WHITE PAPER

Supermicro Software-Defined Storage

Store your most important assets on flexible, scalable, and cost-effective storage solutions that overcome the limitations of traditional storage appliances

Regardless of your organization's size, data is your most important asset, and needs to be treated as such. With a wide range of storage choices available, Supermicro demystifies the process of deciding what storage architecture is appropriate to help your business with its storage needs today and well into the future. We build scalable storage platforms and pair them with the industry's best software-defined storage software to meet all of your storage needs, from those of a small transactional processing system to the largest data lakes.

Data Is Your Most Important Asset

At this point in the evolution of modern business, everyone agrees that data is their most important asset. Transactional data is created when customers purchase what a company sells. Web-tracking data can identify which pages are most frequently visited. Operational data can track factory outputs and automation efficiency. All of this data enjoys a second life when it is used to inform next-generation businesses that want to gain every insight they can from their data—whether to predict customer behavior and change how products are presented, or to predict when a machine might fail and intervene with preventative maintenance or replacement. When artificial intelligence models are used to provide insights, they need to be trained with the largest amount of data possible. And there is the challenge: forward-looking organizations need to retain all of their data for possible future use, and they need a strategy for storage systems that are up to the task.

Supermicro offers a range of storage servers powered by x86-architecture processors that can meet the range of different trade-offs that the type of data dictates. Some of our servers use the fastest and most-dense EDSFF E3.S, E1.S, and U.2 NVMe storage devices to deliver data to applications with low latency and high bandwidth. Others use drawers full of disk drives to maintain massive amounts of data for the long term when performance requirements are comparatively low.

Supermicro partners with the leading software-defined storage vendors to bring our products to life in your organization, making your data available with the access methods, the quantities, the performance, and the availability your organization requires. The solutions we have tested and validated range from supporting a handful of terabytes of data to tens of petabytes. You can start small with a single storage server and create a do-it-yourself solution, or you can work with us or our partners to have a rack full of storage servers with the software and networking necessary for an enterprise-grade, high-capacity, and highly scalable solution.

Benefits of Software-Defined Storage

Software-defined storage decouples your choice of storage software from hardware, opening the door to a wide range of possibilities. First, choose the storage software that provides the features most important to your organization. Second, choose storage servers that deliver the balance of capacity, performance, and economy that best meets business needs. The options are practically limitless. Organizations can:

- Implement block storage on a single server to support your transactional databases
- Implement an object storage cluster to support persistent data for your cloud-native applications
- Store machine-learning training data on a cluster of servers with the fastest NVMe drives that can feed GPU accelerators at the speed at which they consume data
- Establish a data lake with a cluster of high-capacity servers with up to 90 disks each, so you don't ever have to delete information that could hold value in the future

The economics of this approach are compelling: because storage servers from vendors like Supermicro are based on x86-architecture servers, you can take advantage of the lower prices of off-the-shelf hardware compared to purpose-built storage appliances. You can tune the hardware to deliver the performance characteristics that are important to you. For example, you have the flexibility to equip the servers with the I/O you need, ranging from 400 Gb/s InfiniBand all the way down to Gigabit Ethernet.

Traditional storage appliances scale until you reach the maximum number of drives or you run out of storage controller capacity. With storage clusters built with Supermicro storage servers, each increment of storage capacity scale brings along an increment of storage controller scale in the form of the server CPUs. So storage capacity can scale as quickly as an organization's data grows, and it can take advantage of the work that Supermicro does to

test and validate our hardware with our partner software solutions. As faster CPU and storage technology becomes available, some solutions allow you to seamlessly integrate new storage servers into a cluster.

Today's software-defined storage solutions are at the leading edge, and there are choices that support all of the use cases you can imagine. They can implement whatever level of availability you wish, based on technologies ranging from simple replication to sophisticated erasure coding.

As you work with Supermicro or our partners to determine the solution that works best for you, it's important to decide what type of storage you need, and what scalability characteristics it needs to have. Once these architecture-level decisions are made, you can begin to choose the specific storage servers and software that best support the amount of data you have today, with the capability to grow with your anticipated needs.

