



Rethinking SAN Connectivity

A New Economic & Operational Model
for Secondary Storage

A White Paper presented by [Spectra Logic](#)

Executive Overview

For decades, Fibre Channel has been the assumed fabric for enterprise storage networks. Designed for mission-critical block storage, dual-fabric Fibre Channel architectures delivered isolation, deterministic performance, and end-to-end redundancy at scale.

When tape libraries were introduced into those SAN environments, they inherited that architecture by default.

In many cases, secondary storage connectivity — including backup, recovery, and archive — was adapted from existing fabric infrastructure rather than purpose-built for these workflows. As a result, many modern secondary storage environments continue to operate on connectivity models originally optimized for entirely different workload requirements.

What has changed is not the fundamental behavior of tape — but the role of storage fabrics within the modern data center.

Across enterprise infrastructure, primary storage fabrics have steadily shifted toward Ethernet-based architectures and, in high-performance computing environments, InfiniBand. These networks now dominate the environments where most application workloads operate. Fibre Channel remains widely deployed and continues to perform its intended function effectively. However, its role within the broader data center fabric has gradually narrowed.

At the same time, Serial Attached SCSI (SAS) has evolved significantly. Originally deployed primarily as a direct-attach protocol within rack-scale systems, SAS has matured into a high-throughput, bi-directional transport well aligned with streaming storage workloads. The introduction of optical connectivity and switching architectures extends SAS beyond the short-distance limitations that once defined copper deployments.

These developments raise an important architectural question for secondary storage environments: If Fibre Channel is no longer the dominant fabric for primary storage infrastructure, should it remain the default connectivity model for tape?

Switched optical SAS fabrics introduce a different approach — one that aligns connectivity design with the operational behavior, performance characteristics, and economic realities of modern secondary storage infrastructure.

Table of Contents

Executive Overview	2
The Structural Contraction of Fibre Channel	4
The Fabric Inheritance Problem	5
The Economics of Inherited SAN Architecture	5
Traditional Dual-Fabric Fibre Channel Architecture	6
Optical SAS Fabric — Common Archive Deployment	7
Optical SAS Fabric — High-Availability Deployment	7
Operational Alignment & Skill Efficiency	8
Reframing SAS as a Data Center Fabric	8
Why SAS Now?	9
Architectural Implications	9
Optimized SAS Fabric Architectures	9
Simplified Deployment	10
Performance Scaling	10
High Availability	11
Large-Scale Device Expansion	11
Extended Distance Deployments	12
Infrastructure Strategy for Secondary Storage	12
Conclusion	14
About Spectra Logic	15

Table of Figures

Figure 1 — Inherited SAN Architecture vs. Optimized Optical SAS Fabric	6
Figure 2 — Simplified Single-Fabric Optical SAS Deployment	7
Figure 3 — Dual-Fabric Optical SAS High-Availability Deployment	7
Figure 4 — Simplified SAS Server-to-Library Deployment	10
Figure 5 — High-Performance SAS Fabric Topology	10
Figure 6 — Single-Fabric and Dual-Fabric SAS High-Availability Models	11
Figure 7 — Horizontally Scaled SAS Fabric Architecture	11
Figure 8 — Extended-Reach SAS Fabric Architecture	12



The Structural Contraction of Fibre Channel

Fibre Channel is not failing. It is narrowing.

For years, Fibre Channel SANs served as the central fabric for enterprise storage environments. Primary storage arrays, backup systems, and tape libraries were commonly connected through the same SAN infrastructure, creating a shared, but not workload-optimized, architectural model across the data center.

Over time, the role of Fibre Channel has gradually contracted from both directions.

At the application and primary storage layer, Ethernet-based storage architectures — including iSCSI, NVMe-over-TCP, and distributed object storage platforms — have become increasingly common, driven by their flexibility, cost-effectiveness, and alignment with standard data center networking. In high-performance computing environments, InfiniBand often provides the highest-performance fabric for parallel workloads.

