered using different strategies. See Figure 1-5 on page 10. The separate strategy provides for the ability to use different testing and validation methods. Both types of recovery can include some automation. Where MVC workloads might have targeted testing and application-specific tools, the secondary workloads can take a more mass-scale approach using a more common set of generic tools. Proactively building an approach to support both primary and secondary recovery can mean accelerated recovery for both, which reduces overall organization disruption time.

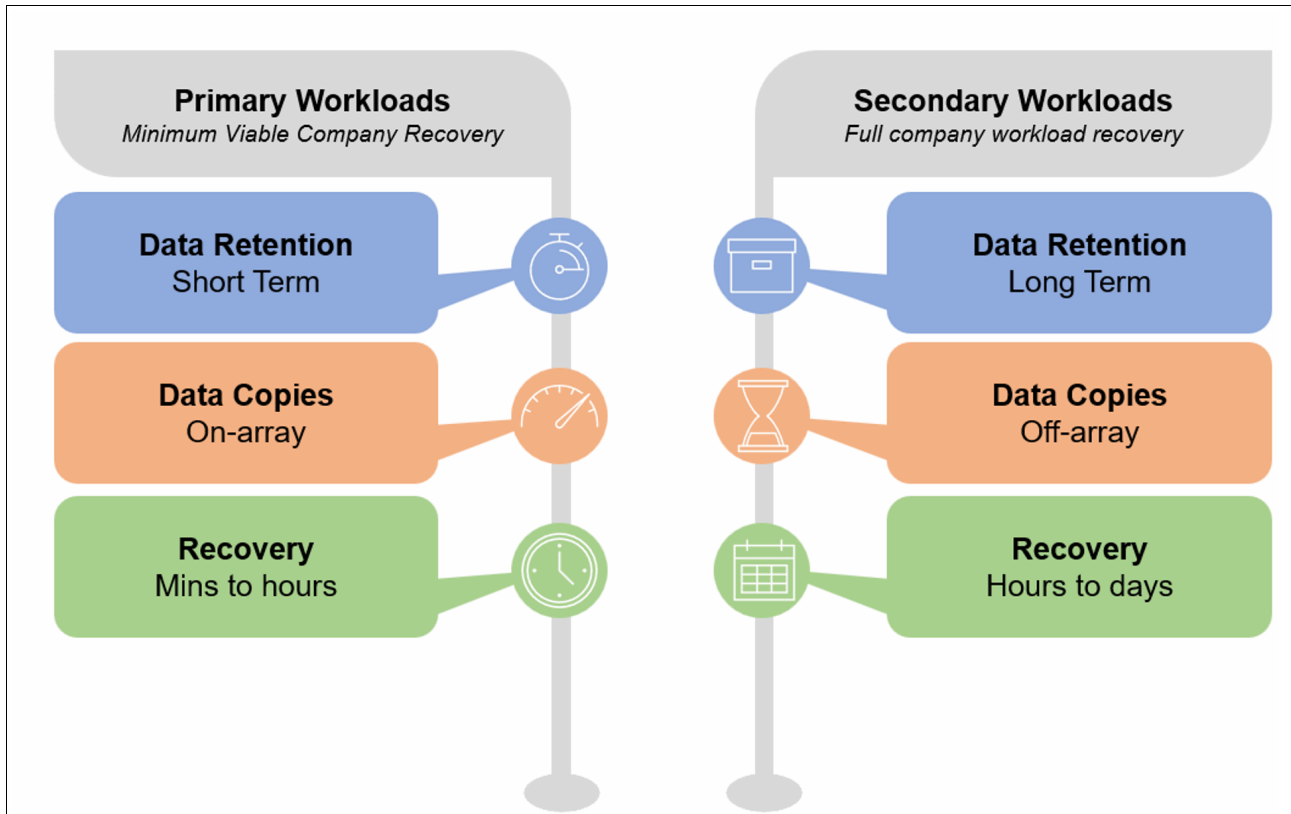


Figure 1-5 Segmentation of workloads

IBM Storage Sentinel is a solution that is designed to help organizations detect ransomware and recover from cybersecurity incidents. It automatically creates immutable copies of data, then will use application-aware artificial intelligence and machine learning to detect possible corruption. It can generate forensic reports to help diagnose problems and find the source of an attack. Because it can isolate infected backups, you can identify which backup copies are verified and which are the most recent ones. This accelerates your time to recovery. IBM Storage Sentinel is available as a component of [IBM Storage Defender](#).

1.2.4 IBM Storage Sentinel as part of IBM Storage Defender

In May of 2023 IBM introduced a new product called IBM Storage Defender. It enabled the ability to leverage the features through some of the existing cyber resiliency solutions like IBM Storage Sentinel and IBM Storage Protect (TSM) as well as some new ones like Defender Data Protect, provided through a flexible licensing structure.

IBM Storage Defender users are able to leverage the different components of the solution through a single standardized licensing metric giving them access to a wide variety of cyber capabilities spanning primary storage, secondary/backup storage, threat detection sensors, security integrations, and many others. IBM Storage Sentinel is one of IBM Storage Defender's components.

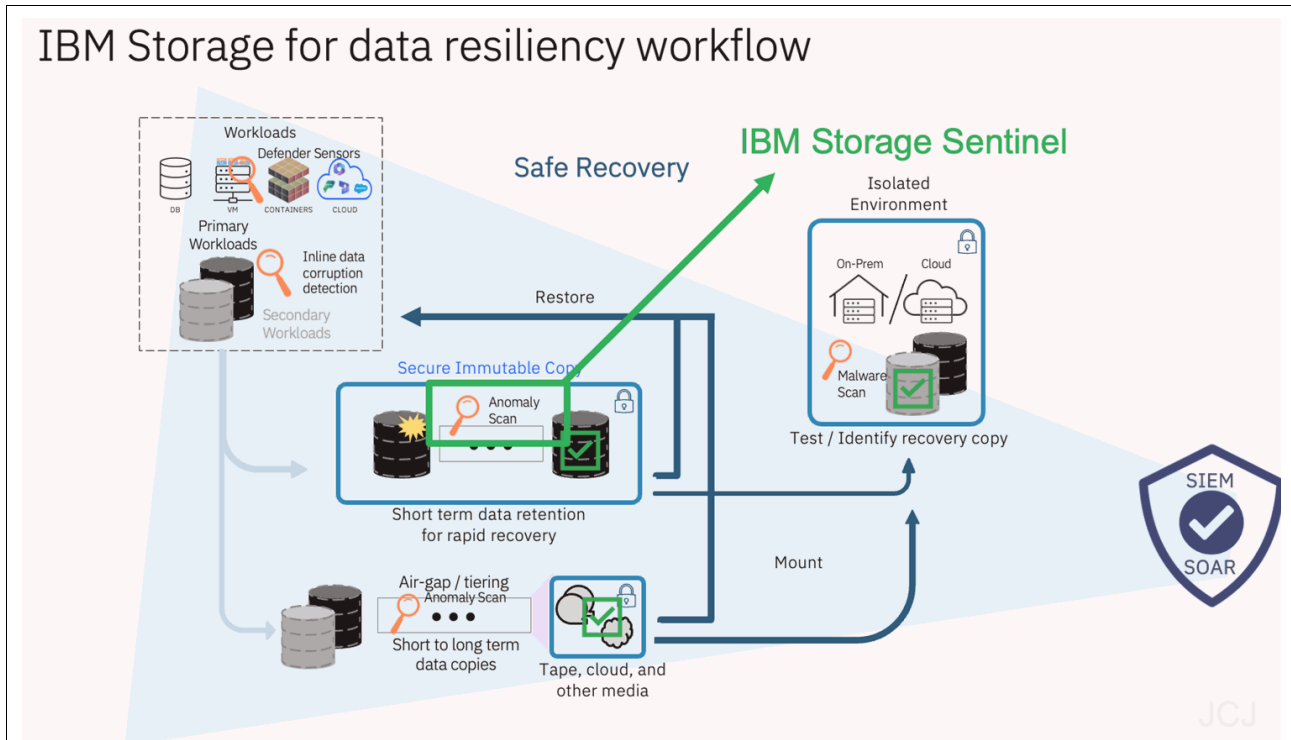


Figure 1-6 How Sentinel fits in IBM's data resiliency workflow and IBM Storage Defender

1.2.5 Identification

A key part of IBM Storage Defender is the ability to provide Early Threat Detection of workload corruption. IBM Storage Sentinel provides an important part of that capability by leveraging a scan engine to determine the state of the IBM FlashSystem Safeguarded Copies. That, however, is just one of the methods that IBM Storage Defender can contribute to storage based threat detection. See Figure 1-7.

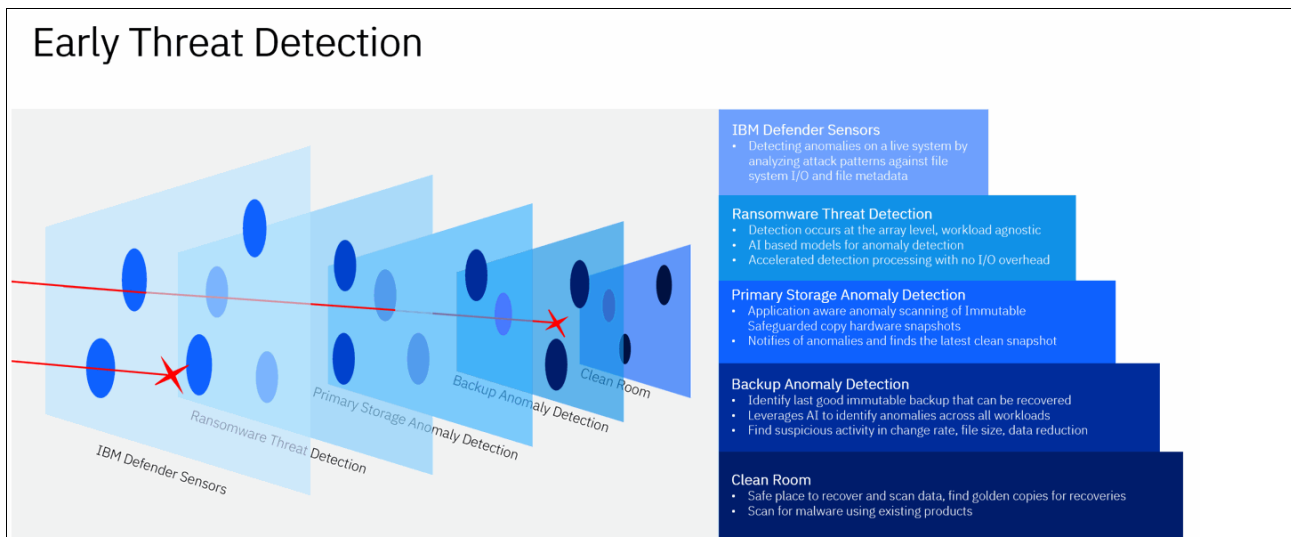


Figure 1-7 SIEM alerts triggered for downstream reporting and automation

There are a variety of threat detection levels and tools that are part of IBM Storage Defender as depicted in Figure 1-7 on page 11.

- ▶ IBM Defender Sensors, which install in the operating system of a virtual machine and are detecting anomalies on a live system by analyzing attack patterns against file system I/O and file metadata.
- ▶ Ransomware Threat Detection, which is part of the IBM Storage FlashSystem family that enables detection via the IBM FlashCore 4 (FCM4) modules.
- ▶ Primary Storage Anomaly Detection, which includes IBM Storage Sentinel which provides application aware anomaly scanning of Immutable Safeguarded copy hardware snapshots and notifies of anomalies and finds the latest clean snapshot.
- ▶ Backup Anomaly Detection, which provides the ability to identify last good immutable backup that can be recovered such as part of IBM Storage Defender Data Protect and can leverage AI to identify anomalies across all workloads. It allows users to find suspicious activity in change rate, file size, data reduction.
- ▶ Clean Room Detection, which provides a safe place to recover and scan data, find golden copies for recoveries as well as scanning for malware using existing products.

1.2.6 What is IBM Storage Sentinel and what are the use cases for Storage Sentinel

IBM Storage Sentinel can be used to detect ransomware in Safeguarded Copy snapshots, and help recover after an attack.

Threat actors will often wait weeks or even months after ransomware is deployed to ensure that it infects all of a business's systems. IBM Storage Sentinel can detect ransomware in snapshots from the primary storage system through a content aware scan of the copy. Leveraging Machine Learning training it is able to be both content aware of the specific applications it is scanning, and is also trained against how ransomware can destroy those type of workloads.

Storage Sentinel can help protect workloads, by identifying copies that are still valid and usable from copies that have been corrupted by malware or ransomware but that infection has yet to be detected by other methods.

Storage Sentinel helps recover after a ransomware attack. It can automatically generate reports listing the files or snapshots that were affected. This helps your organization identify clean copies of data that can be used to restore from.

1.2.7 IBM Storage Sentinel workflow

Figure 1-8 shows where IBM Storage Sentinel fits in an overall cyber resilience strategy. The *IBM Cyber Vault Blueprint* identifies the two main phases of cyber resiliency. The first phase is to protect your data. The second phase is recovering from a cyberattack. IBM Storage Sentinel spans both the protect and recover phases. It automates many of the tasks required for protecting data and can automatically identify safe recover points.

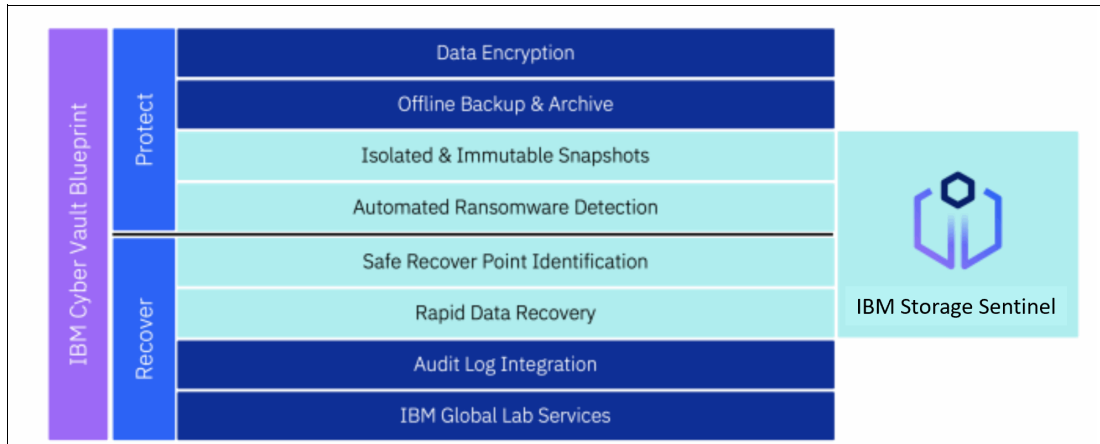


Figure 1-8 Storage Sentinel in Cyber Vault Blueprint

Figure 1-9 on page 13 shows the key points on the Storage Sentinel timeline for responding to an attack:

1. Before an attack begins, IBM Safeguarded Copy creates a series of immutable snapshots, which are proactively scanned for malware by IBM Storage Sentinel.
2. Ransomware begins infecting production files, databases, and systems with Safeguarded Copy, and those snapshots are protected and scanned proactively by Storage Sentinel.
3. Even as the attack is taking place, Storage Sentinel scans snapshots and analyzes changes, file extension mismatch, and other signs of data corruption.
4. Snapshots can now be assessed to verify the snapshots that are free of malware and data corruption. This is the integrity review phase.
5. After the integrity review is complete, IBM Storage Sentinel is the most viable snapshot to restore from is identified.
6. The most recent, uncorrupted snapshot is used to restore data to the production environment so the organization can resume normal business operations.

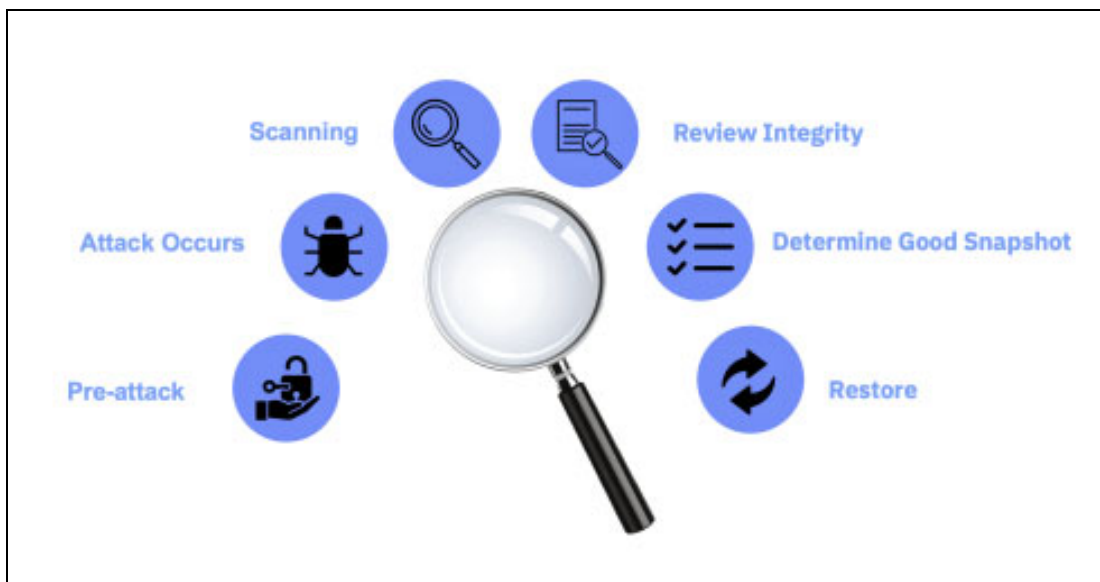


Figure 1-9 Storage Sentinel attack timeline

1.2.8 IBM Storage Sentinel components

IBM Storage Sentinel has the following components.

IBM Safeguarded Copy

IBM Safeguarded Copy is a feature of the DS8000®, IBM Storage FlashSystem, and IBM SVC storage systems that creates immutable snapshots of data to help protect against cyberattacks, malware, acts of disgruntled employees, and other data corruption.

Safeguarded Copy uses the FlashCopy feature available on IBM Storage FlashSystem and IBM SVC storage systems to create special immutable snapshots formerly called FlashCopies that cannot be accessed by hosts. They cannot be mounted by or attached directly to a host. Instead, if recovery from a Safeguarded copy is required, another snapshot is created and that copy is presented to the host.

IBM Storage Sentinel complements IBM Safeguarded Copy (but is not a requirement for IBM FlashSystems) by automatically scanning the copies that are created regularly by Safeguarded Copy and looking for signs of data corruption introduced by malware or ransomware.

IBM Storage Copy Data Management (CDM, also SCDM)

IBM Copy Data Management streamlines the management of Safeguarded Copy snapshots, automating copy processes and workflows to ensure consistency and reduce complexity. By cataloging existing storage, virtual machines, and applications, it provides a comprehensive view of the copy data environment. This powerful tool enhances IT efficiency, reduces costs, and empowers internal customers with self-service access to essential resources.

As a part of this workflow, IBM Storage Copy Data Management is able to create an application aware snapshot of the supported workloads for validation. This ensures that the application is mounted to the IBM Storage Sentinel scan software in an application consistent state each time.

Anomaly scanning and detection engine

Cyber protection solutions are designed to protect from an attack in real-time. However, these solutions are not 100% effective. Scanning data for anomalies adds additional protection to these solutions. It detects and locates corruption that occurs when a successful attack makes it into the data center. Early detection of issues enables IBM Copy Data Management software to start fast application recovery and alert the Security organization. This minimizes downtime and flattens the data resiliency curve.

The scanning software (sometimes called the scan engine) uses statistics about files on the host to identify corrupted files by using a machine learning model (MLM). The MLM is trained using real-world malicious codes. The software identifies malicious code attacks and checks the integrity of databases to detect corruption of the internal database. This corruption might occur because of an attacker; data corruption due to logical or physical causes; damage at the disk or volume level; as a flaw in the process to create a snapshot; or as a flaw in the process to back up the database.

The scanning software examines existing database pages and allocation tables, if they exist, to ensure that all the allocated database pages are present and located in their correct position. Sometimes a data signature is available or enabled by the database administrator, such as a checksum or CRC. The scanning software recalculates the signature based on the current page contents and verifies it against the value found in the page header. Other ancillary fields are also verified within each page, depending upon the database application. The MLM of the scanning software is designed to tolerate a small amount of database

corruption that is commonly observed in production database systems to avoid excessive false-positive alerts. Figure 1-10 shows the five steps to cyber resilience.




Figure 1-10 Five steps to cyber resilience

1.2.9 Supported applications

IBM Storage Sentinel support includes the several applications.

SAP HANA

SAP HANA is a database that stores data in system memory instead of on disk. This enables processing data at speeds that are magnitudes faster than disk-based systems. This allows for advanced, real-time analytics to be performed on the data. SAP HANA can be used on premises, in the cloud or in a hybrid cloud and deployed in both locations. SAP HANA can apply machine learning and AI to data from multiple areas of a business. For example, it can integrate data from several sources:

- ▶ Traditional documents - spreadsheets, contracts, and so forth
- ▶ Emails, website forms and other user experience documents
- ▶ Internet of Things (IoT) - such as data from sensors in warehouses or trucks, security sensors, RFID tags and the many types of sensors in all aspects of a business
- ▶ Mobile - data from the mobile devices of customers and employees

SAP HANA can integrate and analyze the vast amounts of data that sits in data warehouses and provide no value unless it is analyzed to provide more customer value and increase business impact.

SAP HANA can interact with other backup and restore functions.

Backup

The following types of backups are available for SAP HANA:

- ▶ IBM FlashCopy® NoCopy
- ▶ FlashCopy Incremental
- ▶ Global Mirror with Change volumes
- ▶ Safeguarded Copy

There is also an option to back up the log file.

Restore

Copy Data Services Manager creates a temporary volume from a backup, then mounts it to the original server for recovery.

For more information, see Chapter 4, “Protecting SAP HANA databases” on page 41.

Epic

EPIC is electronic healthcare record (EHR) software that covers all functions of healthcare operations. This includes patient records, patient engagement, billing, mobile, clinical data from medical tests, interoperability, specialist care, and even government regulations. EPIC uses two database technologies:

- ▶ An operational database that handles online transactions. This database runs Cache’ from Intersystems Corp.
- ▶ An analytical database that can run on either Microsoft SQL Server or Oracle.

Note: IBM Storage Sentinel currently does not support Microsoft SQL Server.

Refer to Chapter 4, “Protecting SAP HANA databases” on page 41 for more information.

Oracle

In June 2023 (with IBM Storage Sentinel Version V1.1.4), IBM announced support for Oracle DB running on both Linux and IBM AIX®.

Further information, see [What’s new in IBM Storage Sentinel anomaly scan software 1.1.4.](#)

An example of an Oracle Backup job is shown in Figure 1-11 on page 16.

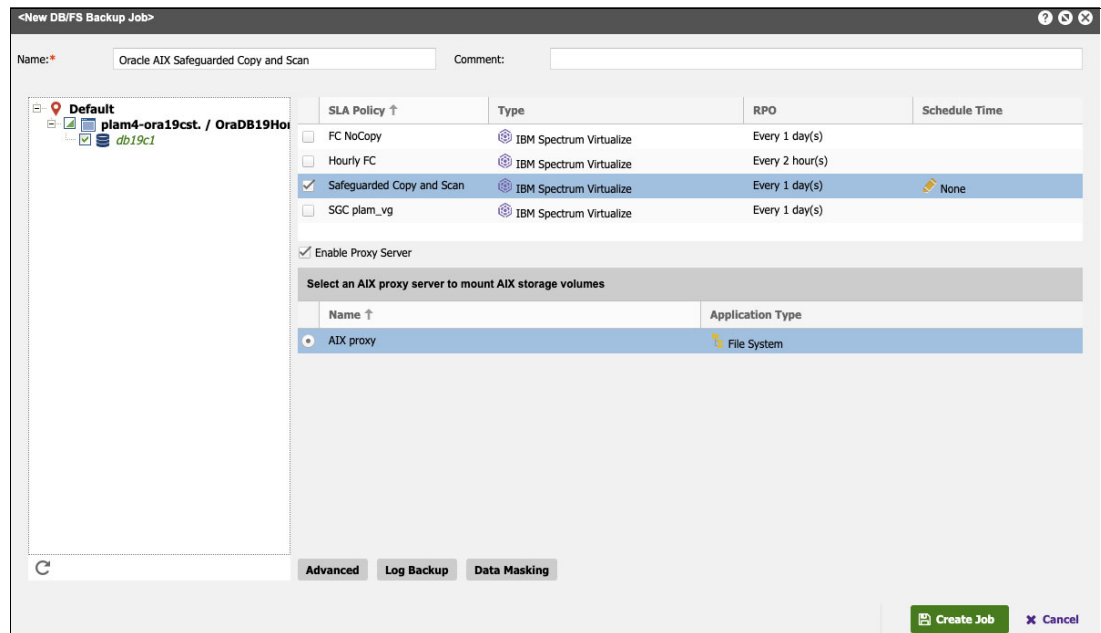


Figure 1-11 An example Oracle Backup job

The joblog shown in Figure 1-12 shows an anomaly detected by IBM Storage Sentinel.

Type	Time ↑	Task...	Message
			4ee56f638608
i	Apr 12 03:01:56 2023	2	Guest tools on 9.11.43.19 already at latest version: 2.15.4
i	Apr 12 03:01:57 2023	2	[9.11.43.19] Unix Agent 2.12.0.5 running as cdadmin for nfsmount (task ID: 678da049-85a0-4795-84c6-4ee56f638608)
i	Apr 12 03:01:57 2023	2	[9.11.43.19] Hostname: index-sle15-9 / Operating System: x86_64
i	Apr 12 03:01:57 2023	2	[9.11.43.19] Collecting list of NFS shares exported from the NFS Server: 9.11.60.88
i	Apr 12 03:01:57 2023	2	[9.11.43.19] Mounting NFS share: /tmp/mounts/10_11_59_121/1107/9_11_62_111/ec5b03cc629cccd6b800d89ec8f
i	Apr 12 03:01:57 2023	2	[9.11.43.19] NFS share mounted successfully
i	Apr 12 03:01:57 2023	2	[9.11.43.19] Completed mount operation in 0s, 1 NFS share(s) mounted successfully and 0 NFS share(s) failed
i	Apr 12 03:01:59 2023	2	Security Scanning of protected databases
i	Apr 12 03:01:59 2023	2	Starting Index job on mount path /tmp/mounts/10_11_59_121/1107 with job name 1107...
i	Apr 12 03:02:05 2023	2	Index job (121) created.
i	Apr 12 03:02:06 2023	2	Index job (121) started.
w	Apr 12 03:09:01 2023	2	Security Scan finished with state: Done. Previous threat detected: false. Number of new threats detected: 1.
i	Apr 12 03:09:01 2023	2	Unmounting database snapshot copies after Security Scanning
i	Apr 12 03:09:01 2023	2	ECX log dir=/data/log/ecxdeployer/2023-04-12/656b7731-c5e6-4608-81c8-7a4654bbbefd
i	Apr 12 03:09:04 2023	2	Guest tools on 9.11.43.19 already at latest version: 2.15.4
i	Apr 12 03:09:05 2023	2	[9.11.43.19] Unix Agent 2.12.0.5 running as cdadmin for cleanup (task ID: 656b7731-c5e6-4608-81c8-7a4654bbbefd)
i	Apr 12 03:09:05 2023	2	[9.11.43.19] Hostname: index-sle15-9 / Operating System: x86_64

Page 3 of 4 Download All Displaying 101 - 150 of 176

Figure 1-12 An example of a Backup joblog containing a detected anomaly

Figure 1-13 shows an example of a Backup joblog with no detected anomaly.

ID	Type	Duration	Status	Message	Type	Time ↑	Task...	Message
1	Resolve	0h 0m 0s	CO...	COMPLETED	i	May 2 06:50:20 2023	2	Guest tools on 9.11.43.19 already at latest version: 2.15.4
2	Protection (Oracle)	0h 13m 33s	CO...	COMPLETED	i	May 2 06:50:21 2023	2	[9.11.43.19] Unix Agent 2.12.0.5 running as cdadmin for nfsmount (task ID: ebf77315-27c5-40ad-a639-f8489daa7949)
	Finding databases to protect:	Done (Total:1)			i	May 2 06:50:21 2023	2	[9.11.43.19] Hostname: index-sle15-9 / Operating System: x86_64
	Finding data and log disks of databases :	Done (Total:1)			i	May 2 06:50:21 2023	2	[9.11.43.19] Collecting list of NFS shares exported from the NFS Server: 9.11.60.88
	Resolving database disks on IBM SVC storage:	Done (Total:1)			i	May 2 06:50:21 2023	2	[9.11.43.19] Mounting NFS share: /tmp/mounts/10_11_59_121/1107/9_11_62_111/ec5b03cc629cccd6b800d89ec8f
	Performing pre snapshot operations:	Done (Total:1)			i	May 2 06:50:21 2023	2	[9.11.43.19] NFS share mounted successfully
	Creating safeguard copies of volumes:	Done (Total:1)			i	May 2 06:50:21 2023	2	[9.11.43.19] Completed mount operation in 0s, 1 NFS share(s) mounted successfully and 0 NFS share(s) failed
	Performing post snapshot operations:	Done (Total:1)			i	May 2 06:50:21 2023	2	Security Scanning of protected databases
	Total databases protected:	1			i	May 2 06:50:23 2023	2	Starting Index job on mount path /tmp/mounts/10_11_59_121/1107 with job name 1107...
	Total databases not protected:	0			i	May 2 06:50:23 2023	2	Index job (131) created.
	Load storage data:	Done (Total:1)			i	May 2 06:50:23 2023	2	Index job (131) started.
	Load host data:	Done (Total:1)			i	May 2 06:50:28 2023	2	Security Scan finished with state: Done. No threats detected.
	Mount snapshot copies:	Done (Total:1)			i	May 2 06:50:29 2023	2	Unmounting database snapshot copies after Security Scanning
	Map LUNs:	Done (Total:1)			i	May 2 06:58:10 2023	2	ECX log dir=/data/log/ecxdeployer/2023-05-02/e1feb6ce-d556-49d2-9705-0ed216eef4d
	Security Scanning of protected databases:	Done			i	May 2 06:58:10 2023	2	Guest tools on 9.11.43.19 already at latest version: 2.15.4
	Dismount snapshot copies:	Done (Total:1)			i	May 2 06:58:13 2023	2	[9.11.43.19] Unix Agent 2.12.0.5 running as cdadmin for cleanup (task ID: e1feb6ce-d556-49d2-9705-0ed216eef4d)
	Cataloging objects:	Done (Total:17)			i	May 2 06:58:15 2023	2	[9.11.43.19] Hostname: index-sle15-9 / Operating System: x86_64
	Condensing catalog:	Done			i	May 2 06:58:15 2023	2	[9.11.43.19] Cleaning up mounted volumes

Page 1 of 1 Displaying 1 - 2 of 2 Page 3 of 4 Download All Displaying 101 - 150 of 176

Figure 1-13 An example of a Backup joblog with no detected anomaly

VMware

IBM announced support for VMware in 2Q 2024 (with IBM Storage Sentinel Version V1.1.9 and CDM version 2.2.24.0).

See [What's new in IBM Storage Sentinel anomaly scan software 1.1.9](#).

Users are now able to protect critical VMs with Safeguarded Copy immutable snapshots and scan those copies with IBM Storage Sentinel. See Figure 1-14.

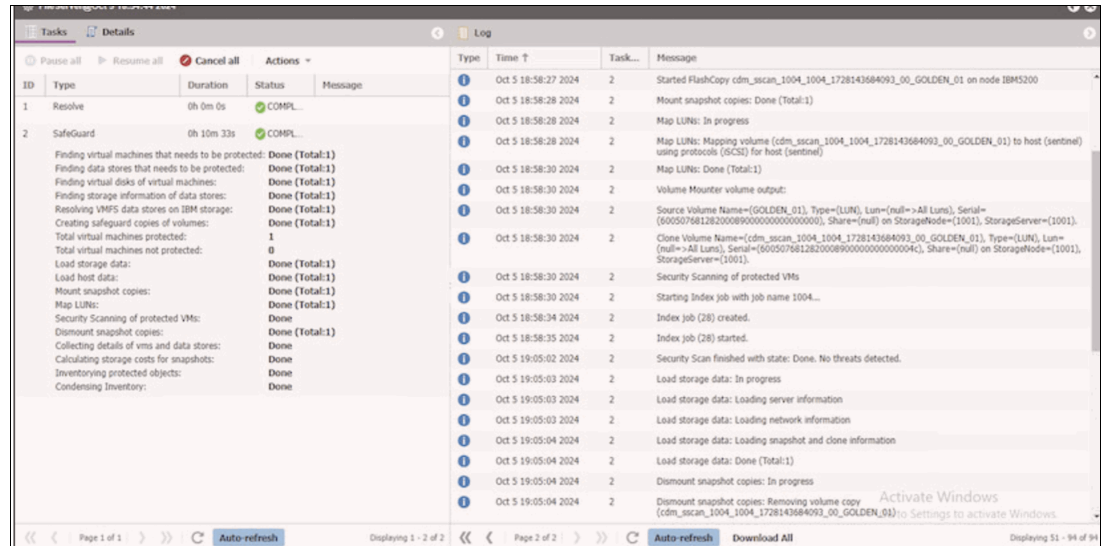


Figure 1-14 VM scanning job details



IBM Storage Sentinel: An end-to-end automated cyber resiliency solution

An end-to-end automated cyber resiliency solution requires hardware, software for the orchestration and software for handling events. IBM Storage Sentinel is a powerful orchestration tool that streamlines the implementation of end-to-end automated cyber resiliency solutions. This chapter discusses components of this solution.

This chapter has the following sections:

- ▶ “An end-to-end automated cyber resiliency solution” on page 20
- ▶ “IBM Storage FlashSystem” on page 20
- ▶ “IBM Storage Safeguarded Copy” on page 21
- ▶ “IBM Storage Copy Data Management” on page 22
- ▶ “IBM Storage Sentinel: Anomaly scan software” on page 24
- ▶ “IBM Security QRadar” on page 25
- ▶ “IBM Security Guardium” on page 27

2.1 An end-to-end automated cyber resiliency solution

By integrating seamlessly with IBM FlashSystems and IBM Copy Data Management, Sentinel offers a robust framework for:

- ▶ **Data protection:** Leveraging IBM Copy Data Management, IBM Storage Sentinel ensures data availability and business continuity through automated data replication and recovery processes.
- ▶ **Data security:** IBM Storage Sentinel's built-in data scanning capabilities proactively identify and mitigate potential security threats, safeguarding your critical data.
- ▶ **Efficient orchestration:** IBM Storage Sentinel automates complex workflows, reducing manual intervention and minimizing the risk of human error.

Figure 2-1 shows this end-to-end automated cyber resiliency solution.

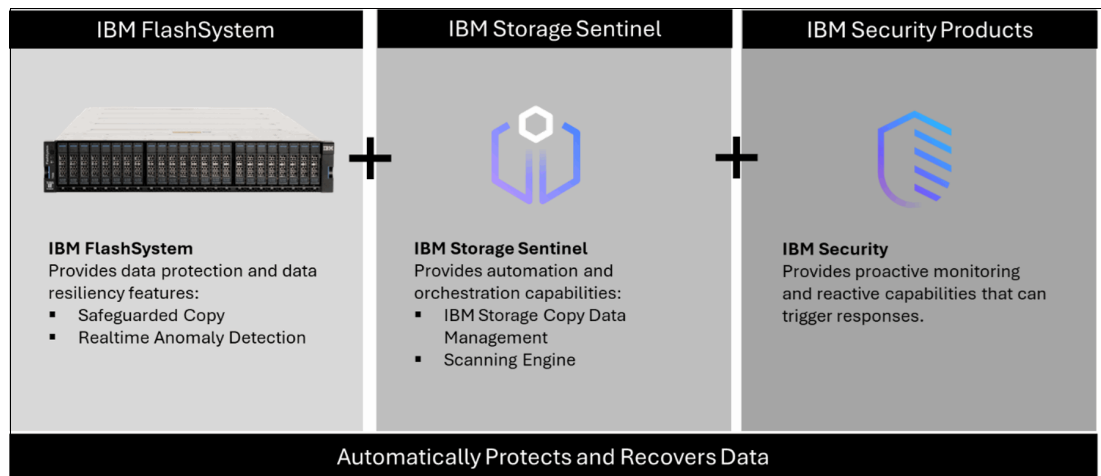


Figure 2-1 End-to-end automated cyber resiliency solution

This section takes a look at each major component of the IBM Storage Sentinel solution describing the high-level values of each and contributions towards an overall comprehensive cyber resiliency solution.

2.2 IBM Storage FlashSystem

IBM Storage FlashSystem® is an enterprise-class multipurpose flash-based storage family of products with integrated data protection, cyber resilience, sustainability and efficiency features.

IBM FlashSystem by design has a number of capabilities focused on data protection, resiliency and availability and provides a foundational capability to the IBM Storage Sentinel solution. The Safeguarded Copy feature is provided by IBM Storage Virtualize, the underlying software stack used not only within IBM FlashSystem products but also with the IBM SAN Volume Controller (SVC) and IBM Storage Virtualize for Public Cloud.

Note: Safeguarded Copy is the name of the function that many readers will know by the term immutable snapshots. Starting with IBM Storage Virtualize V8.7, Safeguarded Copy is also called Safeguarded Snapshots.

Additional features that are part of FlashSystem can enhance data protection, high-availability and resiliency including the IBM FlashSystem fourth generation of FlashCore Modules (FCM4) which are actually computation storage devices and provide [Realtime Ransomware Detection capabilities](#).

Of the two functions Safeguarded Copy is the function that for the purpose of this book is important, so next we will describe it in more detail.

For more information on IBM FlashSystem, refer to the IBM Redbooks *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.

2.2.1 IBM Storage Safeguarded Copy

Safeguarded Copy is available on all current models with the exception of the FlashSystem 5015 model. The feature creates immutable point-in-time copies (snapshots) of a volume's data that cannot be seen nor mounted by any host and are completely inaccessible to users that do not have the correct security privileges - which can include the traditional Storage Administrator managing the IBM FlashSystem.

Safeguarded Copy is designed to protect data against intentional or accidental corruption, modification, or deletion and provides a logical airgap as part of an overall data resiliency strategy.

Figure 2-2 shows IBM FlashSystem with Safeguarded Copy.

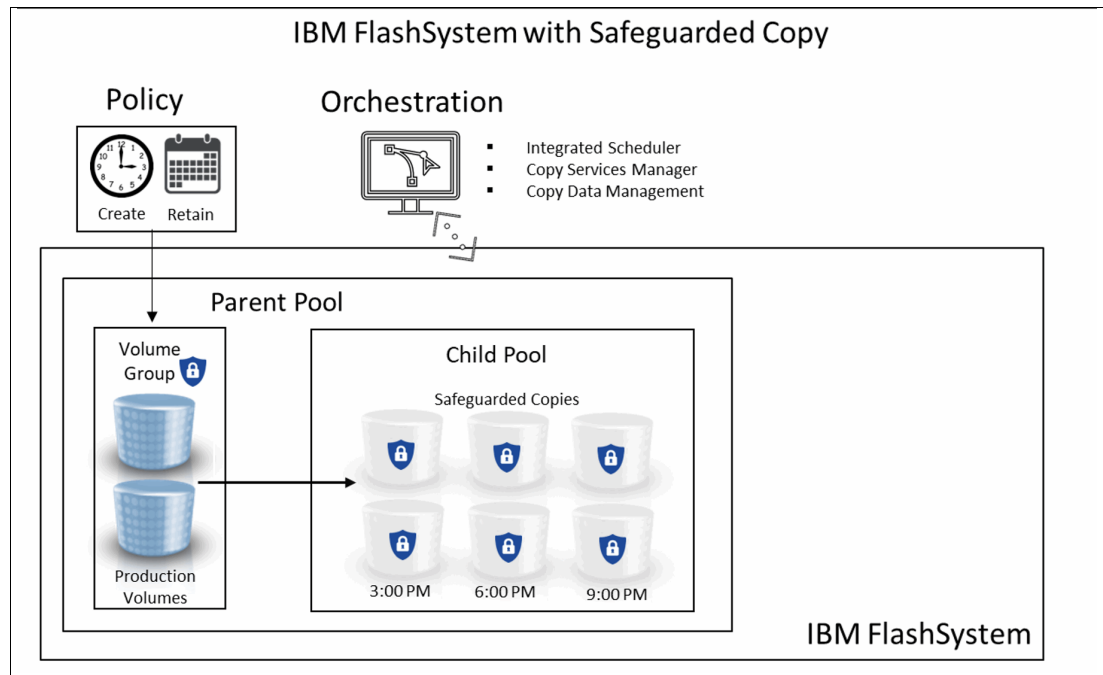


Figure 2-2 IBM FlashSystem with Safeguarded Copy

At a foundational level, snapshots are managed by policies at a volume group level. Such snapshots can be regular snapshots, or a flag can be set, which makes every volume group snapshot created with this policy as an immutable snapshot. Additional policy settings determine the frequency and the retention.

