



SUPERMICRO'S CLOUD DENSITY STORAGE

Modular Top-loading Storage Optimized for TCO for 3 Popular Use Cases for 3.5" HDDs



TABLE OF CONTENTS

Executive Summary	1
Use Cases and Tradeoffs	1
Supermicro High Density and Massive Capacity Storage	2
Server Selection	3
Data Protection	4
Scale-up Storage for Enterprise Backup & Archive	4
Scale-out Storage for Public and Private Cloud Applications	4
High Availability Storage for HPC, Parallel File Systems for AI & ML ...	6
Better, Faster, Greener	8
Summary	9

Executive Summary

The amount of data created every year continues to grow, and it is estimated that 175 zettabytes of data will be generated in 2025. The majority of the data generated will need to be stored for future use in AI, Analytics, and Business Intelligence. While there is a choice of storage technologies today, the most cost effective devices for storing massive amounts of data are hard disk drives (HDD). Innovative servers from Supermicro allow up to 90 drives to be housed in a 4U chassis. When using the latest (early 2022) high capacity drives at 20 TB each, a single server can contain 1,800 TB (or 1.8 PB) of raw capacity, providing an easy way to meet massive storage requirements.

Use Cases and Tradeoffs for Storage Applications

All industries need a large storage pool. Whether installed in a hyper-scale environment where hundreds to thousands of clients require simultaneous data access or a in smaller environments where an enormous amount of data must be stored to meet regulatory compliance, offering the best value - lowest TCO. In an effort to provide optimal configuration to meet a customer's capacity needs, Supermicro has designed

SUPERMICRO

As a global leader in high performance, high efficiency server technology and innovation, we develop and provide end-to-end green computing solutions to the data center, cloud computing, enterprise IT, big data, HPC, and embedded markets. Our Building Block Solutions® approach allows us to provide a broad range of SKUs, and enables us to build and deliver application-optimized solutions based upon your requirements.



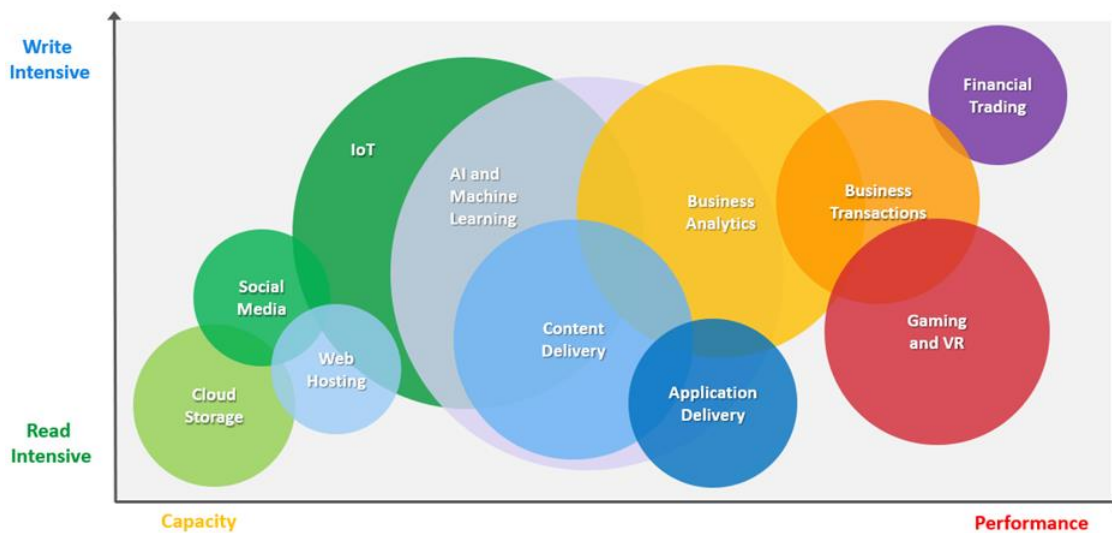
several easy to service large storage systems. These systems, known as Top-loading Storage Servers, offer configurations to meet high capacity storage requirements across application and business segments.

Applications are classified as either read or write intensive and whether capacity or performance is critical. The figure below from Storage Networking Industry Association (SNIA) maps out various applications and where each might fit with respect to the capacity and latency based performance requirements. Typically top loading systems can achieve an economic balance of cost per TB of storage with the best large throughput capabilities delivered by mechanical media at scale and are suitable for the workloads landing in the mid to lefthand side of Figure 1.

