Modernizing Enterprise Data Infrastructure with Cisco Unified Edge and Qumulo

Unlock the full value of your unstructured data—wherever it's created—by unifying edge, core, and cloud into one consistent platform, built on Cisco Unified Edge and the Qumulo Data Platform. Together, they deliver the performance, scalability, and simplicity needed to turn data into real business outcomes.

Benefits

- Ensure durability with local write acknowledgment, eliminating dependence on WAN round-trips for high-value data
- Enable low-latency access by processing and serving data directly at the edge
- Consolidate diverse workloads including analytics caches, collaboration data, video, endpoint backups, VDI, and operational datasets on a single edge platform
- Centralize unstructured data, enterprise application caching, and shared services to support observability, security analytics, and compliance
- Strengthen resilience with multi-availability zone disaster recovery and Tier-4-class protection against regional outages

Overview

Unstructured data is no longer confined to the data center — it's created everywhere, from core systems to distributed edge locations. However, most organizations still manage these environments as disconnected silos, increasing complexity, latency, and risk.

The Cisco Unified Edge and Qumulo solution changes that model by delivering a single, unified data platform across edge, core, and cloud. Built on Cisco's modular edge infrastructure and powered by Qumulo Cloud Data Fabric, the solution enables you to process data locally, maintain global consistency, and scale seamlessly into the cloud.

The result is a simpler, more resilient architecture that lets you turn distributed data into actionable insight without adding operational overhead.

Trends and challenges

Data everywhere, control nowhere

Unstructured data is now the backbone of enterprise operations, and it's growing faster than most infrastructures can handle. High-resolution media, sensor feeds, machine logs, and AI training datasets routinely generate terabytes daily. Analysts project global datasphere growth will exceed 175 zettabytes by 2025 (IDC Global DataSphere), with the majority of it unstructured.

The pressure isn't just about volume. Real-time decision-making, AI inference at the edge, and compliance-driven retention mean enterprises must access, process, and protect data simultaneously across locations that were never designed to work together.

Distributed operations demand more than legacy infrastructure can give

Consider a typical enterprise footprint: two regional data centers, thirty distributed edge or campus locations, and a cloud environment for disaster recovery. Keeping that ecosystem running cohesively requires infrastructure that is:

- Highly available across edge, core, and cloud
- Scalable enough to absorb rapid, unpredictable data growth
- Secure and auditable without compromising operational agility
- Capable of running multiple workloads simultaneously without performance tradeoffs

Legacy storage architectures weren't built for this. They can't deliver local performance at the edge while maintaining global consistency. They can't support AI pipelines and real-time analytics without expensive, fragile workarounds. And they can't keep data available during regional outages or cyber incidents.

The result: enterprises are forced to choose between speed and resilience when the business demands both.

How it works

Solution architecture

The architecture is built for distributed enterprises operating at scale — typically organizations with multiple regional data centers, dozens of edge or campus locations, and a need for cloud-based disaster recovery. Customers can deploy this architecture incrementally, starting at the edge or core and expanding outward.

