Advanced Technology Group

Accelerate with ATG:
Data resiliency with IBM Storage Scale

Lindsay Todd, PhD
Principal Storage Technical Specialist – Storage Scale/SSS for Data/AI/HPC
IBM Advanced Technology Group – Storage Team

Accelerate with ATG Technical Webinar Series

Advanced Technology Group experts cover a variety of technical topics.

Audience: Clients who have or are considering acquiring IBM Storage solutions. Business Partners and IBMers are also welcome.

To automatically receive announcements of upcoming Accelerate with IBM Storage webinars, Clients, Business Partners and IBMers are welcome to send an email request to accelerate-join@hursley.ibm.com.

2023 Upcoming Webinars – click on the link to register for the live event:

August 22 – Introduction to IBM's newest Tape Storage, the IBM Diamondback Tape Library

August 29 – IBM Storage Virtualize 8.6 and Storage Sentinel Technical Update

Important Links to bookmark:

ATG Accelerate Support Site: https://www.ibm.com/support/pages/node/1125513

ATG MediaCenter Channel: https://ibm.biz/BdfEgQ

ATG-Storage Offerings

CLIENT WORKSHOPS

- **IBM DS8900F Advanced Functions – August 29-30, 2023, Virtual**
- IBM Storage Point of View on Cyber Resiliency
- IBM FlashSystem and Storage Virtualize
- IBM Storage for Data and AI
- IBM FlashSystem 9500 Deep Dive & Advanced Functions
- IBM Storage Fusion

Please reach out to your IBM Rep or Business Partner for future dates and to be nominated.

TEST DRIVE / DEMO'S

- North America ATG Storage - IBM Storage Scale and Storage Scale System GUI
- North America ATG Storage - IBM Storage Virtualize Test Drive
- North America ATG Storage - IBM DS8900F Storage Management Test Drive
- North America ATG Storage - Managing Copy Services on the DS8000 Using IBM Copy Services Manager Test Drive
- North America ATG Storage - IBM DS8900F Safeguarded Copy (SGC) Test Drive
- North America ATG Storage - IBM Cloud Object Storage Test Drive - (Appliance based)
- North America ATG Storage - IBM Cloud Object Storage Test Drive - (VMware based)
- North America ATG Storage - IBM Storage Protect Live Test Drive
- North America ATG Storage - IBM Storage Protect Plus Live Test Drive
- North America ATG Storage - IBM Storage Ceph Test Drive - (VMware based)

Please reach out to your IBM Rep or Business Partner for more information.

Accelerate with ATG Technical Webinar Series - Survey

Please take a moment to share your feedback with our team!

You can access this 6-question survey via Menti.com with code 2243 3599 or

Direct link https://www.menti.com/albneqj15g57

Or

QR Code

IBM TechXchange Conference Americas 2023

Must-Attend Technical Learning & Community Event for Tech Experts

September 11–14
MGM Grand Conference Center
Las Vegas

3
full days and activities + pre-conference day

8
technical tracks

- AI
- Data
- Automation
- Security
- Sustainability
- Hybrid Cloud
- Infrastructure
- Quantum

70+
products covered

1,000+
sessions, demos, instructor-led labs, roadmap discussions

5,000
technical peers to engage and network

Level up your skills in IBM tech and products

Get certification testing

“Look under the hood” at the latest updates

Stay ahead, Register Now
secure your spot →

IBM TechXchange

What to expect in the Infrastructure Track for Storage

Register now →

40+
3rd party speakers and SMEs

60
Deep technical keynotes & sessions

15+
Hands-on, instructor-led labs

Featured Themes

- AI/ML/HPC-driven hyper-Data Growth
- Modernization supply chain Edge-Core-Cloud
- Business & Operational Resiliency
- *and more*

Key Learnings

- Practical how-to advice
- Patterns and best practices
- Success stories, IBM PoV, proven techniques

Featured Products

- IBM Storage Defender
- IBM Storage Fusion
- IBM Storage Scale + IBM Storage Ceph
- IBM Storage FlashSystem + IBM Storage DS8000
- IBM Tape + IBM SAN

Amazing Technical Speakers

Goldman Sachs sharing their Cyber Resiliency experience with IBM SafeGuarded Copy

Steve Rapasky
Sr. Vice President
Goldman Sachs

University of Queensland - Fusion HCI - Container based computing in a research environment

Jake Carroll
Research Computing Manager
University of Queensland

Fundamentals of Building a Cyber Resilient Data Protection Infrastructure

IBM Champion
Jeff Helton
Data Resiliency Architect
Horizon Computer Solutions, Inc.

How IBM Storage FlashSystems fights in cyber warfare

IBM Champion
Tony Owens
Solution Architect
Arrow Electronics, Inc.

zHyperLink New Features and Real-World Experiences

Brian Sherman
Distinguished Engineer
IBM

How IBM Storage helps you overcome the data management and security challenges

Sam Werner
Storage Product Management
IBM

Note: Sessions and speakers are subject to change.

Advanced Technology Group

IBM

Accelerate with ATG:
Data resiliency with IBM Storage Scale

Lindsay Todd, PhD
Principal Storage Technical Specialist – Storage Scale/SSS for Data/AI/HPC
IBM Advanced Technology Group – Storage Team

IBM

Agenda

- What is data resiliency?
- What is Storage Scale?
- Storage Scale – features in support of **protect** and **detect**
  - ACLs and immutability
  - Encryption
  - File audit logging
- Storage Scale – features in support of **recovery**
  - Using Storage Scale snapshots
  - Using AFM
  - Using the policy engine with Storage Protect and `mmbackup`
- Storage Scale architectures for cyber-resiliency
  - Using File Audit Logging with IBM Qradar
  - Storage Scale cyber-vault architectures based on AFM
  - Cyber-resiliency through Storage Protect

What is data resiliency?

... and why do I care?

Data Resiliency

Data Resiliency is the ability to keep all the data available that an organization needs to continue functioning, even in the face of disruptions such as:

- Natural disasters
- Human-caused disasters
- “Oops!”
- Cyber attacks

2X
Cyber Attacks YTY

17%
of Cyber Attacks are Ransomware

21%
dormant threats, up from 5% YTY

26%
Clients who paid the ransom still could not recover the data

45%
45% of production data affected

23
days, average recovery after a ransomware attack

Data Resilience

Why should you care?

