

IBM® Storage

Cyber Resiliency Solution for IBM Spectrum Scale

IBM Storage Team



© Copyright International Business Machines Corporation 2019.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this document	1
Executive summary	1
Support for the Blueprint and its configurations	2
Requesting assistance	2
Scope	3
Prerequisites	3
National Institute of Standards and Technology framework	3
Implementing a Cyber Resiliency solution with IBM Spectrum Scale	5
IBM Spectrum Scale	5
IBM Spectrum Protect	5
IBM Tape Storage	5
Use cases considered for a Cyber Resiliency solution for IBM Spectrum Scale	5
Cyber Resiliency solution for IBM Spectrum Scale architectural overview	6
Configuring a Cyber Resiliency solution for IBM Spectrum Scale	7
IBM Spectrum Protect server configuration	7
IBM Spectrum Protect client configuration on IBM Spectrum Scale	8
Phases of an IBM Spectrum Scale Cyber Resiliency solution	10
Initialization Phase	11
Backup Synchronization technique in the Production phase	13
Production phase	14
Detection of Anomaly phase	17
Failover and Failback phase	18
Summary	20
Notices	21
Trademarks	22
Terms and conditions for product documentation	23
Applicability	23
Commercial use	23
Rights	23
Privacy policy considerations	23



About this document

This document is intended to facilitate the deployment of the Cyber Resiliency solution for IBM® Spectrum Scale. This solution is designed to protect the data on IBM Spectrum™ Scale from external cyberattacks or insider attacks using its integration with IBM Spectrum Protect™ and IBM Tape Storage. To complete the tasks that it describes, you must understand IBM Spectrum Scale™, IBM Spectrum Protect, and IBM Tape Storage architecture, concepts, and configuration.

The information in this document is distributed on an as-is basis without any warranty that is either expressed or implied. Support assistance for the use of this material is limited to situations where IBM Spectrum Scale or IBM Spectrum Protect are supported and entitled, and where the issues are specific to a blueprint implementation.

Executive summary

In today's data-driven world, an organization's information and data is considered the most important asset to its business, and it can serve as a key asset for the growth of an organization. As more and more data is collected by businesses, organizations, and companies, data volume is growing at a staggering pace.

With this exponential data growth, there is an increased need to protect the data from various cyberattacks in the form of malware and ransomware. These cyberattacks can have a catastrophic impact on an organization, and can result in devastating financial losses and affect an organization's reputation for many years.

The financial impact of cyberattacks is rising. According to Ponemon's *Cost of a Data Breach Report 2019*,¹ the average cost of a data breach is estimated at a staggering USD 3.92 million. Moreover, that same Ponemon's report also placed the average chance of experiencing a data breach over the next two years at 29.6%. Therefore, it's a matter of when, not if.

These cyberattacks can happen in several forms. They can be in the form of malware or ransomware targeted at stealing confidential data or holding users' information for ransom. Sometimes these attacks are targeted to destroy confidential and critical data to cripple organizations. Moreover, according to Verizon,² 34% of data breaches involved internal actors.

¹ <https://www.ibm.com/security/data-breach?lnk=ushpv18l1>

² <https://enterprise.verizon.com/resources/reports/2019-data-breach-investigations-report.pdf>

Per Wikipedia,³ *Cyber Resiliency* refers to an entity's ability to continuously deliver the intended outcome despite cyber events. Assuming that you already have an infrastructure that uses some of the current data protection techniques, such as backups, snapshots, and replication, the next step is expanding your current infrastructure to add the necessary cyber resiliency focus.

IBM Spectrum Scale is a high-performance, highly available, clustered file system available on a variety of platforms (including the public cloud service providers). It provides concurrent access to a single file system or set of file systems from multiple nodes. IBM Spectrum Scale has multiple data access points (via different protocols) where data in the form of files and objects is directly accessible by end users as well as applications.

Moreover, the product integrates with an organization's external directory services for different types of authentication and authorization. Unlike block storage products, these aspects make the attack vector for such products much wider, making them more vulnerable to cyberattacks and increasing the overall threat index. IBM Spectrum Protect can simplify data protection, whether data is hosted in physical, virtual, software-defined, or cloud environments.

IBM Spectrum Scale provides conventional features, such as snapshots, replication, encryption, and immutability, which can be implemented to reduce cyber security threats. In addition, IBM Spectrum Scale's integration with IBM Spectrum Protect provides an effective backup and restore mechanism. This feature integration can be effectively leveraged to provide a Cyber Resiliency approach for organizations that are planning to improve cyber threat resolution, reduce costs, and deliver quick recovery for the environment.

Adding IBM Tape Storage offers cost-effective, long-term backup and archive WORM storage, with a true physical air gap and total separation from ransomware and cyber-attacks.

Support for the Blueprint and its configurations

The Cyber Resiliency solution for IBM Spectrum Scale provides an integrated support experience for clients. The information in this document (referred to throughout as the *Blueprint*) is distributed on an "as-is" basis without any warranty that is either expressed or implied. Support for the underlying components that make up this solution are provided by way of the standard procedures and processes that are available for each of those components, as governed by the support entitlement that is available for those components. For more information about these components, see "Prerequisites" on page 3.

Requesting assistance

All components of the solutions are part of the unified support structure. Support assistance for the solution that is described in this Blueprint is available by requesting assistance for any of the components in the solution. This is the preferred method.

³ https://en.m.wikipedia.org/wiki/Cyber_resilience

Scope

This Blueprint provides the following information:

- A solutions architecture and related solution configuration workflows, with the following essential software components:
 - IBM Spectrum Scale
 - IBM Spectrum Protect
 - IBM Tape Storage
- Detailed technical configuration steps for building an end-to-end solution

This technical report does not include the following:

- Provide scalability and performance analysis from a user perspective
- Provide claims of creating totally isolated air-gap infrastructure
- Replace any official manuals and documents issued by IBM

Prerequisites

This technical report assumes basic knowledge of the following prerequisites:

- IBM Spectrum Scale installation and configuration
- IBM Spectrum Protect installation and configuration
- IBM Tape Storage installation and configuration

National Institute of Standards and Technology framework

As systems became linked with external networks, organizations adopted a *defense-in-depth* security mode so that if the perimeter was breached, there were additional layers of security to protect critical information from falling into the wrong hands. The focus was on the technical aspects of recovery. However, these measures are no longer sufficient for protection against cyberattacks.

Organizations are beginning to understand that traditional device-centric and technology-centric security measures, such as firewalls, fail to provide security in the cyber ecosystem. Moving forward, you must take a holistic approach across your data, applications, and the entire infrastructure to not only recover, but prevent (or at the very least minimize) the attack.

Some of the following factors are considered for designing a Cyber Resiliency approach:

- Although regulations continue to play an important role, consumers decide the ultimate outcomes for a business.
- For implementing effective Cyber Resiliency approach, it must be changed from a reactive approach to a proactive approach. A repeated cycle of planning, protecting, testing, and learning must be implemented by a Cyber Resiliency team.
- Most organizations' backup and disaster recovery plans are designed around the fact that most disasters are caused by either technical failures or human errors, with secondary concern about natural disasters. Modern data protection approaches must also consider data compromise due to cyber events, and be implemented accordingly.

- As attackers are getting smarter, approaches should consider continuous improvements, innovations, and reengineering to address the newer threats that are challenging organizations.
- Though effort is made to leverage existing infrastructure, modern technologies help automate systems to deal more effectively with cyber threats.

