WHITE PAPER

Data Lakes with Supermicro® Storage Servers

Organizations are increasingly turning to data lakes for storing all of their data with the economics of on-premises infrastructure. Data lakes have three characteristics that make them compelling for this purpose. They typically store data in its native format, including structured, unstructured, and semistructured data for applications including data warehouses, big data, and AI training data. They provide a high level of security including support for immutable files and objects that can protect active data as well as backups, ensuring they can recover from ransomware attacks. Finally, data lakes are elastic and can grow as needed, economically, to store massive amounts of data. Supermicro offers a range of storage servers built to support the needs of the various types of data lakes discussed in this paper.

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.

The release of AMD EPYC 9005 Series processors enables our H14 series servers to deliver even more parallelism, whether you use it for high-speed servers in data lake deployments, or for CPU-based model training and AI inference. AMD's high-frequency CPU products are optimized for AI with high per-core performance.

Whether you need as few as 8 or as many as 192 cores, 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.

Introduction

Data, in and of itself, is not powerful, it's just data. The ability to leverage your data to deliver insights to be more competitive, to bring new products to market faster, to use artificial intelligence (AI) to make better, smarter decisions—this is the power of information that comes from data you already have in your possession. The challenge is storing the vast amounts of data in a single name space as it grows, to extract value in the form of advanced analytics, and to create better AI and ML models.

Legacy storage architectures

Traditionally, all data was tightly coupled with the system and application using the data. The data's location and format was, in effect, hardwired to the computer running the application. This created many, many islands of data, each accessible to only a single application. Then, storage networking emerged, enabling distributed applications to access information as long as the data formats were consistent and well understood—which often wasn't the case. Standards groups such as the Storage Networking Industry Association (SNIA) emerged to solve both access and format disparities. As the cost of storage plummeted, the complexity of storage silos remained, motivating the development of new approaches to handling data.

Contemporary storage architectures

As the cloud opened up new ways to architect and deploy applications, data as objects came into prominence. Public clouds, for example, were built using object storage. One of the compelling reasons for this is that objects can be accessed from anywhere on the Internet, without the scaling limitations that were found in other unstructured storage types. This provided a platform for unlimited growth, making them ideal for data accessed online and by cloud-native applications that need access to persistent data.

Amazon created the ***de facto*** standard Amazon Simple Storage Service (S3) that allowed virtual machines and containers to access storage through URLs and transferred securely using Secure Socket Layer (SSL) protocols. Today, most of the world's data is stored as objects accessible through S3 protocols.

As the amount of data continues to grow, it exposes the complexity of less-scalable storage architectures. Today, many organizations are building limitlessly scalable centralized repositories for all their data, often natively stored as objects and accessed using S3 semantics. Object storage is ideal for much of the data stored in data lakes in that it is written once and read many times.

A modern data lake is comprised of a collection of smart storage servers that are disaggregated from the application servers and networked together to deliver seamless, high-performance access to all data in the data lake. Traditionally, data lakes have been thought of as large archives in which to store stagnant data of all types, including:

- Structured, such as images, spreadsheets, and databases
- Unstructured, such as text, IoT data, and log files
- Semistructured, such as HTML, XML, and email

SIMPLE STORAGE SERVICE (S3)

S3 is a public cloud storage service that was originally developed by Amazon Web Services. It was developed as an interface into data management services for data stored in the cloud. S3 manages data as objects (rather than as files) and comes with an API to control the data. While S3 is the de facto standard for cloud storage, most organizations need on-premises data lakes for performance- and mission-critical data. Fortunately, S3 is also the "easy button" for moving storage around if needed. For example, you may choose to move some data to the Cloud for resiliency or burst capacity. However, the performance demands of machine learning (ML) training requires on-premises storage and is often the most challenging use case for data lakes.

Contemporary data lakes include enterprise-class data features such as support for immutable objects, security features that protect the data including anti-ransomware protection, S3 semantics, and often other file-access protocols including Network File System (NFS) v4.2 and Flex Files.