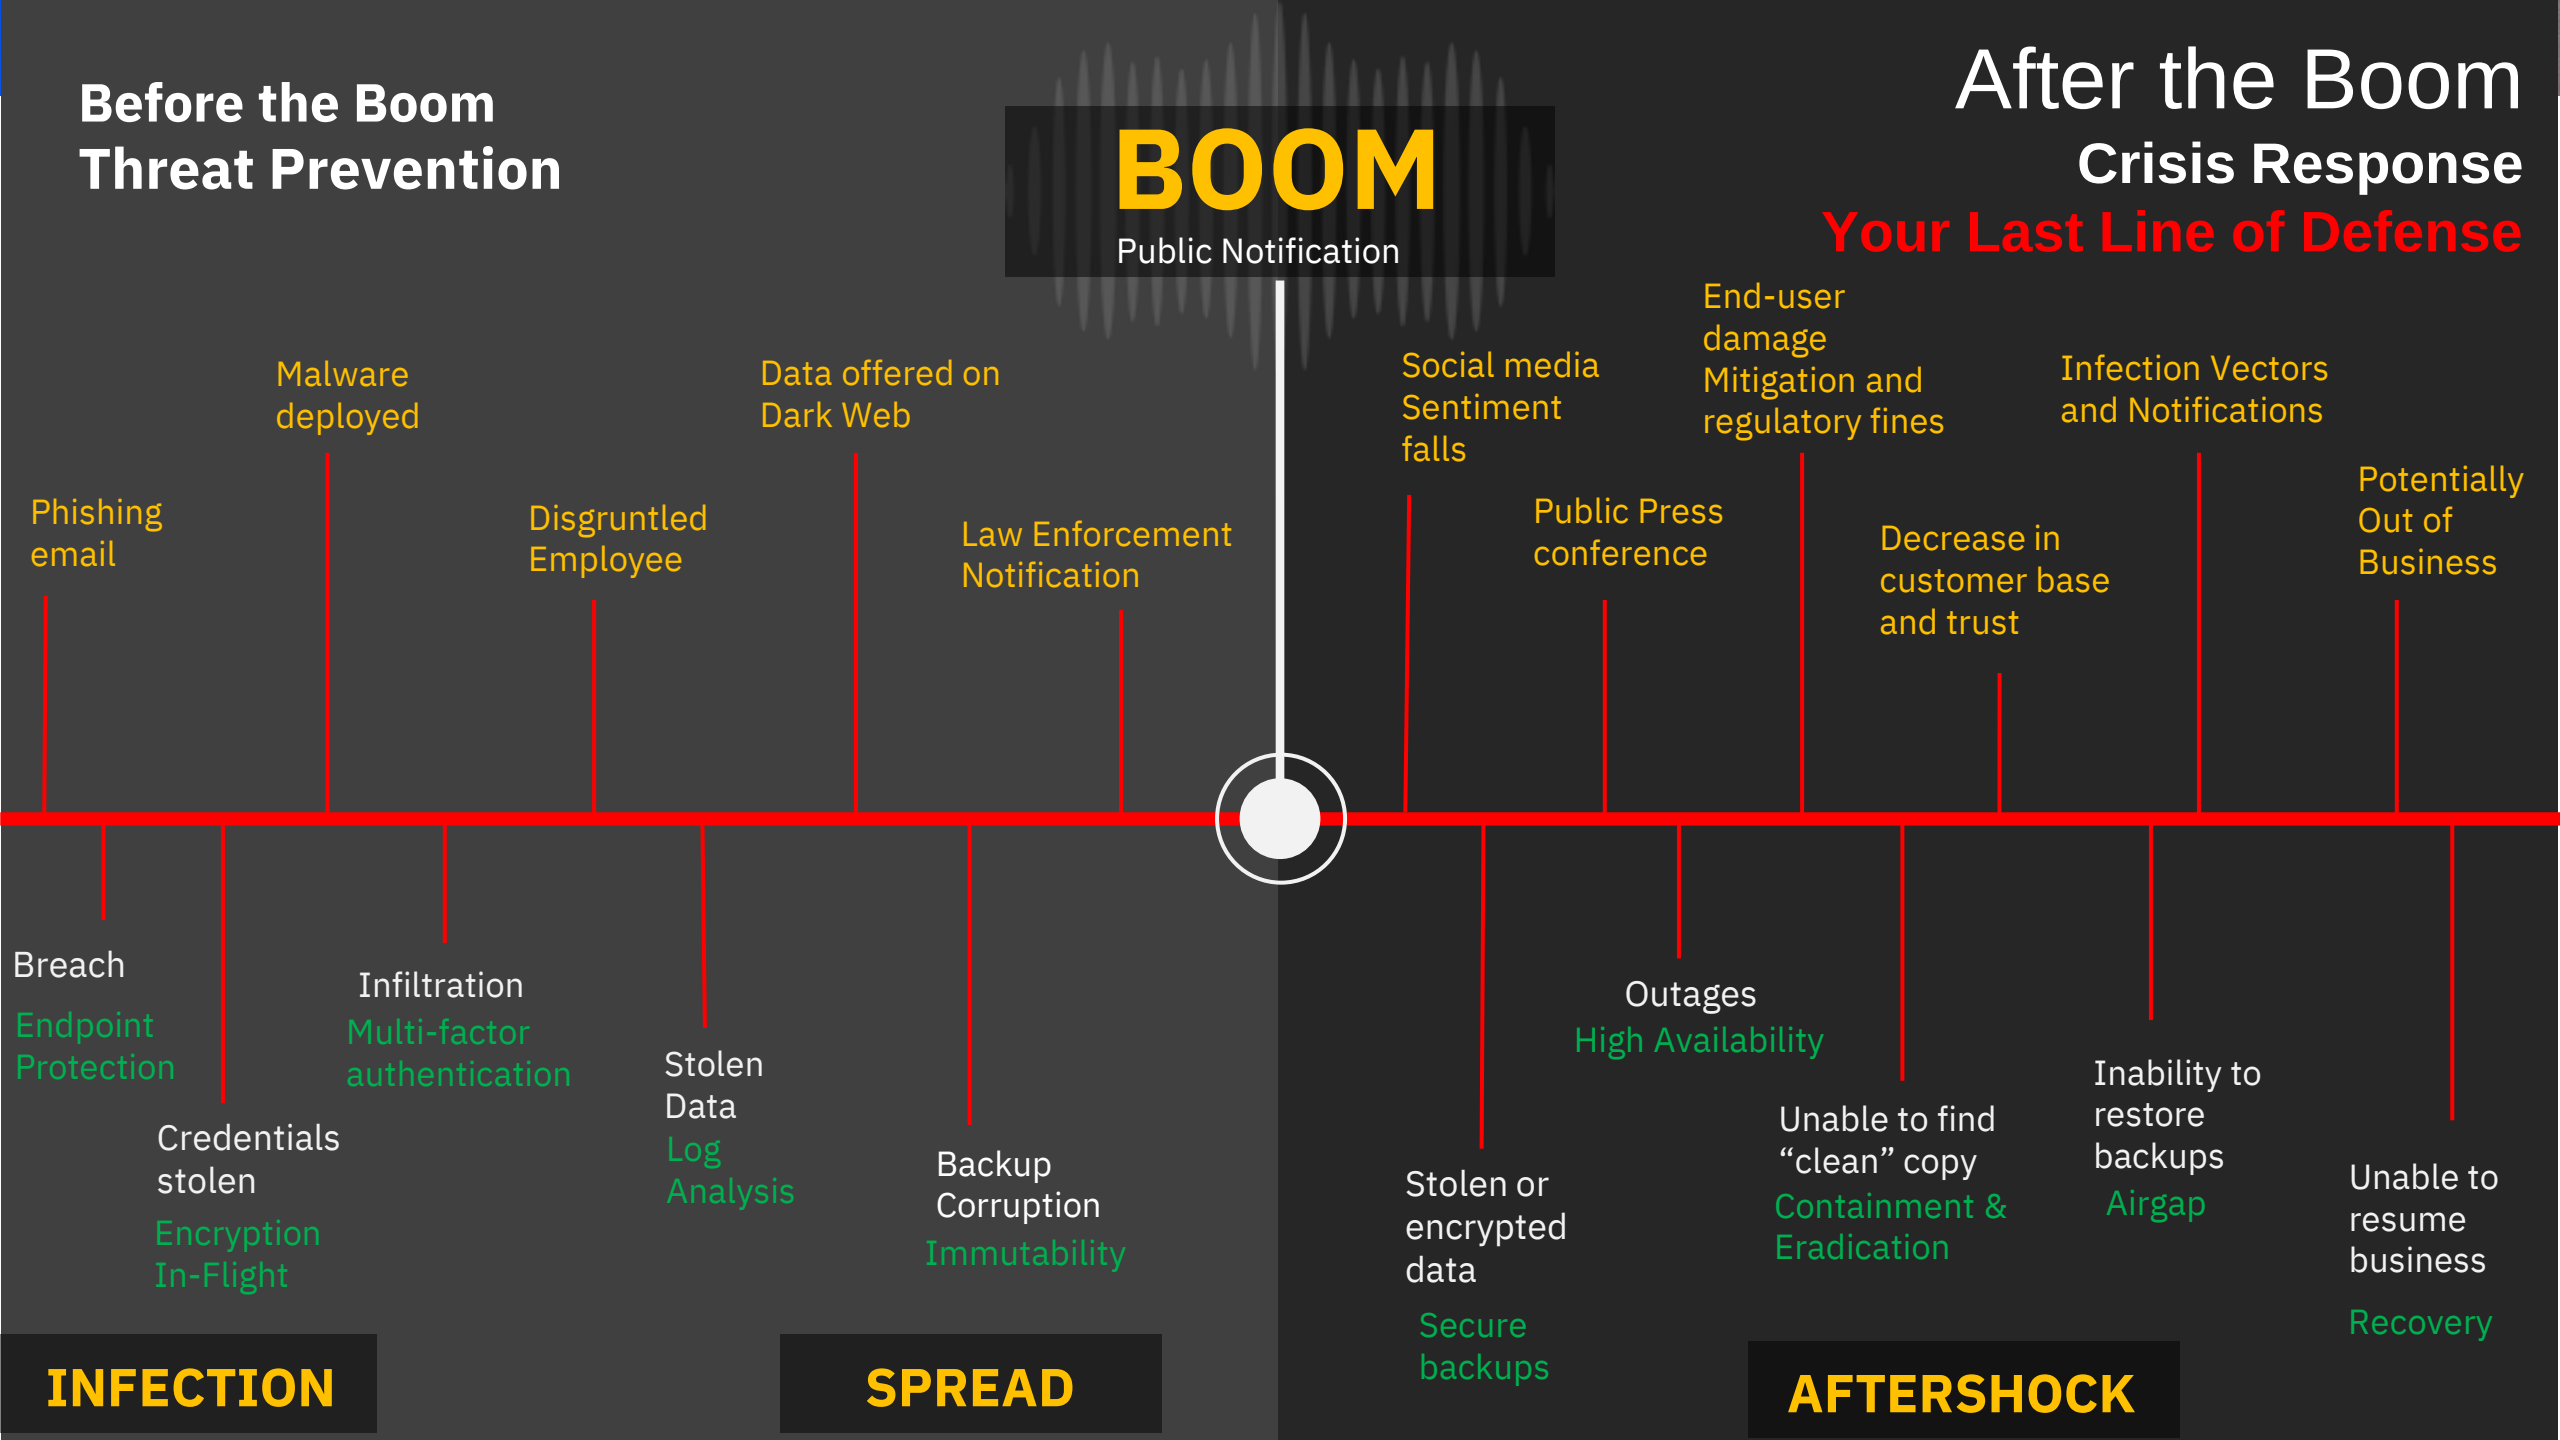
Before the Boom
Threat Prevention
BOOM
Public Notification
After the Boom
Crisis Response
Your Last Line of Defense
Phishing email
Malware deployed
Disgruntled Employee
Data offered on Dark Web
Law Enforcement Notification
Social media Sentiment falls
Public Press conference
End-user damage Mitigation and regulatory fines
Decrease in customer base and trust
Infection Vectors and Notifications
Potentially Out of Business
Breach
Endpoint Protection
Credentials stolen
Encryption In-Flight
Infiltration
Multi-factor authentication
Stolen Data
Log Analysis
Backup Corruption
Immutability
Stolen or encrypted data
Secure backups
Outages
High Availability
Unable to find “clean” copy
Containment & Eradication
Inability to restore backups
Airgap
Unable to resume business
Recovery
INFECTION
SPREAD
AFTERSHOCK

Business Impacts of Cyber Attacks

83%
of organizations have experienced more than one data breach

50%
of breach costs are incurred more than a year after the breach

60%
of organizations raised their product or services prices due to a breach

23
days, average recovery after a ransomware attack

$812k
average ransom payment in 2021

BUSINESS IMPACT

$5.2m

Average annual cost of Ransomware

What is IBM Storage Scale?