Block, File, and Object Storage

Software-defined storage software typically offers a combination of three access mechanisms: block, file, and object. The type of access your applications need, combined with the amount of data you have, influence your overall storage architecture.

- ***Block storage*** provides a sequential series of disk blocks to an application that is in charge of how it structures data within them. Some transactional database management systems, high-frequency trading systems, and media production applications perform best if they can read and write directly to and from the storage devices for top performance without the overhead of a file system structure. In most uses of block storage, a file system is installed to provide a hierarchical file system to applications. Block storage is most often used in small environments such as small businesses, remote, and branch offices.
- ***File storage*** provides a hierarchical file system and is the most familiar mechanism. A file system can be overlaid onto block storage, or it can be provided by software-defined storage software that manages how and where file contents are stored, in addition to managing redundancy for reliability. Applications that typically use file-based storage include high-performance computing, AI and ML, content delivery, and archival storage.

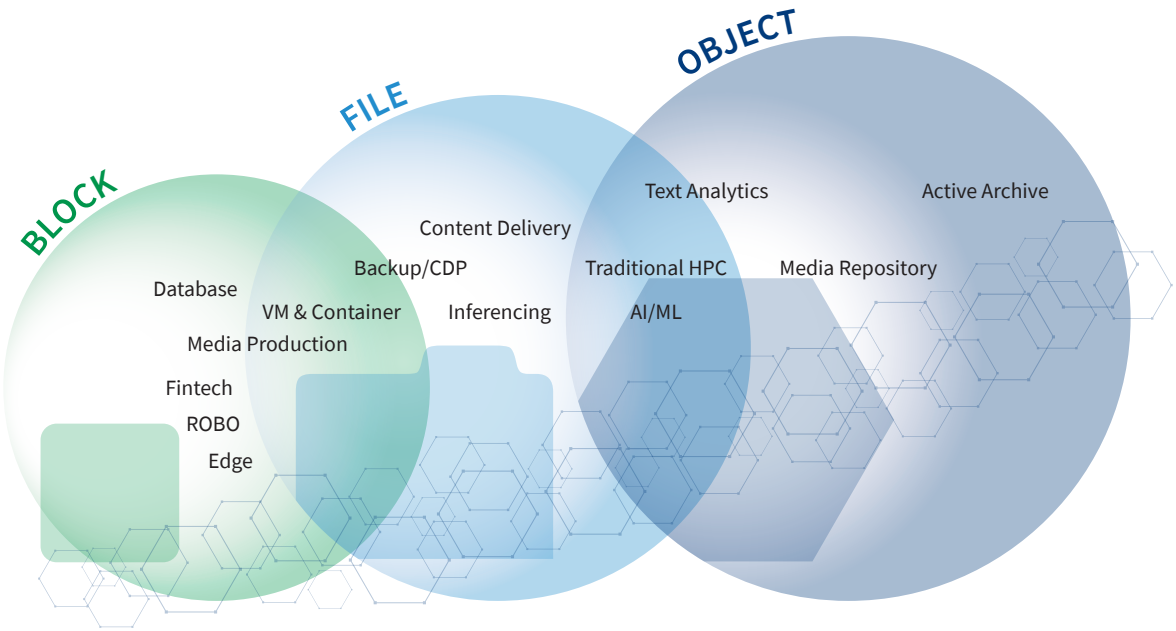

Figure 1. File, block, and object-based storage are the typical access methods, with object storage the most scalable of the three

SUPERMICRO PETASCALE STORAGE SERVERS

Supermicro ASG-1115S-NE316R

Supermicro SSG-121E-NES24R

Supermicro ASG-2115S-NE332R

The Petascale storage servers featured in this paper open the door to the future with high-density and high-performance Enterprise Datacenter Standard Form Factor (EDSFF) NVMe drives, with 16, 24, and 32 drives, respectively. Our 16- and 32-drive servers are powered by a single 4th Gen AMD EPYC™ processor with up to 128 cores for high throughput. These servers are populated with EDSFF E3.S drives. Our 24-drive system is a 1U server populated with ESDFF E1.S drives for high density. Each of these servers support PCIe 5.0 x16 network interfaces for connectivity with the fastest interfaces available, for example the NVIDIA ConnectX®-7 card with 400 Gb/s of InfiniBand bandwidth.