In order to more effectively deal with cyber events, the National Institution of Standards and Technology (NIST) provides a policy framework of computer security guidance regarding how organizations can assess and improve their ability to prevent, detect, and respond to cyberattacks. This framework is an industry-accepted methodology for building a plan to develop and implement safeguards to ensure delivery of critical business services.

As shown in Figure 1, a Cyber Resiliency plan is a continuous process that needs to be repeated in the environment to safeguard data from cyberattacks.



Figure 1 NIST Cyber Security framework

The NIST framework is a set of five Cyber Security functions:

- **Identify:** NIST recommends building organizational understanding during the *Identify* stage so that business IT systems can be confidently restored to their operational state. It is important to identify what must be protected, and then prioritize your protection plan.
- **Protect:** During the *Protect* stage, implement various safeguards, such as identity management, access control, awareness and training, data security, code currency procedures, and data protection technology, to ensure delivery of critical services.
- **Detect:** The best way to reduce costs during an event is to detect it early, and then rapidly recover. The point of the *Detect* stage is implementing activities and technologies to identify anomalies and events that are out of the ordinary. This enables you to quickly respond and limit the damage by containing the event.

- **Respond:** In the *Response* state, develop and implement appropriate activities to take actions regarding a detected cybersecurity incident.
- **Recover:** In the *Recover* stage, develop and implement appropriate activities to maintain plans for resilience, and to restore any capabilities or services that were impaired due to a cybersecurity incident. In this stage, the goal is to get a compromised environment back up and running quickly and efficiently.

Implementing a Cyber Resiliency solution with IBM Spectrum Scale

This section describes the components and solution building blocks used for implementing a Cyber Resiliency solution.

IBM Spectrum Scale

IBM Spectrum Scale is a high-performance, highly available, clustered file system available on a variety of platforms. IBM Spectrum Scale provides concurrent access to a single file system or set of file systems from multiple nodes. On-premise, IBM Spectrum scale nodes can be SAN-attached, network-attached, a mixture of SAN-attached and network-attached, or in a Shared Nothing cluster configuration.

IBM Spectrum Protect

IBM Spectrum Protect can simplify data protection, whether data is hosted in physical, virtual, software-defined, or cloud environments. With IBM Spectrum Protect, organizations can choose the right software to manage and protect their data while also simplifying backup administration, improving efficiency, delivering scalable capacity, and enabling advanced capabilities.

IBM Tape Storage

IBM Tape Storage offers cost-effective, long-term backup and archive WORM storage, with a true physical air gap and total separation from ransomware and cyberattacks. Tape is used to optimize data protection costs, and mitigates the risk of ransomware for data-centric organizations. At a cost of less than a half a cent per GB (gigabyte), it is also an extremely cost-effective solution.

Use cases considered for a Cyber Resiliency solution for IBM Spectrum Scale

The architectural design in this Cyber Resiliency solution addresses the following use cases:

- As a storage architect and administrator, data should be safeguarded from virus attacks, ransomware encryption, or deletion by a malicious user.
- As a storage architect and administrator, data is a most-important asset, and the business of my organization relies on the data on the storage system. Business can continue even if the data on the primary system holding the data has been compromised.
- Multiple copies of data are maintained using multiple features of data protection, even if one or more copies of data are compromised.

- Copies of data are available in an immutable format to avoid overriding valid copies of data. This state provides the ability to restore valid copies of the data at a remote system to validate the authenticity of recovered data.
- Copies of my data are stored in an air-gapped environment where only authorized personnel have access to the data.
- Avoid the human element from accessing and compromising all copies of data, with a provision to store multiple copies of data at different locations, and to separate administrative access for the different copies of data.

Cyber Resiliency solution for IBM Spectrum Scale architectural overview

Figure 2 shows the high-level architectural overview of a Cyber Resiliency solution to achieve protection of data on an IBM Spectrum Scale cluster.

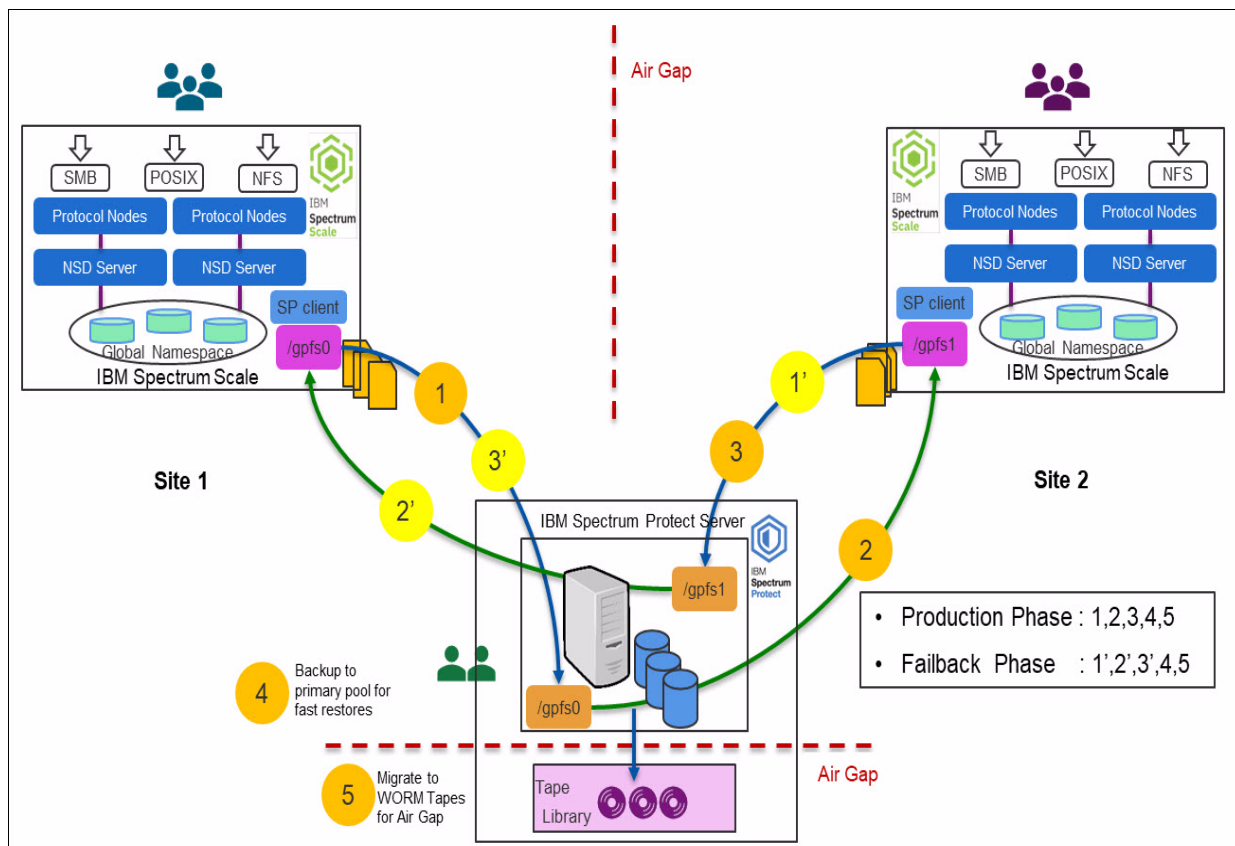


Figure 2 Architectural overview

In this test setup and validation, IBM Spectrum Scale clusters are configured on two sites. One of the sites is a production site, and the other site is an air-gapped site where only authorized personnel and processes have access. Designing an air-gapped system is out of the scope of this document.