... the best General Parallel File System around...

Storage Scale – IBM's flagship file storage for unstructured data

CORAL COLLABORATION
OAK RIDGE • ARGONNE • LIVERMORE

https://www.olcf.ornl.gov/summit/

Pushing the limits

- 2.5 TB/sec single stream IOR
  as requested from Oak Ridge National Lab (ORNL)
- 1 TB/sec 1MB sequential read/write
  as stated in CORAL RFP
- Single Node 16 GB/sec sequential read/write
  as requested from ORNL
- 50K creates/sec per shared directory
  as stated in CORAL RFP
- 2.6 Million 32K file creates/sec
  as requested from ORNL

IBM Storage Scale is a *parallel* file system providing a single namespace over clusters of systems with the same complete file consistency a single system would enjoy, with outstanding performance.

Leading use cases for IBM Storage Scale and Storage Scale System

A Strategic

Storage for AI, Big Data, and Analytics
High performance and scalability

1. Data-intensive technical computing
2. Artificial Intelligence (AI)
3. Big Data and Analytics, Hadoop

B Strategic

Data lake(house), industry applications
Enterprise data architecture

4. Unified storage for “Data lake(house)”
5. Select ISV and industry solutions

C Tactical

Data optimization and resiliency
Enterprise data optimization and data management.
Infrastructure pre-requisites for segments A and B

6. Archive
7. Information Lifecycle Management
8. Back-up / restore

Storage Scale – a POSIX file system, and much more

Storage Scale: POSIX file system features, and:

- **Filesets** – partitioning a file system into pieces, without losing performance but gaining manageability.
  - Filesets can have quotas.
  - Filesets can be made immutable or append-only.
- **Snapshots** – efficient read-only copy of file systems or filesets.
- **AFM** – Caching and remote distribution of data in filesets.
- **Policy engine** – policy-controlled encryption, compression, expiration, and tiering – including to tape and cloud.
- **File Audit Logging** – log all activity

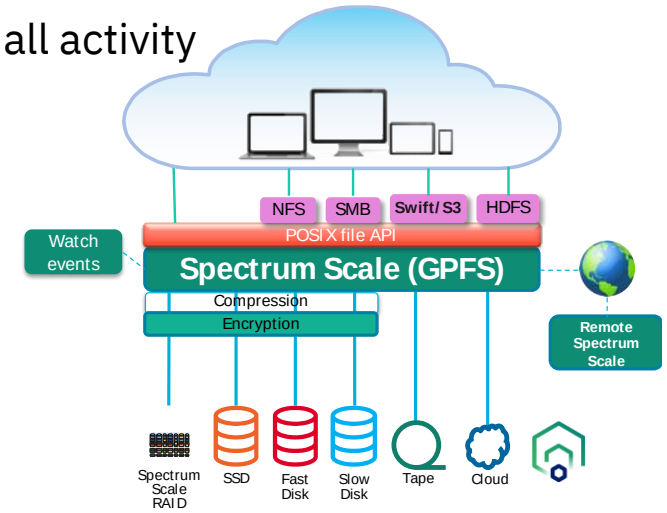

Fileset looks just like a directory

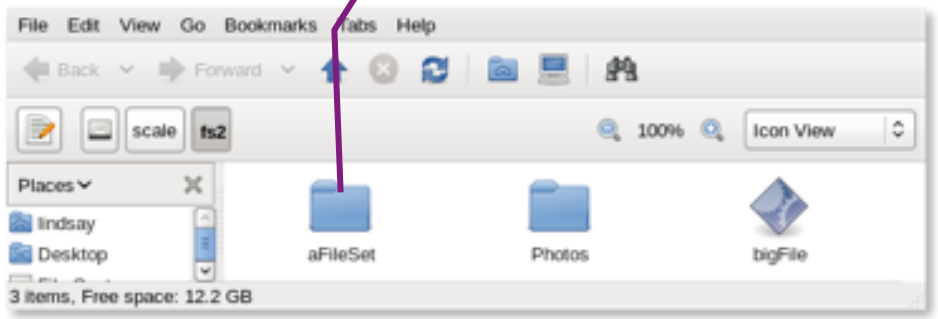


```
$ pwd
/scale/fs2
$ ls -l
total 20608
drwxr-xr-x. 4 root root     4096 Jan 17 10:07 aFileSet
-rw-r--r--. 1 root root 21098464 Jan 13 15:57 bigFile
drwxr-xr-x. 2 root root     4096 Jan 13 15:56 Photos
$ cd aFileSet
$ aFileSet]# ls -l
total 61824
drwxr-xr-x. 2 root root     4096 Jan 17 10:06 aDir
drwxr-xr-x. 2 root root     4096 Jan 17 10:07 anotherFileSet
-rw-r--r--. 1 root root 63295392 Jan 17 10:07 biggerFile
[root@gpfs01 aFileSet]#
```

Containers too – Scale's Container Storage Interface (CSI) driver

The Storage Scale Container Storage Interface (CSI) driver lets Kubernetes (K8s) containers use Storage Scale storage for persistent volumes.

- Both RWX (ReadWriteMany) and RWO (ReadWriteOnce)
- Both static and dynamic provisioning
- Volumes may be directories or filesets

The driver uses the REST management API provided by the Scale GUI (which, by the way, will be getting 2FA in the future).

Container Native Storage Access (CNSA) runs the Storage Scale NSD clients as K8s pods in an OpenShift environment.

