Data Protection for IBM Spectrum Scale

Contrasting data protection methods

Nils Haustein

Version 1.0

EMEA Storage Competence Center

Agenda

- Data protection overview

Data protection techniques

Contrasting

Data Protection overview

- Goal is to protect data against failures that cause data access lost

- Failures can be classified in:
  - Operational failures (someone deleted a file or directory)
  - Component failures in software and hardware (disk, server, network, etc.)
  - System outages (complete systems or site)

Key metrics for Data Protection

- Key metrics to differentiate data protection methods
  - **Recovery Point Objective**: how much data can be lost in case of failures
  - **Recovery Time Objective**: how long does it take to re-establish full data access

Criteria	Operational Recovery	High Availability	Disaster protection
Goal	Quickly recover access to data that has been lost or disappeared	Recover component failures quickly and automatically	Recover system outages with well defined and tested processes
RPO	Close to 0 h	Close to 0 h	1 – 24 h
RTO	0 – 1 h	Close to 0 h	4 – 24+ h
Typical techniques	Snapshots, Backup and restore	RAID protection, synchronous replication, server clustering, redundant networks	Secondary site, replication, backup, well defined and tested recovery process

More Reading ...

- New revision of our book “Speichernetze” explains business continuity and all relevant techniques in the context of systems and solutions

Agenda

Data protection overview

- Data protection techniques

Contrasting

IBM Spectrum Scale - Snapshots

- Snapshots freeze the content of the file system or file set
  - Snapshot reflects the file and directory structure
  - Snapshots are space-efficient
    - Space is allocated when files are changed or deleted
    - Using redirect on write
  - Snapshots are read-only
- Characteristic
  - For operational recovery: a user lost a file
  - Limited disaster recovery because snapshots require the file system and cluster

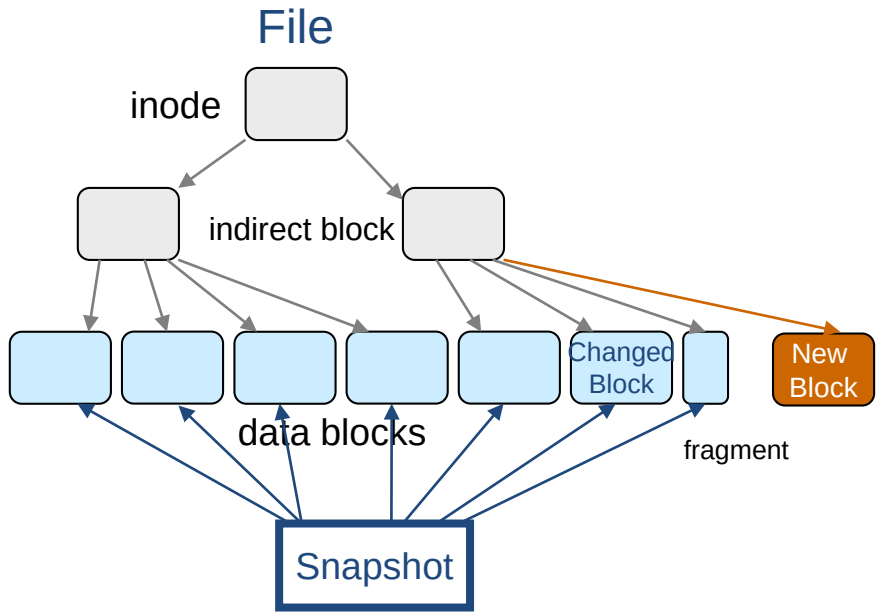

RPO: ~1 – 4 h

RTO: minutes for files, days for file system

IBM Spectrum Scale - Synchronous replication

- IBM Spectrum Scale can write data blocks of files synchronously to different storage system
  - Storage systems must belong to the same cluster
  - Up to three copies are possible
  - Based on failure groups
  - Allows active – active operations
  - If one DC has outage then operations seamlessly continue in other DC