This proposed solution consists of two identically configured file systems of the same capacity and characteristics: /gpf0 (the production file system) and /gpf1 (the failover file system) on two IBM Spectrum Scale clusters at different sites. The IBM Spectrum Protect server is used to facilitate the transmission of data across two IBM Spectrum Scale clusters.

The IBM Spectrum Protect client is used to synchronize the two file systems across the IBM Spectrum Scale clusters. This solution uses IBM Spectrum Scale's **mmbackup** utility, which internally uses the IBM Spectrum Protect backup-archive (BA) client to achieve higher performance and scalability of the synchronization mechanism.

In this solution approach, data from the production IBM Spectrum Scale cluster is backed up incrementally to the IBM Spectrum Protect server. All incremental changes are then synchronized with the air-gapped IBM Spectrum Scale cluster. The transmission of data uses file lists created using the **mmbackup** utility and reused for synchronized processing.

When the data on the production system is compromised, or when there is a complete failure of the production system, applications using the data can be directed to the failover system to maintain business continuity. At this time, the air-gapped IBM Spectrum Scale cluster assumes the role of the production IBM Spectrum Scale cluster.

After the original production IBM Spectrum Scale cluster is recovered, incremental data from the air-gapped IBM Spectrum Scale cluster is synchronized with it. The application can now be redirected back to the original production IBM Spectrum Scale cluster.

Consider the following information:

- File systems on the production and failover IBM Spectrum Scale clusters have different mount points. Therefore, appropriate changes need to be applied to the application accessing the data after the failover.
- When production and failover sites are in multiple locations, then I/O latency and bandwidth for processing backup and restore needs to be considered.
- This solution allows multiple failover systems. This can be implemented based on the feasibility and business need to maintain multiple copies of data. In this example, only one production and one failover IBM Spectrum Scale cluster has been considered.

Configuring a Cyber Resiliency solution for IBM Spectrum Scale

This solution approach is designed using standard IBM Spectrum Scale and IBM Spectrum Protect software components.

IBM Spectrum Protect server configuration

This solution considers a single IBM Spectrum Protect server for backup and restore operations. For setting up an IBM Spectrum Protect server, see the [IBM Spectrum Protect server blueprints](#).

Note: All of the configurations described in the following sections are for illustration only. Actual configurations vary in real environments and need to be planned appropriately.

This document discusses the IBM Spectrum Protect configuration where local disk-based storage was used as a primary storage media on the IBM Spectrum Protect server. However, it is advised to configure a tape pool to be used as an additional air-gapped mechanism to store the data. Tape serves as the proven mechanism and building block of a complete modern data protection plan. It provides a low-cost, long-term data archiving solution.

Tape provides a *write once read many* (WORM) capability that can be used to thwart ransomware attacks. Tape achieves this functionality by ensuring an air gap between live data and recovery data when both the production and another site are compromised in a ransomware attack.

To configure the IBM Spectrum Protect server, complete the following steps:

1. After the storage pools are configured on an IBM Spectrum Protect server, register the production IBM Spectrum Scale cluster nodes on the IBM Spectrum Protect server, as shown in Figure 3.

```
Protect: SSS-IP06>
Protect: SSS-IP06>reg node sss-ip02 password maxnummp=10
ANR2060I Node SSS-IP02 registered in policy domain STANDARD.
Protect: SSS-IP06>reg node sss-ip03 password maxnummp=10
ANR2060I Node SSS-IP03 registered in policy domain STANDARD.
Protect: SSS-IP06>
```

Figure 3 Register the production IBM Spectrum Scale cluster nodes

2. Register the IBM Spectrum Scale cluster nodes from the air-gapped IBM Spectrum Scale cluster on the same IBM Spectrum Protect server, as shown in Figure 4.

```
Protect: SSS-IP06>
Protect: SSS-IP06>reg node sss-ip09 password maxnummp=10
ANR2060I Node SSS-IP09 registered in policy domain STANDARD.
Protect: SSS-IP06>reg node sss-ip10 password maxnummp=10
ANR2060I Node SSS-IP10 registered in policy domain STANDARD.
Protect: SSS-IP06>
```

Figure 4 Register air-gapped IBM Spectrum Scale cluster nodes

3. After the nodes are registered, grant proxy authority of the IBM Spectrum Scale cluster nodes to the IBM Spectrum Scale cluster nodes on the other site. This enables IBM Spectrum Scale cluster nodes on one site to perform restore operations of the data backed up from the IBM Spectrum Scale cluster nodes from another site. Figure 5 shows that nodes on the air-gapped sites sss-ip09 and sss-ip10 were provided authority to perform restores of the data backed up from the production IBM Spectrum Scale cluster node sss-ip02.

```
Protect: SSS-IP06>
Protect: SSS-IP06>
Protect: SSS-IP06>grant proxynode target=sss-ip02 agent=sss-ip09,sss-ip10
ANR0140I GRANT PROXYNODE: success. Node SSS-IP03 is granted proxy authority to node SSS-IP02.
ANR0140I GRANT PROXYNODE: success. Node SSS-IP09 is granted proxy authority to node SSS-IP02.
ANR0140I GRANT PROXYNODE: success. Node SSS-IP10 is granted proxy authority to node SSS-IP02.
Protect: SSS-IP06>
Protect: SSS-IP06>
```

Figure 5 Grant proxy authority to the IBM Spectrum Scale cluster nodes

4. Grant proxy authority to the IBM Spectrum Scale cluster nodes on the production site for the IBM Spectrum Scale cluster nodes on the air-gapped site.

IBM Spectrum Protect client configuration on IBM Spectrum Scale

After installing the IBM Spectrum Protect BA client on the IBM Spectrum Scale cluster nodes, they need to be configured for IBM Spectrum Protect backup:

1. Configure `dsm.sys` on the participating IBM Spectrum Scale cluster nodes to point to the IBM Spectrum Protect server, as shown in Figure 6 on page 9.

```

[root@sss-ip02 bin]#
[root@sss-ip02 bin]# cat dsm.sys
*****
*
* Sample Client System Options file for UNIX (dsm.sys.smp)
*
*****

* This file contains the minimum options required to get started
* using the Backup-Archive Client. Copy dsm.sys.smp to dsm.sys.
* In the dsm.sys file, enter the appropriate values for each option
* listed below and remove the leading asterisk (*) for each one.

* If your client node communicates with multiple servers, be
* sure to add a stanza, beginning with the SERVERNAME option, for
* each additional server.

*****

Servername    sss-ip06
NODENAME      sss-ip02
ASNODENAME    sss-ip09
PASSWORDACCESS GENERATE
  COMMMethod   TCPip
  TCPPort      1500
  TCPServeraddress sss-ip06.tuc.stglabs.ibm.com

[root@sss-ip02 bin]#

```

Figure 6 Sample dsm.sys on the IBM Spectrum Scale cluster node

2. Configure dsm.opt on the participating IBM Spectrum Scale cluster nodes to point to the IBM Spectrum Protect server, as shown in Figure 7.

```

[root@sss-ip02 ~]# cat /opt/tivoli/tsm/client/ba/bin/dsm.opt
*****
*
* Sample Client User Options file for UNIX (dsm.opt.smp)
*
*****

* This file contains an option you can use to specify the
* server to contact if more than one is defined in your client
* system options file (dsm.sys). Copy dsm.opt.smp to dsm.opt.
* If you enter a server name for the option below, remove the
* leading asterisk (*).

*****

* Servername      A server name defined in the dsm.sys file
Servername        sss-ip06
[root@sss-ip02 ~]#