A Diversity of Data Lakes

As organizations change their thinking of data lakes as places to store not just archive, but also active data, they are evolving to support a much wider range of data. Deploying a data lake is rarely a greenfield opportunity. Existing application data silos are being incorporated into data lakes so that they can be routinely accessed by new applications. For example, data analytics silos using Hadoop Distributed File System (HDFS) can be stored as part of a data lake and then used by AI software (Figure 1). As data lakes subsume data from across the organization, they are being called upon to support a much wider range of applications and access protocols. Contemporary data lakes are supporting access protocols including:

- POSIX file system semantics
- Network file system protocols including NFS and SMB
- Container services interface (CSI) to support containers and microservices
- NVIDIA GPU Direct Storage that leverages technologies such as RDMA over Converged Ethernet (RoCE) for direct transfers between data lakes and GPU memory for machine-learning (ML) training

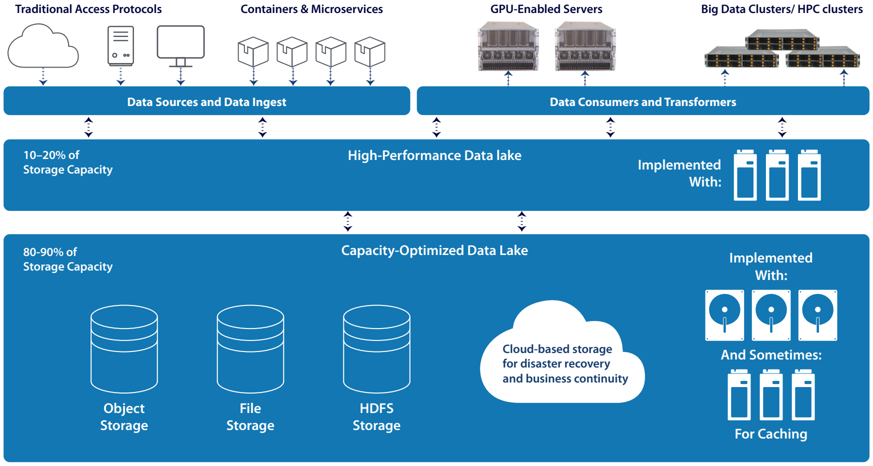

Figure 1. Today's object stores increasingly support more than just S3

With support for many protocols beyond S3, forward-looking companies can knock down silos and use data lakes in some of the following new ways:

- AI data preparation including data normalization and cleaning that uses traditional POSIX interfaces
- Big data mined for business insight with Legacy Hadoop infrastructure storing data in the data lake
- High-performance computing data and models such as weather, geophysical, and financial models that ingest petabytes of data per run
- IoT data migrated from multiple collection points

STORAGE SYSTEMS FOR TRADITIONAL DATA LAKES

TOP-LOADING SYSTEMS

A variety of 4U top-loading storage servers deliver the ultimate in 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 SimplyDouble systems support 24 3.5” hard-disk drives in a compact 2U form factor, double what most 2U systems can support. Powered by AMD EPYC™ processors, you can run storage applications on the server with your choice of number of cores per processor.

- Open table formats like Parquet and Apache Iceberg for storing columnar data optimized for analytical (OLAP) workloads
- Data warehouses
- Content and creative asset development
- Healthcare records repositories

Data Lake Architectures

Originally conceived as large repositories for unchanging data, legacy data lakes were exclusively implemented with hard-drive-based, distributed, scale-out storage clusters. As AI training sets become massive, data lakes have been tasked with storing both historical data—so that every model training set is reproducible—as well as active data used to saturate GPU-intensive servers with data in parallel, high-bandwidth streams. As data lakes are more frequently holding active as well as archived data, the need for both scalability and performance have come to the forefront. Software-defined storage vendors have begun to respond by supporting data lakes completely based on flash storage, and some have implemented caching with flash storage in a data lake that is otherwise comprised of scale-out nodes hosting hard drives.

As storage prices continue to decline, there is no longer a reason to throw away data that can possibly be used in the future, and data lakes are increasingly the repository of choice for both archive and active data, supporting applications illustrated in Figure 2).

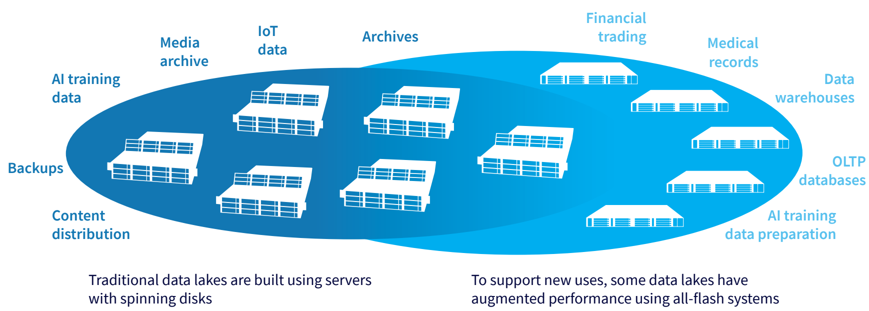

Figure 2. Data lake architectures have typically used servers with many hard drives; as they are brought into supporting additional use cases, all-flash data lakes are starting to be deployed

Legacy data lakes

In legacy data lakes, performance was secondary to capacity. This led to the need for some form of storage tiering, either internal or external to the primary data store. As data becomes less relevant and important, tiering moves it to less-expensive media. For example, moving from NVMe storage to slower and less expensive hard disk drives. Many storage management products can manage and automate tiering to the cloud, to NVMe caches, and to even faster on-premises data stores.