These technologies have reduced the need for specialized SAN domains across much of the infrastructure.

More recently, SAS has matured as a protocol capable of supporting high-throughput streaming storage workloads. Tape drives in particular benefit from predictable, sustained data flows rather than the low-latency, transactional characteristics that Fibre Channel was originally designed to support.

The result is a structural shift in how storage connectivity is deployed across the data center. Fibre Channel remains an effective technology where it is already established, such as legacy SAN environments and certain high-availability block workloads. But the set of workloads requiring a dedicated Fibre Channel fabric has gradually shrunk.

In many environments, tape libraries are among the remaining infrastructure systems connected to a Fibre Channel SAN.

The Fabric Inheritance Problem

Tape infrastructure rarely chose Fibre Channel. In most environments, it inherited it.

When enterprise tape libraries were introduced into data centers that already operated Fibre Channel SANs, attaching those libraries to the existing fabric was the simplest and most logical approach. The SAN infrastructure already existed. The operational expertise was established. The architecture supported the redundancy expectations of enterprise storage environments.

Over time, this inheritance model became normalized. Fibre Channel connectivity for tape libraries became the assumed default.

But the architectural assumptions that once supported that model have changed. Many organizations deploying new infrastructure today are not expanding existing Fibre Channel SAN environments. Instead, they are operating primarily Ethernet-based application infrastructures with specialized connectivity used only where necessary.

Within that context, the continued use of Fibre Channel for tape connectivity is often less a deliberate architectural choice and more a relic of historical precedent.

As secondary storage infrastructure scales and modernizes, that inherited model deserves reconsideration.

The Economics of Inherited SAN Architecture

Most secondary storage environments still rely on Fibre Channel SAN architectures originally designed for high-performance primary storage. While effective in that role, these architectures introduce a cost structure that becomes increasingly misaligned with long-term data retention.

Fibre Channel carries a higher cost profile across the stack, including premium switching infrastructure, specialized HBAs, optics, and cabling. High-availability designs often require dual fabrics, effectively doubling infrastructure and associated costs.

Operationally, Fibre Channel environments add further expense. Design, zoning, and ongoing management require specialized expertise, increasing both labor costs and operational complexity. As environments scale, these requirements compound — driving higher support overhead and longer resolution times.

Scaling the SAN itself introduces additional cost. Expanding connectivity requires more switches, ports, and interconnects, each adding incremental expense. Over time, this creates a cost curve that grows faster than the value of the data being stored — particularly for backup, recovery, and other capacity-oriented workloads.


These workloads do not require the same performance and availability characteristics as primary storage. Continuing to run them on the inherited SAN infrastructure leads to over-provisioning and unnecessary cost.

The result is a higher total cost of ownership driven not only by infrastructure, but by sustained operational overhead. Managing complexity, maintaining specialized skills, and supporting aging SAN environments all contribute to rising costs over time — making legacy architectures increasingly expensive for modern secondary storage use cases.

Traditional Dual-Fabric Fibre Channel Architecture

Fibre Channel environments are typically deployed as dual-fabric architectures to support high availability for primary storage workloads. In contrast, SAS-based archive environments often achieve appropriate resilience using a single switched fabric with dual-port HBAs, with optional dual-fabric designs where additional redundancy is required.

- Two FC switch fabrics
- Dual HBAs per host
- Inter-switch links (ISLs)
- SAN zoning and fabric management
- Redundant pathing across fabrics




(Figure 1 — Inherited SAN Architecture vs. Optimized Optical SAS Fabric)

Optical SAS Fabric — Common Archive Deployment

Common Deployment (Single Fabric)

- Single SAS switching fabric
- Dual-port HBA per host
- No inter-switch links
- Simplified topology (T10-based)
- Host-level path redundancy