```

Figure 7 Sample dsm.opt configuration on the IBM Spectrum Scale cluster nodes

3. After configuration, create a test file on an IBM Spectrum Scale file system. Next, verify the basic backup, query, and restore operations from the IBM Spectrum Scale cluster nodes from both the sites, as shown in Figure 8.

```

[root@sss-ip02 ~]# dsmc sel /ibm/gpfs0/testfile
[root@sss-ip02 ~]# dsmc query ba /ibm/gpfs0/testfile
[root@sss-ip02 ~]# dsmc restore /ibm/gpfs0/testfile -replace=all
[root@sss-ip09 ~]# dsmc sel /ibm/gpfs1/testfile2
[root@sss-ip09 ~]# dsmc query ba /ibm/gpfs1/testfile2
[root@sss-ip09 ~]# dsmc restore /ibm/gpfs1/testfile2 -replace=all

```

Figure 8 Verify test backup and restore operations from IBM Spectrum Scale cluster nodes

Phases of an IBM Spectrum Scale Cyber Resiliency solution

This section describes the various phases of the Cyber Resiliency solution with IBM Spectrum Scale. In this solution approach, data on the production site IBM Spectrum Scale cluster is isolated to an air-gapped IBM Spectrum Scale cluster using the IBM Spectrum Protect backup and restore mechanism.

The data on the production IBM Spectrum Scale cluster is backed up incrementally to the IBM Spectrum Protect server. All incremental changes are synchronized with the air-gapped IBM Spectrum Scale cluster by restoring incrementally backed up files from the production to the air-gapped IBM Spectrum Scale cluster. The transmission of the data uses the *file list* created from the IBM Spectrum Scale's **mmbackup** utility, and this list is reused for the synchronization purpose.

IBM Spectrum Scale file systems can hold huge files and data. For higher parallelism and scalability, the backup processing can happen on the file set level. For more information about scale out backup architecture, see [Petascale Data Protection](#).

IBM Spectrum Protect server achieves the isolation purpose for Cyber Resiliency. It isolates the data continuously from the production IBM Spectrum Scale cluster. This helps protect valid copies of data on the IBM Spectrum Protect server. If data on the production cluster gets compromised, it can be immediately restored from the IBM Spectrum Protect server using the uncompromised data backed up on the IBM Spectrum Protect server.

Then, depending on each organization's policies, this data can be offloaded to the WORM capable tapes to achieve an extra level of protection for tampering with the non-compromised data. It is strongly advised to have separate administrators for IBM Spectrum Scale clusters at both locations and for the IBM Spectrum Protect server, therefore removing the possibility of tampering with all valid copies of data by human intervention.

IBM Spectrum Protect detects anomalies to workload patterns to alert administrators of potential ransomware infections for workloads, enabling clients to be aware of possible attacks and mitigate them before they spread. In this solution, when the anomalies are detected, synchronization from the production cluster to the air-gapped cluster needs to be stopped to prevent compromising data on the air-gapped cluster.

Additionally, point-in-time snapshots created at the air-gapped IBM Spectrum Scale cluster help in keeping multiple snapshot copies of data. In case of compromise, data can be restored to the air-gapped cluster and validated for the good copies of data. In case of a total compromise of the production IBM Spectrum Scale cluster, applications using the data can be failed over to the air-gap IBM Spectrum Scale cluster, assuming that data is not compromised at the air-gapped location.

At this point, the air-gapped IBM Spectrum Scale cluster takes over the role of the production IBM Spectrum Scale cluster. When the recovery of the original production cluster completes, data can again be reversed and synchronized to the original production cluster, and applications can be pointed back to the original production IBM Spectrum Scale cluster.

This solution approach discusses the following phases for this Cyber Resiliency solution.

Initialization Phase

As shown in Figure 9, initial synchronization between the file systems are established across the IBM Spectrum Scale clusters at both the sites in this phase.

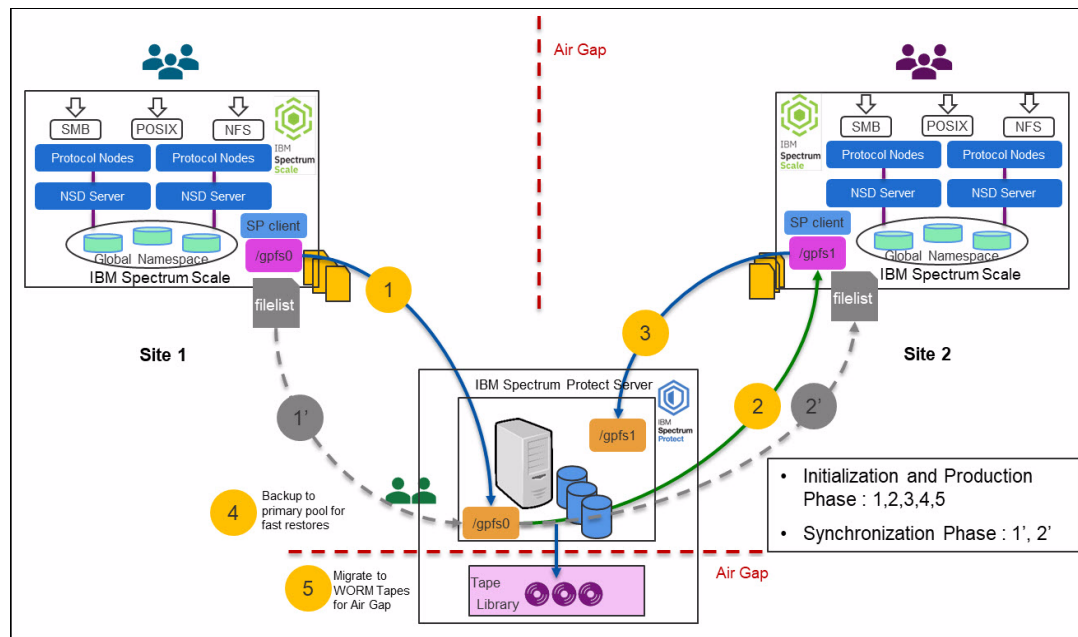


Figure 9 Initialization and production phase

To prepare the initialization phase, complete the following steps:

1. Create an initial full backup of the file system from the production IBM Spectrum Scale cluster. In order to have a consistent backup, create a snapshot of the production IBM Spectrum Scale file system, as shown in Figure 10.

```
[root@sss-ip02 ~]#
[root@sss-ip02 ~]# mmcrsnapshot gpfs0 backup_snap
Flushing dirty data for snapshot :backup_snap...
Quiescing all file system operations.
Snapshot :backup_snap created with id 9.
[root@sss-ip02 ~]#
```

Figure 10 Creating a snapshot of the production IBM Spectrum Scale file system

2. For the first time, take a full backup of the file system from the production IBM Spectrum Scale cluster. After the full backup, delete the temporary snapshot created on the production IBM Spectrum Scale cluster (step 1 shown in Figure 9).

