

Cyber resilient consistent
safeguarded copy for IBM
Storage Protect on IBM
Storage Scale



Table of Contents

- About this white paper 3
- Introduction 3
- IBM Storage Protect 3
- IBM Storage Scale 4
 - Safeguarded copy for IBM Storage Scale..... 4
- Cyber-resilient consistent safeguarded copy for IBM Storage Protect on IBM Storage Scale..... 5
 - Requirements and deployment prerequisites 5
 - REST API and CLI usage..... 5
- IBM Storage Protect and IBM Storage Scale integration topologies 6
 - Topology 1: Storage Protect instances within the storage cluster 6
 - Topology 2: Storage Protect instances in a remote cluster 7
- Consistent SGC for IBM Storage Protect on IBM Storage Scale..... 8
- HUK-COBURG cyber resilient solution requirements 9
 - Existing environment 9
 - Regulatory and recovery requirements..... 10
- IBM Storage solution for HUK-COBURG 10
 - Preparation 11
 - Deployment 13
 - Operations 15
 - Solution outcomes – feedback from the customer 16
- Conclusion 17
- About the authors 17

About this white paper

This IBM® white paper outlines how HUK-COBURG and IBM have partnered to provide cyber resilient backups using IBM Storage Protect and IBM Storage Scale. HUK-COBURG is a German insurance group that serves over 13 million customers and is the market leader in motor insurance. Founded in 1933, the company leverages its digital spearhead HUK24 to maintain its market position.

This paper is intended for storage architects, backup administrators, system engineers, and IT professionals responsible for designing, operating, or securing IBM Storage Protect and IBM Storage Scale environments.

Introduction

In today's increasingly data-dependent world, data has become an important asset as well as a primary target for ransomware attacks and data breaches. With the growth in complexity and frequency of ransomware and cyber threats, organizations must prioritize cyber resilience. Cyber resilience focuses on ensuring the delivery of continuous business operations by being able to anticipate, withstand, recover from, and adapt to cyber threats and disruptions.

Backup data, which is usually the last line of defense, is now being targeted by cyberattacks that aim to encrypt, delete, or corrupt these backups to prevent recovery making resilience essential for backup systems. This requires strategies such as backups, physical or logical isolation, and the ability to restore quickly and accurately. Aligning with frameworks like the National Institute of Standards & Technology (NIST) Cybersecurity framework ensures a comprehensive approach that integrates detection, response, and recovery.

This paper presents an IBM approach that combines consistent, immutable safeguarded copies (SGC) on IBM Storage Scale with IBM Storage Protect automation to deliver fast, predictable recovery points and instance-level restoration.

With cybercrime damages set to exceed \$10.5 trillion annually as of 2025, prioritizing resilience in backup systems is no longer a good practice; rather, it is a necessity to minimize downtime, ensure compliance, and safeguard operational continuity.

This paper outlines how HUK-COBURG and IBM have partnered to provide cyber resiliency for IBM Storage Protect using IBM Storage Scale Consistent Safeguarded Copies solution. This paper also introduces the components, explains integration topologies, details the SGC workflow, and presents a production-tested deployment at HUK-COBURG with measured outcomes.

IBM Storage Protect

IBM Storage Protect, formerly known as IBM Spectrum® Protect, is a comprehensive data protection solution that offers backup and recovery for physical file servers, virtual environments, and a wide range of applications. It enables data migration or copying to tape, cloud services, or on-premises storage, helping organizations reduce backup infrastructure costs with built-in data-efficiency capabilities. The product supports massive data growth, with a single server capable of managing up to 4 petabytes of client data and ingesting up to 100 terabytes of new and changed data per day. IBM Storage Protect also provides incremental forever backups, compression, and deduplication for exceptional storage efficiency.

In terms of security, it offers two-factor authorization for administrator commands, data encryption, proactive security notifications, and native support for tape and object storage. The solution also allows for cost-effective data retention by leveraging cloud and on-premises storage options, including IBM Cloud Object Storage and Amazon Web Services (AWS) S3 Intelligent-Tiering for inactive data.

IBM Storage Scale

IBM Storage Scale (formerly known as IBM Spectrum® Scale) is a high-performance, scalable storage solution designed for managing large amounts of data across distributed environments. It provides a global shared file system that enables fast data access across multiple servers and locations. An IBM Storage Scale cluster can be set up in many ways to address various architectural and performance requirements. One of the most common installations is where all the nodes are directly connected to a common pool of storage drives through a storage area network (SAN). Here, each node has direct and equal access to the same disks, which enables high-speed data operations. Nevertheless, one must ensure that the nodes that access the same disks have the same operating systems for consistency and to prevent conflicts.

Another configuration method is a combination of nodes—some with direct disk access and some without direct access. The nodes that lack direct disk access in this configuration fetch information from other nodes that serve as intermediaries. These intermediate nodes are called Network Shared Disk (NSD) servers, whereas the nodes that depend on them for disk access are referred to as NSD clients. This kind of structure is most suitable for large deployments or for reducing infrastructure costs while maintaining high performance.

IBM Storage Scale also accommodates high-end configurations where data is shared among numerous clusters. These clusters may be installed in a single data center, or they may be geographically distributed across various regions. When a file system in one cluster is mounted on another, users and applications present in both systems can access the data as if they were within the same system. This multiple-cluster arrangement not only optimizes data access between varied locations but also facilitates easier administrative processes by allowing each cluster to be controlled separately or as a component of a combined, integrated infrastructure.

Safeguarded copy for IBM Storage Scale

IBM Storage Scale introduces the Safeguarded Copy feature starting from version 5.1.5, aimed at strengthening data protection within file systems. Built on the foundation of snapshots, SGC offers a resilient mechanism to safeguard data against cyberattacks, human errors, and system failures. These SGC are immutable, meaning they cannot be altered or deleted until their defined retention period expires.