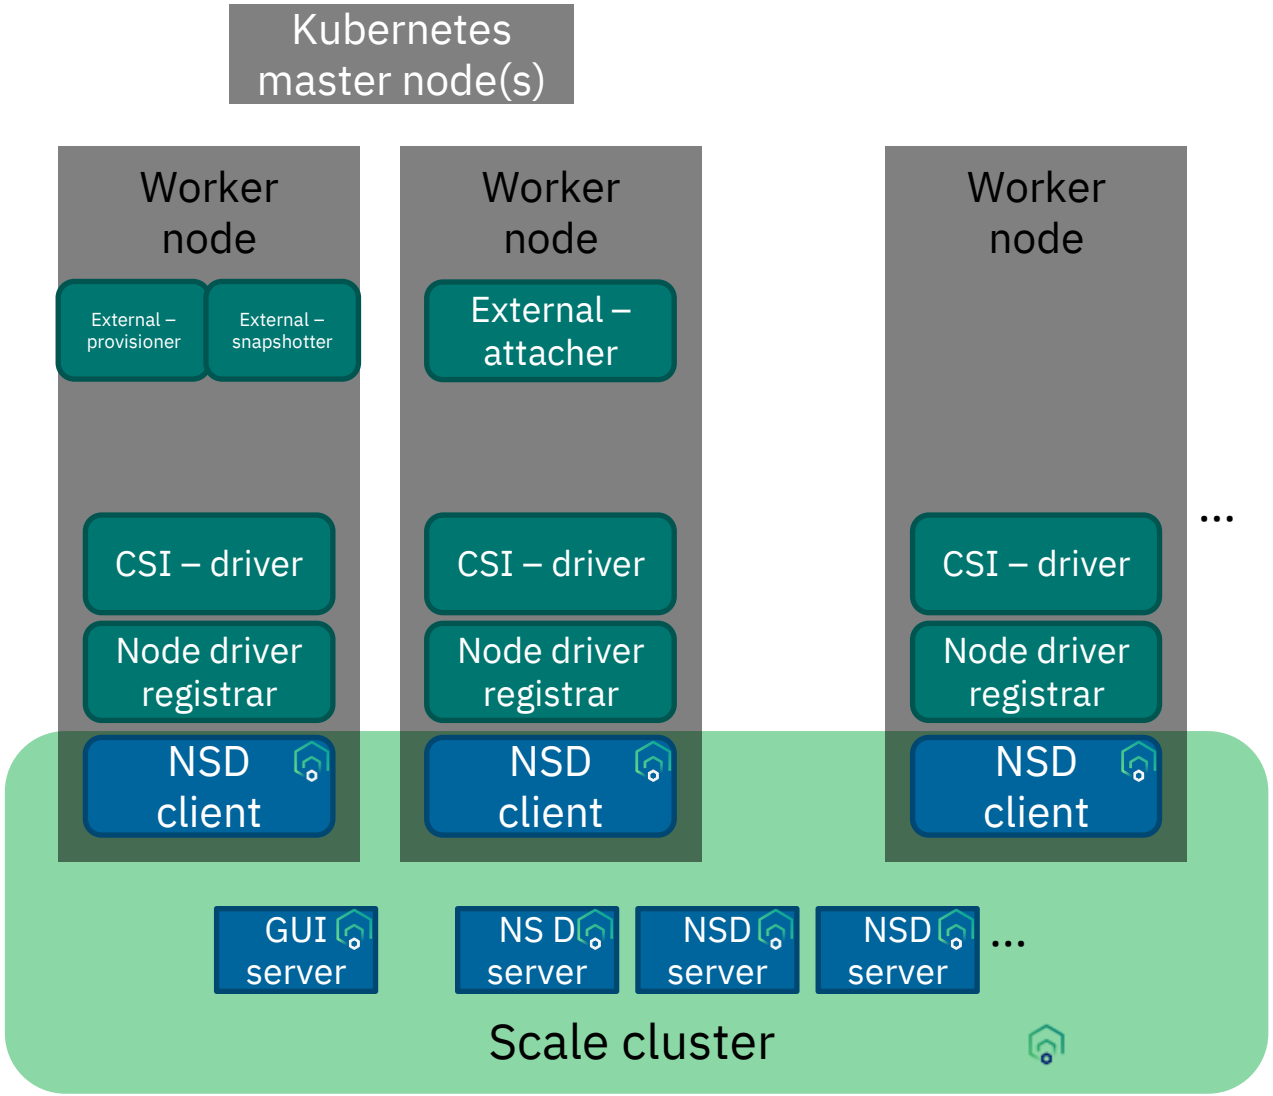

Storage Scale stretched clusters

A *single* Storage Scale cluster can be configured using nodes and storage from two data centers.

- In other words, it is “stretched” between two sites, connected by a WAN.

File systems in such a cluster are available to systems at both sites and may be actively used concurrently by both sites.

The file systems may judiciously use “failure group” replication to ensure both sites have a current replica of all the data.

Careful design can ensure one site remains active, even if the other site (or link) fails.

Result is an **active/active highly available** synchronously replicated Storage Scale file system.

Of course, you’ll need more to defend from cyber-attacks.

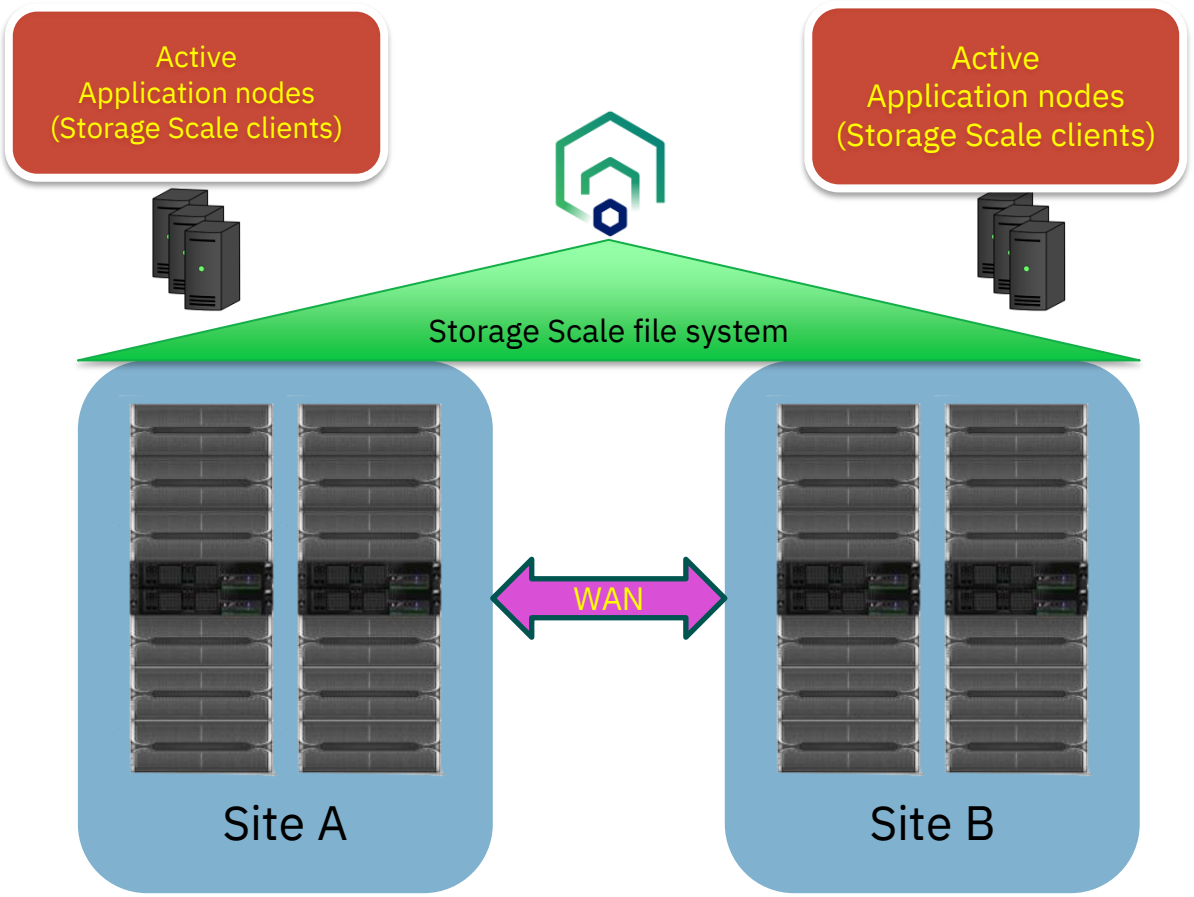

Getting deeper with Scale cyber-resiliency – the NIST Cyber Resiliency Framework

Framework of computer security guidance for how organizations can assess and improve their ability to prevent, detect, and respond to cyber attacks:

Identify:

Defining a organizational understanding to build or improve **cyber resiliency plan** – critical assets and strategy

Protect:

Implementing Safeguards to ensure delivery of critical services – protecting against vulnerabilities before they are exploited

Detect:

Detecting occurrence of cyber security events – timely, continuous monitoring, detection processes

Respond:

Taking action regarding a detected event – analysis, **contain**, mitigation, and communication

Recover:

Restore capabilities and services - recovery, improvements, communications

National Institute for Standards and Technology (NIST)

Security Guidelines for Storage Infrastructure

Storage Scale features to protect and detect

… so attackers can’t get very far

Access Control Lists – Supported by IBM Storage Scale

Storage Scale supports both POSIX and NFS4 ACLs.

- POSIX ACLs compare a process' user and group privileges to those of the file, regulating access to its data.
- NFS4 ACLs are richer but more complicated.

Users are authenticated by the operating system determining:

- privileges of the user,
- groups the user belongs to,
- and privileges of those groups

Other attributes on files and filesets

Notable attributes available include:

- **Immutability** (file or fileset)
  - controls whether or not a file may be modified
  - retention and expiration time attributes control ability to delete immutable files,
  - complies with US SEC17a-4f .
- **Append Only** (file or fileset)
  - Controls whether a file may be modified, or only allow data to be appended to it.
- **Extended attributes** (“xattrs”)
  - Ability to specify user or application-specific attributes on files

```
$ mmlsattr -L -d biggerFile
file name:            biggerFile
metadata replication: 2 max 3
data replication:     1 max 3
immutable:            no
appendOnly:           no
flags:
storage pool name:    system
fileset name:         aFileSet
snapshot name:
creation time:        Tue Jan 17 10:07:03 2017
Misc attributes:      ARCHIVE
Encrypted:            no
security.selinux:     "unconfined_u:unlabeled_t"
user.complicated:     ”you_got_this”
$
```

Storage Scale secure data at rest – Key-based policy-driven encryption

Master Encryption Key (MEK)

- Used to encrypt file encryption keys
- Stored in Remote Key Management (RKM) Servers
- MEK's have a unique key name that combines the name of the key and the RKM server where it resides

File Encryption Key (FEK)

- Used to encrypt sectors of an individual file
- Unique key randomly generated
- Encrypted (or "wrapped") with one or more MEK's and stored in the extended attribute on filesystem.
- FEK must have access to MEK to be decoded
- FEK can be re-wrapped to new MEK(s) in the case of a compromised key

