MODERN MEDIA PRESERVATION AND ACTIVATION AT SCALE

Repatriation Strategies for Building Sustainable, Long-Term Content Lakes

Why Long-Term Media Preservation Demands a Rethink

For studios, broadcasters, sports leagues, and streaming platforms, media content is not temporary data—it is enduring intellectual property. It must be preserved for decades, accessed repeatedly, monetized across evolving formats, and increasingly analyzed by AI. Yet much of today's media infrastructure was never designed with that reality in mind.

Public cloud storage offered an expedient answer during periods of rapid growth. Elastic capacity and on-demand services made it a logical choice as media libraries expanded rapidly. Over time, however, content operations have become more complex and are increasingly AI-driven. The economics and operational constraints of cloud-only models have become more apparent, especially for long-lived, high-resolution media assets that are accessed frequently and retained indefinitely. As libraries grow, organizations encounter rising and unpredictable costs, latency that disrupts production timelines, growing concerns about content sovereignty, and increasing friction when data must be recalled at scale. What was once convenient now limits flexibility.

In response, many media organizations are deliberately repatriating core content back under their own control—not as a step backward, but as a way to establish long-term, low-cost, and sustainable preservation while still supporting modern, cloud-compatible workflows.

This guide explores how content repatriation and modern content lake architectures enable media organizations to preserve content for the long term, control costs and power consumption, and create a foundation for future AI-driven workflows—without sacrificing accessibility or operational agility.

CLOUD AT SCALE: WHERE MEDIA REALITY PUSHES BACK

The current state of many large media environments is the result of years of rational decisions made under pressure—rapid content growth, new formats, expanding distribution channels, and the need to move quickly. Public cloud storage played an important role in enabling that growth. But scale demands a reassessment of these decisions.

As libraries expand and workflows become continuous, data-intensive, and increasingly AI-driven, the assumptions that once made cloud storage attractive are beginning to strain under operational realities. Cost, performance, accessibility, and control—concerns that were once secondary—now shape day-to-day execution.

THE ACCIDENTAL COMPLEXITY OF MODERN MEDIA ESTATES

Most large media organizations operate across fragmented content estates that span public cloud services, on-premises infrastructure, and long-standing archives. This complexity was not intentionally designed. It accumulated over time as content volumes grew and workflows evolved.

As a result, basic operational questions become harder to answer: where should critical assets reside, which versions are authoritative, how quickly can content be retrieved, and what will the access cost be under deadline pressure? This fragmentation is not a failure of discipline or governance. It is the natural outcome of applying generalized infrastructure models to long-lived, frequently reused media assets.

WHY CLOUD WORKED—UNTIL SCALE CHANGED THE EQUATION

Cloud storage earned its early adoption by solving real problems. On-demand capacity, elastic growth, and relief from managing physical infrastructure made it a practical choice for rapidly expanding media libraries.

Over time, however, media workloads diverged from the assumptions underlying public cloud economics. Media content is not ephemeral. High-resolution assets are retained for decades, reused frequently, and accessed across a wide range of workflows—from editorial and remastering to localization, monetization, and AI-driven analysis.

As archives grow, costs rise in ways that are difficult to predict. Latency, once a tolerable inconvenience, becomes a production bottleneck. Teams accustomed to high-throughput local access now wait hours—or longer—when content resides in deep cloud tiers, often leading to unplanned retrieval or egress costs when access becomes urgent.

WHEN EFFICIENCY TURNS INTO FRICTION

At scale, cloud storage also alters the relationship between media organizations and their content. Access patterns become opaque. Workflows grow brittle. Infrastructure decisions increasingly constrain operational flexibility.

Concerns around data and content sovereignty, cyber resilience, and operational control move from theoretical to practical. These pressures are most acute when content must be accessed frequently, at scale, and under time pressure—conditions that define modern media operations.