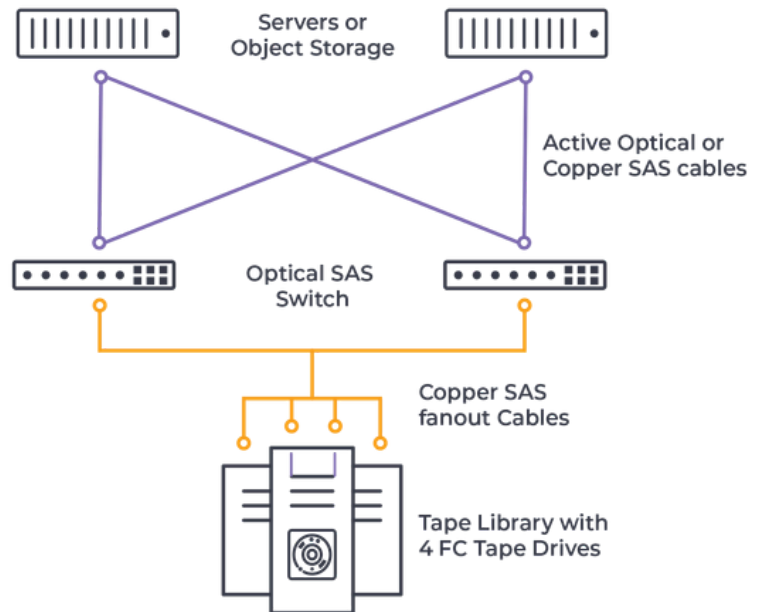


(Figure 2 — Simplified Single-Fabric Optical SAS Deployment)

Optical SAS Fabric — High-Availability Deployment

Dual Fabric Deployment

- Dual SAS switch fabrics
- Dual-port HBA per host
- No inter-switch links
- Reduced component count vs. FC
- Fabric-level redundancy when required



(Figure 3 — Dual-Fabric Optical SAS High-Availability Deployment)

Operational Alignment & Skill Efficiency

Fibre Channel environments require specialized operational expertise. Managing SAN fabrics involve configuring zoning, planning segmentation, coordinating firmware across redundant switches, and managing the fabric infrastructure throughout its lifecycle.

These skills remain common in organizations that maintain large SAN environments. However, as more application workloads migrate toward Ethernet-centric architectures, Fibre Channel expertise is concentrated among a smaller group of specialists.

Modern SAS fabrics align more closely with mainstream server and storage operational models. SAS connectivity involves simpler topology management, fewer configuration layers, and reduced dependency on specialized SAN management tools.

Optical SAS switching extends these advantages across larger areas of the data center without reintroducing the operational complexity of a full SAN fabric.

Reframing SAS as a Data Center Fabric

Bandwidth comparisons between Fibre Channel and SAS are often based on simplified per-lane metrics that do not fully reflect how each technology is implemented in practice.

In Fibre Channel environments, each enterprise tape drive typically consumes a dedicated switch port, requiring eight ports for an eight-drive library — or sixteen in dual-fabric configurations. By contrast, SAS connectivity is typically implemented using x4-wide ports, allowing a single port to service multiple drives. This enables the same eight-drive library to be supported with as few as two ports — or four in dual-fabric configurations — improving port efficiency and reducing overall infrastructure requirements.

Equally important, SAS is a fully bi-directional protocol. Like Fibre Channel, it supports simultaneous transmit and receive operations, enabling deterministic, sustained data flows that closely align with tape-streaming behavior.

Historically, the primary limitation of SAS fabrics was physical distance. Copper-based connections typically supported cable lengths of only a few meters, constraining deployment flexibility.

Optical SAS connectivity removes this constraint. Modern implementations extend connectivity to approximately 100 meters, enabling SAS fabrics to span multiple racks or rows within the data center. Through switched fabric architectures, SAS connectivity can be extended further across larger portions of the data center while preserving native protocol characteristics and maintaining operational simplicity.

Why SAS Now?

SAS has long provided reliable connectivity for storage devices within rack-scale systems, but several developments have expanded its architectural role.