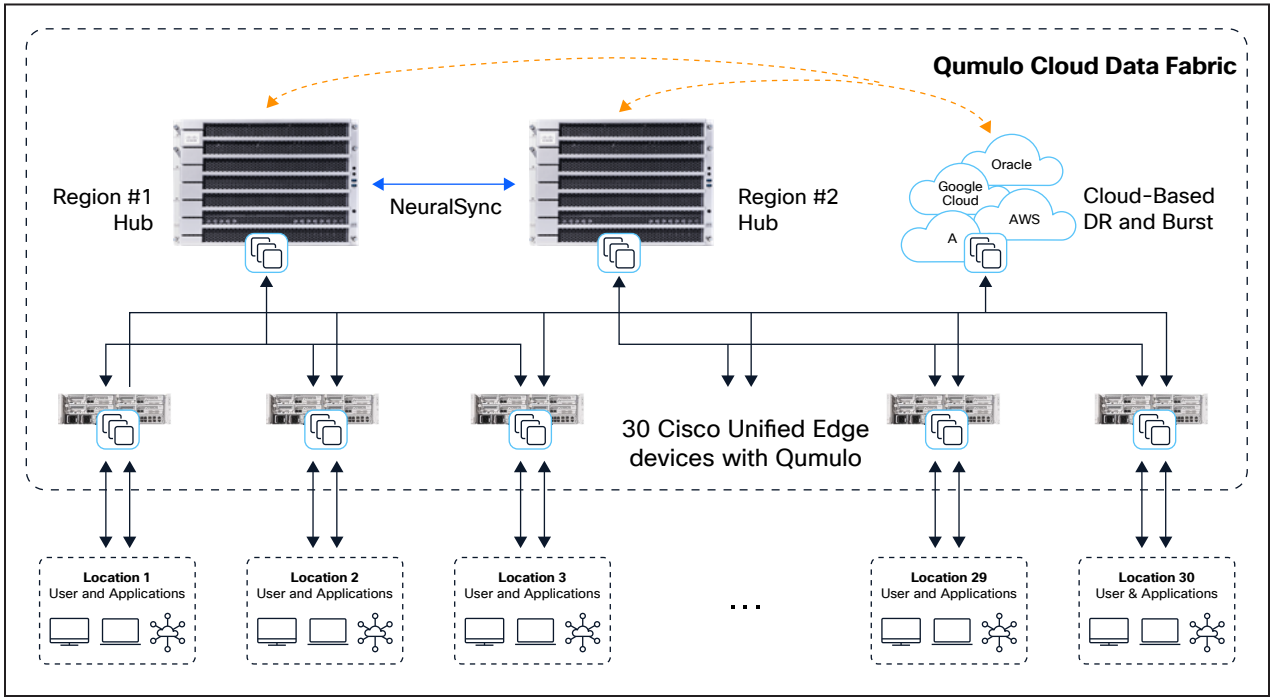

Figure 1. Example reference deployment consisting of two primary data-centers and thirty distributed branch, campus, or operational edge locations, along with a cloud-based disaster recovery and elastic compute environment, selected for geographic isolation.

At the edge: local performance without compromise

At each distributed location, Qumulo runs on Cisco Unified Edge systems with high-speed networking and NVMe storage. Edge deployments are purpose-built to handle the demands of distributed operations without relying on WAN connectivity for every transaction. Key capabilities include:

- **Durable local writes:** Data is acknowledged and protected locally, ensuring resilience even during WAN outages.
- **Local access and processing:** Users and applications get low-latency access to data without round-tripping to a central data center.
- **AI and accelerated compute:** Edge nodes can be configured with GPU acceleration for real-time inference and on-site analytics.
- **Multi-workload consolidation:** A single edge-cluster handles video, collaboration data, endpoint backups, VDI, analytics caches, and operational datasets simultaneously.

At the core: aggregation, intelligence, and control

Two primary data centers serve as dual aggregation hubs, receiving block-level data streams from edge locations over QoS-aware networks. This isn't passive storage — the core is where compute-intensive and cross-domain workloads are concentrated. Core functions support:

- Centralized unstructured data repositories and long-term archives
- GPU-accelerated AI training and large-scale analytics
- Enterprise application caching and shared services
- Centralized observability, security analytics, and compliance reporting

The dual-hub design balances edge autonomy with centralized scale, so organizations don't have to choose between distributed agility and centralized governance.

In the cloud: resilience and elastic scale

The third pillar is Cloud Native Qumulo deployed in the public cloud (for example, Amazon Web Services) across multiple availability zones. This isn't a cold backup — it's an active, policy-governed extension of the same platform running on premises. Cloud capabilities include:

- Tier-4-class resilience for regional outage scenarios
- GPU-based cloud instances for burst analytics and AI workloads
- Strict consistency and data governance enforced through Qumulo Cloud Data Fabric policies
- Elastic scale that activates on demand and contracts when workloads subside

One namespace across all of it

Regardless of where data lives — at the edge, at the core, or in the cloud — Qumulo Cloud Data Fabric presents it through a unified global namespace. There are no silos to manually reconcile, no separate data-management consoles for each tier, and no custom scripts to keep data in sync. Policies set at the platform level govern data movement, access control, encryption, and retention everywhere, automatically.

Workloads this architecture supports

The architecture is designed to support multiple enterprise workloads simultaneously without performance tradeoffs:

- **Analytics and AI:** ingest at the edge, train centrally, scale in the cloud
- **Video and sensor data:** local capture with centralized and cloud retention
- **Observability and security data:** distributed ingestion, centralized analysis
- **Application and collaboration data:** low-latency edge access with global consistency
- **Enterprise backup and recovery:** local protection backed by cloud-tier resilience

The Cisco Advantage

No other vendor can deliver validated compute infrastructure at every tier - edge, core, and cloud with the enterprise-grade networking to connect them. Cisco Unified Edge systems are purpose-built for distributed, high-throughput environments, and are certified with Qumulo to eliminate the guesswork of DIY deployments. That means faster time to value, predictable performance across locations, and a single architecture your team can actually operate at scale.

Use cases

Industry	Description
For security, compliance, and governance	• Encryption at rest and in transit, including FIPS-validated options • Role-based access control and audit logging • Tenant and workload isolation across business domains • Supply-chain-conscious infrastructure, leveraging certified enterprise
For business users	• Fast, local access to critical data, even at the edge • AI-driven insights delivered closer to where work happens • Continuous availability during regional disruptions
For IT operations	• Simplified architecture spanning edge, core, and cloud • Predictable network performance with WAN-aware data movement • Unified management and global namespace • Robust disaster recovery outside regional risk zones

Unstructured data is now the backbone of enterprise operations, and it's growing faster than most infrastructures can handle.

Unlock the full value of your unstructured data wherever it's created by unifying edge, core, and cloud into one consistent platform, built on Cisco Unified Edge and the Qumulo Data Platform. Together, they deliver the performance, scalability, and simplicity needed to turn data into real business

Learn more

See what your data infrastructure could look like

Ready to stop managing data silos and start running a unified edge-to-cloud architecture? Explore how Qumulo and Cisco can transform your distributed operations by visiting our webpage. Then contact your Cisco sales representative

C22-6272052-00 06/26