- Characteristics:
  - High availability solution
    - Can seamlessly recovery data center outage
    - Does not require any restoration of data (if cluster and file system survives)
  - Limited disaster recovery because the if the cluster or file system is down no data can be accessed

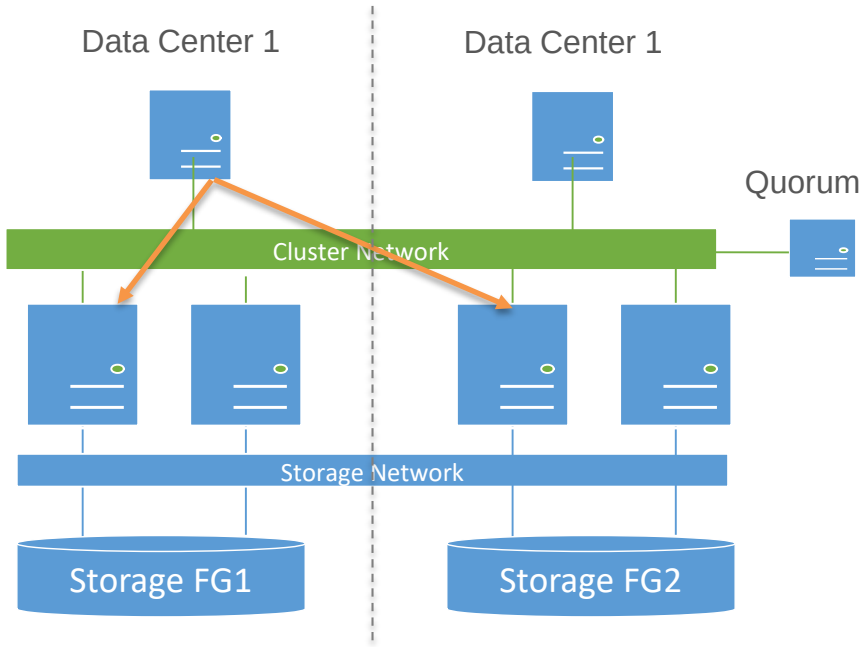

Stretched Spectrum Scale Cluster

RPO: ~0 h

RTO: ~0 h

IBM Spectrum Scale - Asynchronous replication

- Active File Management asynchronously replicates data from primary to DR cluster
  - Peer Snapshots to can be used to provide consistent recovery point
  - Replicates changed sub-blocks
  - If primary site fails then failover to DR site
    - One command to failover from primary to DR site
  - Provides procedure for failback to old primary

- Characteristics:
  - Disaster recovery upon complete cluster or file system failure
  - Simple and quick failover procedure
    - Does not require any restoration of data
  - Limited High Availability because failover is not automated

