IBM® Storage

IBM Spectrum Sentinel protecting SAP HANA® with IBM FlashSystem

Version 1.0

IBM Storage Team

Contents

About this document

This IBM® White Paper focuses on malware detection within an SAP HANA database environment by using IBM Spectrum Sentinel®. It also highlights how to perform a cyber resilience restore in response to a cyberattack. The workflow that is presented here uses IBM Spectrum® Copy Data Management as orchestration software to start IBM FlashSystem® Safeguarded Copy functions. The Safeguarded Copy creates an immutable copy of the data in an air-gapped form on the same IBM FlashSystem for isolation and eventual quick recovery.

Finally, this document outlines the steps that are involved to create a backup and restore Job by using IBM Spectrum® Copy Data Management with various actions.

Scope

This document was developed using the following software tools:

- SAP HANA 2.0 database
- IBM Spectrum Sentinel 1.1.1

This technical report does not:

- Replace any official manuals and documents issued by IBM
- Explain installation and configuration of SAP HANA

Executive Summary

Our client's most critical applications need the best protection. Since even backup data are attractive targets for attackers, we need to ensure that a safe restore point is always available - even if an active attack is underway. This can be achieved by testing data copies against malware signatures and other signs of cyber-attacks, like specific changes incurred by malware, or even data encryption.

IBM Spectrum Sentinel is a cyber resiliency solution designed to help organizations enhance ransomware detection and incident recovery. IBM Spectrum Sentinel automates the creation of immutable backup copies of application data, then uses machine learning to detect signs of possible corruption and generate forensic reports that help you quickly diagnose and identify the source of the attack.

The IBM Spectrum Sentinel offering consists of IBM Spectrum Copy Data Management and an anomaly scanning software. The creation of immutable (“Safeguarded”) copies requires storage systems such as IBM FlashSystems running IBM Spectrum Virtualize.

IBM Spectrum Sentinel generates application consistent, storage-based backups and checks these backups, proving that they do not contain malware. Currently it supports SAP HANA and Epic Cache DB, for 2023 additional support for VMware environments, Microsoft SQL, Oracle and Epic on AIX is planned. If an attack occurs, IBM Spectrum Sentinel helps to identify the best Safeguarded Copy to use. It also automates the process to restore SAP HANA data to online volumes. Because a restore action uses the same snapshot technology, it is much faster than the use of offline or cloud-based copies.

Prerequisites

For a list of all IBM storage systems certified for SAP HANA production please refer to:

- Certified and Supported SAP HANA® Hardware

It is assumed that you are familiar with and have basic knowledge of the following products:

- IBM FlashSystem
- SAP HANA database

IBM FlashSystem

The IBM FlashSystem family combines the performance of flash storage and end-to-end Non-Volatile Memory Express (NVMe) with the reliability and innovation of IBM FlashCore technology, the ultra-low latency of Storage Class Memory (SCM), advanced data services enabled by IBM Spectrum Virtualize and AI predictive storage management and proactive support by Storage Insights. The complete FlashSystem family is shown in Figure 1.

The protocol inside FlashSystem NVM Express (NVMe) is an optimized, high-performance scalable host controller interface designed to address the needs of systems that utilize PCI Express-based solid-state storage. The NVMe protocol is an interface specification for communicating with storage devices. It is functionally analogue to other protocols, such as Serial attached SCSI (SAS). However, the NVMe interface is designed for extremely fast storage media, such as flash-based solid-state drives (SSDs) and other low-latency non-volatile storage technologies.

IBM FlashSystems include IBM Spectrum Virtualize software and introduce remarkable new features in comparison to the predecessor models:

- End-to-end **NVMe** support: NVMe is a logical device interface standard from 2011 for accessing non-volatile storage media that is attached via a PCI Express bus.
- **Lower latencies** through RDMA: Direct memory access from the memory of one node into that of another without involving either one's operating system.
- **Data reduction pools** (DRP) represent a significant enhancement to the storage pool concept. With the introduction of data reduction technology, compression, and deduplication, it has become more of a requirement to have an uncomplicated way to stay “thin”.
- **FlashCore Modules** (FCMs) or industry standard NVMe drives can be used for the IBM FlashSystem. If the FCM option is chosen, the user can take advantage of built-in hardware compression, which will automatically compress the stored data written to the FCM modules.
- **Thin-provisioned** IBM FlashCopy only occupies disk space, if updates are made to the source or target data, and not for the entire capacity of a volume copy.
- The **HyperSwap** feature allows each volume to be presented by two IBM FlashSystems. This high-availability configuration tolerates combinations of node and site failures, using the host multipathing driver that is available for the IBM FlashSystems.
- The IBM FlashSystem supports the new low latency, high speed **Storage Class Memory** (SCM). SCM is a non-volatile memory device that performs faster (~10µs) than traditional NAND SSDs (100µs), but slower than DRAM (100ns).
- **IBM Storage Insights** is an IBM software product that enhances the monitoring capability of the IBM FlashSystem family and supplements the views available in the standard GUI.

- The **Safeguarded Copy** (SGC) function enables the creation of immutable point-in-time copies of volumes that cannot be changed or deleted by user errors, malicious actions, or ransomware attacks. IBM Spectrum Sentinel is using the Safeguarded Copy function for creating immutable copies.

To deliver the highest level of performance along with the ability to provide external storage virtualization, our clients will require Non-Volatile Memory express (NVMe) based solutions. Although these arrays can use industry standard NVMe solid-state drives (SSDs), it is assumed that most clients will opt for IBM FlashCore modules (FCMs) which offer significant differentiation and competitive advantages. FlashSystem 5200, 7300, 9500 and 9500R are truly future proof, since they support essentially all drive types, including storage class memory (SCM), in order to optimize a client's data economics. Up to 12 SCM drives can be placed in any IBM FlashSystem solution that supports NVMe.

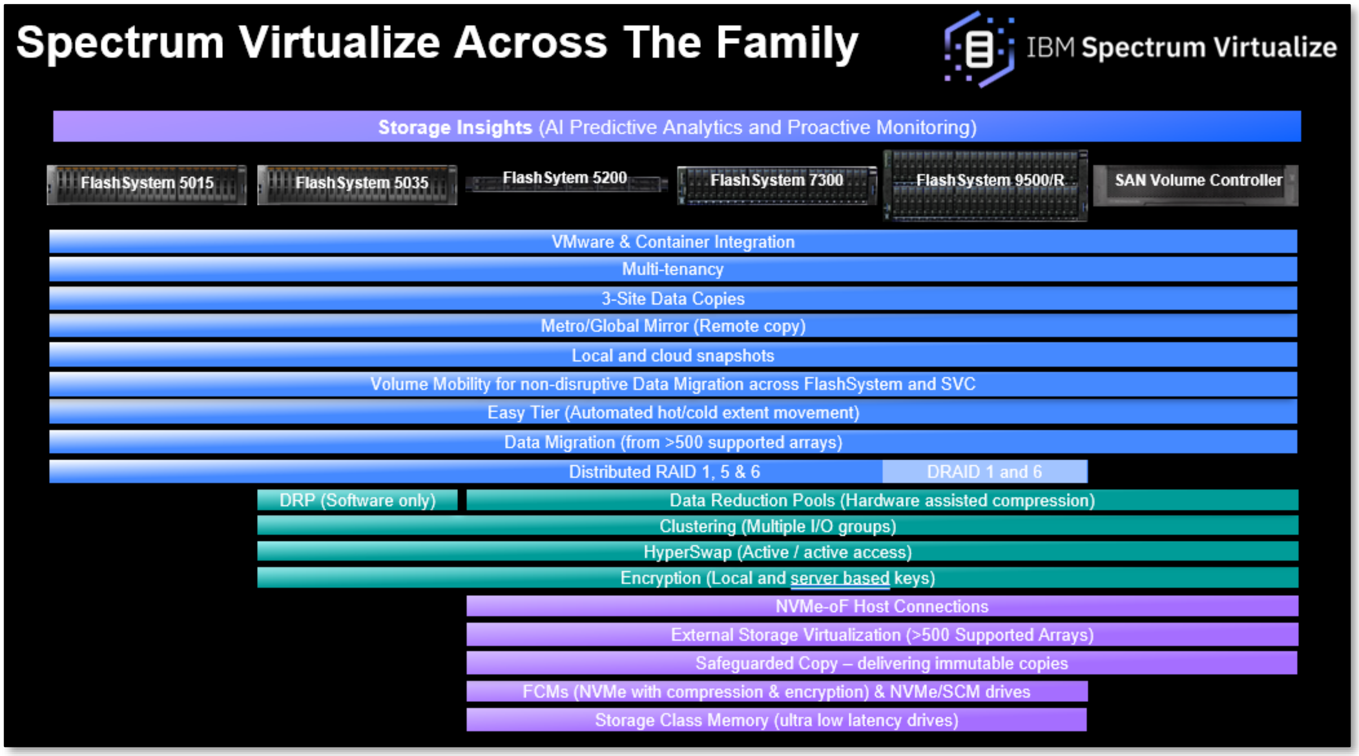

Spectrum Virtualize Across The Family

IBM Spectrum Virtualize

Storage Insights (AI Predictive Analytics and Proactive Monitoring)

FlashSystem 5015 | FlashSystem 5035 | FlashSytem 5200 | FlashSystem 7300 | FlashSystem 9500/R | SAN Volume Controller

VMware & Container Integration

Multi-tenancy

3-Site Data Copies

Metro/Global Mirror (Remote copy)

Local and cloud snapshots

Volume Mobility for non-disruptive Data Migration across FlashSystem and SVC

Easy Tier (Automated hot/cold extent movement)

Data Migration (from >500 supported arrays)

Distributed RAID 1, 5 & 6 | DRAID 1 and 6

DRP (Software only) | Data Reduction Pools (Hardware assisted compression)

Clustering (Multiple I/O groups)

HyperSwap (Active / active access)

Encryption (Local and server based keys)

NVMe-oF Host Connections

External Storage Virtualization (>500 Supported Arrays)

Safeguarded Copy – delivering immutable copies

FCMs (NVMe with compression & encryption) & NVMe/SCM drives

Storage Class Memory (ultra low latency drives)

Figure 1: IBM FlashSystem family and features

For more detailed information about the IBM FlashSystem family please refer to the IBM FlashSystem Best Practices and Performance Guidelines:
http://www.redbooks.ibm.com/redpieces/abstracts/sg248503.html?Open

Safeguarded Copy

The IBM FlashSystem Safeguarded Copy feature creates safeguarded backups that are not accessible by the host system. It also protects these backups from corruption or any other change that can occur in the production environment. A Safeguarded Copy schedule can be defined to regularly create multiple backups (such as hourly or daily).
Safeguarded Copy can create backups with a higher frequency and capacity compared to IBM FlashCopy® volumes. Creating Safeguarded backups also affects performance less than the multiple target volumes that are created by IBM FlashCopy.