IBM FlashSystem comes with several pre-defined Safeguarded Copy policies. Administrators can also define custom policies and use them instead.

Safeguarded Copy policies can be orchestrated using one of three options, using the integrated scheduler introduced with Safeguarded Copy 2.0, using IBM Storage Copy Services Manager, or using IBM Storage Copy Data Management.

The advantage of Safeguarded Copies are they can be used to quickly recover data back to a production environment because the immutable data resides on the same physical storage array as the source data. Data can be restored in as few as seconds, minutes or hours - depending on how much data is being restored and the restore method.

The speed of the restore is what differentiates Safeguarded Copies or normal snapshots from more traditional backup. If you have to restore a single volume, this might not be a big deal, but when lots of volumes need to be restored this is a huge advantage, especially in such cases where a backup environment was never designed to cope with lots of simultaneous restores.

While this approach provides a significant layer of protection, it is crucial to understand that it is not a substitute for a comprehensive backup solution. While Safeguarded Copies are logically protected from deletion, they still reside on the same physical system and storage array. Additionally, due to the copy-on-write mechanism, only modified data is actually copied.

This means that a catastrophic event, such as a physical hardware failure or a fire, could potentially compromise both the original data and the Safeguarded Copies. Therefore, it is essential to maintain a robust backup strategy that includes off-site storage to ensure data durability and resilience against various threats.

For more information on Safeguarded Copies, refer to the IBM Redpaper *Data Resiliency Designs: A Deep Dive into IBM Storage Safeguarded Snapshots*, REDP-5737.

2.2.2 IBM Storage Copy Data Management

Until recently, data had been copied, and those copies have been residing in different physical locations and maybe even having copies stored on different types of physical media. Organizations ended up having entire disaster recovery sites full of copies of data, waiting for something to happen. The approach is wasting a lot of potential in terms of what else you can do with copies of your data for other purposes.

There has been a shift in IT with the introduction of the concept of Copy Data Management (CDM) which supports a wide range of IT-related functions all dependent on data. CDM automates the creation of data copies, facilitates the reuse of those copies and helps organizations do something useful with those copies allowing them to move beyond the traditional role of just supporting data protection, but actually using the same data for multiple purposes.

The premise of CDM is to help drive efficiency and lower overall datacenter costs by allowing data protection to be just one of the many uses of data copies. Other use cases could include DevOps, Test/Dev, analytics, operation automation, and so forth.

IBM's technology supporting the role of Copy Data Management is IBM Storage Copy Data Management (SCDM), which helps bring modernization to existing IT environments and can do so non-disruptively. SCDM can help organizations of all sizes with their modernization initiatives.

As shown in Figure 2-3, IBM Storage Copy Data Management can automate a variety of data copy tasks.



Figure 2-3 IBM Storage Copy Data Management

IBM Storage Copy Data Management delivers "in-place" copy data management to enterprise storage arrays from IBM, NetApp, Dell EMC, and Pure Storage and allows IT to easily use its existing infrastructure and data in a manner that is efficient, automated, and scalable.

While you can create multiple copies of your data at different times and retentions, you might face a new problem: in the event to get hit with a ransomware attack, you might be able to find out when it started, but can you be sure if the last copy you created before the attack does not already include the malicious code, and after the restore everything starts over again?

IBM Storage Copy Data Management, as part of the IBM Storage Sentinel solution acts as an orchestrator adding automation capabilities to not only the process of protecting data, but the process of testing the integrity of data and determining and cataloging good copies of data as well as streamlining the recovery of data when needed.

Figure 2-4 on page 24 shows the concept of the solution:

1. Creation of the Safeguarded copies.
2. Automatic creating of a thin clone, which gets mounted to a scan server.
3. The scan server checks for anomalies.
4. The Safeguarded copies are marked with the scan results in the recovery catalog within the CDM server.
5. In case you need to restore from a snapshot, you know can easily identify if the copy has been scanned and what the outcome was, instead of just hoping the restore will fix everything.

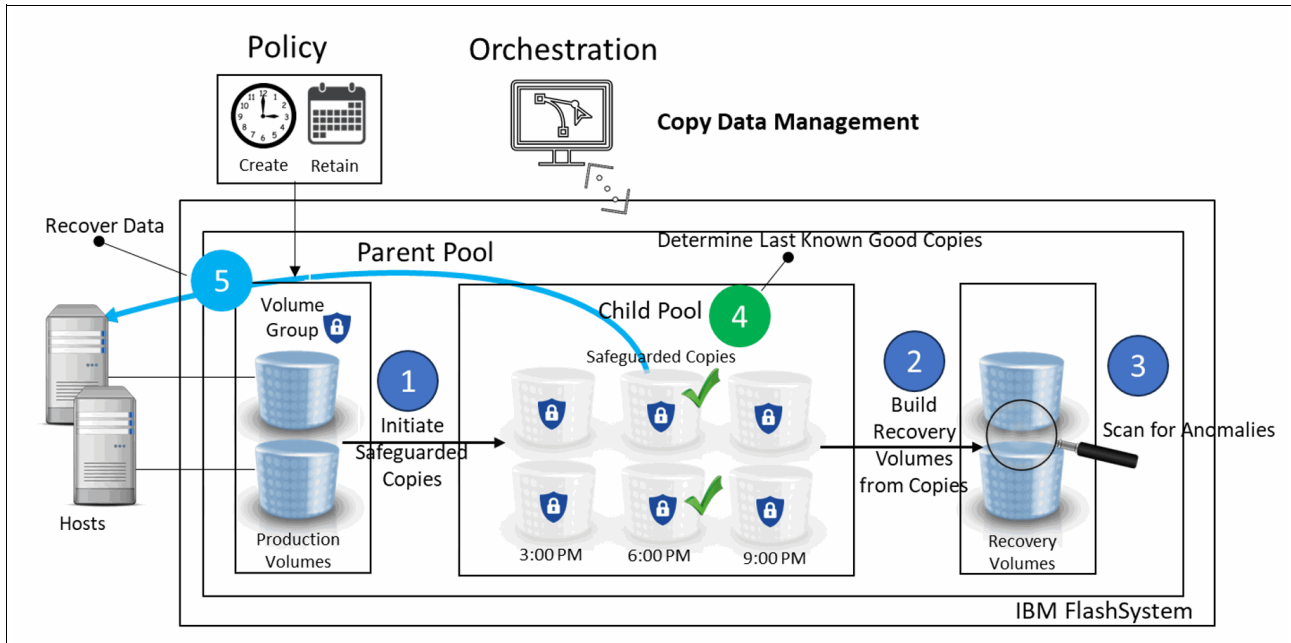


Figure 2-4 IBM Storage Copy Data Management orchestrating and automating

2.2.3 IBM Storage Sentinel: Anomaly scan software

The goal of cybersecurity solutions is to detect, protect, and defend against cyberattacks; however, these solutions are not 100% effective and data can become impacted by way of malware and ransomware. Having the ability to scan data for anomalies adds an additional layer of data protection.

The IBM Storage Sentinel solution uses anomaly scan software to scan data recovered from Safeguarded Copies. The automated scanning process orchestrated by IBM Storage Copy Data Management in conjunction with the anomaly scan software, looks for irregularities to the data based on a continuously updated list of recognized data anomalies.

Figure 2-5 on page 25 shows some of the validations performed by the anomaly scan software.

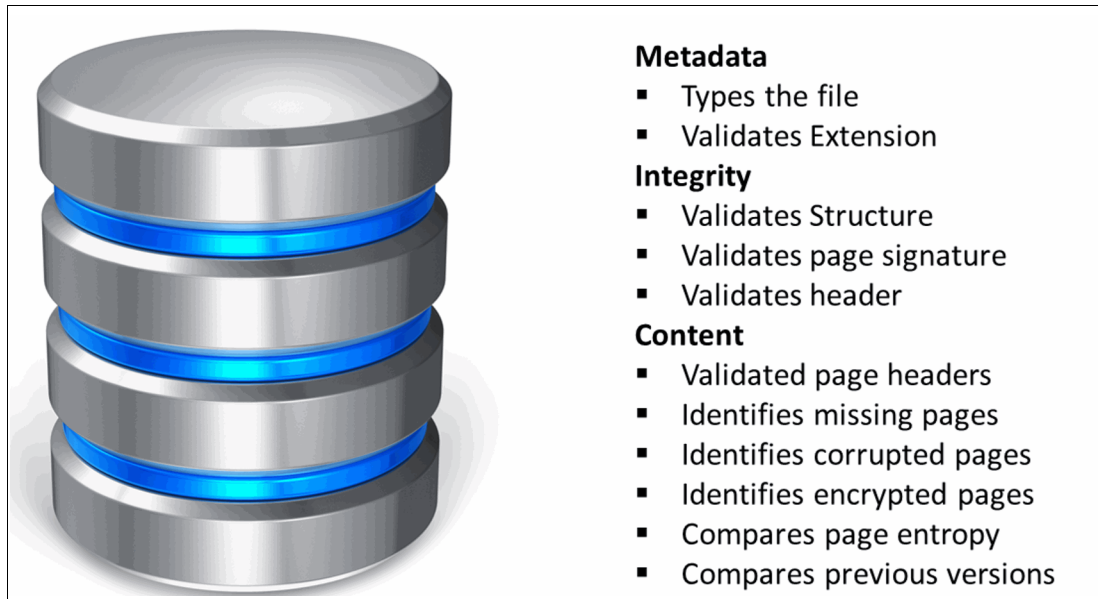


Figure 2-5 Some of the validations performed by the anomaly scan software

The scanning process not only detects malicious code attacks but also verifies the integrity of databases, identifying both intentional and unintentional data corruption.

IBM Storage Sentinel employs AI-powered anomaly detection to proactively identify unusual data patterns. By swiftly recognizing these anomalies, IBM Storage Copy Data Management can promptly initiate and orchestrate a streamlined recovery process for affected data. This proactive approach enhances business resilience and minimizes downtime.

The scanning software uses statistics about files on the host to identify corrupted files by using a machine learning model (MLM). The MLM is trained using real-world malicious codes.

2.2.4 IBM Security QRadar

Detecting security threats as early as possible can help prevent or lessen the impacts on data and speed the recovery process. One method of detecting security threats is to use a Security Information and Event Management (SIEM) solution such as IBM QRadar SIEM. A SIEM specializes in generating alerts in real time by correlating analytics of logs and events, evaluating threat intelligence, monitoring for changes to network and user behavior and helps security analysts stay focused on investigating and remediating threats,

IBM QRadar helps protect data by providing continuous monitoring of activity, enforcing compliance and supporting a zero-trust approach to data management across entire IT environments and can play a part of an overall IBM Storage Sentinel solution by integrating with IBM FlashSystem and or IBM Storage Copy Data Management.

IBM QRadar SIEM can monitor various activities by analyzing logs in real-time looking for signs of unusual behavior which could indicate intruders and or a cyberattack. Some examples could be looking for login attempts by an administrator whose credentials have been revoked, or an administrator user logging in from unusual IP addresses outside business hours.

The integration of IBM Security QRadar with IBM Storage provides a robust approach to detecting and mitigating potential security threats in real-time. Here are some key points regarding this integration:

- ▶ **Proactive threat detection:** QRadar enhances the ability to detect suspicious activities by analyzing logs and network events, helping to identify anomalies, such as unauthorized access attempts outside business hours or from unusual IP addresses.
- ▶ **Automated responses:** When suspicious behavior is detected, QRadar can automatically trigger the Safeguarded Copy feature, which creates secure backups of data, minimizing the risk of data loss or corruption before any damage occurs.
- ▶ **Separation of duties:** Safeguarded Copy ensures that administrative roles are segregated, enhancing security by preventing cyber attackers from easily manipulating backup data. QRadar monitors these activities, providing alerts for any unusual administrative actions.
- ▶ **Real-time analytics:** Utilizing AI and analytical capabilities, QRadar processes vast amounts of data quickly, turning raw logs into prioritized alerts that security teams can address promptly, reducing response time to potential threats.
- ▶ **Comprehensive integration:** With over 500 pre-packaged integrations, QRadar can work seamlessly with various storage hardware vendors, ensuring that organizations can implement robust security measures irrespective of their specific storage solutions.

This combination of threat detection and proactive data protection enables organizations to bolster their cybersecurity posture, helping to mitigate risks associated with cyber threats effectively.

Figure 2-6 shows IBM FlashSystem with Safeguarded Copy Triggered by IBM QRadar SEIM.

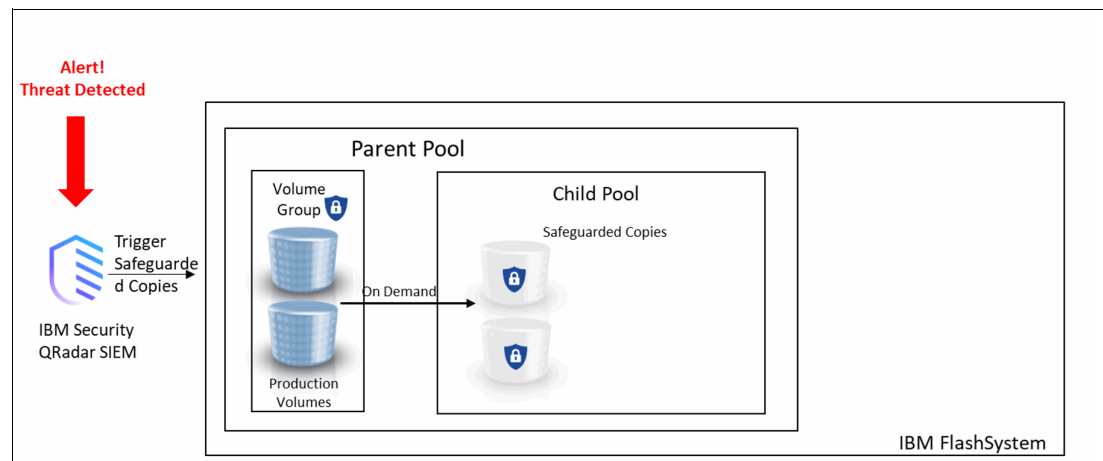


Figure 2-6 IBM FlashSystem with Safeguarded Copy Triggered by IBM QRadar SEIM

When alerted to a suspected threat to data, IBM QRadar SIEM can proactively trigger an action such as initiating the Safeguarded Copy function on an IBM FlashSystem - creating immutable backups of an organization's most critical data as a precaution.

An expanded example, if a security breach is detected that may have compromised data on an IBM FlashSystem, IBM Security QRadar SIEM can automatically trigger a response. It can send a request to IBM Storage Copy Data Management to initiate a data scan and, if necessary, start a recovery process. This integrated approach can significantly reduce the overall impact of the breach and accelerate recovery time.

2.2.5 IBM Security Guardium

IBM Security Guardium (Guardium) is uniquely positioned to offer customers a broad set of security capabilities as they embrace digital transformation initiatives and evolve their IT infrastructure to a hybrid multicloud environment. Let us look at how Guardium does that.

- ▶ **Secure modern data environments:** Guardium assists customers with a centralized view of security data and reporting capabilities from on-premise to cloud data sources. Guardium can ingest data using agent-based collectors, in real-time, from data sources containing sensitive data or they can collect data using from agentless collectors from data sources that may not contain critical data. And to address a range of security and privacy regulations, Guardium encryption can protect data at rest or data in motion, from business applications to back-end storage, across the hybrid multicloud environment.
- ▶ **Proactive security controls:** Guardium stops or contains data security threats by applying different policies across different user roles to ensure the proper checks and balances are in place -- such as an unauthorized user using elevated access privileges to move large amounts of information, after hours, to a private cloud storage location. Guardium also uses AI and algorithms to self-learn the regular logical operations that occur within an environment (such as payroll activities, banking transactions, and the sequences demonstrated by other business process patterns). Then, if the logical operation varies in any way, Guardium would flag an anomaly with a high-risk score, and an alert would also be sent to the data security team for investigation.
- ▶ **Connected data security:** Guardium provides analytics-based, in-depth insight while seamlessly integrating into existing security solutions, such as QRadar, Splunk, Resilient®, CyberArk, and HP ArcSight to name a few. In addition, Guardium provides a modular integration model with existing IT systems, such as data management, ticketing and archiving solutions such as IBM Cloud® Pak for Data, ServiceNow, and Amazon Simple Storage Service (S3). The goal is to streamline IT and security operations by complementing and extending them with data security capabilities.
- ▶ **Simplified compliance, auditing, and reporting:** Guardium is architected to provide data security administrators with a centralized hub where they can store data security and compliance data to improve operational efficiencies and assist with addressing compliance with pre-built compliance templates and workflows to monitor personal data and produce reports in seconds.

For more information see the [IBM Guardium website](#). You can also book a live [demo](#).

2.2.6 IBM Guardium Data Protection

IBM Guardium Data Protection is a data security platform that empowers security teams to safeguard sensitive data through automated and continuous discovery and classification, in-depth vulnerability assessments and advanced threat detection. It extends comprehensive data protection across disparate data environments, including all of the following:

- ▶ Databases
- ▶ Data platforms
- ▶ Data warehouses
- ▶ File shares
- ▶ File systems
- ▶ Mainframes

IBM Guardium offers a centralized control console to help you streamline management and security of your data-both on-premises and in the cloud-to meet security and compliance requirements.

IBM Guardium Data Protection evaluates risk based on severity, helping analysts remain up to date on suspicious user access or connections to data sources. The platform also identifies the presence of sensitive objects in data traffic related to monitored sources. The insights generated from these and other processes inform and enhance the following additional capabilities:

- ▶ User access profiling
- ▶ Near real-time forensics search and analytics
- ▶ Outlier detection algorithms
- ▶ Detailed investigative dashboard

IBM Guardium Data Protection supports

- ▶ Oracle
- ▶ Microsoft SQL
- ▶ Intersystems IRIS

For more information, see the following:

[Supported Platforms and Requirements for Guardium Data Protection 11.5](#)

[IBM Guardium supported data sources](#)



Scanning engine and its technology

This chapter discusses the technology and process of the malware scanning software of IBM Storage Sentinel. It describes the planning process, the different options for implementation, and other considerations.

This chapter has the following sections:

- ▶ “Storage Sentinel architecture” on page 30
- ▶ “The advantage of anomaly scanning versus signature scanning” on page 31
- ▶ “The scanning process” on page 32
- ▶ “Scanning process for databases” on page 33
- ▶ “Machine learning” on page 33
- ▶ “Scanning encrypted data” on page 34
- ▶ “How to recognize and handle alerts” on page 35
- ▶ “Scanning engine planning considerations” on page 37
- ▶ “Administration” on page 38

3.1 Storage Sentinel architecture

A data scan is an important part in maintaining data integrity and ensuring data remains unaltered, particularly in today's complex cyber threat landscape. By verifying the consistency and authenticity of data stored in Safeguarded Snapshot volumes, scanning prevents data corruption and malicious tampering caused by cyberattacks. This process becomes invaluable when recovering from an incident, ensuring that restored data is clean, uncorrupted, and safe to reintegrate into operational environments.

Mitigating Risks in Data Recovery

Effective data scanning helps organizations respond efficiently in the following real-world scenarios:

- ▶ **Manual analysis of restored data:** Following a suspicious cyberattack, organizations may choose to restore data into an isolated environment, commonly known as a “Clean Room”, for analysis. Without automated, intelligent scanning tools like IBM Storage Sentinel, this process can be labor-intensive and prone to human error. IBM Storage Sentinel’s scanning capabilities streamline this process by automatically identifying and isolating potential threats embedded in the restored data.
- ▶ **Restoring multiple data versions:** When restoring compromised data, organizations often need to go through multiple backup versions to identify a clean, uncompromised version. IBM Storage Sentinel simplifies this process by scanning each version for anomalies or malicious behavior, ensuring that only verified, clean versions are restored. This significantly minimizes the time and effort necessary to recover from an attack.
- ▶ **Preventing time-delayed attacks:** One of the most sophisticated techniques used by modern attackers is embedding malware that activates after a delayed period, often lying dormant until after data restoration. This tactic can be particularly damaging if undetected during the initial recovery process. With IBM Storage Sentinel anomaly detection capabilities, organizations can scan for early indicators of dormant malware, such as subtle changes in file metadata or unusual data access patterns. This proactive scanning helps to prevent attackers from executing time-delayed threats and ensures that restored data is truly safe.

Modern workflow for data recovery

The overall process of scanning data involves taking a Safeguarded Snapshot of production data, which is then restored into a recovery volume. This volume is mounted to the scanning engine within a secure clean room environment. The scanning engine analyzes the data for any signs of abnormal activity, focusing on anomalies that suggest malware, corruption, or manipulation of file systems. See Figure 3-1 on page 31.

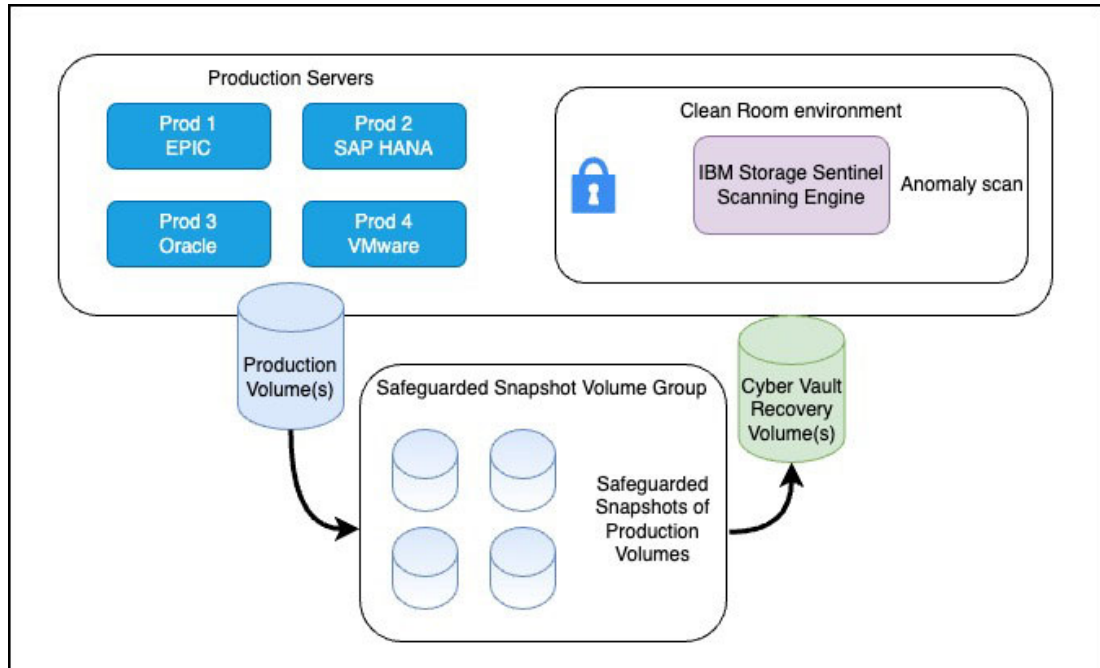


Figure 3-1 IBM Storage Sentinel Scanning workflow

3.2 The advantage of anomaly scanning versus signature scanning

In the face of increasingly sophisticated cyber threats, traditional signature-based scanning methods, which rely on identifying known malware signatures, are becoming less effective. Attackers frequently modify existing malware or develop new variants that bypass signature detection, especially with the rise of ransomware-as-a-service (RaaS) platforms and polymorphic malware, which constantly changes its code to evade detection.

Signature scanning limitations

Signature-based antivirus software operates by comparing files to a database of known malware signatures. While this approach was once effective, it now suffers from several key limitations:

Inability to detect zero-day threats

Zero-day exploits—vulnerabilities that are unknown to the public or software vendors—pose a significant challenge to signature-based tools. Until the exploit is identified and a signature is added to the database, traditional antivirus programs are blind to the threat.

Ransomware evasion tactics

Today's ransomware often employs techniques like partial file encryption or manipulation of metadata, making it hard to detect based on signature scanning alone. Moreover, malware developers continually release new variants, leading to constant delays between the discovery of a threat and the distribution of updated signatures to protect against it.

Anomaly scanning

In contrast, anomaly scanning -such as the approach used by the IBM Storage Sentinel scanning engine- focuses on detecting abnormal patterns or behaviors within data, rather than relying on known signatures. This allows for the detection of both known and unknown threats, including zero-day attacks.

Behavioral analysis

Anomaly scanners look at the behavior of data rather than its exact structure or content. For instance, ransomware generally behaves in predictable ways, such as encrypting files or altering metadata. Even if the specific malware variant is unknown, these behavioral markers remain consistent. IBM Storage Sentinel detects these behavioral anomalies by continuously analyzing more than 200 data points within each file.

Zero-day and advanced threat detection

Because anomaly scanning does not rely on predefined signatures, it excels at detecting new, unknown threats. This includes zero-day vulnerabilities and fileless malware, which often bypass traditional detection methods. For example, Storage Sentinel can detect encryption patterns typical of ransomware even when there are no signatures available for the specific malware variant.

Proactive malware detection

Anomaly scanning can identify the early signs of an attack -such as a sudden increase in data entropy (an indicator of encryption)- before the malware fully executes. This proactive approach significantly reduces the risk of damage.

Detecting intermittent encryption

Some new ransomware variants employ encryption algorithms that partially corrupt files, enough to make them unavailable but not enough to trigger an increase in entropy. Relying on significant entropy changes is not enough, other criteria must be utilized such as inspecting the content to uncover this hidden corruption.

Continuous learning

Unlike static signature-based solutions, the IBM Storage Sentinel anomaly scanning engine utilizes machine learning to improve detection over time. As the system scans more data and observes more threats, it refines its algorithms, reducing false positives and improving detection rates. The use of machine learning enables Storage Sentinel to recognize patterns of malicious behavior that may have been missed by signature-based detection.

3.2.1 The scanning process

The scanning engine performs a full scan during the first time pass, and builds the metadata of what the data looks like now and it's called Index. Subsequent scans are incremental for filesystems, in comparing what the changed data looks like at that point to what it previously looked like. There are several checks that go into building the metadata and comparing the first scan versus subsequent scans.

The scanning engine uses machine learning (ML) technology to find anomalies, which are considered typical signs of malware activity. It analyzes more than 200 different data points of a file, and compares these data points with thousands of malware patterns, their effects, and previous data collected from tens of thousands of affected backups. This helps it to accurately understand if the files, or databases, are compromised by ransomware. One of those data points is entropy, which is one of the tell-tale signs of encryption. The deep level of analysis allows it to identify the most subtle attacks where bad actors are using partial encryption to accomplish two tasks.

The primary goal of an attack is to avoid detection by changing a small portion of the file only, which is undetected by tools that are basing their scans only on metadata or exceeding thresholds. The second goal is to perform the attack as quickly as possible, and by touching only a portion of each file they compromise, move more quickly through a file system.

Unlike many standard antivirus software products, with the scanning process, the scanning engine builds an index containing historical data about previous scans. The comparison between current and older scan data helps it detect any suspicious changes between subsequent scans, improving the accuracy and sensitivity of the scan results.

3.2.2 Scanning process for databases

In active database applications, corruption can go unnoticed until the database is taken offline. It fails only as the database is brought online. Because the scanning engine can detect corruption in a database's snapshot or backup, whether it was from ransomware or from any other issues, even the most critical applications can be validated.

For databases all scans are full and the scanning process differs from the one used for non-database files. The scanning engine looks for signs of corruption due to a ransomware attack. When it scans databases, the scanning engine first identifies the type of a file, verifying that the headers and metadata match the known format of the database file. Second, it begins examining the structure of the file, verifying that it is readable and has integrity. Lastly, it scans at the page level and verifies that the individual pages are intact and corruption free.

3.2.3 Machine learning

As the sophistication of cyber threats like ransomware continues to evolve, traditional detection methods have become insufficient. The IBM Storage Sentinel scanning engine uses advanced machine learning (ML) techniques to detect threats that might otherwise bypass signature-based detection. Machine learning enhances the scanning process by identifying anomalous patterns and behaviors in data rather than relying solely on known malware signatures.

Significant developments in Machine Learning for scanning

IBM Storage Sentinel, has been validated to detect ransomware corruption in Safeguarded Copy Snapshots with an unprecedented 99.99% effectiveness. This high accuracy is achieved through the continuous training of ML models, which are trained on petabytes of real-world attack data and simulated ransomware patterns. These models are rigorously tested and updated to stay ahead of emerging threats and ransomware variants.

Proprietary AI/ML engine:

The heart of accuracy lies in its proprietary AI engine, which uses data from over 7,000 known ransomware variants. The system inspects the full content of files, not just metadata or superficial attributes, to detect the most subtle changes caused by ransomware corruption, including file obfuscation, encryption, or modification.

Comprehensive data analysis:

IBM Storage Sentinel evaluates over 200 different data points per file, including:

- ▶ File Properties: Changes in file size, compression, and entropy levels.
- ▶ Number of Files Added, Deleted, or Modified: Tracks abnormal file operations.

- ▶ **Deep Content Inspection:** Beyond metadata, Sentinel analyzes how the content of copies evolves over time, detecting anomalies that could indicate a ransomware attack.

This in-depth analysis allows the engine to precisely detect even the most sophisticated ransomware attacks, offering an unparalleled level of protection for data stored in backups and snapshots.

Continuous learning and adaptation

Unlike traditional models that rely on static updates, the IBM Storage Sentinel ML engine is continuously trained and updated. The AI models are updated regularly, typically several times a year, to ensure they can recognize new ransomware behaviors as they emerge. The training process involves millions of samples, both from real-world environments and from controlled detonations of ransomware.

Data sources for training

IBM Storage Sentinel draws from a variety of sources to build its ML models, including:

- ▶ **Subscription Services** (for example, VirusTotal) for acquiring the latest ransomware executables.
- ▶ **Public Sources** (for example, academic research, global threat intelligence, dark web monitoring) for emerging threats.
- ▶ **Customer Data:** Anonymized customer data is contributed daily to help train the models and improve detection accuracy.

Reducing false positives with Supervised Learning

False positives can disrupt operations and lead to unnecessary recovery processes. IBM Storage Sentinel, is designed to minimize false positives by continuously improving its ML models. In its most recent validation, the system achieved an estimated 99.99% accuracy, based on an extensive testing process involving more than 125,000 data samples.

Iterative learning

After training on an initial set of 2 million samples, the ML models are tested against 15 million customer samples. Any incorrect predictions are fed back into the training process, ensuring that the models are continuously improving. This iterative learning process ensures that the system adapts to real-world scenarios, refining its ability to detect threats with minimal false alarms.

Threat detection and reporting

Sentinel enables detection of ransomware attacks and issues on safeguarded snapshots and alerts users. This capability allows organizations to respond faster to attacks and speeding up recovery efforts.

Blast radius detection

Upon detecting an attack, Sentinel analyzes the extent of the corruption (for example, the blast radius) and generates detailed forensic reports to aid recovery. These reports help identify which backups are affected and which are safe to restore, reducing downtime and recovery costs.

3.2.4 Scanning encrypted data

Whether encrypted data can be scanned depends on where the encryption occurred:

- ▶ Data that is encrypted at the volume level cannot be read and analyzed by the scanning engine, but fails to mount successfully and triggers an appropriate alert.

- ▶ If data is encrypted at file level, then changes to file name, file type, and file suffix can be discovered and analyzed if the change is suspected to be a malware effect. This type of corruption is more typical of a generic malware attack on a file system and will render a DB inaccessible, which will trigger an alert.
- ▶ If a database is encrypted for security purposes, the scanning engine cannot read the data from each page but can still see the structure and understands the concept of user-based application encryption. It can discover if the file contains corruption due to ransomware.

3.2.5 How to recognize and handle alerts

The first sign of an alert can be seen in the IBM Storage Copy Data Management GUI. In the job list, you see a failed scanning process with the message “*Threat detected*”. An example is shown in Figure 3-2.

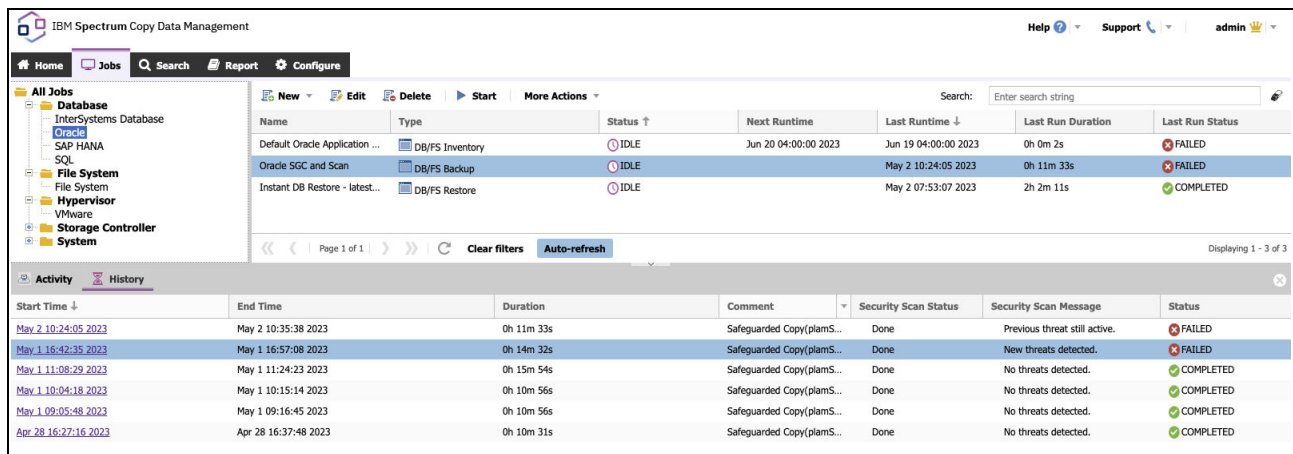


Figure 3-2 Threat detected message

If you analyze the job log, you see exactly when and on which volume the threat was detected, as shown in Figure 3-3.

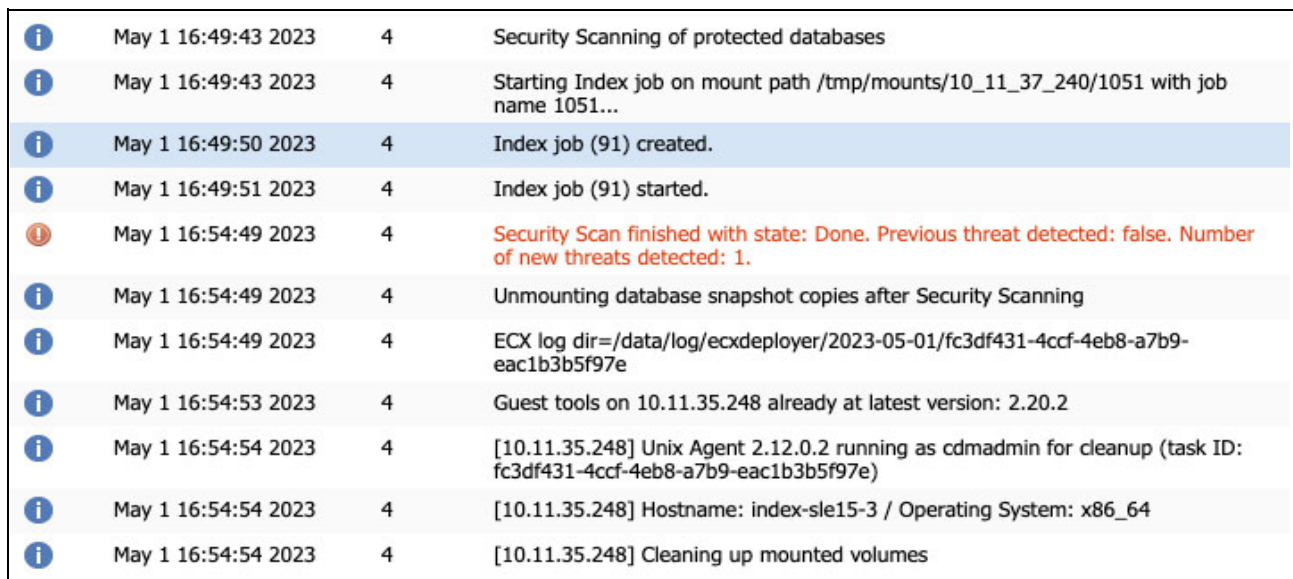


Figure 3-3 Job log showing details on the detected corruption

The IBM Storage Sentinel dashboard contains more details about the nature of the detected threat. See Figure 3-4.

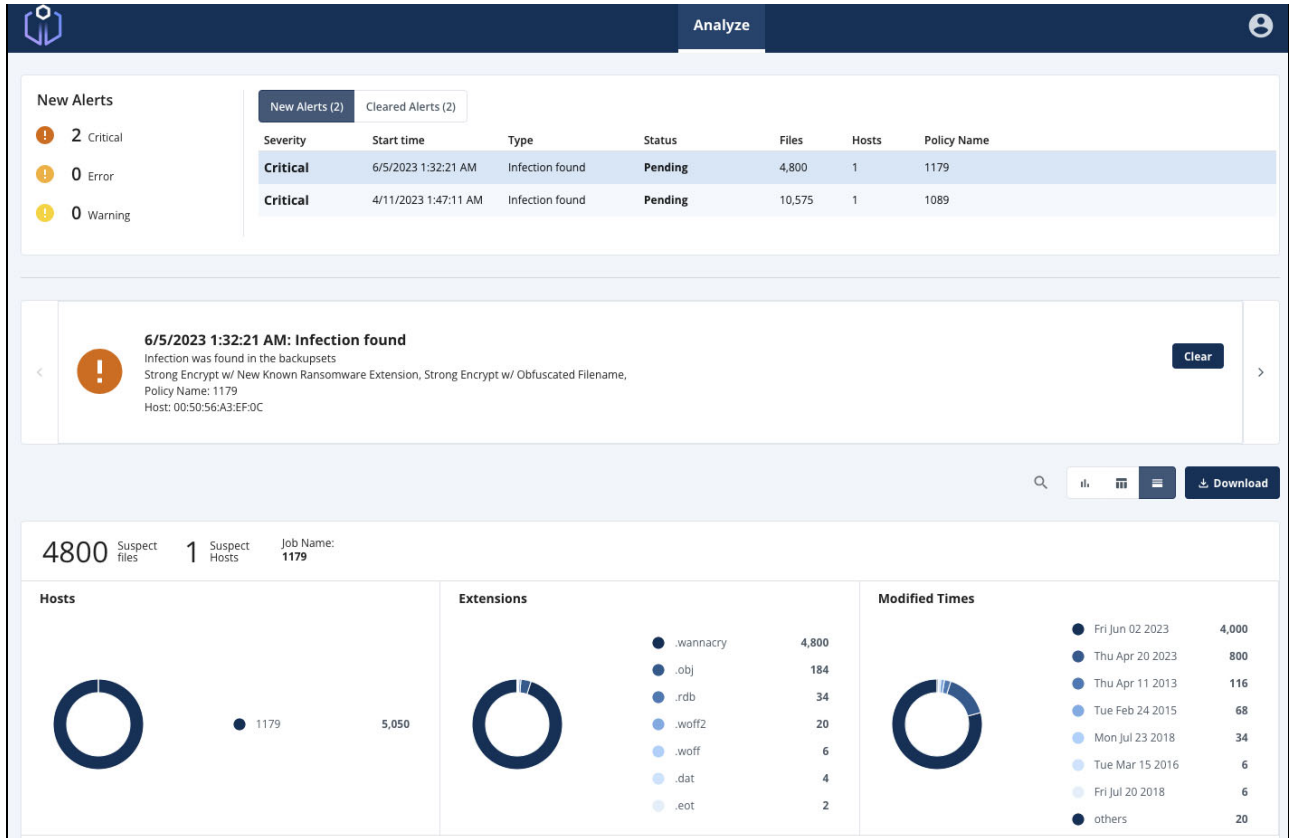


Figure 3-4 IBM Storage Sentinel Scanning Engine dashboard

3.2.6 After alert workflow

IBM Storage Copy Data Management raises an alert if a job fails due to Storage Sentinel detecting corruption. Also, the scanning engine itself can be configured to trigger email notifications and SYSLOG output. The latter can be scripted into any solution that can query logs, such as the typical SIEM and SOAR tools that are available.

3.2.7 What to do when the scanning engine finds an issue

If a backup is flagged for potential corruption due to ransomware, contact IBM support immediately. IBM support teams are trained to help customers and will involve further levels of support up to the development labs as needed. It is critical that a failed scan be analyzed by specialists to verify if an actual corruption is correctly detected, or if the failed scan is a rare false positive.

3.2.8 Dealing with false positives

When an anomaly found by the scanning engine is analyzed but is found to be an intended data change, it is called a false positive. In such a case, the administrator should make a note of what files were involved. If the false positive is caused by a file type that the scanning engine server does not recognize, the information should be submitted to IBM support describing the file type for inclusion in future builds.

3.3 Scanning engine planning considerations

You can regard the scanning part of the solution mostly as a black box, but understanding some of the concepts is important.

3.3.1 Sizing considerations

For sizing recommendation of the scanning engine, see [Server requirements and recommendations](#).