SGC serves as tamper-proof representations of either an entire file system or selected file sets. Administrators can configure SGC to be automatically captured on demand or at regular intervals—such as hourly—and retained for a specified duration. This ensures the availability of a clean recovery point, in case the primary data becomes compromised. Since these SGC are read-only and protected from unauthorized changes or deletion—even by compromised admin accounts—they add a robust layer of security.

The creation and deletion of SGC can be performed by using the Storage Scale command-line interface (CLI) or Representational State Transfer (REST) application programming interface (API). Restoration of SGC is straightforward using commands like `mmrestorefs`, which can revert the file system or fileset to a previous SGC state. A key security aspect of SGC is its ability to prevent root-level breaches from impacting the backup environment. To reinforce this protection, IBM strongly advises that SGC operations be performed by a non-root administrator using the IBM Storage Scale sudo wrapper (`mmrestorefs`, if applicable). This approach minimizes the risk of attackers exploiting elevated privileges to interfere with backup data.

The following section describes how these components integrate to deliver a consistent, cyber-resilient recovery solution.

Cyber-resilient consistent safeguarded copy for IBM Storage Protect on IBM Storage Scale

This section describes the key elements of the validated solution, including its prerequisites, supported interfaces, and overall operational design.

Requirements and deployment prerequisites

A validated solution for managing consistent SGC and SGC with Storage Protect on IBM Storage Scale (version 5.2.1 or later) has been developed and tested on Red Hat® Enterprise Linux (RHEL) 8 and 9 using Storage Protect version 8.1.27 or later. This solution supports deployment on both local Storage Scale clusters and remote clusters. In remote environments, SGC operations are handled through the REST API.

When using the Storage Scale CLI to manage SGC, the Storage Protect instance user must have elevated privileges to execute SGC-related Storage Scale commands. This can be achieved through proper sudo configuration. Additionally, tools such as `curl` and `jq` are required to run the *isnap-scripts*.

Specific configurations for file systems and filesets are necessary for implementation. While SGC are initially space-efficient, their storage usage increases over time depending on the number of SGC and their retention period. Therefore, effective capacity planning is crucial.

For details on how SGC consumes storage space, refer to the section Storage space consumption with SGC: [Consistent SGC for IBM Storage Protect on IBM Storage Scale](#).

REST API and CLI usage

For managing consistent SGC, the *isnap-scripts* file can be configured to use the Storage Scale CLI or REST API. The Storage Scale REST API is used when the parameters `apiServerIP`, `apiServerPort`, and `apiCredentials` are specified in the configuration file. If these parameters are not specified, then the Storage Protect CLI is used for managing SGC.

When the REST API is used, automated restoration of SGC is not supported. Instead, administrators are provided with step-by-step instructions on the console to manually execute the necessary commands. When the CLI is used for managing SGC, administrators can decide whether SGC is restored automatically by setting the configuration parameter `AutoRestore` accordingly. For details on the restoration process, refer to [Restore safeguarded copy](#).

The following section introduces the deployment topologies that show how the solution is implemented across different environments.

IBM Storage Protect and IBM Storage Scale integration topologies

IBM Storage Protect can be deployed in conjunction with IBM Storage Scale to provide a resilient, scalable backup infrastructure. When integrated, multiple IBM Storage Protect server instances can share the same set of file systems offered by IBM Storage Scale. Each instance operates within its own dedicated fileset in the file system, ensuring logical separation of data and enabling consistent, SGC-specific protection for each instance. A fileset is a directory within the file system that allows SGC to be taken only for the fileset directory. This fileset structure plays a crucial role in supporting fast and precise recovery for individual instances in the event of failure or data corruption.

There are two key deployment topologies for this integration.

Topology 1: Storage Protect instances within the storage cluster

In the first configuration (see Figure 1), the Storage Protect instances run directly on the nodes of the Storage Scale storage cluster. These nodes are responsible for providing both flash and disk storage. Each instance of Storage Protect maintains its own database (for metadata), logs, backup storage pools, and configuration—all stored in separate file systems within the Storage Scale file systems. By assigning a unique fileset per instance, it becomes possible to take consistent, SGC-specific safeguarded copies independently for each server instance.

Automation scripts—referred to as *isnap-scripts*—can be deployed on the same cluster nodes to manage SGC creation, restoration and deletion of expired SGC. When using the Storage Scale CLI, SGC can be restored automatically; however, manual restoration is recommended. Alternatively, if the Storage Scale REST API is used, restoration becomes a manual process, offering flexibility depending on the system administrator's preference and the environment's requirements.

