



# SUPERMICRO® ATEMPO DATA MIGRATION SOLUTION FOR UNSTRUCTURED DATA

*Miria™, a data migration software platform for Petascale Unstructured File Sets Excels on Supermicro servers*

## TABLE OF CONTENTS

Solution Overview.....	1
Solution Features.....	1
Solution Architecture.....	2
Use Cases.....	3
Networking Options.....	4
Supermicro Miria Ready Server Platforms .....	5
Conclusion and Summary.....	5
Further Information .....	5



## Solution Overview

Unstructured data growth is accelerating at an ever-increasing pace, with petascale file sets becoming commonplace. To store these large file sets, companies have typically used heterogeneous storage platforms from multiple vendors—cloud, NAS, distributed, or parallel file systems—which contributes to data sprawl and adds management complexity. As a result, many organizations are turning to next-generation software-defined storage platforms to consolidate into a single namespace and simplify management. However, migrating unstructured file data between these disparate storage platforms is complex and risky. Solutions such as rsync or Robocopy are not scalable, require many time-consuming checks, and can lead to data loss. Organizations require a cost-effective migration solution that can quickly and reliably move hundreds of millions or billions of files from existing storage to these next-gen platforms.

Atempo’s Miria, built on Supermicro SuperServers, offers a high-performance, scalable, multi-platform solution for migrating petabytes of unstructured file data with millions/billions of folders and files.

## Solution Features

Miria is an open solution for multi-storage sources (cloud, object storage, NAS & Scale-out NAS, storages, and shared or parallel filesystems) with a wide range of target storages: disk, object, optical disk, tape, and/or cloud. It provides migration, archive, and disaster recovery capabilities for organizations of all sizes.

## Key Benefits

- Multi-Storage Platforms: Object, Cloud, SAN/NAS, distributed and parallel filesystems
- Multi-Storage Protocols: NFS, SMB/ CIFS, PFS, S3, and more
- Scalable: Adjust performance by simply adding or removing a Data Mover
- Incremental and automatic file migration between heterogeneous architectures
- Minimum impact on production and a fast cutover to the new storage
- Automatic integrity checks on all migrated files
- Simple and user-friendly web interface to supervise the migration process
- Evolutive solution: New storage can be backed up using the same solution if necessary

## Migration Usages

- Local to Cloud: Migrate data from local storage to the cloud.
- Cloud to Cloud: Migrate data from provider A to provider B to change to a new cloud provider or move data from public to private cloud.
- Cloud to Local or back to on-prem: Migrate data from cloud storage to local storage, a “reverse cloud” migration.
- Local to Local: Migration between different storage technologies or different vendors, for instance, from Lustre to GPFS or from GPFS to Lustre, from NAS to GPFS, etc.

## Solution Architecture

The infrastructure for Miria is designed to be scalable, flexible, and efficient. Miria’s infrastructure is composed of the following software components:

- Miria Server
- Miria Data Movers

### Miria Server

The Miria Server hosts a central repository to monitor the complete Miria infrastructure and to provide reporting and statistical information. This catalog stores all file information and all relevant configuration details, such as backup workflows and associated backup policies, data mover pools, and associated target pools.

The primary role of the Miria Server is to maintain a centralized catalog based on a relational database.

### Data Mover

The role of a data-mover is to simultaneously read data from a source storage and write data to the target storage at the maximum speed.