3.3.2 Federation setup and scaling of scan workloads

In the latest release of the software, CyberSense licenses are no longer constrained to the confines of an engine. In the new multi-federation licensing, one engine can act as a license server for several other engines by dynamically distributing its license among all engines at the site to meet usage needs. Engines can automatically request and return license counters as their active data changes over time, ensuring the most efficient distribution of the licenses among all the engines registered with the engine acting as the license server.

Federation helps to scale the scanning process by allowing multiple engines to collaborate and distribute tasks. Each scanning engine can run multiple jobs simultaneously. The scanning application is multi-threaded to improve performance.

The history of scans is one of the things that makes the Storage Sentinel scanning so powerful. However, when you change the scanner that is used for a specific instance of your application, the history is not transferred, and the new scan engine creates a new history. If you are running a single server, the best practice is to stagger the workloads throughout the window of processing. If you are running multiple servers, balance the scanning jobs across the available scanning engines for optimum usage of the processing window. Moving a scanning job from one server to another requires a new initial scan, and you lose the scan history from the previous scans.

The implementation of Storage Sentinel creates one job per snapshot. Multiple snapshots run as multiple jobs. If you have more jobs than a single server can handle, these should be split across two or more scan servers.

Storage Sentinel has a federated licensing scheme, so licenses of all scanning engines can be managed from a central instance. For more information, see [IBM Storage Sentinel anomaly scan software 1.1.9](#).

3.3.3 Virtual versus physical servers

If the scanning engine is installed on a physical server, it can scan volumes that are used by physical application servers only. If the scanning engine is installed on a virtual machine, it is able to scan volumes used by both physical and virtual application servers.

When scanning virtual application servers, there are times that you might want to limit the number of volumes that require a Safeguarded Snapshot. For example, if the application uses volumes that are virtual disks on a VMFS datastore, you can create dedicated datastores for the VMs that contain an instance of the application. More commonly, customers can place the protected applications on dedicated disks that are presented as physical raw device mapped (pRDM) volumes or by using iSCSI. At the time of writing, Storage Sentinel does not currently support vSphere vVols.

Note: At the time of writing, scanning software supports SUSE Linux Enterprise Server 15.4 - 15.5 and Red Hat Enterprise Linux 9.2.

This means that the file systems on volumes on application servers running x64 Linux can be directly mapped and mounted to the scanning engine. However, Storage Sentinel also supports protected applications running on an IBM Power server and AIX operating system, which cannot be mounted directly to the Storage Sentinel scanning engine. Copy Data Management requires an AIX proxy machine running on the Power platform, so that Copy Data Management can mount the volumes to be scanned to the proxy, and share the file systems over NFS for the scanning engine to access. See Figure 3-5.

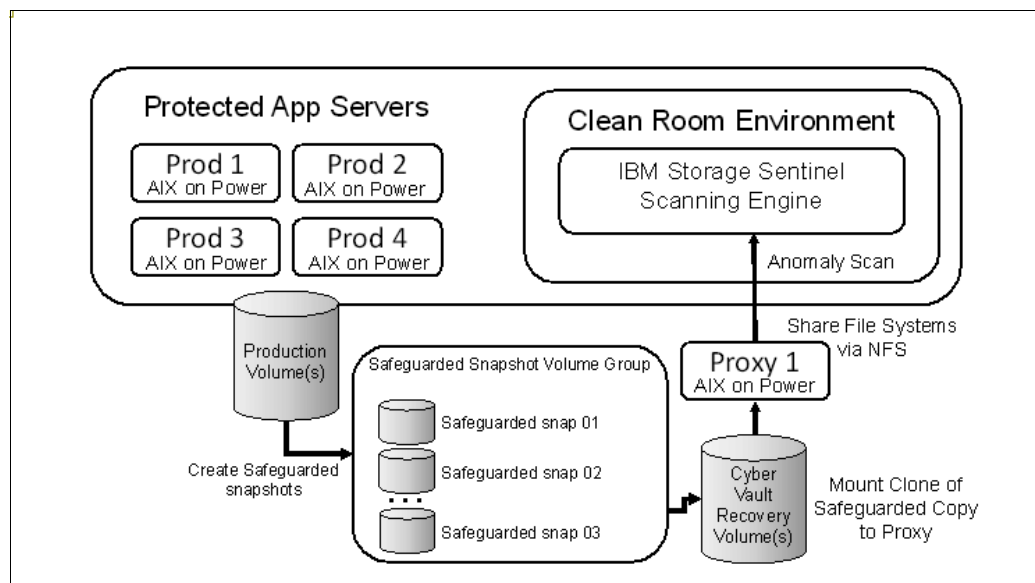


Figure 3-5 IBM Storage Sentinel configuration with an AIX proxy machine on the Power platform

3.4 Administration

This section includes discussion of the administration of the scanning engine.

3.4.1 Monitoring the scanning engine

The basic day to day monitoring task is checking the index size and the amount of front-end data that the system scans to make sure that the user is not exceeding the licensed amount. In addition, back up the configuration regularly.

IBM Storage Sentinel engine status and service status can be checked at `https://<hostname>/admin`. See Figure 3-6 and Figure 3-7.

```

Engine Status

Software Version: 8.7.0 - Build 1.16
Build Date: Aug-08 at 6:09 pm
Platform: SUSE 15

Current Date: Oct 6, 2024 11:10 am Europe/Istanbul
System Uptime: 8 days, 23 hours, 2 minutes
CPU Load Averages: 15-min: 0.08, 5-min: 0.14, 1-min: 0.55
Disk Status: Disk1 - Unconfigured
                Disk2 - Unconfigured
                Disk3 - Unconfigured
                Disk4 - Unconfigured
Disk Space: 2.2% of main file system is in use. (976.384GB is free.)
Sensors: N/A
HTTP: Disabled (Secured)

Safe Mode: Custom Mode
Index Segments: 29 (138,935 files – 14.850GB)
Installed Packages: ransomsgdb-20240802-01
                        ie_vmware_vddk-6.7.1-10362358
                        indexengines-8.7.0-1.16.suse15
                        qemu-tools-ie-6.2.0-150400.37.26.1
                        kernel-5.14.21-150400.24.100-default
    
```

Figure 3-6 IBM Storage Sentinel engine status

Service	Status	Information
Application Server	Status: Running	
Archive Configuration Manager	Status: Running	
Catalog Ingestion	Status: Running	
Extraction	Status: On Line Started: Sep-27 at 7:27 pm	State: Available
Index Manager	Status: On Line Started: Sep-27 at 5:54 pm	Role: Manager <input type="checkbox"/> Engines: Joined: 1 Indexes: Active: 1, Deactivated: 1, Total: 2 <input type="checkbox"/> Selected: Index Name: SCDM_2d0801f0-7ced-11ef-9f32-005056bab7e3, Id: 5
Indexing	Status: On Line Started: Sep-27 at 7:26 pm	Indexing Mode: Full Content <input type="checkbox"/>
Location Manager	Status: On Line Started: Sep-27 at 5:53 pm	Location Checks: Successful: 1
Post-Processing	Status: On Line Started: Sep-27 at 5:54 pm	Segments: Total: 29, Correct Version: 29
Query	Status: On Line Started: Sep-27 at 5:54 pm	
Scheduler	Status: On Line Started: Sep-27 at 5:54 pm	LAN Crawler: Available Pause Query Service: Available Pause Post-Processing Service: Available Pause
Security Domain Manager	Status: Running	
Signature Set	Status: Running	
Tape Manager	Status: On Line Started: Sep-27 at 7:26 pm	Tape Caching: Disabled <input type="checkbox"/> SQL Server: Status: Running, Server Host: localhost, Database: imtapedb5

Figure 3-7 IBM Storage Sentinel service status

3.4.2 Backing up and restoring the scanning engine components

Storage Sentinel includes utilities to allow backup software to take clean backups of the key application files. If you are running a federated Storage Sentinel environment, back up the member instances first, then the manager. At recovery time, you reverse the order and recover the manager first, then the members.

The default backup location of Storage Sentinel is `/opt/ie/backup`. However, you can specify a different location by modifying `/opt/ie/var/backup` location. The backup directory must reside in the same file system as `/opt/ie/var`.

For IBM Storage Protect (previously called Tivoli Storage Manager), follow these steps:

1. Add the following lines to `dsm.sys`:

```
PRESCHEDULECMD /opt/ie/bin/tsm_presched  
POSTSCHEDULECMD /opt/ie/bin/tsm_postsched
```

Note: The files `/opt/ie/bin/tsm_presched` and `/opt/ie/bin/tsm_postsched` already exist on the Index Engines system.

2. On the Storage Protect server, specify the copy group for the policy domain, policy set, and management class that the client belongs to and set Copy Serialization to Dynamic.
3. Define a scheduled job to back up the client directory `/opt/ie/backup` (or the alternate directory if you created one). Either an Incremental or Selective backup works. There are multiple ways to specify what you want to back up with Storage Protect. If you want to have a dedicated backup schedule for protecting the Storage Sentinel backup location, you can specify that location in the `objects` parameter in the schedule definition. You need to include any subdirectories, so specify **Subdir Yes** in the `dsm.sys` stanza or include it in the schedule settings in the `Options` parameter. For example:

```
UPDATE SCHEDULE examplePD BackupSentinel type=client action=incremental  
objects='/opt/ie/backup/*' options=-subdir=yes startdate=TODAY starttime=NOW  
dayofweek=any
```

If you are running a federated environment, be sure to back up the members before you back up the manager, so you need at least 2 schedules.

The Storage Sentinel interface is used to recover data from the current contents of the backup directory (Select **Administration** → **System** → **Recovery** → **Recover From Backup**).

If this is a pristine environment, you need to install Storage Sentinel, restore the backup directory from the backup software, and then use the Storage Sentinel GUI to rebuild the instance from the backup files. If you are running a federated Storage Sentinel environment, restore the manager first, then the members.

For successful recovery requirements, see [IBM Docs for Storage Sentinel](#).

3.4.3 Adding new applications

Register new applications and define their scanning process in IBM Storage Copy Data Management. As you add workloads to be scanned, it is important to monitor your Storage Sentinel license information so that you do not exceed your licensed capacity. The number of protected applications or the number of scan servers is not monitored or restricted by the license.

3.4.4 Adding new scanning engines

Added scanning engines must be registered in the GUI of IBM Storage Copy Data Management. The scanning process for multiple applications should be spread across the scan servers for load balancing.



Protecting SAP HANA databases

Cyber resilience is an enterprise-wide concept that is aimed at ensuring business continuity in the face of threats and failures. While this concept encompasses a wide range of areas, in this chapter we will focus on protecting SAP HANA databases using IBM cyber resilience solutions.

This chapter has the following sections:

- ▶ “Protecting SAP-HANA: A comprehensive approach” on page 42
- ▶ “Protecting SAP-HANA: Architecture and step by step implementation” on page 43
- ▶ “IBM Storage Sentinel” on page 57
- ▶ “Integrating IBM Storage CDM, IBM Storage FlashSystem and IBM Storage Sentinel” on page 65
- ▶ “IBM Security Guardium for SAP HANA” on page 72

4.1 Protecting SAP-HANA: A comprehensive approach

Cyber resilience is a comprehensive approach to ensuring business continuity in the face of threats and failures. To safeguard SAP HANA databases, IBM offers a robust solution combining IBM Storage Copy Data Management, IBM Storage Sentinel, and IBM Storage FlashSystem. This integrated approach delivers advanced protection, automation, and management capabilities by consistently creating, protecting, and analyzing data copies. Figure 4-1 illustrates a typical infrastructure leveraging these solutions.

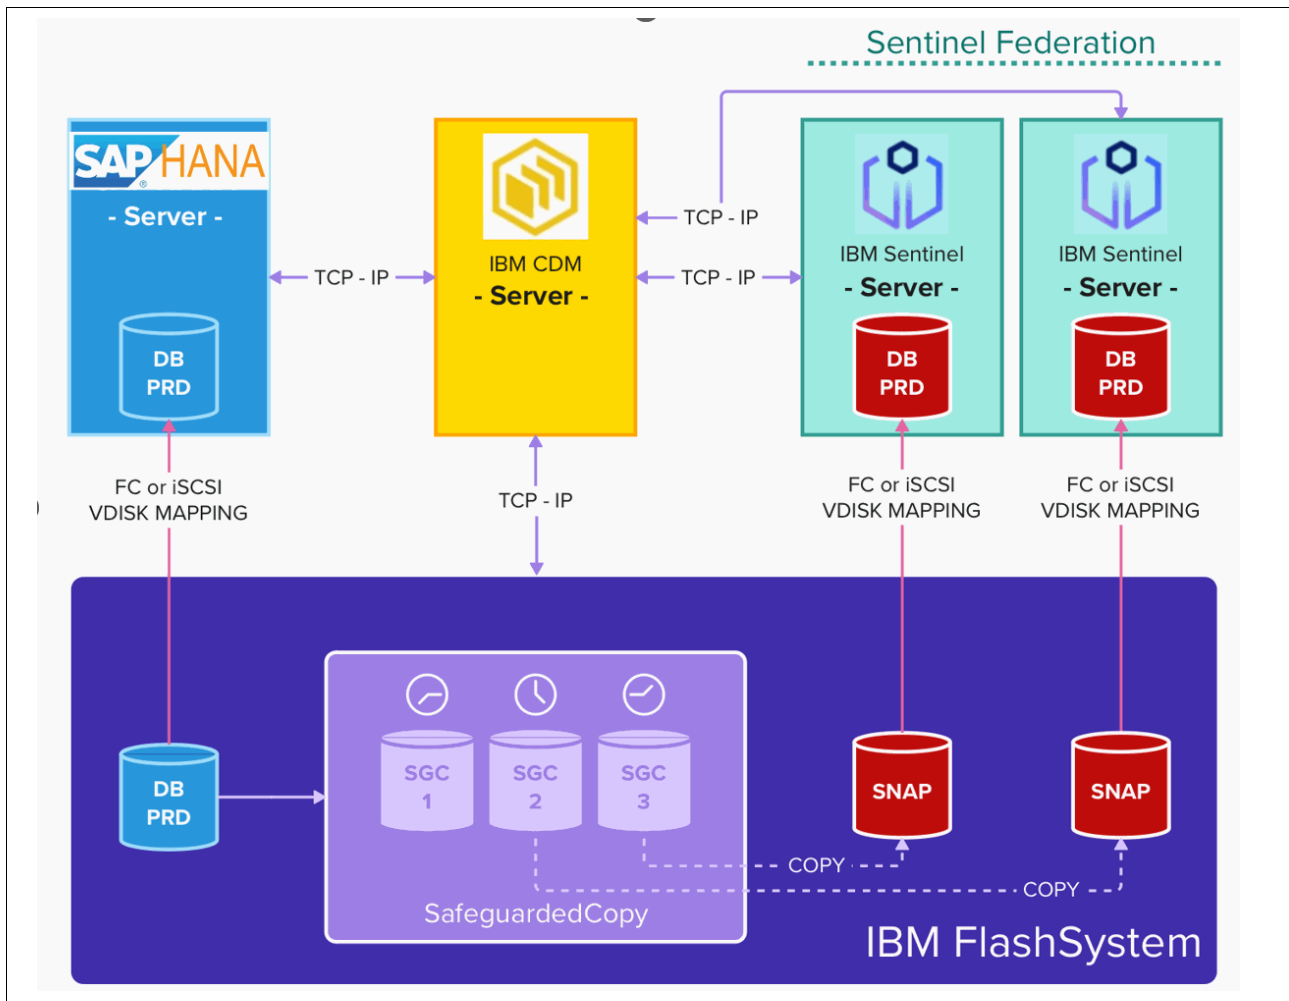


Figure 4-1 Common infrastructure to protect SAP-HANA

4.1.1 Components

The IBM Storage CDM software, running on a VMware virtual machine, acts as the central hub for automating database protection and managing the lifecycle of copied data for SAP HANA servers. Backups happen automatically, and you can monitor their progress and status through a user-friendly dashboard or receive email alerts. These SAP HANA servers simply act as clients to the CDM system, requiring no additional software installation.

CDM utilizes SSH to remotely control backup and restore processes for the SAP HANA servers, including the critical `/hana/data` directory and others.

Backup copies created by CDM are stored on the IBM FlashSystem, leveraging the benefits of Safeguarded Copy technology: immutability (protection from accidental deletion), compression, and deduplication (reducing storage footprint). These copies are guaranteed to meet your defined Service Level Agreement (SLA) and become unalterable once created.

IBM Storage Sentinel is responsible for scanning Safeguarded Copy snapshots for potential threats. CDM automates the process of creating these snapshots from the SGC data, preparing them for use. This involves mounting volumes, importing the crucial `/hana/data` volume groups, and creating necessary filesystem mount points on Sentinel servers.

While Sentinel can operate on a single server, it can be scaled to multiple servers for demanding workloads, known as *Sentinel federation*. Importantly, these servers share a single license, measured in terabytes.

The combined capabilities of these IBM appliances ensure the highest levels of data protection and resilience. Each software component plays a specific role in this process, working together seamlessly to safeguard your data, as outlined in 4.2, “Protecting SAP-HANA: Architecture and step by step implementation” on page 43.

4.2 Protecting SAP-HANA: Architecture and step by step implementation

We can break down the process into 10 steps, demonstrating how CDM, in tandem with IBM FlashSystem and IBM Storage Sentinel, orchestrates each phase of SAP HANA data protection. See Figure 4-2 on page 44.

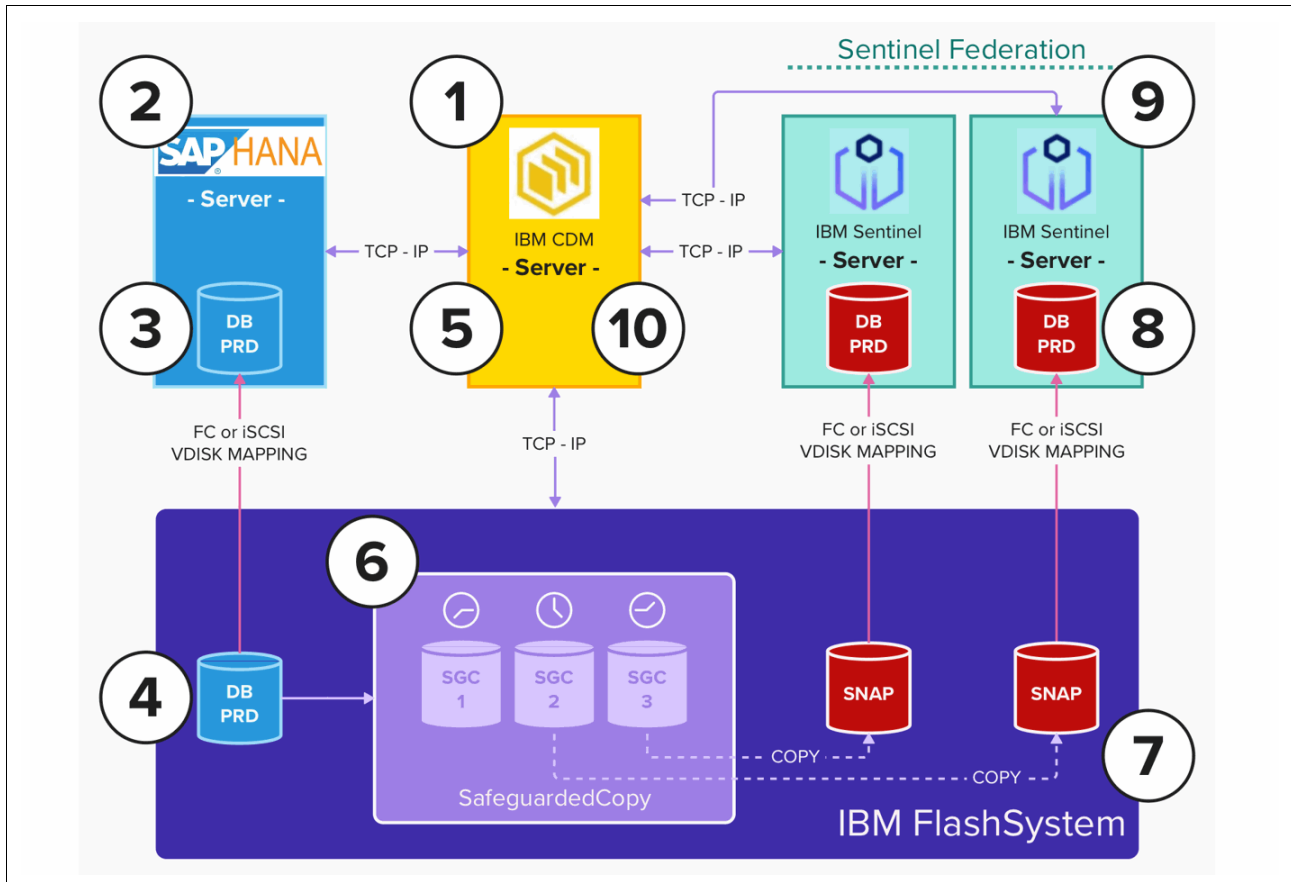


Figure 4-2 CDM snapshot process

1. CDM Server establishes communication with SAP HANA servers through SSH (port 22) and hdbclient (port 33013), leveraging pre-configured operating system and database credentials.
2. CDM remotely copies a `cdm_guestapps_cdmadmin` directory, containing essential backup scripts, to the `/tmp` folder. In parallel, CDM assesses the health of the `/hana/data/DB` filesystem and confirms the online status of the database.
3. CDM collects detailed information about the storage subsystem disks that underpin the `/hana/data/DB` filesystem. For virtual machine environments, CDM also identifies the datastore hosting the VMDK disks. To achieve application consistency during the backup process, CDM quiesces the `/data/hana/DB` filesystem by executing the `xfs_freeze -f /data/hana/DB` command.
4. CDM inventories and correlates the collected disk information to prepare a Safeguarded Copy operation, capturing a point-in-time copy of the data. A consistency group is created, encompassing the productive vdisks that constitute the `/hana/data/DB` filesystem. Subsequently, a Safeguarded Copy operation is triggered for all vdisks within this group, leveraging Storage Virtualize's `add snapshot` command. All communication between CDM and the storage subsystem is facilitated through SSH (port 22).
5. After the Safeguarded Copy is successfully captured, CDM releases the freeze on the `/hana/data/DB` filesystem by executing the `xfs_freeze -u` command.
6. To optimize storage utilization, CDM creates thin-provisioned snapshot copies of the recently taken Safeguarded Copy.

7. These snapshot copies are subsequently mapped to an IBM Storage Sentinel server, with all disks belonging to the database assigned to a single server. It is crucial to have pre-configured storage host mapping, SAN zoning, or iSCSI networking to facilitate this mapping process.
8. CDM establishes an SSH connection (port 22) to the Sentinel server to initiate disk mapping and mount the hanadatavg and hanadatalv volumes. Additionally, CDM renames hanadatavg to facilitate the concurrent mounting of multiple SAP HANA instances on the same IBM Storage Sentinel server.
9. The filesystem `/tmp/mounts/"sap-hana_server_ip"/"hana_db_name"/...` is mounted on the Sentinel Server, initiating the scanning process.
10. The progress of the threat-scan process can be monitored through CDM's GUI interface. Once the process concludes, CDM initiates a cleanup of all structures created on the Sentinel Server and Storage Virtualize. The results of the scan can be viewed in both the CDM and Sentinel GUIs.

4.2.1 Validating the SAP-HANA server requirements for CDM

IBM CDM communicates with SAP HANA servers via TCP/IP. Two credentials are required to perform backup and restore operations: one for operating system-level tasks and another for database-level tasks.

1. Before configuring HANA credentials in CDM, certain prerequisites must be met to ensure successful backups. A symbolic link must be created at the `/opt/hana` folder. This guarantees that IBM CDM will connect to the hana database as needed.

Example: `ln -s /hana/shared/ALB/hdbclient/ /opt/hana/`

(where ALB is the hana database name)

2. SAP-HANA also requires a client installed in the same server as the database lives (HANA-Client).

hdbuserstore list is a useful command to validate if HANA-Client is properly installed and running, it can also tell the port used by HANA-Client to listen to external communications (30013 is the default). See Example 4-1.

Example 4-1 hdbuserstore list

```
[ibmadm@myserver:/usr/sap/ALB/HDB00> hdbuserstore list
DATA FILE: /usr/sap/ALB/home/.hdb/myserver/SSFS_ALB.DAT
KEY FILE: /usr/sap/ALB/home/.hdb/myserver/SSFS_ALB.KEY
KEY IBMSAPDBCTRL
```

```
ENV : myserver:30013
USER: IBMSAPDBCTRL
```

```
ACTIVE RECORDS : 7
DELETED RECORDS : 21
NUMBER OF COMPLETE KEY: 1
```

```
Operation succeeds.
ibmadm@myserver:/usr/sap/ALB/HDB00>_
```

3. If the HANA client is not installed use the following command to install it.

Example 4-2 Installing the HANA client

```
cd /opt/hana
```

pip install hdbcli-x.xx.xx.tar.gz (where the x.xx.xx is the current version of hana-client)

- Now, a user with proper rights must be created in the operating system and SAP-HANA to permit control. Example 4-3 describes how to add a user in Linux machines.

Example 4-3 Create the cdmgroup group in a Linux server

```
/usr/sbin/groupadd -g 1192 cdmgroup
```

- Create the user named cdmagent that belongs to cdmgroup. See Example 4-4.

Example 4-4 Create the cdmgroup group in a Linux server

```
/usr/sbin/useradd -u '1193' -g 'cdmgroup' -m -s '/bin/bash' -c 'Copy Data Management OS user' 'cdmagent'
```

- To configure sudoers for the cdmagent user, place the following lines at the end of your sudoers configuration file, typically /etc/sudoers. If the existing sudoers file is set to import configuration from another directory (for example, /etc/sudoers.d), you can also place the lines in a new file in that directory. See Example 4-5.

Example 4-5 Configuring sudoers for the cdmagent user

```
Defaults:cdmagent !requiretty  
cdmagent ALL=(ALL) NOPASSWD:ALL
```

- For SAP-HANA database, the example in Example 4-6 can be used to create a user to interact with CDM. See Example 4-6.

Example 4-6 Creating a user to interact with CDM

```
CREATE USER IBMCDMUSR PASSWORD "xxxxxx" NO FORCE_FIRST_PASSWORD_CHANGE;  
ALTER USER IBMCDMUSR DISABLE PASSWORD LIFETIME;
```

Figure 4-3 shows the SAP-HANA database and IBM CDM Server interaction.

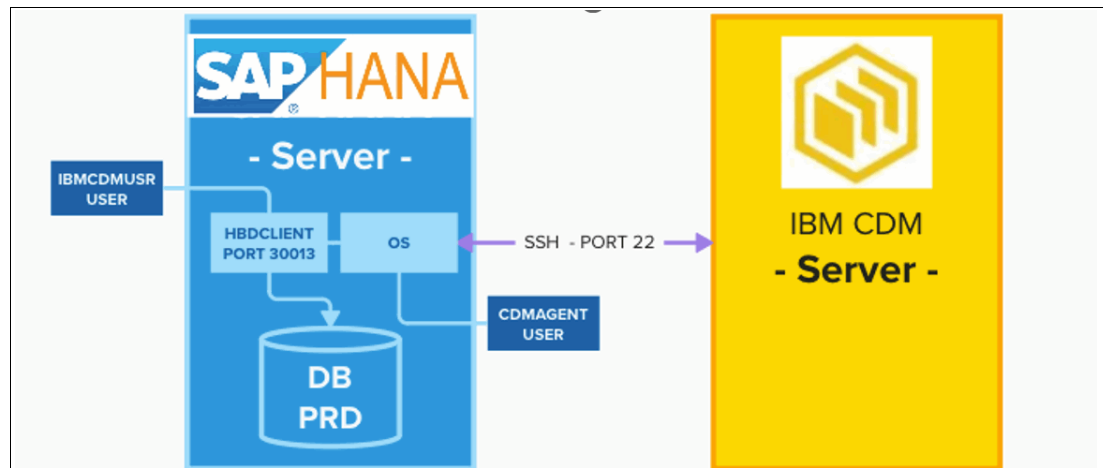


Figure 4-3 SAP-HANA database and IBM CDM Server interaction

Additional details how HANA-Client is installed and configured can be found at [Installing SAP HANA HDB Client \(Linux\)](#).

Tip: The complete guidelines to set CDM to work with SAP-HANA, can be found at [SAP HANA requirements: IBM Storage Copy Data Management 2.2.24](#).

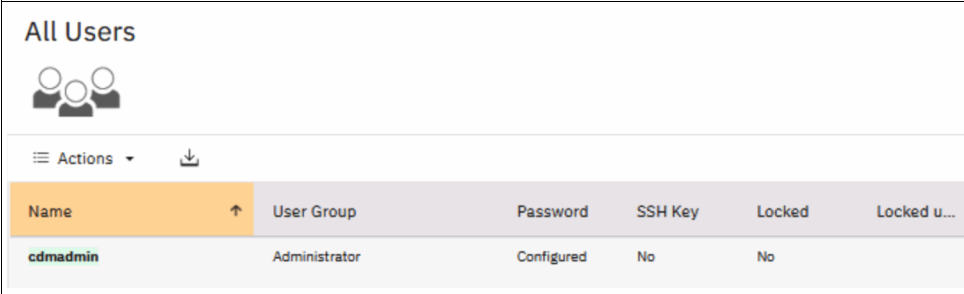
4.2.2 Configuring IBM FlashSystem credentials

Once the prerequisites are met, adding SAP HANA to CDM is straightforward. Configuring a CDM SAP HANA job for Safeguarded Backup and IBM Storage Sentinel data validation is crucial.

1. In the IBM FlashSystem, create a storage administrative user to manage all CDM operations. This can be set by GUI (Figure 4-4) or CLI (Example 4-7).

Example 4-7 mkuser

```
mkuser -name cdmin -usergrp Administrator -password xxxxxx
```



Name	User Group	Password	SSH Key	Locked	Locked u...
cdmin	Administrator	Configured	No	No	

Figure 4-4 Create a storage administrative user through the GUI

2. Log in to the IBM FlashSystem before adding the credential to the CDM to validate access.

Tip: Do not use superuser credentials for CDM operations; this poses an unnecessary security risk.

3. Test SSH access (port 22) from the CDM server to the IBM FlashSystem (Example: cdmin@fs9500.ibm1ab.com).
4. Log in to the CDM GUI (<https://CDM-SRV-IP:8443/portal>) and register storage subsystems and their credentials under the **Sites and Providers** tab. See Figure 4-5 on page 48.

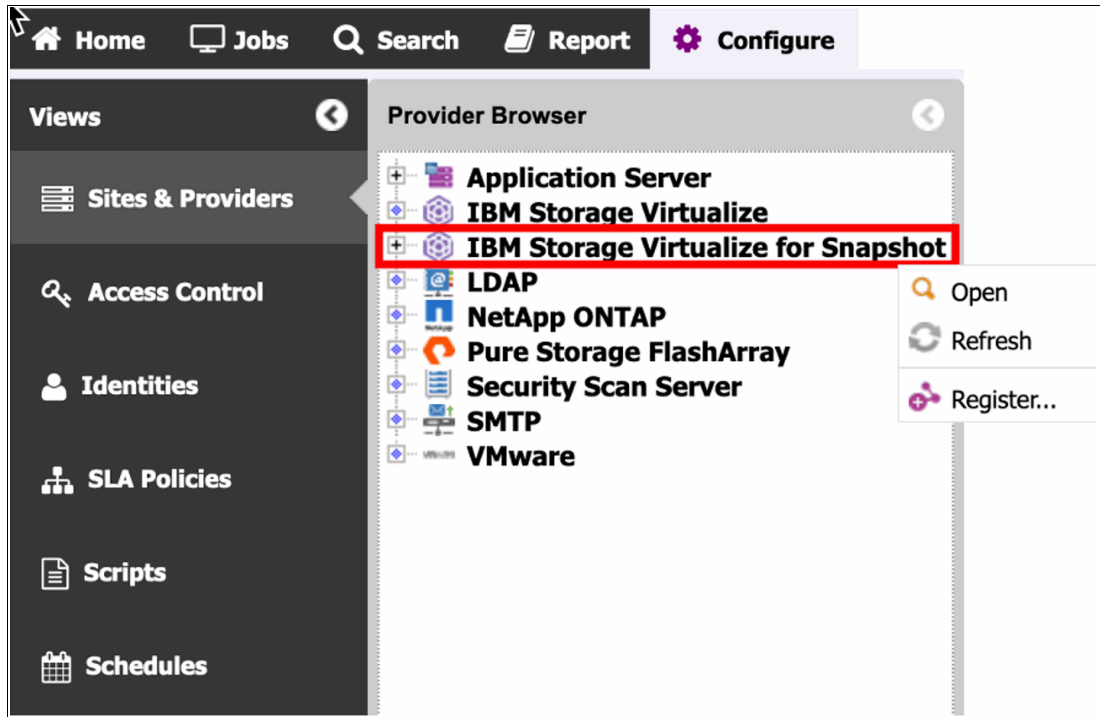


Figure 4-5 Register storage subsystems and their credentials

5. Configure CDM to connect to the storage subsystem using the specified IP address. In our lab example we used FQDN (fs9500.ibm1ab.com). See Figure 4-6 on page 49.

Register IBM Storage Virtualize for Snapshot provider

Site: Default

Name: IBM FS9500

Host Address: fs9500.ibmmlab.com

Comment: This storage holds SAP-HANA PRD and SGC copies

Run Inventory job after registration

Select **New**

Name	Username	Type

OK Cancel

Figure 4-6 Configure CDM to connect to the storage subsystem

6. Click **New** to set the storage credentials to this configuration. See Figure 4-7.

Create Credential

Name: FS9500_CDM_UsrKey

Username: cdmadmin

Password:

Comment: stg-admin user for FS9500

[Need help?](#)

Create Discard

Figure 4-7 Set the storage credentials

7. Once the configuration is completed, the credentials will be seen in the **Identities** tab. See Figure 4-8 on page 50.

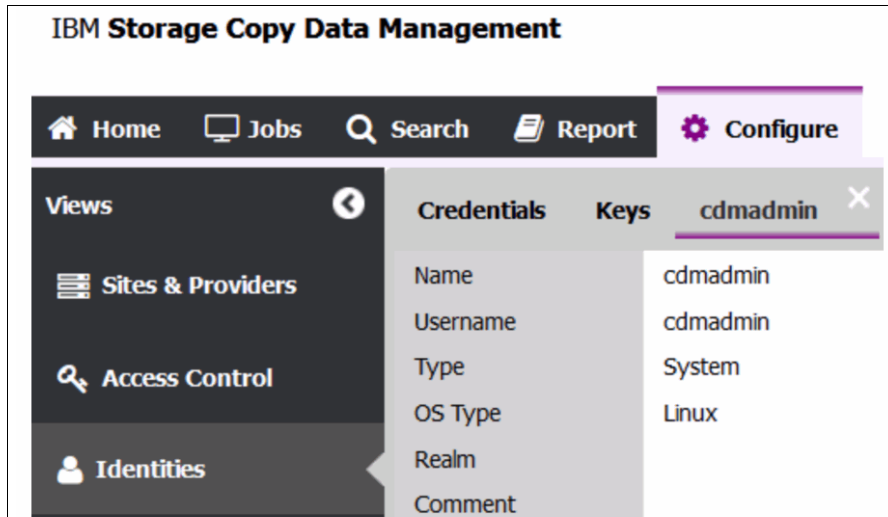


Figure 4-8 Storage credentials

4.2.3 Defining storage SLA policies

A minimum of one policy must be configured in CDM's SLA Policies tab to enable backup job execution. IBM Storage offers two distinct protection methods, each tailored to specific IBM Spectrum Virtualize code levels.

The CDM SLA tab offers two distinct protection options: IBM Storage Virtualize and IBM Storage Virtualize for Snapshot. The selection of the appropriate option is influenced by factors such as storage hardware generation, software version, and the desired number of safeguarded backups. Both technologies are capable of safeguarding data against cyber threats.

The IBM Storage Virtualize option is directly associated with the first-generation copy services, which comprise Global Mirror with Change Volumes and FlashCopy 1.0. This method necessitates the creation of FlashCopy mappings for each copy and Safeguarded Copy, requiring SGC backup locations, typically child or parent pools.

Starting with 8.5.4, IBM Storage Virtualize introduced a second generation of copy services (2.0), offering significant improvements for safeguarded copies. This generation enables snapshot management at the volume group level, eliminates the requirement for SGC backup locations, and automatically applies immutability to all snapshots. Additionally, GMCV is being replaced by PBR (policy-based replication), although PBR was not yet supported by CDM at the time of this writing.

1. Create a new SLA using the **IBM Storage Virtualize for Snapshot** option. See Figure 4-9.

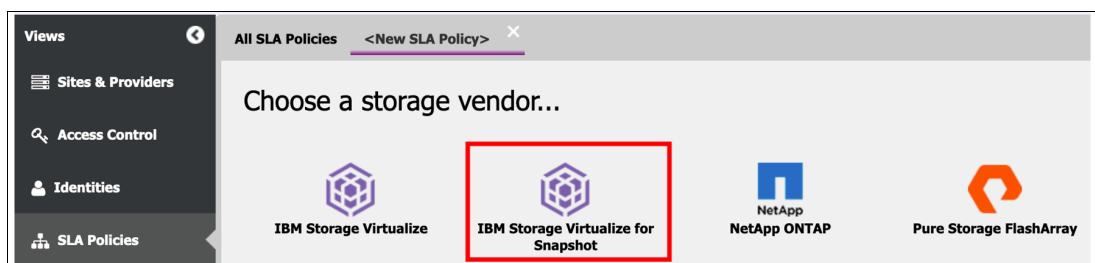


Figure 4-9 IBM Storage Virtualize for Snapshot option

2. At the wizard, choose the **Add Safeguarded Copy** option. See Figure 4-10.

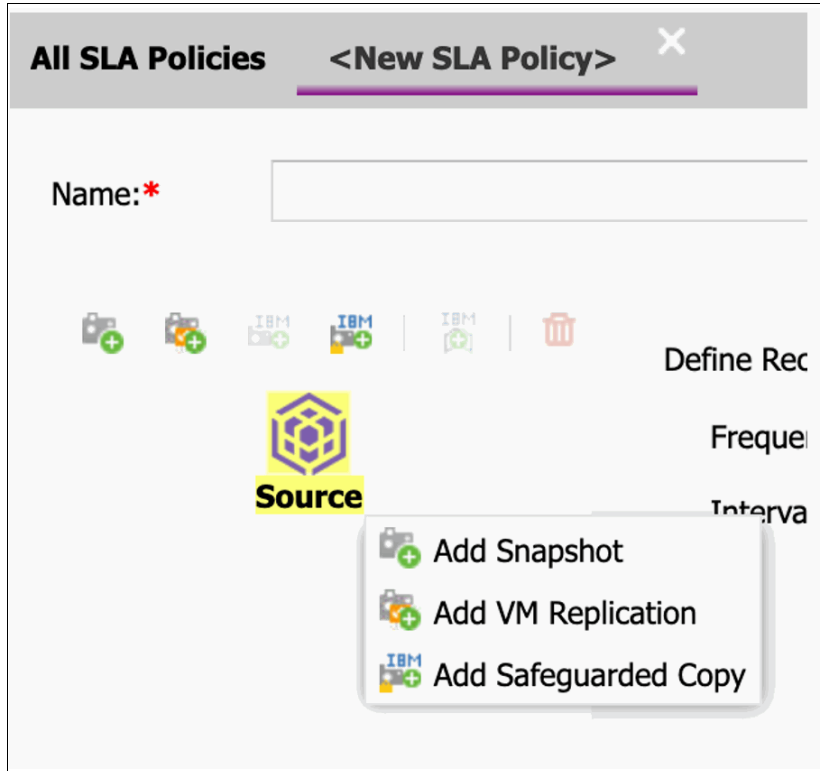


Figure 4-10 Add Safeguarded Copy option

3. In the **<New SLA Policy>** tab we are going to set the copy frequency. This can vary from minutes to hours, days, weeks until months. The minimum space between one copy and another are 5 minutes. See Figure 4-11.

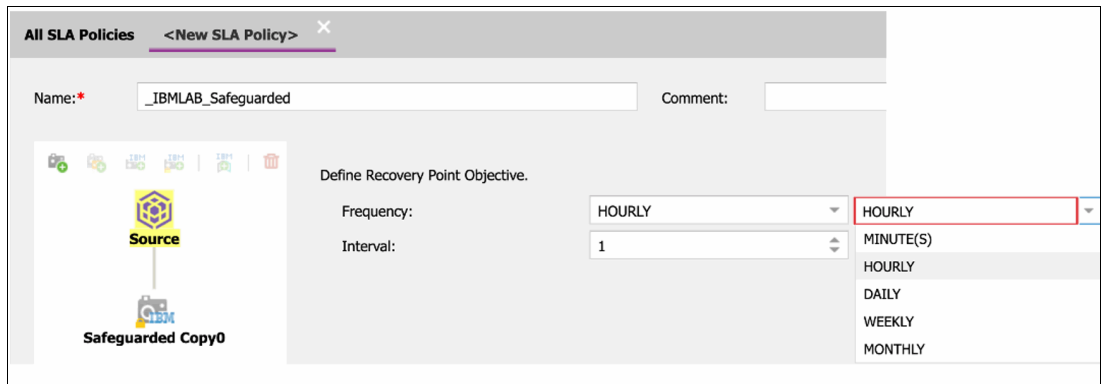


Figure 4-11 Set the copy frequency

4. By clicking the **Safeguarded Copy0** icon, we can specify the storage pool destination. While the image shows only the main pool, child pools can also be used to store and track the space utilized by Safeguarded Copies.

