


TURN ENTERPRISE DATA INTO AI-READY FUEL ON AWS

Amazon FSx for NetApp ONTAP® is a high-performance, multiprotocol data layer that lets customers run AI and ML directly on their existing data without duplication or refactoring.



AI is now central to how enterprises innovate, but most teams struggle less with models than with data: feeding GPUs and AI services with massive, diverse datasets without endless copying, pipeline sprawl, or runaway costs. Amazon FSx for NetApp ONTAP enables AWS sellers to turn existing enterprise data into AI-ready fuel by combining ONTAP's data management strengths with native access to AWS AI, ML, and analytics services.



[The Business Value of Amazon FSx for NetApp ONTAP](#)




[Reduce Costs and Increase Efficiency with FSx for ONTAP | NetApp](#)

Buyer pain: “What’s breaking AI and ML projects?”

Most organizations have big ambitions for AI, but they’re blocked by the data layer: moving, preparing, and governing massive volumes of data across file, object, and databases.

Sellers should recognize these pain points:

-  **Weeks wasted on data wrangling and duplication**
Teams spend most of their AI project time re-copying files, exporting NAS data into S3, and rebuilding feature sets instead of training and deploying models—every extra copy adds cost, delay, and risk.
-  **Cost overruns from unmanaged data growth**
AI and ML workloads generate explosive data growth, and without efficient storage and tiering, infrastructure costs quickly exceed forecasts and erode ROI and confidence in the AI program.
-  **Silos between file data and AI services**
High-value data sits in file shares and NAS systems while AI tools expect S3 or other object interfaces, forcing teams to build and maintain complex, fragile sync pipelines before any insight can be delivered.
-  **Security and governance headaches**
AI projects increasingly touch sensitive data. When that data is copied into multiple uncontrolled environments, poor governance leads to failed audits, breaches, and elevated model risk.
-  **Quality bottlenecks that stall models before production**
Inconsistent, incomplete, or biased datasets make it impossible for AI to deliver impact, and without a consistent data platform, organizations struggle to keep training data clean and pipelines repeatable across teams and regions.



[The Business Value of Amazon FSx for NetApp ONTAP](#)

The Seller’s Advantage: Positioning FSx for ONTAP as the AI data layer

FSx for ONTAP gives AWS sellers a way to connect enterprise data directly to AI and ML services—without massive data migrations or re-platforming.

-  **Expose ONTAP file data to AWS AI services in place**
With Amazon S3 Access Points for FSx for NetApp ONTAP, customers can present existing file datasets as S3 sources to AWS AI, ML, and analytics services without copying data into separate buckets. This unlocks access to more than 100 exabytes of enterprise data stored on NetApp systems for generative AI and analytics on AWS.
-  **Collapse data silos across file, block, and object**
FSx for ONTAP supports NFS, SMB, iSCSI, NVMe/TCP, and S3, so the same data can serve traditional applications, databases, and AI pipelines. Sellers can position a single platform that supports both day-to-day operations and advanced AI workloads.
-  **Turn storage into an accelerator for AI experiments**
ONTAP snapshots and cloning let teams spin up instant, space-efficient copies of datasets and checkpoints for experimentation, testing, and rollback. This means AI builders are working with real data, not manufactured data, shortening iteration cycles and reducing GPU idle time in training pipelines.
-  **Make AI economics more predictable**
Compression, deduplication, and automatic tiering combined with AWS elasticity reduce the cost of storing and serving AI data at scale, helping customers keep infrastructure spend aligned with business value.

AI/ML scenarios to lead with

Sellers win when they turn FSx for ONTAP capabilities into clear AI/ML patterns. Use these scenarios as conversation anchors, then bring your NetApp counterpart in to design the details.