The Safeguarded Copy function provides backup copies to recover data if a logical corruption occurs or primary data is destroyed. Safeguarded Copy uses capacity for the safeguarded backups, and for recovery volumes that are needed to access the data copies, as shown in **Figure 2**.

A recovery volume is used to restore a backup copy for host access while production continues to run on the production volume. The recovery volume is the target volume for a Safeguarded Copy recovery, which enables a previous backup copy to be accessed by a host that is attached to this volume. The recovery volume is thin provisioned by default.

Managing Safeguarded Copy for SAP HANA is supported by IBM Spectrum Copy Data Management version 2.2.17 or later in combination with Spectrum Virtualize 8.5.1. The management software helps to create and recover backups and define policies for expiration. Safeguarded Copy is supported by the IBM FlashSystem 5200 or larger models, as shown in **Figure 1**.

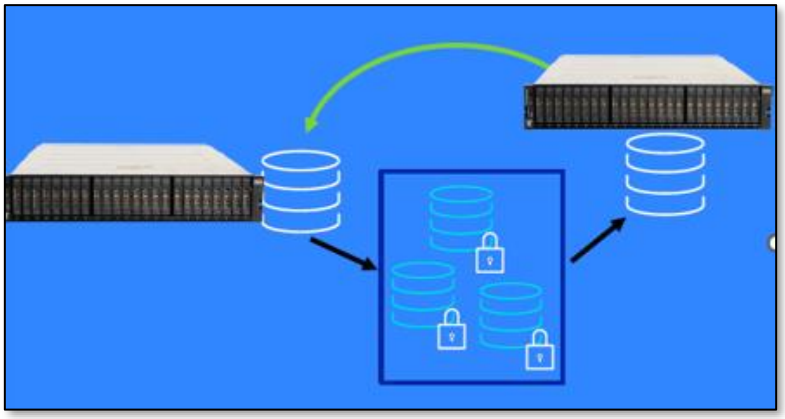

Figure 2: Safe Guarded Copy workflow

SAP HANA

SAP HANA (High-performance ANalytic Appliance) is a multi-model database that stores data in the server’s main memory instead of keeping it on a disk. This results in data processing that is magnitudes faster than that of disk-based data systems, allowing for advanced, real-time analytics. Serving as a platform for enterprise resource planning (ERP) software and other business applications, SAP HANA can be placed on premises, in the cloud, or in a hybrid cloud system.

SAP HANA’s in-memory, multi-model data management engine takes full advantage of the capabilities of the server hardware to minimize data movements, thus increasing application speed and agility as it analyzes real-time data.

Depending on the needs of an enterprise, SAP HANA can be deployed on premises, in the cloud, or as a hybrid system, blending the privacy and control of an on-premises system with the lower cost, greater memory, and increased access of the cloud. Its ability to efficiently process enormous amounts of data makes it easily scalable to suit a growing business without sacrificing security or stability.

On the SAP HANA platform, developers can build their own tools and applications that integrate business logic, control logic, and the database layer with unprecedented performance.

IBM Spectrum Sentinel

IBM Spectrum Sentinel is a cyber resiliency solution designed to help organizations enhance ransomware detection and incident recovery. IBM Spectrum Sentinel automates the creation of immutable backup copies of your data, then uses machine learning to detect signs of possible corruption that help you quickly diagnose and identify the source of the attack. Because IBM Spectrum Sentinel can intelligently identify infected backups, an organization can quickly identify the most recent verified and validated backup copies and minimize recovery time after a cyber-attack.

Building on the capabilities of IBM Safeguarded Copy, IBM Spectrum Sentinel uses Safeguarded Copy snapshots to create a secure and isolated backup. It periodically checks the data copies for evidence of malware or ransomware and sends the scanning results back to Copy Data Management which records them in its database. Since the copies are immutable by design, any malware or ransomware cannot remove, alter, or encrypt Safeguarded Copy snapshots.

Technically, IBM Spectrum Sentinel is part of the IBM Cyber Vault Blueprint. It uses the component database and the backup functionality of IBM Spectrum Copy Data Management for orchestrating the creation of SAP HANA application consistent IBM FlashSystem Safeguarded Copies and leverages a ransomware scanning engine for automated ransomware detection.

IBM Spectrum Copy Data Management

IBM Spectrum Copy Data Management provides a powerful suite of copy management services that can simplify copy management and provide data protection solutions for multi-cloud container environments through its leading-edge snapshot capabilities. It leverages copy services within an enterprise’s existing storage and hypervisor infrastructure for visibility of data copies in almost any storage environment. It provides rapid, self-service data access for operational use cases, DR, Test/Dev, and business analytics.

IBM Spectrum Copy Data Management is deployed as a preconfigured VMware® virtual machine appliance. It works agentless. All orchestrated environments like storage, application, hosts and VMware environment are managed by using their own interfaces.

Environments are being added into a catalog, which is used by IBM Spectrum Copy Data Management to orchestrate backups and maintenance operations. These are the base of further business use cases such as automated disaster recovery, DevOPs or business analytics.

Anomaly Scanning Engine

Real-time cyber protection solutions are designed to protect from an attack. However, these solutions cannot prevent 100% of attacks. Anomaly scan software adds a layer of protection to real time solutions. It will find corruption that occurs when an intruder has successfully penetrated the data center.

The scanning software component of IBM Spectrum Sentinel provides scanning for signatures of known malware and will identify corruption or encryption due to malicious code. It will feed back the scan results to IBM Spectrum Copy Data Management for inclusion into its backup catalog, enabling clients to automate recovery after an event.

The scanning engine identifies corruption and malware signature by analyzing the file system metadata with a Machine Learning Model (MLM), which is trained using information from real world cyber-attacks. In addition, it also checks the integrity of databases by examining database pages and allocation tables to ensure that all allocated database pages are present and located in their correct position. It will also verify if checksums or CRC codes are correct, and ancillary fields within database pages. To avoid excessive false positive alerts, Machine Learning Model (MLM) has been designed to tolerate a small amount of database corruption that is commonly observed in production database systems.

Ransomware

Ransomware is an online attack perpetrated by cybercriminals 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 cybercriminals' 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 up.

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

- **Phishing emails and other social engineering attacks:** Phishing emails manipulate users into downloading and running a malicious attachment (which contains the ransomware 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***:* Cybercriminals often exploit existing vulnerabilities to inject malicious code into a device or network.

There are companies in the underground economy, that offer Ransomware as a Service (RaaS) and this market is growing.

Business impact

Ransomware victims and negotiators are reluctant to disclose ransom payment amounts. However, according to the report Definitive Guide to Ransomware 2022 (PDF, 966 KB) , ransom amounts have grown to seven-figure and eight-figure amounts. In more extreme cases, companies may pay as much as USD 40-80 million to get back access to their data. Another aspect is the reputational damage to the enterprise if the cyber-attack gets public. **Figure 3** shows an example of a ransomware message that a customer found in their system.

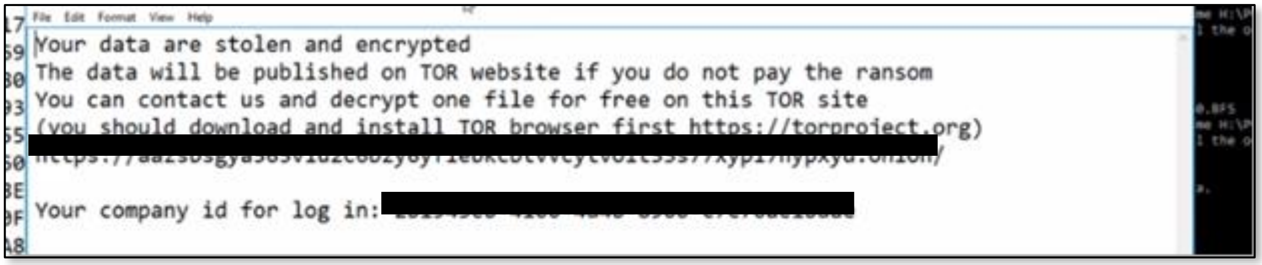

```
Your data are stolen and encrypted
The data will be published on TOR website if you do not pay the ransom
You can contact us and decrypt one file for free on this TOR site
(you should download and install TOR browser first https://torproject.org)
[redacted]

Your company id for log in: [redacted]
```

Figure 3: Ransomware attack message

SAP HANA Ransomware protection with IBM Sentinel

Building on the capabilities of IBM Safeguarded Copy, IBM Spectrum Sentinel periodically checks SAP HANA data copies for evidence of data damage caused by malware or ransomware. IBM Spectrum Sentinel then uses Safeguarded Copy snapshots to create a secure and isolated backup. Ransomware cannot remove, alter, or encrypt Safeguarded Copy snapshots, even with administrator capabilities.

