



Solution Brief

Silk AI Enablement: Training Workloads on Azure

Training AI models, especially large language models (LLMs), requires a storage platform that can handle vast datasets, high-speed data access, and seamless scalability. Traditional cloud storage often struggles to meet these demands, creating bottlenecks that slow training cycles, increase costs, and reduce operational efficiency.

Silk for Training Workloads on Azure

The Silk Platform provides the high-performance storage backbone required for training AI models on Azure. By delivering ultra-low latency, high IOPS, and frictionless scaling, Silk ensures faster training cycles, reduced costs, and a smoother experience for AI practitioners. This technical data sheet outlines how Silk optimizes AI training workflows, detailing its architecture, benefits, and real-world use cases.

Key Features and Benefits

1. High-Speed Data Access

- **Low Latency:** Silk enables microsecond latency for accessing large training datasets, ensuring that nodes are never waiting for data.
- **High Throughput:** With up to 20GB/sec per workload, Silk minimizes data transfer times, even for large-scale AI models.
- **Optimized Checkpointing:** Silk accelerates saving and loading model checkpoints, ensuring quick recovery and resumption of training processes.

2. Scalability for Large Datasets

- **Dynamic Scaling:** Automatically adapts to increasing dataset sizes and compute demands during training.
- **Real-Time Compression:** Drastically reduces storage requirements for datasets and checkpoints, lowering overall cloud costs without impacting performance.
- **Distributed Access:** Enables parallel access to training data and intermediate outputs with a system designed for multi-node training architectures.

3. Data Resiliency and Reliability

- **Symmetric Active-Active Architecture:** Ensures continuous availability of training data, even in the event of infrastructure failures.
- **Zero-Footprint Snapshots:** Take instant snapshots of training datasets and checkpoints for backups or version control without consuming additional storage.

Silk Ensures AI Reliability, Security, and Scalability:

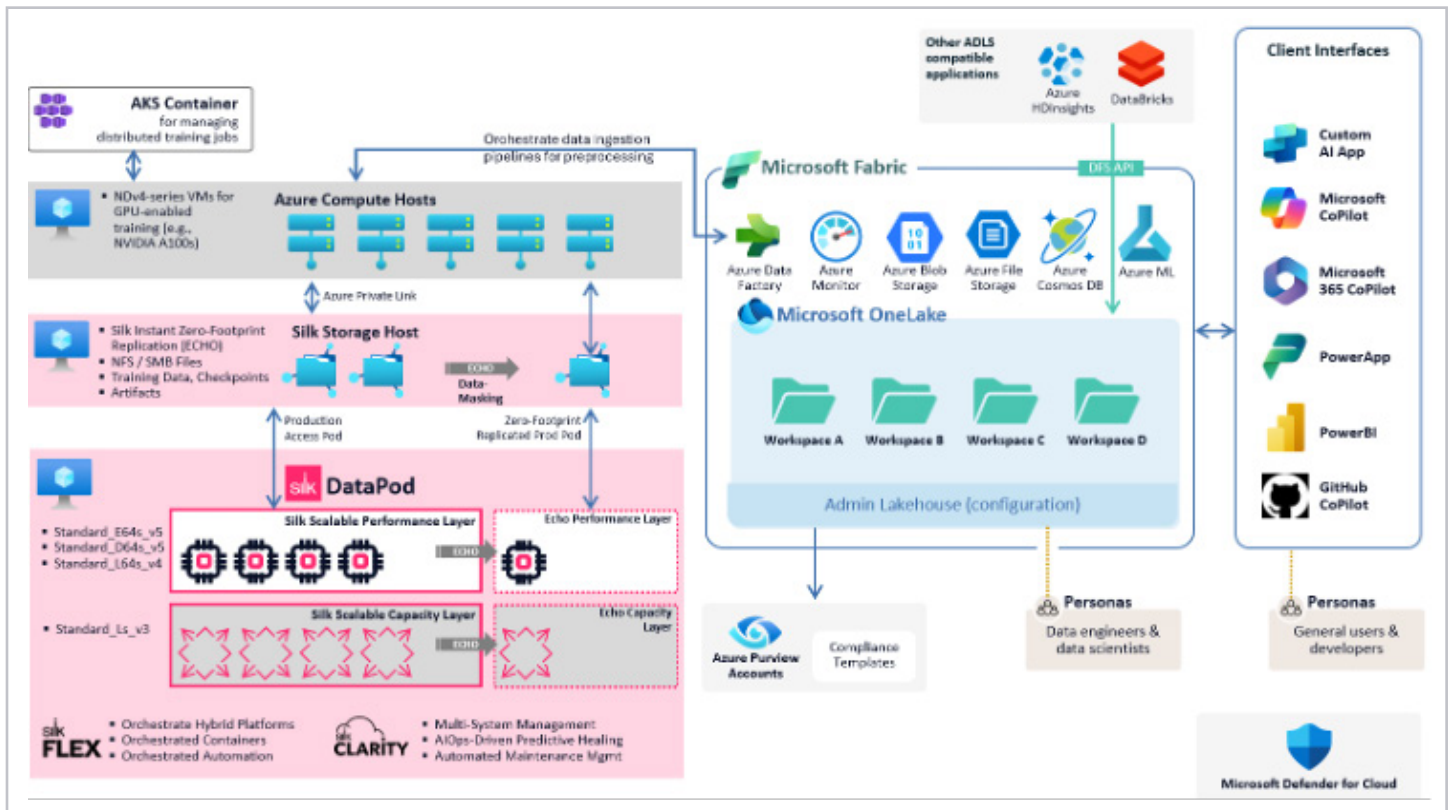
- Sub-millisecond latency for vectorized queries.
- Dynamic scaling for large language models (LLMs).
- Direct integration with Azure AI tools like Copilot and Azure OpenAI.
- Data masking with Redgate to protect sensitive information in AI training and SQL pipelines.