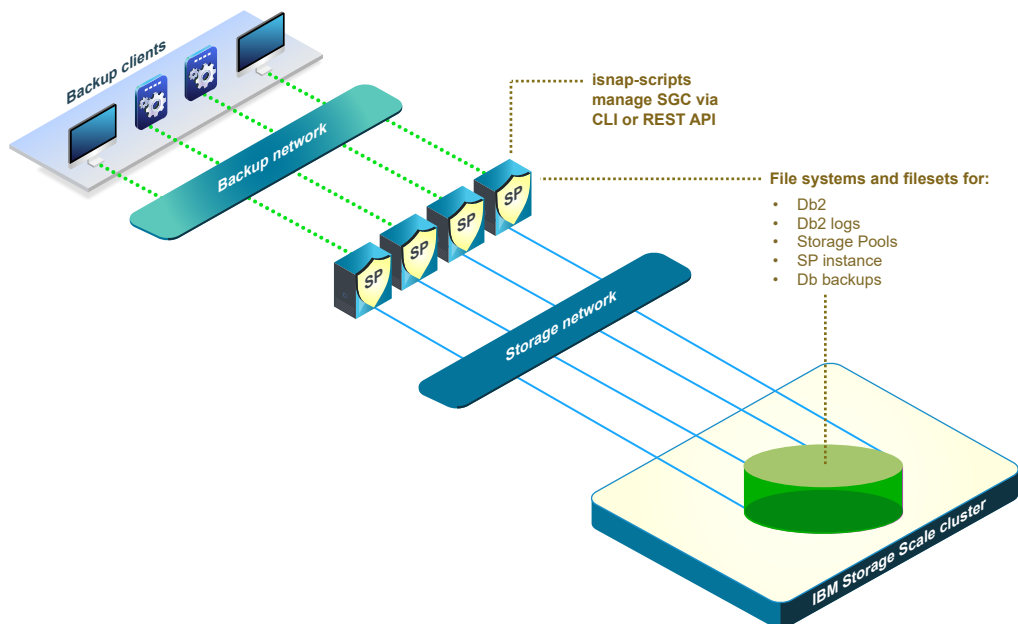


Figure 1: Storage Scale topology in storage cluster

The advantage of the solution with Storage Protect on Storage Scale is that all Storage Protect servers use the same storage provided as highly available Storage Scale file systems. This eliminates storage islands, fosters better storage utilization, and improves operational efficiency. In addition, Storage Scale allows scalability of performance and capacity in multiple dimensions.

Topology 2: Storage Protect instances in a remote cluster

In the second configuration (see Figure 2), Storage Protect instances are installed on a separate, storage-less Storage Scale nodes forming a Storage Protect cluster. This cluster accesses file systems remotely, which are hosted and maintained by a distinct *storage cluster*. Even though the file systems are mounted remotely, each Storage Protect instance is still assigned a unique fileset within the file system provided by the storage cluster. This again enables instance-level safeguarded copies and restorations.

In this setup, *isnap-scripts* are run from the remote cluster and use the REST API service of the storage cluster to manage SGC. In this setup, the restoration of SGC is performed manually, whereby the restore script provides guidance on the console.

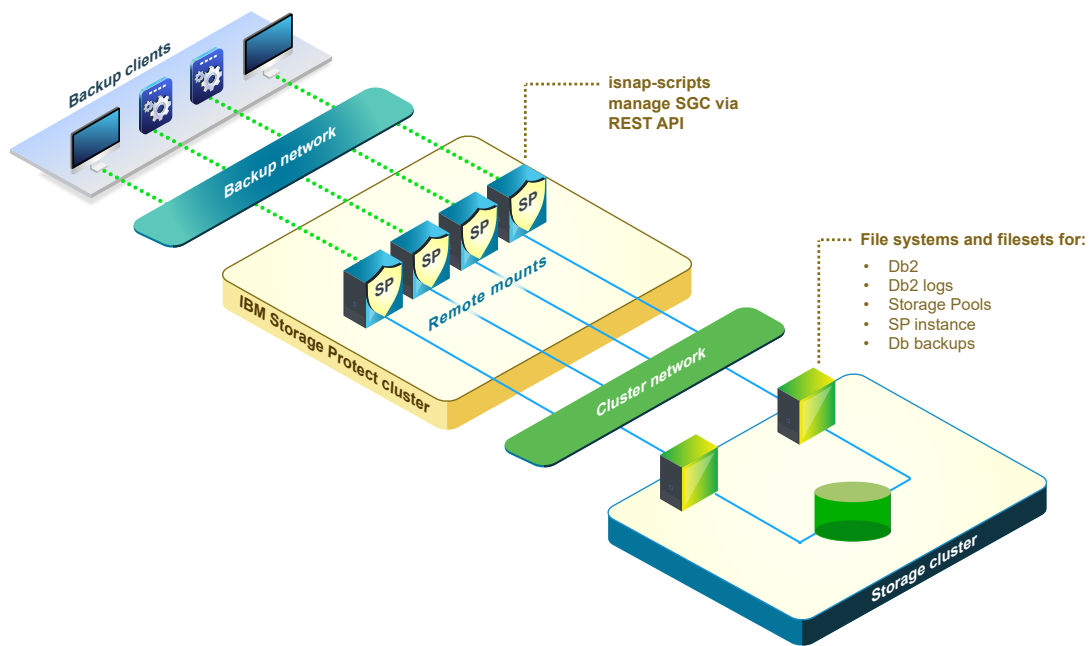


Figure 2: Storage Scale topology with remote and storage cluster

The advantage of the architecture shown in Figure 2 is the separation of administrative domains. The servers in the Storage Protect cluster are administered by the Storage Protect administrators, while the Storage Scale file systems and storage are administered by the Storage Scale administrators. The Storage Protect administrators have no access to the file-system disk in the storage cluster, only to the file system itself.

The Storage Protect servers are more exposed to cyberattacks because they serve backup clients in different networks. If there is an attack on the Storage Protect servers, then the attacker cannot tamper with the file system and SGC infrastructure, because this infrastructure is decoupled from the Storage Protect cluster.

After understanding the supported deployment topologies, it is important to examine how consistent safeguarded copies are created and restored within these environments. The following section describes the detailed workflow that underpins the solution's consistency and recoverability guarantees.

Consistent SGC for IBM Storage Protect on IBM Storage Scale

The provided solution consists of UNIX® shell scripts that handle SGC operations, supporting *IBM Storage Scale version 5.2.1 or later*. Each IBM Storage Protect instance operates within its own dedicated *fileset* in a Storage Scale file system, which ensures logical separation and enables instance-specific SGC. These SGC play a pivotal role in disaster recovery and cyberattack mitigation, allowing data to be restored quickly and precisely for any affected instance.

A key benefit of this approach is that **multiple IBM Storage Protect server instances can coexist on a shared set of file systems**, each with isolated data and metadata areas. This architecture is especially effective in large-scale environments and supports fast, targeted recovery by enabling SGC and restorations at the instance level.

