

High-Performance Backup and Disaster Recovery for Petascale/Exascale Unstructured Data File Storages

Some challenges for protecting very large volumes and processing millions or billions of files:

- The daily rate of file changes for your large storage is no longer compatible with your operational backup windows
- Protection costs are much higher if you attempt to reduce service downtime and the I/O impact on your storages






Miria Reduces:


- The dependence on single vendor solutions and their budgetary impacts (Snapshots, replications, etc.)
- Costly investments and complex management for storage Disaster Recovery solutions (replication, synchronization, etc.)
- The complexity and the time required to securely restore complete storage backups either on premise or from the cloud

Key Questions from Storage Backup Managers

- How do I successfully perform full backups when my incrementals already exceed my backup window?
- I have legal obligations to back up our data, is your solution powerful and reliable enough to handle hundreds of terabytes of data with millions/billions of folders/files?
- How can I back up my data stored in the cloud or on object storage?
- How do I restore very large volumes rapidly?
- How can I restore data from one type of storage to another?
- How can I absolutely guarantee continuity of access to large volumes of critical data in the event of an incident on my storage devices?

Miria Backup Key Benefits

-  Minimize backup windows.
-  Perform regular and efficient backups with successive increments.
-  Restore simply & with full confidence a few files or petascale/exascale data sets including to another storage platform.
-  An open backup and multi-storage solution (cloud, disk, object, tape).
-  Ensure ongoing access to data in case of incident on primary storage (minimal production impact) with Miria's unique Snapstor feature.



Traditional SAN/NAS and NDMP Backup Limitations:

- A slow file by file approach with never-ending filesystem scans to identify files to backup (added, modified or deleted) reduce performances and storage availability and actual backups start well after the job launch
- Very slow massive data recovery due to the absence of parallelization
- Slow NDMP protocol which is designed for tape storage
- Locked solutions designed for identical platforms restoring
- Lack of scalability due to traditional Full plus Incrementals backup approaches poorly adapted to very large volumes (>100 TB)

Key Takeaways of the Miria Backup Solution:

Rapid detection of created, modified or deleted files

The FastScan feature⁽¹⁾ collects a list of added, modified or deleted files on file storage servers. No need for lengthy and extensive storage crawling, the backup starts very early on and uses powerful parallel processing which adapts to all files sizes.

⁽¹⁾ Now available for Lustre and EXAScaler

Incremental forever protection

Miria Backup implements an "incremental forever" technology when the target is an object or cloud storage. The initial backup is a full, followed only by incrementals. Miria reconstructs the full to restore on demand.

"Agnostic" solution enables restore to a different platform

Miria Backup collects files with their ACLs and adapts their storage in a neutral and open format. During restore, the data and associated ACLs are formatted for the target protocol and storage.

Disaster Recovery and continuity of data access

Miria enables the recovery of selected files as well as priority folders or volumes. If a major incident disrupts your high-capacity storage, Miria Snapstor is activated to support production by ensuring ongoing data access while waiting for maintenance or replacement storage.

All-in-One Data Management Platform Any workflow. Any file storage.

5 Data Services, 1 Platform



Analytics

Assess your data landscape with SAN and file system discovery, streamline file placement and reduce storage costs



Backup

Rapidly protect file storages from damage and loss. Guarantee continuity of access to valuable datasets and provide fast production recovery



Migration

Efficiently migrate data volumes and billions of files between dissimilar storage and file systems



Archiving

Free up expensive storage space on primary storages and ensure data preservation compliancy



Mobility

Move data files wherever necessary and synchronize datasets across heterogenous storages

Compatible Storages



<https://links.atempo.com/DM-compatibility>

Contact an Atempo expert



<https://links.atempo.com/DM-contact-us>