- ***Object storage*** is the modern approach used by cloud-native applications. It is used where a massive number of objects need to be stored, and/or objects of massive size, such as for backups of other storage. Without the structure of a hierarchical file system, objects are stored and retrieved with unique identifiers. Whereas file storage saves metadata outside of the actual file (such as file type, creation date), metadata is stored inside of an object. Applications using object storage includes containerized applications, archives, and media repositories.

Block storage is usually used for the smallest data set sizes, with object storage used for the largest. As you move up the ladder from block to file to object storage, there is less contention for locks, so larger numbers of clients can use the storage simultaneously.

Scale Up vs. Scale Out Storage

From the simplest file system to the most feature-rich software-defined storage products, they all rely on underlying storage that grows by scaling up or scaling out—or both (Figure 2).

- ***Scale-up storage*** is typically used in smaller-scale environments or specific applications where the number of clients using the data is low. It uses one or two x86-architecture servers with optional RAID hardware or software that adds redundancy to the data. This enables the server to withstand the loss of a single drive. RAID is also the primary mechanism that aggregates multiple disks into a block-storage device. Such a RAID configuration can withstand disk failures, but server-level problems, can interrupt storage availability. You can mitigate these issues with a two-node, dual-port architecture, or by choosing servers with dual power supplies, ECC memory, and redundant network interfaces. Any architecture can grow by adding drives up to the capacity of the storage server, and from there by adding a just-a-bunch-of-disks (JBOD) device that enables more drives to be accessed by the server. At some point in the growth cycle, the server CPU or RAID controller becomes a bottleneck, and to continue growing storage means adding more scale-up devices. This is problematic because now you have multiple name spaces on two or more distinct devices.
- ***Scale-out storage*** offers virtually limitless possibilities. Using a storage server as a unit of increment, it uses a scale-out model that increases storage and controller capacity in unison. Software-defined storage software uses the servers to build a distributed, scale-out storage system

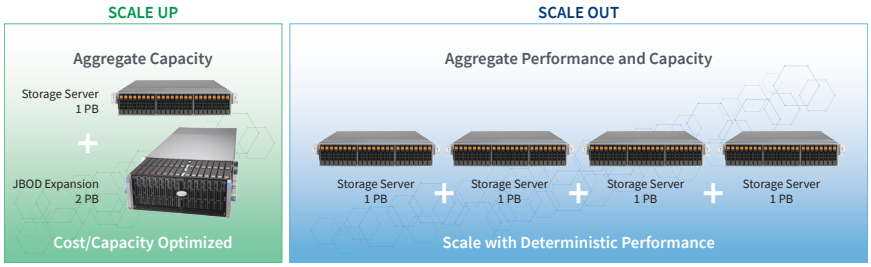

Figure 2. Scale-up storage aggregates capacity into one system, while scale-out storage aggregates capacity from multiple servers

SUPERMICRO BIGTWIN® AND GRANDTWIN® MULTI-NODE SYSTEMS

AS -211GT-HNTR GrandTwin

AS -2124BT-HNTR BigTwin

Supermicro twin servers are multinode systems that support four nodes in a compact 2U rackmount form factor with redundant, shared power supplies for reliability and up to two AMD EPYC processors per node for utmost performance. The GrandTwin can be configured with nodes having four or six 2.5" NVMe or SATA3 drives per node. The BigTwin system can support nodes having six 2.5" NVMe or SATA3 drives or three 3.5" SATA3 hard-disk drives.