Advances in switching technology allow SAS fabrics to scale across larger device counts while maintaining predictable throughput. Optical connectivity removes the distance limitations that historically confined SAS to direct-attach deployments. At the same time, the growth of large-scale secondary storage environments has increased demand for fabrics optimized for sustained streaming workloads.

Together, these developments allow SAS to function not only as a device interconnect but as a practical data center fabric for secondary storage infrastructure.

Architectural Implications

As data center connectivity evolves, secondary storage infrastructure can increasingly be designed to optimize tape's operational behavior rather than be constrained by the inherited assumptions of legacy SAN architectures.

Connectivity fabrics can be simplified to match the streaming performance characteristics of tape workloads. Secondary storage connectivity can scale using switching fabrics optimized for high-throughput sequential data movement rather than low-latency transaction processing.

Extending SAS connectivity through optical switching also allows secondary storage fabrics to span meaningful distances within the data center while preserving the operational simplicity that has historically defined SAS deployments.


These developments enable a new category of infrastructure: optimized SAS fabrics, designed specifically for large-scale data preservation environments.

Optimized SAS Fabric Architectures

As switched optical SAS fabrics extend SAS connectivity across larger portions of the data center, several architectural patterns emerge that align closely with the operational behavior of secondary storage environments. These deployment models illustrate how switched optical SAS fabrics enable infrastructure designs purpose-built for high-throughput streaming, large data volumes, and operational simplicity.

Simplified Deployment

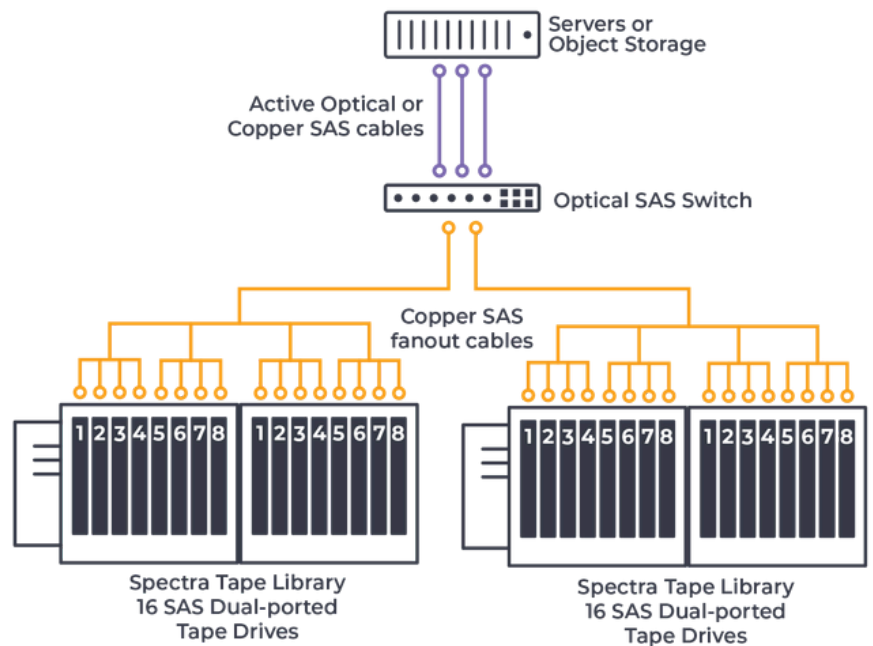
SAS connectivity aligns closely with mainstream server and storage deployment models, allowing the infrastructure to be integrated without the specialized configuration layers typically associated with SAN environments. As a result, new secondary storage systems can often be deployed more quickly while reducing integration complexity and operational overhead.