In the Options tab, set the Safeguarded Copy retention (**Keep Snapshots**). Once the specified retention period is reached, IBM FlashSystem will automatically delete the copies, freeing up storage space and maintaining optimal pool utilization.

Optionally a prefix can be added to the snapshots created by the SLA policy. See Figure 4-12.

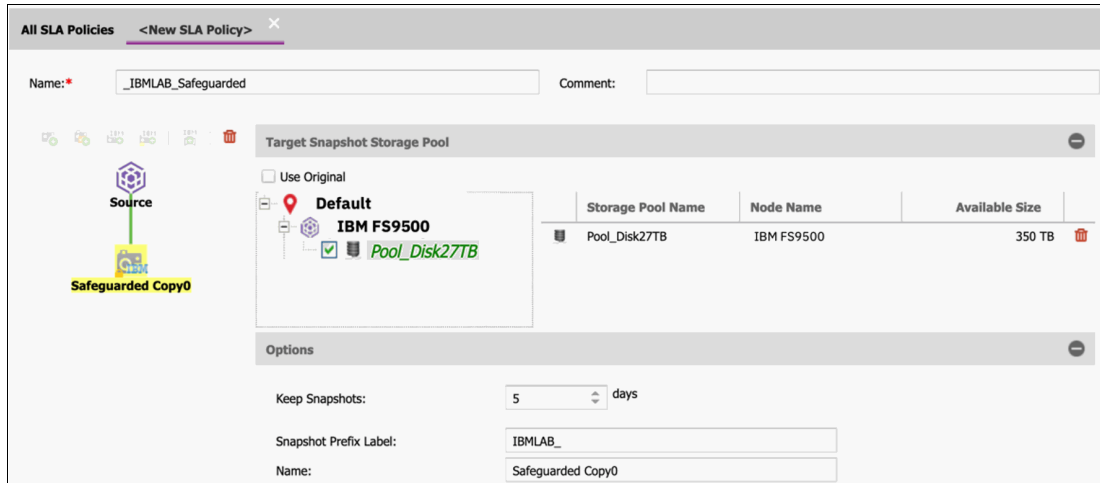


Figure 4-12 Options tab

5. The created SLA policy will be available at the **Configure** → **SLA Policies** tab. See Figure 4-13.

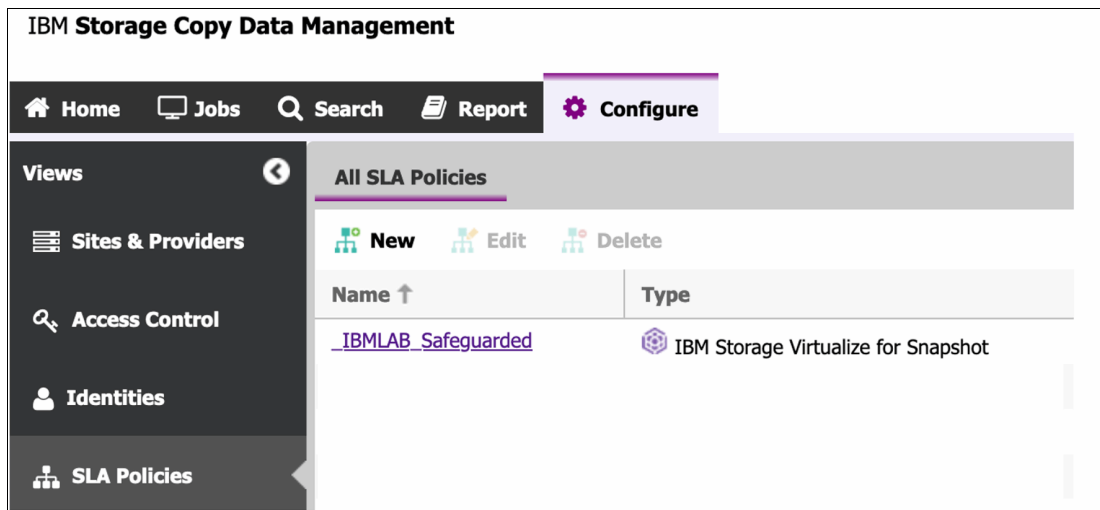


Figure 4-13 SLA Policies tab

4.2.4 Creating SAP-HANA backup jobs

Once the SLA policy has been created, proceed to the Sites and Providers tab to register one or more SAP HANA machines and their corresponding credentials.

1. Select the **Register** option. See Figure 4-14 on page 53.

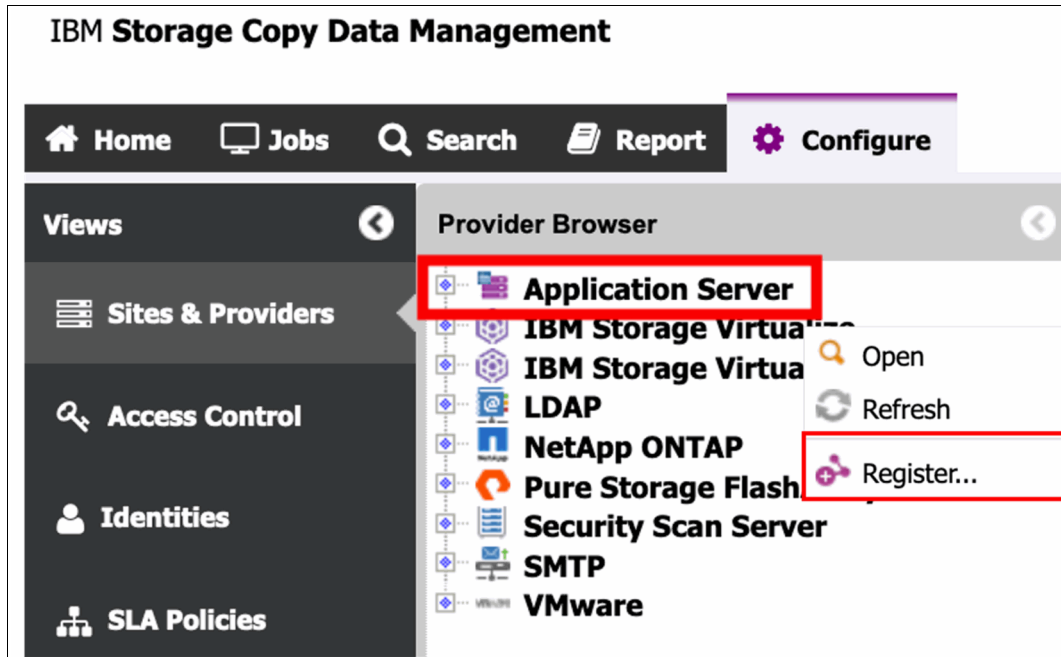


Figure 4-14 Select the Register option

2. Register one or more SAP-HANA machines and their credentials to the same Sites and Providers tab. Select **SAP HANA** as the server type. Figure 4-15.

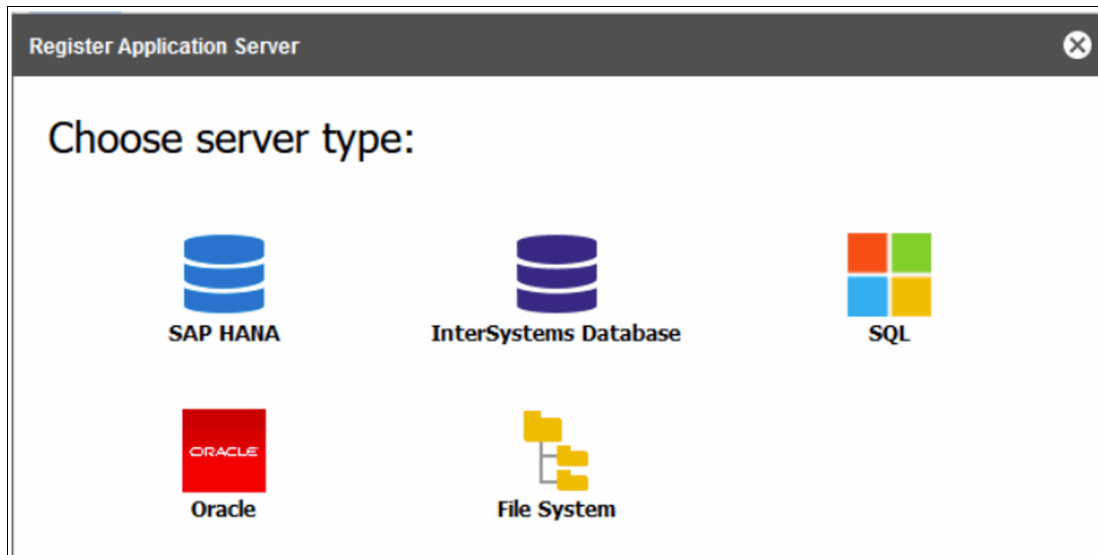


Figure 4-15 Select SAP HANA as the server type

3. The registration can be performed by adding the server FQDN (Fully Qualified Domain Name) or IP address to the host address field.

The port used is the same as shown in the SAP-HANA `hdbuserstore list` command output (See Example 4-1 on page 45). If the **Run Inventory Job After Registration** option is selected, CDM will connect to the SAP HANA server to gather database configuration information.

If this initial registration step is omitted, it can be executed at a later time using the default SAP HANA application job.

When registering the SAP HANA machine, select either **Physical** or **Virtual** based on its type. For virtual machine environments, it is essential to have the vCenter, where the SAP HANA machines reside, previously added to the CDM Sites and Providers. See Figure 4-16.

Register Application Server (SAP HANA)

Site:

Name:

Host Address:

Port:

Run Inventory job after registration

Type: Virtual Physical

System Credential:

Select **New**

Name	Username	Type

Database Credential(s):

Select **New**

Name	Username	Type	Instance

OK Cancel

Figure 4-16 Register Application Server

4. Create the system credentials using the SAP-HANA operating system user and password. See Figure 4-17 on page 55.

Create Credential [X]

Name: saphana-srv01-cdmagent

Username: cdmagent

Password:

Comment: OS user to connect to sap-hana operating system

[Need help?](#)

[Create] [Discard]

Figure 4-17 Create the system credentials

5. Create the database credentials using the SAP-HANA database user and password. See Figure 4-18.

Create Credential [X]

Name: saphana-srv01-ibmcdmusr

Username: ibmcdmusr

Password:

Comment: SAP HANA DB user

[Need help?](#)

[Create] [Discard]

Figure 4-18 Create the database credentials

6. Once the server has been registered an entry will appear under the Application Server Provider Browser. See Figure 4-19 on page 56.



Figure 4-19 Application Server Provider Browser

7. Go to the **Jobs** tab, select **SAP HANA**. Under the button **NEW** choose **Backup**. See Figure 4-20.

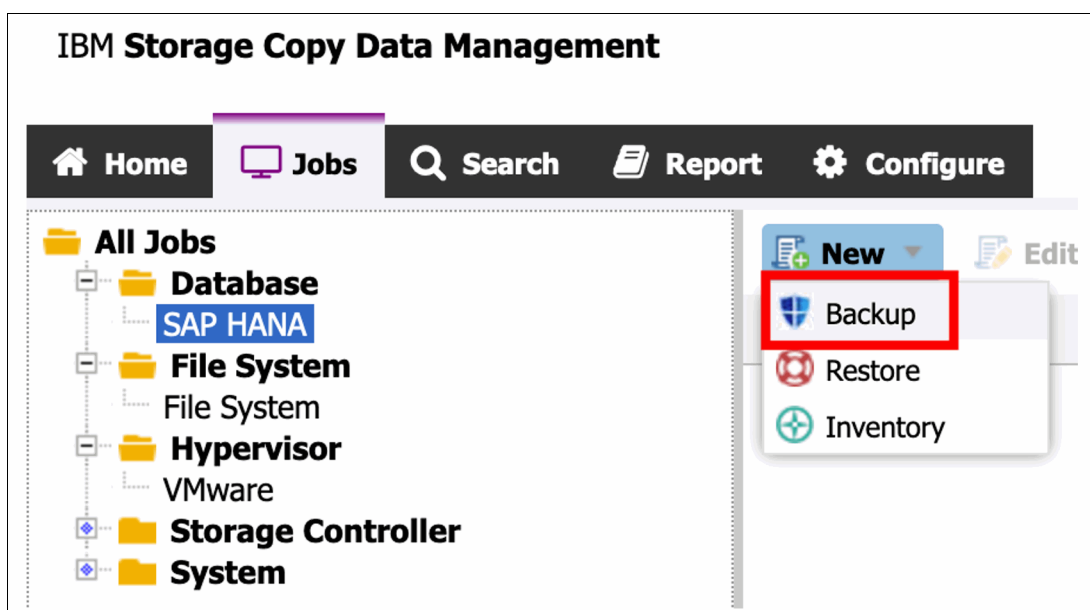


Figure 4-20 Choose Backup

8. Select the database to be protected. If the desired database is not shown on the list run the default SAP HANA application catalog job. Attach the SLA policy created and click the **Schedule Time** icon. See Figure 4-21.

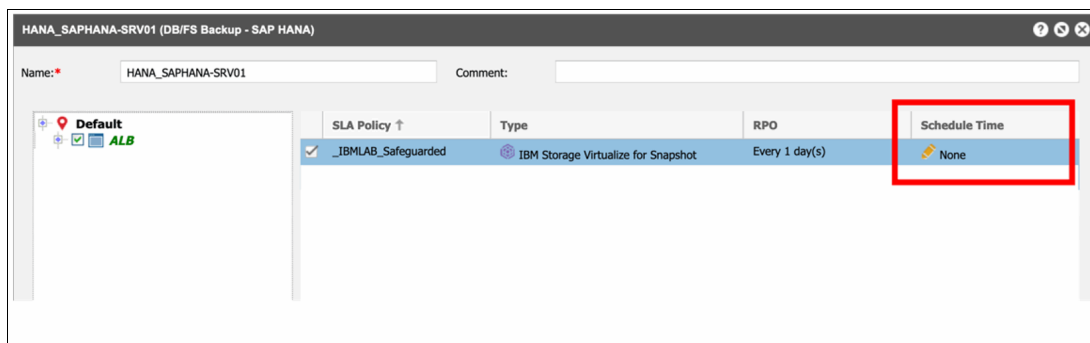


Figure 4-21 Choose the database to be protected and the SLA policy

9. Select **Enable Schedule** and choose the date and time. Once it is enabled the schedule will follow the RPO selected during the SLA Policy creation.

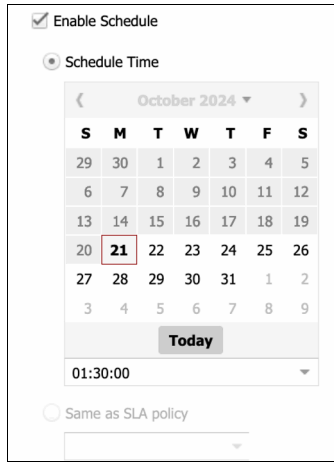


Figure 4-22 Enable Schedule

10. When the schedule is configured the column **Next Runtime** will be populated. See Figure 4-23.




Figure 4-23 Next Runtime

4.3 IBM Storage Sentinel

IBM Storage Sentinel is a powerful cyber-resilience tool that leverages machine learning to analyze snapshot backups for database integrity, including data structure, metadata, and content. Integrated with IBM CDM and IBM FlashSystem, Sentinel automates and simplifies data protection workflows.

4.3.1 Architecture

The cyber-resilience solution architecture offers flexibility to accommodate different SAP HANA deployment models, such as physical (IBM Power with NPIV and x86_64) and virtual (VMware). The deployment of Sentinel servers can be either physical or virtual, aligned with the SAP HANA infrastructure. Figure 4-24 on page 58 depicts a fundamental configuration where both SAP HANA and Sentinel servers are deployed physically, allowing for the scanning of physical storage LUNs.

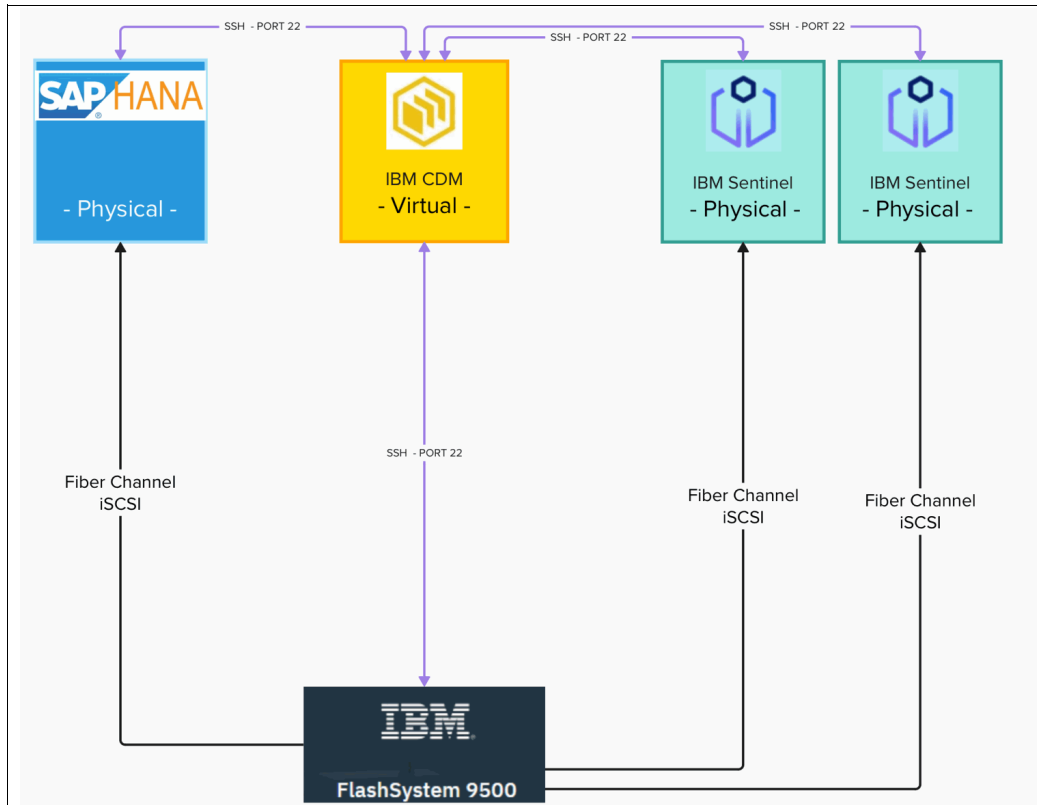


Figure 4-24 Both SAP HANA and Sentinel servers are deployed physically

The second scenario offers flexibility by accommodating both physical and virtual SAP HANA servers. To effectively conduct scan activities, Sentinel servers must be deployed as virtual machines. Disk mapping can be accomplished through pRDM devices for physical LUNs or by mounting datastores containing Safeguarded Snapshot disks. In either case, VMware vCenter integration is essential for managing LUN mapping at the hypervisor layer. See Figure 4-25 on page 59.

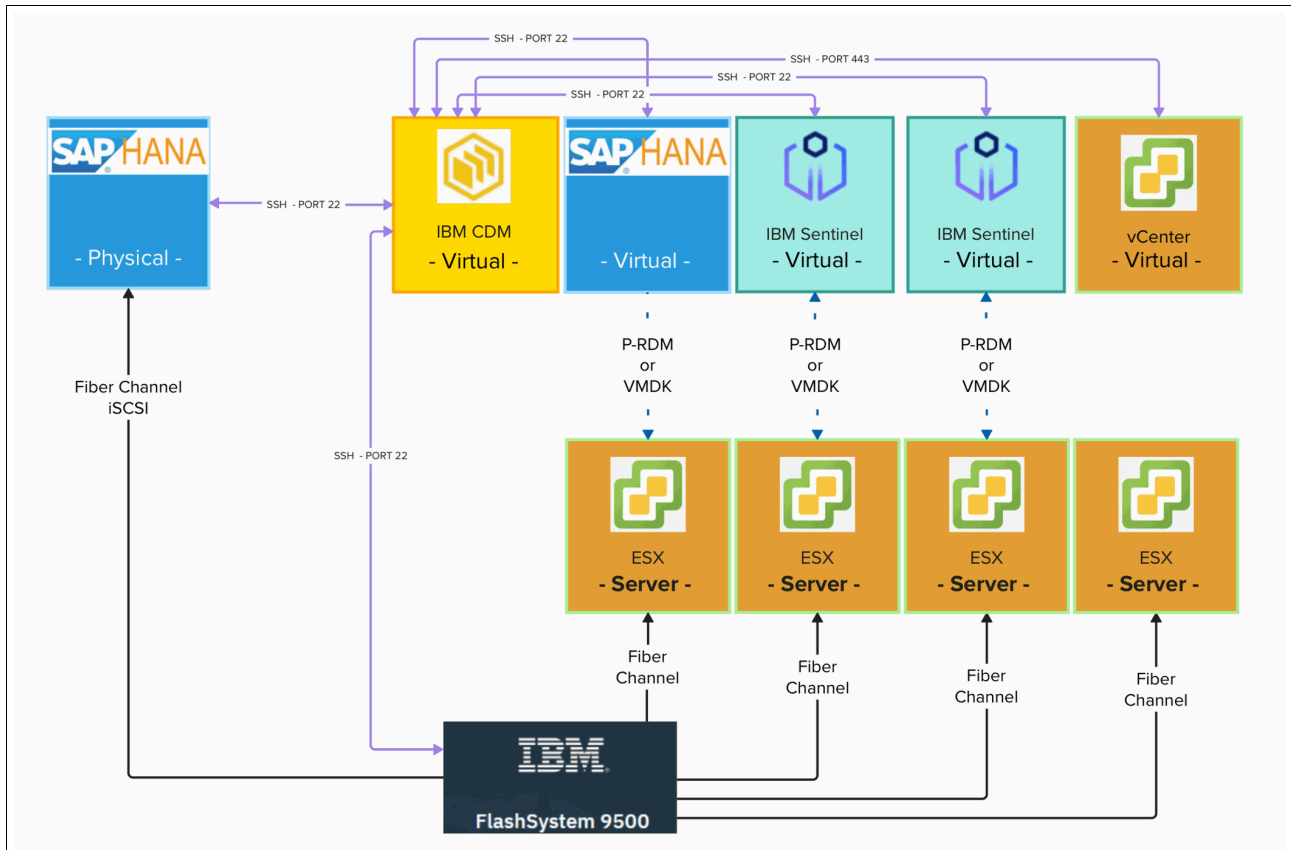


Figure 4-25 Architecture with physical and virtual SAP HANA servers

The third scenario leverages the iSCSI protocol for communication. Virtual IBM Sentinel servers are installed and registered in CDM as physical devices. Disk mapping is done directly to the virtual machines, eliminating the need to add vCenter to the CDM configuration. See Figure 4-26 on page 60.

Ensure you have a proper storage iSCSI configuration in place. To configure your IBM FlashSystem for iSCSI, follow these [guidelines](#).

Benefits of iSCSI deployment:

- ▶ **Reduced complexity:** This approach avoids the need for vCenter integration, simplifying the setup process.
- ▶ **Scalability and management:** Sentinel servers remain virtual, offering flexibility and ease of management.

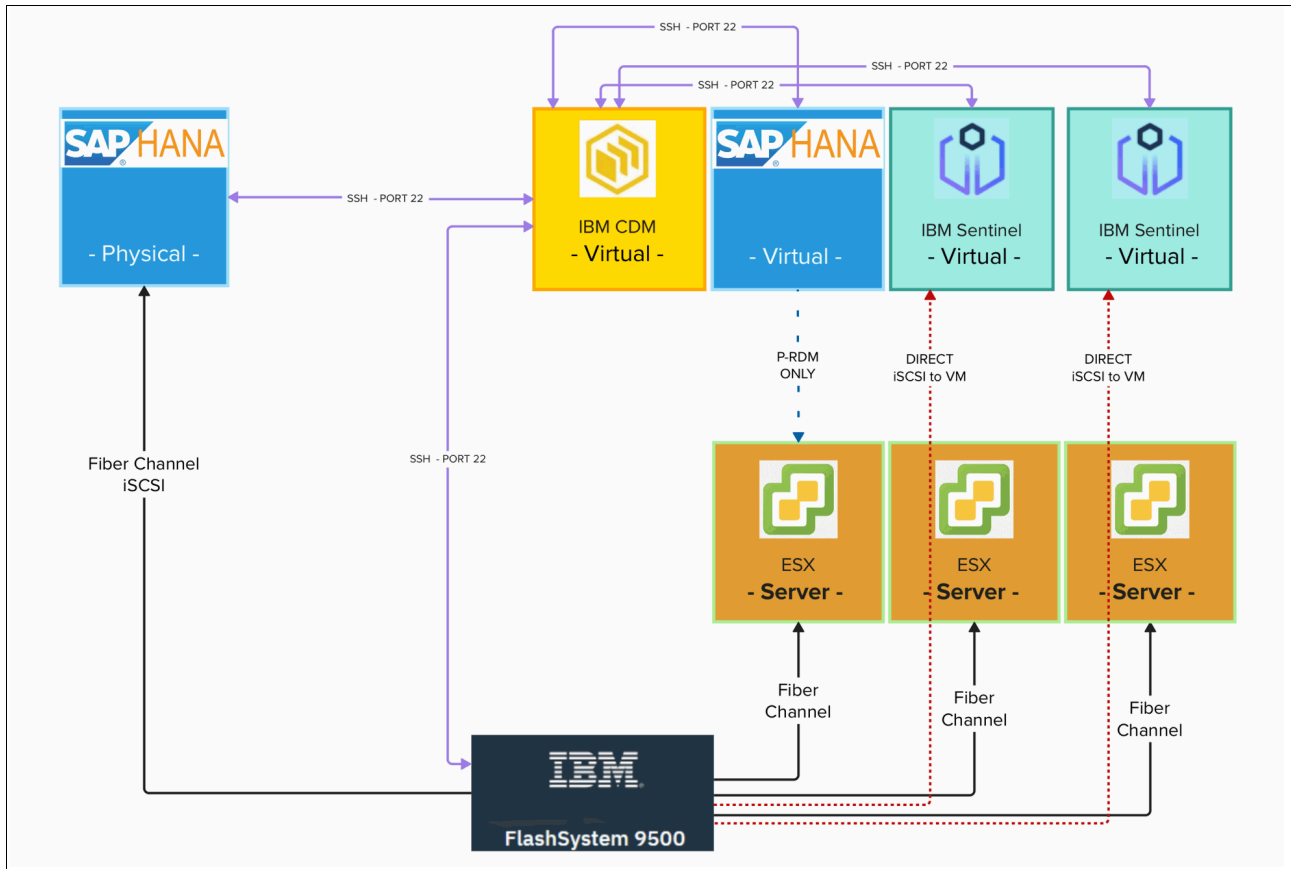


Figure 4-26 The third scenario with the iSCSI protocol for communication

4.3.2 Server requirements

IBM Storage Sentinel is compatible with the following operating systems on x86_64 architecture: SUSE Linux Enterprise Server (SLES) 15.5

For SUSE Linux Enterprise Server 15 SP4 (SLES/15.5/x86_64), the following module is required: Base System Module (`sle-module-basesystem/15.5/x86_64`).

Operating system packages or licensing are not part of IBM Storage Sentinel. Check the [compatibility page](#) for more details.

An internet connection is required to perform Sentinel installation. Make sure that the firewall service is disabled to avoid any failure during installation. See Example 4-8.

Example 4-8 `systemctl status firewalld`

```

Ibmsentinel1:/ # systemctl status firewalld
firewalld.service - firewalld - dynamic firewall daemon
Loaded: loaded (/usr/lib/systemd/system/firewalld.service; disabled; vendor
preset: disabled)
Active: inactive (dead)
Docs: man:firewalld(1)

```

To stop firewall service, issue `systemctl firewalld stop` command or `systemctl firewalld disable` to keep the firewall disabled across reboots.

Ports used for Sentinel are listed in Table 4-1.

Table 4-1 Ports used for Sentinel

Open Ports	To Allow	When
80 or 443	Inbound connection	To access the anomaly scan software
All appropriate NDMP, CIFS, and NFS ports	Outbound connection	To access the corresponding sources
22	Inbound and outbound connections on all engines	For SSH (Secure Shell) communications
25	Inbound connection	To enable remote command line access (recommended)
5432	Inbound connection on Federation Manager; outbound connection on each Federation Member	To participate in a federation
22, 7776, 7779, 7781, 7785, 7795, 7799, and 8476	Inbound and outbound connections on all engines	To participate in a federation

To configure firewall exceptions, use these set of commands on each Linux host. See Example 4-9.

Example 4-9 Configuring firewall exceptions

```
firewall-cmd --permanent --add-service=ssh --add-service=http --add-service=https
firewall-cmd --permanent --add-port=7776/tcp --add-port=7779/tcp
--add-port=7781/tcp
firewall-cmd --permanent --add-port=7785/tcp --add-port=7795/tcp
--add-port=7799/tcp
firewall-cmd --permanent --add-port=8476/tcp --add-port=5432/tcp
firewall-cmd --reload
firewall-cmd --list-all
```

Check if Mail Transfer Agent (MTA) is running. See Example 4-10.

Example 4-10 Check if Mail Transfer Agent (MTA) is running

```
ibmsentinel1:/ # mail root
Subject: test
.
EOT
```

IBM Sentinel servers can be deployed in both physical and virtual (VMware) environments. However, physical Sentinel servers cannot scan VMware virtual disks (VMDK).

When manually partitioning a Sentinel server, consider the following layout:

LVM logical volumes:

- ▶ /dev/sentinelvg/sentinel1v (700 GB): XFS formatted for /opt/1e

- ▶ /dev/swapg/swap1v (832 GB):

For swap space GPT partitioning:

- ▶ /dev/sda1 (512 MB): FAT32 formatted for /boot/efi
- ▶ /dev/sda2 (149.5 GB): LVM physical volume

LVM volume group rootvg:

- ▶ /dev/rootvg/root1v (119.5 GB): Btrfs formatted for /
- ▶ /dev/rootvg/var1v (15 GB): Btrfs formatted for /var
- ▶ /dev/rootvg/home1v (15 GB): Btrfs formatted for /home

Validate if [hardware and software requirements](#) are met.

Run the **getenforce** command to check if SELinux is in permissive mode:

The expected output should be "disabled". If it is not, use the following command to switch it to permissive mode: **sudo setenforce 0**. The results can be validated in /etc/selinux/config file.

Set the hostname in the /etc/hostname file.

Set the /etc/hosts/ for sentinel servers if no DNS is provided. See Example 4-11.

Example 4-11 /etc/hosts/

```
ibmsentinel1:~ # cat /etc/hosts
127.0.0.1localhost
172.26.18.16ibmsentinel1.ibm1ab.com ibmsentinel1
172.26.18.17ibmsentinel2.ibm1ab.com ibmsentinel2
```

4.3.3 Sentinel installation

IBM Storage Sentinel packages can be downloaded from:

- ▶ [IBM Passport Advantage® \(PAO\)](#): For initial installations.
- ▶ [IBM Fix Central](#): For upgrades and updates.

Note: Only authorized users with IBM-ID can download installation packages from IBM Passport Advantage.

Follow these [instructions](#) to install Sentinel.

4.3.4 Sentinel federation

A Sentinel federation is a group of two or more Sentinel servers, referred to as engines, that work collaboratively to perform indexing tasks. During configuration, one engine is designated as the manager, responsible for creating the index repository and managing licenses. The remaining engines act as members of the federation. See Figure 4-27 on page 63.

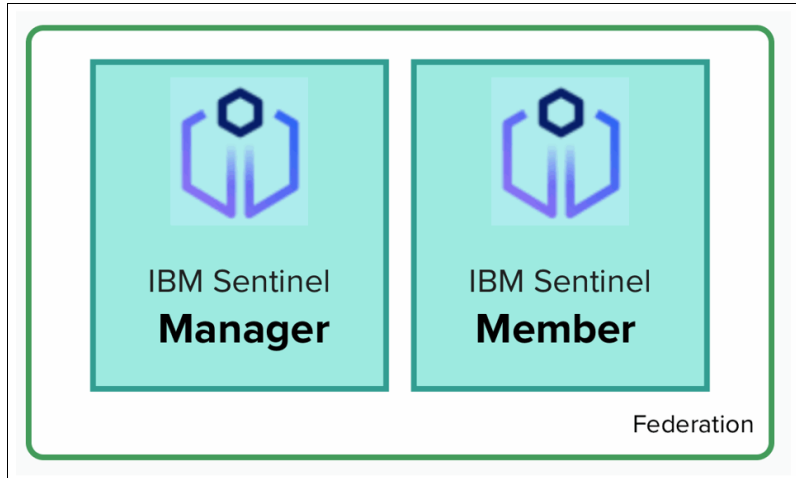


Figure 4-27 Sentinel federation

4.3.5 Configuring the manager engine - Sentinel federation integration

To configure the manager engine with the Sentinel federation, follow these steps:

1. After installation finishes open the URL `https://sentinel_node_FQDN/admin`. Use the same admin password set by the `iepasswd` command. See Figure 4-28.

The screenshot shows a 'Security Scan Engine Sign In' window. It contains a 'Username' field with 'admin' entered, a 'Password' field with masked characters, and a 'Domain' dropdown menu set to '(engine)'. There is a 'Remember me on this computer' checkbox and a 'Sign In' button.

Figure 4-28 Sign in panel

2. Accept the End User Licensing Agreement: **Administration** → **System** → **Licenses** → **EULA**.
3. Go to **Administration** → **System** → **Licenses** → **Setup License**. Take note of engine ID, it must be added to the IBM Storage Sentinel license fulfillment template. Only the Engine ID from manager engine is necessary to request the permanent license. See Figure 4-29.

The screenshot shows the 'Setup License' page in the administration console. The breadcrumb trail is 'Administration > System > Licenses'. There are tabs for 'EULA', 'Setup License', 'Download System Information', and 'Counter History'. The 'Upload License' radio button is selected. A red box highlights the 'EngineID 14:7d:da:b4:de:a9'. Below this, there is a 'License File' section with a 'Browse...' button and the text 'No file selected.' and 'Upload File' button. At the bottom, there is a summary of system statistics:

Active Tapes Under Management: 0	Active Backup Image Files Under Management: 0
Archived Tapes Under Management: 0	Archived Backup Image Files Under Management: 0
Total Tapes Under Management: 0	Total Backup Image Files Under Management: 0
Total Bytes Under Management: 10240750639324	

Figure 4-29 Setup License

- Submit the IBM Storage Sentinel license fulfillment template to the email described in the document. A response will be sent with the permanent license that must be added to the manager engine.
- To add the license, go to **Administration** → **System** → **Licenses** → **Setup License**. Select the **Upload License** radio button. Browse to the license file and upload it.
- Once the upload is complete, verify the installed licenses under **Administration** → **System** → **Licenses**.
- Log out of the manager engine.
- Sentinel is licensed based on the scanned capacity. A single license can be shared across multiple engines within a federation.
- To access the Sentinel monitoring interface, log in to https://sentinel_node_FQDN/sentinel using the same administrative credentials. This interface allows you to manage and monitor the Sentinel engine. See Figure 4-30.

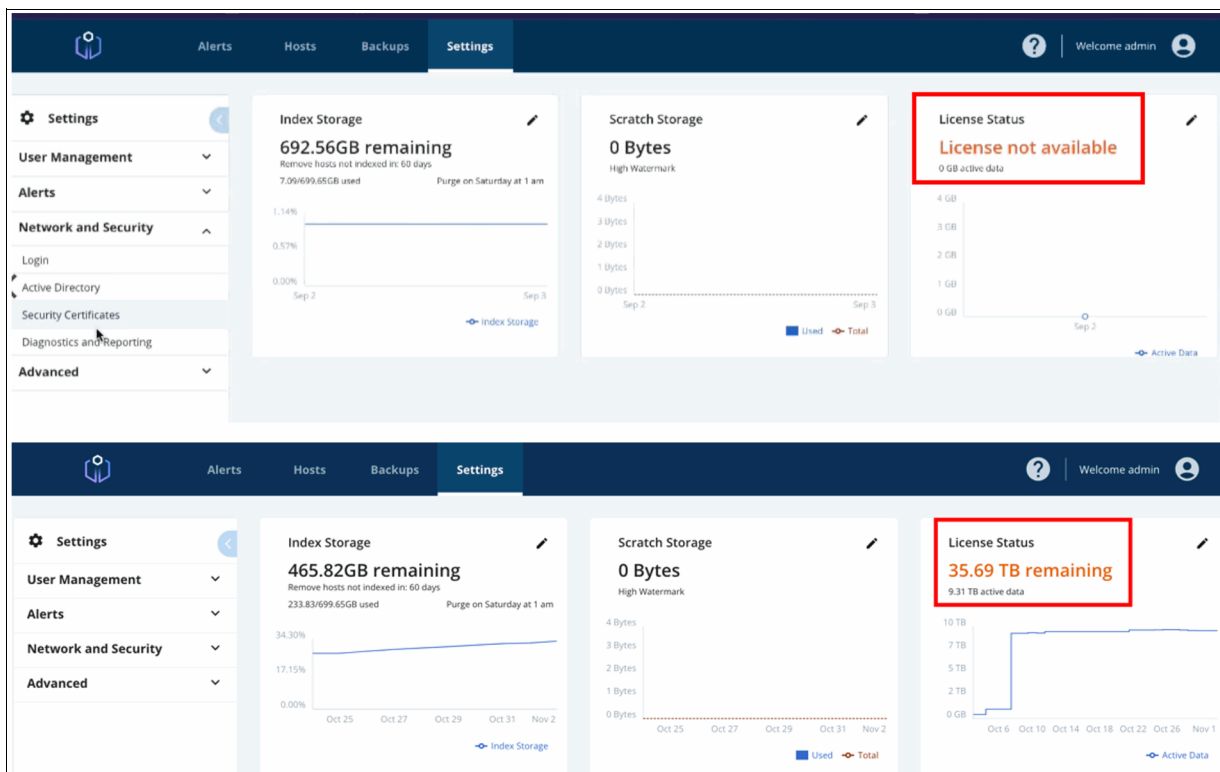


Figure 4-30 Sentinel monitoring interface

4.3.6 Adding member engines to the Sentinel federation

Perform the following steps to add member engines to the Sentinel federation.

- After installation finishes open the URL https://sentinel_node_FQDN/admin. Use the same admin password set by the `iepasswd` command. See Figure 4-31.

The screenshot shows the 'Security Scan Engine Sign In' panel. It contains the following fields and controls:

- Username: admin
- Password: [masked]
- Domain: (engine)
- Remember me on this computer
- Sign In button

Figure 4-31 Sign in panel

2. Accept the End User Licensing Agreement: **Administration** → **System** → **Licenses** → **EULA**.
3. No additional licenses are required for member nodes. Each Sentinel node includes a built-in manager engine, which automatically creates a default index. To remove this default index, navigate to **Administration** → **Home** → **Index Manager**, select the **default index**, and deselect it. See Figure 4-32.




Figure 4-32 Deselect default index

4. Go to **Administration** → **Home** → **Index Manager** → **Setup** and select **Join Federation**. See Figure 4-33.




Figure 4-33 Join Federation

5. When the operation completes, the manager engine name will be registered at the Index Management Setup box. See Figure 4-34.




Figure 4-34 Manager engine name is registered

4.4 Integrating IBM Storage CDM, IBM Storage FlashSystem and IBM Storage Sentinel

CDM is the central management application that automates all activities related to snapshot data protection, validation and recovery. These activities include integration with IBM FlashSystem and IBM Sentinel. Sentinel engines must be added to their credentials

1. In the **CDM configure** tab, browse to the Security Scan Server and select **Register**.

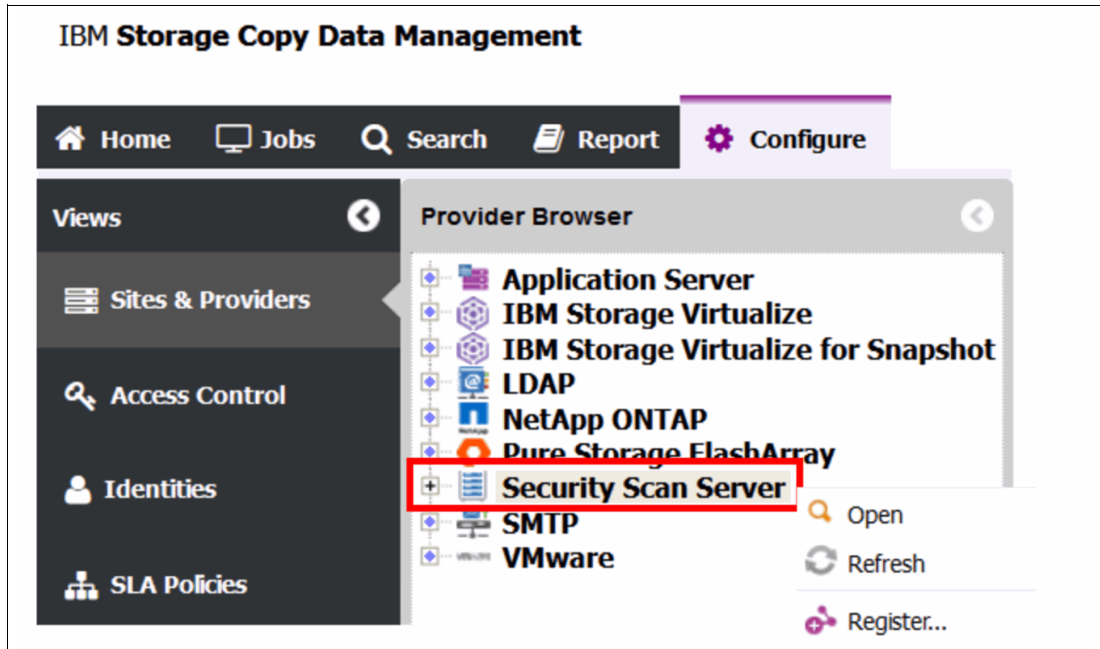


Figure 4-35 Select Register

Fill in the **Site**, **Name**, and **Host Address** fields with the appropriate information. Place the engine scan in the same site as the SAP HANA servers. See Figure 4-36 on page 67.