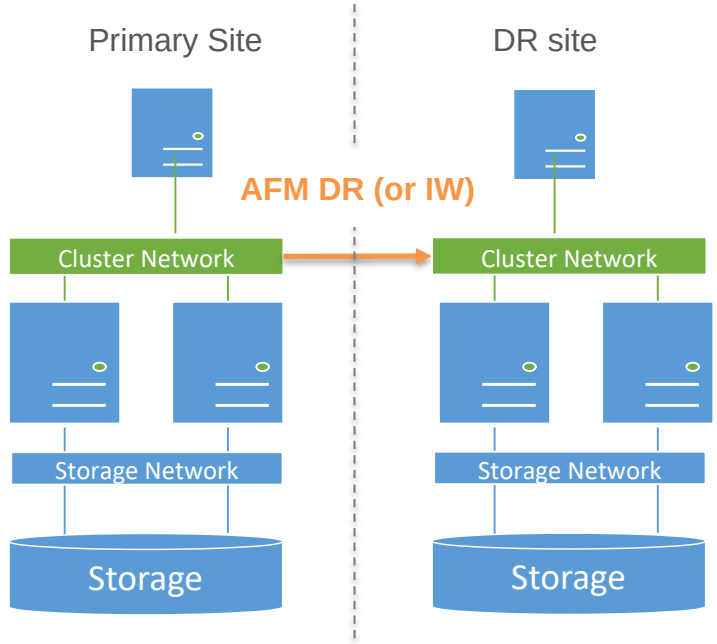

Separated IBM Spectrum Scale clusters

RPO: ~15 – 60 min

RTO: ~ 1 h

IBM Spectrum Scale - Replication with Aspera

- IBM Aspera Sync can be used to continuously replicate files from one IBM Spectrum Scale cluster to another
  - Can be integrated with policy engine to identify candidates
  - Aspera Sync is specially suitable for long distance file transfer
  - Aspera Sync transfers whole files (no sub-blocks)
    - Transfers all file attributes including ACL
  - If primary site fails then bring DR site online
  - Failback requires to identify the delta and copy it back

- Characteristics:
  - Disaster recovery upon complete cluster or file system failure
  - Simple and quick failover
    - Does not require any restoration of data
  - Limited High Availability because failover is not automated

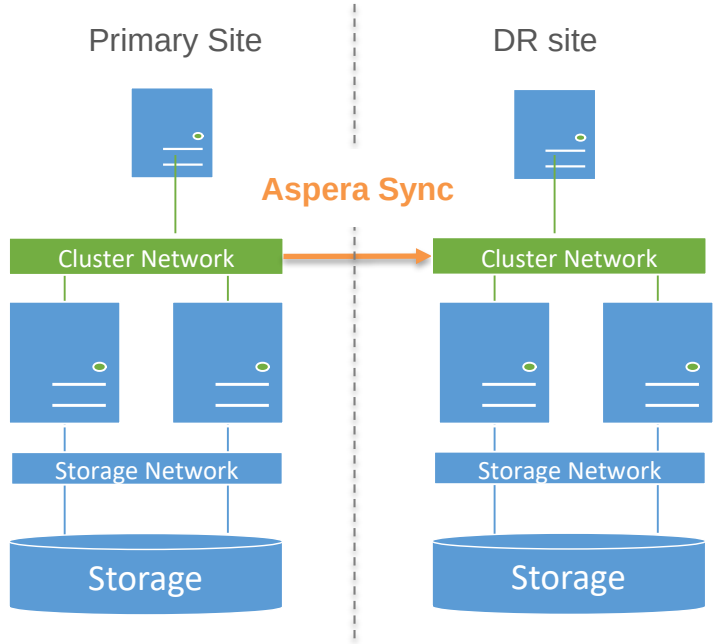

Separated IBM Spectrum Scale clusters

RPO: ~15 – 60 min

RTO: ~ 1 h

IBM Spectrum Scale - Backup (mmbackup)

- IBM Spectrum Scale Backup integrates with IBM Spectrum Protect server
  - Utility `mmbackup` uses policy engine to identify candidates for backup and expiration
  - Invokes Spectrum Protect backup client to perform the backup operation
    - Scalable approach because multiple nodes can participate in backup operation
  - Restore files using SP B/A client (`dsmc restore`)

- Characteristics:
  - Operational recovery if files have been corrupted or deleted
    - Leverage versioning in backup
  - Limited disaster recovery upon complete cluster or file system failure
    - Requires to restore all files
  - No High Availability because data has to be restored manually