High-performance data lakes

The high data demands of machine learning training has required object storage vendors to address ways in which to build high-performance data repositories that can feed data to multi-GPU servers fast enough to keep them hydrated with data.

STORAGE SYSTEMS FOR HIGH-PERFORMANCE DATA LAKES

SUPERMICRO PETASCALE STORAGE SERVERS

ASG-1115S-NE316R

SSG-121E-NES24R

ASG-2115S-NE332R

Supermicro Petascale storage servers are ideal for supporting all-flash data lakes. They 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.

We have tested and validated a three-tier architecture described in the Supermicro white paper *Accelerating AI Data Pipelines*. With this architecture, both performance and capacity are necessary. Flash drives are used as a high-performance tier and either HDDs or quad-level cell (QLC) flash drives are used for the persistent storage tier, which is implemented as a data lake. Although QLC has lower endurance than triple-level cell (TLC) flash, in read intensive environments such as data lakes it may strike a perfect balance between performance, cost, and reliability.

Some software-defined storage vendors have designed products to deliver sufficient bandwidth to feed GPU memory directly from data lakes. These high-performance data lakes use flash drives exclusively and can store all enterprise data on solid-state storage but at a significantly higher infrastructure cost than HDD-based storage.

Hybrid Data Lakes

There are two types of hybrid data lakes that are being used today. Hybrid cloud and hybrid device data lakes.

Hybrid-Cloud Data Lakes

Many organizations have used cloud services to support their data lakes, as they are already S3 compatible and have virtually limitless storage. But the promise of cutting capital costs has given way to higher-than-expected monthly fees that have prompted a rethinking regarding where to house data lakes. With on-premises data lakes, you have ultimate flexibility, data located where it is needed, and security and data sovereignty that is under your control. Many software-defined storage vendors support hybrid cloud lakes, which can shift data between an on-premises data lake and the cloud depending on business priorities, which could include cost considerations and data locality issues, having data located nearest the applications using it.

Hybrid-Device Data Lakes

These data lakes can use multiple types of storage devices to create an effective tiering scheme (Figure 3). For example, you can build a three-tier architecture that has some fast TLC (Triple Level Cell) flash for your most demanding applications. For your less demanding applications, you can use fast HDDs or QLC Flash and for your write-once-read-many applications like archival or backup, you can use high capacity, slower HDDs. While our architecture supporting machine learning training uses multiple tiers, each with their own software-defined storage products, vendors are beginning to offer products that manage tiered storage within a single data lake.

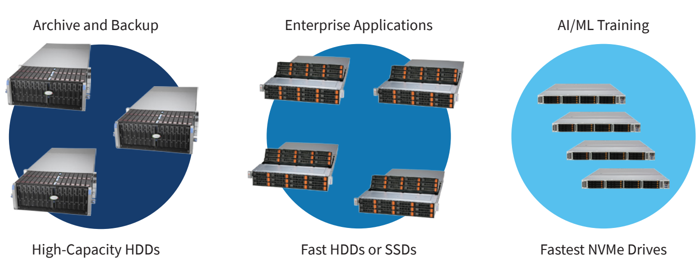

Figure 3. Hybrid-device data lakes can support a range of applications using storage systems with different performance and longevity characteristics

HIGH-PERFORMANCE DATA LAKES, CONTINUED

Depending on your overall requirements, creating data lakes with more, yet smaller servers can deliver high performance through parallelism, and high availability by reducing the impact of a single server failure.

HYPER 1U

[AS -1125HS-TNR]

The 1U Hyper DP is a 2-socket server is powered by the latest AMD EPYC™ 9005 Series processors and is a flexible platform that hosts up to 12 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.

CLOUDDC

[AS -1116CS-TNR]

The 1U CloudDC system is a single-socket, Open Compute Project (OCP)-compliant server that also supports hosts up to 12 NVMe, SAS, and SATA3 drives. With a single AMD EPYC 9005 Series processor with up to 192 cores, the CloudDC system strikes a balance between CPU power and storage capacity that is often ideal for software-defined storage deployments.