Notes:

- a. Depending on the snapshot retention policy of your organization, snapshots of the production IBM Spectrum Scale file system can be kept for operational recovery. If kept, these snapshots can serve as an operational recovery mechanism from any infections, or from accidental deletion of data from the file system.
- b. All use of backup and restore commands, such as **mmbackup**, in this document is for illustration purposes only (Figure 11 on page 12). Actual use can vary in production environments.

```

[root@sss-ip02 ~]#
[root@sss-ip02 ~]# mmbbackup gpfs0 -t full -S backup_snap
-----
mmbbackup: Backup of /ibm/gpfs0 begins at Tue Aug 6 21:21:47 MST 2019.
-----
Tue Aug 6 21:21:52 2019 mmbbackup:Scanning file system gpfs0
mmbbackup: TSM Summary Information:
Total number of objects inspected:      6766
Total number of objects backed up:      6485
Total number of objects updated:        0
Total number of objects rebound:        0
Total number of objects deleted:        0
Total number of objects expired:        0
Total number of objects failed:        0
Total number of objects encrypted:      0
Total number of bytes inspected:        2544768122
Total number of bytes transferred:      2544768122
-----
mmbbackup: Backup of /ibm/gpfs0 completed successfully at Tue Aug 6 21:22:44 MST 2019.
-----
[root@sss-ip02 ~]#
[root@sss-ip02 ~]#
[root@sss-ip02 ~]# mmdelsnapshot gpfs0 backup_snap
Invalidating snapshot files in :backup_snap...
Deleting files in snapshot :backup_snap...
100.00 % complete on Tue Aug 6 21:23:35 2019 ( 655360 inodes with total 28 MB data processed)
Invalidating snapshot files in :backup_snap/F/...
Delete snapshot :backup_snap successful.
[root@sss-ip02 ~]#

```

Figure 11 Initial full backup of the production IBM Spectrum Scale file system

3. Initiate full synchronization of the production IBM Spectrum Scale file system (gpfs0) to the air-gapped IBM Spectrum Scale file system (gpfs1). Log in to the air-gapped IBM Spectrum Scale cluster and use a no-query-restore for this to enable massive parallel restore processing (Step 2 shown in Figure 9 on page 11).

After the full restore of the file system, delete the shadow DB generated by the **mmbbackup** utility on the production IBM Spectrum Scale system that was restored during a full restore process, as shown in Figure 12.

```

[root@sss-ip09 ~]#
[root@sss-ip09 ~]# dsmc restore /ibm/gpfs0/ /ibm/gpfs1/ -replace=all -subdir=yes -tapeprompt=no
IBM Spectrum Protect
Command Line Backup-Archive Client Interface
Client Version 8, Release 1, Level 7.0
Client date/time: 08/06/2019 21:29:58
(c) Copyright by IBM Corporation and other(s) 1990, 2019. All Rights Reserved.

Node Name: SSS-IP02
Session established with server SSS-IP06: Linux/x86_64
Server Version 8, Release 1, Level 7.000
Server date/time: 08/06/2019 21:30:04 Last access: 08/06/2019 21:27:17

Restore function invoked.

ANSI247I Waiting for files from the server...
Restoring          19 /ibm/gpfs0/worm_fset01/worm_test2 --> /ibm/gpfs1/worm_fset01/worm_test2 [Done]
Restoring      1,200,227 /ibm/gpfs0/.mmbbackupShadow.1.sss-ip06.filesys --> /ibm/gpfs1/.mmbbackupShadow.1.sss-ip06.filesys [Done]

Restore processing finished.

Total number of objects restored:      6,491
Total number of objects failed:        0
Total number of bytes transferred:     2.37 GB
Data transfer time:                    9.74 sec
Network data transfer rate:            254,743.14 KB/sec
Aggregate data transfer rate:          79,174.32 KB/sec
Elapsed processing time:                00:00:31
[root@sss-ip09 ~]#
[root@sss-ip09 ~]#
[root@sss-ip09 ~]# rm -f rm /ibm/gpfs1/.mmbbackupShadow.1.sss-ip06.filesys
[root@sss-ip09 ~]#

```

Figure 12 Initiate a full restore of backed up data to the IBM Spectrum Scale file system at the air-gap location

4. Initiate a full backup of the IBM Spectrum Scale file system at the air-gap location on the same IBM Spectrum Protect server. This acts as a baseline for the reverse production workflow when a failover occurs (step 3 shown in Figure 9 on page 11).

Follow the previous procedure to create a snapshot of the file system and then initiate a full backup of the IBM Spectrum Scale file system at the air-gap location using the **mmbbackup** utility, as shown in Figure 13.

```
[root@sss-ip09 ~]#
[root@sss-ip09 ~]# mmcrsnapshot gpfs1 backup_snap_airgap
Flushing dirty data for snapshot :backup_snap_airgap...
Quiescing all file system operations.
Snapshot :backup_snap_airgap created with id 7.
[root@sss-ip09 ~]#
[root@sss-ip09 ~]# mmbbackup gpfs1 -t full -S backup_snap_airgap
-----
mmbbackup: Backup of /ibm/gpfs1 begins at Tue Aug 6 21:38:58 MST 2019.
-----
Tue Aug 6 21:39:22 2019 mmbbackup:Scanning file system gpfs1
mmbbackup: TSM Summary Information:
      Total number of objects inspected:      7359
      Total number of objects backed up:      6493
      Total number of objects updated:         0
      Total number of objects rebound:        0
      Total number of objects deleted:         0
      Total number of objects expired:         0
      Total number of objects failed:          0
      Total number of objects encrypted:       0
      Total number of bytes inspected:      2555505541
      Total number of bytes transferred:     2544768122
-----
mmbbackup: Backup of /ibm/gpfs1 completed successfully at Tue Aug 6 21:40:00 MST 2019.
-----
[root@sss-ip09 ~]#
```

Figure 13 Initiate a full backup of the IBM Spectrum Scale file system at the air-gap location

Note: It is advisable to keep the snapshots at the IBM Spectrum Scale cluster in the air-gap location per your snapshot retention policy. These snapshots can serve as a recovery mechanism on an air-gapped IBM Spectrum Scale cluster to validate the authenticity of data at a remote location.

5. As explained previously, Cyber Resiliency at an IBM Spectrum Protect server can be further achieved by migrating the data to WORM capable tapes (Steps 4 and 5 from Figure 9 on page 11).

Backup Synchronization technique in the Production phase

As explained earlier in this document, this approach uses the **mmbbackup** utility from IBM Spectrum Scale to initiate backups to the IBM Spectrum Protect server. The **mmbbackup** utility provides scalable backup processing and creates file lists with backup candidates. In a normal scenario, these file lists are deleted from the file system by the **mmbbackup** utility after finishing the backup processing.

In this approach, the generated file lists are used to synchronize file systems across both locations. To keep these file lists, debug level 2 mode (DEBUGmmbbackup=0x002) must be enabled on the **mmbbackup** utility. After finishing the backup, these file lists will be backed up separately to the IBM Spectrum Protect server. On the air-gapped cluster, these file lists are restored first. The file lists generated in an **mmbbackup** utility format are then converted to the IBM Spectrum Protect recognized file lists format.

Then, incrementally backed up files from the production file system are restored on the air-gapped IBM Spectrum Scale file system using these converted file lists. This is also shown in Figure 9 on page 11.

