

Accelerate Deep Learning with Edge to Core to Cloud Data Management Solution

Cisco and SwiftStack collaborate for Distributed, large scale AI / ML data pipelines

Digital Transformation and the emergence of AI/ML workloads have created a new set of challenges for Line-of-Business owners to Data Engineers and IT staff. Data has become an organization’s most important asset, and with the right strategy and system, AI/ML has created an unprecedented opportunity to :

- Monetize new business models
- Gain predictive insights
- Enable competitive differentiation
- Improve business processes

The compute layer is transforming with the rise of GPU computing enabling massive parallelism and several petaFlops of computational power. Storage needs to be designed to provide the same degree of scale and parallelism.

Additionally, AI/ML workflows have distinct storage I/O requirements. Stages exist within data pipelines that require massive ingest and training bandwidth, mixed read/write handling, metadata labeling and lifecycle management. SwiftStack has unique, differentiated capabilities to meet these requirements. The SwiftStack and Cisco partnership brings together best-of-breed products - making it easy to consume a complete AI/ML solution stack, so organizations can focus on deriving business value.

HIGHLIGHTS

- Accelerate and scale your data pipelines whether it’s 15PB and 100 GBytes/sec today or smaller with plans to grow
- Leverage 1space to bring Cloud and AI (AWS, GCP) to your data, and build multi-cloud workflows across edge, core and cloud
- Use rich metadata tagging and search provided by SwiftStack middleware and SwiftStack Client
- Access the same namespace with both file protocols and object APIs
- Depend on a best-of-breed solution stack

USE CASES

- Autonomous Vehicles
- Healthcare and Life Sciences
- Retail and Smart Cities

AI/ML/DL storage workflow phases and their I/O challenges



Artificial intelligence in the data center at scale

Cisco® machine learning computing solutions ease the challenges faced by IT organizations and data scientists: supporting the needs of machine learning (ML) workloads while making them part of the enterprise data center. With Cisco solutions, you can power artificial intelligence (AI) workloads at scale and help extract more intelligence out of data to make better decisions.

SwiftStack has collaborated with Cisco to put together a reference design with Cisco UCS S3260 storage servers, Cisco 100 GbE network switches and Cisco C480ML GPU dense servers for compute. This on-premises stack deploys a variety of AI/ML frameworks from the NVIDIA GPU Cloud (NGC) repository of easy-to-consume containers, designed for Kubernetes orchestration. In addition, this environment can be extended seamlessly with SwiftStack 1space to AWS or GCP for cloud-bursting and multi-cloud workflows. Businesses can quickly start their AI/ML journey, expand to leverage multi-cloud workflows and frameworks, and scale to petabytes of storage and 100 GB's of throughput to meet their requirements.

SwiftStack Cisco | AI Reference Architecture

AI/ML Frameworks

- On-premises, container-based NVIDIA GPU Cloud (NGC) Deep Learning stack with the ability to extend to AWS and GCP
- Kubeflow for Deep Learning workflow management

GPU Compute


- The Cisco UCS C480 ML M5 Rack Server is designed for the most compute-intensive phase of the AI and ML lifecycle: deep learning.
- This server integrates GPUs and high-speed interconnect technology combined with large storage capacity and up to 100-Gbps network connectivity.

Networking

- Two Cisco Nexus 3232C 100 GbE switches

Storage

- SwiftStack software for scale-out multi-cloud storage and data management
- Cisco UCS S3260 storage servers with 1.5PB of capacity. Erasure coding or replication-based data protection policies
- Scale: Increase Cisco UCS C480ML for additional PFLOPS and UCS 3260s for PB's of storage and GB's of throughput



To try **SwiftStack for free**, go to <https://www.swiftstack.com/try-it-now/>.
For additional assistance or to learn more, always feel free to contact us. We're here to help.

Phone - (415) 625-0293 Email - contact@swiftstack.com

Chat - Just go to [swiftstack.com](https://www.swiftstack.com) and look for the chat pop-up in the bottom right