There are several possible scenarios of cyber-attacks which can be protected by IBM Spectrum Sentinel.

Encrypting the database

This type of attack is a ransomware attack that targets the SAP HANA data. In this use case, critical data is encrypted by the attacker, which renders database operations useless. Then, ransom is demanded against the release of the data to bring back the database in operational mode.

Hiding malware in the database

Database users can easily and unknowingly download malicious or unauthorized software. What happens if malware successfully infects your environment and is not detected in time to prevent data corruption? In the past, data corruption was often detected because data was quickly found to be inaccessible or unusable. However, as malware has grown in sophistication, the infection has become more difficult to detect immediately and lies hidden until a later time.
IBM Spectrum Sentinel can discover infected backups that contain encrypted data. For example, immutable snapshots of Epic Cache or SAP HANA databases can be scanned, and potential corruption identified. With an accurate detection rate above 99%, Spectrum Sentinel can analyze the internal data structures of the database to detect corruption that would otherwise stay hidden.

Once ransomware is detected by IBM Spectrum Sentinel, it is critical to take immediate action to resolve any potential issues. You need to immediately contact your internal Security team and IBM Spectrum Sentinel Support to investigate the extent of the potential threat and fix it if necessary.

Configuring IBM Spectrum Copy Data Management

Datacenter setup

Figure 4 illustrates the environment used to perform the setup and tests with IBM Spectrum Sentinel and SAP HANA which are described in the following chapters.

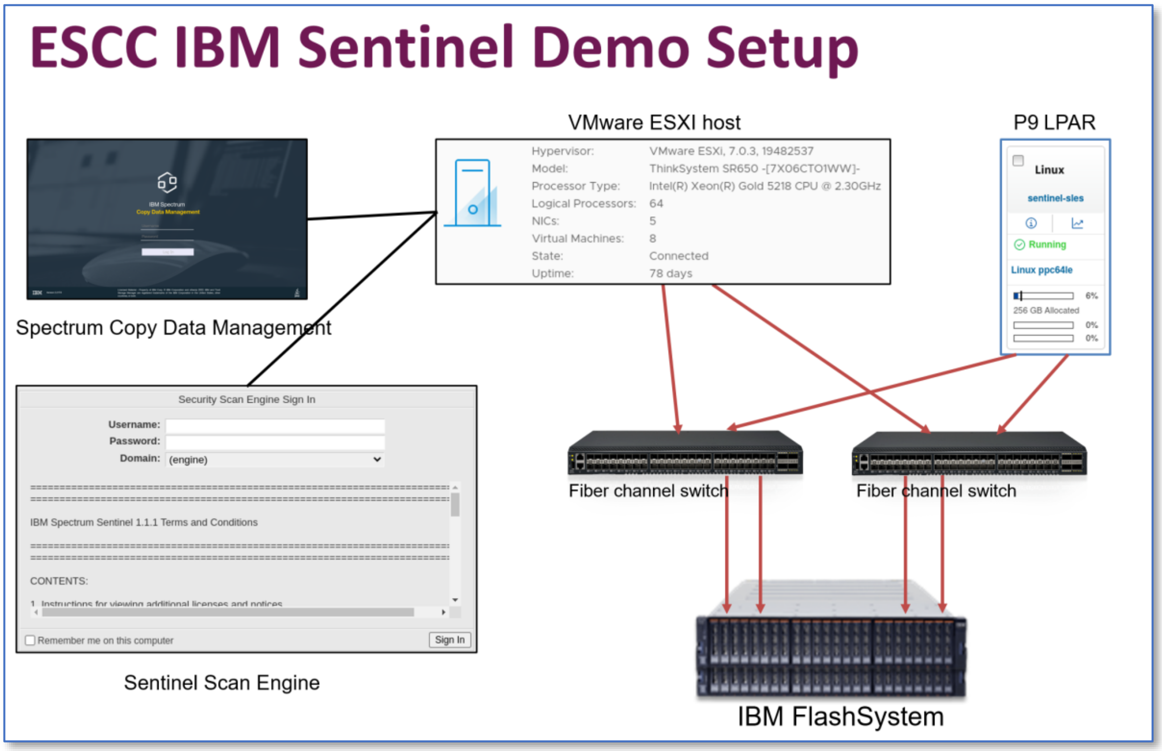

Hypervisor: VMware ESXi, 7.0.3, 19482537
Model: ThinkSystem SR650 -[7X06CTO1WW]-
Processor Type: Intel(R) Xeon(R) Gold 5218 CPU @ 2.30GHz
Logical Processors: 64
NICs: 5
Virtual Machines: 8
State: Connected
Uptime: 78 days

Security Scan Engine Sign In

Username:
Password:
Domain: (engine)

```
=====================================================================
=====================================================================

IBM Spectrum Sentinel 1.1.1 Terms and Conditions

=====================================================================
=====================================================================

CONTENTS:

1. Instructions for viewing additional licenses and notices
```

□ Remember me on this computer
Sign In

Figure 4: Demo setup for IBM Sentinel

Hardware components:

- IBM FlashSystem 7200 HyperSwap with 24x 4,36Tb FCM2 modules
- IBM Power9 H922 System with 2048Gb RAM
- Lenovo ThinkSystem SR650 (used as VMware ESXI host)

Logical components:

- IBM Spectrum Copy Data Management appliance (on ESXI host)
- Security Scan Engine host (on ESXI host)
- vCenter appliance (on ESXI host)
- SAP HANA database host (LPAR on Power9)

Connectivity:

- Ethernet hardware connectivity is 1 Gb Ethernet.
- Internal VMware network
- Internal Power9 network

Register components in IBM Spectrum Copy Data Management

All components of the environment which will be orchestrated by IBM Spectrum Copy Data Management must be registered first, including all permissions required to perform the appropriate tasks.

SAP HANA Database

For backing up the SAP HANA database, two different types of access are used.

1. An SAP HANA database user which needs the BACKUP OPERATOR or ADMINISTRATOR privilege to run the database backup and restore tasks.

2. An operating system user for managing the SAP HANA database from the operating system perspective (<SID>adm). Please read the IBM Spectrum Copy Data Management Documentation for further information.

Figure 5 shows the SAP HANA registration window with the required entry fields.

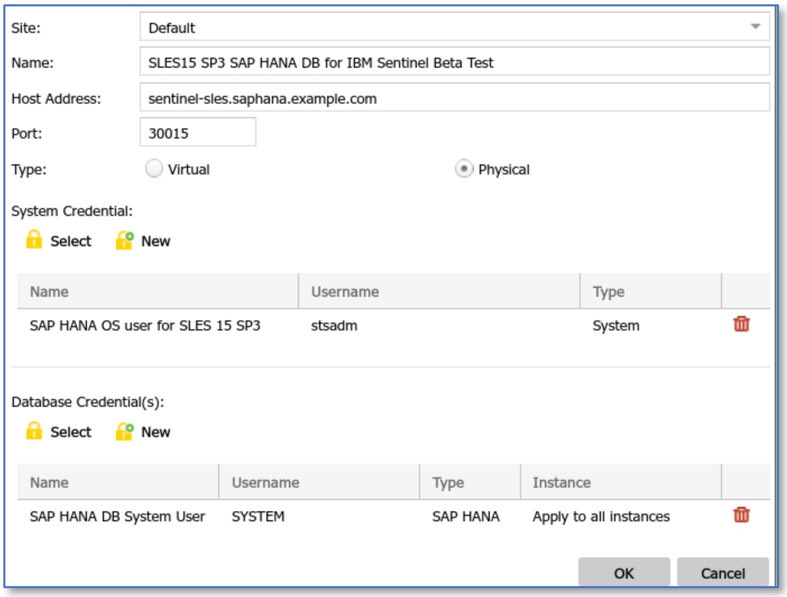

Site:	Default
Name:	SLES15 SP3 SAP HANA DB for IBM Sentinel Beta Test
Host Address:	sentinel-sles.saphana.example.com
Port:	30015
Type:	○ Virtual ● Physical

System Credential:

Name	Username	Type
SAP HANA OS user for SLES 15 SP3	stsadm	System

Database Credential(s):

Name	Username	Type	Instance
SAP HANA DB System User	SYSTEM	SAP HANA	Apply to all instances

Figure 5: Register SAP HANA database in IBM Spectrum Copy Data Management

IBM FlashSystem

For registering the IBM FlashSystem, it is recommended to create a dedicated user on the system. It needs to be assigned to the user group “Administrator”, since IBM Spectrum Copy Data Management requires administrative privileges on the storage system.

Note:
Do not add the user to the “SecurityAdmin” group! Members of this group are allowed to delete SafeGuardedCopy secured volumes, which would render the security approach unusable.

Scanning Engine

To perform malware scanning, IBM Spectrum Copy Data Management requires two users:

1. A dedicated Scanning Engine user to communicate with the scanning engine
2. An operating system user for mapping and mounting storage devices to the Scanning Engine host as part of the scanning process.

VMware vCenter

IBM Spectrum Copy Data Management needs access to the VMware vCenter with appropriate privileges which controls the ESXI host where the scanning engine is hosted. Since the scanning engine does not run on a bare metal machine, nor uses NPIV virtualization, all storage related operations need to be done using VMware APIs.

Prepare IBM FlashSystem Safe Guarded Copy for IBM Spectrum Copy Data Management

IBM Spectrum Sentinel requires the following preparation on the IBM FlashSystem or IBM SAN Volume Controller storage device:

To use the Safe Guarded Copy feature, the volumes which should be protected by IBM Spectrum Sentinel need to be added to a volume group. A volume group is a container for managing a set of related volumes as a single object. The volume group provides consistency across all volumes in the group and can be used for various consistency related functions such as policy-based replication, the recently introduced snapshot function and of course for managing Safeguarded copies. Safeguarded copy works like a regular flash copy, but the target volumes of the flash copy relations must reside in a specific child pool which is flagged as safeguarded. Only one child pool per parent can be flagged in such a way. Create the child pool as shown in **Figure 6**.

