

Solution Brief

Extend Your Cloud AI Experience to the Premises

Scale your AI with Core Scientific's NetApp AI solutions and NVIDIA GPU offerings

Key Benefits

Eliminate Complexity

Get an AI environment that is ready to run out of the box.

Start Small and Scale as Needed

Eliminate large capital expenditures with best-in-class scalability.

Leverage a Cloud-Connected Data Pipeline

Easily create a hybrid multicloud AI environment for data scientists.

The Challenge

There is no doubt that AI is empowering today's businesses in almost every major industry. However, data scientists and data engineers are facing various challenges on their journey toward successful machine learning (ML) and deep learning (DL) initiatives. One challenge is choosing the right platform. Concerns about data center readiness, data privacy, escalating cost, scalability, and many other considerations complicate decision making, which usually doesn't happen overnight. This leaves many data scientists struggling to start their ML/DL journey on time. Data-sensitive organizations want the simplicity of a cloud-style service for their AI initiatives but are obliged to remain on the premises.

All these challenges necessitate combining data science tools, DevOps, GPUs, and data pipeline management in a simple operational expenditure (opex) offering that enables organizations to accelerate their ML/DL projects.

The Solution

NetApp, NVIDIA, and Core Scientific offer a cloud-style "as-a-service" solution, combining powerful data science tools and DevOps functionality with GPU compute power and data pipeline management in an opex consumption model.

The offerings range from an aiLab for data scientists who want to start their AI experiments with one of the tools in the aiLab catalog, to hosting of dedicated hardware consisting of highly available GPU compute and all-flash storage systems.

The aiLab is built on highly available containers designed for data scientists to take advantage of already installed Kubernetes with Jupyter Notebook and run aiLab catalog tools such as OmniSci, RStudio, TensorFlow, and Fastdata.io. The aiLab deploys both NetApp® ONTAP® AI and FlexPod® AI reference architectures, allowing choice of GPU acceleration delivered either by NVIDIA DGX-1/DGX-2 supercomputers or by Cisco UCS 480 ML servers with NVIDIA V100 GPUs. You can scale your deployments as needed, starting as small as two GPUs with 5TB of flash storage.

If you want to manage your own data science clusters and tools but need access to GPUs and data storage, you can take advantage of the dedicated, highly available GPU platform offering. This offering starts at two GPUs and 5TB of flash storage and can scale as needed. There's also an option to own the equipment or deploy as a service, using an opex consumption model.



Figure 1) NetApp FlexPod and ONTAP AI.

Ready-to-Run AI Environment

Assembling and integrating off-the-shelf DL compute, storage, networking, and software components can increase complexity and lengthen deployment times. As a result, valuable data science resources are wasted on systems integration work. When you choose the aiLab offering, you simply eliminate the complexity of setting up and managing the right environment for your data scientists.

NetApp AFF systems keep data flowing to DL processes with the industry's fastest and most flexible all-flash storage, featuring the world's first end-to-end NVMe technologies. Trident, the NetApp storage provisioner for Kubernetes, further accelerates your ONTAP AI deployment by seamlessly moving your NVIDIA GPU Cloud (NGC) container images onto NetApp enterprise-grade flash storage, allowing end-to-end platform management from Kubernetes. This integration also facilitates data versioning along with code versioning to enable true comprehensive ML versioning for your data scientists.

The aiLab data center is equipped with fully redundant power and networking, 24/7 TechOps, security, and time-scheduling software.

Start Small and Scale as Needed

Many organizations are starting their AI work in a simple pilot project. Moreover, DL best practices suggest that organizations should start small and scale as they go. That's why Core Scientific is offering an opex model; it can help you eliminate large capital expenditures, start small with your AI proof of concept (POC) or pilot project, and scale as needed, using the best-in-class scalability features offered by ONTAP AI and FlexPod AI.

Cloud-Connected Data Pipeline

By taking advantage of NetApp solutions and Core Scientific's direct private network connection to the major cloud hyperscalers, you can easily create a hybrid cloud and multicloud AI environment for your data scientists. In this environment, data can seamlessly flow from and to Core Scientific data centers without incurring egress or ingress fees.

NetApp offers best-in-class data management and cloud integration features to help you accelerate DL while managing and protecting your critical data. With its industry-leading data services capabilities, ONTAP helps you manage and protect your data with a single set of tools, regardless of where it resides, and freely move data to wherever it's needed, from edge to core to cloud. And the NetApp StorageGRID® solution provides greater data management intelligence on a simplified platform for your object data. Because StorageGRID uses S3, it painlessly bridges hybrid cloud workflows and enables your data to be fluid to meet your business demands.

Getting Started

We want you to start your AI project as early as today. Get access to two to four GPUs and 10TB of all-flash storage for two to four weeks. Start your POC by using the currently available libraries in the aiLab catalog, such as Fastdata.io, or use our White Glove setup feature for libraries that are less common. For more information, simply contact your Core Scientific, NetApp, or NVIDIA sales representative.

About Core Scientific

Core Scientific is a leader in artificial intelligence and blockchain hosting, transaction processing, and application development. Led by a team that has a 10+ year AI success story, Core Scientific provides custom hosting solutions at scale. Core Scientific is pioneering new innovations and best practices in the AI and blockchain landscape, with advanced capabilities operating infrastructure at scale. Our platform is trusted by large-scale partners around the world to deliver reliable solutions that quickly adapt to dynamic market conditions.