Typical Applications:

- a. Object Storage for Both Cold or Warm Tier Storage
- b. Cloud / Enterprise Block File and Object storage
- c. AI/ML Data Repository and Big Data Analytics
- d. HPC Parallel File Storage
- e. CryptoCurrency & Data mining and Backend Storage

Data Center Storage Applications Characteristics



Source : <https://www.snia.org/educational-library/tiered-storage-deployments-24g-sas-2020>

Figure 1 - Mapping of Applications to Requirements

Supermicro High Density and Massive Capacity Storage Systems

Supermicro Top Loading systems support up to 2 server nodes within 4U. Each node features two 3rd Gen Intel Xeon Scalable processors per node. The modular design allows Supermicro to offer multiple configurations utilizing a resource efficient building block approach. The two base chassis feature 60 or 90 drive bays and internal and rear bay options and can use the latest available 3.5" and 2.5" storage media updated regularly and published on the approved vendor list online.

Capacities: As of early 2022, the maximum capacity of an HDD is 20 Terabytes (TB). A 90 bay storage server (4U) can hold a total of 90 * 20 TB = 1,800 TB per system. With the ability to install 10 of these servers in a single standard rack, over 18 Petabytes (PB) can be easily accessed in one rack.

This paper identifies common applications to help users select the best Top-Loading server configuration to meet their storage requirements.

- Scale-up storage for low cost enterprise backup & archive
- Scale-out storage for Public and Private Cloud Applications
- Highly Available storage for Traditional Parallel File systems

There are several features common to all Top-loading models: modular architecture with a drive drawer design, all active components are easily serviceable.




Figure 2 - Supermicro Top-Loading Storage System - Single and Dual Node Comparisons

These servers are designed so that maintenance can be performed without special tools (Tool-less design)

Server Selection

If you are new to top-loading storage, selecting the best model starts by identifying the deployment scale and the application's best data protection method.

There are deployment at scale considerations that need to be considered. The size of the deployment and the application use case required are critical considerations in determining if Top-loading storage is the best fit for your application. As an established storage category, top-loading systems have found their way into most fortune 1000 companies, so these considerations and many others may already be established for your IT organization.

- How much data needs to be stored & how fast will it grow?
- Is there a performance requirement, and does this need to scale with capacity?
- What is the availability requirement? For example, does this need to have 99.999 availability?
- Will a customer application run on the server (converged) or with this be disaggregated - storage serving other application servers?
- Will Top-loading Servers fit in your racks?

- If this is a hardware refresh, what are some of the key features you want to keep or improve from your current deployment?

Data protection and Resiliency Considerations

All digital storage needs protection. This includes protection from corruption, device failure, and protection from malicious attacks. The level of protection can significantly impact both the initial cost and the operational expense of a deployment, which is outside the scope of this paper (CAPEX and OPEX). There are two basic ways data can be protected at a high level. First, at the HDD Hardware level, RAID can be used to protect data from HDD/SSD device failure. This type of data protection (scale-up) has its advantages at local/small scale but typically yields to software-defined storage (SDS) methods when a deployment reaches a larger scale.

SDS approaches were developed from the start to solve the many problems associated with scale. SDS allows multiple systems to be pooled together, presenting a large federation of resources abstracted and provisioned dynamically to meet each client's requirement. Data protection schemes are usually separated from the physical hardware, and access protocols and storage resources (capacity and performance) can be driven by IT policy as opposed to physical limitations. Because of the extreme level of hardware abstraction, SDS architectures are typically deployed using Scale-out topologies, and there are hundreds of SDS vendors today, both open source and commercial.

Scale-out Storage using Top-loading storage servers tends to be disaggregated, meaning the server role is only to deliver storage to other computer clients or compute clusters.

Scale-up Storage for Enterprise Backup & Archive

In a typical backup application, the deployment must meet capacity and Cost per GB targets and still be able to serve data fast enough to restore critical services in the event of failure. In today's large enterprises, these services are typically delivered with On-premises hardware and utilize cloud-based services providing a balanced mix of resiliency and cost effectiveness that CIO/CTO's demand (Hybrid Cloud).

In Remote Office/Branch Office (ROBO) environments, selecting top-loading systems featuring a single high capacity node and utilizing built-in RAID offer an advantage over many software-defined strategies. The main factor is that data protection is built into the server, and the backup application can then be run as a converged application on the Top-loading server itself. For example, the high spindle count of a 60 Bay server with a hardware RAID controller offers more than enough throughput performance to backup and restore local servers at the network speeds deployed in these environments. In addition, the powerful dual processor server platform has more than enough horsepower to replicate data to S3 asynchronously to the cloud for complete site level disaster recovery.

