


AFM Update

Storage Scale User Group Meeting @ SC23
Denver, CO – Nov 12, 2023

Venkat Puvvada (IBM)






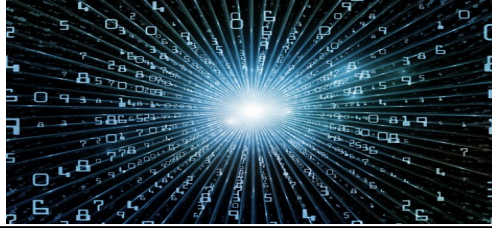
Storage Scale

Disclaimer

- IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code, or functionality. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.
- 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.



Data Caching Services (Active File Management) Use cases

			
<h3>Data Virtualization</h3>	<h3>Data Collaboration</h3>	<h3>Data Resilience</h3>	<h3>Hybrid cloud / Bursting</h3>
<ul style="list-style-type: none"> • Integrate legacy file and object data stores into single file system to breakdown legacy data silos • Migrate data to new storage or continue to use legacy stores • Create a High-Performance Tier for analytics for legacy data with transparent data access 	<ul style="list-style-type: none"> • Geo-distributed collaboration on data transparently shared between data centers, the cloud and edge sites • Coalesce data to a home site from the edge and redistribute it to all sites 	<ul style="list-style-type: none"> • Provide an asynchronous Disaster Recovery solution for business continuity over WAN distances • Supports analytics and archival access to passive data 	<ul style="list-style-type: none"> • Dynamically increase computation resources in the cloud and optimally make required data available for Cloud bursting • Process data consolidated on S3 Cloud Storage on with high performance tier in the Cloud Compute Cluster • Archive data to S3 Object storage
<p>Public Cloud Services</p> <p>Use case:</p> <p>Enables end user service to upload large amount of data via Object interface that can be analysed on high performance file system</p>	<p>Research / University</p> <p>Use case:</p> <p>Generate 100's of TB per day across multiple silos, leveraged AFM to provide common namespace with transparent multiprotocol data access</p>	<p>Multinational financial services</p> <p>Use case:</p> <p>Disaster Recovery, retention and compliance data with FileNet and ESS</p>	<p>Research Biopharmaceutical</p> <p>Use case:</p> <p>Multi site / public cloud bursting for collaboration</p>

Native Azure support

AFM to Azure feature enables placement of files or objects in an IBM Storage Scale cluster to Microsoft Azure Blob storage from IBM Storage Scale 5.1.9.

Application interface uses an AFM fileset of a preferable mode that connects to a Microsoft Azure Blob storage container.

In the background, data is seamlessly transferred through AFM between these filesets and Microsoft Azure Blob containers, which ensures high-performance and robust support for object applications.

Existing filesets replicating to Microsoft Azure Blob storage using minio gateway node can be converted to native Azure support filesets.

Online “TCT to AFM” conversion

Migration of a TCT-enabled fileset or file system is only for the tiering cloud-service-type. It is not supported for sharing Cloud-Service-Type.

Before conversion, ensure that a TCT fileset does not have any file or data structure that an AFM to cloud object storage fileset in the MU mode does not support.

AFM converts metadata of only TCT-enabled files to AFM supported file metadata because TCT supports tiering of only files and not directories.

Symlinks, which were migrated to the cloud tier, are not supported during the migration. Upload symlinks manually after the promotion by issuing the `mmafmcosctl upload` command.

All resident files, such as files that TCT did not tier to a cloud, remain only at an IBM Storage Scale location. After the promotion, user can upload these files by using the `mmafmcosctl upload` command.

Dependent fileset replication support

online creation and linking of a dependent fileset ensures that the dependent fileset with the same name is also created on the secondary for data movement.

Applications can create data on the AFM-DR fileset path and the linked dependent fileset path. AFM synchronizes the created data to the secondary fileset.

When a fileset is unlinked and deleted from a primary fileset, it is unlinked and deleted from the AFM-DR secondary fileset.

All I/Os generated in the linked dependent fileset use the same gateway node that is assigned to the AFM-DR fileset.

Replication of fileset quotas are not supported

Out-of-band meta-data population

Populate metadata without checking with home

Supported only for NFS backend

Run policy scan at home (mmafmctl getOutbandList)

Run prefetch at cache.

Quick metadata population on high latency networks

getOutbandList

Generates list-files of files or directories and their metadata, which exists on the specified path. The generated files can be used to pre-populate an AFM fileset on the cache site by using the prefetch --outband option.

Example:

```
mmafmctl Device getOutbandList -j FilesetName --path Path
```

```
mmafmctl device prefetch -j fileset --list-file /gpfs/fs1/homeFset/.mmmigrateCfg/.mmmigrate.changed.files
```

```
mmafmctl device prefetch -j fileset --delete --list-file /gpfs/fs1/homeFset/.mmmigrateCfg/.mmmigrate.deleted.files
```

5.1.9 new config options

`afmFastLookup` – Enables fast lookup on the AFM filesets without acquiring the fileset locks. Default value is “no”. Supported on all Linux distributions. This option is used on application (client) nodes. This option will be enabled by default in the future releases.

`afmLookupMapSize` - Enables lookup map at the gateway node for lookup optimization. Supported values are 0 to 32. Each lookup entry takes 16 bytes of memory. Lookup map of size $(2^{\text{afmLookupMapSize}}) * 16$ is allocated. For example, `afmLookupMapSize=20` allocates 512MB of memory ($2^{25} * 16$).

Future AFM enhancements (Intended)

AFM recovery improvements

Rules based exclusion of files

ACL mapping – fileset bucket

500+ million files support in replication modes

Patch string support

Multi-target replication.



Thank you for using
IBM Storage Scale!