HIGH AVAILABILITY DUAL PORT STORAGE SERVER

SSG-221E-DN2R24R

The 2U Storage SuperServer is also a multinode sever but with two nodes each of which share access to the 24 front-panel-accessible U.2 NVMe drives. This versatile server can be used for redundant scale-up storage.

that lays out data across the cluster so that, for example, even blocks of a singe file might be distributed across multiple servers so the workload is evenly distributed. Starting with four nodes, clusters can grow to be very large, and they can also serve a large number of clients because there is less contention for shared resources. Data is replicated or otherwise protected through techniques such as erasure coding so that a single disk or server failure cannot interrupt services. Because the storage software manages data protection on its own, RAID controllers are often not used in scale-out storage. When architecting a scale-out storage cluster, you can finely tune the relationship between IOPS performance and capacity to meet the price point that business needs dictate.

- ***Hyperconverged clusters*** are basically scale-out storage clusters that also host applications in virtualized environments. This technology offers an easy on-ramp to to using scale-out clusters with the capability to grow both storage and processing capacity together. Hyperconverged solutions are common in small businesses as they can scale easily as they grow into larger operations.

Most of these approaches to software-defined storage provide features that are expected of enterprise storage systems, such as integration with backup software, the capability to burst into the cloud, asynchronous and remote replication, and anti-ransomware features such as immutable storage,.

Software-Defined Storage with Supermicro

Supermicro has a history of industry innovation in the realm of storage servers. We were the first to offer a server with 36 disk drive bays, while other vendors were accomplishing higher capacities only with external JBOD devices. We pioneered the concept of running applications on storage servers, a precursor to today's hyperconverged clusters. Today, we offer a wide range of storage servers that address different needs along the capacity vs. performance spectrum. Our storage-optimized systems include:

- ***Petascale systems*** that host the latest high-performance EDSFF E1.S or E3.S drives with options for 16, 24, or 32 drives per system
- ***BigTwin and GrandTwin multinode systems*** that support a cluster in a box with multiple server nodes each having disk drives dedicated to them.
- ***High Availability Dual Port Storage Servers*** with two nodes dual ported to 24 front-panel-accessible drives; this configuration can support highly available storage that continues to be accessible in the event of a single-node failure.
- ***Hyper systems*** that deliver flexibility, scalability, and serviceability in demanding environments, with choices ranging from 8 to 24 drives per rack-mount chassis.
- ***Simply Double servers*** that support 12 3.5" front-panel drives and 12 more midplane-accessible drives for a total of 24 hot-swappable, large-form-factor drives
- ***Top-loading systems*** that support 60 or 90 high-capacity, 3.5" drives; these can be purchased as single-node systems with the full capacity through a single integrated server node, or as multinode systems with the disk drive capacity split between two integrated nodes, each with either 30 or 45 drives

HYPER 1U

AS -1125HS-TNR

The 1U Hyper server is powered by AMD EPYC processors and is a flexible platform to host up to12 NVMe, SAS, and SATA3 drives. The platform is optimized for performance and reliability, including a tool-less physical design that makes it quick and easy to repair. These factors contribute to the Hyper 1U server being the foundation of economical scale-out storage solutions.

These systems form the basis for a set of scale-up and scale-out solutions for small, medium, large, and extra-large storage requirements that are discussed in the following sections. These are tested and validated solutions that you can put to work in your data center, small business, or edge environment with the support of Supermicro or our partners. Table 1 on the following page summarizes the options for tested combinations of servers and software-defined storage products. A summary of our software partners is provided in Table 2.

Small Solutions (<500 TB)

What we consider small storage solutions may be in support of small organizations, or to meet the needs of latency-sensitive, business-critical applications such as online transaction processing or high-frequency trading. These needs often can be supported by a single or a fully redundant scale-up solution either with a set of internal disk drives with or without additional storage through JBODs. Since these are small environments with a limited set of users, hardware RAID can be used to create a single logical volume that can be offered as block-based storage or, file-based storage with a resilient file system such as ZFS. GRAID can support RAID plus NVMe over Fabrics, which can speed data from the drive to the client memory with no CPU intervention. OSNexus ZFS software provides a supported solution of this open-source file system for businesses wanting more of a plug-and-play solution.