Tip: If Sentinel servers are VMs, their vCenter must be added to the CDM configuration beforehand.

Register Security Scan Server

Site: Default

Name: ibmsentinel1

Host Address: ibmsentinel1.ibm1ab.com

Type: Virtual Physical

vCenter: vCenter

Security Scan API Credential:

Select New

Name	Username	Type

System Credential:

Select New

Name	Username	Type

OK Cancel

Figure 4-36 Register Security Scan Server

- To configure the Security Scan API Credential, click **New** and input the Sentinel administrative username and password. These credentials are identical to those used to access the Sentinel monitoring GUI. See Figure 4-37.

Create Credential

Name: ibmsentinel1_admin

Username: admin

Password:

Comment: IBM sentinel administrator

[Need help?](#)

Create Discard

Figure 4-37 Security Scan API Credential

- In the System Credential, add the Sentinel server root user and password. See Figure 4-38 on page 68.

Create Credential

Name:

Username:

Password:

Comment:

[Need help?](#)

Figure 4-38 Create Credential

- Follow the same procedure to add any remaining Sentinel engines to CDM. This process automatically provisions the required indexes on each engine to store scan activity data. See Figure 4-39.

General	
Site	Default
Name	ibmsentinel1
Host Address	Ibmsentinel1.ibmlab.com
Catalog Eligible	no
Server Type	Virtual
vSphere Id	1001
Application Type	securityscanserver
Version	8.6.0-1.24
OS Type	linux
Use Key Authentication	no
Fed Id	5d761b7e-f720-4882-a6f6-3b57b6fff110
Fed Manager	
Fed Node Type	MANAGER
Index Name/Id	E9f9df52-66ce-11ef-9fec-0050568b6564

Figure 4-39 required indexes are provisioned on each engine

- The information is available in the Index Manager section of the Sentinel administration console for any engine in the federation (**Administration** → **Home** → **Index Manager**.) See Figure 4-40.

Administration > Home > Index Manager : [Indexes](#) | [Select Index](#) | [Add Index](#) | [Joined Engines](#) | [Setup](#) | [Locally Managed Indexes](#)

Index Manager

Role: Member

Engines: Joined: 2

Federation Manager: Host: ibmsentinel1
Engine ID: 14:7d:da:b4:de:a9

Indexes: Active: 3
Total: 3

Selected: Index Name: e9f9df52-66ce-11ef-9fec-0050568b6564
Id: 20

Figure 4-40 Index Manager section of the Sentinel administration console

- In CDM, the added nodes will appear under the **Security Scan Server** tab. See Figure 4-41.

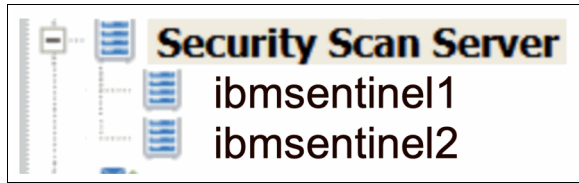


Figure 4-41 Security Scan Server tab

- The subsequent step involves configuring security scan activities within SLA policies. This can be achieved by editing existing or creating new SLA policies. Within the SLA tab, enable the **Perform Security Scan every** option and define the desired scan frequency, such as after each Safeguarded Snapshot or a specified number of backups. See Figure 4-42.

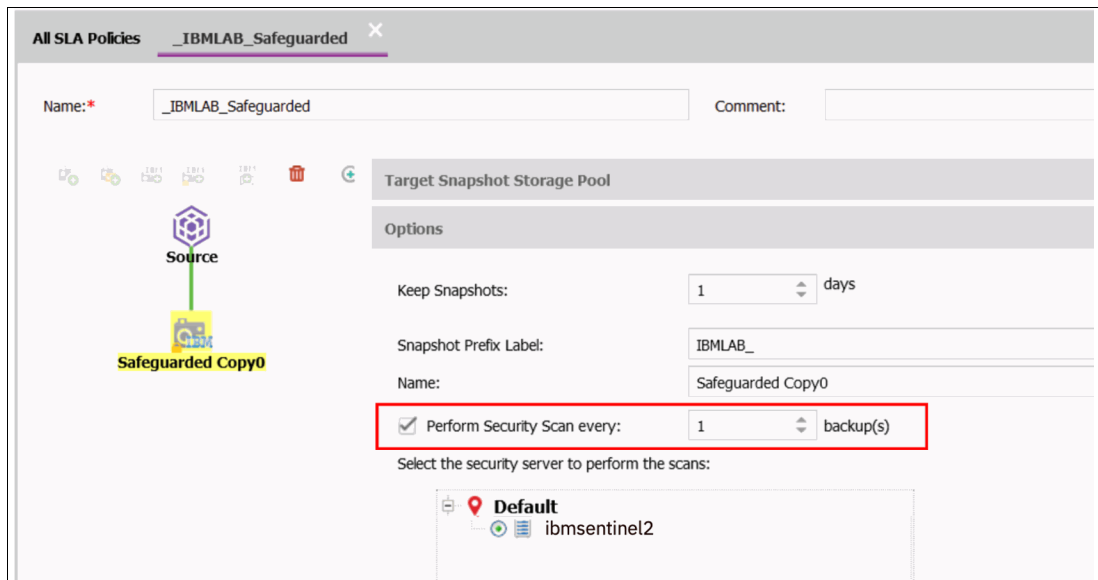


Figure 4-42 Perform Security Scan every option

- Security scans will run after each Safeguarded Snapshot for those SAP HANA servers. You will see the results in the **Security Scan Status** and **Security Scan Message** columns. See Figure 4-43.

Start Time ↓	End Time	Duration	Comment	Security Scan Status	Security Scan Message	Status
Oct 31 02:30:01 2024	Oct 31 02:57:55 2024	0h 27m 54s	Safeguarded Copy(...)	Done	No threats detected.	COMPLETED
Oct 30 16:46:19 2024	Oct 30 17:08:22 2024	0h 22m 2s	Safeguarded Copy(...)	Done	No threats detected.	COMPLETED

Figure 4-43 Security scans will run after each Safeguarded Snapshot

- Before proceeding with the configuration, it is essential to establish appropriate SAN zoning or iSCSI connectivity between IBM FlashSystem, IBM Storage Sentinel servers, and host systems. See Figure 4-44 on page 70.

Name	Status	Host Type	# of Po...	# of iSCSI or iSE...	Host Mappings	Host Cluster ID	Host Cluster Name	Protocol Type
ibmsentinel1	✓ Online	Generic	8	0	Yes			SCSI
ibmsentinel2	✓ Online	Generic	8	0	Yes			SCSI

Figure 4-44 iSCSI connectivity between IBM FlashSystem, IBM Storage Sentinel servers, and host systems

10. In the event that Sentinel servers are deployed as virtual machines on VMware ESXi, it is necessary to establish SAN zoning and host connections for the ESXi hosts. See Figure 4-45.

Name	Status	Host Type	# of Po...	# of iSCSI or iSE...	Host Mappings	Host Cluster ID	Host Cluster Name	Protocol
esxibmlab01	✓ Online	Generic	8	0	Yes			SCSI
esxibmlab02	✓ Online	Generic	8	0	Yes			SCSI

Figure 4-45 SAN zoning and host connections for the ESXi hosts

4.4.1 IBM FlashSystem best-practices for administrative users

Proper storage user-policy configuration is crucial for cybersecurity. IBM storage systems, including those using Safeguarded Snapshots, benefit from security tools and processes proven to protect against unauthorized access.

Note: As of this writing, Storage Virtualize 8.7 is not certified with the latest CDM releases (2.2.24). Consider using Storage Virtualize 8.6 for compatibility with CDM.

Dedicated user and group for Cloud Data Management (CDM)

Create a unique user group named CDM_SGC_GRP with the Administrator role using `mkusergrp`.

Disable GUI and Storage API access for this group (`-disablegui yes & -disableapi yes`). This restricts CDM to secure CLI commands.

Assign administrative privileges to this user. It is important to note that storage administrators do not possess the authority to delete Safeguarded Snapshots, as their expiration is governed by the retention period defined in Safeguarded Snapshot policies.

Tip: Separating CDM access with a dedicated user group minimizes the attack surface should unauthorized access occur. Disabling GUI and API access further restricts potential misuse. You can use the command in Example 4-12.

Example 4-12 `mkusergrp` command

```
mkusergrp -name CDM_SGC_GRP -role Administrator -multifactor no
-passswordkeyrequired yes -disablegui yes -disableapi yes
Modifying the authentication setting for this user group will affect logins for
all users in the group.
Are you sure you want to continue? (y/yes to confirm) yes
User Group, id [6], successfully created
```

Establish a privileged security administrator account to manage other administrators. Ensure this account is securely stored.

Lockout policy

Establish security policies that enforce login failure restrictions and lockout periods for administrator accounts.

Use `svctask chsecurity` to configure:

- ▶ `-maxfailedlogins 5`: Maximum failed login attempts before lockout.
- ▶ `-lockoutperiod 30`: Duration (minutes) a user is locked out after failed attempts.

This prevents brute-force attacks by locking out users after exceeding the allowed failed login attempts.

Multi-Factor Authentication (MFA)

Enable MFA for all administrative users except the CDM user using `svctask chsecurity`. Keep in mind that MFA must be integrated to IBM Security Verify or Duo Security. Instructions are provided in [IBM Documentation](#).

MFA adds an extra layer of security by requiring a secondary verification code after entering the password.

Superuser lockdown

Upon completion of the installation process, consider locking the superuser account. Alternatively, enable multi-factor authentication (MFA) for this account.

Use the following command if you want to maintain the ability to log in as the superuser but add an extra layer of security with MFA: `svctask chsecurity -superuserlocking enable`.

Use the following command if you want to completely disable the superuser account, because you have other administrative users with sufficient privileges or you no longer need the superuser account: `svctask chuser -lock superuser`

Tip: Locking or enabling MFA on the superuser account is crucial as it has unrestricted access.

Strong password policy

Use `svctask chsecurity` to configure a strong password policy. The following settings enforce complex and unique passwords for enhanced security:

- ▶ `-minpasswordlength 15`: Minimum password length (consider 15 characters minimum).
- ▶ `-passwordspecialchars 3`: Require at least 3 special characters.
- ▶ `-passworduppercase 3`: Require at least 3 uppercase characters.
- ▶ `-passwordlowercase 3`: Require at least 3 lowercase characters.
- ▶ `-passworddigits 3`: Require at least 3 digits.
- ▶ `-checkpasswordhistory yes`: Prevent password reuse.
- ▶ `-maxpasswordhistory 3`: Remember the last 3 used passwords.
- ▶ `-passwordexpiry 365`: Password expiry in days (consider regular password changes).
- ▶ `-expirewarning 30`: Days before receiving a password expiry notification.

Audit logging

Utilize [IBM FlashSystem Audit logs](#) to track and analyze CDM activities. These logs provide valuable information for security audits, including:

- ▶ User access details (date/time).

- ▶ User issuing the command.
- ▶ Command syntax executed.
- ▶ IP address and hostname where the command originated.

See Figure 4-46.

Date and Time	User Name	Command
11/3/2024 11:46:22 PM	cdmadmin	svctask rmvdisk -force 475
11/3/2024 11:46:22 PM	cdmadmin	svctask rmvolume group 117
11/3/2024 11:46:21 PM	cdmadmin	svctask rmvdisk -force 474
11/3/2024 11:46:21 PM	cdmadmin	svctask rmvdisk -force 467
11/3/2024 11:46:20 PM	cdmadmin	svctask rmvdisk -force 457

Figure 4-46 Audit log showing IP address and hostname where the command originated

The same user information can be found in the CDM logs `/opt/ECX/virgo/serviceability/logs/log.log` file.

4.5 IBM Security Guardium for SAP HANA

IBM Security Guardium is the flagship offering from IBM for Data Security. It automatically discovers and classifies sensitive data from across the enterprise, providing real-time data activity monitoring and advanced user behavior analytics to help discover unusual activity around sensitive data.

4.5.1 Data security concerns for SAP HANA environments

There are several data security concerns for SAP HANA environments:

- ▶ **Dispersed data:** Sensitive information may be stored in hundreds of different database columns, making it extremely difficult to conduct column-level monitoring or encryption.
- ▶ **Unauthorized privileged user access:** Insiders with privileged access can potentially view, change or exfiltrate SAP-stored data outside of business policy, without their actions being tracked and detected.
- ▶ **Data security blind spots:** Data access behavior and activities change over time, so there needs to be an easy way to distinguish between normal business-as-usual activities versus malicious activity
- ▶ **Database vulnerabilities:** Vulnerabilities such as missing patches, weak passwords, unauthorized changes and misconfigured privileges put SAP data stores at risk

Manual data management processes often lead to:

- ▶ **Inaccurate and inconsistent data:** For example, spreadsheets can be prone to human error, resulting in data inconsistencies and inaccuracies.
- ▶ **Poor data visibility:** Lack of a centralized system makes it difficult to track data locations, ownership, and access rights.

- ▶ Inefficient workflows: Email-based communication and manual processes can slow down data management tasks and increase the risk of errors.
- ▶ Compliance risks: Without a robust data management framework, organizations may struggle to meet data security and compliance requirements.

By addressing these challenges, organizations can improve data quality, enhance operational efficiency, and mitigate security risks.

4.5.2 Data security controls for SAP HANA

Guardium for SAP HANA (Systems, Applications, and Products - High Performance Analytic Appliance) is an accelerator provided by Guardium specifically for meeting the data security controls for SAP environment. Like the measures provided to other data sources, Guardium for SAP provides the following features:

- ▶ Database discovery.
- ▶ Data discovery and classification.
- ▶ Vulnerability assessment.
- ▶ Database activity monitoring.
- ▶ Realtime alerts.
- ▶ Data redaction.
- ▶ User quarantine and blocking.
- ▶ Various out-of-the-box reports related to database (DB) activities, Sessions, Object access, client/server details and so forth.
- ▶ Auditing workflows to track and fix suspicious activities.
- ▶ Integration with SIEM (Security information and event management) and ticketing tools such as IBM QRadar, ServiceNow.

While the solution is primarily designed for SAP HANA databases, it can also be seamlessly applied to other SAP enterprise applications using custom databases like Oracle, ensuring consistent protection across the entire IT landscape.

Out-of-the-box templates for discovery, vulnerability assessment, data protection policies, and reports can serve as a solid foundation for evaluating data protection controls in both pre-production and production environments.

Guardium's non-intrusive solution ensures minimal impact on production environments. By installing an agent on the database host and offloading the heavy security analysis to a separate Guardium Appliance, it maintains minimal overhead (typically 2-3%) and allows uninterrupted business operations.

Figure 4-47 on page 74 shows the data security proactive approach provided by IBM Security.

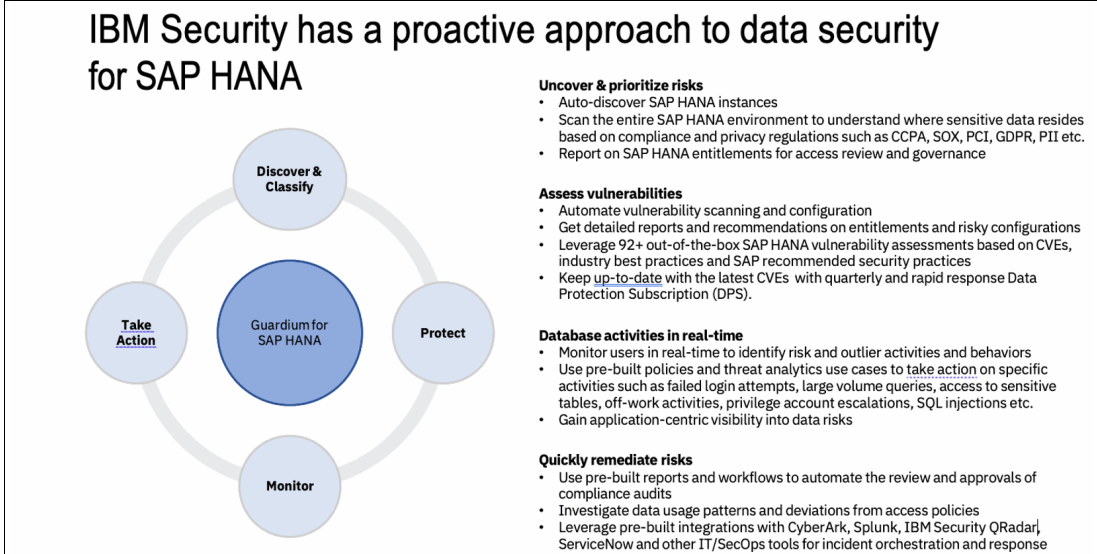


Figure 4-47 Data security proactive approach provided by IBM Security

Figure 4-48 shows the deployment options for using IBM Guardium and SAP-HANA.

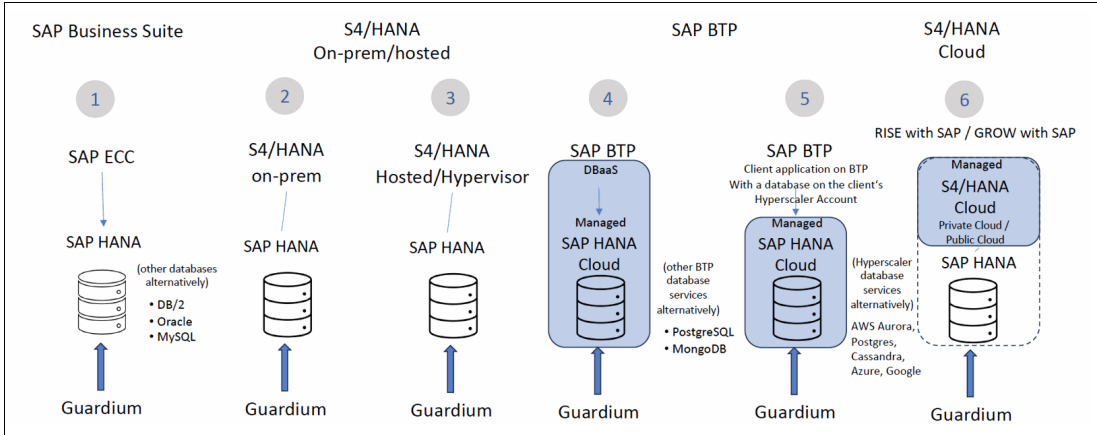


Figure 4-48 Deployment options for using IBM Guardium and SAP-HANA

4.5.3 Demo: IBM Guardium and QRadar Preventing Attackers with Safeguarded Copy and Copy Services Manager

This demo showcases three real-world examples of how IBM Guardium can effectively thwart attacks on Oracle Databases protected by IBM FlashSystem Safeguarded Copy immutable snapshots. Additionally, QRadar can be configured to automatically trigger additional safeguards, including creating extra immutable snapshots, in response to suspicious activity. By prioritizing prevention over remediation, this solution empowers the IT ecosystem to proactively identify and respond to internal threats, mitigating potential outages and ransomware attacks.

This proactive security framework for Oracle Databases leverages three key tools: Guardium, QRadar, and CSM. Guardium and QRadar work in tandem to monitor user access and data integrity.

When a user attempts to log into the Oracle Database, the system verifies the user's

authorization and assesses the sensitivity of the data they are accessing to determine the appropriate level of monitoring.

The system continuously monitors all database activities in real-time through Guardium, while QRadar oversees security events and potential threats. Two user accounts, Reninder and Andrew, are observed as they log into the database, with their actions closely tracked.

In the background, IBM Flash Storage utilizes Copy Services Manager to create safeguarded database copies for recovery purposes. If Guardium detects a significant threat, it can intervene to stop the unauthorized action and trigger QRadar to assess the situation. QRadar may then initiate the creation of additional safeguarded copies to safeguard the database.

By combining these technologies, organizations can proactively respond to security threats, safeguarding the integrity and availability of sensitive Oracle Database data.



Configuring IBM Storage Sentinel for VMware

This chapter describes how to configure IBM Storage Sentinel to scan safeguarded snapshots containing virtual machine data running in a VMware environment.

This chapter has the following sections:

- ▶ “VMware configuration on Copy Data Management” on page 78
- ▶ “Safeguarded Snapshot of critical VMs” on page 81
- ▶ “Scanning process” on page 82
- ▶ “Monitoring” on page 83
- ▶ “Restore and recovery” on page 85

5.1 VMware configuration on Copy Data Management

To protect critical VMs with Safeguarded Snapshots and scan those copies with IBM Storage Sentinel, the first step is to register the vCenter in the IBM Storage Copy Data Management (SCDM). Once vCenter is successfully registered, SCDM initiates the inventory process automatically, scanning the environment to discover ESXi hosts and virtual machines (VMs) under the registered vCenter.

5.1.1 Registering VMware vCenter

Complete the following steps to register the VMware vCenter:

1. Click the **Configure** tab. In the **Views** window, select **Sites & Providers** and then select the **Providers** tab.
2. In the Provider browser window, select **VMware**.
3. Right-click **VMware**. Then, click **Register**. The Register dialog window opens.
4. Complete the fields in the dialog window. Select **New** to add credentials if they are not yet added through identities. See Figure 5-1.

For the required VMware vSphere privileges, see [VMware vSphere Privileges](#).

Name	Username	Type	
vcadm	administrator@vsph...	System	

Figure 5-1 Registering VMware vCenter

5.1.2 Registering Storage Sentinel Security Scan Server

Complete the following steps to register the Storage Sentinel security scan server:

1. Click the **Configure** tab. In the **Views** window, select **Sites & Providers** and then select the **Providers** tab.
2. In the Provider Browser window, select **Security Scan Server**.
3. Right-click **Security Scan Server**. Then, click **Register**. The Register dialog window opens.
4. Complete the fields in the dialog window. Select **New** to add credentials if they are not yet added through identities. If the security scan server is a virtual machine on VMware, select the pre-registered VMware vCenter and enter the credentials for the Sentinel server. See Figure 5-2.

The screenshot shows the 'Register Security Scan Server' dialog box. It contains the following fields and options:

- Site: Default
- Name: Sentinel
- Host Address: 192.168.152.100
- Type: Virtual, Physical
- vCenter: PentiaVCenter
- Security Scan API Credential: Select, New
- System Credential: Select, New

There are two empty tables below the credential sections, each with columns for Name, Username, and Type. At the bottom right are 'OK' and 'Cancel' buttons.

Figure 5-2 Registering Sentinel Security Scan Server

During the registration process, SCDM creates a new index on the scan server and makes it the active index for corresponding scans. See Figure 5-3

The screenshot shows the 'Indexes' window with the following data:

Index	Engines	Actions
Index Name: Index-VmLab Created: Sep-26 at 6:14 pm State: Deactivated Tapes Under Management: 0 Backup Image Files Under Management: 0 Bytes Under Management: 0 Id: 1		Activate Remove Migrate... Edit
Index Name: SCDM_2d0801f0-7ced-11ef-9f32-005056bab7e3 Created: Sep-27 at 7:25 pm Tapes Under Management: 0 Backup Image Files Under Management: 0 Bytes Under Management: 31823014126 Id: 5 Active Tapes Under Management: 0 Active Backup Image Files Under Management: 0 Archived Tapes Under Management: 0 Archived Backup Image Files Under Management: 0 Total Tapes Under Management: 0 Total Backup Image Files Under Management: 0 Total Bytes Under Management: 31823014126	Host: sentinel, Engine ID: 00:50:56:BA:B7:E3, Version: 8.7.0-1.16	Deselect Edit...

At the top right, there are 'Index States' filters: Active, Deactivated, Removed, and an 'Apply Filter' button. At the bottom right is an 'Add New Index...' button.

Figure 5-3 Sentinel Indexes

Important: Ensure that an iSCSI host is created for the Sentinel security scan server on IBM Storage Virtualize which can be used to map snapshots to the scan server. IBM Storage Sentinel uses the iSCSI protocol to access data, and the bandwidth of the iSCSI connection directly affects the performance and duration of the scanning process. For iSCSI host creation steps, see “Linux iSCSI host attachment” in [IBM Documentation](#).

5.1.3 Configuring SLA policies

The use of SLA policies allows for the customization of templates by administrators for the primary processes that are involved in the creation and use of Backup jobs. These policies configure copy types, destinations, and parameters that can be reused in future Backup jobs.

During the configuration of a Backup job, suitable SLA policies appear in the job creation wizard. The listed policies are tailored to the specific type of Backup job being created.

1. Click the **Configure** tab. On the **Views** window, select **SLA Policies**. The All SLA Policies page opens.
2. In the All SLA Policies page, click **New**. The New SLA Policies page opens.
3. Select a type of policy to create based on your storage provider. Select **IBM Storage Virtualize** to create an IBM Backup policy.
4. Add a sub-policy (SLA Policy) to an IBM Storage Virtualize SLA policy.
 - a) Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the Frequency field, select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The smallest interval available is five minutes.

Note: If changes are made to the frequency and interval of an SLA Policy, those modifications impact all job schedules that are linked to it.

- b) Click **Add Safeguarded Copy**.
- c) In the Associated Safeguarded Volume Group page, expand the storage device and select the volume group that you want to back up by using Safeguarded Copy. Any volume that you want to back up by using Safeguarded Copy must belong to a volume group. If it is not a member of any of these groups, then it is not backed up as a Safeguarded Copy.

Note: The Associated Safeguarded Volume Group lists only those volume groups that have the Safeguarded Copy policy applied on the storage array side.

- d) In the Options page, set the **Safeguarded Copy** sub-policy options.

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the Days field.

Name

Enter an optional name to replace the default FlashCopy sub-policy name displayed in IBM Storage Copy Data Management. The default name is Safeguarded Copy0.

FlashCopy Volume Prefix

Enter an optional label to identify the FlashCopy. This label is added as a prefix to the FlashCopy volumes name created by the job.

Tip: FlashCopy labels must contain only alphanumeric characters and underscores.

Perform Security Scan

Enable the scan and select your security scan servers, so you can scan for every backup number that you have specified.

e) Enter a name for the new sub-policy (SLA Policy).

f) Click **Finish**. See Figure 5-4.

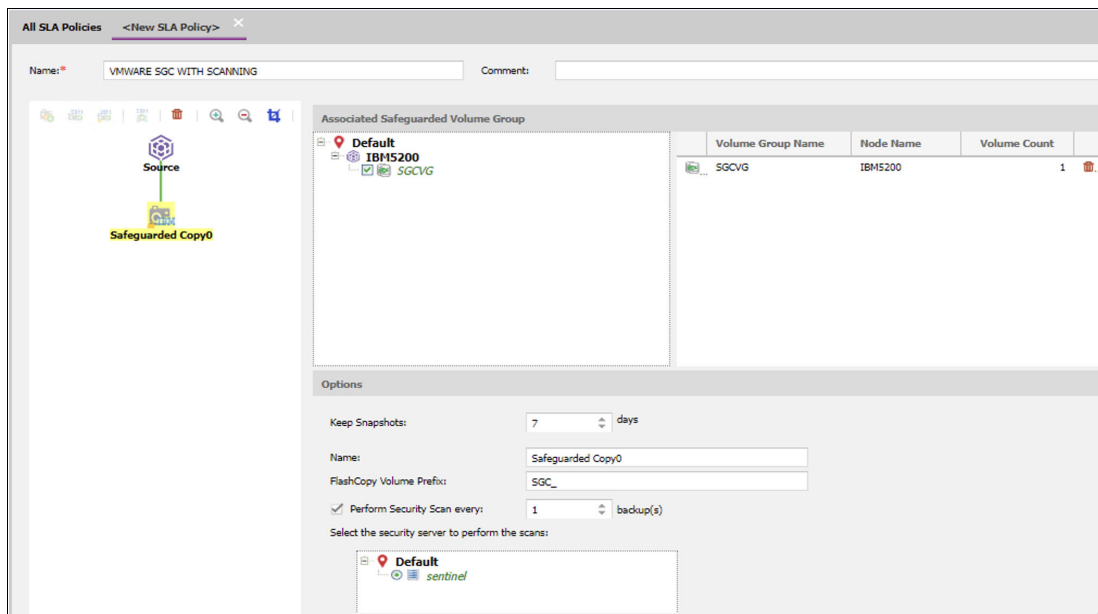


Figure 5-4 Configuring SLA policy and Storage Sentinel Scan frequency

Note: When setting the interval for the safeguarded snapshots with scanning, make sure that the scanning process will be finished within the specified interval, and set an interval accordingly.

5.1.4 Safeguarded Snapshot of critical VMs

The backup jobs create Safeguarded Snapshot of your selected virtual machines and dependent volumes, according to rules defined in an SLA policy. See Figure 5-5 on page 82.

1. Click the **Jobs** tab. Expand the **Hypervisor** and click on **VMware**.
2. Click **New** and select **Backup**.

3. Give this job a meaningful name and comment and select previously defined SLA and set schedule for the job.

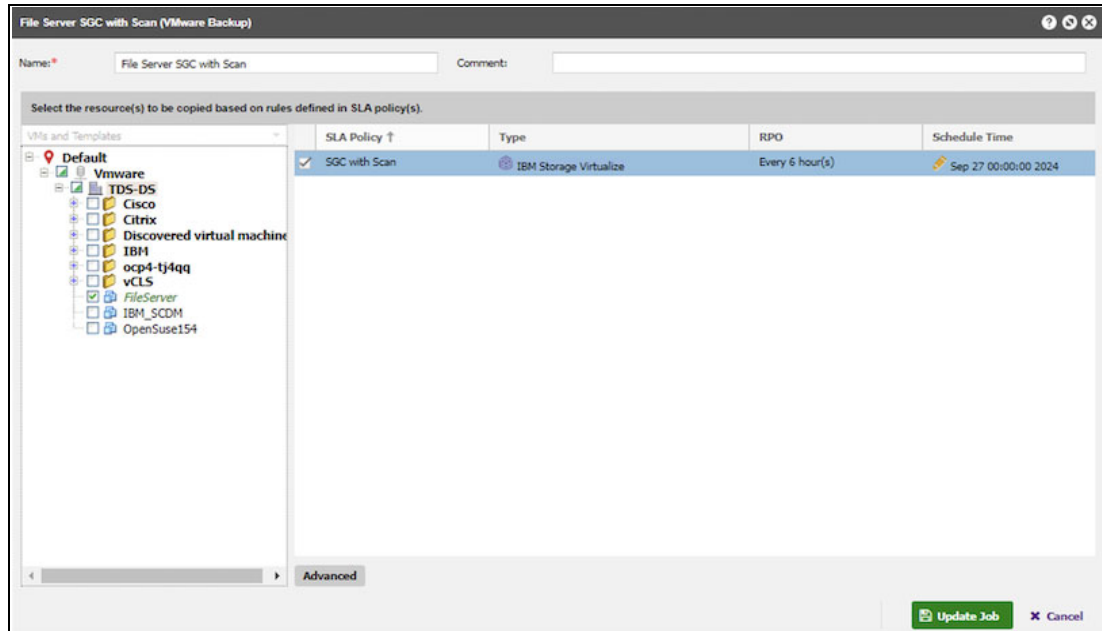


Figure 5-5 Defining a backup

5.2 Scanning process

Once you have created the backup job and its schedule, safeguarded snapshots run at the scheduled time and the scanning process automatically runs as defined in the SLA policy. IBM Storage Virtualize creates a copy of the protected snapshot and maps it to the scanning engine. The scanning engine initiates full indexing and scanning.

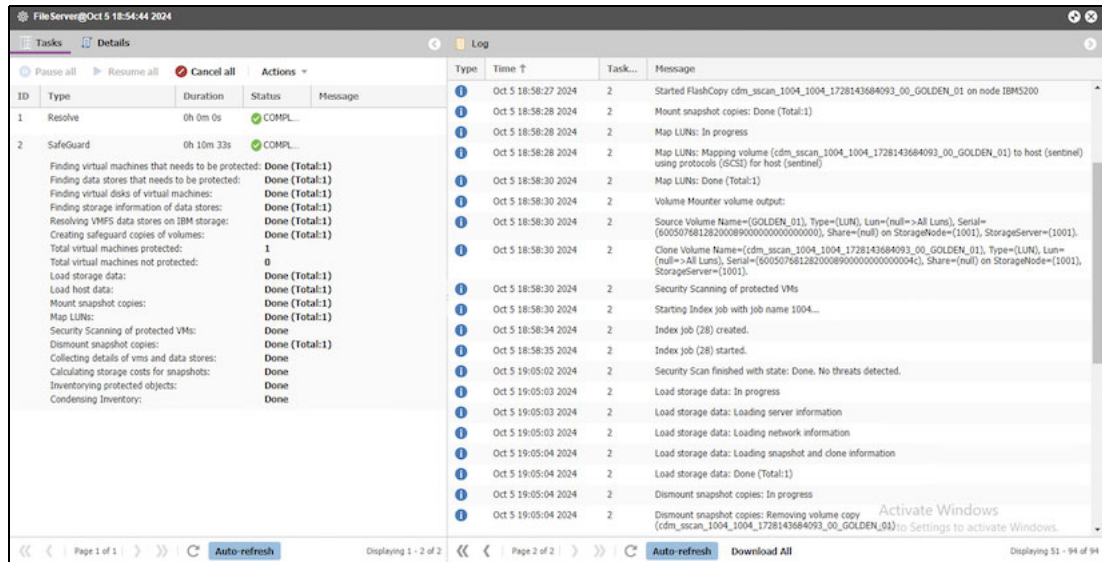


Figure 5-6 Job details on CDM

5.3 Monitoring

The SCDM job log can be considered the first step in monitoring sentinel scans. Successful completion of the job means that the safeguarded snapshot for the relevant VM was properly created, the created snapshot was mounted to the sentinel, the scan process was successful, no threat was found in the relevant copy as a result of the scan process, and the entire process was successful. To check backup jobs, See Figure 5-7.

1. Click the **Jobs** tab. Expand the **Hypervisor** and click **VMware**.
2. Click on your job and click **History** to see the latest job status.
3. Click a job to see detailed logs about this particular job.

The screenshot shows the 'Jobs' tab in the IBM Storage Sentinel interface. The left sidebar shows a tree view with 'All Jobs' expanded to 'Hypervisor' and then 'VMware'. The main area displays a table of jobs:

Name	Type	Status	Next Runtime	Last Runtime	Last Run Duration	Last Run Status
FileServer	VMware Backup	IDLE	Oct 5 22:00:00 2024	Oct 5 20:00:00 2024	0h 14m 6s	COMPLETED
Default VMware Inventory	VMware Inventory	IDLE	Oct 6 11:00:00 2024	Oct 5 11:00:00 2024	0h 0m 10s	COMPLETED
FileServer_Restore	VMware Restore	IDLE		Sep 27 12:44:44 2024	0h 2m 44s	COMPLETED

Below the jobs table, the 'History' tab is selected, showing a detailed log of completed jobs:

Start Time	End Time	Duration	Comment	Security Scan Status	Security Scan Message	Status
Oct 5 20:00:00 2024	Oct 5 20:14:07 2024	0h 14m 6s	Safeguarded Copy(SGC_)	Done	No threats detected.	COMPLETED
Oct 5 18:54:44 2024	Oct 5 19:05:18 2024	0h 10m 33s	Safeguarded Copy(SGC_)	Done	No threats detected.	COMPLETED
Oct 5 18:29:12 2024	Oct 5 18:39:56 2024	0h 10m 43s	Safeguarded Copy(SGC_)	Done	No threats detected.	COMPLETED
Oct 5 18:00:00 2024	Oct 5 18:10:52 2024	0h 10m 52s	Safeguarded Copy(SGC_)	Done	No threats detected.	COMPLETED
Oct 5 16:00:00 2024	Oct 5 16:10:21 2024	0h 10m 21s	Safeguarded Copy(SGC_)	Done	No threats detected.	COMPLETED
Oct 5 14:00:00 2024	Oct 5 14:10:31 2024	0h 10m 30s	Safeguarded Copy(SGC_)	Done	No threats detected.	COMPLETED
Oct 5 12:00:00 2024	Oct 5 12:10:18 2024	0h 10m 17s	Safeguarded Copy(SGC_)	Done	No threats detected.	COMPLETED
Oct 5 10:00:00 2024	Oct 5 10:10:22 2024	0h 10m 21s	Safeguarded Copy(SGC_)	Done	No threats detected.	COMPLETED

Figure 5-7 VMware job history

11. The scan jobs that are orchestrated by SCDM can be monitored in the Backups section of the Sentinel web interface at <https://<hostname>/sentinel>. See Figure 5-8.

The screenshot shows the 'Backups' tab in the IBM Storage Sentinel interface. The main area displays a table of scan jobs:

STATUS	HOST	ID	TYPE	EXCEPTIONS	START OF SNAPSHOT	START OF SCAN	SCAN DURATION
Clean	VMFS Volumes	VMFS Volumes-10.06.2024 at 08:...	Incremental	No	2024-10-06 20:07:01	2024-10-06 20:07:01	1m 53s
Clean	FileServer	VWVM-FileServer-00000006702c...	Full	No	2024-10-06 20:07:03	2024-10-06 20:07:03	1m 51s
Clean	VMFS Volumes	VMFS Volumes-10.06.2024 at 06:...	Incremental	No	2024-10-06 18:07:01	2024-10-06 18:07:01	1m 25s
Clean	FileServer	VWVM-FileServer-00000006702a...	Full	No	2024-10-06 18:07:03	2024-10-06 18:07:03	1m 24s
Clean	VMFS Volumes	VMFS Volumes-10.06.2024 at 04:...	Incremental	No	2024-10-06 16:07:01	2024-10-06 16:07:01	1m 26s
Clean	FileServer	VWVM-FileServer-00000006702...	Full	No	2024-10-06 16:07:03	2024-10-06 16:07:03	1m 24s
Clean	VMFS Volumes	VMFS Volumes-10.06.2024 at 02:...	Incremental	No	2024-10-06 14:07:00	2024-10-06 14:07:00	5m 25s
Clean	FileServer	VWVM-FileServer-00000006702...	Full	No	2024-10-06 14:07:02	2024-10-06 14:07:02	5m 23s
Clean	VMFS Volumes	VMFS Volumes-10.06.2024 at 12:...	Incremental	No	2024-10-06 12:07:01	2024-10-06 12:07:02	1m 24s
Clean	FileServer	VWVM-FileServer-00000006702...	Full	No	2024-10-06 12:07:03	2024-10-06 12:07:03	1m 24s

Figure 5-8 IBM Storage Sentinel scan list

12. The IBM Storage Sentinel Web UI Hosts screen provides detailed information about the scanning processes of the hosts whose data is being scanned. This screen can be used to monitor the Cyber Sensitivity Index, daily file modification information, statistics on the types of files modified, statistics on the number of files deleted per day, and entropy value information of the files. See Figure 5-9.

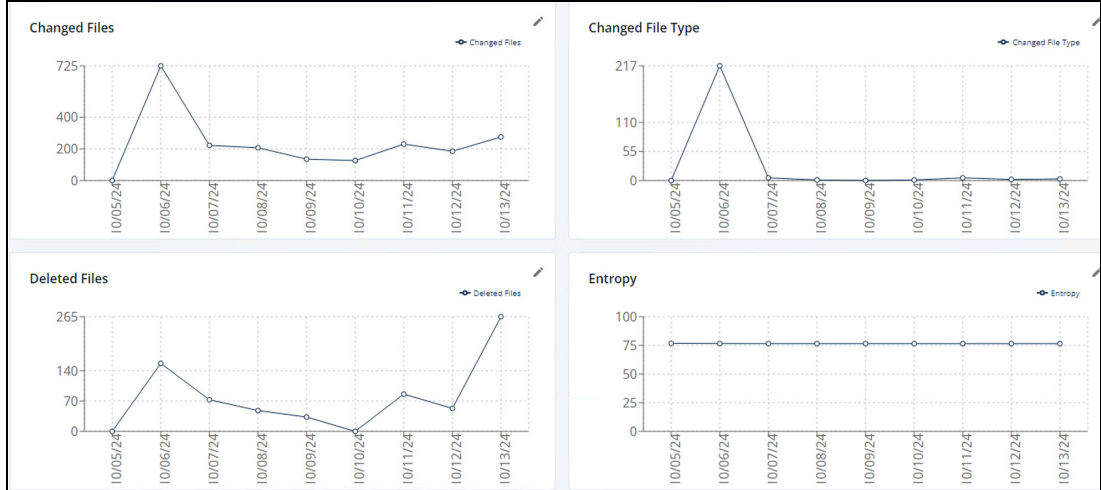


Figure 5-9 IBM Storage Sentinel Hosts scanned file statistics

13. The ability to customize the monitoring system, allowing users to set thresholds and receive automated notifications. This would help ensure proactive monitoring and faster response to potential issues. To create alerts from the statistics on this screen, you can define your own sub-values for each section and automatically receive notifications when values drop below defined levels. See Figure 5-10.