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

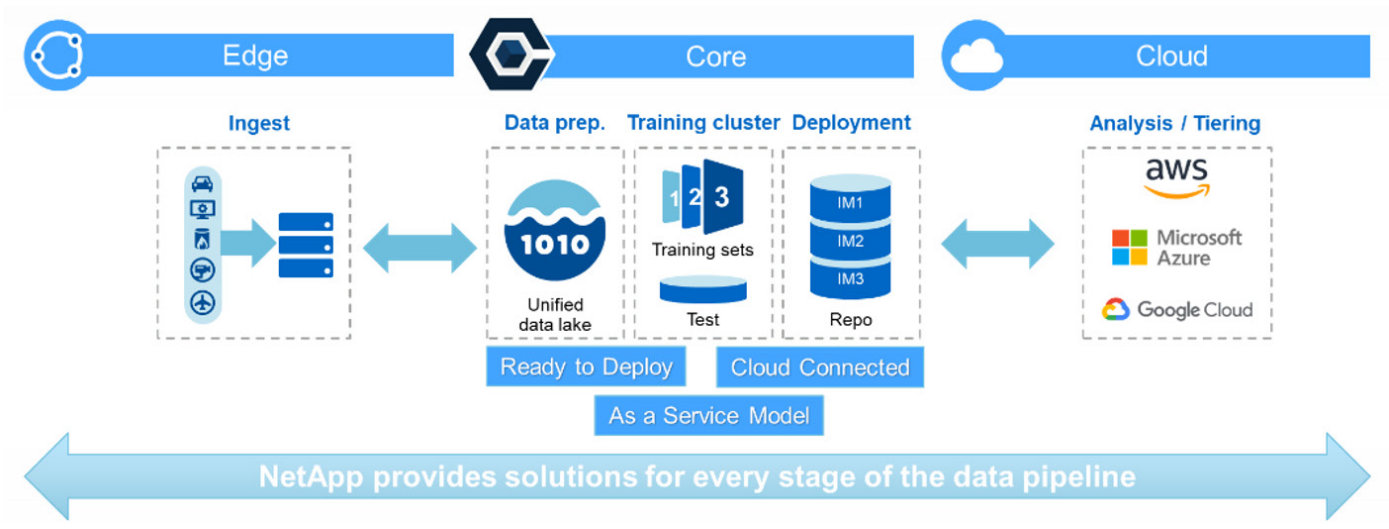


Figure 2) Data management for AI.

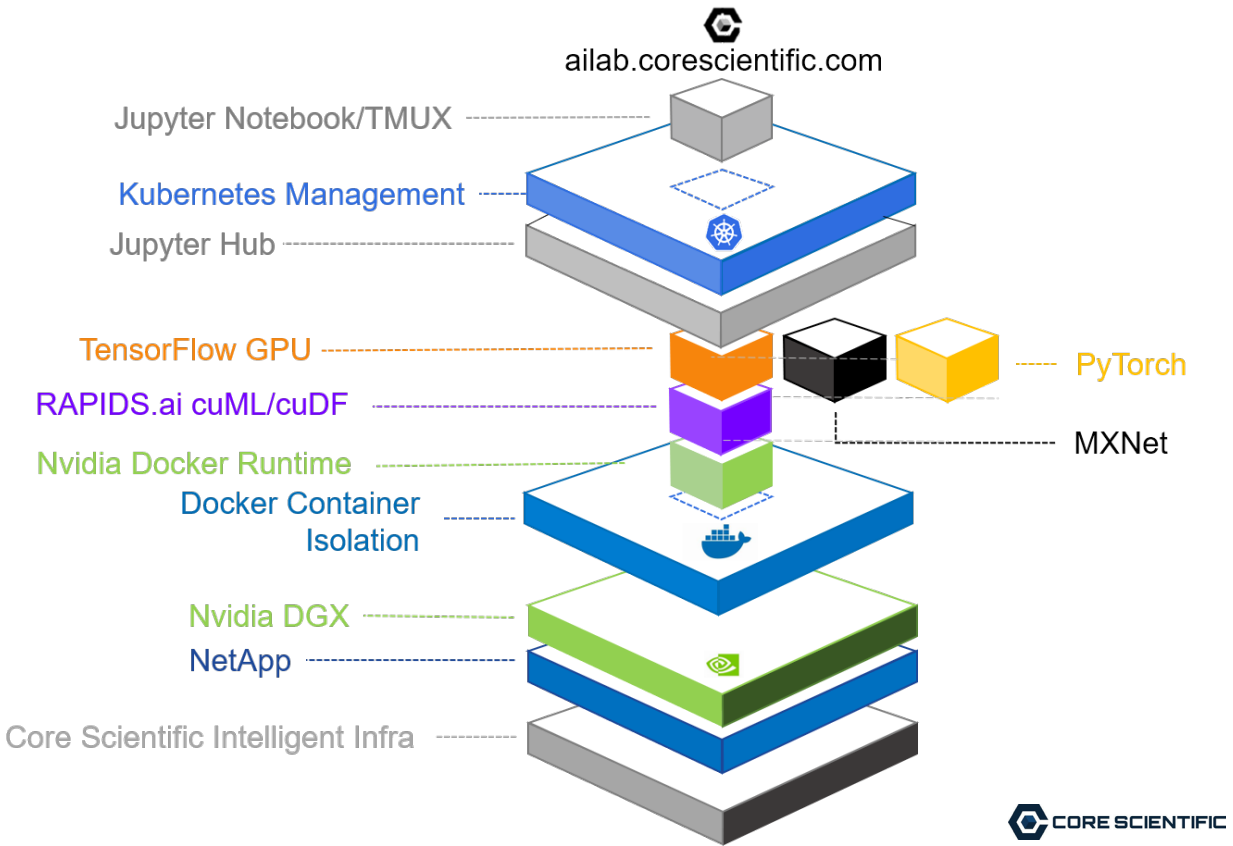


Figure 3) aiLab Components.