Scale-out Storage for Public and Private Cloud Applications

Scale-out Storage offers an organization the ability to grow dynamically with their business requirements. With the growing dependency on the public cloud in recent years, many enterprises have wanted to bring some portion of their IT resources back on-premises. The cloud storage category of S3 / Object Storage is probably the most common cloud segment quickly repatriated by enterprises today. For this segment, top-loading servers are the go-to solution to accomplish this goal.

The quick reaction is to select the most significant cheapest server you can find to throw at this problem, but in most cases, this is a mistake. Looking at the best practices for scale-out storage (provided by most SDS vendors), you will discover that most recommend a minimum of 6-8 nodes to build a resiliency cluster, the more nodes, the better. Fewer nodes may be used, but the data protection and availability may suffer. Protection may only be at the drive level, and in some cases, with a small cluster, the storage may go into read-only mode until the failed or missing component is replaced. This is not the best choice if trying to match cloud-level availability.

Another major factor to consider when selecting a server is the ratio of HDD per node. Many SDS vendors will require a maximum number of drives per node, depending on the cluster size. Because of these two factors, Supermicro offers dual node configurations on the Top-loading servers. Dual nodes allow for extremely high density configurations to be built with a 4U footprint, and at the same time, the ability to meet the maximum design requirements of the software being used.

Scale out Storage utilizing Erasure code

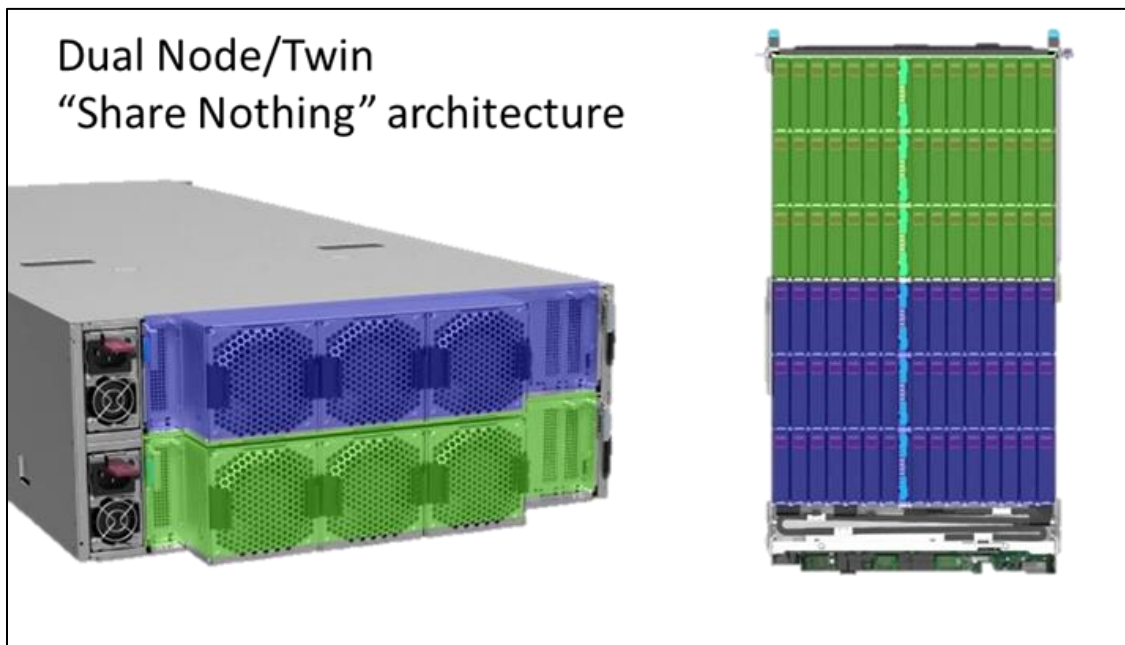


Figure 3 - Dual Node Share Nothing Architecture

The example below shows a cluster using the dual node configuration and typical network layout (6 nodes in a 12RU configuration)