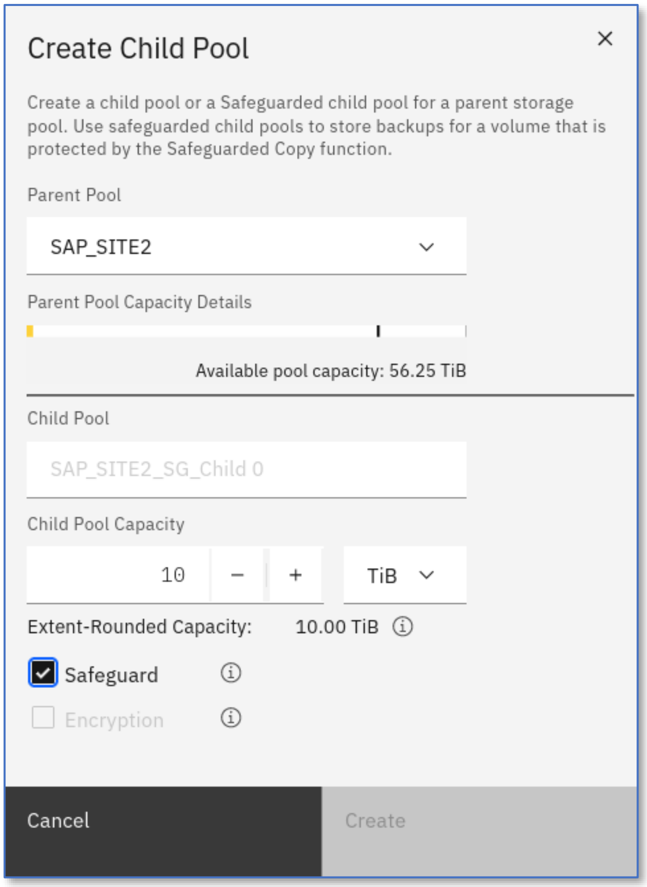

Create Child Pool

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

Parent Pool: SAP_SITE2

Parent Pool Capacity Details

Available pool capacity: 56.25 TiB

Child Pool: SAP_SITE2_SG_Child 0

Child Pool Capacity: 10 TiB

Extent-Rounded Capacity: 10.00 TiB

☑ Safeguard

☐ Encryption

Figure 6: Safeguarded Child Pool creation

Next, the Volume Group needs to be created. Volume groups can be created using existing volumes, or the volumes can be added later. In case a safeguarded child pool of the volume's parent pool already exists, the Volume Group identifies this Safeguarded Copy pool as the Safeguarded Backup Location, as shown in **Figure 7.**

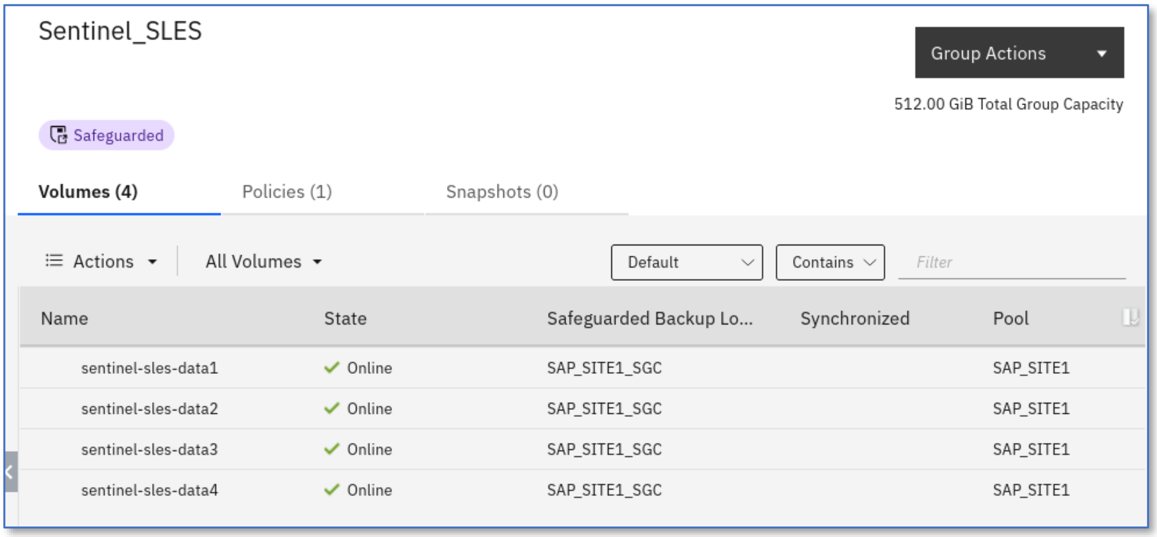

Sentinel_SLES

Safeguarded

512.00 GiB Total Group Capacity

Name	State	Safeguarded Backup Lo...	Synchronized	Pool
sentinel-sles-data1	✓ Online	SAP_SITE1_SGC		SAP_SITE1
sentinel-sles-data2	✓ Online	SAP_SITE1_SGC		SAP_SITE1
sentinel-sles-data3	✓ Online	SAP_SITE1_SGC		SAP_SITE1
sentinel-sles-data4	✓ Online	SAP_SITE1_SGC		SAP_SITE1

Figure 7: Volume Group with already configured Volumes

Finally, a Safeguarded Backup Policy must be assigned to the volume group. The backup policy defines the retention parameters of the Volume Group backups. The Safeguarded Backup allows to define the retention policy parameters by an external instance like IBM Copy Data Management.
Now this Volume Group will be visible in IBM Spectrum Copy Data Management for further usage.

Note:
Do not define a Snapshot policy for your Volume Group. Snapshot policies are under control of the IBM Spectrum Virtualize administrative user and cannot be managed by IBM Spectum Copy Data Management.

The storage system is now ready for use with IBM Spectrum Sentinel. The next step is configuring backup operations, which are defined by two different objects inside IBM Spectrum Copy Data Management:
The job definition and the SLA policy. The backup job definition defines what must be protected, while the SLA policy defines how to protect it.

SLA Policy: Specify rules and protection type

The SLA policy defines how to protect a database or file system application. It uses specific features of the underlying Storage system; therefore, separate policies are predefined for different Storage systems. Since many Storage systems offer multiple backup features, additional policies may exist for a specific storage system. For IBM Spectrum Virtualize, four different types of SLA policies are available:

- VM Replication (VMware only feature)
- FlashCopy
- Safeguarded Copy
- Global Mirror with Change Volumes

Next to the backup type, the SLA policy also defines the backup source, its Recovery Time Objective (RPO) and the backup retention time.
The backup source is the system where SAP HANA stores its savepoints and snapshots.
With SAP HANA, a FlashCopy policy and a Safeguarded Copy policy is supported. The Safeguarded Copy policy supports adding the scanning engine and a security scan job to the policy.
Creating a Safeguarded Copy SLA policy is easy when the involved objects as described in chapter “**IBM Spectrum Copy Data Management”** are correctly defined.

Spectrum Copy Data Management has an integrated wizard for creating SLA policies. It requires a unique name for the policy and the desired RPO target. Next, the FlashSystem Volume Group for SAP HANA needs to be selected. Finally, the retention rules and the security scanning server of IBM Sentinel need to be set, as shown in **Figure 8.**

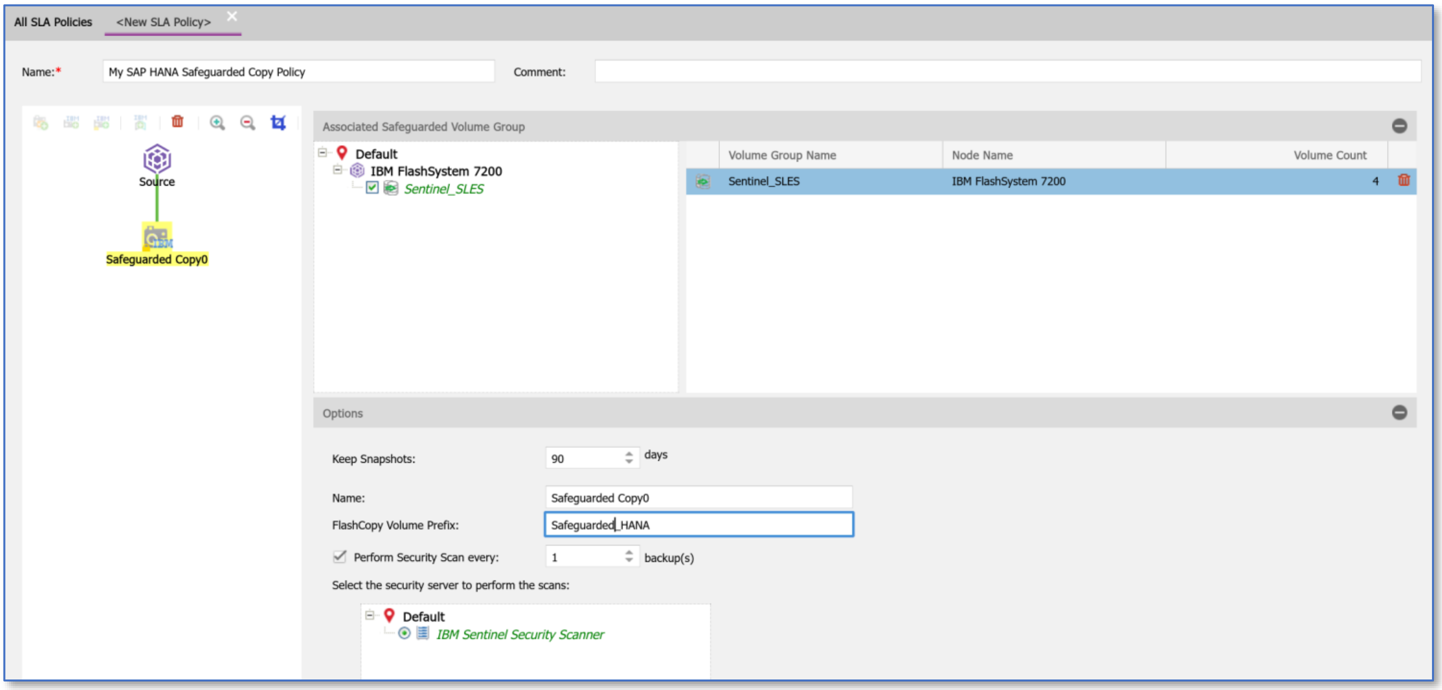

Name:* My SAP HANA Safeguarded Copy Policy

Comment:

Associated Safeguarded Volume Group

- Default
  - IBM FlashSystem 7200
    - ☑ Sentinel_SLES

Volume Group Name	Node Name	Volume Count
Sentinel_SLES	IBM FlashSystem 7200	4

Options

Keep Snapshots: 90 days

Name: Safeguarded Copy0

FlashCopy Volume Prefix: Safeguarded_HANA

✓ Perform Security Scan every: 1 backup(s)

Select the security server to perform the scans:

- Default
  - ● IBM Sentinel Security Scanner

Figure 8: Creating a Safeguarded Copy SLA policy with security scanner

Create backup jobs

When one or more SLA policies are in place, a backup job definition for the SAP HANA database must be created. The backup job definition connects the rules specified by the SLA policy with the protection target, and it adds a schedule for running the job. While the SLA policy defines how to run a backup, the backup job definition says what to backup (and when). **Figure 9** shows an example of a backup job definition.

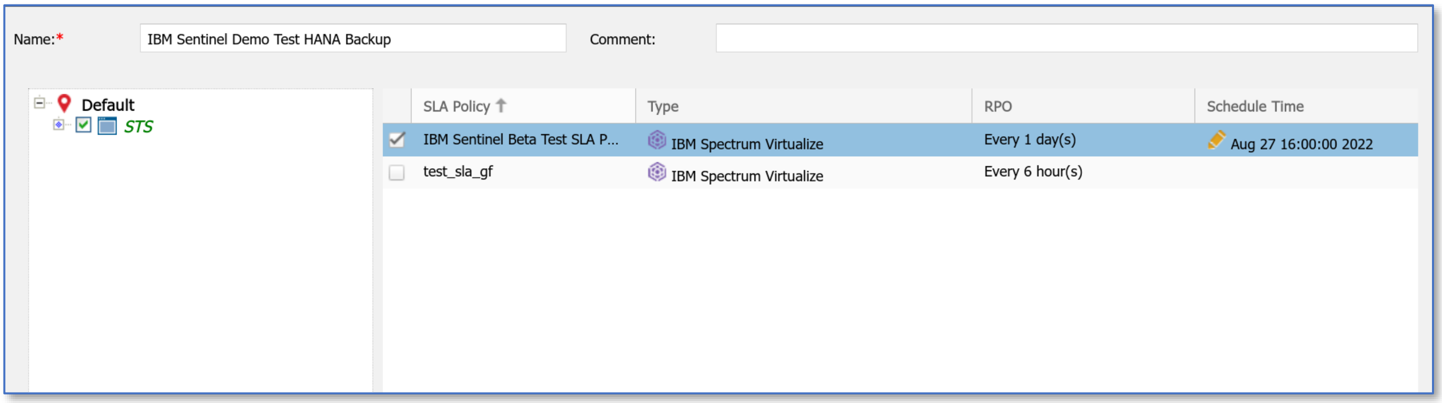

Name:* IBM Sentinel Demo Test HANA Backup

Comment:

- Default
  - ☑ STS

	SLA Policy	Type	RPO	Schedule Time
☑	IBM Sentinel Beta Test SLA P...	IBM Spectrum Virtualize	Every 1 day(s)	Aug 27 16:00:00 2022
□	test_sla_gf	IBM Spectrum Virtualize	Every 6 hour(s)	

Figure 9: Backup Job definition for SAP HANA

In the example above, the SAP HANA database “STS” is backed up using the “IBM Sentinel Beta Test SLA Policy” which defines the RPO. In addition, a schedule defines the start time and frequency of the backup. In this example, the RPO is 1 day, and the job starts at 16:00. This means that the backup job will be repeated every day at the same time.

A backup job may also contain pre- and post-scripts, which are called before and after running the backup job. SAP HANA consistent backups are already implemented in Spectrum Copy Data Management, but other applications might need some preparation before a backup, or additional clean up steps afterwards.

Note:
Spectrum Copy Data Management also allows to protect the SAP HANA transaction log backups. Because the log backup cannot be taken independently from the snapshot, using this feature will not shorten the RPO. However, backing up the logs is possible and an appropriate way to do this is described in the IBM practical guide "Protecting SAP HANA with IBM Spectrum Protect and IBM Spectrum Copy Data Management".

Backup / restore workflow with IBM Spectrum Sentinel

This chapter describes the overall picture of the backup and restore workflow for SAP HANA with IBM Spectrum Sentinel.
In a first step the registered components are scanned by IBM Spectrum Copy Data Management. The information collected by this scan is stored in the internal database of IBM Spectrum Copy Data Management. This allows IBM Spectrum Copy Data Management to perform backup and restore jobs fast, without the need to re-scan the systems before and after running a job.
To back up an SAP HANA database, IBM Spectrum Copy Data Management needs to know the complete data path from the SAP HANA data volumes down to the storage volumes where this data is stored. To configure the storage for SAP HANA properly, IBM has released the "IBM System Storage Architecture and Configuration Guide for SAP HANA Tailored Datacenter Integration. A schematic view of the data path is shown in **Figure 10**.

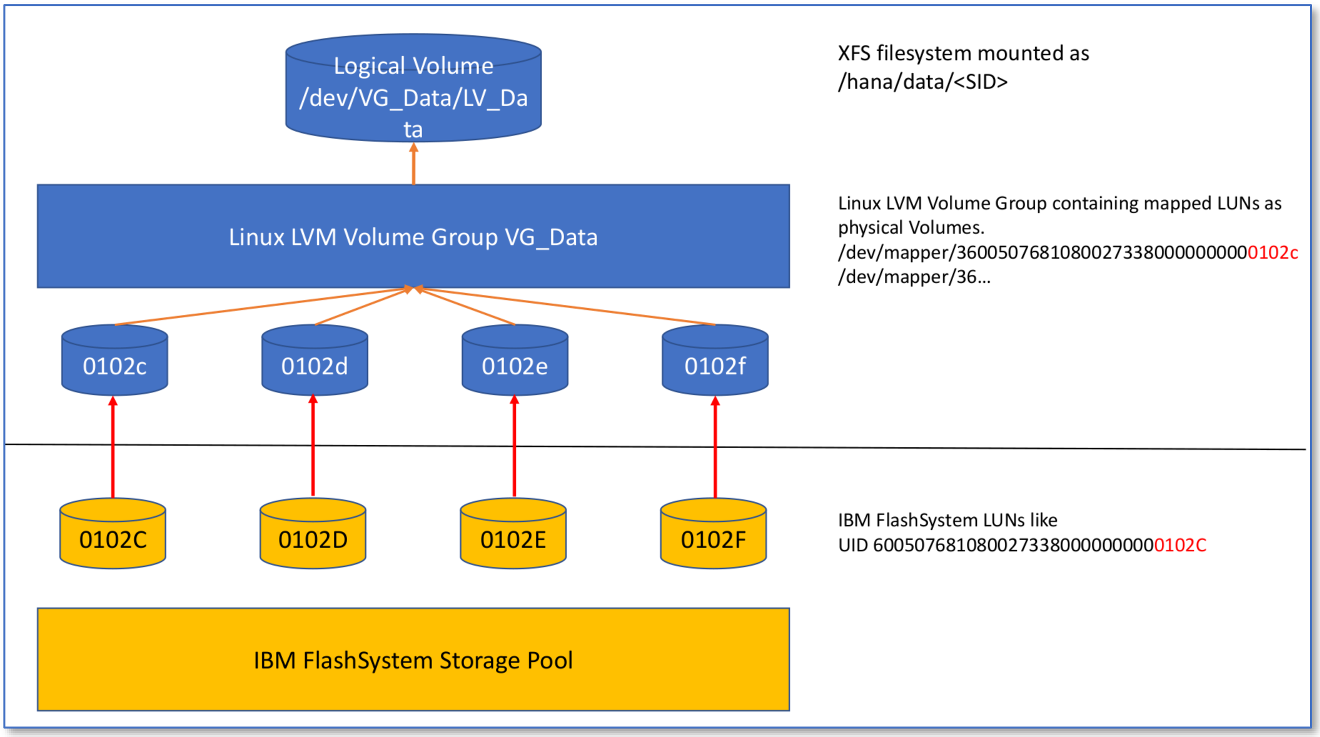
Logical Volume
/dev/VG_Data/LV_Da
ta
XFS filesystem mounted as
/hana/data/<SID>
Linux LVM Volume Group VG_Data
Linux LVM Volume Group containing mapped LUNs as physical Volumes.
/dev/mapper/36005076810800273380000000000102c
/dev/mapper/36…
0102c
0102d
0102e
0102f
0102C
0102D
0102E
0102F
IBM FlashSystem LUNs like
UID 6005076810800273380000000000102C
IBM FlashSystem Storage Pool

Figure 10: SAP HANA data volume: data path

In this example, IBM Spectrum Copy Data Management will identify the four Spectrum Virtualize Volumes 0102C – 0102F as the LUNs holding the SAP data. These volumes will be backed up using the IBM Spectrum Virtualize FlashCopy feature.