To ensure consistency during SGC operations, the **SGC creation workflow** is divided into three phases:

- **Phase 1:** Connect to the instance's IBM Db2® and initiate a write suspend (executed by the instance user).
- **Phase 2:** Create safeguarded copies for all data areas (can be performed by a privileged user).
- **Phase 3:** Resume Db2 operations (executed by the instance user).

These operations are executed by the script *isnap-create.sh*, which must be run on the server hosting the active Storage Protect instance. The *isnap-create* script creates SGC for all data areas (filesets) of the IBM Storage Protect instance, including database, logs, and storage pools.

Restoration of SGC follows a similarly structured workflow:

- **Phase 1:** Restore the safeguarded copies for all data areas (requires appropriate privilege).
- **Phase 2:** Restart and resume Db2 (must be executed by the instance user).
- **Phase 3:** Start the Storage Protect instance (can be executed by a privileged user).

These operations are executed by the script *isnap-restore.sh*. It is essential that the instance is stopped prior to restoration. If the script is mistakenly run on a different server while the instance remains active elsewhere, it may lead to system inconsistency or unavailability. The script *isnaprestore.sh* restores SGC for all data areas (filesets) of the IBM Storage Protect instance, including database, logs, and storage pools. Manual restoration of SGC is recommended, whereby the script *isnap-restore.sh* provides guidance for the restoration of SGC and does not perform the restoration automatically.

SGC that have expired are not deleted automatically. To accommodate deletion of expired SGC, the script *isnap-del.sh* is used. This script allows deletion of SGC older than a specified number of days by using the parameter `-g age`. The `age` parameter corresponds to the retention time configured for the SGC.

By combining robust SGC mechanisms with fine-grained automation, this integration offers a scalable, secure, and resilient architecture for data protection. It is especially relevant in scenarios demanding rapid recovery from malicious activity or unexpected data loss.

HUK-COBURG cyber resilient solution requirements

HUK-COBURG is a major German mutual insurance company headquartered in Coburg, specializing in auto, liability, legal protection, and life insurance. Founded in 1933, it primarily serves civil servants but has expanded to a broader customer base. The company operates on a cooperative model, meaning its policyholders are also its members. It distributes its products through both physical branches and digital platforms. HUK-COBURG is known for its strong market position in motor insurance and its commitment to customer-oriented service.

Existing environment

The IT department of HUK-COBURG provides company-wide backup services currently leveraging 15 IBM Storage Protect server instances. The IBM Storage Protect server instances use the IBM Storage Scale file system to store backup data and metadata.

There are three IBM Storage Protect clusters, each hosting a set of IBM Storage Protect instances. The IBM Storage Protect instances run in storage-less IBM Storage Scale clusters and remotely mount the file systems provided by the storage cluster (see the section [Topology 2: Storage Protect Instances in a Remote Cluster](#)). These IBM Storage Protect clusters are remotely connected through InfiniBand® network to the storage cluster that consists of IBM Elastic Storage System (ESS) 3000 All-Flash, IBM ESS 5000 NL-SAS (nearline serial-attached SCSI), and IBM SSS6000 Non-Volatile Memory Express (NVMe) storage systems. The IBM Storage Protect clusters and the Storage cluster are stretched across two locations. Figure 3 illustrates the IBM Storage Protect and IBM Storage Scale ESS environments.