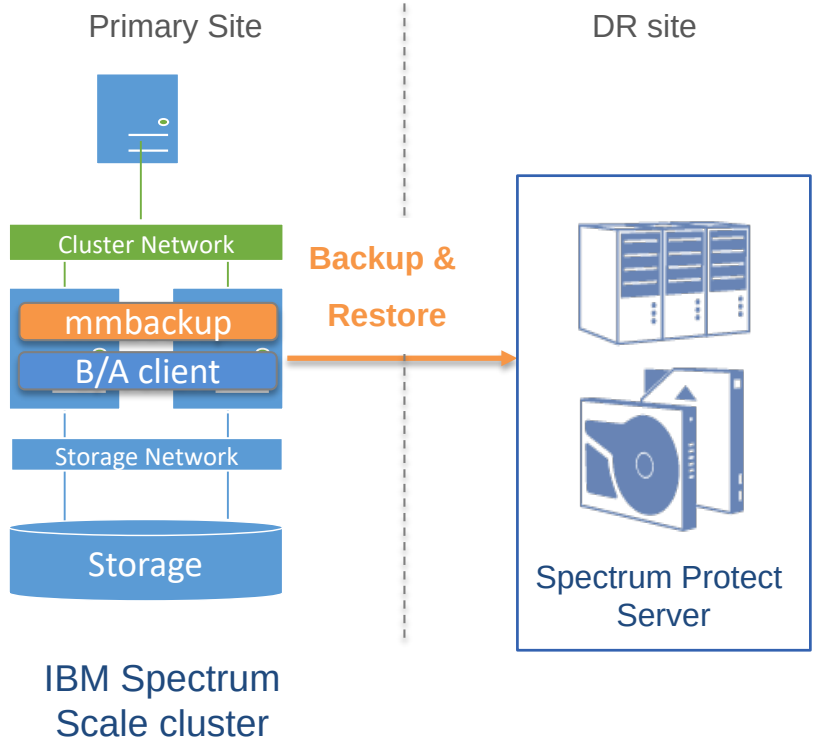

RPO: ~4 – 24 h

RTO: minutes for files, days for file systems

IBM Spectrum Scale – Transparent Cloud Tiering

- IBM Spectrum Scale Transparent Cloud Tiering provides backup and restore function
  - TCT (pre)migrates files to object storage
    - Can use policy engine to identify file candidates
    - TCT can utilize versioning provided by object storage
  - Restore files using TCT utility (`mmcloudgateway restore`)
    - Can restore different versions
    - Does not restore file EA and ACL.
    - Shall only be used for emergency restores
- Characteristics:
  - Operational recovery if files have been corrupted or deleted
    - Leverage versioning in object storage
    - ACL and EA are not restored
  - Limited disaster recovery upon complete cluster or file system failure
    - Requires to restore all files
  - No High Availability because data has to be restored manually

