

Solution Brief

NetApp ONTAP AI and ScaleMatrix

Deploy data center technologies at AI scale

Key Benefits

Simple, Fast Deployment

Eliminate guesswork and get started faster.

Integrated Solutions from the Edge to the Core to the Cloud

Unify data management across the pipeline with a single system.

Limitless Compute and Storage for AI

Start small and grow as needed.

Greater Density, Increased Efficiency

Maximize the potential of your ONTAP AI infrastructure and run at peak performance at all times.

Extreme Security and Reliability

Keep your solution up and running and ensure your data is protected wherever it lives or wherever it moves to.

The Challenge

Artificial intelligence (AI) is transforming business in virtually every industry, including healthcare, retail, manufacturing, financial services, automotive, and government. Designed to answer some of the world's biggest questions, today's AI and analytics workloads require cutting-edge hardware and software. These new, more powerful, high-density architectures come with unprecedented demands for airflow, power, and cooling. These demands leave many organizations struggling to support the latest high-performance AI computing architectures in their own data centers and drive the need to bring IT infrastructure closer to the edge.

The Solution

NetApp, NVIDIA, and ScaleMatrix with its Dynamic Density Control (DDC) Cabinet Technology have partnered to provide powerful, end-to-end solutions. These joint solutions deliver massive AI exploration and development capacity while meeting power, cooling, airflow, and density requirements wherever the solutions are deployed.

The NetApp® ONTAP® AI proven architecture is powered by NVIDIA DGX supercomputers and NetApp AFF all-flash storage inside a DDC cabinet. You get a single system that is easy to install, is simple to maintain, and can scale seamlessly as your data requirements grow. The innovative cabinet design of DDC enables you to add more high-density compute and storage units per rack. You can take advantage of leading compute, storage, and power efficiencies without worrying about overheating or ballooning data center costs.

ScaleMatrix high-density colocation data centers are built with DDC Cabinet Technology as their backbone and can thus host any hardware, scaling to virtually any density, with availability across all regions of the United States. DDC S-Series cabinets are purpose-built to scale and can be deployed on site in your data center to future-proof your investment and to support growth without needing to rearchitect your data center. You can also quickly deploy DDC R-Series enclosures to support demand at the edge or where you might have resource constraints. Ruggedized features enable you to deploy AI-ready IT infrastructure in some of the most challenging environments imaginable, including outdoors.



Figure 1) ONTAP AI inside the DDC trademark green cabinet in a ScaleMatrix data center.

Simple, Fast Deployment

Designing and deploying an AI-ready infrastructure can be a long and complex process. With a validated reference architecture that detangles design complexity, NetApp ONTAP AI enables you to eliminate guesswork and get started faster. When you combine ONTAP AI with DDC cabinets, you can be up and running in your data center, in a ScaleMatrix high-density data center, or in any covered location in just days. Simply connect a chilled water source and power to deliver high-performance data center capacity virtually anywhere.

Integrated Solutions from the Edge to the Core to the Cloud

Traditional machine learning (ML) and deep learning (DL) pipelines begin with data collection and ingest; move on to preparation, training, and deployment; and finish with analysis and tiering. As AI gains traction, organizations increasingly require IT resources to process data closer to the edge where the data is created. It is common—and inefficient—to have to transition formats between CPUs and GPUs and various storage types as you move through the data pipeline.

ONTAP AI with DDC Cabinet Technology unifies data management across the pipeline with a single system. You use the same tools to securely control and protect your data in flight, in use, or at rest and meet compliance requirements with confidence. With high-density edge, data center, colocation, and cloud solutions, ScaleMatrix and DDC can support your AI workloads wherever they run, including warehouses, airport terminals, military command centers, autonomous vehicle garages, and hospital wards. Today, DDC ruggedized R-Series cabinets can be found everywhere, from a Southeast Asia mine 9,000 feet underground to an Arizona mountain more than 8,000 feet above sea level.

Limitless Compute and Storage for AI

ONTAP AI enables you to start small and to grow as needed. You can add compute, storage, and networking to clustered configurations without disrupting your ongoing operations. You can start with a 1:1 storage-to-compute configuration and scale out as your data grows to a 1:9 configuration and beyond. With NetApp's rack-scale architecture, you can scale from hundreds of terabytes to tens of petabytes with all flash.

As your ONTAP AI infrastructure scales, DDC Cabinet Technology enables you to mix and match cabinet and chiller units to easily meet your capacity or redundancy requirements. If your project requires spikes of compute power, you can tie into a ScaleMatrix multitenant cloud on demand. Or you can connect to AWS, Microsoft Azure, or Google Cloud Platform by using Megaport software-defined networking (SDN) direct connections at ScaleMatrix.

Greater Density, Increased Efficiency

DL training routines demand massive amounts of compute power. Just one DGX-2 server provides over 2 petaflops (PFLOPS) of AI computing power, the equivalent of an entire data center of traditional CPU-based servers. At the same time, that DGX-2 draws a tremendous amount of heat—12kW for each unit. Hardware heating issues can throttle processing speeds and limit your return on investment. DDC cabinets at ScaleMatrix have the power to cool the highest-density compute and storage so that you can maximize the potential of your ONTAP AI infrastructure and run at peak performance at all times.