[Figure 4 — Simplified SAS Server-to-Library Deployment]

Performance Scaling

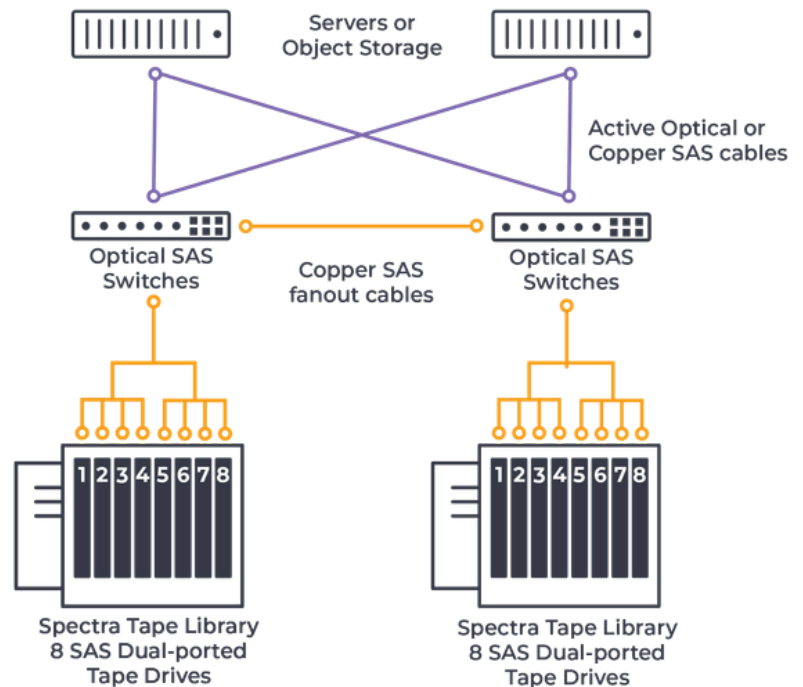
High-density switched SAS fabrics enable parallel aggregation of multiple tape drives within a single switching domain. By distributing streaming workloads across multiple drives, these architectures deliver predictable, high-throughput performance while avoiding the infrastructure duplication typically associated with dual-path SAN environments to distribute load.



[Figure 5 — High-Performance SAS Fabric Topology]

High Availability

Single-fabric SAS deployments, paired with multi-port host bus adapters, provide host-level resilience suitable for many secondary storage environments. Where additional redundancy is required, dual-fabric SAS architectures can also be implemented, enabling higher availability while maintaining a simpler design with fewer infrastructure components than traditional Fibre Channel SANs.




[Figure 6 — Single-Fabric and Dual-Fabric SAS High-Availability Models]

Large-Scale Device Expansion

Switched SAS fabrics can scale across large numbers of hosts and storage devices within a single fabric. By enabling horizontal expansion across multiple switches, these architectures allow environments to grow from tens to thousands of connected endpoints without introducing the zoning complexity or domain segmentation typically associated with Fibre Channel SANs.

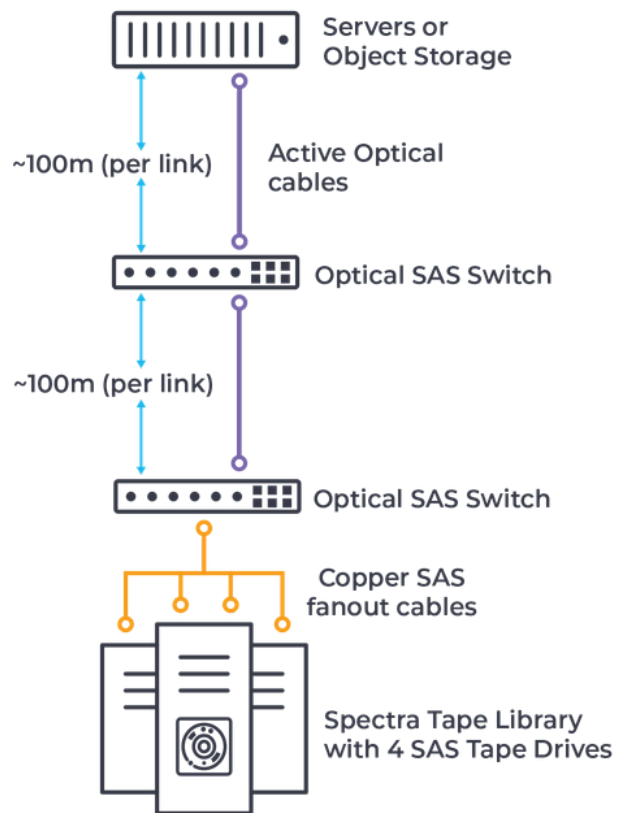
This model supports the natural growth of secondary storage environments, where additional hosts, libraries, and storage devices are incrementally added over time while maintaining a consistent and manageable fabric architecture.