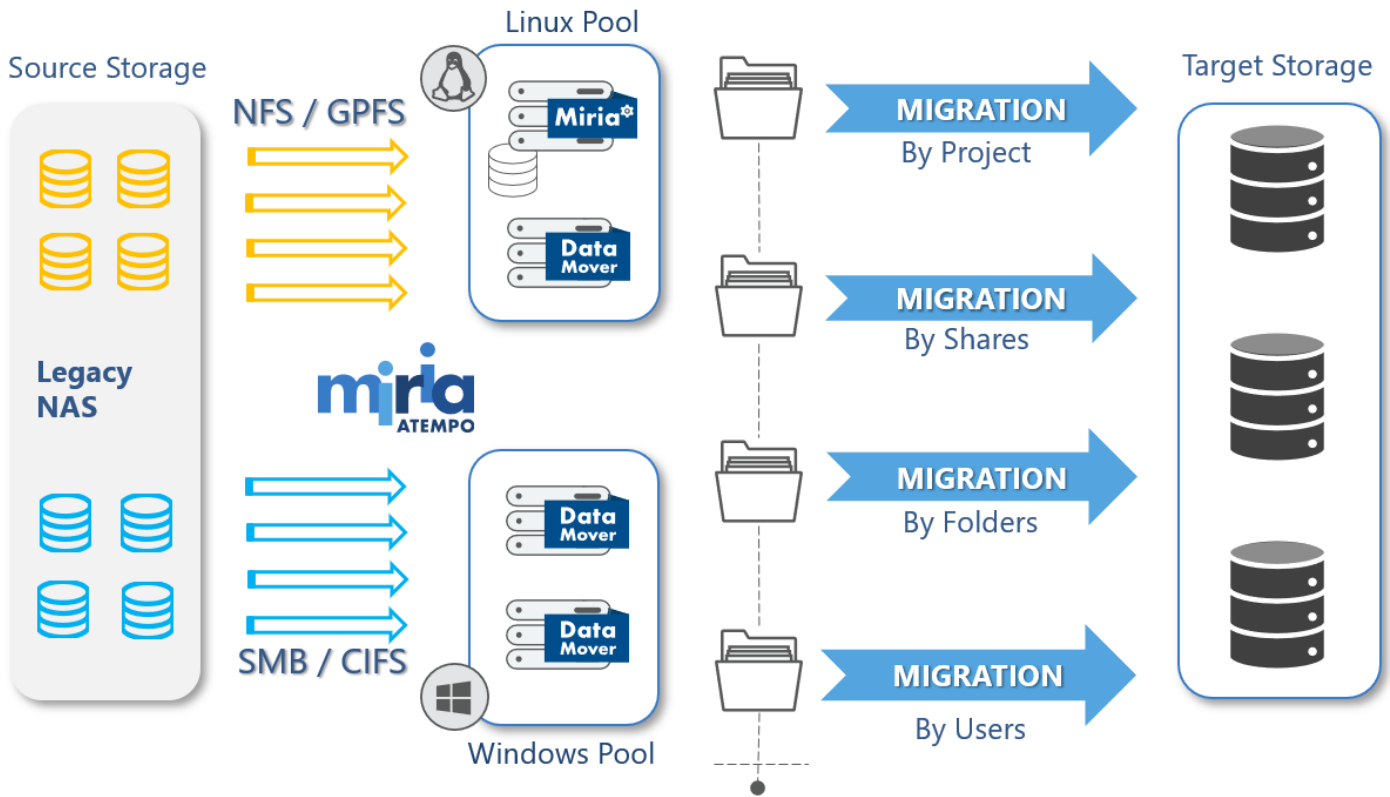


Figure 1 - Miria Architecture

## Use Cases

The Supermicro and Atempo solution can be used in many industries, and the type of servers needed depends on the volume and type of data.



### M&E

Media Editing, Archive



### Banking

High-Performance Block Storage and Remote Replication



### Finance & Insurance

Long Term Archive, Data Lakes



### Healthcare

PACS Images, Data Lakes



### EDU/Research

Physics/Astronomy, Biology/Microscopy, Departments



### Industry

Manufacturing, Analysis, Supply Chain

Figure 2 - Miria Use Cases

## Supermicro Configurations

The following configurations are provided for the Miria Server:

Solution Name	UP SuperServer for Miria for Migrating up to 400TB	UP SuperServer for Migrating up to 3 PB	CloudDC for Migrating up to 20 PB
Base Server SKU	<a href="#">SYS-510P-WTR</a>	<a href="#">SYS-110P-WTR</a>	<a href="#">SYS-120C-TR</a>
Form Factor	1U x 4x 3.5" Drive Bays (4x hybrid NVMe)	1U x 10x 2.5" Drive Bays (4x hybride NVMe)	1U x 8x 2.5" Drive Bays
System Dimensions	17.2" (437mm) x 1.7" (43mm) x 25.6" (650mm)		
CPU	Single 12-Core Intel® Xeon® Scalable processor	Single 16-Core Intel® Xeon® Scalable processor	Dual 12-Core Intel® Xeon® Scalable processors
Memory	64GB (Supports up to 2TB per node)		128GB (Supports up to 4TB per node)
Storage Controller	3908-based Hardware RAID controller		
Networking*	1x 1GbE RJ45 IPMI (onboard) 2x RJ45 10GBase-T LAN ports (onboard) 2x 10GbE SFP+ (Supermicro AOC-STGF-i2S-O)		1x 1GbE RJ45 IPMI (onboard) 2x 10GbE RJ45 & 2x 10GbE SFP+ (Supermicro AOC-ATG-I2T2SM)
Storage Devices	Boot/OS: 2x 800GB SAS SSD (RAID1) Capacity: 2x 1.92TB SAS SSD (RAID1)	Boot/OS: 2x 800GB SAS SSD (RAID1) Capacity: 4x 1.92TB SAS SSD (RAID5)	Boot/OS: 2x 800GB SAS SSD (RAID1) Capacity: 6x 1.92TB SAS SSD (RAID6)
Servers Required	One Physical Server	One Physical Server	Two Physical Servers in hot-standby mode. Hot-Standby configurations require shared storage (FC or iSCSI)