When using FlashCopy, the storage system must ensure that all involved volumes are in a consistent state when the FlashCopy operation starts. IO-operations are not allowed during this time to ensure consistency of the copied data. Spectrum Virtualize uses consistency groups to guarantee time consistent flash copies. With Safeguarded Copy, the write consistency is managed by the IBM Spectrum Virtualize Volume Group. However, currently IBM Spectrum Copy Data Management uses legacy Consistency Groups when it flashes the volumes, even with Safeguarded Copy. The FlashCopy operation takes just a few microseconds, so there is no measurable impact to the host IO performance.

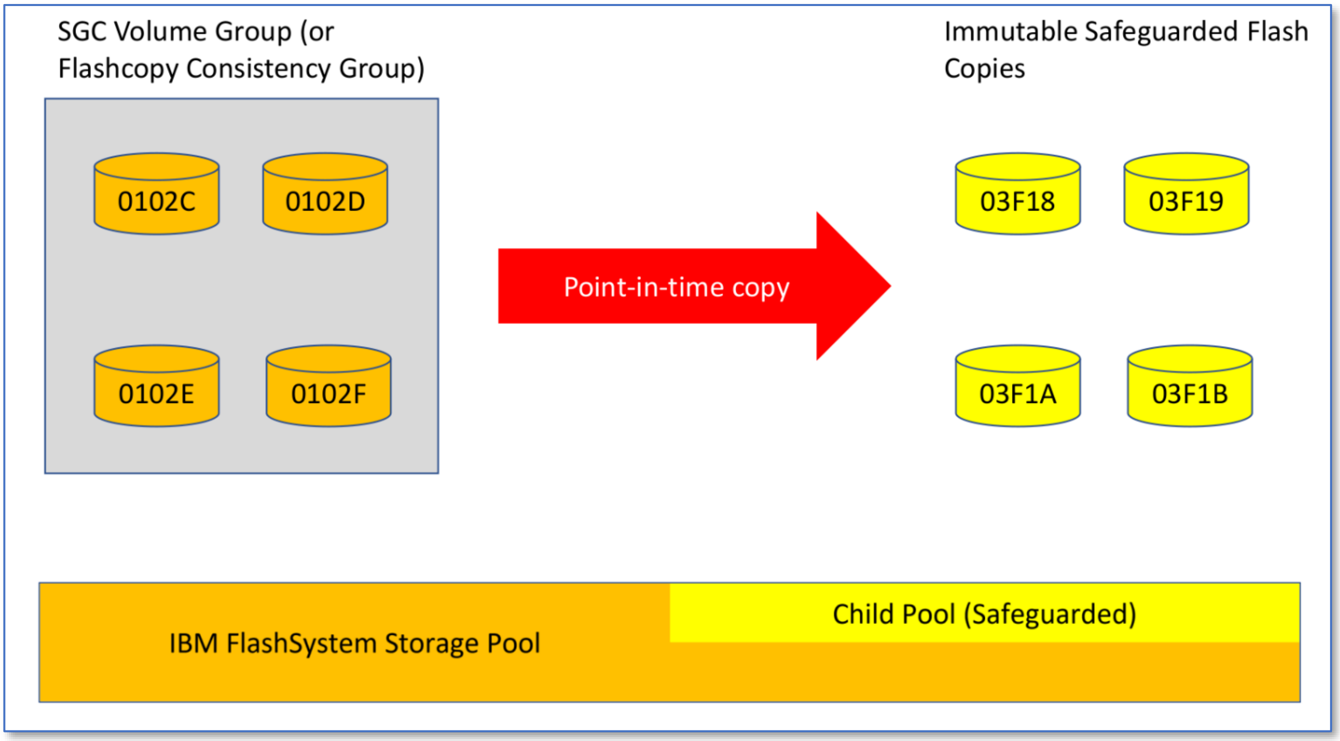

Figure 11: FlashCopy using Safeguarded volumes

The SAP HANA database needs to be prepared before flashing the data volumes. IBM Spectrum Copy Data Management does this by running a set of SQL statements as a database Backup operator. An SAP HANA snapshot, which is an online backup of the SAP HANA database, is written to the SAP HANA data volume. As soon as the snapshot has been created, I/O operations to the XFS filesystem are suspended, and the volumes are being flashed. When the flash copy operation is finished, I/O operations are resumed. As a last step, additional SQL statements are issued to inform the database that the snapshot has now been saved, a.k.a. “committed”. SAP HANA stores the backup information into its backup catalog and removes the snapshot file from the data area.

Running Backups

After a backup job has been defined in IBM Spectrum Copy Data Management, it can be started either manually or by using the built-in scheduler. In this example we start the backup job manually.

1. Click the Jobs tab.
2. Select the job to run by clicking in the row containing the job name, as shown in **Figure 12**.
3. Click Start, or right-click the job name and select *Start.* A confirmation dialog box opens.
4. Click Yes. The job session runs.

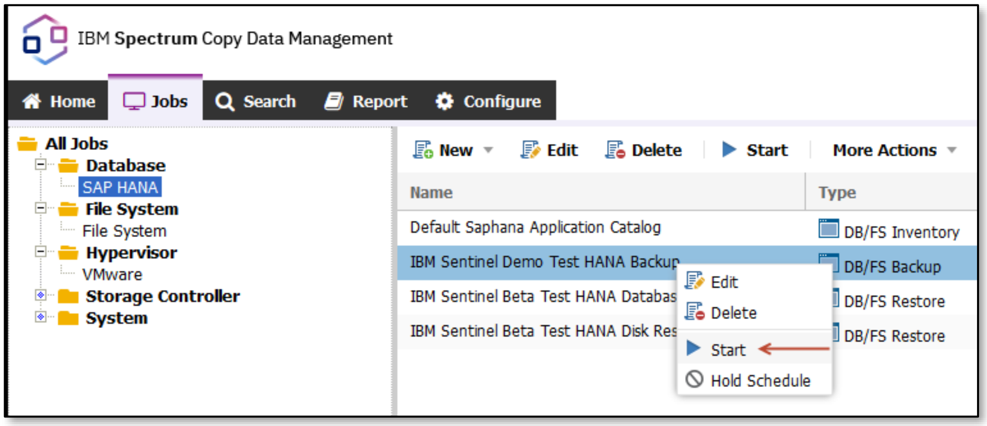

Name	Type
Default Saphana Application Catalog	DB/FS Inventory
IBM Sentinel Demo Test HANA Backup	DB/FS Backup
IBM Sentinel Beta Test HANA Databas	DB/FS Restore
IBM Sentinel Beta Test HANA Disk Res	DB/FS Restore

Figure 12: Start the HANA backup job

When the backup is completed successfully, the backup job reflects the status as *COMPLETED*. If the backup has the Status FAILED, it could either be due to a setup error or the scanning engine has detected an infected backup. To check for the cause of a failed backup, view the backup activity log.

In the lower left of the Spectrum Copy Data Management GUI is the *Activity* pane. Click the job name to view the activity log for the specific job, including the job session's start date and time, duration, description, status, and associated messages, as shown in **Figure 13**. In this backup log example, the scanning engine has detected a possible malware threat in the backup data. For this reason, the backup job has the status *FAILED*.

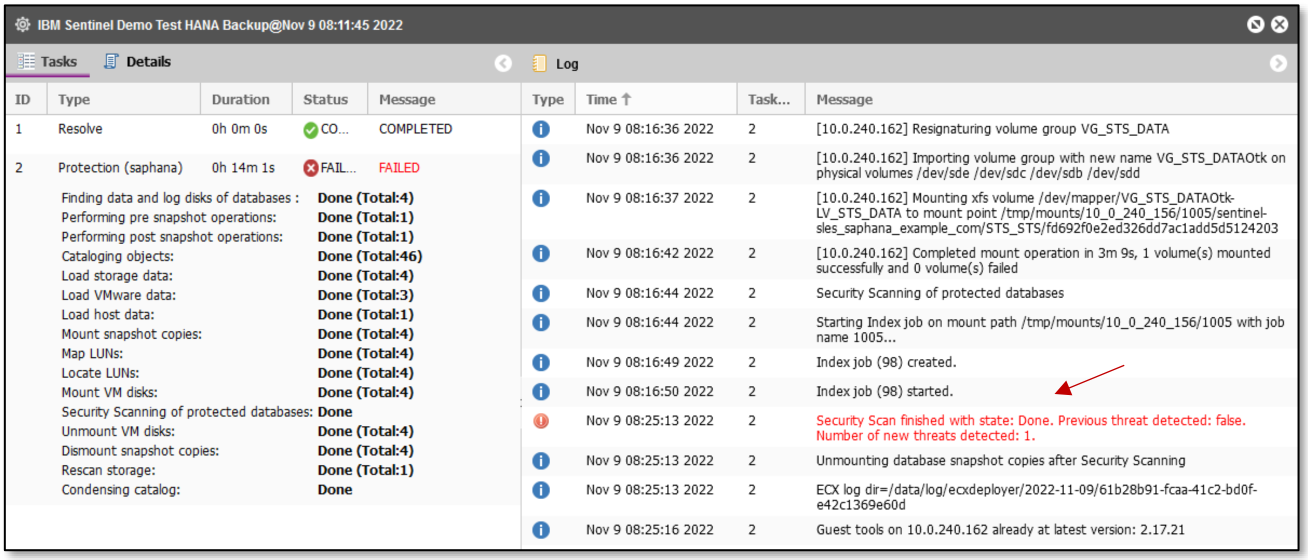

ID	Type	Duration	Status	Message
1	Resolve	0h 0m 0s	✓ CO...	COMPLETED
2	Protection (saphana)	0h 14m 1s	✗ FAIL...	FAILED

```
Finding data and log disks of databases : Done (Total:4)
Performing pre snapshot operations:       Done (Total:1)
Performing post snapshot operations:      Done (Total:1)
Cataloging objects:                       Done (Total:46)
Load storage data:                        Done (Total:4)
Load VMware data:                         Done (Total:3)
Load host data:                           Done (Total:1)
Mount snapshot copies:                    Done (Total:4)
Map LUNs:                                 Done (Total:4)
Locate LUNs:                              Done (Total:4)
Mount VM disks:                           Done (Total:4)
Security Scanning of protected databases: Done
Unmount VM disks:                         Done (Total:4)
Dismount snapshot copies:                 Done (Total:4)
Rescan storage:                           Done (Total:1)
Condensing catalog:                       Done
```