[Figure 7 — Horizontally Scaled SAS Fabric Architecture]

Extended Distance Deployments

Optical SAS connectivity removes the traditional copper distance limitations of SAS, extending native connectivity to approximately 100 meters. In switched fabric architectures, these deployments can scale across racks, rows, or broader data center environments, extending well beyond single-link distances while preserving native protocol characteristics and maintaining operational simplicity.



[Figure 8 — Extended-Reach SAS Fabric Architecture]

These architectural designs illustrate how switched optical SAS fabrics can support a wide range of secondary storage deployment models — from high-performance streaming environments to multi-tenant storage platforms — while maintaining simpler infrastructure than traditional SAN architectures. As organizations expand their secondary storage infrastructure, these design patterns provide a practical foundation for implementing optimized SAS fabrics that align connectivity architecture with the operational behavior of modern data preservation systems.

Infrastructure Strategy for Secondary Storage

The architectural patterns described above illustrate how switched optical SAS fabrics can support a range of secondary storage deployment models — from high-throughput streaming environments to segmented multi-tenant storage platforms. These designs demonstrate how connectivity can be designed around the operational characteristics of tape-based storage systems.

Tape remains one of the most durable and cost-efficient technologies for long-term data preservation. As secondary storage infrastructures evolve, the connectivity architectures that support them must evolve as well.

Switched optical SAS fabrics represent a meaningful step in that evolution. By extending SAS beyond rack-level connectivity and enabling scalable switching architectures, optical SAS allows environments to deploy fabrics optimized for their specific workload characteristics.

Rather than inheriting connectivity models from traditional SAN infrastructure, organizations can now match secondary storage design more directly with the operational behavior of tape-based systems.

Conclusion

Tape remains one of the most durable and cost-efficient technologies for long-term data preservation. As secondary storage infrastructures evolve, the connectivity architectures that support them must evolve as well.

Switched optical SAS fabrics represent a meaningful step in that evolution. By extending SAS beyond rack-level connectivity and enabling scalable switching architectures, optical SAS allows environments to deploy fabrics optimized for their specific workload characteristics.

Rather than inheriting connectivity models from traditional SAN infrastructure, organizations can now match secondary storage design more directly with the operational behavior of tape-based systems.

Platforms such as the [Spectra Logic OSW-2400 Optical SAS Switch](#) demonstrate how this approach can be implemented in modern environments, enabling scalable SAS fabrics that extend across the data center while preserving the simplicity and efficiency of SAS connectivity.



As organizations continue to expand their secondary storage infrastructure, the question is no longer which fabric works.

It is about which fabric best aligns with the architecture of modern data preservation systems — and whether secondary storage infrastructure should continue to inherit connectivity models designed for entirely different workloads.

About Spectra Logic

Spectra Logic modernizes IT infrastructures to preserve, protect, and defend data from days to decades, whether on-premises, in a single cloud, across multiple clouds, or in all locations simultaneously. Our cost-effective solutions enable organizations to manage, migrate, and store long-term data efficiently, ranging from terabytes to exabytes, with features that make them resilient against ransomware.

To learn more, visit www.SpectraLogic.com



Copyright ©2026 Spectra Logic Corporation. All rights reserved worldwide. Spectra and Spectra Logic are registered trademarks of Spectra Logic. All other trademarks and registered trademarks are property of their respective owners. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. All opinions in this white paper are those of Spectra Logic and are based on information from various industry reports, news reports and customer interviews.