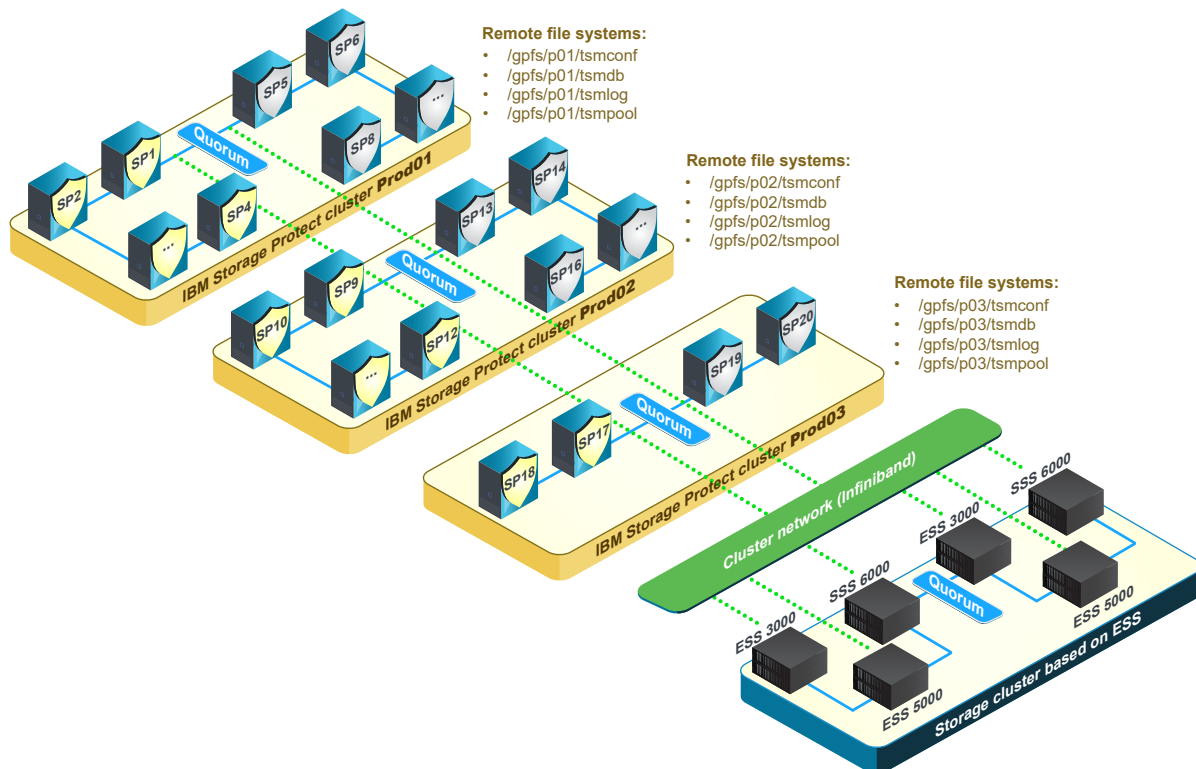


Figure 3: Storage Scale topology with remote and storage cluster based on ESS

As shown in Figure 3, there are three IBM Storage Protect clusters, whereby cluster Prod01 hosts six IBM Storage Protect instances, cluster Prod02 hosts seven instances and cluster Prod03 hosts two instances. Each IBM Storage Protect cluster mounts its own set of file systems for the backup data and metadata from the Storage cluster based on ESS.

HUK-COBURG has a fourth IBM Storage Protect cluster that is configured as a classical IBM Storage Scale cluster using Hitachi Data System storage. This cluster is not shown in Figure 3 because it does not use ESS-based storage. However, *isnap-scripts* are deployed on this cluster as well.

Each IBM Storage Protect instance has one directory in each file system remotely mounted by the Storage Protect cluster. These directories are independent filesets that allow SGC to be created and restored. For example, the Storage Protect instance SP1 has the following directories in the relevant file systems:

```
/gpfs/p01_tsmconf/sp1, /gpfs/p01_tsmdb/sp1, /gpfs/p01_tsmlog/sp1 and  
/gpfs/p01_tsmppool/sp1.
```

With this configuration, all Storage Protect instances of a cluster store their data and metadata in the same file system in different directories (filesets).

The IBM Storage Protect instance use primarily directory-container storage pools and some file storage pools to store the backup data. The backup client workloads include large databases, file servers and virtual machines.

The IBM Storage Protect environment is managed and monitored by the [TSMCluster](#) software provided by [eXstor](#).

The following requirements reflect the regulatory, operational, and recovery objectives that guided HUK-COBURG's adoption of the solution.

Regulatory and recovery requirements

As an insurance company, HUK-COBURG is subject to strict regulatory requirements of the European Union and German authorities. Some of the legal requirements address cyber resilience of the backup environment, demanding a restoration of an IBM Storage Protect instance within hours and not days. Unfortunately, the restoration methods included in IBM Storage Protect can require several days to restore a medium-sized IBM Storage Protect instance. Refer to [IBM documentation](#).

HUK COBURG was looking for a solution to restore an IBM Storage Protect instance within a few hours.

IBM Storage solution for HUK-COBURG

The solution providing consistent SGC for Storage Protect on Storage Scale addresses the cyber resilience requirements of HUK-COBURG. HUK-COBURG co-created this solution with IBM and conducted comprehensive testing of the *isnap-scripts*. HUK-COBURG demonstrated that the restoration of a small-sized IBM Storage Protect instance (200 GB database, 20 TB storage pools) took only 10 minutes.

Based on this testing, HUK-COBURG estimates that the restoration of a medium-sized Storage Protect instance takes up to 4 hours. This is more than eight times faster than the restoration time with legacy Storage Protect methods.

In this section, an example is provided of how the solution was prepared, deployed and operated in HUK-COBURG's environment that is shown in Figure 3.

Preparation

For background on independent filesets, dependent filesets, and the Storage Scale copy command (mmxcp), see the earlier sections describing SGC prerequisites and file system structure.

Before the *isnap-scripts* can be deployed, some preparation steps must be completed (for more information refer to the project documentation at [Storage capacity planning](#)). Briefly, the important steps are:

1. Create an independent fileset for each instance in the relevant file system to allow creation and restoration of SGC for each instance individually (see the section [File system configuration](#)).
2. Ensure that the remote cluster running the Storage Protect instance can communicate with the REST API of the storage cluster. The default port of the REST API is 443.
3. Create a Storage Scale REST API user in the group *snapAdmin* when using the REST API.
4. Adjust the sudo configuration; provide instance users with the required privileges to manage SGC when using the command-line interface.
5. Ensure that sufficient storage capacity is available to hold the SGC copies (see the section [SGC capacity estimation](#)).

File system configuration

To create and restore SGC for each Storage Protect instance individually, all directories used by a Storage Protect instance in the relevant file systems must be independent filesets.

If one or more directories of a Storage Protect instance are not independent filesets, then independent filesets must be created in the relevant file systems, and data must be copied from the ordinary directory into the independent filesets. An independent fileset can only be created for a new directory. The copy process requires a downtime for the Storage Protect instance and sufficient storage capacity in the relevant file systems. Briefly, the copy process can include the following steps:

1. Stop the Storage Protect instance.
2. Rename the existing ordinary directories.
3. Create an independent fileset at the original path of the ordinary directories.
4. Copy files from the ordinary directory to the independent fileset path using either the standard copy command (cp) or with the Storage Scale copy command mmxcp.
5. Ensure that all files were copied, for example by using the UNIX command `find`.
6. Start the Storage Protect instance.
7. Remove the ordinary directories.