Production phase

In this phase of the Cyber Resiliency approach, files on the production IBM Spectrum Scale cluster are backed up incrementally. IBM Spectrum Scale typically holds a large amount of data. Therefore, it is a best practice to back up IBM Spectrum Scale file systems incrementally at a regular interval. The frequency of the backup depends on the file system usage and the change rate of files.

In this phase, there can be two types of files that are under consideration.

Newly generated files or modified files

The following steps illustrate the synchronization of the files that are newly generated or modified on the IBM Spectrum Scale cluster at the production site:

1. To back up incrementally, create a snapshot of the production file system and enable the **mmbbackup** utility debug parameter so that the file lists are kept in the IBM Spectrum Scale file system, as shown in Figure 14.

```
[root@sss-ip02 ~]#
[root@sss-ip02 ~]# mmcrsnapshot gpfs0 backup_snap
Flushing dirty data for snapshot :backup_snap...
Quiescing all file system operations.
Snapshot :backup_snap created with id 10.
[root@sss-ip02 ~]# export DEBUGmmbbackup=0x002
[root@sss-ip02 ~]#
```

Figure 14 Creating snapshot and enabling mmbbackup debug parameter for incremental backup

2. Initiate incremental backup of the production IBM Spectrum Scale file system to the IBM Spectrum Protect server. Then, delete the snapshot if it is not required under the snapshot retention policy, as shown in Figure 15 (also step 1 shown in Figure 9 on page 11).

```
[root@sss-ip02 ~]#
[root@sss-ip02 ~]# mmbbackup gpfs0 -t incremental -S backup_snap

-----
mmbbackup: Backup of /ibm/gpfs0 begins at Tue Aug 6 23:32:14 MST 2019.
-----
Tue Aug 6 23:32:19 2019 mmbbackup:Scanning file system gpfs0
mmbbackup: TSM Summary Information:
      Total number of objects inspected:      2471
      Total number of objects backed up:      2144
      Total number of objects updated:        0
      Total number of objects rebound:        0
      Total number of objects deleted:         0
      Total number of objects expired:         0
      Total number of objects failed:          0
      Total number of objects encrypted:       0
      Total number of bytes inspected:      850132992
      Total number of bytes transferred:     846924349
-----
mmbbackup: Backup of /ibm/gpfs0 completed successfully at Tue Aug 6 23:32:35 MST 2019.
-----
[root@sss-ip02 ~]#
[root@sss-ip02 ~]# mmdelsnapshot gpfs0 backup_snap
Invalidating snapshot files in :backup_snap...
Deleting files in snapshot :backup_snap...
 100.00 % complete on Tue Aug 6 23:33:03 2019 ( 655360 inodes with total      28 MB data processed)
Invalidating snapshot files in :backup_snap/F/...
Delete snapshot :backup_snap successful.
[root@sss-ip02 ~]#
```

Figure 15 Incremental backup from the IBM Spectrum Scale cluster at the production site

3. Selectively backup the **mmbbackup** utility generated file list and then delete the file list locally from the production IBM Spectrum Scale file system as shown in Figure 16 and Step 1 in Figure 9 on page 11.

```
[root@sss-ip02 ~]#
[root@sss-ip02 ~]# dsmc sel /ibm/gpfs0/.mmbbackupCfg/updatedFiles/.list.1.sss-ip06
IBM Spectrum Protect
Command Line Backup-Archive Client Interface
  Client Version 8, Release 1, Level 7.0
  Client date/time: 08/06/2019 23:35:07
(c) Copyright by IBM Corporation and other(s) 1990, 2019. All Rights Reserved.

Node Name: SSS-IP02
Session established with server SSS-IP06: Linux/x86_64
  Server Version 8, Release 1, Level 7.000
  Server date/time: 08/06/2019 23:35:13 Last access: 08/06/2019 23:32:39

Selective Backup function invoked.

Normal File-->          453,246 /ibm/gpfs0/.mmbbackupCfg/updatedFiles/.list.1.sss-ip06 [Sent]
Selective Backup processing of '/ibm/gpfs0/.mmbbackupCfg/updatedFiles/.list.1.sss-ip06' finished without failure.

Total number of objects inspected:          1
Total number of objects backed up:          1
Total number of objects updated:            0
Total number of objects rebound:            0
Total number of objects deleted:            0
Total number of objects expired:            0
Total number of objects failed:            0
Total number of objects encrypted:          0
Total number of objects grew:              0
Total number of retries:                   0
Total number of bytes inspected:           442.62 KB
Total number of bytes transferred:         442.80 KB
Data transfer time:                        0.06 sec
Network data transfer rate:                7,055.12 KB/sec
Aggregate data transfer rate:              420.90 KB/sec
Objects compressed by:                     0%
Total data reduction ratio:                 0.00%
Elapsed processing time:                    00:00:01
[root@sss-ip02 ~]#
[root@sss-ip02 ~]# rm -f /ibm/gpfs0/.mmbbackupCfg/updatedFiles/.list.1.sss-ip06
[root@sss-ip02 ~]#
```

Figure 16 Selective backup of a file list generated by mmbbackup utility

4. Restore the file list generated by the **mmbbackup** utility on the IBM Spectrum Scale file system at the air-gap location (Figure 17 and step 2 in Figure 9 on page 11).

```
[root@sss-ip09 ~]#
[root@sss-ip09 ~]# dsmc restore /ibm/gpfs0/.mmbbackupCfg/updatedFiles/.list.1.sss-ip06 /ibm/gpfs1/incre_restore_list -asnodename=sss-ip02
IBM Spectrum Protect
Command Line Backup-Archive Client Interface
  Client Version 8, Release 1, Level 7.0
  Client date/time: 08/07/2019 00:23:01
(c) Copyright by IBM Corporation and other(s) 1990, 2019. All Rights Reserved.

Node Name: SSS-IP09
Session established with server SSS-IP06: Linux/x86_64
  Server Version 8, Release 1, Level 7.000
  Server date/time: 08/07/2019 00:23:07 Last access: 08/07/2019 00:18:46

Accessing as node: SSS-IP02
Restore function invoked.

Restoring          453,246 /ibm/gpfs0/.mmbbackupCfg/updatedFiles/.list.1.sss-ip06 --> /ibm/gpfs1/incre_restore_list [Done]

Restore processing finished.

Total number of objects restored:          1
Total number of objects failed:            0
Total number of bytes transferred:         442.69 KB
Data transfer time:                        0.00 sec
Network data transfer rate:                232,022.13 KB/sec
Aggregate data transfer rate:              142.54 KB/sec
Elapsed processing time:                    00:00:03
[root@sss-ip09 ~]#
```

Figure 17 Restore the file list generated by the mmbbackup utility at the air-gap location

5. Convert the file list generated from the **mmbbackup** utility format to the IBM Spectrum Protect recognized file list format for the restore. Figure 18 shows the example of converting the **mmbbackup** utility format to the IBM Spectrum Protect recognized file list format. (Figure 18 is for illustration purposes only.)

```
[root@sss-ip09 ~]#  
[root@sss-ip09 ~]# awk -F: '{print $9}' incre_restore_list | grep -v "^$" | sed "s/[0-9]*\\!\\/ibm\\gpfs0\\/g" >> incre_restore_list  
[root@sss-ip09 ~]#
```

Figure 18 Example of converting a file list generated from the mmbbackup utility format