4. Integration with Azure AI Ecosystem

- **Seamless Deployment:** Deploy Silk-backed storage for Azure NDv4-series VMs or Kubernetes clusters (AKS) to power AI training workloads.
- **Native AI Framework Support:** Works seamlessly with PyTorch, TensorFlow, and Hugging Face Transformers, optimizing training pipelines.
- **Efficient Orchestration:** Integrates with Azure Machine Learning for experiment tracking, automated pipelines, and lifecycle management.

Reference Architecture

Silk-Optimized Training Workloads on Azure



Silk-Optimized AI Training Workflow on Azure

This architecture outlines a typical AI training workflow using Silk as the storage backbone:

1. Data Ingestion Layer:

- Raw datasets are ingested into Silk using high-speed connections (e.g., Azure ExpressRoute).
- Data preprocessing pipelines can directly write outputs to Silk storage.

2. Storage Layer (Silk):

- Hosts training datasets, preprocessed data, model checkpoints, and artifacts.
- Provides shared access via NFS/SMB mounts to distributed training nodes.

3. Compute Layer:

- Training Nodes: Azure NDv4-series VMs with NVIDIA GPUs for distributed model training.
- Orchestration: Azure Kubernetes Service (AKS) manages scaling and workload distribution.

4. Networking Layer:

- High-speed networking ensures seamless data flow between Silk and Azure compute resources.

5. Monitoring and Logging:

- Use Silk's built-in analytics and Azure Monitor to track performance metrics and detect bottlenecks.

Use Cases

1. Large-Scale Language Model Training

Challenge: Training LLMs like GPT or BERT requires seamless access to massive datasets and frequent checkpointing without bottlenecks.

Solution: Silk ensures high-speed, low-latency access to datasets and accelerates checkpointing, significantly reducing training time.

2. Distributed Model Training

Challenge: Multi-node training frameworks like Horovod or DeepSpeed require synchronized data access for optimal performance.

Solution: Silk provides consistent, distributed access to data, ensuring efficient scaling across all training nodes.

3. Experiment Tracking and Model Iteration

Challenge: AI experiments require frequent dataset and checkpoint snapshots for reproducibility and debugging.

Solution: Silk's zero-footprint snapshots make it easy to manage dataset versions and checkpoint iterations.

Integration Workflow Example

1. Data Preparation

- Raw datasets are ingested into Silk and preprocessed using Azure Data Factory or Azure Databricks.
- Preprocessed datasets are stored in Silk and indexed for efficient access.

2. Training Execution

- Training jobs are launched on Azure NDv4-series VMs or AKS nodes with direct access to Silk storage.
- Model checkpoints are saved periodically to Silk, ensuring fast recovery and iteration.

3. Post-Training Workflow

- Trained models and metrics are exported from Silk to Azure Blob Storage for deployment.
- Logs and results are archived on Silk for reproducibility.

Technical Advantages of Silk for Training

Workloads:

- **Low-Latency Access:** Accelerates data loading and checkpointing for faster training cycles.
- **Massive Scalability:** Handles petabyte-scale datasets and multi-node training environments seamlessly.
- **Cost Efficiency:** Reduces storage costs with real-time compression and deduplication.
- **High Reliability:** Guarantees data availability with active-active architecture

Performance

Test Configuration

- **Model:** GPT-style model with 6.7 billion parameters
- **Dataset:** 100TB of preprocessed text data
- **Compute:** Azure ND96asr_v4 nodes with 8x NVIDIA A100 GPUs
- **Storage:** Silk Platform connected via NFS

Results			
Metric	Native Azure Storage	Silk Platform	Improvement %
Data Access Latency (ms)	1-10ms	0-2ms	733%
Checkpoint Save Time (s)	120	15	700%
Training Throughput (tokens/sec)	100k	160k	60%
Storage Footprint (TB)	100	70	30% Reduction

Conclusion

The Silk Platform is the perfect solution for training AI models on Azure. By delivering high-performance storage, seamless scalability, and advanced data services, Silk eliminates storage bottlenecks, accelerates training workflows, and optimizes costs. Whether you're training large language models, running distributed AI workloads, or iterating on experiments, Silk ensures your AI training pipelines operate at peak efficiency.

About Silk

Silk fuels AI innovation by enabling real-time access to production data in the cloud. Seamlessly integrating high-performance cloud storage into AI workflows, Silk empowers organizations to enhance innovation while maintaining security, reliability, and control over trusted enterprise data. With Silk, organizations can mitigate and run their most complex business-critical applications in the public cloud, continuously optimizing performance, reliability, and costs. Silk's agile data delivery eliminates the need to copy production data for Dev/Test teams, enhancing flexibility and enabling production data to be leveraged for Generative AI. Backed by over 20 technology patents,

Silk helps customers unlock the full potential of the public cloud with speed and ease. Silk is headquartered outside of Boston, MA.

To learn more, visit www.silk.us