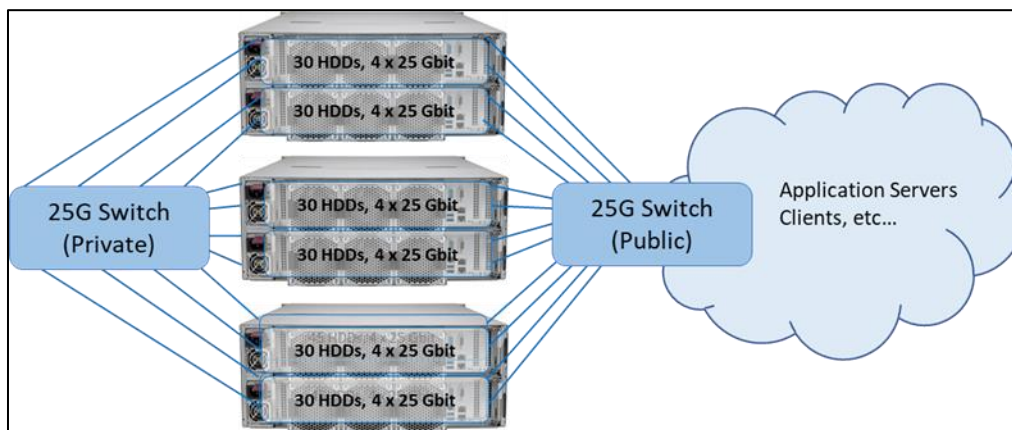


Figure 4 - Example Configuration of a six node Storage Cluster using a Dual node 60-bay

With the dual node 60 bay chassis, an HDD:Node ratio of 30:1 will meet the design maximum for a fully supported SDS architecture using Red Hat Ceph Storage Subscription. In the case of Red Hat, a production deployment would start at eight nodes, with up to 36 HDD per node being their published maximum. Typically the ratio can be relaxed at Red Hat’s discretion if they determine the cluster size is large enough to absorb device failures to meet customers' expected SLAs (>100 nodes, for example).

Other SDS/software vendors might have a six node minimum without a set drive maximum. In these cases, we can do a quick paper comparison/exercise of a six node cluster using the 20TB drive for raw capacity and a rough estimate for usable capacity. Usable capacity will vary depending on each software vendor's Erasure Code method (EC). In this exercise, we can assume each node has the same compute and networking configuration. In practice, the server’s CPU/Memory and Networking configurations would be tuned to support the use case and software recommendations and requirements.

Unit count: SSG Model	Server Type	Number of HDD	Ratio	Raw Capacity	Usable Capacity (EC 4+2)	Rack Units Required
QTY3: SSG-640SP-DE1CR60	Dual Node 60	180	30:1	3.6 PB	2.4 PB	12
QTY3: SSG-640SP-DE1CR90	Dual Node 90	270	45:1	5.4 PB	3.6 PB	12
QTY6: SSG-640SP-E1CR60	Single node 60	360	60:1	7.1 PB	4.8 PB	24
QTY6: SSG-640SP-E1CR90	Single node 90	540	90:1	10.8 PB	7.2 PB	24

Table 1 - Supermicro Top-Loading Models

In the 6 Node comparison/exercise, the capacity of the solutions ranges from 3.6 Petabytes to 10.8 Petabytes of raw capacity. Still, with six nodes being the standard metric, the resiliency of the cluster will be relatively similar, sharing the same node based failure domain. In a system level benchmark, the aggregate sequential drive throughput performance is higher in a node with the most drives. When an SDS architecture is deployed, dual node systems offer more IOP per GB of Capacity due to the compute per HDD ratio. Networking bandwidth and total network port count would be equal across all 4 clusters, so in this example, dual node models offer greater performance per Rack Unit. Software efficiency also becomes a key

consideration on a small cluster with Node ratios like 90:1 in the highest capacity example. They are making the single node models a better choice for larger clusters.

In many cases selecting the dual node model with 45:1 offers the same storage density with better overall performance per stored TB. For extremely large clusters with 100+ nodes, the 90:1 ratio is more tenable because the additional nodes help distribute and balance the client load. The hyper-scale cloud vendors commonly use top-loading storage servers, as they rely on a massive scale to deliver their guaranteed SLAs.

High Availability Storage for HPC, Parallel File Systems for AI & ML

Long before the public cloud became the largest storage consumer, Top-loading storage arrays were used In HPC environments to store the massive datasets needed for large scientific projects and commercial oil and gas exploration. They typically attached through SANs to compute host (OSS) that stripped all the arrays together to present a large parallel file system (PFS) to compute clusters. Data stored on the PFS would be shared by hundreds-to-thousands of client computers, allowing for massive supercomputing power. With so many clients depending on the storage infrastructure, No Single Point of Failure (NSPOF) architectures became the de facto standard. This means fully redundant designs were the only choice. Unfortunately, all this redundant infrastructure started to add up to costly environments =SANs+RAID Arrays+ OSS computers. Then SDS came to the rescue; with Linux+ZFS running on a Dual/Redundant Top-loading storage server, all the expensive Storage Area Networking hardware and software layers could be converged on a large Dual Node server/ with redundant access to the same HDDs (dual port SAS). Eliminating the SAN and other proprietary storage Hardware and combining all the data across a single fast fabric instead of multiple vendor fabrics provides better overall manageability and economy of scale.