If your organization is taking an S3 approach, Scality ARTESCA provides object storage with enterprise-level features including asynchronous replication, backup integration, and immutable storage that serves as anti-ransomware protection. All of this can be accomplished with a single scale-up storage server, and as the business grows, it can expand to support scale-out solutions with multiple nodes and distributed storage.

An economical way to start with scale-out storage when compute resources are also being updated is in the form of hyperconverged solutions such as VMware vSAN. While the cost of entry is the purchase of multiple servers, vSAN solutions can grow well in to the medium range of 500 TB to 1 PB. This is a solution to consider if your growth rate is expected to be high.

Server choices for this size solution includes Supermicro high-performance Petascale storage systems with low-latency EDSFF E.3S and E.1S storage. Petascale solutions are targeted to areas where high performance is important: in this size category the database management systems and financial technology. Our High Availability Dual Port server offers two nodes in a single chassis, each with access to 24 drives. Configured with hardware RAID, this storage server can provide a highly available solution. The Hyper 1U server provides a good amount of storage at an economical price, and the Supermicro Simply Double server, based on 3.5" disk drives, provides a lower-performance yet high-capacity solution.

Medium Solutions (500 TB – 2 PB)

Medium solutions need to store data for a wide range of enterprise workloads, from virtualized and cloud-native applications to content delivery, AI/ML inferencing, and high-performance computing. These workloads need high

Table 1. Summary of solution sizes and applicable storage servers

		Storage Server Relative Capacity											
		Small					Medium				Large		
		BigTwin/GrandTwin	Hyper 1U	Petascale	High Availability Dual Port	Enterprise 36 Bay	Petascale	Simply Double	Top Loading Dual	Top Loading Dual	Petascale	Top Loading Single	Top Loading Single
	Rack Units		1U	1U	2U	4U	1U	2U	4U	4U	2U	4U	4U
	Nodes	4	1	1	2	1	1	1	2	2	1	1	1
	Drives per Node	4	12	16	24	36	24	24	30	45	32	60	90
	Format	U.2	U.2	E3.S	U.2	3.5"	E1.S	3.5"	3.5"	3.5"	E3.S	3.5"	3.5"
	Type	NVMe	NVMe	NVMe	NVMe	HDD	NVMe	HDD	HDD	HDD	NVMe	HDD	HDD
Solution Capacity	Dual Port Support				Yes								
Small (<500 TB)			✔	✔	✔			✔					
Medium (500 TB - 1 PB)		✔	✔		✔			✔			✔		
Large (1-10 PB)			✔						✔	✔	✔		
CloudScale (XXL) 10+ PB			✔								✔	✔	✔

Storage Software	Block	File	Object
GRAID (Hardware RAID)	✔		
GRAID + ZFS		✔	
OSNexus with ZFS	✔	✔	
VMware vSAN	✔	✔	
Scality ARTESCA			✔
Quantum ActiveScale			✔
Cloudian			✔
Lightbits	✔		
OSNexus with Ceph	✔	✔	✔
Qumulo		✔	✔
Scality RING		✔	✔
VMware vSAN	✔	✔	
Quantum ActiveScale			✔
Cloudian		✔	
Hammerspace		✔	
OSNexus with Ceph	✔	✔	✔
Qumulo		✔	✔
Scality RING		✔	✔
VAST	✔	✔	✔
VMware vSAN	✔	✔	
WEKA		✔	✔
Quantum ActiveScale			✔
Cloudian			✔
Hammerspace	✔	✔	✔
OSNexus with Ceph		✔	✔
Scality RING		✔	✔
VAST	✔	✔	✔
WEKA		✔	✔