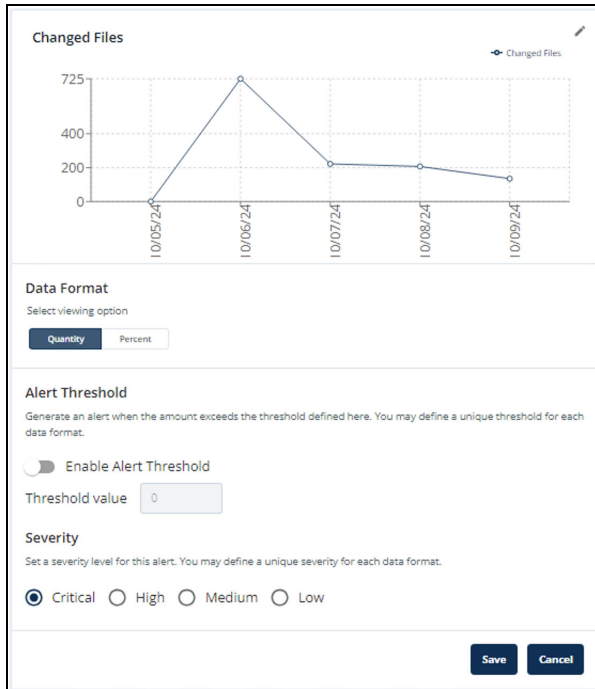


Figure 5-10 Edit settings for creating alerts

14. When IBM Storage Sentinel detects a threat during scanning, it displays it as an alert on the Sentinel screen and sends an alert email to the appropriate users as defined.

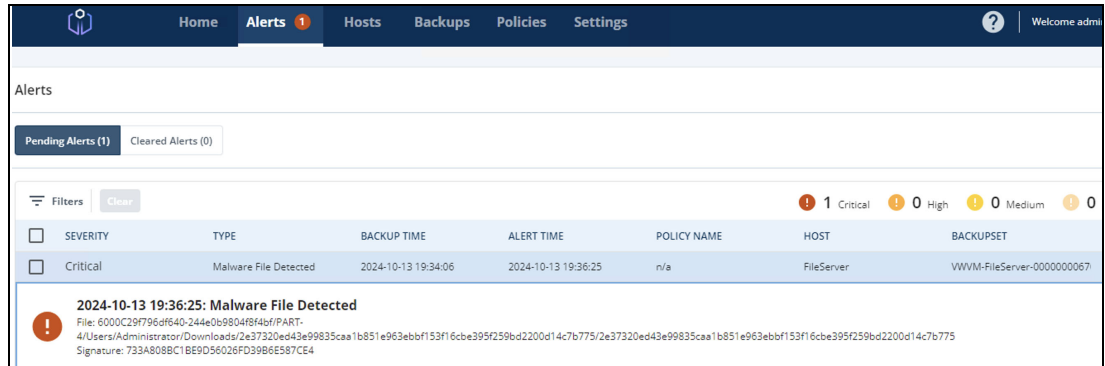


Figure 5-11 IBM Storage Sentinel alerts

5.4 Restore and recovery

IBM Storage Copy Data Management uses Copy Data Management technology for testing and cloning use cases, instant recovery, and full disaster recovery. If a threat alert is received during the scan process, SCDM can be used to restore and recover the affected VM.

VMware restore options

For the VMware environment, you can choose **Instant Disk Restore**, **Instant VM Restore** or **Instant VM Restore (Long Distance)**. See Figure 5-12.

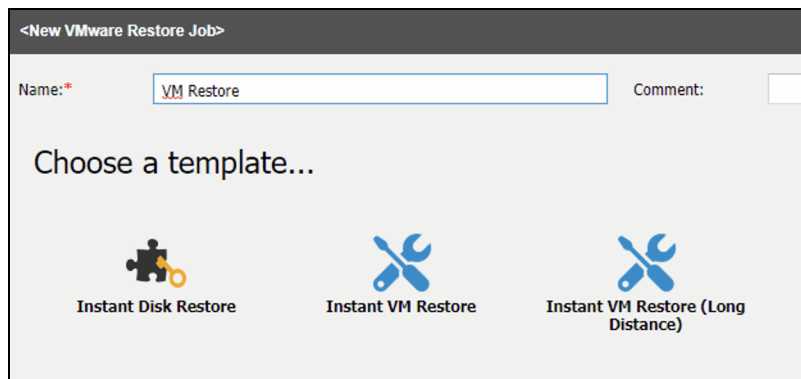


Figure 5-12 VMware restore options

Test mode, clone mode, and production mode can be selected during the restore process. For detailed restore options, see [IBM Documentation](#).



Protecting EPIC Cache and IRIS use case (on AIX)

This chapter describes how to configure IBM Storage Sentinel to protect and scan your EPIC databases on AIX. Many of the tasks are described in more detail in other chapters in this book.

Note: For detailed instructions on configuring IBM Storage Sentinel to protect EPIC databases on X86 and X64 platforms, consult IBM Redbooks *Cyber Resiliency with IBM Storage Sentinel and IBM Storage Safeguarded Copy*, [SG24-8541](#).

This chapter has the following sections:

- ▶ “Introduction” on page 88
- ▶ “Supported configurations for IBM Storage Copy Data Management and IBM Storage Sentinel for Epic databases” on page 88
- ▶ “Setting up a CDM and Storage Sentinel environment to scan Epic databases” on page 88
- ▶ “Scanning process” on page 99
- ▶ “Performing a restore of an Epic database backup” on page 101

6.1 Introduction

This chapter describes creating application-consistent, immutable snapshots of Intersystems Cache or IRIS databases and then having them scanned with IBM Storage Sentinel to detect possible malware corruption.

Epic Systems databases often use Intersystems database technology to process health records. The Epic health records management solution is an industry leader in this space. Because cyber criminals often target healthcare organizations, the Epic databases were the first applications supported by IBM Storage Sentinel.

6.2 Supported configurations for IBM Storage Copy Data Management and IBM Storage Sentinel for Epic databases

Before implementation, verify your application, OS, platform and storage requirements against the most current product documentation for IBM Storage Copy Data Management and IBM Storage Sentinel. Requirements might have changed after this book was written. The current manuals and other documentation may be more up to date than this book due to the passage of time. As the time of writing, the listed supported configurations are for IBM Storage Copy Data Management v2.2.25 (CDM) and IBM Storage Sentinel 8.7.0-1.16.

Important: The product documentation is the official source of information. Always verify any design against the most current release of documentation before implementation.

The supported Intersystems database applications include the following releases:

- ▶ Intersystems Cache 2015, 2016, 2017, 2018 or later.
- ▶ Intersystems IRIS 2021, 2022 or later.

CDM and Storage Sentinel support the Epic databases being hosted in vSphere virtual machines or on physical hosts. Configuration and operation details are described in this chapter.

- ▶ For physical machines, AIX 6.1 TL9 [5]. AIX 7.1 [5] [8], AIX 7.2 [5] [8] (beginning with 2.2.19), AIX 7.3 [5] [8] (beginning with 2.2.19), CentOS 7.0.
- ▶ Storage Configuration: Fibre Channel, iSCSI.

6.3 Setting up a CDM and Storage Sentinel environment to scan Epic databases

This section will describe the high-level steps needed to design and deploy a CDM and Storage Sentinel environment to protect and scan your Epic databases on AIX. Many of the tasks are described in more detail in other chapters in this book.

1. Plan and implement your supported server and storage deployment, as outlined earlier in this chapter. This will need to include predefining Safeguarded Copy volume groups.
2. Plan and implement the security settings and user accounts needed for creating the Safeguarded Copies, for integrating with vSphere, and for logging into CDM and Storage Sentinel, and any other necessary accounts. Decide to either use local accounts or use LDAP or Active Directory (AD). Decide the scope of authority for each account.

- Plan and implement your Storage Sentinel configuration at the scale you need to be able to scan your Epic databases. Plan for other workloads that you plan to support with the Storage Sentinel configuration. Plan how to distribute your scanning workloads across the configuration, and plan how often to scan your application servers.
- Deploy a new CDM virtual appliance, if needed. If using LDAP or AD for authentication, configure the security directory. To configure CDM to use LDAP or AD accounts, first register the LDAP server on the Provider Browser in the Configure Page. See Figure 6-1.

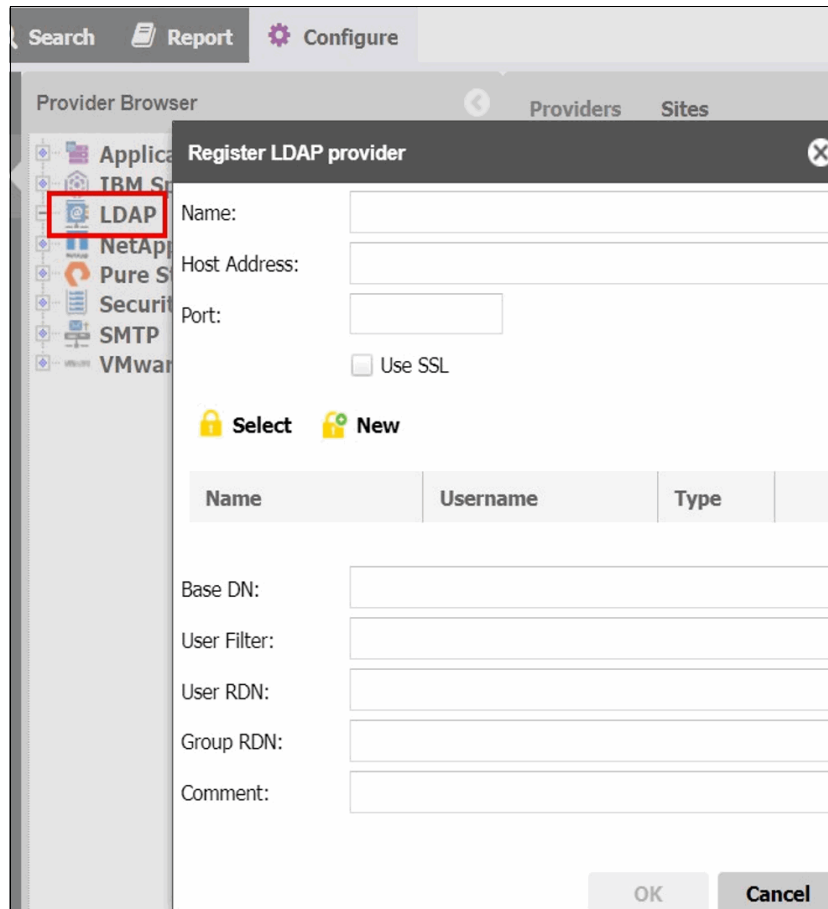


Figure 6-1 Register the LDAP server

- Import the LDAP group by going to the **Access Control panel** in the Configure Page, add a New User and Import the LDAP Group. See Figure 6-2 on page 90.

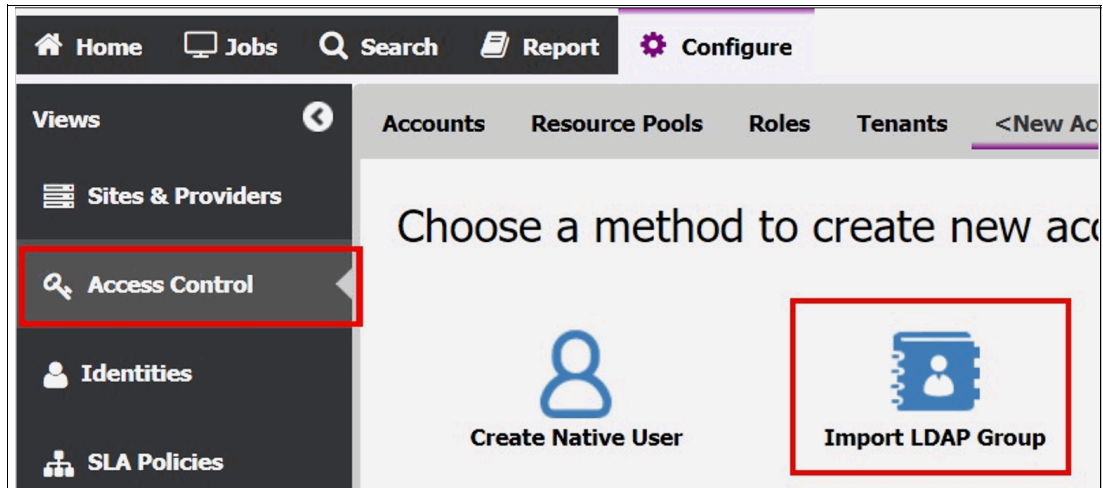


Figure 6-2 Import LDAP Group

- Define a CDM Site that will contain your storage, vSphere, application server and Storage Sentinel components. You configure Sites on the **Sites and Providers** section of the Configuration tab. See Figure 6-3.




Figure 6-3 Configure Sites

- Register your storage components using the appropriate Storage Virtualize user accounts. This automatically adds the storage to the daily Storage Virtualize inventory job and starts an inventory of the newly registered storage. Register the storage under the IBM Spectrum Virtualize node in the Provider Browser on the Configuration page. See Figure 6-4.

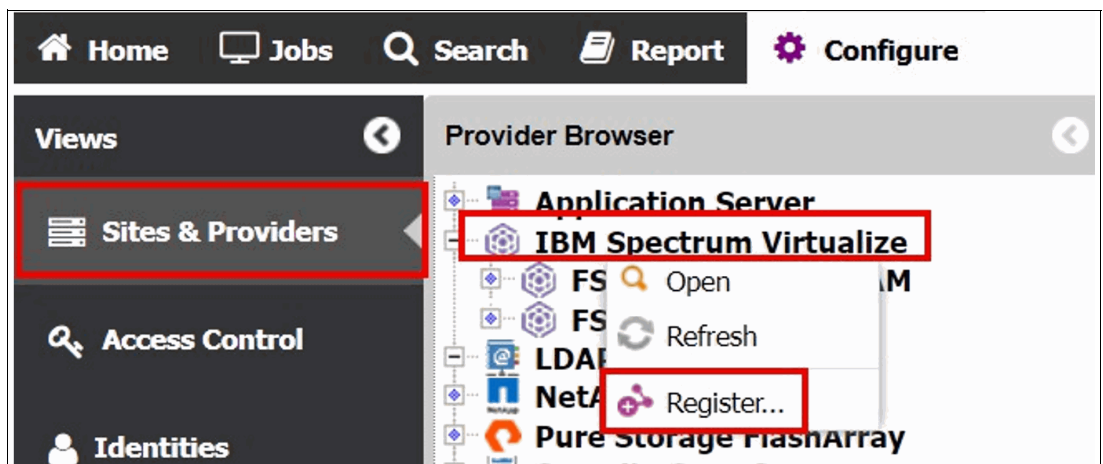


Figure 6-4 Register your storage components

8. If your Storage Sentinel server is a VM, register your vCenter server(s) to CDM using the appropriate account or certificate. This automatically adds the vSphere environment to the daily vSphere inventory job and starts an inventory of the newly registered components. Register any vCenter servers on the VMware node of the Provider Browser on the Configuration page. See Figure 6-5.

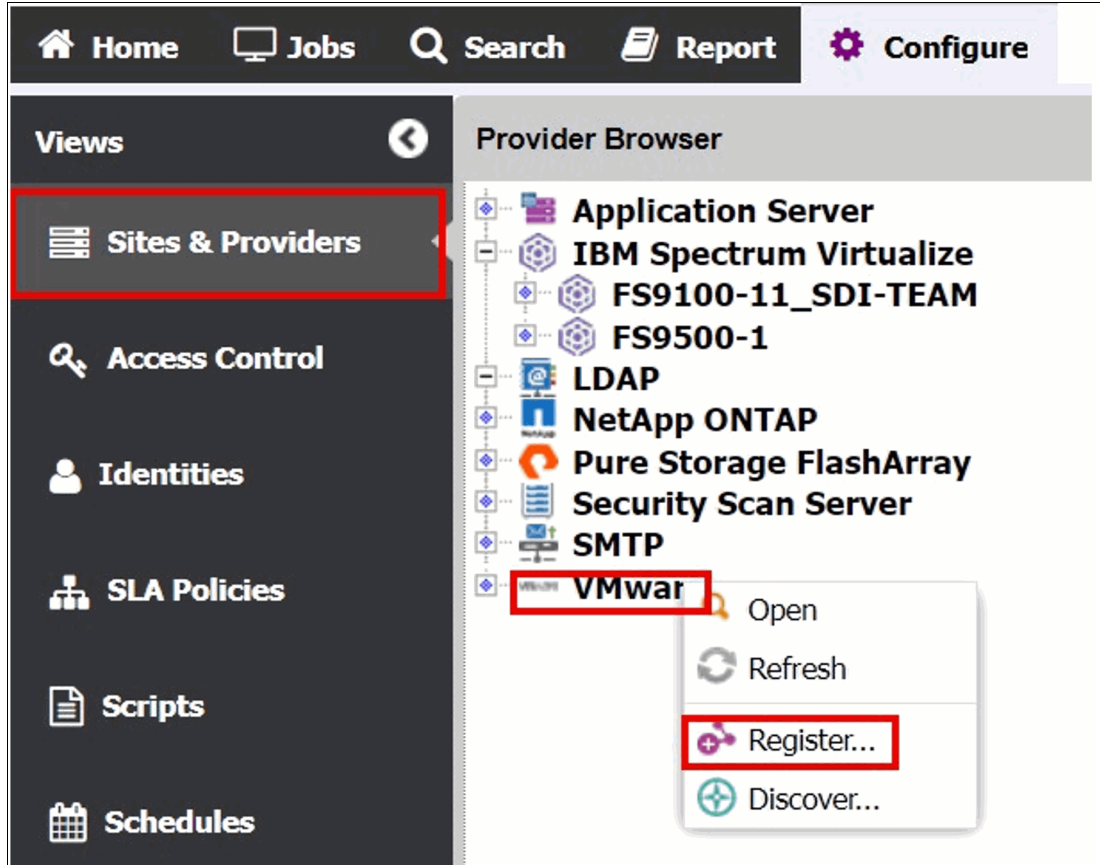


Figure 6-5 Register your vCenter server(s)

9. Register the Epic DB server in the Application Server node on the Provider Browser on the Configuration Page. See Figure 6-6 on page 92.

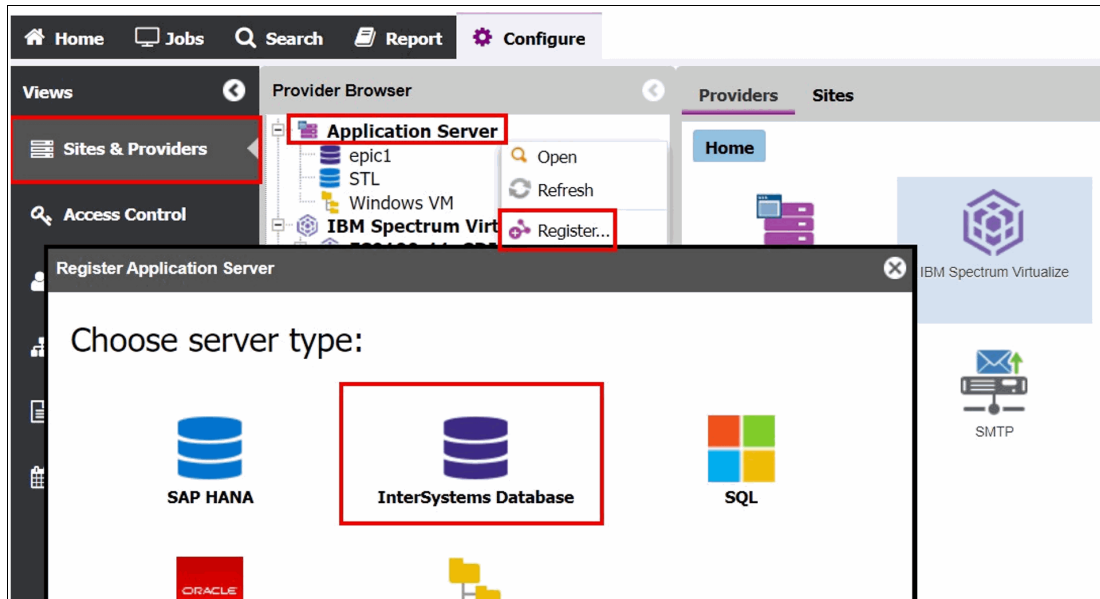


Figure 6-6 Register your Epic DB application servers

10. Fill the Epic DB name, Host Address, OS Type, System Credentials. See Figure 6-7.

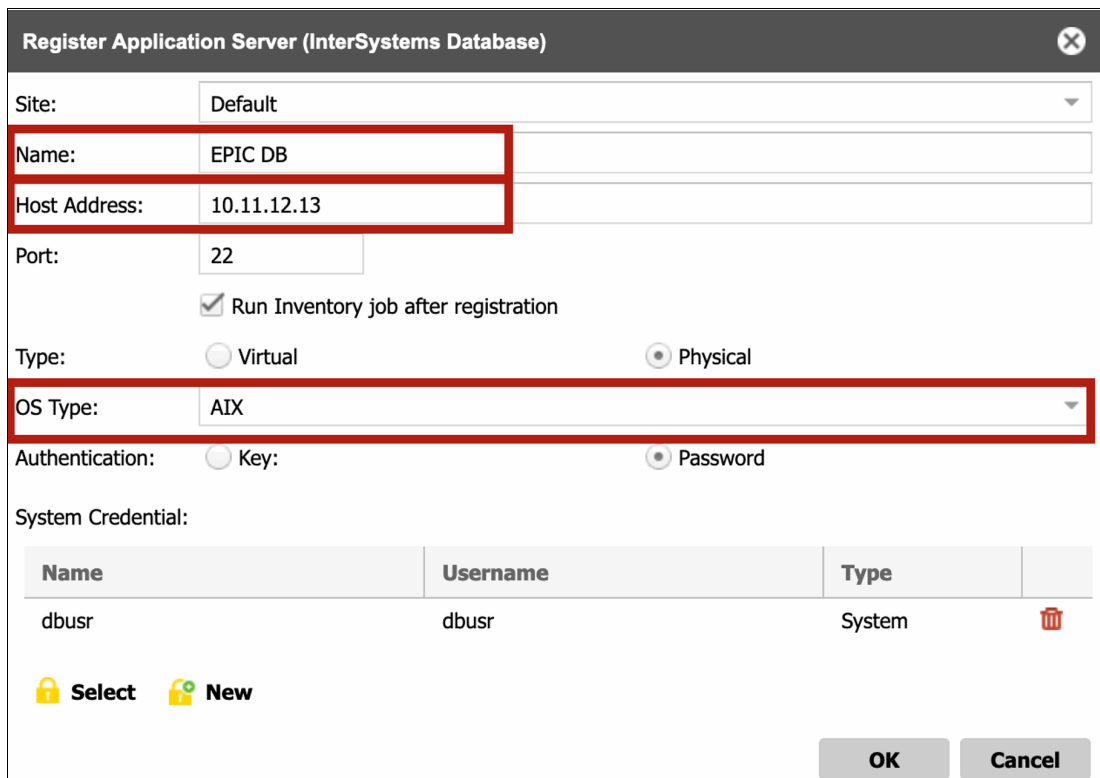


Figure 6-7 Register your Epic DB application servers

11. Register your Storage Sentinel server(s) using the Security Scanner node in the Configuration tree. If your Storage Sentinel servers are VMs, wait until the vCenter inventory has been completed and you have validated the Storage Sentinel VMs that were found during the inventory by expanding the navigation tree in the Configuration tab where you registered the vCenter server.

You should see the ESXi hosts and the VMs defined to vCenter in that tree, including the Storage Sentinel servers you need to register. You register the Storage Sentinel server under the **Security Scan Server** node on the Provider Browser in the Configuration page. See Figure 6-8.

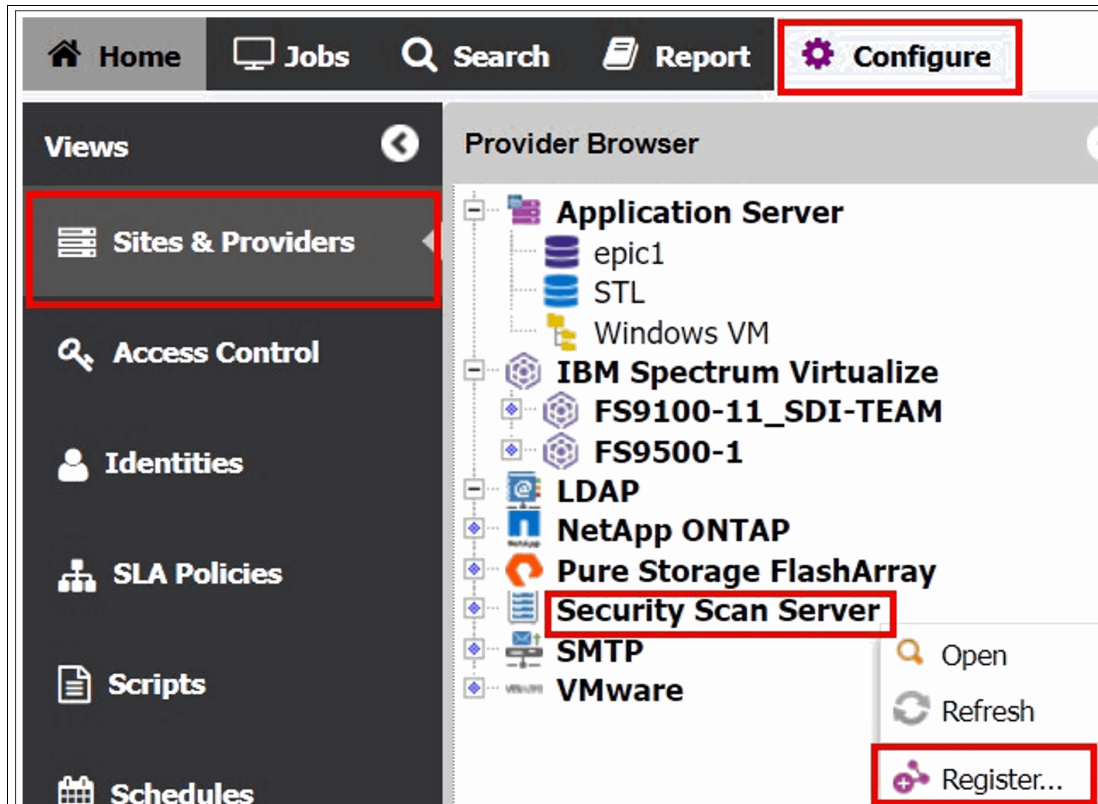


Figure 6-8 Register your Storage Sentinel server(s)

12. Validate that each of your components has registered correctly. See Figure 6-9 on page 94.
 - a. View the results of the Epic inventory job to validate that the Cache or IRIS instances were protected.
 - b. View the information on the application detected and the number of databases cataloged.
 - c. Navigate to the Jobs page, expand the application servers and select Intersystems.
 - d. Click the **Default Inventory Job**.
 - e. On the bottom panel, click **History** and then click on the job log hyperlink to open the job log page.

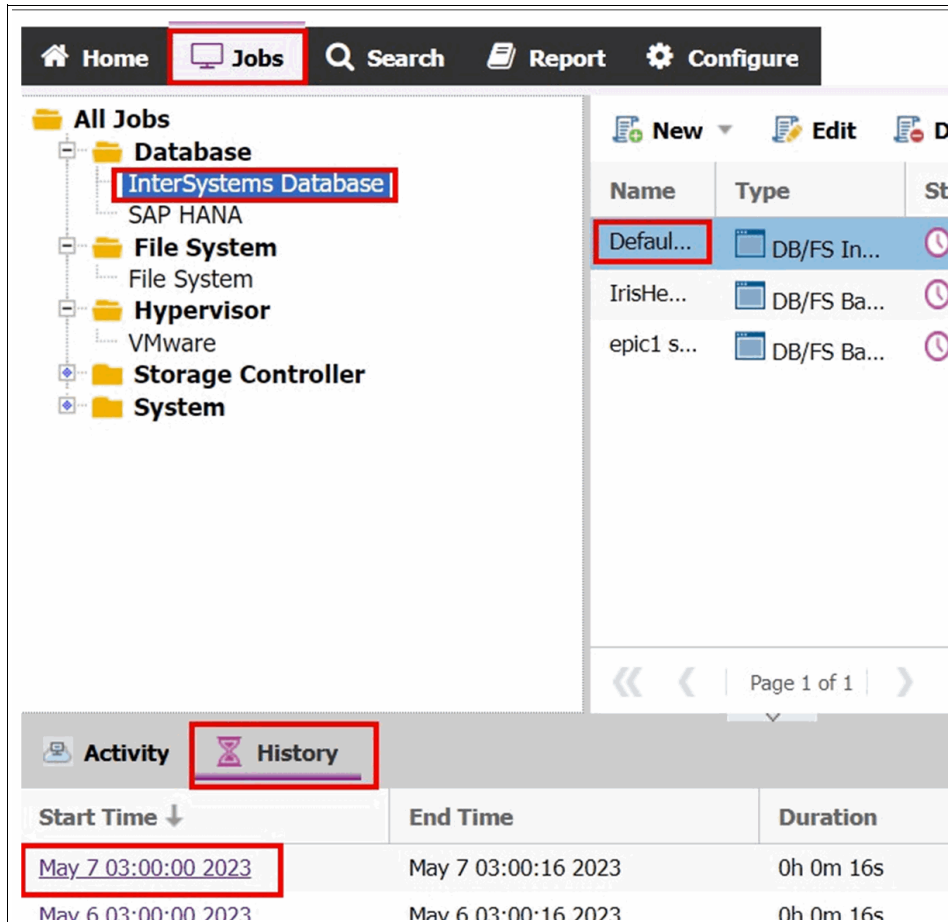


Figure 6-9 Click on the job log hyperlink to open the job log page

Figure 6-10 shows the job log page.

	May 7 03:00:16 2023	2	Cataloged 1 instance(s), 0 database group(s) 1 database(s), 8 disk(s)
--	---------------------	---	---

Figure 6-10 Job log page

13. Define an SLA for your Epic DB data protection. You will need to navigate to the SLA Policies section of the Configuration page, click **New** and then select **IBM Spectrum Virtualize**. See Figure 6-11 on page 95.

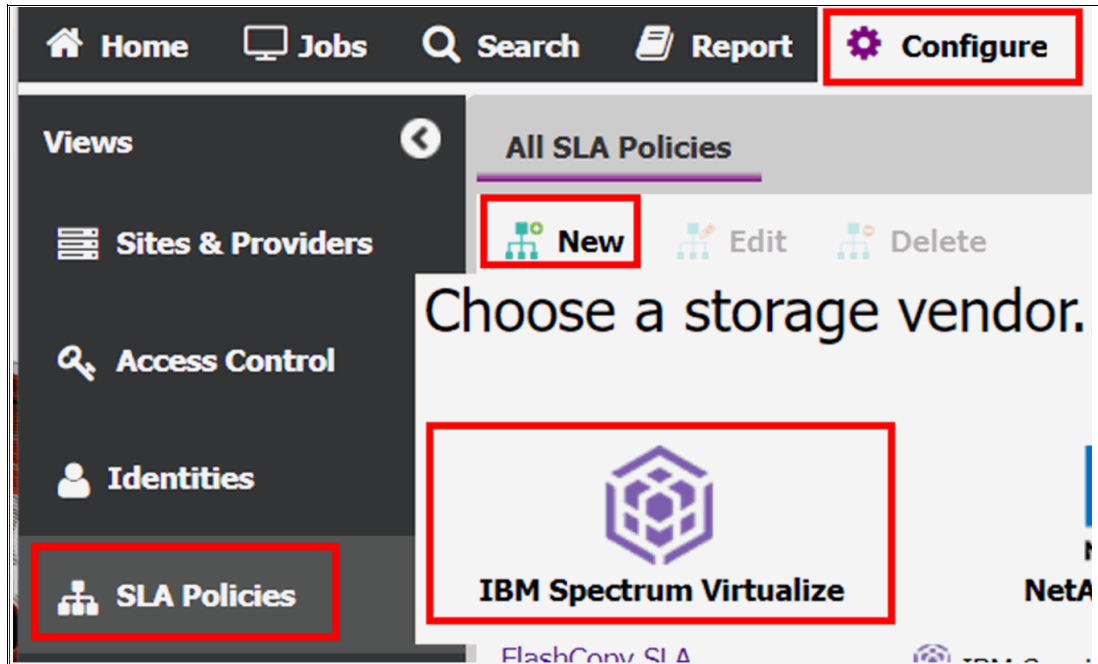


Figure 6-11 Define an SLA for your Epic DB data protection

- a. Give the SLA a unique name (ideally one that identifies the type of protection the SLA manages). Add a meaningful comment.
- b. Select the **Source icon** and enter the Frequency and Interval for this SLA.
- c. Right-click on **Source** and add **Safeguarded Copy**. See Figure 6-12.

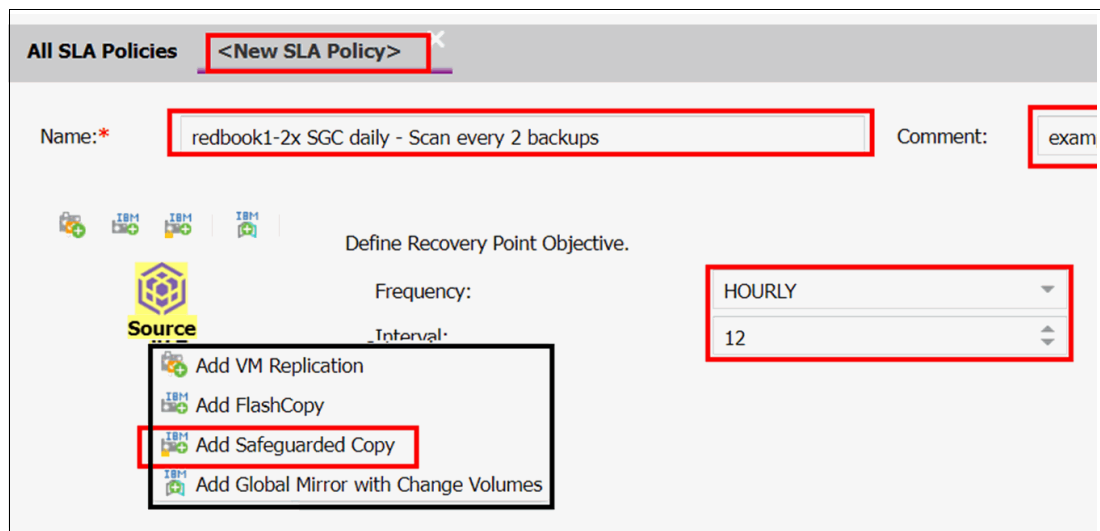


Figure 6-12 Add Safeguarded Copy

- d. Select the **Safeguarded Copy** Volume Group to associate with this SLA.
- e. Identify the number of days to retain the Safeguarded Copies.
- f. Give this set of Safeguarded Copies a unique and meaningful name.

- g. If desired, define a volume prefix for the Safeguarded Copies. It is recommended that you use a unique and meaningful prefix to identify the CDM instance and SLA name, to help correlate Safeguarded Copies back to what created those copies.
- h. Select the checkbox to perform security scanning, identify how many backups between scans and select the Storage Sentinel instance you have previously registered for this workload. See Figure 6-13.

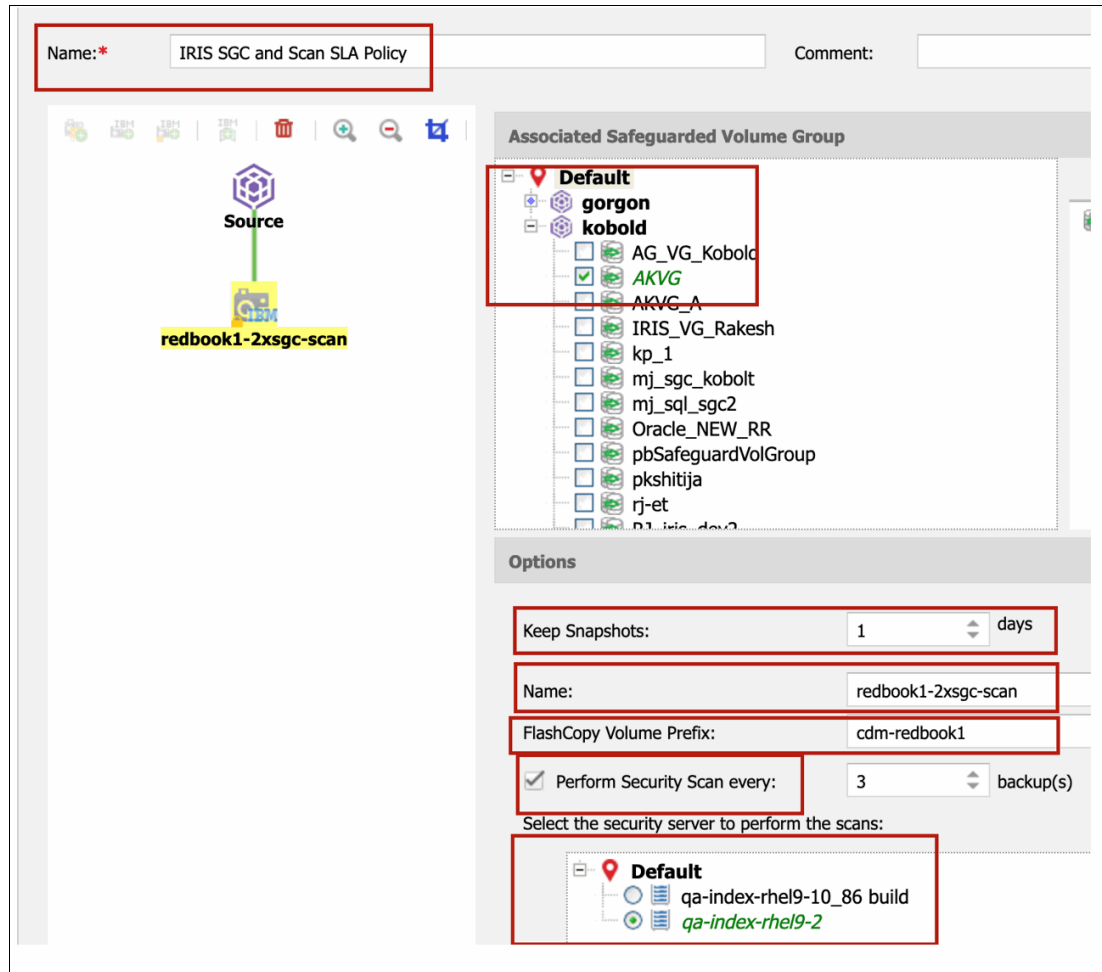


Figure 6-13 Completing the SLA configuration

- i. Save your SLA.
14. EPIC database scanning (on AIX OS) requires an AIX-based file system host registered in SCDM, which the backup workflow will make use of to mount JFS volumes. Hence, register the AIX filesystem in the Application Server node on the Provider Browser on the Configuration. See Figure 6-14 on page 97 and Figure 6-15 on page 97.