The following table shows the configuration of independent filesets for Storage Protect instance *sp1* (see Figure 3):

File system name	File system path	Fileset path for sp1
Tsmconf	/gpfs/p01_tsmconf	/gpfs/p01_tsmconf/sp1
Tsmdb	/gpfs/p01_tsmdb	/gpfs/p01_tsmdb/sp1
Tsmlog	/gpfs/p01_tsmlog	/gpfs/p01_tsmlog/sp1
Tsmpool	/gpfs/p01_tsmpool	/gpfs/p01_tsmpool/sp1

The independent fileset sp1 for the storage pools in the path `/gpfs/p01_tsmppool/sp1` can be further partitioned to accommodate multiple storage pools. HUK-COBURG uses file and directory-container pools and sets quota for each storage pool. This prevents one storage pool from exceeding the capacity limits. To accommodate this, a dependent fileset for each storage pool is created underneath the independent fileset. The following table shows an example of the partitioning of the storage-pool fileset sp1:

Storage pool fileset	Storage pool name	Dependent fileset path
sp1 under <code>/gpfs/p01_tsmppool/sp1</code>	file-pool	<code>/gpfs/p01_tsmppool/sp1/file-pool</code>
	container-pool	<code>/gpfs/p01_tsmppool/sp1/container-pool</code>

As shown in the table, the storage pool *file-pool* has a dependent fileset under `/gpfs/p01_tsmppool/sp1/file-pool`, and the storage pool *container-pool* has a dependent fileset under `/gpfs/p01_tsmppool/sp1/container-pool`. Block and inode quota are defined individually for the dependent fileset representing the storage area for the respective storage pool.

SGC capacity estimation

Consistent SGC are space-efficient right after creation and consume storage capacity over time as files are changed and deleted. Because HUK-COBURG needed predictable sizing for large-scale deployments, the following section quantifies the additional storage capacity required for SGC based on measured workload patterns. Depending on the Storage Protect instance workload in the file system or fileset, the additional capacity required for SGC can be significantly high.

The total storage capacity allocated by SGC for a given Storage Protect instance depends on the number of SGC created per day and the retention time. The number of SGC per day and their retention time define the number of SGC retained for a Storage Protect instance. As more SGC are retained, more storage capacity is consumed.

The IT team of HUK-COBURG conducted tests to estimate the additional storage capacity required for SGC. The tests were conducted for one Storage Protect instance that backed up large volumes of data every day. The results derived from this test were applied to all other Storage Protect instances.

The test revealed that the additional storage capacity depends on the type of Storage Protect workload. Storage Protect workloads include the database, active and archive logs, storage pools, and configuration data. The following table summarizes the additional storage capacity that is required when one SGC is created per day and kept for eight days:

Workload	SGC capacity relative to live data	SGC capacity factor	Used capacity for live data	Additional SGC capacity
Database	60% per SGC	4.8	1 TB	4.8 TB
Logs	83% per SGC	6.6	0.6 TB	4 TB
Storage pools	5% per SGC	0.4	338 TB	130 TB
Configuration	2% per SGC	0.16	0.015 TB	0.002 TB
Total capacity	478.4 TB		339.6 TB	138.8 TB

As shown in the table, each SGC taken for the database fileset (for example, /gpfs/p01_tsmdb/sp1) consumes 60% of the capacity allocated in the fileset. With eight SGC in total, the SGC capacity factor for the database fileset is 4.8 TB. This means that with a database size of 1 TB, eight SGC consume an additional 4.8 TB. With this SGC configuration (one SGC per day and retained for eight days), the additional storage capacity for SGC is 41% higher.

The next table shows the additional storage capacity when one SGC is created per day and retained for 16 days:

Workload	SGC capacity relative to live data	SGC capacity factor	Used capacity for live data	Additional SGC capacity
Database	60% per SGC	9.6	1 TB	9.6 TB
Logs	84% per SGC	13.4	0.6 TB	8.0 TB
Storage pools	4.5% per SGC	0.72	338 TB	243.4 TB
Instance home	2% per SGC	0.32	0.015 TB	0.004 TB
Total capacity	600.6 TB		339.6 TB	261.0 TB

With this SGC configuration (one SGC per day retained over 16 days), the additional storage capacity required for SGC is 77% higher.

As shown in the tables, most storage capacity is consumed for the storage-pool fileset. Tests have shown that the additional storage capacity for SGC in the storage-pool fileset is equivalent to the daily change rate in the Storage Protect instance. The daily change rate applies to all SGC created per day. For example, if the daily change rate is 5% and two SGC are created per day, then two SGC share the 5% daily change rate.

Deployment

After the preparation, the *isnap-scripts* can be installed and configured on all servers running Storage Protect instances. Detailed instructions can be obtained from the project documentation available at [Installation and configuration](#).

The *isnap* project includes the following scripts:

File name	Description
snapconfig.json	Describes the Storage Protect instance configuration and REST API connection
isnap-create.sh	Creates consistent SGC on Storage Scale filesets used by the Storage Protect instance
isnap-delete.sh	Deletes expired consistent SGC for Storage Protect instance
isnap-restore.sh	Restores consistent SGC for the Storage Protect instance
isnap-list.sh	Lists all consistent SGC for the Storage Protect instance
isnap-fscap.sh	Prints the file system and SGC capacity allocation
isnap-wrapper.sh	Wraps the other <i>isnap-scripts</i> and can be used with schedulers

To deploy the solution, copy the content scripts and the configuration file to a common directory (such as */usr/local/bin*) on each server running Storage Protect instances.

Note: Do not copy the scripts into a directory where SGC are created and restored. The restore process may revert to an older version of the scripts.