Table 2. Supermicro software-defined storage partners

CLOUDIAN®	Graid Technology Inc.	HAMMERSPACE DATA IN MOTION
Cloudian is an enterprise-class, S3-compatible object storage system that provides a combination of a data lake and data warehouse dubbed a "Data Lakehouse." It provides immutable storage, safe from hacker encryption, and it is built to support data sovereignty regulations.	***SupremeRAID*** is a software-defined solution that is deployed on a GPU and uses unique out-of-path RAID technology. Data travels directly from CPU to storage to deliver maximum SSD performance, comprehensive data protection, and unmatched flexibility from the cloud to the desktop.	***Hammerspace*** delivers a global data environment that spans across data centers, AWS, Azure, and Google cloud infrastructure. With origins in Linux, NFS, open standards, flash, and deep file system and data management technology leadership, Hammerspace connects users with their data and applications, anywhere.
lightbits	OSNEXUS STORAGE MADE SIMPLE.	Quantum®
Lightbits offers software-defined, disaggregated block storage for private, edge, and public clouds. This high-performance, low-latency storage system accelerates cloud-native applications and reduces TCO with increased flash endurance, standard networking, and increased utilization.	***OSNexus*** is a software-defined storage solution that supports scale-out object storage solutions and provides a single pane of glass management system across data centers and sites worldwide. It is deployed using a variety of Supermicro hybrid, all-flash and all disk storage servers.	***Quantum ActiveScale*** provides S3-compatible object storage that is durable, secure, and supports archiving and long-term retention of cold data. It provides unlimited scale, automatic scaling, high performance, and always-available data access.
Qumulo	SCALITY	VAST
Qumulo offers a Scale Anywhere platform that makes it easy for hybrid enterprises to efficiently store and manage file and object data at the edge, in the core, and in the cloud. It can power exabyte-scale, geographically-distributed, multi-site, and multi-cloud data storage deployments.	***Scality*** is enterprise-class software-defined storage with distributed file and S3 object storage with the capability to tier into the cloud. It provides hacker-resilient storage to help protect data from ransomware attacks.	***VAST*** offers service providers, hyperscale technology companies, and large, data-centric enterprises with a highly performant and resilient AI solution. It helps organizations automatically store and apply context to unstructured data, accelerate data analytics, and eliminate data silos.
vmware® by Broadcom	WEKA	
VMware vSAN is software-defined, enterprise storage solution that supports HCI. Part of the ESXi hypervisor, vSAN aggregates local or direct-attached data storage devices to create a single storage pool that is shared across all hosts in a vSAN cluster.	The ***WEKA Data Platform*** is a distributed parallel file system that stores, processes, and manages data virtually anywhere with cloud simplicity and on-premises performance. It supports clusters with thousands of nodes and integrated tiering to deliver scalable performance for next-generation workloads like AI and HPC.	

TOP-LOADING SYSTEMS

We offer a variety of 4U top-loading storage servers for the ultimate in scale-up storage, or for massive capacity in large or extra-large scale-out configurations. Each system boasts 60 or 90 3.5" hard-disk drives that are connected to a single node, or which are split between two nodes. This enables flexible configurations that optimize performance. Choose two-node systems when you need the highest performance or bandwidth.

SIMPLY DOUBLE SYSTEMS

Supermicro Simply Double systems offer 24 3.5" hard-disk drives in a compact 2U form factor with 12 front-panel-accessible drives, and 12 more accessible in a midplane. These servers offer a choice of AMD EPYC or Intel® Xeon® processors

performance, but also more capacity than a single storage server can provide, so scale-out clusters are often the most appropriate solutions. Virtualized environments need access to block and file storage. Most enterprise workloads, AI/ML and HPC use file storage, while cloud-native applications and content delivery often use object stores. Hyperconverged environments fill the need for block and file storage, and a medium-sized cluster might have started as a small solution having scaled into the medium range. Indeed, growth is handled in medium solutions by increasing the number of nodes, the speed of storage in each node, and the number of network interfaces per node.

The servers we recommend for medium solutions range from four to 32 drives per node. Our 2U Petascale server delivers the highest performance in this category, and the 2U choice provides higher capacity compared to the 1U version recommended for small environments. Stepping down to systems with 24 drives per node, the High Availability Dual Port server supports 24 U.2 NVMe drives, and they are dual ported between two nodes in the system, adding a second path to each drive. The Simply Double system supports 24 3.5" hard-disk drives, delivering high capacity with lower cost and performance. The Hyper 1U system holds 12 U.2 NVMe drives. Finally, for environments requiring a high amount of parallelism, the BigTwin and GrandTwin systems provide four nodes per chassis however they are limited to four drives per node.