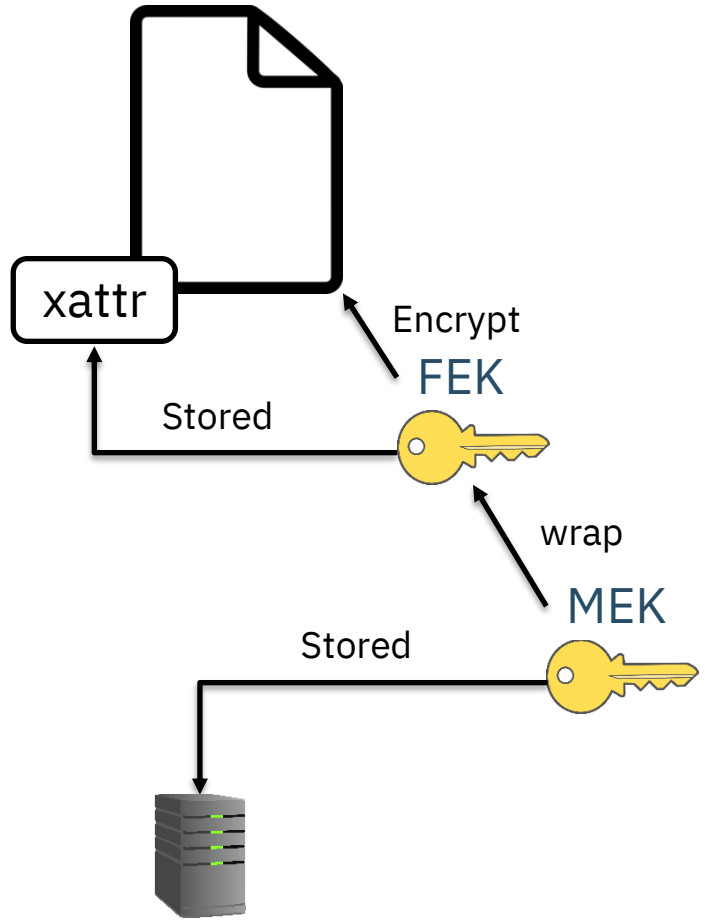

External Key Manager Server
(IBM GKLM or Vormetric DSM Key Server)

File Audit Logging

Captures the most common types of file activity:

- open, close, delete, rename, POSIX permission changes, ACL changes, etc., configurable with `mmaudit`.
- Doesn't capture internal operations (e.g., restripe).

Events are captured within the Storage Scale `mmfsd` daemon, representing attributes of filesystem action at that point.

Interfaces to IBM Guardium and Varonis.

- Each file system enabled has a dedicated fileset where the audit logs will go.
  - Default option is `.audit_log` at the root of the file system.
- Audit log files are nested within `/FSNAME/.audit_log/topic/year/month/date/*`
- Log files are append-only with default retention of 365 days.

```
{"LWE_JSON": "0.0.1", "path": "/newfs/1Kfile2.restore", "oldPath": null,
"clusterName": "pardie.cluster", "nodeName": "c6f2bc3n10", "nfsClientIp": "",
"fsName": "newfs", "event": "OPEN", "inode": "26626", "openFlags": "32962",
"poolName": "sp1", "fileSize": "0", "ownerUserId": "0", "ownerGroupId": "0",
"atime": "2017-10-25_12:36:22-0400", "ctime": "2017-10-25_12:36:22-0400",
"eventTime": "2017-10-25_12:36:22-0400", "clientUserId": "0", "clientGroupId":
"0", "processId": "10437", "permissions": "200100644", "acls": "u::rwc, g::r,
o::r, ", "xattrs": null }
```

Example audit log entry

Storage Scale features to recover

... so you can prepare for the worst

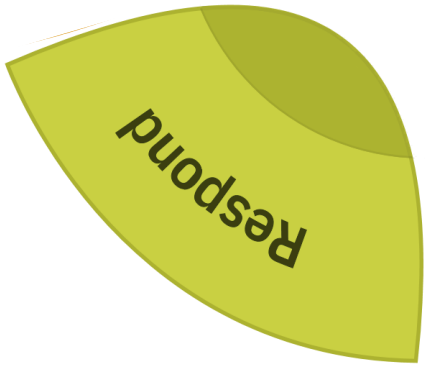

Protecting live data with Storage Scale snapshots

Snapshots are efficient ***read-only*** copies of either entire *file systems* or individual *filesets*.

- File systems may have 256 global snapshots.
- Additionally, each fileset may have 256 snapshots.

Snapshots use ***copy-on-write*** semantics:

- Snapshot creation is fast – only inodes are copied.
- Snapshots are space-efficient – the only space taken is by changes to the source file system or fileset.

Individual files can be copied out of the snapshot directory.

- This can be useful as a simple “self-serve” recovery mechanism.

Filesets or file systems can be restored to the state of a snapshot.

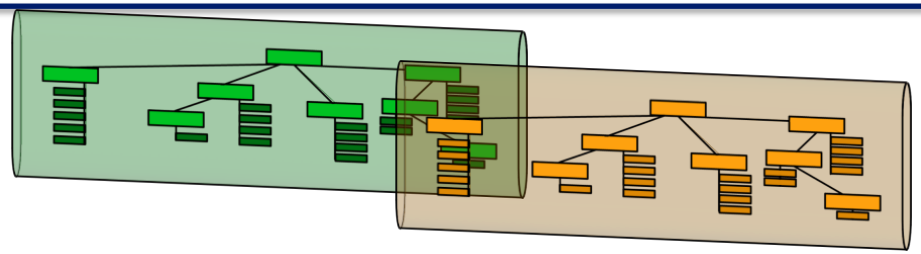

```
prodadmin$ ls demosnap/
hill  world

prodadmin$ sudo mmcrsnapshot fs1 snap1 -j demosnap
Flushing dirty data for snapshot demosnap:snap1...
Quiescing all file system operations.
Snapshot demosnap:snap1 created with id 38.

prodadmin$ rm demosnap/hill
prodadmin$ ls demosnap/
world

prodadmin$ ls demosnap/.snapshots/snap1/
hill  world
prodadmin$ cp demosnap/.snapshots/snap1/hill demosnap/
prodadmin$ ls demosnap/
hill  world
```

“Immutable” snapshots for Storage Scale Safeguarded Copy

A “Safeguarded Copy” (SGC) is a snapshot with a retention time.

- Like all other snapshots, the data is a read-only copy of a file system or independent fileset.
- Data from an SGC can be restored using `mmrestorefs`, or by directly copying affected files.
- SGCs cannot be deleted with `mmdelsnapshot` before the retention time (set with `--expiration-time` flag).
  - The file system version must be at least 5.1.5.0 to set a retention time.
- Snapshots and SGCs can both be automatically created on a schedule set in the Storage Scale GUI.

Storage Scale can be configured (using the “sudo-wrappers feature”) to require the collaboration of a separate security administrator to delete a SGC snapshot before its expiration time (with **`mmrestrictedctl`**).

```
prodadmin$ sudo mmcrsnapshot fs1 sgc-2 -j home --expiration-time 2022-11-05-12:00
Flushing dirty data for snapshot home:sgc-2...
Quiescing all file system operations.
Snapshot home:sgc-2 created with id 8.

prodadmin$ mmlssnapshot fs1 -j home
Snapshots in file system fs1:
Directory                SnapId    Status  Created
ExpirationTime           Fileset
psnap0-rpo-15DA000A6358816F-7 5         Valid   Fri Oct 28 17:52:27 2022  Fri Oct
28 17:52:27 2022  home
sgc-1                    6         Valid   Fri Oct 28 19:46:17 2022  Sat Oct 29
12:00:00 2022  home
sgc-2                    8         Valid   Mon Oct 31 10:35:56 2022  Sat Nov  5
12:00:00 2022  home

prodadmin$ date
Mon Oct 31 10:47:19 EDT 2022

prodadmin$ sudo mmdelsnapshot fs1 sgc-1 -j home
Invalidating snapshot files in home:sgc-1...
Deleting files in snapshot home:sgc-1...
 100.00 % complete on Mon Oct 31 10:50:02 2022  (     81920 inodes with total
8 MB data processed)
Invalidating snapshot files in home:sgc-1/F/...
Delete snapshot home:sgc-1 successful.

prodadmin$ sudo mmdelsnapshot fs1 sgc-2 -j home
The snapshot home:sgc-2 cannot be deleted since its snapshot retention time is not
yet expired.
Delete snapshot Error: home:sgc-2, err = 1.
mmdelsnapshot: Command failed. Examine previous error messages to determine cause.

prodadmin$ sudo mmrestrictedctl fs1 deleteSnapshot sgc-2 -j home
Sorry, user prodadmin is not allowed to execute '/usr/lpp/mmfs/bin/mmrestrictedctl
fs1 deleteSnapshot sgc-2 -j home' as root on prodb01-wan.local.
```

Using Storage Scale AFM and active/passive asynchronous DR

Active File Management (AFM) is Storage Scale's file system caching feature.

- A fileset can be configured as a cache for a remote target file share.
- Anything written into the cache will be "replicated" to the remote target.
- The **AFM-DR** feature extends AFM to be a full DR solution, we the secondary site kept read-only until failover is needed.

AFM can support "peer snapshots" – snapshot of cache fileset will create identical snapshot at the target site.

- AFM-DR can regularly take a peer snapshot (but only one automatic snapshot is retained), or additional peer snapshots can be taken (manually or out of `cron`) – these are retained.
- Single-site snapshots (including SGCs) can be made of either the primary or secondary fileset.

The secondary site is a **separate** Storage Scale cluster, potentially with different administrators.