Make the scripts executable and allow all Storage Protect instance users to execute them. Usually, the Storage Protect instance users are in a common group. For example, if the common Storage Protect instance user group is *tsmsrvrs*, then apply the following permissions on the scripts:

```
-rw-r--r--. 1 spl tsmsrvrs
```

Now create the SGC configuration for each Storage Protect instance in the configuration file *snapconfig.json*. The configuration parameters are available; only a subset of parameters is required, as shown in the following table:

Parameter	Description
<code>instName</code>	Specifies the instance user name (required).
<code>dbName</code>	Specifies the Db2 database name, default is TSMDB1.
<code>snapPrefix</code>	Specifies the SGC name prefix used to create and restore SGC (required).
<code>dirsToSnap</code>	Specifies the file system name and fileset name to be snapped; the syntax is: <code>fspath+fsetname</code> . If fileset name is not provided, global SGC are used (required).
<code>snapRetention</code>	Sets the SGC retention time in days; default is 0 days. SGC cannot be deleted during retention time.
<code>serverInstDir</code>	Specifies the server instance directory (where <i>dsmserv.opt</i> resides); required only if it is different from the instance user home directory. Default is instance user home directory.
<code>sudoCommand</code>	Specifies the full path of the command used by the instance user to run commands with root privileges; default is <i>/usr/bin/sudo</i> .
<code>apiServerIP</code>	Specifies the IP address or host name of the REST API server; required when the REST API is used instead of the CLI.
<code>apiServerPort</code>	Specifies the IP port of the REST API server; required when the REST API is used; default is 443.
<code>apiCredentials</code>	Provides the obfuscated authentication string for the REST API.
<code>AutoRestore</code>	Controls whether restore is performed automatically or manually when using the CLI; when the REST API is used, restore is always manual regardless of <code>autoRestore</code> . Default is <code>false</code> (manual restore).

Following is an example of the SGC configuration for Storage Protect instance sp1:

```
[
  {
    "instName": "sp1",
    "dbName": "TSMDB1",
    "snapPrefix": "sp1",
    "snapRetention": "8",
    "serverInstDir": "/path-to-server-instance-dir",
    "sudoCommand": "/usr/bin/sudo",
    "apiServerIP": "api-server-ip",
    "apiServerPort": "api-server-port",
    "apiCredentials": "Obfuscated authentication string",
    "dirsToSnap": ["tsmdb+sp1", "tsmlog+sp1", "tsmpool+sp1", "tsmconf+sp1"],
    "autoRestore": false
  }
]
```

The configuration for Storage Protect instance sp1 is defined in one JSON object encapsulated by curly brackets. The configuration for more Storage Protect instances can be added to the same configuration file as separate JSON objects.

After the configuration file has been adjusted, it is copied to all Storage Protect servers in the same directory (for example, */usr/local/bin*). Ensure that the Storage Protect instance-user group has read access to the configuration file.

Finally, the scripts are tested in the Storage Protect instance-user environment. In the following example, the instance-user invokes the scripts *isnap-list.sh* and *isnap-create.sh*:

```
# su - sp1
$ /usr/local/bin/isnap-list.sh
$ /usr/local/bin/isnap-create.sh -r
$ /usr/local/bin/isnap-list.sh
```

To make the invocation of the *isnap* scripts easier for the instance-user, add the path */usr/local/bin* to the PATH variable in the instance-user profiles:

```
# su - sp1
$ vi .profile
PATH=$PATH:/usr/local/bin
export PATH
```

If the scripts operate correctly, schedules can be created for the creation and deletion of SGC.

Operations

After successfully deploying and testing the *isnap* scripts, the creation and deletion of SGC can be scheduled. The different methods for scheduling are:

- Use a system scheduler (cron)
- Use a Storage Protect client schedule in the Storage Protect instance
- Use an external scheduler

HUK-COBURG uses an external scheduler provided by the TSMCluster software that is developed by eXstor GmbH. The creation of SGC is scheduled individually for each IBM Storage Protect instance and executes the script `isnap-create.sh` as the Storage Protect instance user. The scheduled start time depends on the workload for each instance. Usually, the SGC creation for an IBM Storage Protect instance is scheduled after the daily backup workload has finished. The deletion of SGC is scheduled after the SGC creation by executing the script `isnap-del.sh`.

The TSMCluster software provides a sensor script to schedule SGC operations and monitors the SGC operations. If SGC creation or deletion fails, or the number of available SGC is not as expected, then TSMCluster software raises an alert with the HUK-COBURG IT operations team.

HUK-COBURG also collects the SGC storage allocation. For this purpose, the TSMCluster software periodically executes the script `isnap-fscap.sh`, which captures SGC storage allocation in a numeric format. These numbers are used to analyze the SGC storage allocation over time and to make predictions for future capacity needs.

Furthermore, HUK-COBURG monitors the storage capacity consumption in the IBM Storage Scale file systems and fileset for all Storage Protect instances. It uses thresholds configured in Storage Scale that generates alerts when the capacity used in file systems crosses a threshold.

Solution outcomes – feedback from the customer

HUK-COBURG successfully tested the solution delivering consistent SGC for IBM Storage Protect on IBM Storage Scale over the course of one year. The evaluation confirmed that the solution had no negative impact on backup workloads or Storage Protect housekeeping operations.

Throughout the testing phase, HUK-COBURG appreciated the high-quality code and responsive support provided by IBM. With assistance from IBM Expert Labs, the solution was seamlessly deployed into HUK-COBURG's production environment.

The solution meets HUK-COBURG's core requirements for cyber resilience in its backup infrastructure. It significantly improves recovery capabilities, reducing the recovery time of a complete IBM Storage Protect instance following a cyberattack from 34 hours to less than 4 hours. The creation of SGC takes only a few minutes.

IBM's solution exceeded our expectations. The dramatic reduction in recovery and backup times has strengthened our cyber-resilience strategy and given us peace of mind.

— Renar Grunenber, IT Department HUK-COBURG

HUK-COBURG found that the solution providing consistent SGC for IBM Storage Protect on IBM Storage Scale does not replace traditional IBM Storage Protect database and storage-pool backups. Instead, it offers an additional, rapid method to recover an IBM Storage Protect instance completely, including the database and all storage pools, after a cyberattack. The solution is easy to deploy and use, and it provides an additional layer of protection for IBM Storage Protect instances.

Another insight is that the manual restore procedure provides more flexibility in case of unforeseen issues. The manual restore procedure is executed by the `isnap-restore.sh` script with the configuration parameter `autoRestore` set to `false`, or when the REST API is used for managing SGC.