When architecting scale-out storage for medium-sized applications, the most important choice you'll make is in choosing the increment of scale. This is where application requirements should dictate the storage architecture, and you'll need to consider how the server's CPU power and I/O bandwidth matches your need for storage capacity. Smaller increments of scale—the choice of storage servers—are balanced for more controller (CPU) and I/O performance per storage capacity. Choose toward this end of the spectrum if performance is an overriding requirement. The Simply Double server balances CPU, I/O, and storage capacity, and is a favorite for servers having the economy of 3.5" drives.

It's easy to see how making the wrong choice can lead to problems. Choose too large of a scaling increment and you can end up with the same problem as legacy storage systems: not enough controller capacity for the amount of storage in the appliance. Choose too small of an increment and you will incur the cost of more servers than you need to meet your performance requirements. That's why choosing Supermicro storage servers gives you piece of mind. You gain access to a wide spectrum of storage server choices backed up by the expertise to architect and size your clusters. Between us and our partners, you can purchase complete integrated solutions delivered to you in the rack.

As an example, consider the voracious data appetite of our 8U GPU servers as described in the white paper *Accelerating AI Data Pipelines* (Figure 3). In AI/ML applications, these can host eight GPUs per server, each of which can be connected to a dedicated 400 Gbps Ethernet or InfiniBand connection. To provide an all-flash tier capable of providing sufficient bandwidth to feed ML training, we architected a 10-server, large-sized cluster that uses even a smaller Petascale server than is typically used in this size cluster—the 1U system with 16 EDSFF E.3S drives. This choice was made in order to match the bandwidth requirements of the application with that of the storage system.

AMD EPYC™ PROCESSORS

If organizations aren't careful, massive amounts of data can result in high energy use. Efficiency is key for data-intensive applications, and AMD EPYC processors power the most energy-efficient servers available. In addition to delivering overall better performance per watt, AMD EPYC processors make it possible to closely match CPU resources with application requirements, creating even greater efficiency.

For example, some analytic applications do not scale well to high core counts. Using high-frequency AMD EPYC processors can increase per-core performance to speed these applications without the burden of carrying additional cores not essential to the mission. Some technical computing applications operate best when processors are equipped with large L3 caches. AMD EPYC processors with AMD 3D V-Cache™ technology free CPUs to process data with fewer cache misses and therefore unimpeded performance.

Whether you need as few as 8 or as many as 128 cores, or specialized processors, AMD EPYC processors offer the freedom to choose. All core features—including memory capacity, I/O bandwidth, and security features—are consistent within each processor family.

