SOLUTION OVERVIEW

Accelerate AI model training on AWS

Build efficient and secure data pipelines for AI workloads with advanced traffic management for data ingestion from F5 and AWS.

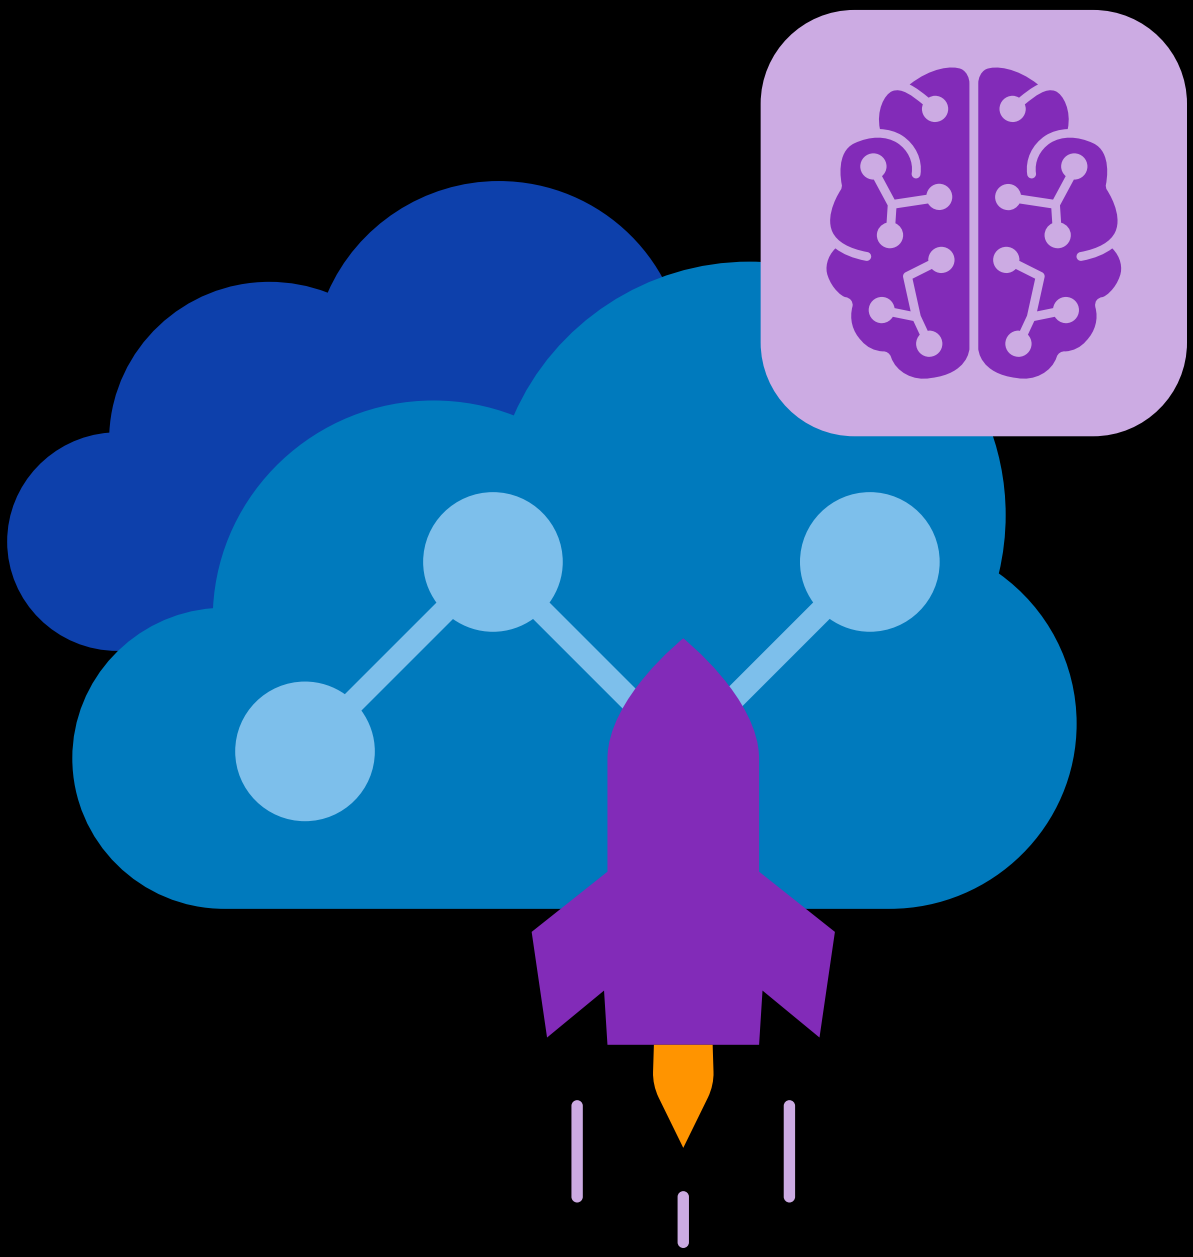

Key benefits

Maximize infrastructure value
Optimize GPU utilization and reduce costs with efficient traffic management that also frees up resources with SSL offloading.

Enhance security posture
Protect valuable training data with consistent security controls and end-to-end visibility that eases compliance efforts.

Simplify and scale complex environments
Reduce network complexity, enable consistent operations, and scale pipelines confidently for diverse data sources.

Organizations need sophisticated traffic management to move data efficiently, keep GPU queues full, and accelerate model training, all while protecting sensitive or proprietary data.

The data challenge in AI model training

Training or refining an AI model requires moving massive amounts of data to your AI infrastructure. This data often includes text, images, audio, and video, which increases network traffic volume and complexity. Organizations source this data from multiple environments, including on-premises data centers, public and private cloud storage, and edge locations that may be geographically diverse.

Bottlenecks in the data ingestion pipeline can cause expensive GPU resources to sit idle, resulting in high costs and slower training. Efficient GPU utilization requires a constant data stream. Organizations need sophisticated traffic management to move data efficiently, keep GPU queues full, and accelerate model training, all while protecting sensitive or proprietary data.

AWS offers multiple services designed for efficient AI data ingestion. However, these capabilities do not provide a complete solution for complex scenarios involving hybrid environments and multiple external sources. Organizations need to add complementary services to satisfy network infrastructure and security requirements and keep costs low as data volumes increase.