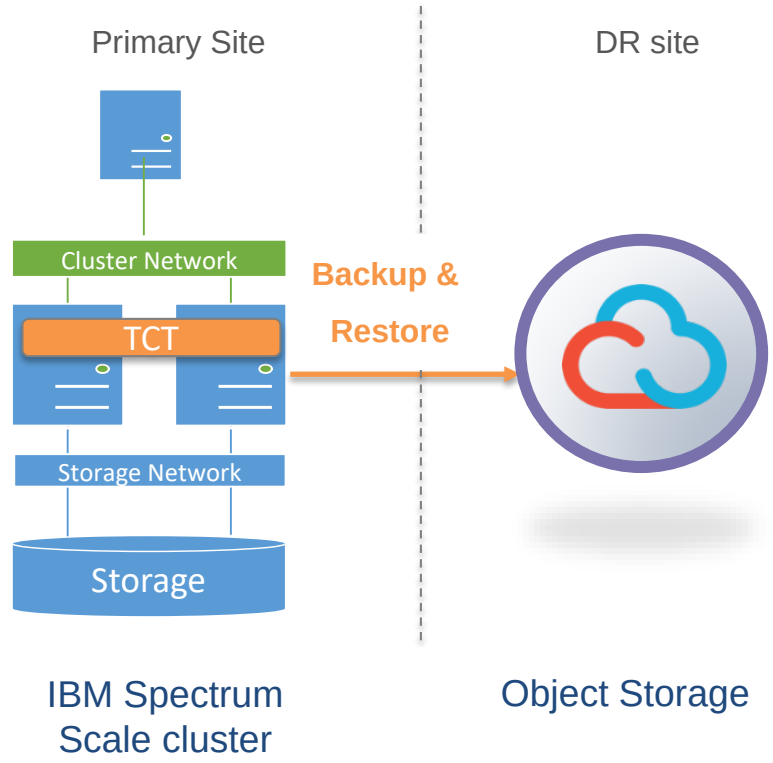

RPO: ~4 – 24 h

RTO: minutes for files, days for file systems

IBM Spectrum Scale – Scale out Backup and Restore

- IBM Spectrum Scale integrates with IBM Spectrum Protect for Space Management to provide Scale out Backup and Restore (Sobar)
  - HSM client continuously (pre)migrates files to SP server
  - Spectrum Scale Sobar periodically backs up all inodes
  - If file system fails then inode-backup can be restored
    - No file data is transferred, only inodes are applied in file system
    - All (pre)migrated files have reference to SP server and can be recalled on access or in bulks
- Characteristics:
  - Disaster recovery upon complete cluster or file system failure
    - Only file and directory inodes are restored
    - File data is restored on-demand using recalls
  - No High Availability because data has to be restored manually

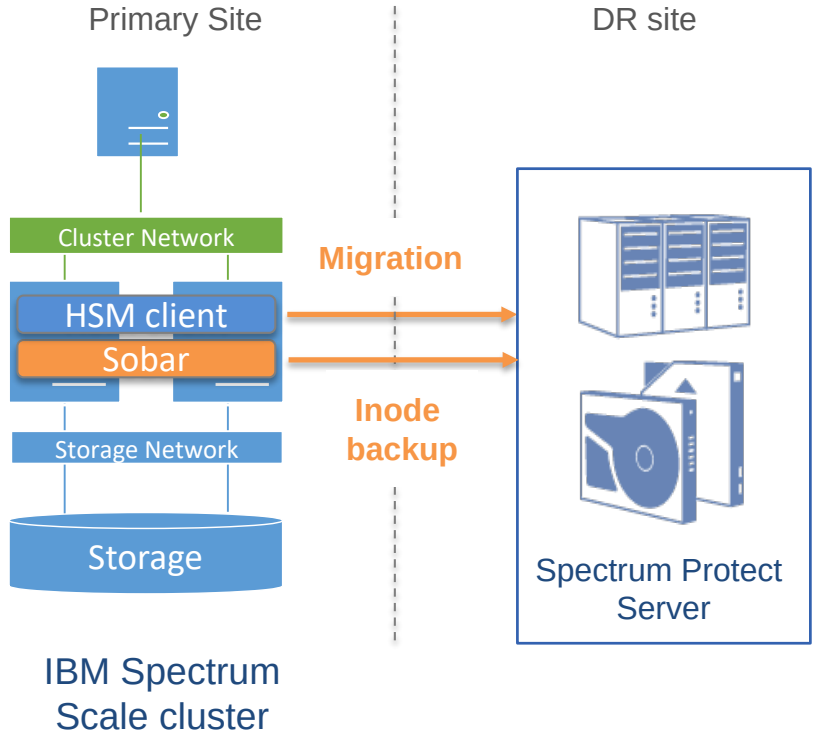

RPO: ~8 – 24 h

RTO: some hours for initial access

Agenda

Data protection overview

Data protection techniques

- Contrasting

Contrasting

Criteria	Sync. Replication	Async. Replication	Aspera Sync	Sobar	Snapshots	mmbackup	TCT
RTO	~0 h	~1 h	~1 h	Some hours for initial access	Minutes for files Days for file system	Minutes for files Days for file system	Minutes for files Days for file system
RPO	~0 h	~15 – 60 min	~15 – 60 min	~8 – 24 h	~1 – 4 h	~4 – 24 h	~4 – 24 h
Multi-site	Yes	Yes	Yes	Yes	No	Yes	Yes
HA	Yes	Limited	Limited	No	No	No	No
DR	Limited	Yes	Yes	Yes	No	Limited	Limited
Restore	N/A	No	No	No	Yes	Yes	Yes