The DDC liquid-air-cooled cabinet technology combines the efficiency of water with the flexibility of air, cooling up to 52kW of power load in a 45U cabinet. Denser cabinets provide greater compute power in a smaller footprint, so you can rent or buy less space for your data center and reduce operational expenditures, such as power, cooling, and management.

Extreme Security and Reliability

Together, NetApp and ScaleMatrix with DDC Cabinet Technology enable you to integrate, to protect, and to secure your data pipeline from the edge to the core to the cloud. With NetApp data encryption at rest and in flight, your data is protected wherever it lives or wherever it moves to. ScaleMatrix offers unparalleled security at the cabinet level with biometric cabinet access, video security, and physical security checkpoints throughout its data centers. ScaleMatrix DDC-enabled data centers have a cleanroom environment that protects the equipment inside from dust, preventing hardware failures and extending the life of the equipment. And 24/7 worldwide services and single-point-of-contact support from NetApp, NVIDIA, and ScaleMatrix help keep your solution up and running.

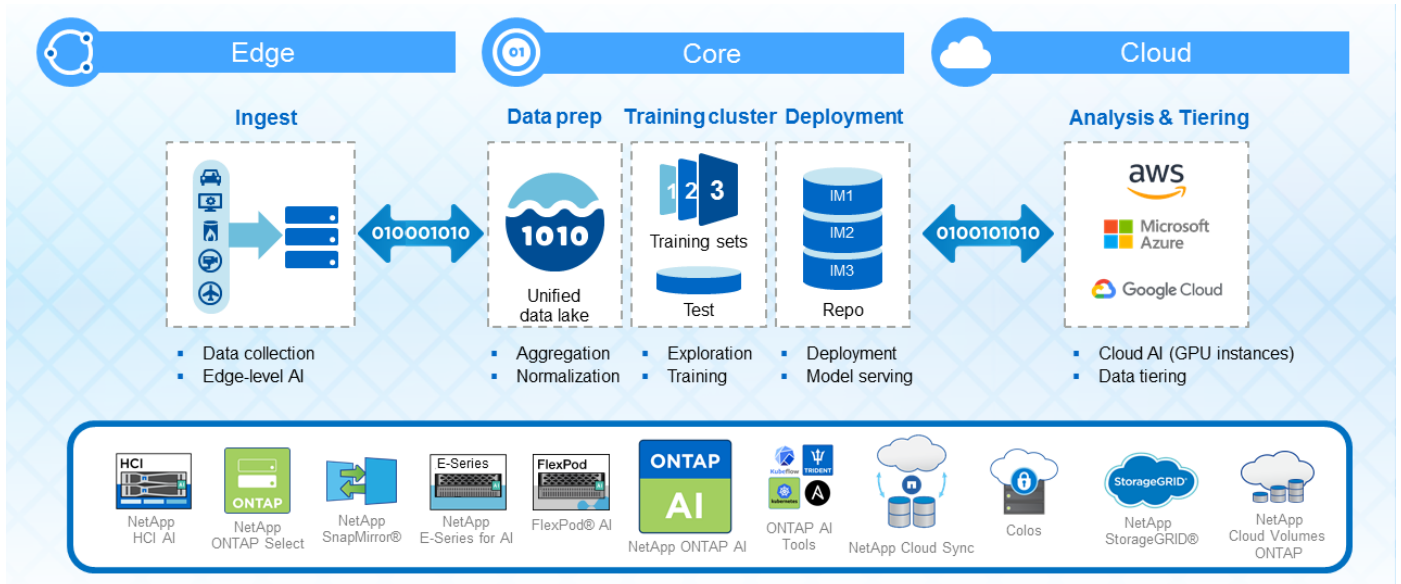


Figure 2) Seamless data management across edge, core, and cloud.

About ScaleMatrix

ScaleMatrix delivers colocation, cloud, backup, disaster recovery, and professional support from national high-density data centers that leverage the future-proof Dynamic Density Control (DDC) cabinet enclosure platform. With power density and efficiency significantly impacting IT costs, these specialized data centers enable ScaleMatrix to deliver competitively priced high-density colocation and high-performance cloud hosting which provide significant differentiation in today's service provider market. By leveraging our innovation and technology, ScaleMatrix clients gain a competitive edge and can scale more efficiently as their business grows. Visit our website at www.scalematrix.com.

About DDC Cabinet Technology

The Dynamic Density Control (DDC) platform is designed to solve density, efficiency, and cooling challenges presented by compute-intensive applications. ScaleMatrix' variable density data centers are DDC-enabled with colocation and cloud hosting capabilities for AI and deep learning workloads. DDC supports any rack-mountable hardware configuration, at nearly any density, without the need for customization or modification to the hardware being deployed. DDC enclosures address thermal management issues, noisy neighbor challenges, and provide significantly improved physical security; enabling them to be deployed in existing or new data center sites with ease. Visit www.ddcontrol.com for more information.

About NetApp

NetApp is the data authority for hybrid cloud. We provide a full range of hybrid cloud data services that simplify management of applications and data across cloud and on-premises environments to accelerate digital transformation. Together with our partners, we empower global organizations to unleash the full potential of their data to expand customer touchpoints, foster greater innovation and optimize their operations. For more information, visit www.netapp.com. #DataDriven