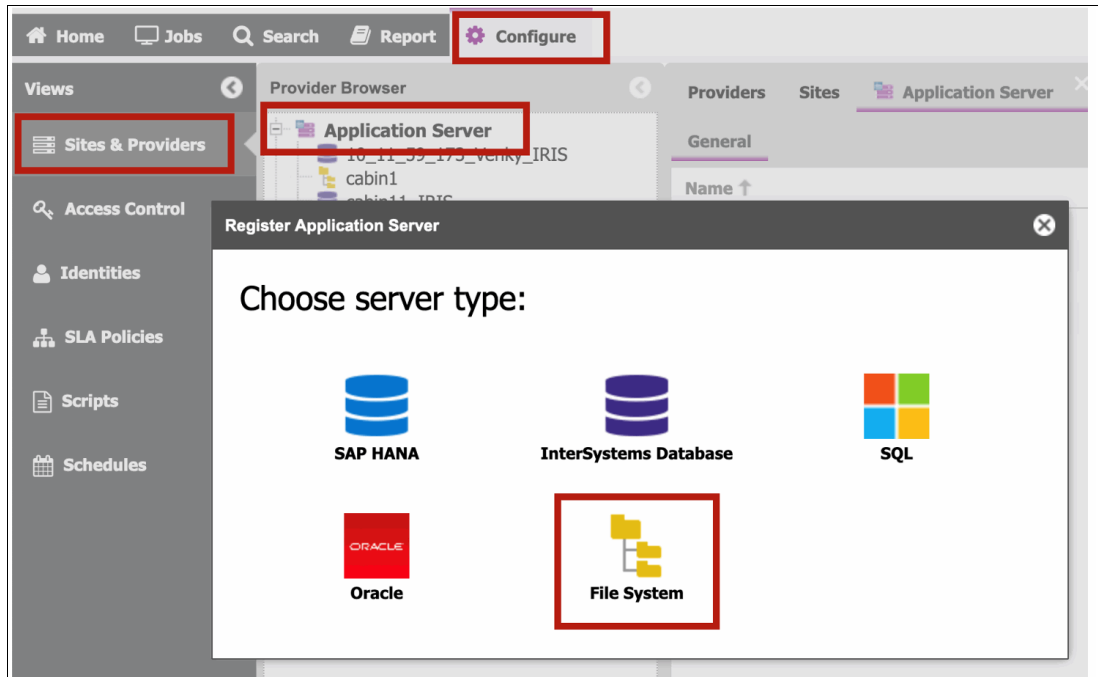


Figure 6-14 Register the AIX proxy FileSystem -1

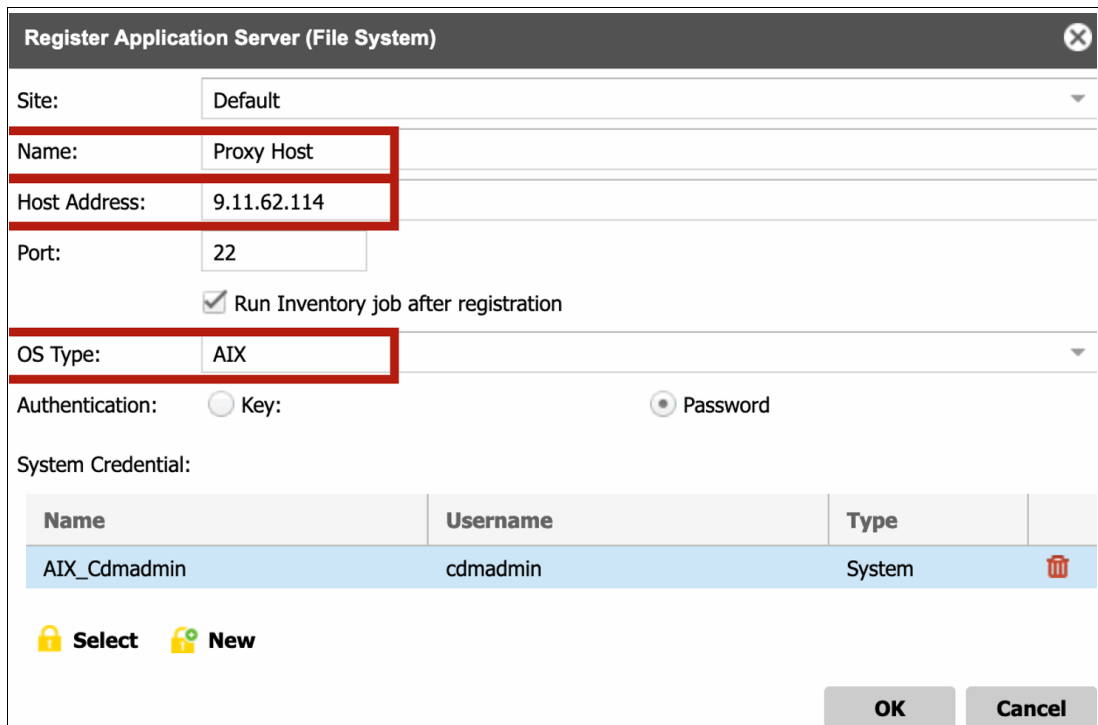


Figure 6-15 Register the AIX proxy FileSystem -2

15. Define a backup job for your Epic application server(s).

- a. Navigate to the Jobs page and expand the Database node and click on InterSystems Database.
- b. Click **New** and select **Backup**. See Figure 6-16 on page 98.

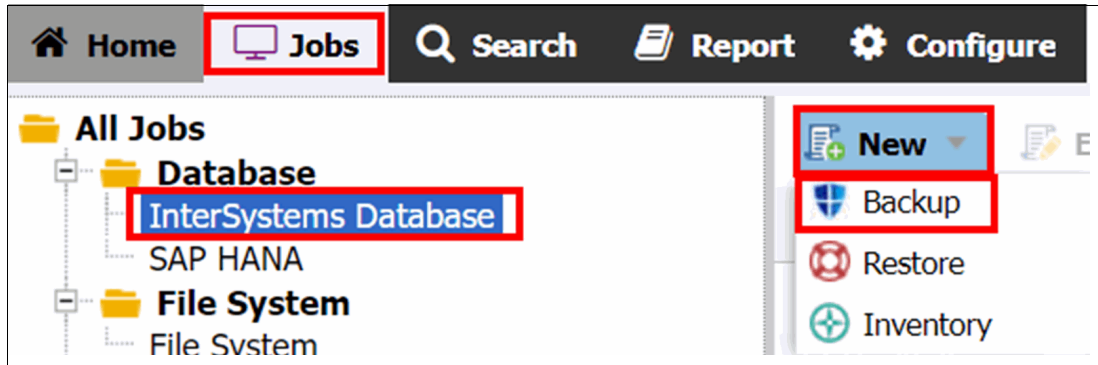


Figure 6-16 Defining a backup job

- c. Give this job a meaningful name and comment.
- d. Expand the node next to the site your Epic instance resides within and expand the Epic server's node. Check the box next to the instance or database, as you wish.
- e. Select your SLA. See Figure 6-17.
- f. Select **Enable Proxy Server**.
- g. Choose a previously configured AIX proxy host, which was shown and registered in the previous step (step 14).

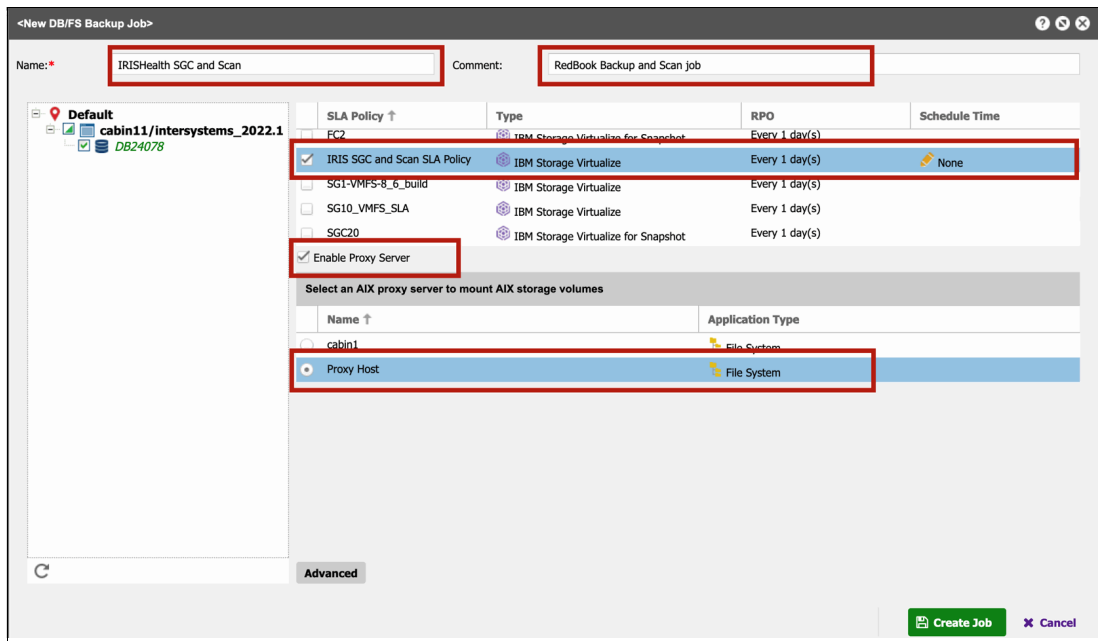


Figure 6-17 Select your SLA

- h. Click the **Schedule Time** column to define a schedule and set a start time. The SLA defines how often the job will run. For example, if you have defined the SLA to run every 12 hours and schedule the backup to start at midnight, it will run at midnight and at noon.
- i. If you wish to define pre-job, post-job, pre-snapshot or post-snapshot scripts to run as part of this job, click the advanced button and identify the scripts you wish to run.

16. Monitor your data protection as shown in Figure 6-18. You can manually start the backup job at any time by right-clicking on the job definition, but you might affect your production Epic instance if you run the job at a time of high activity.

Also, running a security scan might affect any running backups. After the job starts, review the job log for the backup job. Verify the instance was quiesced, a snapshot was created, and the instance was unquiesced. If a security scan is scheduled to be run for this job, the job log lists activity showing copies made of the Safeguarded copies and then assigning and mounting those copies to the security scanner.

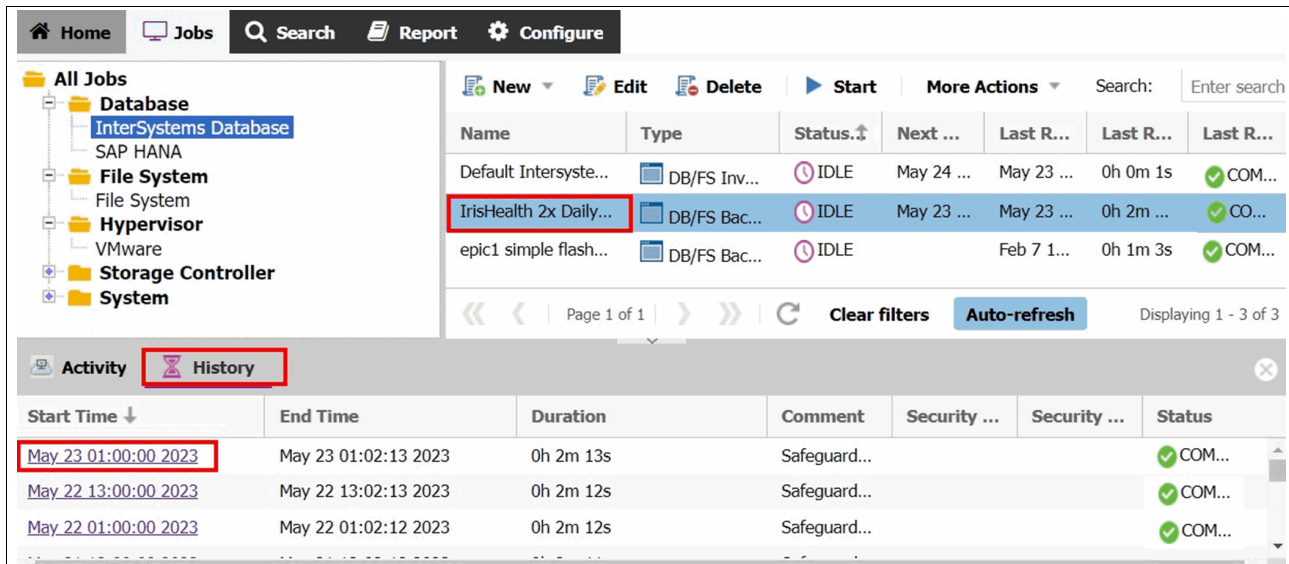


Figure 6-18 Monitor your data protection

6.4 Scanning process

Once the backup job is created and its scheduled, safeguarded snapshots run at the scheduled time and the scanning process automatically runs as defined in the SLA policy.

When the job is finished, view the job logs in the history panel, and if the security scan was successful and no ransomware threats are detected by the job, the job logs can be found stating, "Security scan finished with state: Done. No threats detected." See Figure 6-19.

i	Oct 29 18:48:29 2024	3	Security Scanning of protected databases
i	Oct 29 18:48:30 2024	3	Starting Index job on mount path /tmp/mounts/10_11_59_153/1007 with job name 1007...
i	Oct 29 18:48:34 2024	3	Index job (263) created.
i	Oct 29 18:48:34 2024	3	Index job (263) started.
i	Oct 29 18:50:29 2024	3	Security Scan finished with state: Done. No threats detected.

Figure 6-19 No threats detected

If the job detects ransomware threats, the SCDM would fail the Job with job logs "Security Scan finished with state: Done. Previous threat detected: false. Number of new threats detected: 1". See Figure 6-20 on page 100.

	Oct 16 13:26:23 2024	2	Security Scanning of protected databases
	Oct 16 13:26:23 2024	2	Starting Index job on mount path /tmp/mounts/10_11_59_175/1027 with job name 1027...
	Oct 16 13:26:30 2024	2	Index job (187) created.
	Oct 16 13:26:31 2024	2	Index job (187) started.
	Oct 16 13:32:24 2024	2	Security Scan finished with state: Done. Previous threat detected: false. Number of new threats detected: 1.
	Oct 16 13:32:24 2024	2	Unmounting database snapshot copies after Security Scanning

Figure 6-20 Threats detected

17. The scan jobs that are sent by SCDM can be monitored in the Backups section of the Sentinel web interface at <https://<hostname>/sentinel>. See Figure 6-21.

The screenshot shows the 'Backups' section of the IBM Storage Sentinel web interface. The page title is 'Snapshots'. There are filters and a download icon at the top left. The table below lists various snapshots with columns for Status, Host, ID, Type, Exceptions, Start of Snapshot, Start of Scan, and Scan Duration. The status for all listed snapshots is 'Clean'. The pagination at the bottom indicates 'Rows per page: 10', 'Page 1 of 3', and the current page is '1'.

STATUS	HOST	ID	TYPE	EXCEPTIONS	START OF SNAPSHOT	START OF SCAN	SCAN DURATION
Clean	1029	1029-10.30.2024 at 07:09 AM-211-1	Incremental	No	2024-10-30 07:09:29	2024-10-30 07:09:40	17s
Clean	1029	1029-10.28.2024 at 07:08 AM-210-1	Incremental	No	2024-10-28 07:08:22	2024-10-28 07:09:09	4s
Clean	1029	1029-10.27.2024 at 07:08 AM-209-1	Incremental	No	2024-10-27 07:08:44	2024-10-27 07:08:53	17s
Clean	1029	1029-10.26.2024 at 07:08 AM-208-1	Incremental	No	2024-10-26 07:08:11	2024-10-26 07:08:16	13s
Clean	1029	1029-10.25.2024 at 07:10 AM-207-1	Incremental	No	2024-10-25 07:10:37	2024-10-25 07:10:41	23s
Clean	1027	1027-10.25.2024 at 03:26 AM-206-1	Incremental	No	2024-10-25 03:26:04	2024-10-25 03:26:06	40s
Clean	1027	1027-10.25.2024 at 02:56 AM-205-1	Incremental	No	2024-10-25 02:56:11	2024-10-25 02:56:36	13s
Clean	1027	1027-10.25.2024 at 02:05 AM-204-1	Incremental	No	2024-10-25 02:05:31	2024-10-25 02:05:36	24s
Clean	1032	1032-10.25.2024 at 01:52 AM-203-1	Incremental	No	2024-10-25 01:52:24	2024-10-25 01:52:33	2m 46s
Clean	1029	1029-10.23.2024 at 07:08 AM-202-1	Incremental	No	2024-10-23 07:08:36	2024-10-23 07:08:39	18s

Figure 6-21 IBM Storage Sentinel scan list

When IBM Storage Sentinel detects a threat during scanning, it displays it as an alert on the Sentinel screen. See Figure 6-22 on page 101.

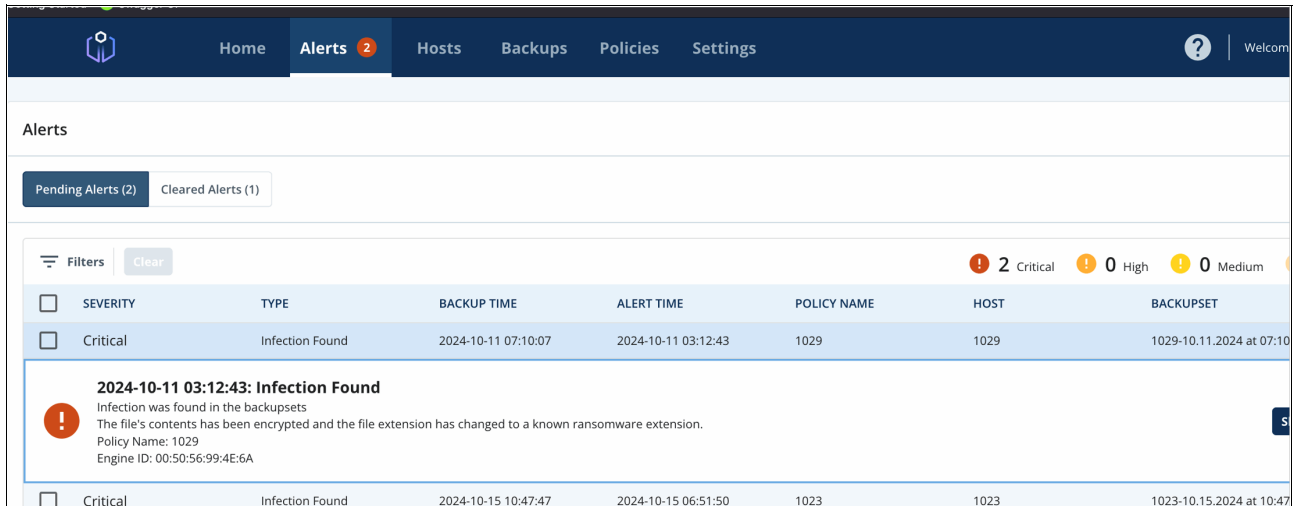


Figure 6-22 IBM Storage Sentinel alerts

If the job is executed correctly, you can define a restore job and select a recovery point, as outlined in 6.5, “Performing a restore of an Epic database backup” on page 101.

6.5 Performing a restore of an Epic database backup

After a scheduled job creates Safeguarded Copies of the Epic databases and scans them for malware corruption, you can perform a restore as needed.

Perform the following steps to define a Restore job.

1. On the Jobs page, expand the Databases node and select InterSystems Database. Click the **New** button and select **Restore**. See Figure 6-23.

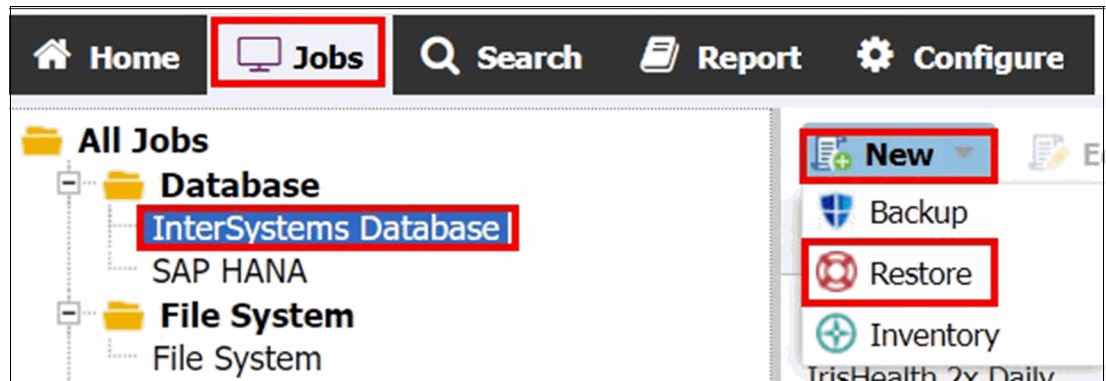


Figure 6-23 Select Restore

2. Enter a meaningful name and comment, select to perform an Instant Disk Restore or Instant Database Restore. An Instant Disk Restore mounts the file systems to the target host but does not define or start the database. An Instant Database restore defines and starts the database to an Epic database instance. For this example, select Instant Database Restore. See Figure 6-24 on page 102.

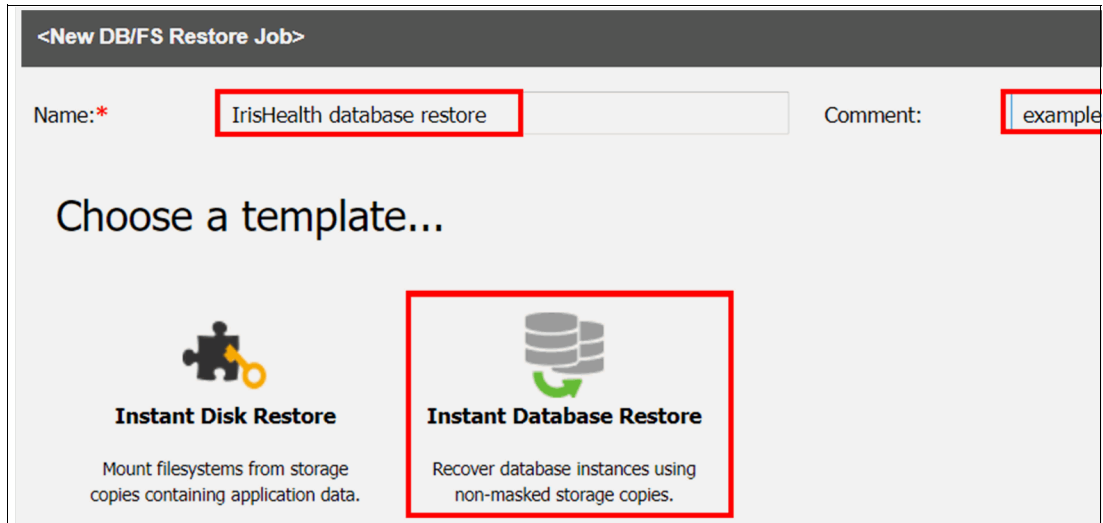


Figure 6-24 Instant Database Restore

3. The source icon will be automatically selected. Use the **Application Browser tree** to navigate and select the database to be restored. See Figure 6-25.



Figure 6-25 Select the database to be restored

4. Click the **Copy** icon. Click the **Select Specific Version** or **Use Latest Successful Scan** button. Here you can click the **Use Latest** or **Use Latest Successful Scan** buttons to select those recovery points. If you wish to select a specific version, click the **Use Latest text** in the Version column. See Figure 6-26.



Figure 6-26 Click the Copy icon

5. You can then select a specific version. See Figure 6-27.

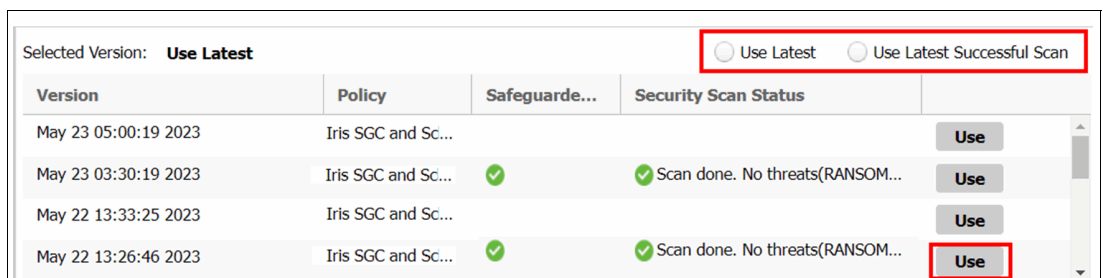


Figure 6-27 Select a specific version

6. Click the **Destination** icon. You can then select the database instance as the recovery target. Click **Create a Job**. You can then start the job and monitor until completion. See Figure 6-28.

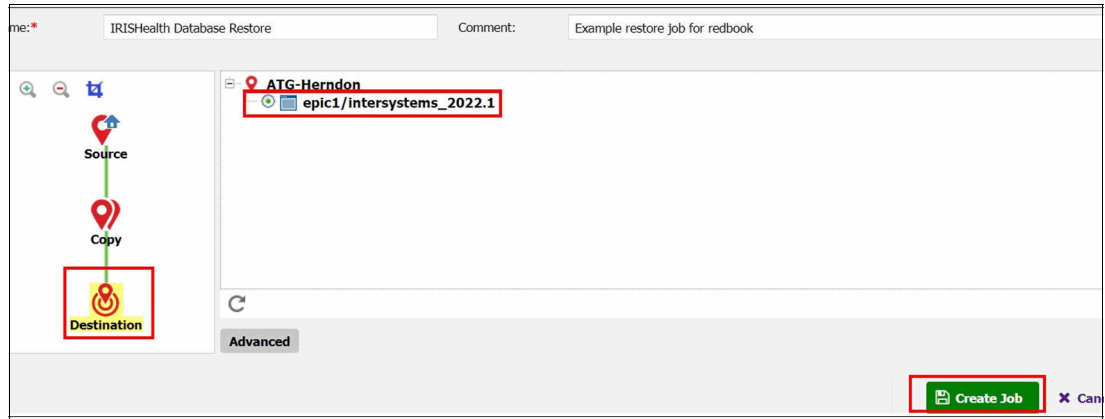


Figure 6-28 Create a Job



Secure and resilient AI

In this chapter, we discuss how the IBM Guardium platform integrates with IBM watsonx. governance to enhance governance, security, and monitoring of artificial intelligence (AI) initiatives while promoting cross-functional collaboration. The following topics are discussed in this chapter:

- ▶ “Exploring LLMs: Applications, risks, and security in AI Systems” on page 106
- ▶ “Examples of risks associated with Large Language Models” on page 106
- ▶ “Moderation and instruction tuning” on page 108
- ▶ “Understanding system instructions, fine-tuning, and vector databases” on page 110
- ▶ “Effective risk management” on page 112
- ▶ “Protecting sensitive information” on page 112
- ▶ “Securing Generative AI: Threat vectors, data protection, and advanced applications” on page 117
- ▶ “Security measures for LLMs” on page 119
- ▶ “Integrated security technologies for AI management” on page 120
- ▶ “IBM Guardium AI Security: Manage data model security risk demo overview and key capabilities” on page 122

7.1 Exploring LLMs: Applications, risks, and security in AI Systems

As artificial intelligence continues to evolve, Large Language Models (LLMs) have emerged as pivotal tools in various applications. With their ability to understand and generate human-like text, LLMs are increasingly being integrated into sectors such as customer service, healthcare, and content creation. However, with great technology comes great responsibility. Alongside the innovative applications of LLMs, concerns about their risks and vulnerabilities in AI systems have emerged. This article explores the applications of LLMs, the inherent risks, and security measures essential for their deployment.

7.1.1 Applications of Large Language Models

The primary strength of LLMs lies in their versatility. Organizations are deploying these models for various use cases, often classified into two main categories: question-answering systems and task completion systems.

Question-answering systems

Chatbots powered by LLMs can answer queries from users in real-time. For instance, a call center for a bank may use an LLM to facilitate customer inquiries about account balances or transaction statuses. These systems aim to reduce the volume of queries that require human intervention, allowing human resources to focus on more complex issues.

Task completion systems

Beyond mere questions, LLMs can also be leveraged to perform tasks. For instance, users may delegate certain functions to LLMs, such as drafting emails or generating reports based on input parameters. This distinction between question-answering and task completion is vital for designing effective AI systems.

7.1.2 Grounding and context

Successful interactions with LLMs depend on their understanding of context. When users pose questions, systems must process these queries through a structured framework that equips the LLM with the necessary data to produce valid responses. This is where grounding mechanisms come into play—where the model integrates contextual information, like user identities and specific data access rights, to improve relevance and accuracy.

7.2 Examples of risks associated with Large Language Models

While LLMs offer multiple advantages, they are not without risks. Understanding these risks is critical to ensure the safety and efficacy of their implementations. Here are some examples of such risks.

7.2.1 Adversarial risks across AI models

There are three main risk points in the AI pipeline. Data Collection and handling, Model Development and training, and Model Inference and live use. Positioning proper security controls to each point is essential to establish a safe and productive AI environment. See Figure 7-1 on page 107.

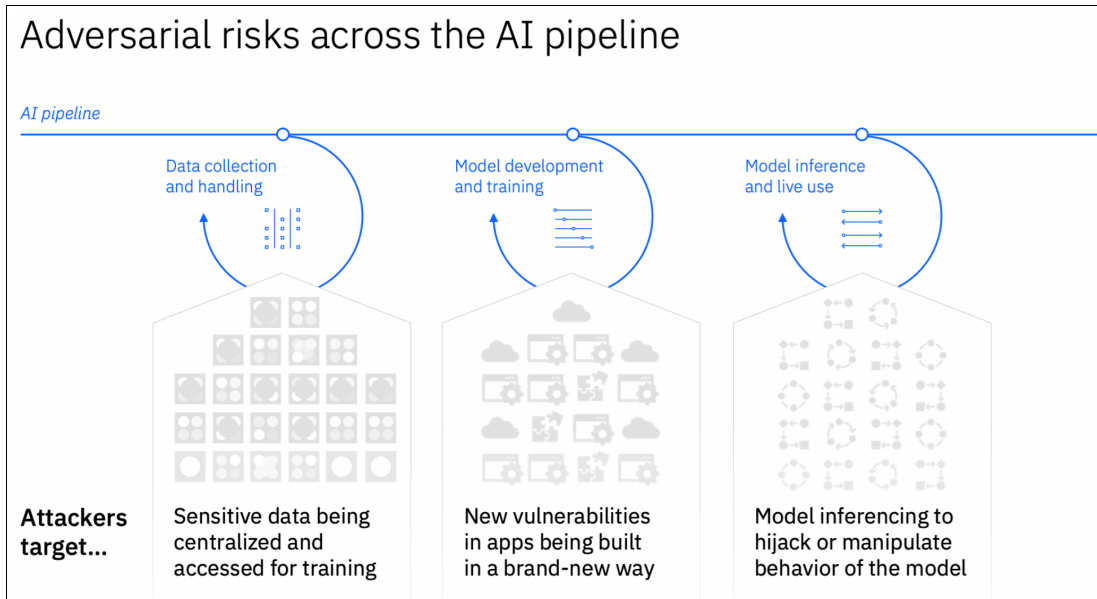


Figure 7-1 Adversarial risks across AI models

7.2.2 Prompt injection attacks

One of the paramount concerns is prompt injection, where malicious actors provide harmful or deceptive inputs that can exploit vulnerabilities in the model. This could lead to the generation of inappropriate content or leakage of sensitive information. See Figure 7-2.

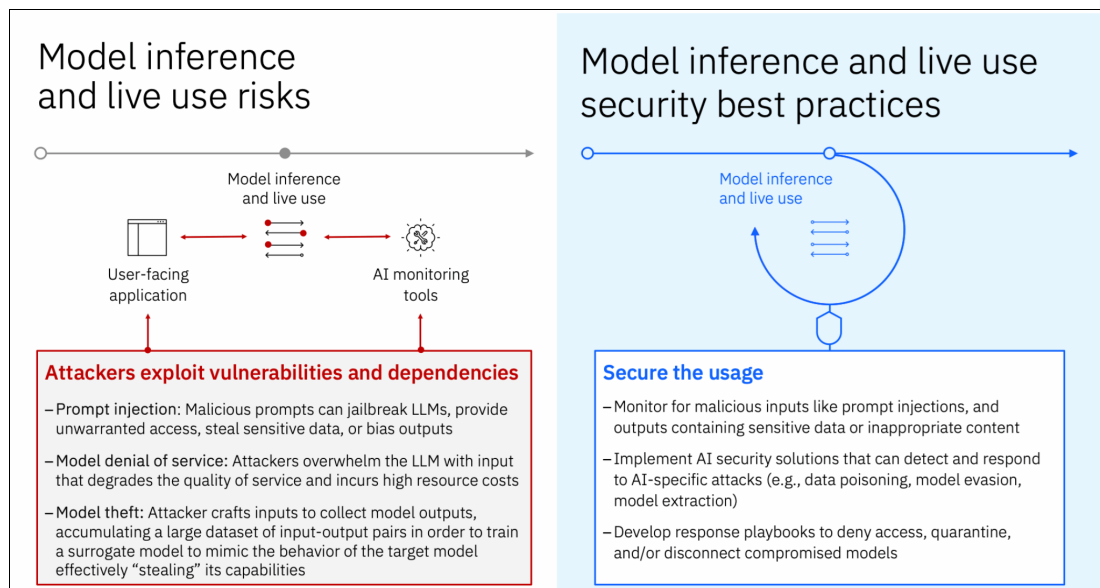


Figure 7-2 Model inference

LLMs are known to produce *hallucinations* - incorrect or fabricated responses that users might mistakenly take as facts. These inaccuracies can have severe implications, especially in sensitive sectors such as finance, government, and healthcare where misinformation could lead to significant errors, critical infrastructure concerns, and public safety.

7.3 Moderation and instruction tuning

Models can be pre-tuned with moderation protocols designed to identify and block inappropriate queries. Additionally, instruction tuning enables the models to better understand and follow directives, improving their reliability in sensitive applications.

Implementing robust grounding mechanisms ensures that models have access to relevant user data while protecting against unauthorized access. Appropriately structuring access controls based on user identities can further mitigate risks associated with data misuse.

7.3.1 Identity management

Managing and Protecting user identities ensures you are running a secure AI environment. IBM Security Verify solutions effectively protects from unauthorized access. IBM Verify Identity Management Suite provides organizations with security through risk-based authentication, access, privilege escalation, control, and monitoring. See Figure 7-3.


IBM Security Verify SaaS 		Multi-factor authentication (MFA) for cloud and on-prem apps	Proactively protect users from risk and fraud by using AI-informed access policies
User lifecycle management and governance identity analytics provides a 360-degree view of access risks and the ability to act based on those risk insights	User lifecycle management and governance identity analytics provides a 360-degree view of access risks and the ability to act based on those risk insights	SSO and access management for cloud and on-prem apps	Risk-based authentication and adaptive access control

Figure 7-3 IBM Verify Identity Management Suite

7.3.2 Identity Threat Detection and Response

Identity Threat Detection and Response (ITDR) is designed to address specific identity-focused threats that traditional Security Information and Event Management (SIEM) systems often struggle to identify. While SIEMs aggregate vast amounts of data to determine user activity, they typically overlook the crucial login activities within identity systems. ITDR not only detects these identity-related threats but also protects the underlying Identity and Access Management (IAM) systems. See Figure 7-4 on page 109.

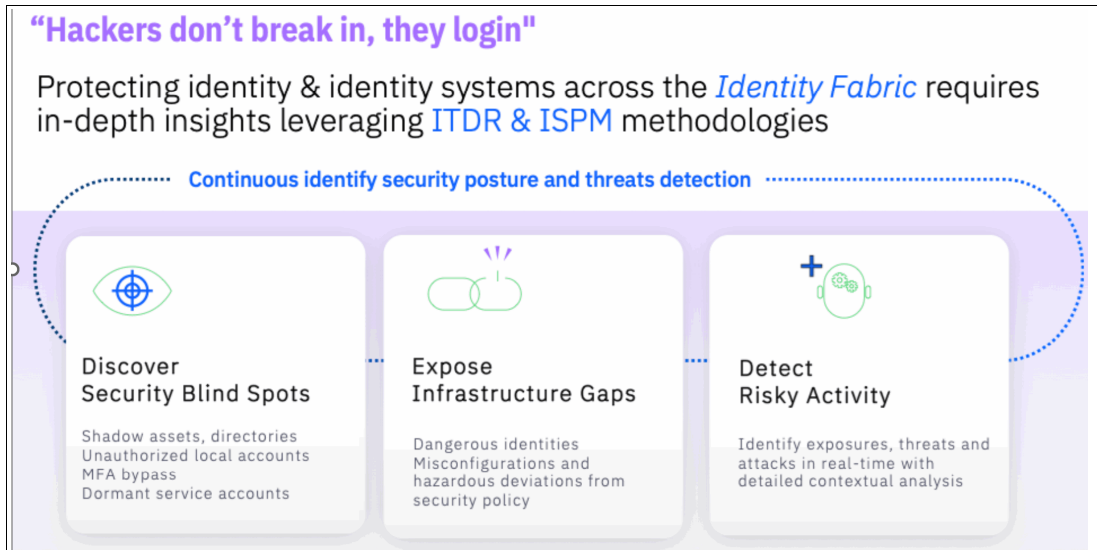


Figure 7-4 Identity Threat Detection with Identity Fabric

By identifying critical and suspicious identity activities, the IAM system sends streamlined insights to the SIEM, allowing for immediate analysis and response. This integration reduces the volume of inputs the SIEM must process, thereby enhancing the overall value and efficiency of threat detection.

ITDR helps organizations answer several key questions:

- ▶ Who is attacking or accessing assets?
- ▶ From where are these attempts originating?
- ▶ How are these attacks being conducted?
- ▶ What are the underlying exposures contributing to these risks?

Verify Identity Protection Platform helps organizations provide threat visibility to blind spots in AI deployments. See Figure 7-5.

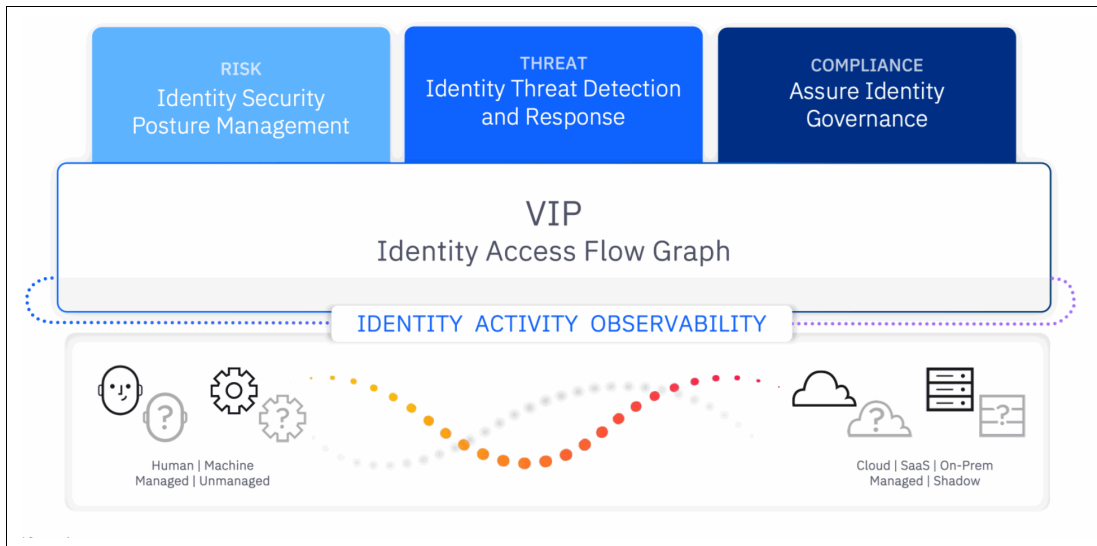


Figure 7-5 Verify Identity Protection Platform

This comprehensive approach enables organizations to enhance their security posture and respond more effectively to identity-related threats.

Verify Identity Protection provides organizations with comprehensive visibility across both managed and unmanaged identities. Figure 7-6 shows the features of Verify Identity Protection and out-of-the-box connectors.

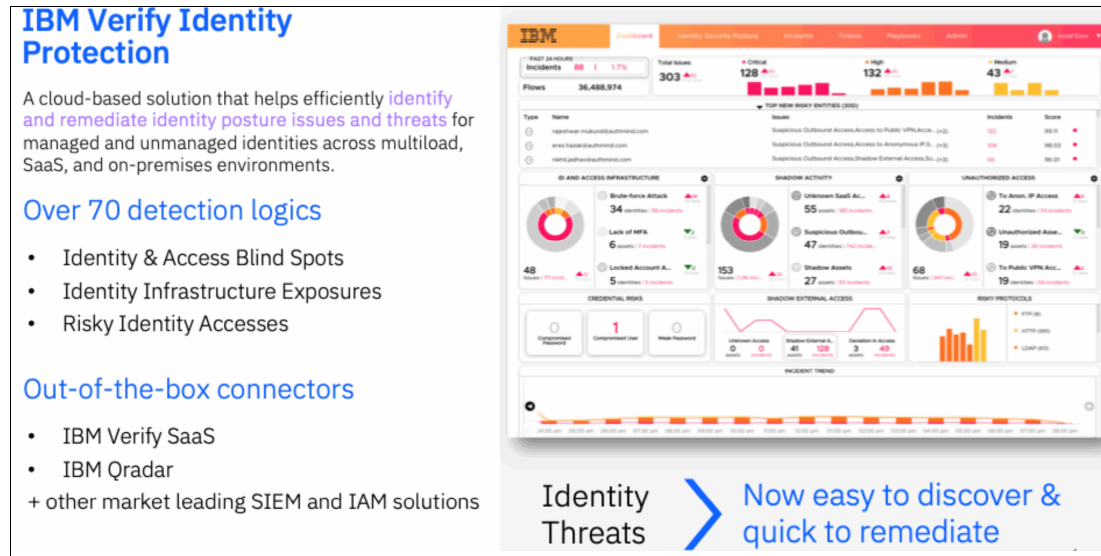


Figure 7-6 Verify Identity Protection

7.3.3 Continuous monitoring and feedback loops

Establishing continuous monitoring processes and introducing feedback loops allows organizations to refine their LLMs based on real-world interactions. This adaptive learning approach can improve model accuracy over time, while users can assess the model's performance through feedback mechanisms.