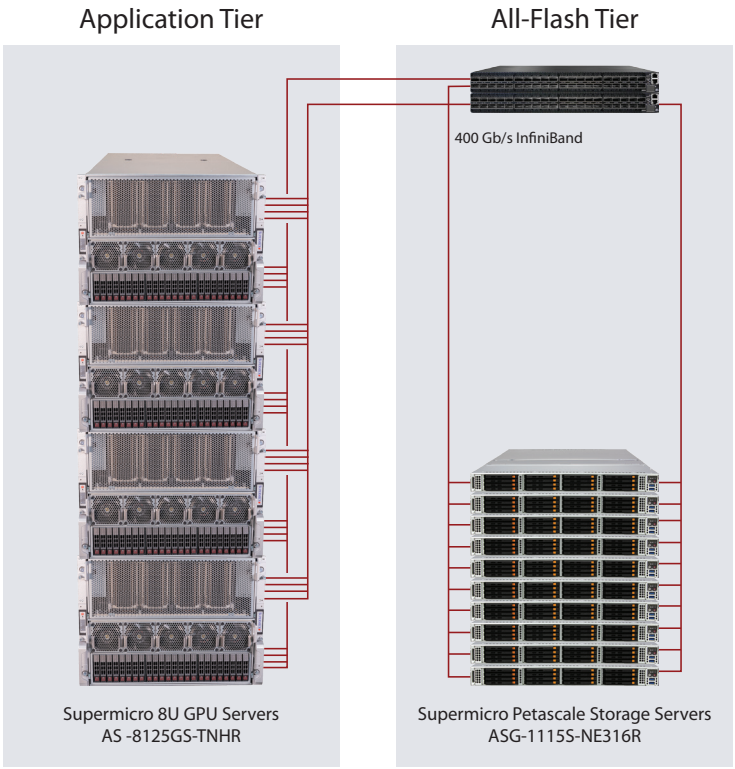

Figure 3. Example architecture designed to feed data to GPU servers with eight 400 Gb/s network interfaces per server

Large and Cloud-Scale Solutions (1-10+ PB)

Large solutions support data lakes, active archives, media repositories, and long-term storage for HPC and AI/ML data. Usually the domain of object storage protocols such as S3, these solutions always use scale-out clusters. Software solutions typically provide both file and object-level access to storage, and include features such as immutable storage as a way to enforce data retention policies and protect against ransomware. Since they serve as large repositories, software-defined storage features often include the capability to manage tiers, including moving long-term data to cloud storage, or staging data onto faster clusters for immediate consumption at high bandwidth, such as ML training.

Large and cloud-scale solutions can be comprised of servers such as the 12-drive Hyper 1U server, and the 2U, 32-drive Petascale server. These servers are small compared to the solution size, and are used to optimize large solutions where high performance is of paramount importance.

When large solutions are used for long-term storage, repositories, and data lakes, top-loading storage servers with up to 90 drives per chassis are often favored for the massive amount of data each one can support. We offer top-loading solutions with 60 or 90 3.5" hard-disk drives, and with one or two CPU nodes. This yields four solutions: a 90-drive system with two nodes connects 45 drives to each node, while the single-node option uses one node for all 90 drives. A similar set of choices is available for the 60-drive top-loading systems. The key difference between large and cloud-scale solutions is the number of drives per

node. cloud-scale solutions have such a massive amount of data, with relatively fewer I/O operations per second, that creating solutions from single-node, top-loading systems yields adequate performance. In large solutions, where performance is more important, we recommend the two-node-per-system choice.

Store Your Most Important Assets on Supermicro Storage Servers

Everyone agrees that data is your most important asset. So choose to store that asset on tested and validated storage solutions from Supermicro. We work closely with the leading software-defined storage software vendors to ensure compatibility with our broad range of storage server solutions, so you can trust that they will run their best when integrated into a solution using Supermicro systems. Supermicro, or our partners, can work with you to size and design a storage solution that meets your specific needs, and that can serve you far into the future. We architect our solutions with scalability as a foremost concern so that you aren't constrained by the limitations and the fork-lift upgrades that are common with traditional storage appliances. When you choose to work with Supermicro, you have access to a broad product line and solutions that are cost effective, highly flexible, and high performance.

Get Started

You can count on better business outcomes when you choose software-defined storage solutions from Supermicro. To learn more, visit supermicro.com/en/products/storage or contact us at www.supermicro.com/en/contact.

Supermicro (NASDAQ: SMCI) is a global leader in Application-Optimized Total IT Solutions. Founded and operating in San Jose, California, Supermicro is committed to delivering first to market innovation for Enterprise, Cloud, AI, and 5G Telco/Edge IT Infrastructure. We are a Total IT Solutions manufacturer with server, AI, storage, IoT, switch systems, software, and support services. Supermicro's motherboard, power, and chassis design expertise further enables our development and production, enabling next generation innovation from cloud to edge for our global customers. Our products are designed and manufactured in-house (in the US, Taiwan, and the Netherlands), leveraging global operations for scale and efficiency and optimized to improve TCO and reduce environmental impact (Green Computing). The award-winning portfolio of Server Building Block Solutions® allows customers to optimize for their exact workload and application by selecting from a broad family of systems built from our flexible and reusable building blocks that support a comprehensive set of form factors, processors, memory, GPUs, storage, networking, power, and cooling solutions (air-conditioned, free air cooling or liquid cooling).