

SwiftStack for AI & Deep Learning

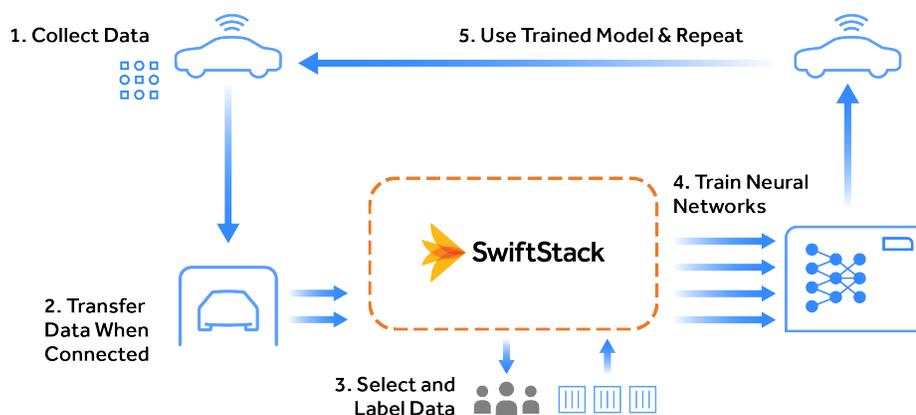
The Platform for Intelligent Data

Only SwiftStack's shared-nothing distributed architecture was designed from its inception to scale linearly to infinite levels of throughput, concurrency, and capacity, which means only SwiftStack can ingest, retain, and feed training data fast enough to fully utilize thousands of parallel GPUs working to produce today's most intelligent DNNs and machine learning algorithms.

Featured Use Case: Autonomous Vehicle Development

SwiftStack powers the data factories and Deep Neural Network (DNN) training for leaders in autonomous vehicle development. Training a single complex DNN requires at least three million images, and refining it to be production-ready will take countless more. Vehicles aimed at SAE autonomy level 2+ will need at least 10 DNNs, and level 5 will require 20 or more.

Autonomous Vehicle Development Workflow



To generate these training images, survey cars drive thousands of hours per year, and each car will produce petabytes of data that must be rapidly ingested and served to GPUs for curation, training, and replay, and which must be retained immutably for validation. Reference architectures for autonomous vehicle development specify usable storage capacities over 15 PB and throughput over 100 GB/second.

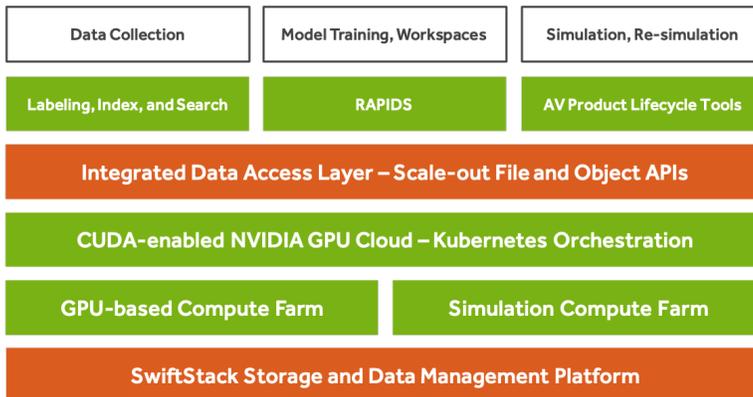
HIGHLIGHTS

- Ultra-Scale Performance**
 A distributed shared-nothing architecture scales linearly to infinite levels of throughput, concurrency, and capacity.
- Elasticity from Edge-to-Core-to-Cloud**
 1space enables applications to ingest and access data on edge devices, in core data centers, and/or in public cloud infrastructure.
- Data Immutability**
 Unique features ensure data is retained as it was originally written—for traceability, accountability, and safety.
- Optimal TCO**
 Savings— $\frac{1}{3}$ the cost of public cloud storage and 90% less than flash—mean greater investment in top-level AI initiatives.
- Real-World Confidence**
 SwiftStack is trusted to manage hundreds of petabytes in the world's largest companies—including NVIDIA and major automakers doing autonomous vehicle development today.

"Infrastructure challenges are the primary inhibitor for broader adoption of AI/ML workflows. SwiftStack's multi-cloud data management solution is the first of its kind in the industry and effectively handles storage I/O challenges faced by edge-to-core-to-cloud, large-scale AI/ML data pipelines."

Amita Potnis, Research Director at IDC's Infrastructure Systems, Platforms and Technologies Group

SOLUTION BRIEF



The SwiftStack AI Architecture is an engineered and customer-proven edge-to-core-to-cloud storage and data management solution.

Core architectural capabilities like ultra-scale performance, distributed flash-based cache, data immutability, standards-based APIs, and workflow integration make deep learning at scale possible.

SwiftStack ingests, retains, and provides data to thousands of GPUs working in parallel—“feeding the beast” to train and retrain the deep neural networks necessary for AV.

“We have about a PB of data—more than a PB of data—that is coming in per month and about 15 PB that is curated. For any large-scale company, this is not a huge amount of data, but what makes it tricky and hard is that all this 15 PB is the active training data set, so all of the internal teams are accessing it all of the time every day; every model that you’re working on is going to be using that data.”

Senior Engineer, AI Infrastructure

- 1 AI Frameworks** - Many of the leading AI frameworks and toolsets—including TensorFlow—can now leverage public- and private-cloud data directly with the S3 API.
- 2 Accelerated Compute** - DNN training is accelerated with NVIDIA’s industry-leading GPUs in NVIDIA-branded DGXs or OEM systems like the Cisco C480ML or Dell DSS 8440.
- 3 Scalable Storage** - SwiftStack software deploys easily on industry-standard servers from partners like Cisco, Dell, HPE, and Supermicro, supports the industry-standard S3 API, and seamlessly merges on-premises and public-cloud storage to enable GPUs in any location.



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 - Visit [swiftstack.com](https://www.swiftstack.com), and look for the chat pop-up in the bottom right