6. After restoring the **mmbbackup** utility generated file list from the IBM Spectrum Protect server, expire the file list from the IBM Spectrum Protect server. This file list is backed up again after the next incremental backup, as shown in Figure 19.

```
[root@sss-ip09 ~]#  
[root@sss-ip02 ~]# dsmc expire /ibm/gpfs0/.mmbbackupCfg/updatedFiles/.list.1.sss-ip06  
[root@sss-ip09 ~]#
```

Figure 19 Expire file list from the IBM Spectrum Protect server

7. Restore the files that are backed up incrementally from the production IBM Spectrum Scale file system to the IBM Spectrum Scale file system at the air-gap location using the file list generated in the previous step (Figure 20 and step 2 from Figure 9 on page 11).

```
[root@sss-ip09 ~]# dsmc restore -filelist=incre_restore_list2 /ibm/gpfs1/ -replace=all -preser=complete -tapeprompt=no -asnodename=sss-ip02  
IBM Spectrum Protect  
Command Line Backup-Archive Client Interface  
Client Version 8, Release 1, Level 7.0  
Client date/time: 08/07/2019 02:29:10  
(c) Copyright by IBM Corporation and other(s) 1990, 2019. All Rights Reserved.  
  
Node Name: SSS-IP09  
Session established with server SSS-IP06: Linux/x86_64  
Server Version 8, Release 1, Level 7.000  
Server date/time: 08/07/2019 02:29:16 Last access: 08/07/2019 02:27:04  
  
Accessing as node: SSS-IP02  
Restore function invoked.  
...  
ANS1898I ***** Processed 2,000 files *****  
  
Total number of objects restored: 2,144  
Total number of objects failed: 0  
Total number of bytes transferred: 807.47 MB  
Data transfer time: 3.87 sec  
Network data transfer rate: 213,250.78 KB/sec  
Aggregate data transfer rate: 13,708.73 KB/sec  
Elapsed processing time: 00:01:00  
[root@sss-ip09 ~]#
```

Figure 20 Incremental restore at the air-gap location

8. In order to keep backups synchronized on the IBM Spectrum Protect server, perform an incremental backup from the IBM Spectrum Scale file system at the air-gap location to the IBM Spectrum Protect server (Figure 21 and step 3 from Figure 9 on page 11).

```
[root@sss-ip02 ~]#  
[root@sss-ip09 ~]# mmcrsnapshot gpfs1 backup_snap_airgap1  
[root@sss-ip09 ~]# mmbbackup gpfs1 -t incremental -S backup_snap_airgap1  
[root@sss-ip02 ~]#
```

Figure 21 Incremental backups from the IBM Spectrum Scale file system at the air-gap location

9. As explained previously, IBM Spectrum Protect server-level Cyber Resiliency can be further achieved by migrating the data to WORM capable tapes (steps 4 and 5 from Figure 9 on page 11).

Deleted files

This is an optional approach to recover deleted files on the production IBM Spectrum Scale cluster, or keep the IBM Spectrum Scale clusters synchronized when the files are deleted purposefully from the cluster. In some scenarios, files are deleted from the production IBM Spectrum Scale file system accidentally, or they are deleted by some virus attacks. In such scenarios, when the next incremental backup is issued against a file system, deleted files are also expired from the IBM Spectrum Protect server.

If the appropriate retention policy is configured for the objects backed up from the IBM Spectrum Scale clusters, those accidentally deleted files can be restored from the IBM Spectrum Protect server. Also, such files could be recovered from the IBM Spectrum Scale file system at the air-gap location using the procedure explained in the previous section.

In some cases, files on the IBM Spectrum Scale clusters are deleted purposefully for releasing the space from the file system on the production site. In such cases, files on the IBM Spectrum Scale cluster can be retained for quick recovery purposes. If IBM Spectrum Scale file systems need to be kept synchronized at all times, those files must also be deleted from the IBM Spectrum Scale cluster at the air-gap location.

Consider the following procedure for keeping the files synchronized on the IBM Spectrum Scale cluster at the air-gap location:

1. When an incremental backup is issued on the IBM Spectrum Scale cluster in debug mode, the **mmbackup** utility also creates a file list containing a list of files marked for expiry. Back up this file list containing a list of the objects marked for expiry using a selective backup to the IBM Spectrum Protect server. (In this example, the file list for expired files is `/ibm/gpfs0/.mmbackupCfg/expiredFiles/.list.1.sss-ip06`).
2. Restore this **mmbackup** utility-generated file list of expired objects on the IBM Spectrum Scale cluster at the air-gap location.
3. Convert this **mmbackup** utility-generated file list into the IBM Spectrum Protect file list format.
4. Then remove the files mentioned in the converted file list from the IBM Spectrum Scale air-gap location.
5. Whenever the next incremental backup is run from the IBM Spectrum Scale cluster at the air-gap location, the deleted files are also expired from the IBM Spectrum Protect server file space of the IBM Spectrum Scale cluster at the air-gap location.

Detection of Anomaly phase

IBM Spectrum Protect detects anomalies in workload patterns to alert administrators of potential ransomware infections for workloads. This detection can depend on significant increased backup rates compared to the historical values, or significant decreased deduplication and compression rates compared to the historical values.

Anomaly detection enables administrators to be aware of possible attacks and mitigate them before they spread to other locations. The alert system can be used to notify the administrators that something malicious happened on the production IBM Spectrum Scale cluster. This system can also stop the synchronization processing to the IBM Spectrum Scale cluster at the air-gap location to prevent corrupted data from being synchronized to the air-gap location. For more information, see [IBM Knowledge Base](#) for IBM Spectrum Protect.

Failover and Failback phase

If the IBM Spectrum Scale cluster at the production site is compromised, applications are failed over to the IBM Spectrum Scale cluster at the air-gap location, and it assumes the role of a new production cluster. At this time, the original production cluster that was compromised needs to be recovered.

Depending on the organization policies, there are various methods of recovering the compromised data on the original production IBM Spectrum Scale cluster:

- If only live data on the production IBM Spectrum Scale cluster was compromised (for example, with ransomware) and immediately recognized, and if an organization has a policy of maintaining snapshots of the production site, that data can quickly be restored using local snapshots on the production Spectrum Scale cluster.
- If non-compromised data copies are not located on the local snapshot copies, that data can be restored from the IBM Spectrum Protect server. Older copies of data can be restored from the IBM Spectrum Protect server using data backed up on WORM tapes.
- If compromised data is not yet synchronized with the IBM Spectrum Scale cluster at the air-gap location, it can be restored back from the IBM Spectrum Scale cluster at the air-gap location.
- If compromised data is synchronized with the IBM Spectrum Scale cluster at the air-gap location, data can be restored at the air-gap location first from the snapshots, and then restored back from the IBM Spectrum Scale cluster at the air-gap location.
- If the entire IBM Spectrum Scale cluster at the production site is compromised by known events, unknown events, or human intervention, applications are failed over to the IBM Spectrum Scale cluster at the air-gap location. This cluster assumes the role of the production site.

After the original production IBM Spectrum Scale cluster is restored, then it needs to be synchronized with the IBM Spectrum Scale cluster at the air-gap location. Finally, appropriate applications are failed back to the original production IBM Spectrum Scale cluster.

Figure 22 on page 19 shows the failback phase of data movement and synchronization from the IBM Spectrum Scale file system at the air-gap location to the original production IBM Spectrum Scale file system.