7.3.4 Conclusion

As LLMs become more integral to various sectors, understanding their applications and inherent risks is essential for organizations. By leveraging robust security measures and maintaining governance standards, businesses can harness the full potential of Large Language Models while minimizing risks associated with their deployment. A proactive approach that integrates technology, security, and human oversight will ultimately enable organizations to navigate the complexities of this evolving landscape of artificial intelligence.

7.4 Understanding system instructions, fine-tuning, and vector databases

In the realm of artificial intelligence, particularly in natural language processing (NLP), chatbots like ChatGPT offer users a conversational interface that mimics human interactions. However, the underlying mechanisms that enable such interactions are complex and multifaceted. This article explores the nuances of system instructions, the importance of fine-tuning, and the role of vector databases in enhancing ChatGPT interactions.

7.4.1 The nature of user interactions

When users engage with a chatbot, they expect natural and seamless dialogue, much like texting a friend. This expectation creates a demand for flexibility in response types, allowing chatbots to adjust their tone or style based on user input. However, that flexibility must be balanced with clear guidance on how the model is expected to behave. This is where system instructions come into play.

7.4.2 System instructions: The foundation of interaction

System instructions provide a framework for how the model should respond to user queries. These instructions delineate the expected behavior, tone, and guidelines for responding appropriately. For instance, a system instruction might state: *"You are a helpful assistant. Respond to the user's questions with honesty and integrity."* Such instructions serve as the first layer of processing that guides the AI's responses.

Moreover, system instructions also encompass ethical considerations, determining what content the model should avoid sharing. For instance, the model is programmed to filter profanity and illegal or hateful content.

7.4.3 The role of fine-tuning

While system instructions establish baseline behavior, fine-tuning further refines the model's capability to adhere to these instructions. Fine-tuning involves training the model on specific datasets that reinforce desired behaviors and responses. This dual-layer approach ensures the model not only understands the instructions but also has practical experience in applying them.

For example, if a user attempts to manipulate the model into ignoring previous instructions (a technique known as prompt injection), fine-tuning can equip the model with responses to such tactics, leading it to maintain appropriate boundaries.

7.4.4 Vector databases: Managing context

As conversations progress, maintaining context becomes increasingly important. ChatGPT models have a limit on how much information they can process at once, depending on the architecture used. For instance, different models have varying token limits for input; some might allow up to 100,000 tokens, while others may handle only a fraction of that.

To manage extensive conversation histories while staying within token limits, vector databases come into play. These databases store embeddings-numerical representations of the input text that capture semantic meaning. By embedding both previous messages and new user queries, the system can assess relevance and filter out unnecessary information, ensuring the most pertinent context is retained for interaction.

7.4.5 Retrieval-Augmented Generation

The combination of embeddings and vector databases facilitates a process known as Retrieval-Augmented Generation (RAG). In RAG, the model synthesizes previous relevant messages through a distance function (like cosine similarity) to retrieve the most applicable context for generating responses.

This means that instead of feeding an entire conversation history into the model, which could exceed its token capacity, only the most relevant segments are retained. This maintains coherent and contextually aware responses.

7.5 Effective risk management

As AI systems become increasingly integral to various industries, securing AI pipelines has emerged as a critical concern. Effective risk management involves ensuring the security of three key components: data, models, and infrastructure. We will explore the interconnectedness of these components, the risks they face, and strategies for securing AI pipelines against potential vulnerabilities and attack vectors.

Attacker will exploit vulnerabilities, target APIs, and engage into agent exploits to elevate access permissions. Organizations are encouraged to continuously scan for vulnerabilities, harden API and plugin integrations and enforce policies.

Figure 7-7 shows how to secure the AI pipeline.

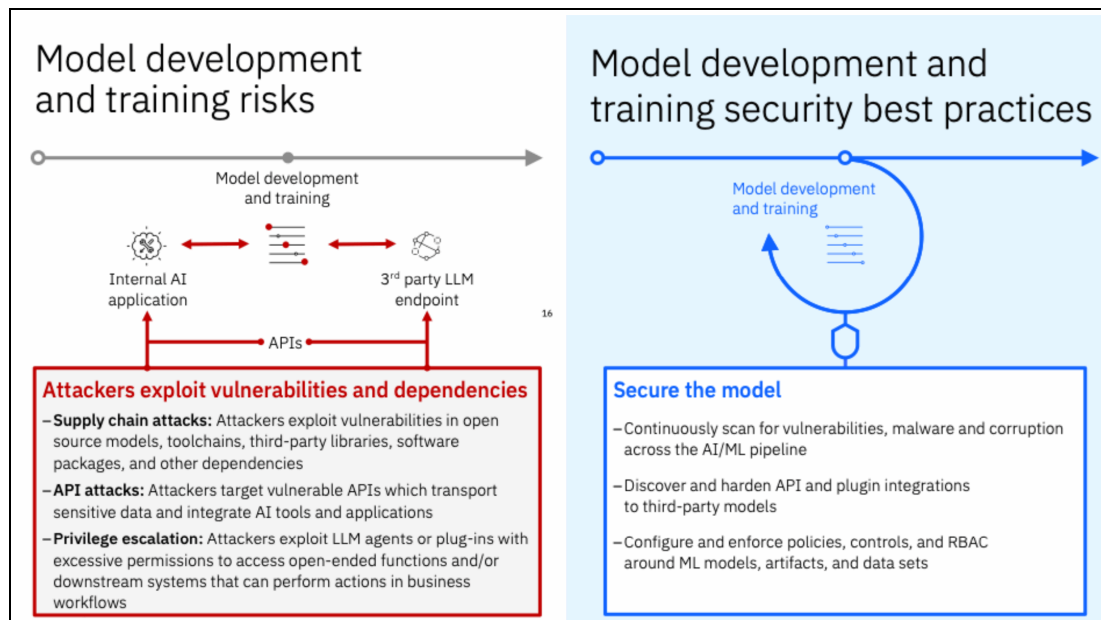


Figure 7-7 Securing AI pipeline - Attacker versus security perspective

7.6 Protecting sensitive information

In scenarios such as customer service, such as insurance company chatbot, ensuring the privacy and security of sensitive data is paramount. Implementing a monitoring system, such as Guardium, can help oversee data integrity and protect against accidental data exposure. This monitoring ensures that user-specific data is handled safely, preserving confidentiality while still allowing the model to deliver personalized responses.

7.6.1 IBM Guardium for AI - Discovering and protecting sensitive data

IBM Security Guardium has powerful capabilities to establish data protection for AI systems.

Alongside its ability to provide data *guardrails*, it is also capable of getting full visibility into AI deployments, discover, classify, protect, manage risk, and comply with regulations and frameworks. Figure 7-8 shows the IBM Guardium AI Security overview.

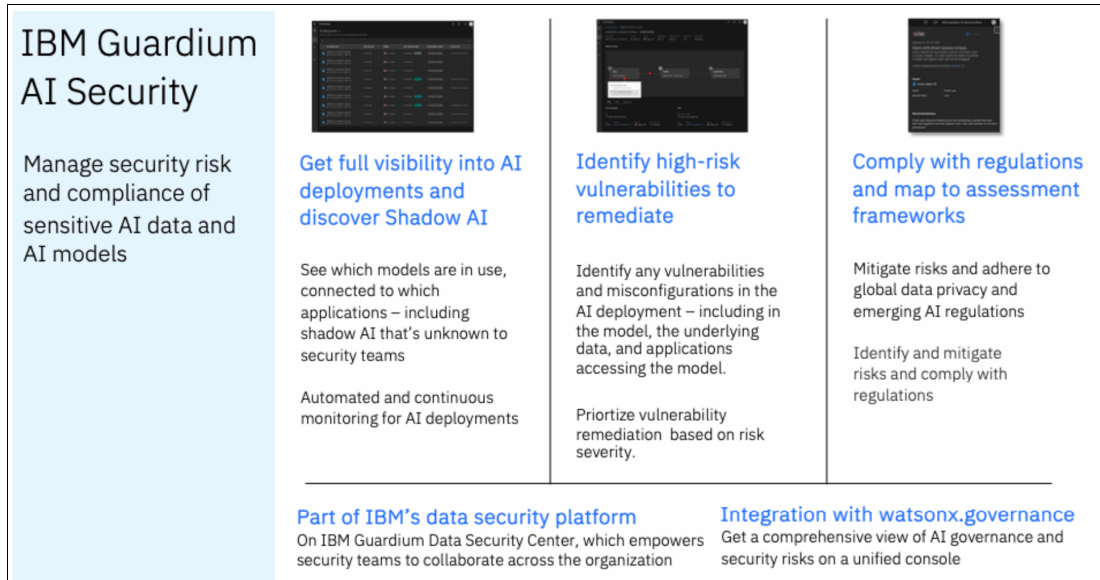


Figure 7-8 IBM Guardium AI Security overview

Figure 7-9 shows a RAG customer example securing data with Guardium Data Protection (Protect data at Ingestion - Source and Use - Vector Database).

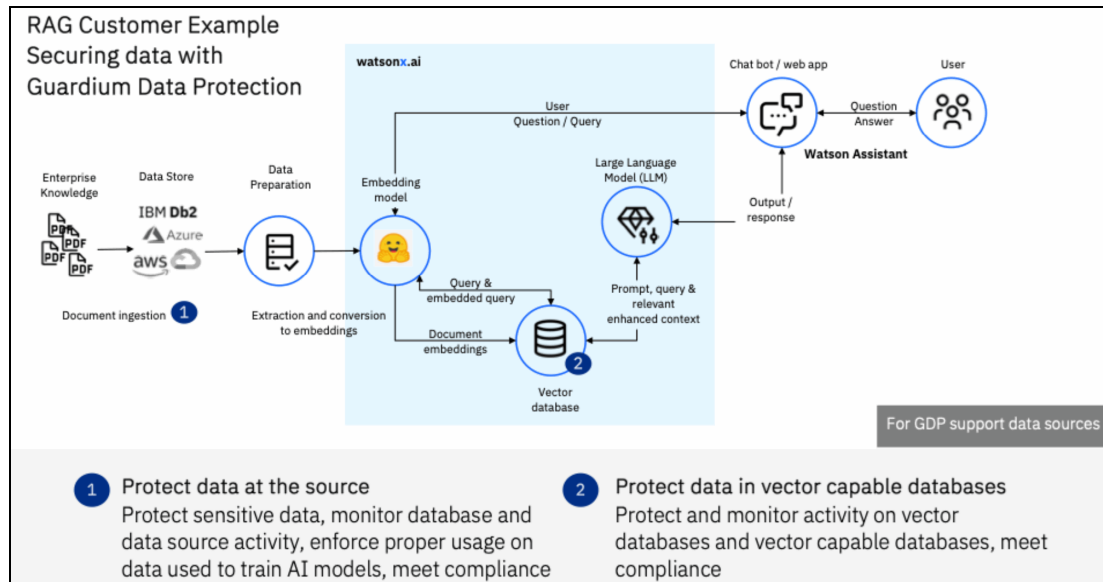


Figure 7-9 RAG customer example securing data with Guardium Data Protection

Using this methodology ensures that data is protected through its lifecycle. Figure 7-10 on page 114 shows IBM Guardium AI-specific use cases.

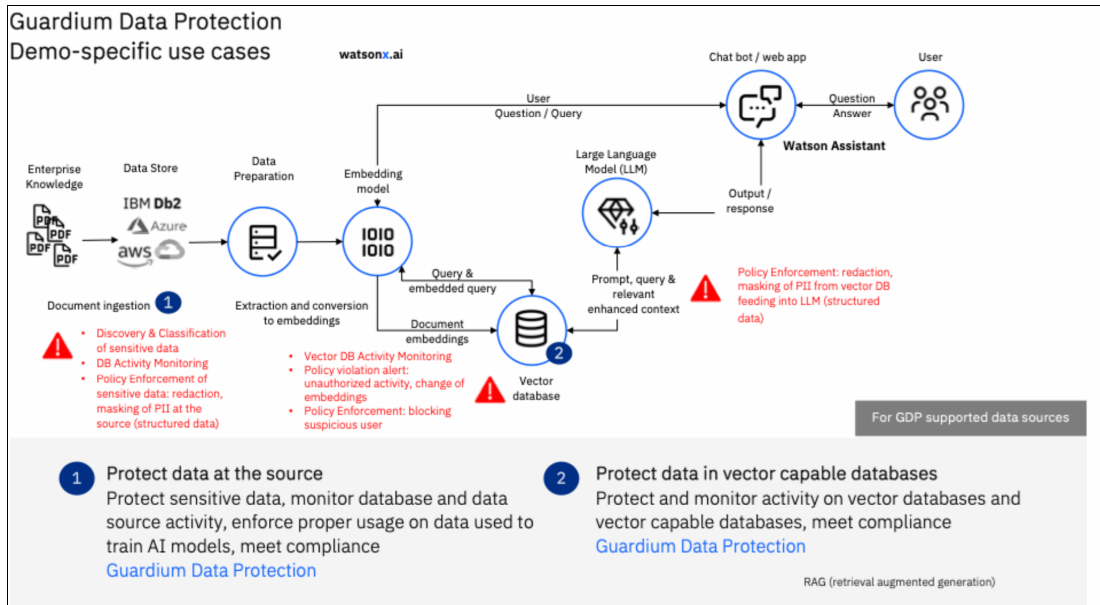


Figure 7-10 IBM Guardium AI-specific use cases

You can use Guardium to monitor and protect the data sources you plan to use to train your AI models to ensure the data is not tampered with prior to use. See Figure 7-11.

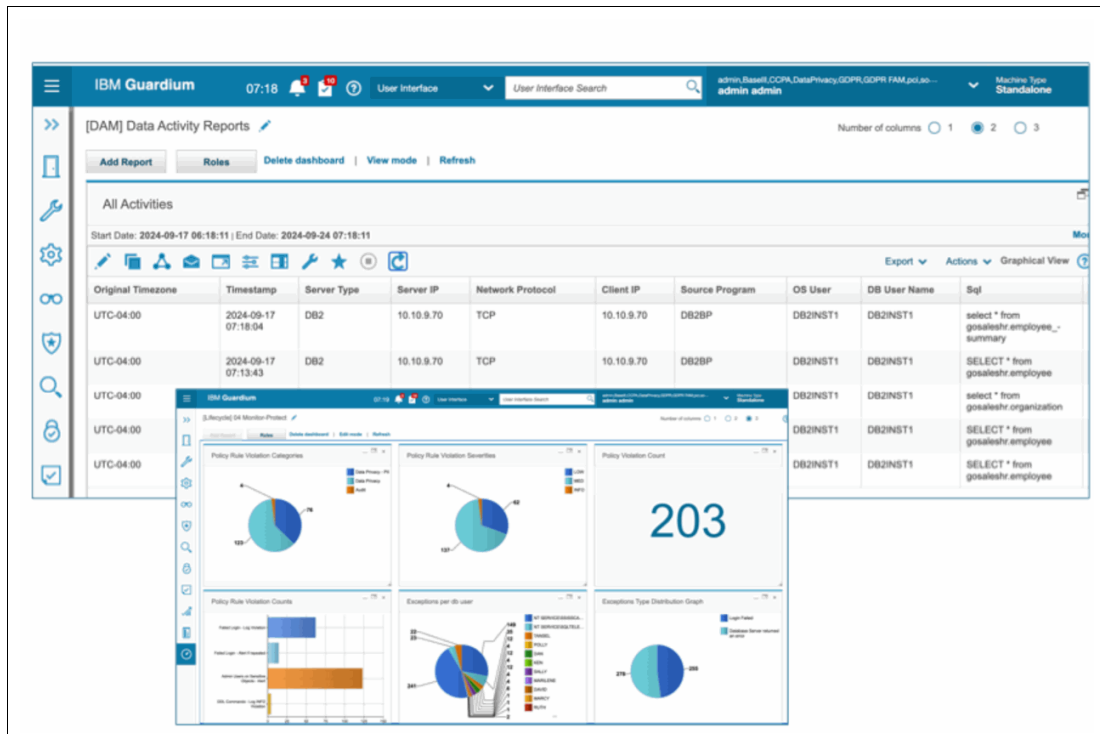


Figure 7-11 Guardium for AI Data Monitoring and Protection

Data privacy concerns: Training LLMs involves the processing of vast amounts of data, and this raises several privacy concerns. In the case of a model that inadvertently retains and utilizes sensitive user data for fine-tuning or future training, organizations run the risk of exposing confidential information which could lead to fines and violation of regulatory compliance standards.

7.6.2 The interrelationship of data and models

At the foundation of any AI system lies data and the models built from that data. A model is a function of the data used during its training process. LLMs, for instance, are trained on vast datasets consisting of billions of tokens. The probability distribution these models generate reflects the information encoded in their training data. Thus, understanding the datasets utilized in training is crucial for maintaining model integrity.

Risks associated with data

Data used for training AI models is not without its risks. If a dataset contains sensitive information such as personally identifiable information (PII) or illegal content, and this data goes undetected during preprocessing, it poses a significant liability. For example, if a model is trained on malicious data, there is a risk of that model producing harmful outputs or disclosing sensitive information during normal operation. This scenario highlights the concept of data poisoning, where harmful data is intentionally injected into the training dataset.

Securing data in AI systems

To secure data used in AI pipelines, organizations must adopt comprehensive data curation and filtering processes that precede the model training phase.

- ▶ **Data curation:** Implementing strict data sourcing policies to ensure that only quality, relevant, and legal datasets are utilized for training models.
- ▶ **Data filtering:** Establishing mechanisms for filtering out potentially harmful information-especially sensitive data-before training begins. This process should verify that any PII or illegal content is eliminated.
- ▶ **Transparency:** Utilizing transparent models (like the IBM approach with its Blue Pile) allows organizations to trace back model performance to the datasets used, enabling easier identification of data-related issues.

The role of models: Once data is curated, models can be built based on this information. However, the model itself can become an avenue for attack if it has unfettered access to sensitive resources or if users can manipulate how it interacts with systems.

7.6.3 Securing models: Addressing LLM injection

In this section we cover the security models for LLM injection.

LLM injection and prompt injection

After a model is deployed, it can still be vulnerable to attacks like LLM or prompt injection. This type of attack occurs when a malicious user finds ways to manipulate the model's functionality by crafting prompts that lead to inappropriate or harmful outputs.

User identity and access management

Businesses should implement strict identity verification mechanisms to ensure that only authorized users can invoke the model and access sensitive data. Effective identity management provides a layer of protection against unauthorized access.

Input validation

Implementing filtering mechanisms-such as checking for regular expressions that identify sensitive data-can help mitigate risks associated with bad input.

- ▶ **Moderation tuning:** Models should be pre-tuned to detect and reject harmful prompts or instructions. This requires extensive testing to ensure that the model can discern between acceptable and unacceptable requests effectively.
- ▶ **Classifiers:** Utilizing classifiers based on embeddings to analyze the semantic similarity of inputs can allow organizations to filter out harmful prompts before they reach the model.

Infrastructure security

While securing data and models is paramount, securing the infrastructure that supports AI pipelines is equally important. This includes protecting the underlying architecture where models are trained and deployed.

Access controls

Implementing strong access control measures ensures that only authorized personnel can modify systems or access sensitive data.

Continuous monitoring

Employing continuous monitoring tools can help organizations detect any unusual behavior or unauthorized access in real time, mitigating potential breaches.

Encrypting data

Encrypting sensitive data both at rest and in transit can provide an additional layer of security, reducing the risk of data theft or exposure.

Patch management

Regularly updating and patching systems to safeguard against known vulnerabilities is critical for maintaining a secure infrastructure.

7.6.4 Conclusion

Securing AI pipelines is an ongoing challenge that requires a multifaceted approach to address vulnerabilities inherent in data, models, and infrastructure. By instilling rigorous data curation processes, applying stringent model security measures, and fortifying the underlying infrastructure, organizations can significantly reduce their risk exposure. As AI systems continue to evolve, proactive security practices will play a pivotal role in ensuring that the transformative potential of AI is harnessed safely and responsibly.

7.7 Securing Generative AI: Threat vectors, data protection, and advanced applications

As generative AI technology continues to evolve and integrate into various business applications, the need to secure these systems against potential threats becomes increasingly important. Understanding the various attack vectors, implementing data protection measures, and navigating the complexities of advanced applications like agents are essential for organizations leveraging generative AI. In this section we explore these security aspects, offering insights into how to best protect generative AI systems.

7.7.1 Advanced applications: The role of agents in relevance to security

In this section we describe the role of agents in relevance to security.

7.7.2 Understanding agents in generative AI

In the context of generative AI, agents are complex systems that leverage multiple AI models and APIs to perform tasks autonomously. Agents enhance the capabilities of LLMs by combining their results with real-time data access, allowing them to execute specific functions and generate nuanced responses.

While agents present powerful opportunities for automation and efficiency, they also expand the attack surface for potential abuses. Given their ability to connect to various external systems, if not safeguarded properly, they could serve as conduits for unauthorized data requests or actions. The following approaches can be used as a mitigation strategy:

- ▶ **Zero Trust architecture:** Implement a zero-trust approach to ensure that every request for data or action is authenticated, monitored, and validated regardless of where it originates.
- ▶ **Permission control:** Enforce strict permissions on the actions an agent can execute in connection to databases and APIs, ensuring that the system can generate only outputs that fall within predetermined security boundaries.

7.7.3 Threat vectors in Generative AI

In this section we discuss threat vectors in Generative AI.

Prompt injection attacks

One of the most significant threats to generative AI systems is prompt injection, where users can manipulate inputs to elicit harmful or unintended responses from the AI model. Because of the open-ended nature of inputs allowed in many generative AI interfaces, users may craft queries designed to bypass established safety protocols, leading to outputs that could be abusive, hateful, or otherwise inappropriate. The following approaches can be used as a mitigation strategy:

- ▶ **Input validation:** Implement rigorous validation processes that can identify potentially harmful prompts. Regular expressions and classifiers can help categorize and filter out undesirable inputs before they reach the AI model.
- ▶ **User identity controls:** Employ strong identity management to ensure only authorized users are interacting with the AI, reducing the risk of manipulation from malicious actors.

Data poisoning

Data poisoning occurs when malicious data is injected into training datasets, compromising the integrity and safety of the model. This can happen during the pre-training phase, where bad data may lead a model to generate results that expose standards or private information. The following approaches can be used as a mitigation strategy:

- ▶ **Data curation:** Establish strict data sourcing and curation processes to filter out harmful data before it is used for training. Transparency in the training dataset is critical to identifying potential risks.
- ▶ **Continuous monitoring:** Regularly examine datasets even after initial training to ensure no new harmful data has been introduced.

Output induced data leakage

Generative AI models can inadvertently output sensitive information that was used in their training datasets. This is a significant risk, particularly for organizations handling confidential or regulatory-sensitive data. The following approaches can be used as a mitigation strategy:

- ▶ **Data masking:** Utilize tools to obfuscate sensitive information in both the training and operational datasets. Implement access controls to ensure that sensitive data is not inadvertently exposed through AI outputs.
- ▶ **Guardium:** Tools like IBM Guardium can help monitor transactions and flag any outputs that contain sensitive information, enhancing data security through structured oversight.

Figure 7-12 summarizes how Guardium Data Protection works for vector capable databases.

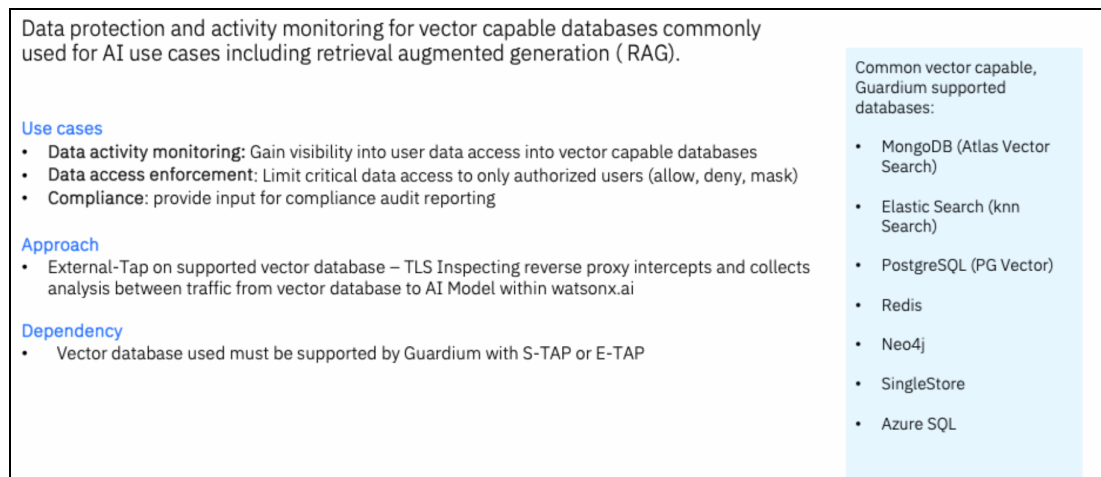


Figure 7-12 Guardium Data Protection for vector capable databases

7.7.4 Data protection mechanisms

Protecting the data that underpins generative AI systems is essential for ensuring both compliance and operational integrity. Here are key strategies to consider:

- ▶ **Secure data storage:** Use encrypted storage solutions for both raw data and processed data, ensuring that unauthorized access is prevented.
- ▶ **Access controls:** Implement role-based access controls (RBAC) to restrict who can view and modify sensitive datasets. This minimizes the risk of data exposure.
- ▶ **Monitoring and auditing:** Continuous real-time monitoring of data access and modifications helps identify and respond to suspicious activities swiftly.

Figure 7-13 shows how to extend exiting security across the underlying AI infrastructure.

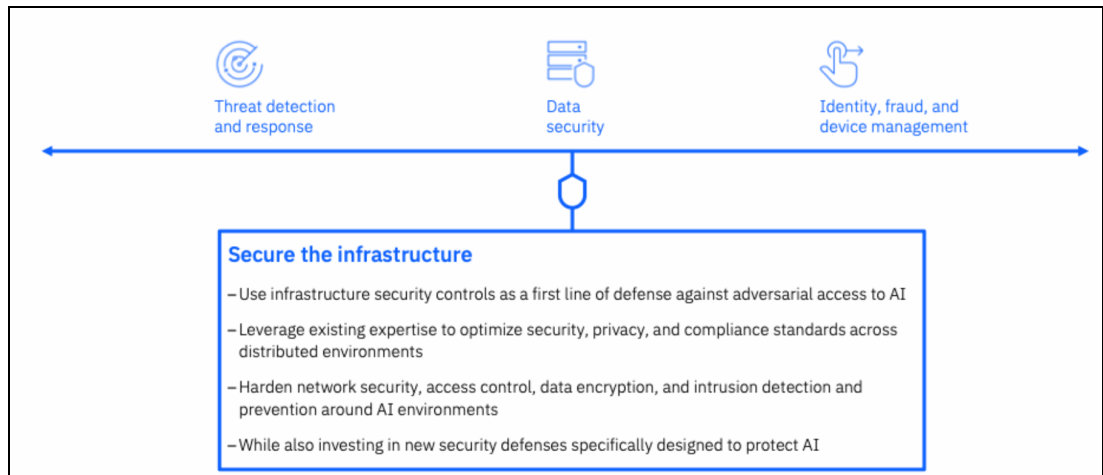


Figure 7-13 Extend existing security across the underlying AI infrastructure

As organizations increasingly deploy generative AI systems, understanding and mitigating the associated risks is crucial. From addressing threat vectors such as prompt injections and data poisoning to ensuring robust data protection measures and navigating advanced applications like agents, companies must adopt a comprehensive security strategy. By harnessing tools like IBM's Guardium and implementing best practices for data management and access control such as IBM Verify, organizations can better secure their generative AI initiatives and maximize the benefits of this transformative technology.

7.8 Security measures for LLMs

Addressing the risks posed by LLMs involves a multi-faceted approach to security and governance. Organizations that implement AI require robust governance and security frameworks to ensure ethical, effective, and compliant usage. In this section we will discuss what is necessary in AI governance and security.

AI governance has the following requirements:

- ▶ **Monitoring and evaluations:** Predictive model monitoring is the process of regularly evaluating models to ensure ongoing fairness, accuracy, and stability, detecting any performance degradation over time.
- ▶ **Explainability:** Model metrics tracking is the process of gathering and analyzing key performance indicators and feature importance scores to improve the understanding of model decisions and identify potential biases.
- ▶ **Risk management and compliance:** AI activity management is the process of establishing systems to monitor and manage AI deployments, from initial requests to post-deployment evaluation, ensuring adherence to regulatory standards and risk mitigation strategies.

AI security has the following requirements:

- ▶ **Visibility:** Model deployment tracking entails maintaining a comprehensive inventory of AI models deployed within the organization, including their locations, use cases, and potential risks, even in cases of unauthorized or undocumented deployments.

- ▶ **Data security:** Protection of sensitive data entails implementing robust security measures to safeguard sensitive data used in AI models, including encryption, access controls, and regular security assessments.
- ▶ **Model and application security:** Security controls involve applying advanced security techniques, including threat modeling, intrusion detection systems, and secure coding practices, to safeguard AI applications and models from cyber threats and data breaches.

7.9 Integrated security technologies for AI management

Security management for AI can be structured into three key activities:

- ▶ **Discover:** Utilize automated tools for visibility into all AI models in use, ensuring awareness of both known and unknown instances.
- ▶ **Assess:** Conduct vulnerability assessments for AI environments, focusing on models, data, and applications to identify any misconfigurations or risks.
- ▶ **Secure:** Implement industry-standard security frameworks (for example, Open Web Application Security Project (OWASP) Top 10 list for LLM) and ensure compliance with relevant data privacy and AI regulations.

Figure 7-14 shows how IBM Guardium AI Security helps organizations manage security risk and compliance of sensitive AI data and AI models.

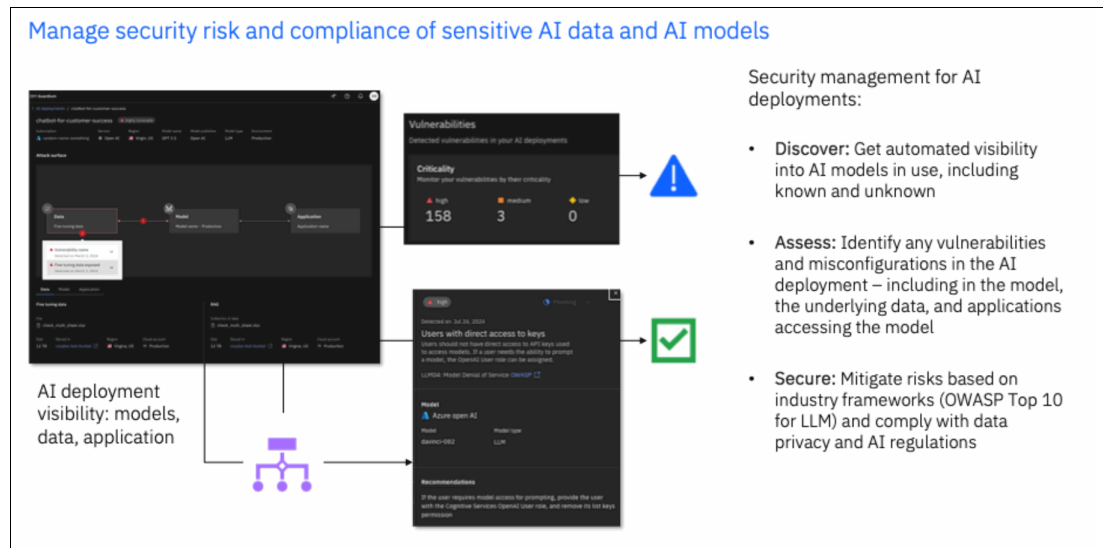


Figure 7-14 IBM Guardium AI Security

7.9.1 Toolkit for AI Governance

An effective toolkit should include:

- ▶ **Compliance management:** Ensure adherence to emerging global safety and transparency regulations, potentially offering a *nutrition label* for AI that outlines compliance with standards.
- ▶ **Risk monitoring:** Establish proactive measures for monitoring and managing issues related to fairness, bias, drift, and other performance metrics to protect organizational reputation.

- **Lifecycle management:** Develop governance frameworks for overseeing the operation of AI models sourced from various providers (for example, IBM, open-source) to ensure they are utilized effectively and confidentially.

This holistic approach allows organizations to effectively govern, secure, and monitor their AI initiatives, ensuring ethical use and compliance with relevant regulations.

Figure 7-15 shows how IBM watsonx.governance helps organizations accelerate responsible and explainable AI workflows.

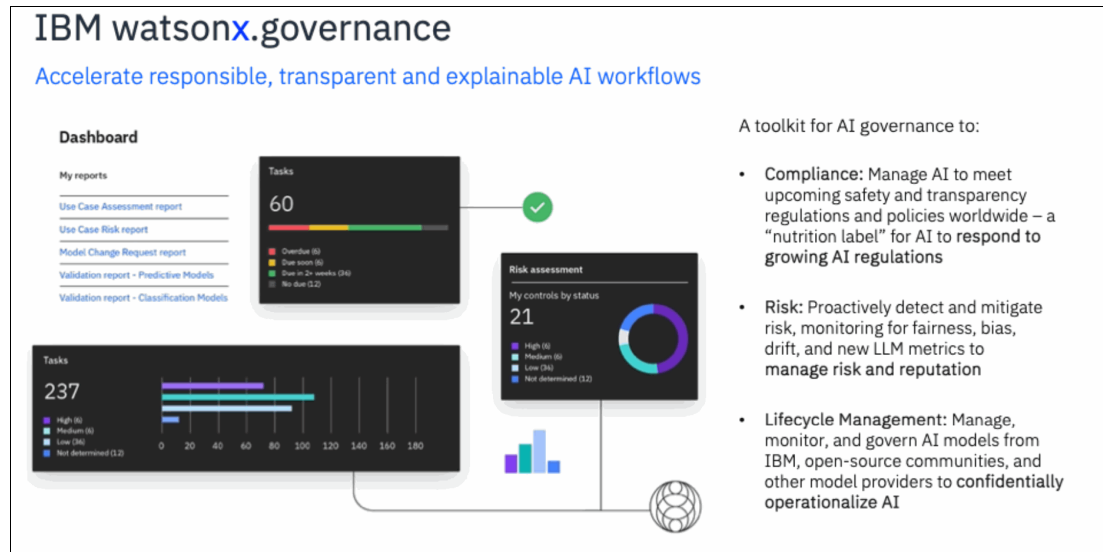


Figure 7-15 IBM watsonx.governance

Figure 7-16 illustrates how IBM watsonx.governance and IBM Guardium AI Security collaborate to provide a unified approach to AI governance, security, and monitoring.

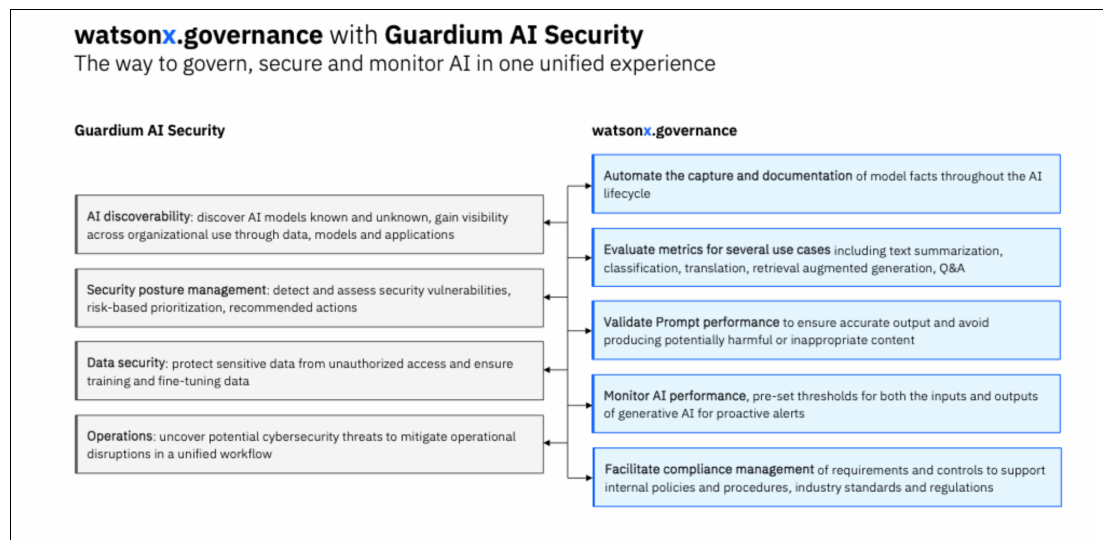


Figure 7-16 IBM watsonx.governance and IBM Guardium AI Security: A unified approach

Figure 7-17 on page 122 further details the integration use cases IBM watsonx.governance with IBM Guardium AI Security.

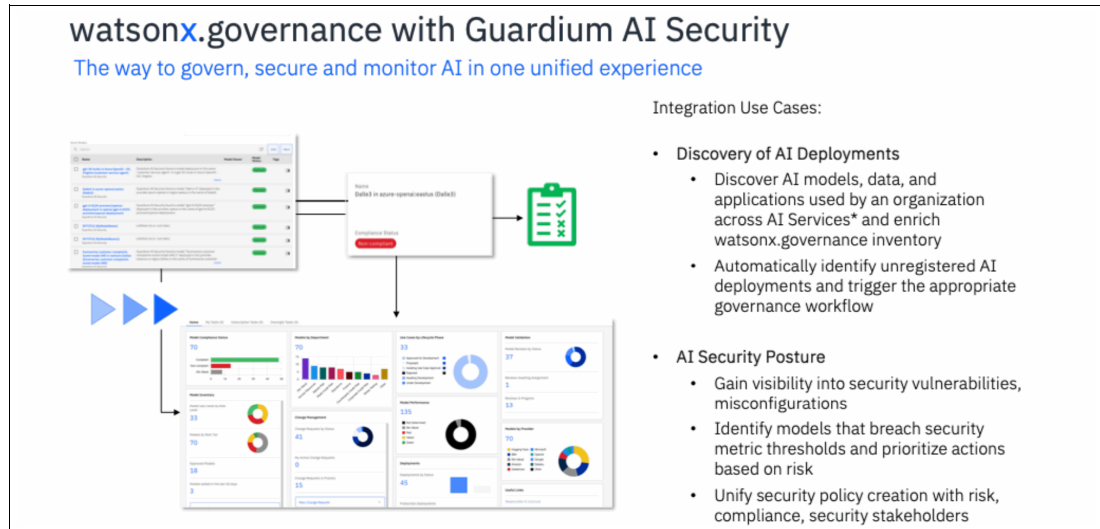


Figure 7-17 Integration use cases IBM watsonx.governance with IBM Guardium AI Security

7.10 IBM Guardium AI Security: Manage data model security risk demo overview and key capabilities

As discussed in this chapter, IBM Guardium AI Security empowers you to discover generative AI deployments, mitigate vulnerabilities in AI models, protect sensitive data, and address regulatory requirements—all on a single dashboard. To have a closer look at the platform's capabilities see this demo: [IBM Guardium AI Security: Manage data model security risk](#).

The following are some of the capabilities shown in this demo:

- ▶ **Continuous scanning:** Automatically detects all AI deployments, including unknown *shadow* AIs, when cloud accounts are connected.
- ▶ **Vulnerability assessment:** Identifies vulnerabilities in AI models and associated applications, assigning criticality scores for prioritization and enabling easy export for reporting.
- ▶ **Detailed insights:** Provides in-depth details on vulnerabilities, including significant concerns like direct access to sensitive keys, along with links to the OWASP Top 10 for further investigation.
- ▶ **Recommended actions:** Offers actionable recommendations to address identified vulnerabilities.

For more information, you can also visit the [IBM Guardium product page](#).

Related publications

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

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
- ▶ *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
- ▶ *Data Resiliency Designs: A Deep Dive into IBM Storage Safeguarded Snapshots*, REDP-5737

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 Storage Sentinel anomaly scan software overview
<https://www.ibm.com/docs/en/storage-sentinel/1.1.9?topic=product-overview>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



Cyber Resiliency with IBM Storage Sentinel and IBM

SG24-8562-00

ISBN 073846192X



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



Cyber Resiliency with IBM Storage Sentinel and IBM Security

SG24-8562-00

ISBN 073846192X



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

Redbooks

Cyber Resiliency with IBM Storage Sentinel and IBM Security

SG24-8562-00

ISBN 073846192X



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

Redbooks

Cyber Resiliency with IBM Storage Sentinel and IBM Security

(0.2" spine)

0.17" <-> 0.473"

90 <-> 249 pages

(0.1" spine)

0.1" <-> 0.169"

53 <-> 89 pages



Cyber Resiliency with IBM Storage Sentinel and IBM

SG24-8562-00

ISBN 073846192X

(2.5" spine)
2.5" <-> mmm.n"
1315 <-> mmm pages



Cyber Resiliency with IBM Storage Sentinel and IBM Security

SG24-8562-00

ISBN 073846192X

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





SG24-8562-00

ISBN 073846192X

Printed in U.S.A.

Get connected