Type	Time ↑	Task...	Message
ⓘ	Nov 9 08:16:36 2022	2	[10.0.240.162] Resignaturing volume group VG_STS_DATA
ⓘ	Nov 9 08:16:36 2022	2	[10.0.240.162] Importing volume group with new name VG_STS_DATAOtk on physical volumes /dev/sde /dev/sdc /dev/sdb /dev/sdd
ⓘ	Nov 9 08:16:37 2022	2	[10.0.240.162] Mounting xfs volume /dev/mapper/VG_STS_DATAOtk-LV_STS_DATA to mount point /tmp/mounts/10_0_240_156/1005/sentinel-sles_saphana_example_com/STS_STS/fd692f0e2ed326dd7ac1add5d5124203
ⓘ	Nov 9 08:16:42 2022	2	[10.0.240.162] Completed mount operation in 3m 9s, 1 volume(s) mounted successfully and 0 volume(s) failed
ⓘ	Nov 9 08:16:44 2022	2	Security Scanning of protected databases
ⓘ	Nov 9 08:16:44 2022	2	Starting Index job on mount path /tmp/mounts/10_0_240_156/1005 with job name 1005...
ⓘ	Nov 9 08:16:49 2022	2	Index job (98) created.
ⓘ	Nov 9 08:16:50 2022	2	Index job (98) started.
!	Nov 9 08:25:13 2022	2	Security Scan finished with state: Done. Previous threat detected: false. Number of new threats detected: 1.
ⓘ	Nov 9 08:25:13 2022	2	Unmounting database snapshot copies after Security Scanning
ⓘ	Nov 9 08:25:13 2022	2	ECX log dir=/data/log/ecxdeployer/2022-11-09/61b28b91-fcaa-41c2-bd0f-e42c1369e60d
ⓘ	Nov 9 08:25:16 2022	2	Guest tools on 10.0.240.162 already at latest version: 2.17.21

Figure 13: Backup job log with error message for infected backup

Schedule Backups and Scans

Backup jobs can be scheduled by the built-in IBM Spectrum Copy Data Management scheduler, or you can use job automation tools like e.g., UC4 to start IBM Spectrum Copy Data Management backup jobs using REST-API calls.

E-mail notification for backup jobs

For email notifications, at least one SMTP server must be configured in IBM Spectrum Copy Data Management, as shown in **Figure 14**. The SMTP server must be added to IBM Spectrum Copy Data Management, before defining a backup job.

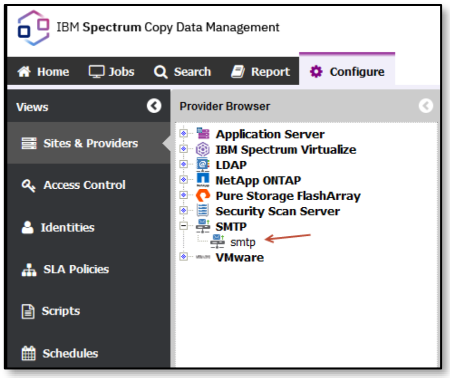

Figure 14: Configure SMTP in IBM Spectrum Copy Data Management

IBM Spectrum Copy Data Management backup jobs can be configured to create status emails for every backup job. An example of an SAP HANA backup job status email is shown in **Figure 15**. This email also includes the job log.

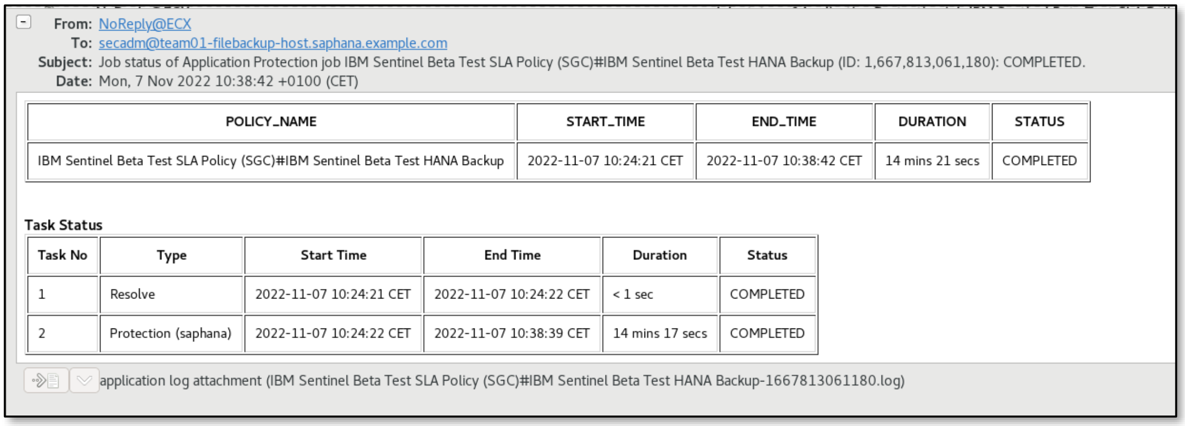

From: NoReply@ECX
To: secadm@team01-filebackup-host.saphana.example.com
Subject: Job status of Application Protection job IBM Sentinel Beta Test SLA Policy (SGC)#IBM Sentinel Beta Test HANA Backup (ID: 1,667,813,061,180): COMPLETED.
Date: Mon, 7 Nov 2022 10:38:42 +0100 (CET)

POLICY_NAME	START_TIME	END_TIME	DURATION	STATUS
IBM Sentinel Beta Test SLA Policy (SGC)#IBM Sentinel Beta Test HANA Backup	2022-11-07 10:24:21 CET	2022-11-07 10:38:42 CET	14 mins 21 secs	COMPLETED

Task Status

Task No	Type	Start Time	End Time	Duration	Status
1	Resolve	2022-11-07 10:24:21 CET	2022-11-07 10:24:22 CET	< 1 sec	COMPLETED
2	Protection (saphana)	2022-11-07 10:24:22 CET	2022-11-07 10:38:39 CET	14 mins 17 secs	COMPLETED

application log attachment (IBM Sentinel Beta Test SLA Policy (SGC)#IBM Sentinel Beta Test HANA Backup-1667813061180.log)

Figure 15: backup job completion e-mail

Post attack actions

In case the scan result of a backup copy is 'FAILED', there is a high probability that the source environment is under attack. The details can be investigated using the log of the scanning engine. Note that IBM Spectrum Sentinel is not a tool for forensic analysis. Such incidents need to be investigated according to the internal cyber response plan of the attacked entity. The IBM X-Force team is also available to help and can be contacted here: IBM Security X-Force https://www.ibm.com/x-force.

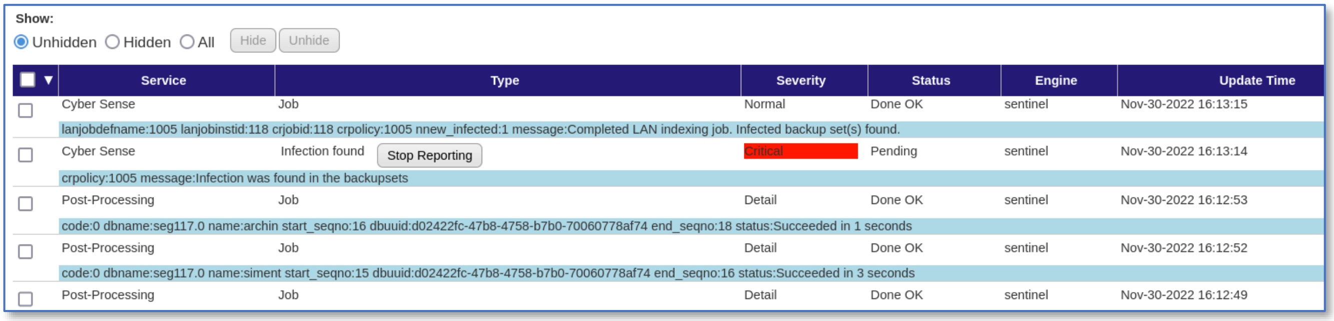

Show:
● Unhidden ○ Hidden ○ All

□ ▼	Service	Type	Severity	Status	Engine	Update Time
□	Cyber Sense	Job	Normal	Done OK	sentinel	Nov-30-2022 16:13:15
	lanjobdefname:1005 lanjobinstid:118 crjobid:118 crpolicy:1005 nnew_infected:1 message:Completed LAN indexing job. Infected backup set(s) found.					
□	Cyber Sense	Infection found Stop Reporting	Critical	Pending	sentinel	Nov-30-2022 16:13:14
	crpolicy:1005 message:Infection was found in the backupsets					
□	Post-Processing	Job	Detail	Done OK	sentinel	Nov-30-2022 16:12:53
	code:0 dbname:seg117.0 name:archin start_seqno:16 dbuuid:d02422fc-47b8-4758-b7b0-70060778af74 end_seqno:18 status:Succeeded in 1 seconds					
□	Post-Processing	Job	Detail	Done OK	sentinel	Nov-30-2022 16:12:52
	code:0 dbname:seg117.0 name:siment start_seqno:15 dbuuid:d02422fc-47b8-4758-b7b0-70060778af74 end_seqno:16 status:Succeeded in 3 seconds					
□	Post-Processing	Job	Detail	Done OK	sentinel	Nov-30-2022 16:12:49

Figure 16: Scanning Engine alert after ransomware attack

In some cases, you may need to recover the server operating system and application installation before IBM Spectrum Sentinel can recover the last known good backup. It is a good practice to have standby servers that can be activated in an emergency. If you have any questions about the detection status, contact IBM Spectrum ® Sentinel Support. IBM may assist with questions about the solution, troubleshoot recovery needs, and may refer you to IBM Services for more extensive threat mitigation assistance.

