

Multi-Cloud Data Management for at-scale AI/ML

Distributed storage for Artificial Intelligence (AI) and Machine Learning (ML) data pipelines

Digital transformation and data-driven applications are driving profound changes in the underlying IT infrastructure. The compute layer is transforming with the rise of GPU computing, which enables massive parallelism and several petaflops of compute power. Storage services need to be re-thought to match this parallelism and scale. SwiftStack has unique, differentiated capabilities to meet these requirements.

Put valuable GPU cycles to work; Accelerate your data pipelines with SwiftStack



SwiftStack provides massive parallelism to match GPUs and accelerated compute

SwiftStack's cloud-native architecture for storing unstructured data is purpose-built to meet the emerging requirements of applications leveraging AI and ML workflows. This enables non-blocking, line-rate ingest (write), training (read) and mixed read/write storage I/O bandwidth.

Massive concurrency to cater to edge and IoT devices

SwiftStack Storage can match the massive concurrency required to handle data streams from several edge and IoT devices.



Best GPU utilization and improved Data Scientist productivity

Neural net training times are one of the primary contributors to deep learning time-to-value. SwiftStack's massive parallelism enables significantly lower training times.

Improved AI/ML inferencing latency with read-only caching

SwiftStack's massive parallelism and throughput, coupled with low-latency caching, enables very effective inferencing workflows.



Independently scale throughput and capacity

Throughput performance can be scaled independently of capacity by separating the access and storage roles in the scale-out cluster.

Bring cloud and AI to your data; Use a multi-cloud workflow to extend on-premise footprint

SwiftStack provides a multi-cloud, global namespace across the edge, core or cloud

SwiftStack 1space provides a global namespace. This enables leveraging the best available infrastructure resources as needed by AI/ML workflows. Distributed GPU jobs orchestrated by Kubernetes can leverage SwiftStack's global namespace no matter where they execute.



SwiftStack versioning enables AI/ML workflow experimentation

Different experiments can be tried out, while retaining the exact data copy corresponding to a particular experiment.

Policy-driven multi-cloud AI/ML workflows

Policy-based workflows using metadata tagging can move only relevant data across multiple clouds or make data visible from one cloud to another without moving it at all.



Don't drown in your data; Leverage metadata labeling and search to extract more value
Enrich your data with custom metadata labeling using SwiftStack middleware

Data can be labeled or tagged inline on ingest by leveraging SwiftStack middleware. Supervised learning workflows often use this metadata in the learning phase. 1space policies can be triggered using this tagging to place specific subsets of data for cloud-bursting workflows as well.



Client
SwiftStack Client to augment AI/ML workflows

Data scientists can leverage the SwiftStack Client to easily discover data to use in their AI/ML workflows by searching metadata. Metadata search is integrated into the SwiftStack Client.

Think Data First; Universal Access enables File and Object API's to same namespace

File and object access to the same data

AI/ML workflows can ingest using object APIs, while training and inferencing can leverage file access protocols for broad AI/ML framework support.

This capability is enabled by a new, purpose-built, cloud-native file system called ProxyFS and is built into the core of SwiftStack.

Object API flexibility

SwiftStack has best-in-class S3 API compatibility and leads the development of the most popular open API, Swift. Workflow integration can be quickly achieved using the common S3 API, additional unique features of the Swift API can be leveraged, or both APIs can be used interchangeably with the same data.


Don't break your bank building at-scale AI; Best economics with elastic resources

SaaS subscription pricing

SwiftStack provides a SaaS subscription pricing, which allows customer to start small and grow to petabyte scale with their data.

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