- Scenario 1: Training and tuning ML models on existing enterprise file data**
Customers have years of high-value data sitting in ONTAP-based file systems. With S3 Access Points for FSx for ONTAP, these datasets can be exposed directly to AWS AI services (such as SageMaker and Bedrock) as S3 sources, enabling model training and tuning without exporting or duplicating data.
- Scenario 2: Gen-AI and RAG use cases with FSx for ONTAP as the knowledge base**
Enterprises want chatbots and Q&A experiences grounded in their own data. FSx for ONTAP can act as the knowledge store, with S3 Access Points feeding Bedrock or custom RAG pipelines from existing file content without building complex ETL into S3.
- Scenario 3: Lakehouse analytics and AI with Databricks + FSx for ONTAP**
For customers running Databricks on AWS, FSx for ONTAP provides high-performance storage for data lakes and feature stores, combining file semantics with S3 access for lakehouse patterns. Position this as a way to accelerate analytics and ML workloads while simplifying storage management.
- Scenario 4: AI/ML pipelines with Domino Data Lab and FSx for ONTAP**
When customers use platforms like Domino Data Lab for MLOps, FSx for ONTAP provides a shared, high-throughput data layer for training, experiment tracking, and model artifacts, enabling data science teams to scale experiments without changing how they access data.
- Scenario 5: High-performance AI and EDA workloads at scale**
FSx for ONTAP scale-out file systems deliver very high throughput and IOPS for demanding workloads, like AI training pipelines, EDA, and media rendering—scaling to tens of GB/s throughput, millions of IOPS, and up to 1 PiB of SSD storage in a single managed service.

Benefits your customers will love

FSx for ONTAP delivers an AI-ready data foundation that lets teams focus on models and outcomes instead of plumbing and copy scripts. Position these benefits to line up with buyers’ AI goals.



High-performance, integrated data pipelines



Bring AI to your data—no massive data movement



Real-time access and governance



Elastic scalability for growing AI workloads



Optimized economics for AI operations

The win: Why it pays to sell FSx for ONTAP for AI/ML

Enterprises are racing to embed AI everywhere, but many are stuck at the data layer. When you position FSx for NetApp ONTAP as the AI-ready data foundation, you turn that bottleneck into your biggest selling advantage on AWS.

With FSx for ONTAP, you can:

- Unblock stalled AI initiatives by removing data-movement, performance, and governance barriers—connecting existing datasets directly to AWS AI and analytics services
- Drive high-impact deals around AI and ML workloads, backed by credible performance and cost-efficiency stories from FSx for ONTAP deployments
- Position yourself as the trusted advisor for AI data strategy by using FSx for ONTAP as the entry point to broader NetApp solutions for protection, resilience, and intelligent data infrastructure

And you don’t have to do it alone. Working with NetApp storage specialists gives you access to proven architectures, best-practice patterns, and joint account support so every FSx for ONTAP-backed AI opportunity lands with confidence.

Ready to help customers turn data into AI-ready fuel?

Every stalled or slow-moving AI initiative is an opportunity to introduce FSx for NetApp ONTAP as the data layer that makes AI pipelines faster, safer, and more economical. Sellers who act now earn a seat at the table for the next wave of AI-driven transformation on AWS.

Partner with your NetApp counterpart to identify in-flight or upcoming AI and analytics projects where FSx for ONTAP can counterenterprise data directly to AWS AI services, simplify pipelines, and improve outcomes.

Get started now →



Contact Us

About ONTAP

For over 30 years, NetApp, the Intelligent Data Infrastructure company, has helped leading organizations turn data into a catalyst for innovation, resilience, and growth. At our core is the NetApp data platform—a unified, intelligent foundation built on proven solutions—that connects, protects, and activates data across every cloud and on-premises environment. With a unique native design that integrates with the world’s largest clouds, NetApp helps build an intelligent data infrastructure with our customers to provide the freedom to modernize without vendor lock-in, ensuring consistent performance and governance everywhere. That’s why the world’s most forward-thinking enterprises trust NetApp to maximize your data into a true driver of transformation. Our AI-ready foundation translates complex data into proactive insights and automated actions, simplifying customer operations and freeing IT teams to focus on driving business innovation.

Learn more at [www.netapp.com](#) or follow us on [X](#), [LinkedIn](#), [Facebook](#), and [Instagram](#).



NETAPP, the NETAPP logo, and the marks listed at [www.netapp.com/TM](#) are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners. NA-1289-0726