Restore DB from last good known copy

A main advantage of IBM Spectrum Sentinel compared to other backup solutions is the shortened RTO after a cyber-attack has occurred. IBM Spectrum Copy Data Management has a list of all backups it manages, and if they have successfully passed the security scan they are flagged as clean. Therefore, those backups can be picked up to recover infected applications rapidly.
Like a data backup, in IBM Spectrum Copy Data Management data restores are also handled by a job which needs to be defined before the restore operation starts.

For SAP HANA, IBM Spectrum Copy Data Management offers two different restore job templates.

- Instant Disk Restore
  This is a crash consistent recovery of the underlying data area of the SAP HANA database, restoring the automatically created savepoint and additionally the snapshot which has been created by the backup job. Further recovery actions need to be done by the SAP HANA database administrator. This approach gives the database administrator full control of the recovery process, e.g., running roll-forward operations using the transaction log.

- Instant Database Restore
  This is an application consistent recovery which runs fully automated. No further action by the SAP HANA administrator is required. The database recovers automatically to the time the selected snapshot has been taken. Roll-forward operations will not be done, and the existing transaction log will be cleared. This means the Instant Database Restore is restricted to the RPO predefined by the backup job's schedule.

The restore job definition in IBM Spectrum Copy Data Management is guided by a wizard. After selecting the appropriate template, the database(s) which should be recovered need to be selected. Then, the copy which should be used for recovery is specified. A backup date range can be defined here, so that a list of eligible backups will be created, as shown in **Figure 17**. Alternatively, either the latest backup or the latest successfully scanned backup can be used.

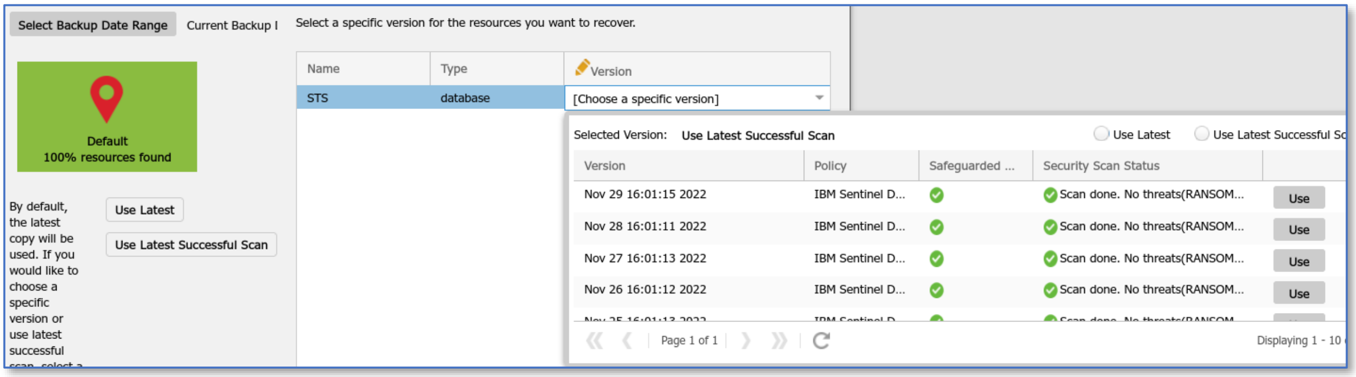

Select a specific version for the resources you want to recover.

Name	Type	Version
STS	database	[Choose a specific version]

By default, the latest copy will be used. If you would like to choose a specific version or use latest successful

Selected Version: Use Latest Successful Scan

○ Use Latest ○ Use Latest Successful Sc

Version	Policy	Safeguarded ...	Security Scan Status	
Nov 29 16:01:15 2022	IBM Sentinel D...	✓	✓ Scan done. No threats(RANSOM...	Use
Nov 28 16:01:11 2022	IBM Sentinel D...	✓	✓ Scan done. No threats(RANSOM...	Use
Nov 27 16:01:13 2022	IBM Sentinel D...	✓	✓ Scan done. No threats(RANSOM...	Use
Nov 26 16:01:12 2022	IBM Sentinel D...	✓	✓ Scan done. No threats(RANSOM...	Use

Figure 17: Choosing a specific backup version for Restore Job

In the next step the destination for the restore operations needs to be defined. SAP HANA copy, clone or refresh operations are currently not supported by IBM Spectrum Copy Data Management. The restore destination must be identical with the backup source. For specific recovery scenarios, advanced options can be adjusted using the “advanced” button.

- Rename mount points
  IBM Spectrum Copy Data Management can rename the target mount points using rewriting rules. However, the existing mount point can also be recovered, which is often used as the default in a disaster recovery scenario,
- Storage options
  - Make permanent
    Making the restored volumes permanent means to restore them by using a FlashCopy with a background copy of all data. Once the background copy has finished, the target volume will be independent of the source.
  - Revert
    This option utilizes the so-called reverse restore FlashCopy feature. It switches the source and target volumes and writes all changes back to the source. The source is useable immediately, even if the restore is still ongoing in the background. A reverse restore can be performed while the volume is mapped to the host, this can further reduce the recovery time. Since the data on the block devices is being changed by this operation outside of the host control, IBM Spectrum Copy Data Management will perform a full unmap and remap cycle of the volumes for better security.

Summary

The IBM Spectrum Sentinel offering that is described in this IBM blueprint publication shows the integration of IBM FlashSystem, IBM Spectrum Copy Data Management, and SAP HANA to perform threat detection of infected database backups. IBM Spectrum Sentinel is a cyber resiliency solution designed to help businesses enhance ransomware detection and incident recovery. IBM Spectrum Sentinel combines IBM Spectrum Copy Data Management with an anomaly scanning software to coordinate file and database corruption scanning with snapshot management and recovery orchestration.

Because IBM Spectrum Sentinel can intelligently isolate infected SAP HANA backups, your organization can identify the most recent verified and validated backup copies, greatly accelerating your time to recovery.

Authors

This White Paper was developed by a team of specialists from the EMEA Storage Competence Center.

Axel Westphal is an IBM Certified IT Specialist at the IBM EMEA Storage Competence Center (ESCC) in Frankfurt, Germany. He joined IBM in 1996, working for IBM Global Services as a Systems Engineer. His areas of expertise include setting up and demonstrating IBM System Storage products and solutions in various environments. He has written several storage White Papers and co-authored several IBM publications.

Thomas Gerisch is an IT Specialist at the IBM EMEA Storage Competency Center (ESCC) in Frankfurt, Germany. He joined IBM in 1999, working as an instructor and specialist for open systems. He is the technical focal point for IBM customers running SAP HANA on IBM Storage. He has authored several IBM White Papers in the SAP HANA and IBM storage environment.

The authors want to thank **Gerd Franke**, IBM Technology Services, Senior Storage Architect, for his valuable contributions and support of this project.

Resources

For further information, please refer to:

IBM FlashSystem 7300 system overview:
https://www.ibm.com/docs/en/flashsystem-7x00/8.5.x?topic=to-flashsystem-7300-system-overview
´

IBM FlashSystem Cyber Vault:
https://www.ibm.com/downloads/cas/ODKXBLR9

IBM Spectrum Sentinel Overview:
https://www.ibm.com/products/spectrum-sentinel

IBM Spectrum Copy Data Management user’s guide:
https://www.ibm.com/docs/en/SS57AN_2.2.18/pdf/b_cdm_guide.pdf

Certified and Supported SAP HANA® Hardware Directory:
https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/#/solutions?filters=v:deCertified;storage;ve:6

IBM® System Storage™ Architecture and Configuration Guide for SAP® HANA™ Tailored Datacenter Integration
https://www.ibm.com/support/pages/node/6355415

IBM Definitive guide to ransomware 2022:
https://www.ibm.com/downloads/cas/EV6NAQR4

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

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

- FlashCopy®
- IBM®
- IBM AIX®
- IBM FlashCore®
- IBM FlashSystem®
- IBM Storage Insights
- IBM Spectrum™
- IBM Spectrum Storage™
- IBM Spectrum Sentinel™
- IBM Spectrum Virtualize™
- POWER®
- Power Systems™
- POWER9™
- PowerLinux™
- Redbooks (logo) ®
- System Storage®

The following terms are trademarks of other companies:

SAP HANA is a trademark or registered trademark of SAP SE or its affiliates in Germany and in several other countries.

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

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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

InterSystems, Caché, and ECP are trademarks or registered trademarks of InterSystems Corporation.

Epic is a registered trademark of Epic Systems Corporation.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Oracle®, Java, and MySQL are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Microsoft SQL Server is a trademark of Microsoft Corporation in the United States, other countries, or both.

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

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

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at http://www.ibm.com/privacy and IBM's Online Privacy Statement at http://www.ibm.com/privacy/details in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at http://www.ibm.com/software/info/product-privacy.

December 2022

Please recycle