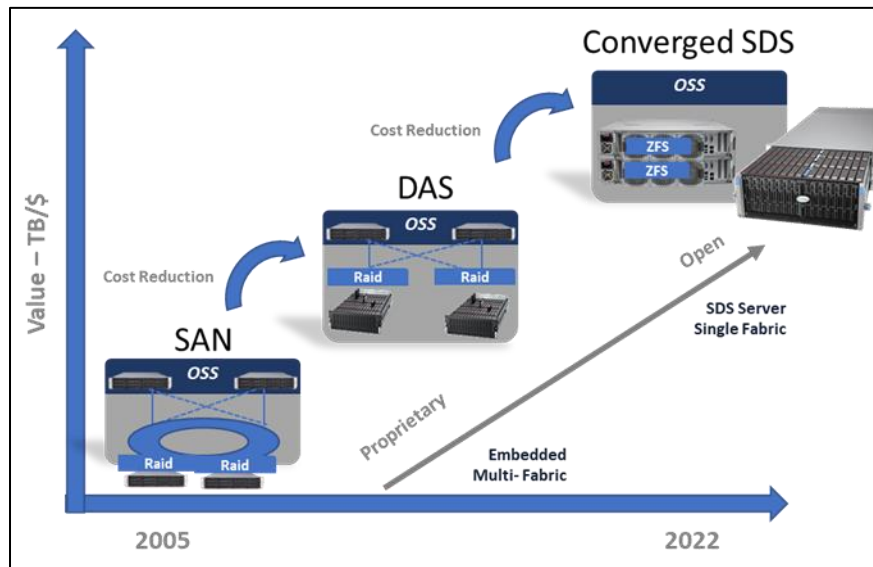


Figure 5 - Storage Architectures Over Time



Figure 6 - 90 Bay Top-Loading Server

The dual node high availability Top-loading models offer the NSPOF design perfect for PFS environments while offering a 100% software-defined architecture, giving complete control to developers and organizations that dominate the HPC and AI/ML domain. Furthermore, as more All flash HPC storage solutions enter the market, high capacity Top-loading systems retain their critical role for long term storage of the massive datasets being analyzed today.

Supermicro SuperStorage Systems are Simply Better, Faster, and Greener.

Better – The 60 and 90 bay system features a modular design allowing easy service for all active components. The Field Replaceable Units (FRU) is the same for the different top-loading configurations, with each personality configured at the factory. In addition, the internal drive-drawer design eliminates the need for a cable management arm and allows for quick and easy disk drive replacement.

Faster – These systems are all based on single or dual Intel 3rd Gen Intel Xeon Scalable CPUs. With PCI-E 4.0, bandwidth is twice the performance of the previous generation of systems. Also, the 3rd Gen Intel Xeon Scalable processors have more cores, and a faster clock rate than previous generations and thus can pass the data to the storage devices faster than before.

Greener – Higher density means less real estate that has to be devoted to certain IT investments. For example, with up to 1.62 PB in a single 4U enclosure, less square footage is needed for cloud-scale storage. In addition, the 3rd Gen Intel Xeon Scalable processors produce more work per watt than previous generations of Intel processors, reducing the power needed for a server, and thus lessening greenhouse gas emissions.

Supermicro SuperStorage Top-Loading Specs:

- 60/90 Hot-Swap 3.5" SAS3/SATA3 bays
- 3rd Gen Intel® Xeon® Scalable processors
- Dual Socket LG-4189 (Socket P+) supported
- TDP up to 205W CPU; up to 11.2GT/s

Key Applications:

- Immutable Storage - Data Protection

- Content Repositories Media & Image Archives
- Financial Services & Healthcare Compliance
- Telco & Cloud Service Providers
- HPC and AI/ML Workloads
- Big Data & Analytics, Data Lake

Outstanding Features:

- Single Node with High Density in 4U rack space

Summary

Supermicro SuperStorage top-loading systems give customers a state-of-the-art system with the highest density available today. In addition, the flexibility of the Supermicro SuperStorage top-loading systems allows organizations to customize this system based on their user demands and IT infrastructure.

New storage capacities per disk drive continue to increase, and innovative methods that transfer data from the disk drive to the CPU continue to evolve. Supermicro, known for its first to market advantage, continues to work with leading vendors to deliver the most powerful and cost-effective servers in the market.

Additional Information

https://www.supermicro.com/manuals/brochure/Brochure_Storage_systems.pdf

<https://www.supermicro.com/en/products/top-loading-storage>

<https://www.youtube.com/watch?v=6yp7sHGfCYw>