AWS data ingestion and migration services:

Amazon SageMaker Pipelines is a serverless workflow orchestration service designed to automate model training.

Amazon Data Firehose is designed to deliver real-time steams into data lakes, warehouses, and Amazon Simple Storage Service (S3).

AWS Database Migration Service (DMS) can maintain a continuous flow of data from source databases to AWS.

AWS DataSync moves datasets between on-premises storage and AWS and works with **AWS Transfer Family** to support standard file transfer protocols.

Key features

Deliver powerful multicloud connectivity
Enable high-performance connectivity between data sources and AI services with F5® Distributed Cloud Network Connect and the F5 Global Network.

Reduce data transfer costs
Optimize traffic management with advanced load balancing, server health monitoring, intelligent routing, and more.

Secure data and workloads end-to-end
Protect valuable data with encrypted traffic inspection, distributed denial-of-service (DDoS) attack protection, web application firewalls, and API security across environments.

Using NetApp storage? F5 securely connects NetApp data stores in the cloud and on premises through multiple protocols, including NFS, SMB, and Amazon S3 APIs, providing flexible and protected options for data access.

Accelerate AI data transfers with secure multicloud connectivity

If you need to connect multiple external sources to your AWS infrastructure, F5 can extend networking, load balancing, and security across the hybrid environment. The F5 Application Delivery and Security Platform provides multiple capabilities that work together to create secure, efficient, and optimized data pipelines.