Figure 3 - Supermicro Miria Server Configurations

The following configurations are provided for the Data Mover:

Solution Name	CloudDC SuperServer for Miria Data Mover for Large Files	CloudDC SuperServer for Miria Data Mover for Small Files
Base Server SKU	SYS-120C-TNR	
Form Factor	1U x 8x 2.5" Drive Bays	
System Dimensions	17.2" (437mm) x 1.7" (43mm) x 23.5" (597mm)	
CPU	Dual 16-Core Intel® Xeon® Scalable processors	Dual 20-Core Intel® Xeon® Scalable processors
Memory	64GB (Supports up to 4TB per node)	128GB (Supports up to 4TB per node)
Storage Controller	3908-based Hardware RAID controller	
Networking*	1x 1GbE RJ45 IPMI (onboard) 2x 100GbE QSFP28 (Supermicro AOC-A100G-M2CM) 2x 10GbE SFP+ (Supermicro AOC-STGF-I2S)	
Storage Devices	Boot/OS: 2x 400GB SAS SSD (RAID1)	
Servers Required	Two Physical Servers for Migrations up to 3PB For Migrations beyond 3PB and up to 20PB requires up to 26 physical servers.	




\*See Networking Options below

Figure 4 - Supermicro Data Mover Configurations

## Networking Options

The design includes all core components to support migration over a dedicated 100GbE QSFP28 network with the addition of a 10GbE SFP+ management network. Flexibilities to accommodate a specific customer use case comes from the number of available PCIe slots in the selected server. For example, optional I/O cards for SAS, Ethernet, or Fiber Channel connectivity require a free PCIe slot in the server being considered. Additional Ethernet card options can be substituted to replace the 100GbE networking provided.

## Supermicro Miria Ready Server Platforms

		
<a href="#"><u>SYS-510P-WTR</u></a>	<a href="#"><u>SYS-110P-WTR</u></a>	<a href="#"><u>SYS-120C-TR</u></a>
Single Socket P+ (LGA-4189) 3rd Gen Intel® Xeon® Scalable processors. Up to 270W TDP.	Single Socket P+ (LGA-4189) 3rd Gen Intel® Xeon® Scalable processors. Up to 270W TDP.	Dual sockets P+ (LGA-4189) 3rd Gen Intel® Xeon® Scalable processors. Up to 270W TDP.
8 DIMMs; Supports 3DS DDR4-3200: RDIMM/LRDIMM/Intel® DCPMM	8 DIMMs; Supports 3DS DDR4-3200: RDIMM/LRDIMM/Intel® DCPMM	16 DIMMs up to 6TB 3DS ECC DDR4-3200: LRDIMM/RDIMM/Intel® DCPMM
2 PCIe 4.0 x16 (FHFL) slots 1 PCIe 4.0 x16 (LP) slot	2 PCIe 4.0 x16 (FHFL) slots 1 PCIe 4.0 x16 (LP) slot	2 PCIe 4.0 x16 (FHHL) slots 2 PCIe 4.0 x16 AIOM (OCP 3.0) 2 PCIe 3.0 x2 NVMe M.2
Intel® Ethernet Controller X550 2x 10GbE RJ45	Intel® Ethernet Controller X550 2x 10GbE RJ45	Networking via Slim AIOM
4x hot-swap 3.5" Hybrid SATA3/NVMe drive bays, SAS3 with additional SAS controller card; Onboard 1x NVMe/SATA M.2	10x hot-swap 2.5" SATA3 drive bays with 4 hybrid NVMe/SATA drive bays, SAS3 with additional SAS controller card; Onboard 1x NVMe/SATA M.2	8x 2.5" hot-swap SATA/SAS drive bays
Redundant Platinum 500W Power Supply	Redundant Platinum 750W Power Supply	Redundant Platinum 800W/860W Power Supply

## Conclusion and Summary

Supermicro is a high-performance server and storage solution provider, while Atempo is a software company specializing in data migration and management. Together, these two companies can help an organization meet its data migration needs by providing hardware and software solutions that are optimized for the task. Supermicro's high-performance servers and storage systems can be used to quickly and efficiently move large amounts of data, while Atempo's data migration and management software can be used to ensure that the data is adequately protected and managed during the migration process. By combining the strengths of these two companies, an organization can ensure that its data migration needs are met in a reliable and efficient manner.

## Further Information

[www.supermicro.com](http://www.supermicro.com)