Conclusion

The solution addresses the growing risk of ransomware and destructive cyber events by strengthening the recoverability of the backup estate, using IBM Storage Protect on IBM Storage Scale. One of the most significant benefits is the reduction in restore times for a complete IBM Storage Protect instance—from 34 hours down to less than 4 hours, adding substantial value to IBM Storage Protect by enhancing its recovery capabilities and overall cyber resilience. By enabling consistent, immutable safeguarded copies and instance-level restoration, the approach helps ensure clean recovery points and operational continuity in the face of cyber disruption.

HUK-COBURG's extensive efforts in testing and collaboration played a key role in making the solution generally available to all IBM Storage Protect customers. Refer to [Consistent SGC for IBM Storage Protect](#).

The successful implementation has sparked interest among other customers who have now begun testing the solution. These results demonstrate a practical, validated path to meeting stringent recovery objectives, while complementing existing database and storage-pool backups. As organizations continue to harden their cyber-resilience posture, this solution provides a scalable, operationally simple option to reduce recovery time and risk across diverse Storage Protect deployments.

About the authors

HUK-COBURG

Renar Grunenberg joined HUK-COBURG in 1993 with a Master of Science degree in electrical engineering. He has led the storage-and-backup team since 1996. For decades, Renar has gained project-managing and operational experience on all relevant storage and backup topics. Renar integrated IBM storage solutions like Storage Protect, IBM Tape System Library Manager (TSLM), Storage Archive, Storage Scale, Scale Protocol Services and InfiniBand connected ESS-Systems in HUK-COBURG's IT landscape. He managed projects with fiber network attached enterprise disk systems from Hitachi Data systems as well as IBM tape-library solutions for all operating-system platforms including IBM z/OS®. Renar is HUK-COBURG's principal Storage Scale specialist and therefore an attendee of the sponsor-user program for Storage scale. His current focus is to enhance the cyber-resilience strategy related to all storage and backup processes.

Christian Buettner works for HUK-COBURG since 2001. After completing his training as an IT specialist, he initially worked as a java programmer for two years before joining HUK-COBURG 's storage-and-backup team in 2004. Christian is a Storage Protect veteran. Starting with IBM Tivoli® Storage Manager (TSM) version 3.1 on z/OS, Christian played a key role in migrating HUK-COBURG 's Storage Protect landscape to the IBM Power platform and later to Red Hat Enterprise Linux running on Advanced Micro Devices (AMD) powered servers. In doing so, Christian gained extensive experience with Storage Protect related networks such as Ethernet, Fibre, and InfiniBand as well with several generations of IBM tape libraries and enterprise disk systems from Hitachi Data Systems, IBM, or Dell. After testing and implementing the software and hardware for SGC in 2025 with the HUK SP-Team, he now plans to implement isolated-recovery environments and recovery assurance automation to develop Storage Protect as a holistic cyber resilience platform.

IBM

Nils Haustein is a **Senior Technical Staff Member** in IBM Client Engineering Storage EMEA, based in IBM Germany. Nils joined IBM in 1994 after successfully achieving a Master of Science degree in electrical engineering and computer science at the Technical University, Chemnitz. As an IT specialist, he has worked in various technical positions at IBM in manufacturing, technical support, development, technical pre-sales, and lab services. In 2023, he joined the IBM Client Engineering Storage EMEA team where he developed critical assets for successful pilots. Nils focuses on file and objects storage solutions for backup, archiving, hybrid cloud, AI, and cyber-resilience. Nils co-authored a book titled *Storage Networks Explained (Speichernetze)*. As a leading master inventor at IBM, Nils has filed more than 160 patents and guides other inventors at IBM.

Nilesh Bhosale is a Lead Architect with IBM's Storage Protect and Copy Data Management development teams, based at the IBM India Systems Development Lab in Pune. With over two decades of expertise in enterprise storage, distributed filesystems, cloud storage, data protection, security, and resiliency, Nilesh has been instrumental in designing and delivering innovative solutions for global customers. He holds multiple patents, has authored technical articles and blogs, and is passionate about collaborating with clients to understand their challenges and drive meaningful product enhancements. As a leading master inventor at IBM, Nilesh has filed more than 28 patents, and he mentors other inventors at IBM.

Hemanand Gadgil serves as the Product Manager for Storage Solution Engineering within IBM Storage Solutions, based at the IBM India Systems Development Lab in Pune. He joined IBM in 2025, and he brings over 23 years of expertise in enterprise storage, cloud storage, and data protection, with extensive experience in data-center consulting, implementations, and migrations for large-scale infrastructure programs. Hemanand has co-authored multiple solutions focused on enterprise storage, data protection, business continuity, and resiliency.

Satyamev Vijay Choudhary is an accomplished Solution Architect and Consultant with an experience of over 17 years in enterprise IT infrastructure. He specializes in designing and implementing solutions for banking, financial services, and telecom sectors, with expertise in storage (block, file, object), virtualization (VMware), cyber recovery, and hybrid cloud platforms (IBM, Red Hat® OpenShift®, AWS, Azure). He is skilled in high-level design/low-level design documentation, request for proposal (RFP)/request for price quotation (RPQ) responses, and automation technologies like Python, PowerShell, and robotic process automation (RPA).

© Copyright IBM Corporation 2026

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the
United States of America
February 2026

IBM, IBM Spectrum, IBM Db2, IBM z/OS, and the IBM logo are trademarks or registered trademarks of International Business Machines Corporation, in the United States and/or other countries. IBM Tivoli is a trademark of International Business Machines Corporation. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on ibm.com/trademark.

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

InfiniBand is a trademark or registered trademark of the InfiniBand Trade Association.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates. All client examples cited or described are presented as illustrations of the manner in which some clients have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions. Generally expected results cannot be provided as each client's results will depend entirely on the client's systems and services ordered. It is the user's responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs. THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.