Typical HA

Typical DR

Typical operational recovery

Good to know

Sync. Replication	Async. Replication	Aspera Sync	Sobar	Snapshots	mmbackup	TCT
Requires high bandwidth and low latency network between sites	• Not well suited for many random file writes on primary site (e.g. databases) • Does not work with immutable files	Requires additional scripting to automate and monitor operations	• Requires Spectrum Protect for Space Management or TCT • Requires additional scripting	Does not work with HSM and TCT	• Requires IBM Spectrum Protect server • Can integration with Spectrum Archive	Does not restore ACLs

Summary

- Snapshots and backup for operational recovery
  - Backup allows restore if the file system is broken
  - Be aware that snapshots do not integrate with HSM
- Stretched cluster with synchronous replication for HA
  - Can be combined with snapshots, backup and AFM DR
- AFM DR or Aspera Sync for DR
  - Aspera Sync integrates with immutable filesets
  - AFM DR transfers changed blocks

Questions

Restore vs. Recovery

- Restore and Recovery are used when a copy of the primary data exists
  - For example a copy created by backup or snapshots

- **Restore** copies individual files or directories back to primary system
  - For example restore files that have been deleted or corrupted
  - Versioning helps to address data corruption (malware)
  - **Operational recovery**: restore individual files and directories quickly

- **Recovery** re-establishes the primary files system and copies back all data
  - For example, recover a file system that has been lost
  - Recovery time is key, typically the latest version is recovered
  - **Disaster recovery:** restore the entire system after a crash

Backup vs. Replication

- **Replication** replicates files from primary system to remote system
  - Files that are deleted on primary system are deleted in remote system
  - No versioning, only the last version of the file is available

- **Backup** copies files from primary system to remote system
  - File that are deleted on primary system are not (immediately) deleted in remote system
  - Supports versioning: earlier versions of the file are available
  - Requires restoration to grant access to files

Disclaimer

This information is for **IBM Spectrum Scale user days 2020**, publication beyond this scope is forbidden

This information is provided on an "AS IS" basis without warranty of any kind, express or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Some jurisdictions do not allow disclaimers of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

This information is provided for information purposes only as a high level overview of possible future products. PRODUCT SPECIFICATIONS, ANNOUNCE DATES, AND OTHER INOFORMATION CONTAINED HEREIN ARE SUBJECT TO CHANGE AND WITHDRAWAL WITHOUT NOTICE.

USE OF THIS DOCUMENT IS LIMITED TO SELECT IBM PERSONNEL THIS DOCUMENT SHOULD NOT BE GIVEN TO A CUSTOMER EITHER IN HARDCOPY OR ELECTRONIC FORMAT.

Important notes:

IBM reserves the right to change product specifications and offerings at any time without notice. This publication could include technical inaccuracies or typographical errors. References herein to IBM products and services do not imply that IBM intends to make them available in all countries.

IBM makes no warranties, express or implied, regarding non-IBM products and services, including but not limited to Year 2000 readiness and any implied warranties of merchantability and fitness for a particular purpose. IBM makes no representations or warranties with respect to non-IBM products. Warranty, service and support for non-IBM products is provided directly to you by the third party, not IBM.

All part numbers referenced in this publication are product part numbers and not service part numbers. Other part numbers in addition to those listed in this document may be required to support a specific device or function.

MHz / GHz only measures microprocessor internal clock speed; many factors may affect application performance. When referring to storage capacity, GB stands for one billion bytes; accessible capacity may be less. Maximum internal hard disk drive capacities assume the replacement of any standard hard disk drives and the population of all hard disk drive bays with the largest currently supported drives available from IBM.

IBM Information and Trademarks

The following terms are trademarks or registered trademarks of the IBM Corporation in the United States or other countries or both: the e-business logo, IBM, IBM Spectrum Scale, IBM Spectrum Archive, IBM Spectrum Protect

The SNIA logo is a registered trademark of Storage Networking and Industry Association.

The ISO logo is registered trademark of International Organization for Standardization

Microsoft Windows is a trademark or registered trademark of Microsoft Corporation.

Linux is a registered trademark of Linus Torvalds.

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

IBM

ESCC
EMEA Storage Competence Center

https://escc.w3ibm.mybluemix.net/#