Data Lakehouses

Data lakehouse is a contraction of the terms *data warehouse* and *data lake*. They combine the horizontal scalability of data lakes with the transactional semantics of a data warehouse. Storage scales indefinitely, with data typically stored in columnar formats like Parquet, and metadata on top of it in open table formats like Delta Lake or Apache Iceberg. Unlike a traditional data warehouse, where compute and storage are on the same servers, data lakehouses decouple storage from compute, so that compute resources can be scaled on demand and queried by different tools that support the open formats.

Data lakehouses have the following characteristics to help them support high-performance, hybrid workloads.

- The data lake itself is completely unstructured
- You can build structures, for example for a database management system, on top of a data lake, or even implement retrieval-augmented generation (RAG) of AI training data
- They support atomic, two-phase commits to keep databases 100% consistent; they are able to provide transactional performance
- Several data management vendors have added features beyond data management that shift them raise them to being data platforms
- Some products work up the value chain beyond storage and data management to help customers gain a richer understanding and ability to introspect and/or process data natively
- Some data lakehouses focus on delivering an object store performance tier, optimizing for AI and ML workloads where both reads and writes have significant velocity

Why Supermicro?

Business moves too fast to anticipate all your storage workloads. When you start with a flexible foundation based on software-defined storage running on Supermicro storage servers, you can mold your infrastructure as your future needs dictate. Our agile product line can support you today and as you evolve into the future.

Choosing Supermicro gives you a simplified purchase process—with our storage partners listed in Table 1 (next page), you can purchase a storage solution from Supermicro including sizing and sever selection, installation, rack-level delivery, and ongoing software and hardware support with one number to call for support. Supermicro has a broad product line that can complete your solution, including GPU-dense servers for artificial intelligence, high-performance, enterprise-grade systems, dense multi-node servers for high-performance computing, and the networking technology to tie your solution together.

For more information

- Data lakes with Supermicro
- AI deep learning solutions
- AI storage solutions
- Storage solutions
- Software-defined storage

Table 1. Supermicro software-defined storage partners supporting data lakes

CLOUDIAN	ddn	EDB
Cloudian Hyperstore is an on-premises, S3-compatible data lake platform, particularly suited for AI/ML applications. HyperStore can scale to exabyte levels without disruption to operations.	*DDN* offers the EXAScaler parallel file system to support a wide range of data-intensive workflows in a hyperconverged model. It supports POSIX, NFS, SMB, HDFS, and S3 access to storage. Its Infinia product supports file and object stores.	*EDB Postgres AI* solves common enterprise challenges by extending enterprise-grade Postgres with native AI vector processing, an analytics lakehouse, and a unified platform for observability and hybrid data management.
HAMMERSPACE DATA IN MOTION	IBM	MINIO
Hammerspace Global Data Platform unifies unstructured data across edge, data centers, and clouds. It provides extreme parallel performance for AI, GPUs, and high-speed data analytics, and orchestrates data to any location and supports a wide range of protocols including NFS, pNFS, SMB, S3 and CLI.	*IBM Storage Ceph* is a software-defined scale-out enterprise storage platform that provides block, file, and object capabilities. Its best in class S3 capabilities support the largest and most intensive data applications.	*MinIO Enterprise Object Store* is an ultra high-performance object store that is used to deliver against AI/ML, analytics and archival workloads - all from a single platform. Software-defined, MinIO is ideal for a broad class of SuperMicro servers.
OSNEXUS STORAGE MADE SIMPLE.	Quantum	Qumulo
OSNexus QuantaStor is a scale-out shared-storage solution that delivers multi-tenant, S3-compatible, object storage with unique dynamic tiering features to optimally place objects for maximum performance and cost efficiency.	*Quantum ActiveScale* merges flash, disk, and optionally, tape libraries, to build data lakes, storage clouds, and computing clusters at any scale with outstanding performance, efficiency, availability, and durability at up to 80% lower cost than alternative solutions.	*Qumulo* is a single platform for all unstructured data, no matter where that data resides. Qumulo is built for managing geographically distributed file and object data on-premises, at the edge, in the core, and in the cloud.
SCALITY	VAST	WEKA
Scality solves organizations' biggest data storage challenges — security, performance, and cost. The world's most discerning companies trust Scality so they can grow faster and execute AI data-driven ideas quicker — while increasing efficiency and avoiding lock-in.	*VAST Data* enables data-intensive enterprises to effortlessly capture, catalog, refine, and enrich data through an API-driven data pipeline, providing real-time insights for applications like agentic AI. VAST empowers organizations to challenge conventional thinking and unlock transformative possibilities through its innovative data platform.	*WEKA* delivers a software-defined data platform that helps get your data out of silos and islands and into streaming data pipelines to fuel information workloads like AI and HPC.