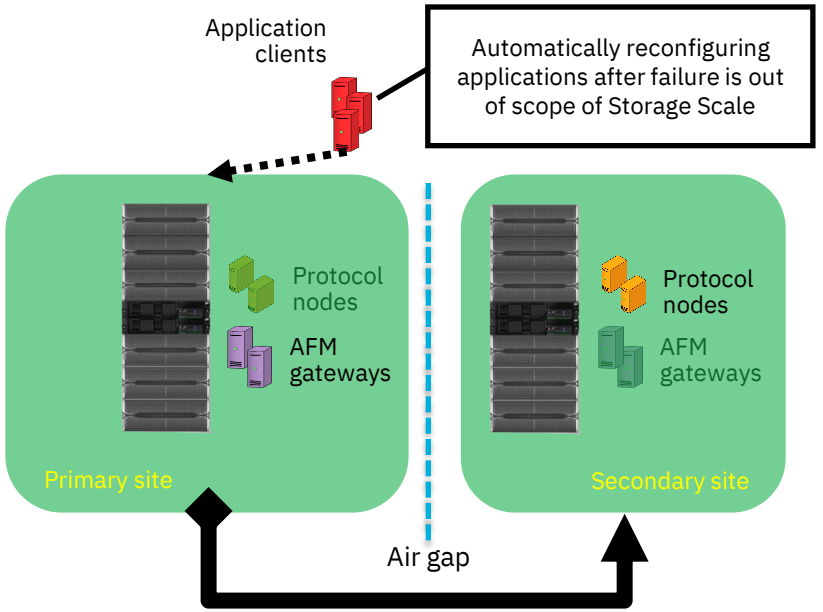

Backing up the Scale File System

Backing up the user data of a Scale file system requires also backing up the metadata associated with each file and directory:

- Create, access, and modify times
- Size of the file
- Mode bits and owner, owning group
- Immutability attributes
- Access Control Lists
- Extended attributes

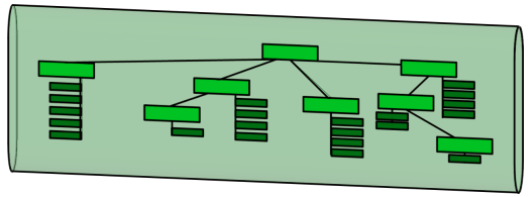

IBM Storage Protect does incremental backups and stores all these metadata. (Other vendors' backup tools may or may not handle all metadata.)

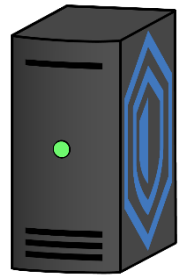

Backups take time, during which files may change – so it is best practice to backup snapshots rather than live files.

Using the Storage Scale policy engine to backup and/or tier to tape or object storage

Besides supporting conventional backup, Storage Scale's policy engine can incorporate tape and object storage as a storage tier.

Tape is extremely cost effective, long-lived, and multiple copies can be kept offline at multiple physical locations.

Hierarchical Storage Management (HSM) to tape can be done using IBM Storage Protect or IBM Storage Archive, and to Object storage using either IBM Storage Protect.

- Storage Scale has a "Scale-out Backup and Recovery"(SOBAR) feature to quickly make files co-resident.
- With SOBAR, a new file system of stubs can be generated, and files restored on demand.

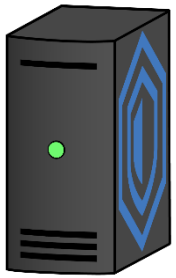

Backup-Archive Engine Storage Tier

Tape Storage Tier

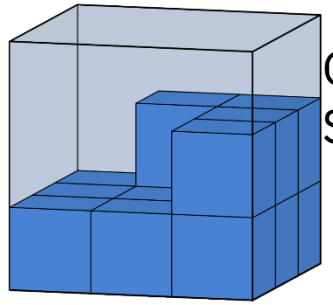

Object Storage Tier

Cyber-resilient architectures

... with Storage Scale in front and center

Using IBM Storage Scale's File Audit Logging with IBM QRadar

Storage Scale's File Audit Logging feature records file access events, with audit information, into a JSON files stored in an immutable (append-only) fileset.

These files can then be parsed by other tools, such as IBM QRadar. These tools can report and react to unusual activity.

- For example, QRadar could be configured to automatically take snapshots of filesets, or perhaps change file system ACLs, if an unexpected file access is detected.

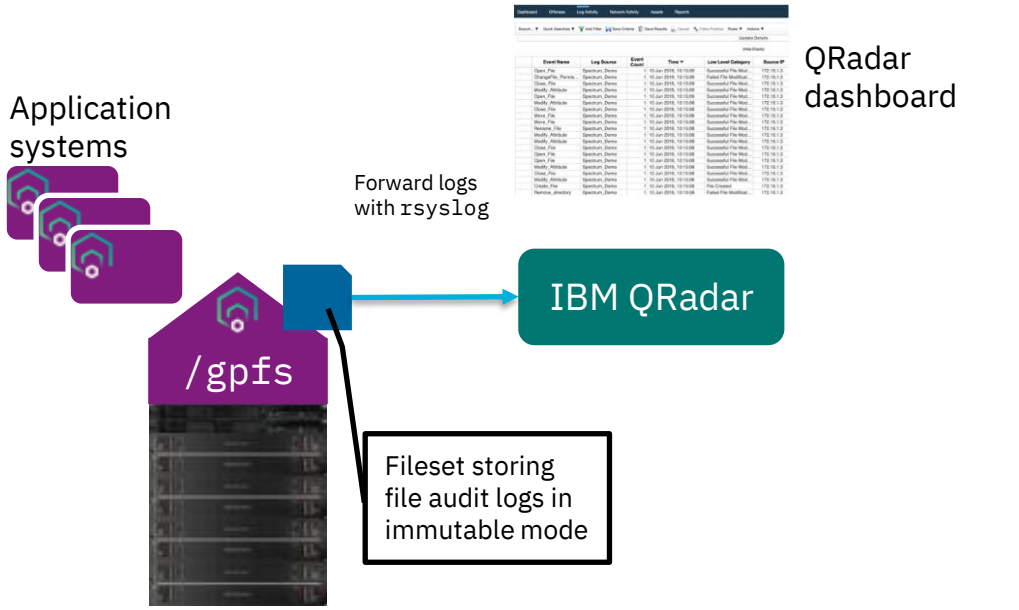

Creating a writable copy of a Safeguarded Copy for a cyber-vault

Snapshots, including Safeguarded Copies, are accessible through the file system as directory structures mirroring the original directory structure.

- Application verification can be done directly on the snapshot.
- But… all snapshots are read-only, including SGCs, and sometimes verification will require a writable instance of the data (for example, to roll back incomplete database transactions).
- Active File Management (AFM) caches from a data source into a cache fileset.
  - Only referenced data is cached.
  - LocalUpdate mode allows updates to the cached data, without trying to send back changes.
- Applications can use the LU cache instead of the read-only snapshot.

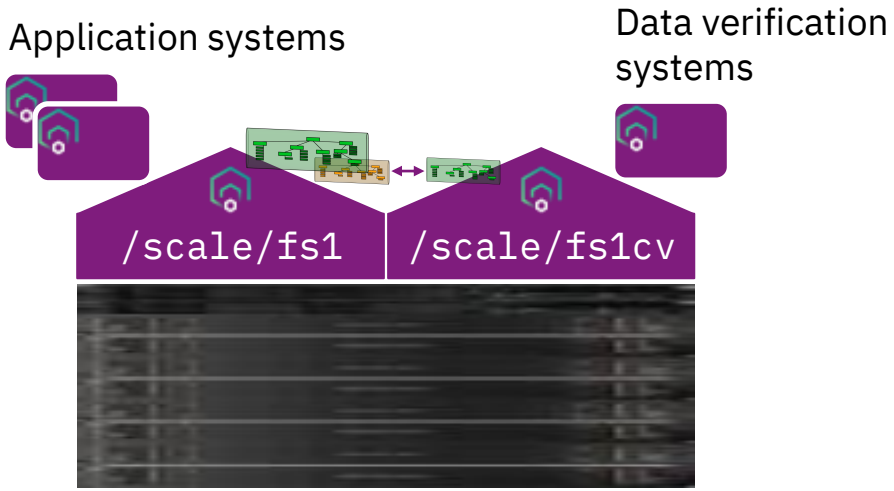

Two file systems in the same Scale cluster:

- `/scale/fs1` has application filesets and safeguarded copies of those filesets
- `/scale/fs1cv` has AFM LU (“LocalUpdate”) caches of the SGCs

Using AFM-DR to build an air-gapped cyber-vault

AFM-DR creates a read-only copy of the primary data at the secondary site.

AFM-LU caches can be made of peer snapshots of the secondary fileset to give effective an effective cyber-vault.