Robust connectivity at the network and application layers

F5 Distributed Cloud Services simplify data movement across multicloud environments. Distributed Cloud Network Connect establishes secure layer 3 connectivity between training sources and AWS, supporting one-click provisioning, site-to-site encryption, and security service insertion with end-to-end visibility. F5 also enables layer 7 service networking through Distributed Cloud App Connect.

Organizations can enable multicloud networking by connecting sites over the Internet, through a private backbone, or by leveraging the private F5 Global Network, which features DDoS protection and 15+ terabits per second (Tbps) of peered capacity.

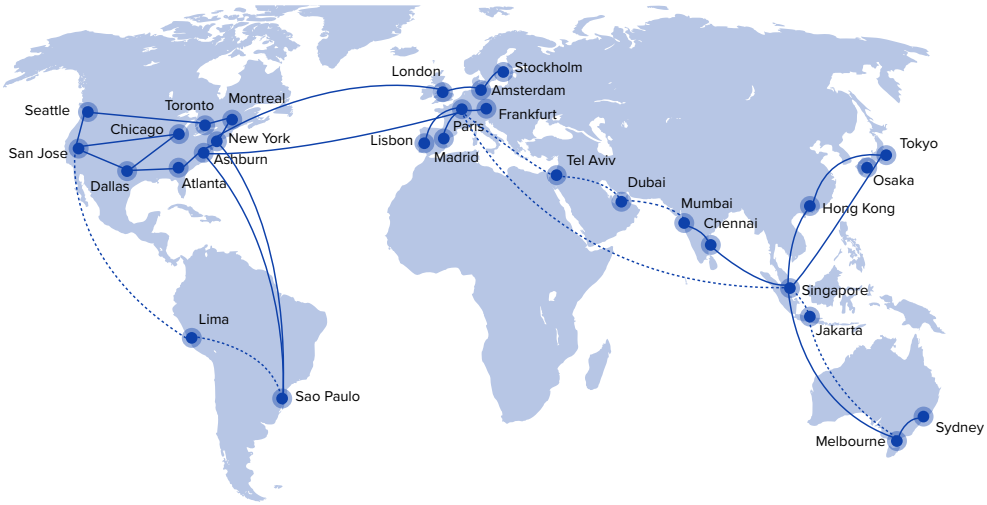

Figure 1: Points of presence in the F5 Global Network.

Sophisticated and streamlined traffic management

F5 BIG-IP® Virtual Edition (VE) provides advanced load balancing and layer 4-7 traffic optimization capabilities on AWS to continuously enhance data pipelines. TCP optimization, server health monitoring, and intelligent routing improve data throughput, ensure reliable data flow, and maximize GPU utilization. Other features, such as SSL offloading and FastL4 and Fast HTTP configuration profiles, optimize traffic processing even further.

Secure data transfers in hybrid environments

F5 security solutions in BIG-IP and Distributed Cloud Services provide broad protection that complements AWS security capabilities. Integrated controls deliver end-to-end security including DDoS mitigation, web application and API protection, full SSL/TLS inspection, and encryption for data in transit. F5 provides a zero-trust access control model with centralized authentication, authorization, and policy enforcement across environments. Audit and logging capabilities make it easier to demonstrate compliance for regulated data and workloads.

Streamline AI model training with F5 and AWS

Backed by a robust portfolio and over a decade of joint engineering and innovations, F5 and AWS deliver solutions to efficiently and securely ingest data for AI model training with simplified deployment through the AWS Marketplace. Leverage these integrated services to accelerate AI initiatives, control costs, and empower your business.

Contact F5 to get started or learn more at f5.com/aws.

©2025 F5, Inc. All rights reserved. F5, and the F5 logo are trademarks of F5, Inc. in the U.S. and in certain other countries. Other F5 trademarks are identified at f5.com.
Any other products, services, or company names referenced herein may be trademarks of their respective owners with no endorsement or affiliation, expressed or implied, claimed by F5, Inc.
DC 08.2025 | JOB-CODE-123456789