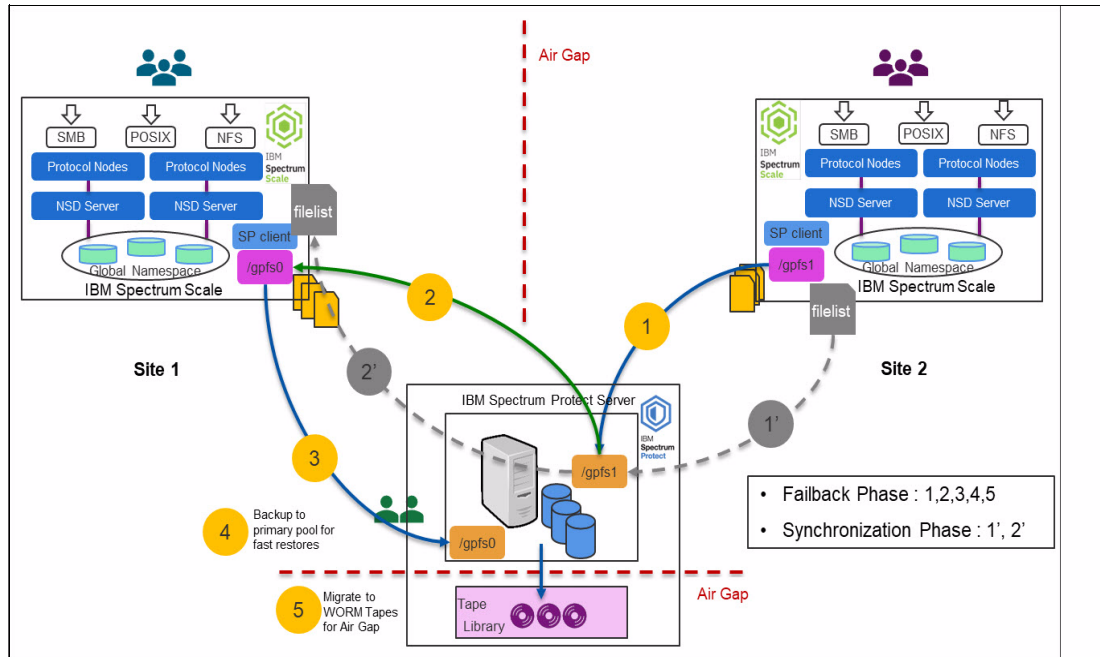


Figure 22 Failback phase

This phase uses the same **mmbackup** utility-generated file lists, and synchronizes the data on the original production IBM Spectrum Scale file system using the file lists:

1. In order to fail back to the original production IBM Spectrum Scale cluster, take a snapshot of the IBM Spectrum Scale file system at the air-gap location and export the **mmbackup** utility debug parameter. Next, take an incremental backup of the IBM Spectrum Scale file system at the air-gap location (Figure 23), and back up the **mmbackup** utility-generated file list to the IBM Spectrum Protect server (Step 1 from Figure 22).

```
[root@sss-ip09 ~]#
[root@sss-ip09 ~]# mmcrsnapshot gpfs1 backup_snap_airgap1
[root@sss-ip09 ~]#
[root@sss-ip09 ~]# export DEBUGmmbackup=0x002
[root@sss-ip09 ~]#
[root@sss-ip09 ~]# mmbackup gpfs1 -t incremental -S backup_snap_airgap1
[root@sss-ip09 ~]#
[root@sss-ip09 ~]# dsmsel /ibm/gpfs1/.mmbackupCfg/updatedFiles/.list.1.sss-ip06
[root@sss-ip09 ~]#
[root@sss-ip09 ~]# rm -f /ibm/gpfs1/.mmbackupCfg/updatedFiles/.list.1.sss-ip06
[root@sss-ip09 ~]#
```

Figure 23 Incremental backup from the IBM Spectrum Scale file system at the air-gap location

2. Depending on the data recovery required on the original IBM Spectrum Scale cluster, either full or incremental recovery of the IBM Spectrum Scale file system needs to be performed.

If a full restore of an IBM Spectrum Scale file system is needed, perform a full restore of the file system from IBM Spectrum Protect. Next, use a no-query-restore for this system to enable massive parallel restore processing, as shown in Figure 24.

```
[root@sss-ip02 ~]#
[root@sss-ip02 ~]# dsms restore /ibm/gpfs1/ /ibm/gpfs0/ -replace=all -subdir=yes -tapeprompt=no
[root@sss-ip02 ~]#
```

Figure 24 A full restore of IBM Spectrum Scale file system at the production site

3. In the case where only incremental data needs to be restored from the **mmbackup** utility-generated file lists, first restore the file lists from the IBM Spectrum Protect server on the original production cluster (Step 2 from Figure 22 on page 19).

Convert the file lists into an IBM Spectrum Protect recognized format and then restore the incrementally backed up data from the IBM Spectrum Scale file system at the air-gap location to the IBM Spectrum Scale file system at the original production site (Figure 25 and step 2 from Figure 22 on page 19).

```
[root@sss-ip02 ~]#  
[root@sss-ip02 ~]# dsmc restore /ibm/gpfs1/.mmbackupCfg/updatedFiles/.list.1.sss-ip06 airgap_restore_list -asnodename=sss-ip09  
[root@sss-ip02 ~]#  
[root@sss-ip02 ~]# dsmc restore -filelist=airgap_restore_list2 /ibm/gpfs0/ -replace=all -preser=complete -tapeprompt=no -asnodename  
[root@sss-ip02 ~]#
```

Figure 25 Incremental restore at the original production site

4. To keep backups synchronized on the IBM Spectrum Protect server, perform an incremental backup from the production IBM Spectrum Scale file system to the IBM Spectrum Protect server (step 3 from Figure 22 on page 19.)

After the data on both the IBM Spectrum Scale clusters is synchronized, applications can be failed back to the IBM Spectrum Scale cluster at the original production site.

Summary

Cyberattacks are likely to remain a significant risk for the foreseeable future. Attacks on organizations can be external and internal. Investing in technology and processes to prevent these cyberattacks is the highest priority for these organizations. Organizations need well-designed procedures and processes to recover from attacks.

The NIST framework provides standards, guidelines, and best practices to manage cybersecurity related risks. Adoption of the NIST framework, the proper discipline of risk management, and IBM Storage offerings can be used to create and implement recovery plans that ensure the safety of business-critical data.

IBM Spectrum Scale is a software-defined storage system for high performance, large-scale workloads, on-premises or in the cloud. IBM Spectrum Protect is a data protection platform that gives enterprises a single point of control and administration for backup and recovery. IBM Tape Storage offers cost-effective, long-term backup and archive WORM storage, with a true physical air gap and total separation from ransomware and cyberattacks.

The IBM Spectrum Scale existing feature integration with IBM Spectrum Protect and IBM Tape Storage enables organizations to implement an effective, easy-to-manage, and automated Cyber Resiliency solution. This solution provides robust protection against external cyber events, such as malware, ransomware attacks, and human elements.

The configurations and examples of commands described in this document are for illustration purposes only. The real-world implementation, configuration of IBM Spectrum Scale clusters, and execution of commands might vary (per the number of IBM Spectrum Scale nodes) to achieve better scalability and performance.

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 <http://www.ibm.com/legal/copytrade.shtml>

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

Redbooks (logo) ®
IBM®

IBM Spectrum™
IBM Spectrum Protect™

IBM Spectrum Scale™

The following terms are trademarks of other companies:

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>.



© Copyright IBM Corporation

September 2019

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



Please recycle

ISBN 0738457965

REDP-5559-00