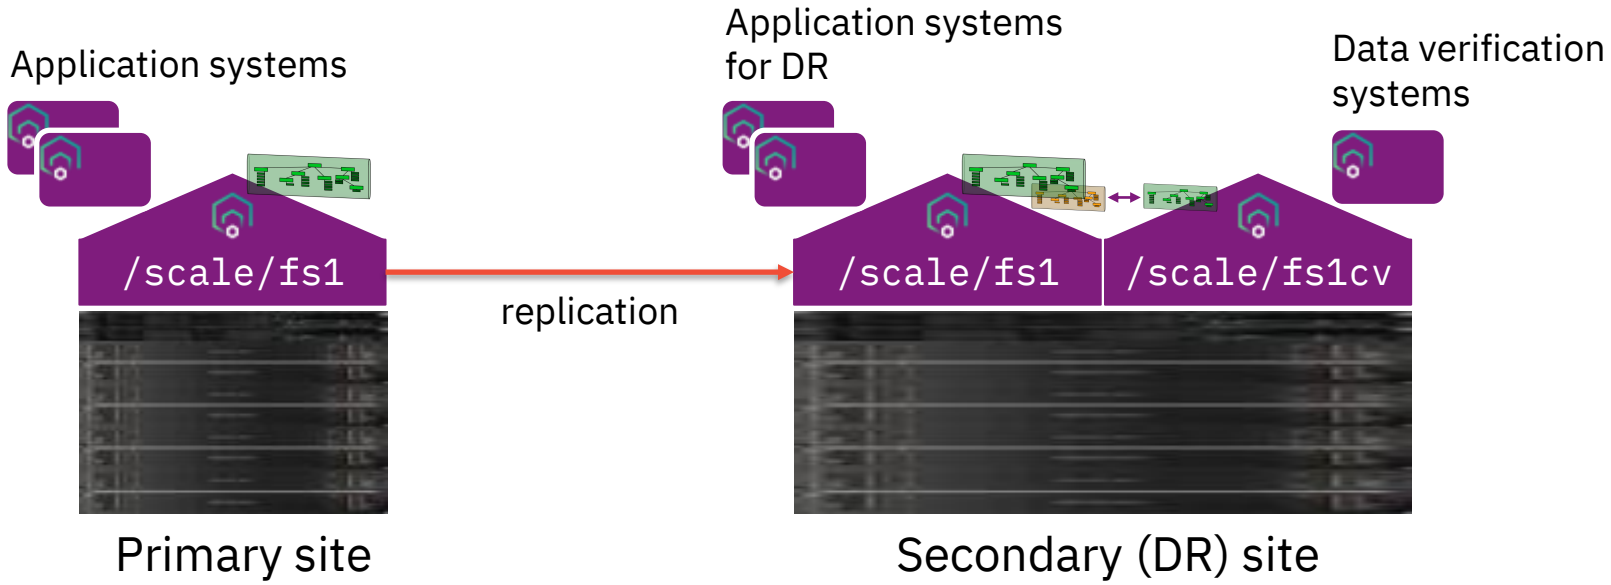

Using Storage Protect to build an air-gapped cyber-vault

After initial synchronization of `/gpfs0` and `/gpfs1`:

- At the primary site, the `mmbackup` command is used to regularly back up `/gpfs0`.
- The file list generated by `mmbackup` is also backed up.
- At secondary site, the file list is regularly restored, and used to choose files to restore into `/gpfs1`.

Results:

- The secondary site's `/gpfs1` is kept nearly synchronized with the primary site's `/gpfs0`, ready for an emergency.
- A complete backup of `/gpfs0` is maintained.

Optional:

- Additional snapshots of `/gpfs1` can be retained to provide additional fast restoration points.
- Backups can be migrated to WORM tape or object storage.

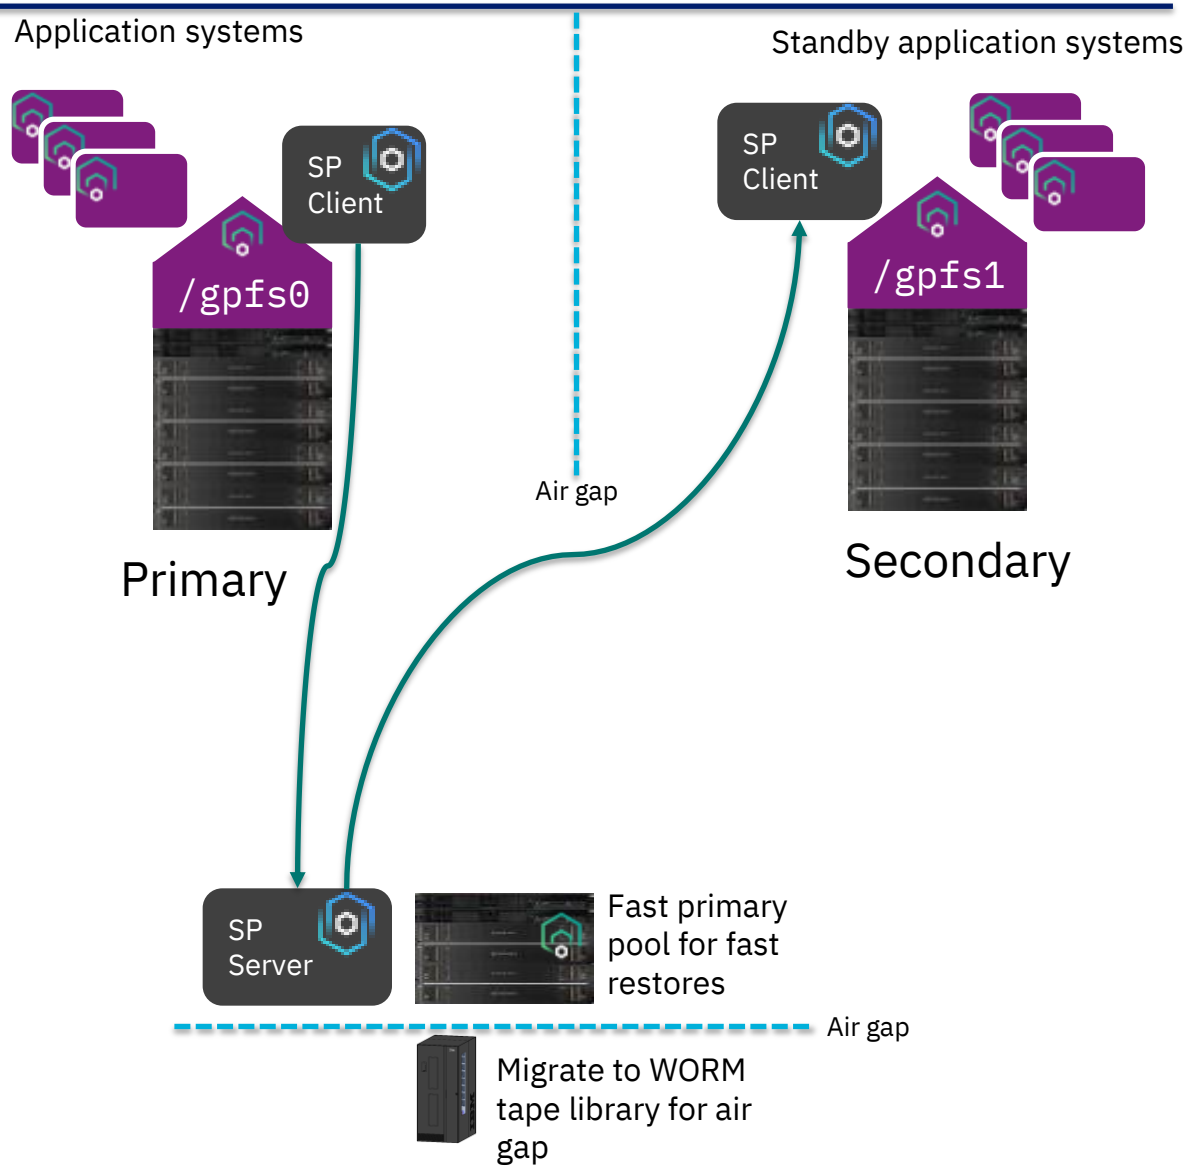

Demo
Air-gapped using AFM-DR

... combining snapshots with replication

Logical diagram for AFM demo

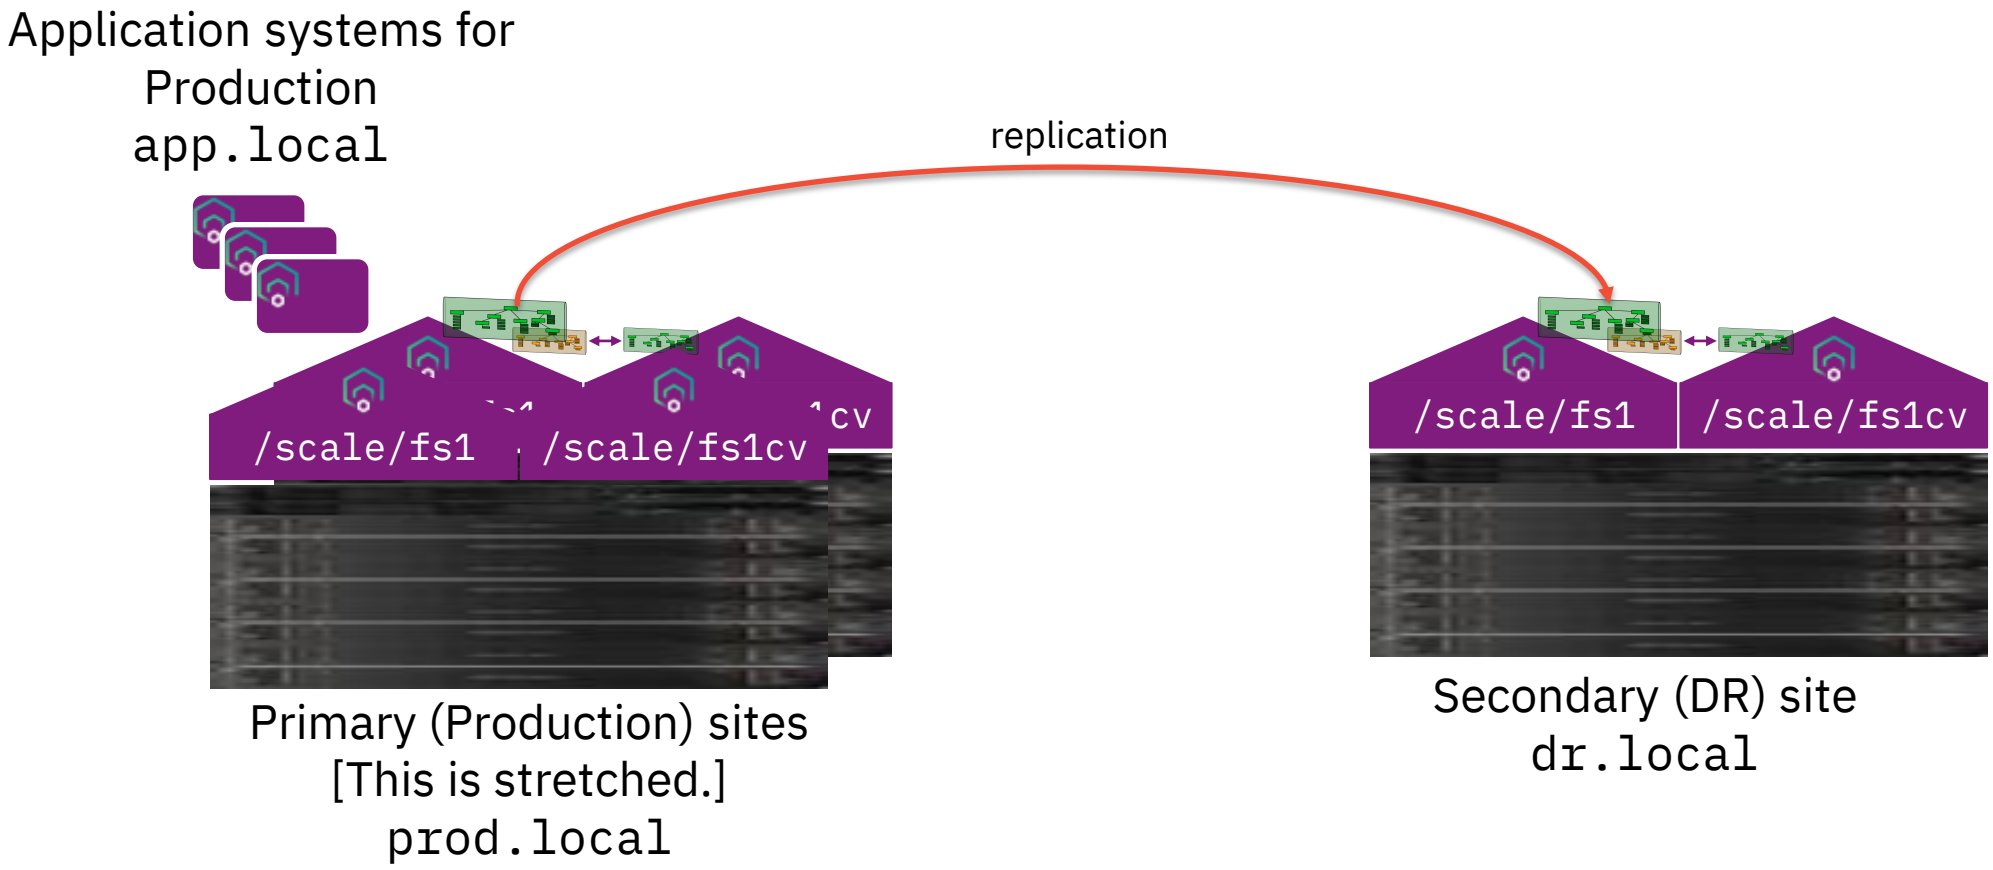

Taking peer snapshots with AFM-DR

You can still create separate snapshots on both primary and secondary sites, e.g.:

```
mmcrsnapshot dr secFileset:b4hack
```

Also, the mmpsnap command can be used to create and delete peer snapshots (“psnap”).

- AFM replicates constantly but updates may not be completely in order – this is fine for many applications, but a psnap will guarantee all updates are completed before the snapshot is generated.
- The psnap on the primary side has a corresponding snapshot on the secondary side – the two snapshots should be identical.
- Thus, the psnap serves as a recovery point for applications that require in-order updates.

```
## On Prod – make psnap and list snapshots
prodadmin$ sudo mmpsnap fs1 create -j repl
prodadmin$ mmlssnapshot fs1 -j repl

## On DR – list snapshots:-
dradmin$ mmlssnapshot fs1 -j repl

## On Prod – restore from a snapshot
prodadmin$ sudo mmrestorefs fs1 SNAPNAME –j repl

## On Prod - remove psnap
prodadmin$ sudo mmpsnap fs1 delete –s SNAPNAME –j repl
```

Build a writable cache copy of a snapshot

A local-update cache is similar to a “read only” cache – accessed files are cached, but they can be updated.

- Updates are not propagated to the target.
- Unmodified files can still be expired – so the local-update cache is a space-efficient writable copy of a space-efficient snapshot.

This writable copy can be used for:

- Failover testing of AFM-DR, without actually taking the primary offline.
- Validating the integrity of the data (which may require writing – such as for validating a database) – this is the basis of a cyber vault architecture.

```
## On DR – make cache of fileset psnapshot
dradmin$ sudo mmcrfileset fs1cv repl \
    --inode-space new -p afmMode=LU \
    -p afmTarget=gpfs:///scale/fs1/repl/.snapshots/candidate
dradmin$ sudo mmlinkfileset fs1cv repl -J /scale/fs1cv/repl
dradmin$ sudo mmafmctl fs1cv prefetch -j repl \
    --directory /scale/fs1cv/repl --metadata-only
```

Failover to the AFM-DR secondary

If a cyber-security event is detected at the primary site, unlink the fileset to stop damage!

- QRadar could automatically trigger these steps.
- Or perhaps found in routine scan, or through scanning of a cyber-vault
- With NFS as the transport, an alternative strategy would be to stop exporting that secondary.

Then convert the secondary site to acting primary.

We might have to roll back snapshots to get a clean replica.

Then applications can start on the secondary site.

```
## On DR – Unexport!
dradmin$ sudo mmnfs export remove /scale/fs1/repl

## On DR – Convert secondary to acting primary
dradmin$ sudo mmafmctl fs1 failoverToSecondary -j repl

## On DR - Look for damage…
dradmin$ sudo ls -l /scale/fs1/repl

## On DR - Look for a snapshot to restore from.
dradmin$ mmlssnapshot fs1 –j repl

## On DR – restore from that snapshot
dradmin$ sudo mmrestorefs fs1 SNAPNAME -j repl
# ls -l /scale/fs1/repl/
```

Further resources

IBM Documentation

- for Storage Scale – https://www.ibm.com/docs/en/storage-scale
- for Storage Scale System – https://www.ibm.com/docs/en/ess

Storage Scale FAQ – https://www.ibm.com/docs/en/STXKQY/gpfsclustersfaq.html

Spectrum Scale User Group – https://www.spectrumscale.org/

Storage Scale Resource Collection (a Box folder) – https://ibm.biz/Scale-Resource-Collection

Some IBM Redbooks:

- Cyber Resiliency Solution for IBM Spectrum Scale – pub#: REDP-5559
- Securing Data on Threat Detection Using IBM Spectrum Scale and IBM QRadar: An Enhanced Cyber Resiliency Solution – pub#: REDP-5560

Thank you!

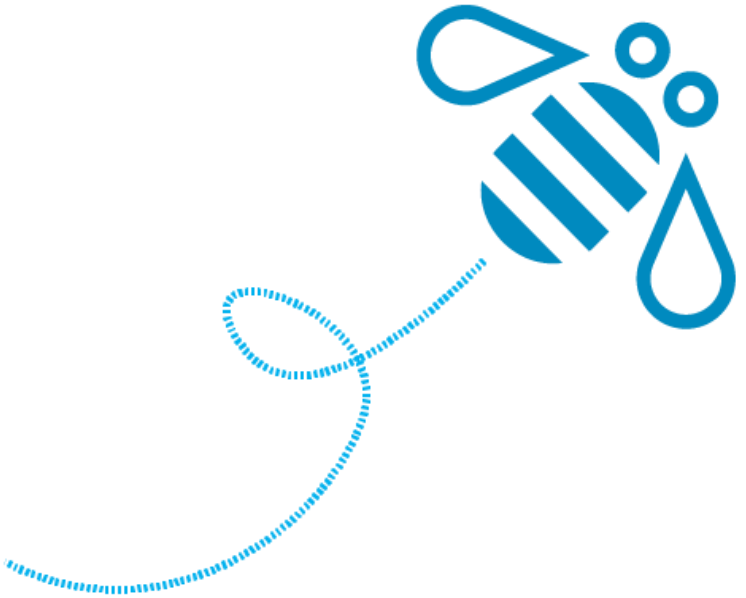

Accelerate with ATG Technical Webinar Series - Survey

Please take a moment to share your feedback with our team!

You can access this 6-question survey via Menti.com with code 2243 3599 or

Direct link https://www.menti.com/albneqj15g57

Or

QR Code