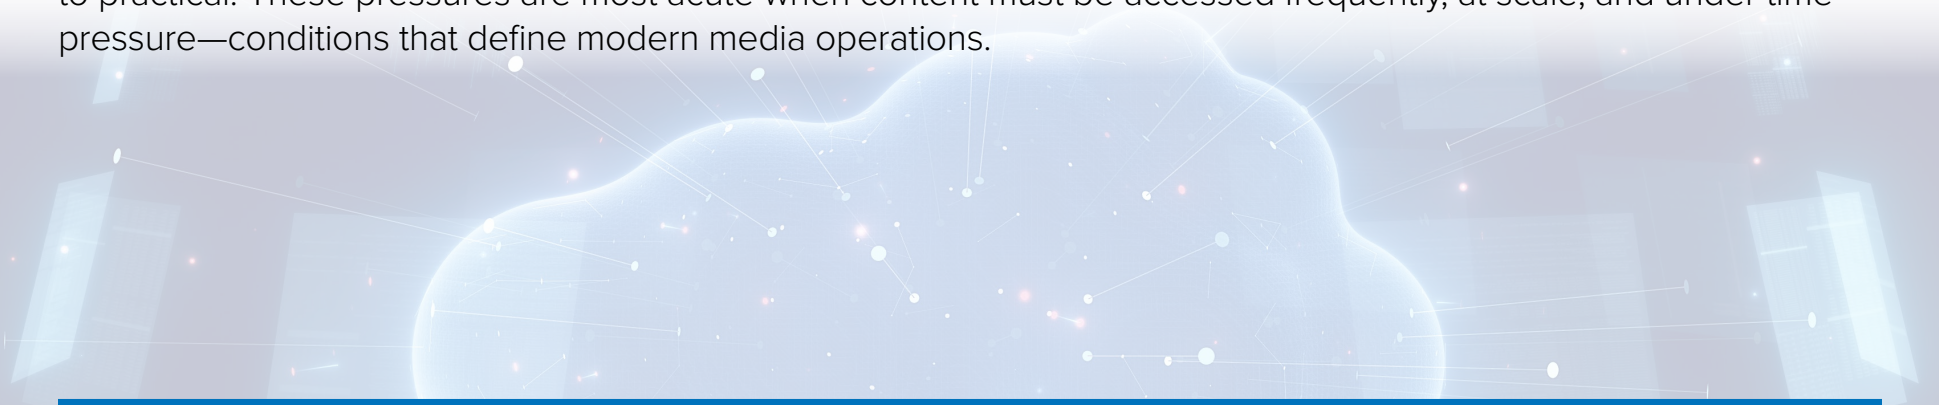

When every interaction with an asset has a cost, creativity and exploration become constrained.

Public cloud storage works best for workloads that are bursty, distributed, or short-lived. Media archives are none of these things. They are permanent, continuously growing, and foundational to both daily operations and long-term monetization. Understanding where cloud models strain under media reality is the first step toward designing architectures that restore clarity, control, and accessibility—without abandoning flexibility.

RE-CENTERING CONTENT AROUND CONTROL, CERTAINTY, AND WORKFLOW

The shift underway in media infrastructure to answer these challenges is not a retreat from modern platforms, nor a rejection of cloud technologies. It is instead a move toward more intentional design. Media organizations are re-evaluating where their most valuable content should reside, under whose control, and with what workflow agility and cost benefits.

When preservation and access are re-centered within architecture owned and operated by the organization, the effects are immediate and tangible: highly predictable storage costs, fewer public cloud access overruns, and most importantly, performance aligns with production needs.

This approach is often implemented through a hybrid content lake that combines cloud-compatible object storage, integrated cold tiers, and policy-driven data movement across different performance and cost tiers. The result is not a return to legacy infrastructure, but a modern, cloud-oriented environment explicitly created for long-term, high-value media assets.

THE OPERATIONAL BENEFITS ARE IMMEDIATE

Once content is placed within infrastructure designed for media realities, several advantages come into focus:

Predictable economics: Storage costs stabilize and are no longer directly tied to access frequency, retrieval urgency, or opaque pricing models.

Consistent, low-latency access: Restores become quick and local, supporting editorial workflows, quality control, re-mastering, and time-sensitive production needs.

Improved visibility and management: Content is easier to organize, enrich, and govern as part of a unified environment rather than across disconnected platforms.

Freedom to experiment with AI: AI and automation initiatives can operate against large content sets without triggering punitive retrieval or egress fees.

Sustainable long-term preservation: Infrastructure optimized for durability and energy efficiency supports retention measured in decades, not billing cycles.

Beyond these operational gains, a more fundamental shift is underway. By placing preservation and access under direct control of your operational teams, decisions about archive growth, catalog enrichment, digitization initiatives, and content reuse are no longer constrained by short-term cost calculations or platform limitations. They become strategic investments in intellectual property, creative flexibility, and future monetization.

IMPLEMENTING A MODERN CONTENT LAKE

The notion of a modern content lake is not simply another name for a larger archive or a different class of storage. It is a deliberately engineered, multi-tiered ‘content operations backplane’ designed to align storage, preservation, and workflow into a single, coherent system. Its purpose is straightforward: to unite your production workflows and give media organizations the scalability, performance, and the economics required for long-term preservation—without forcing tradeoffs between them.

At the core of a content lake is object storage that works like public cloud but is on-premises. Purpose-built for massive unstructured datasets, object storage enables scale without the complexity of traditional file systems. It supports parallel access, can deliver multi-site resiliency, integrates cleanly with media asset management platforms and AI pipelines, and provides a cloud-compatible interface that allows workflows to evolve without re-architecting the underlying storage.

Critically, this content lake should integrate a cold tier that combines the benefits of tape and object storage, delivering exceptionally low cost per terabyte and unmatched energy efficiency for long-term retention. In a modern content lake, this cold tier is not a separate system or workflow; it operates within the same logical environment and namespace as your active storage, allowing content to move between performance and cost tiers through policy or in response to workflow actions, without requiring manual intervention.

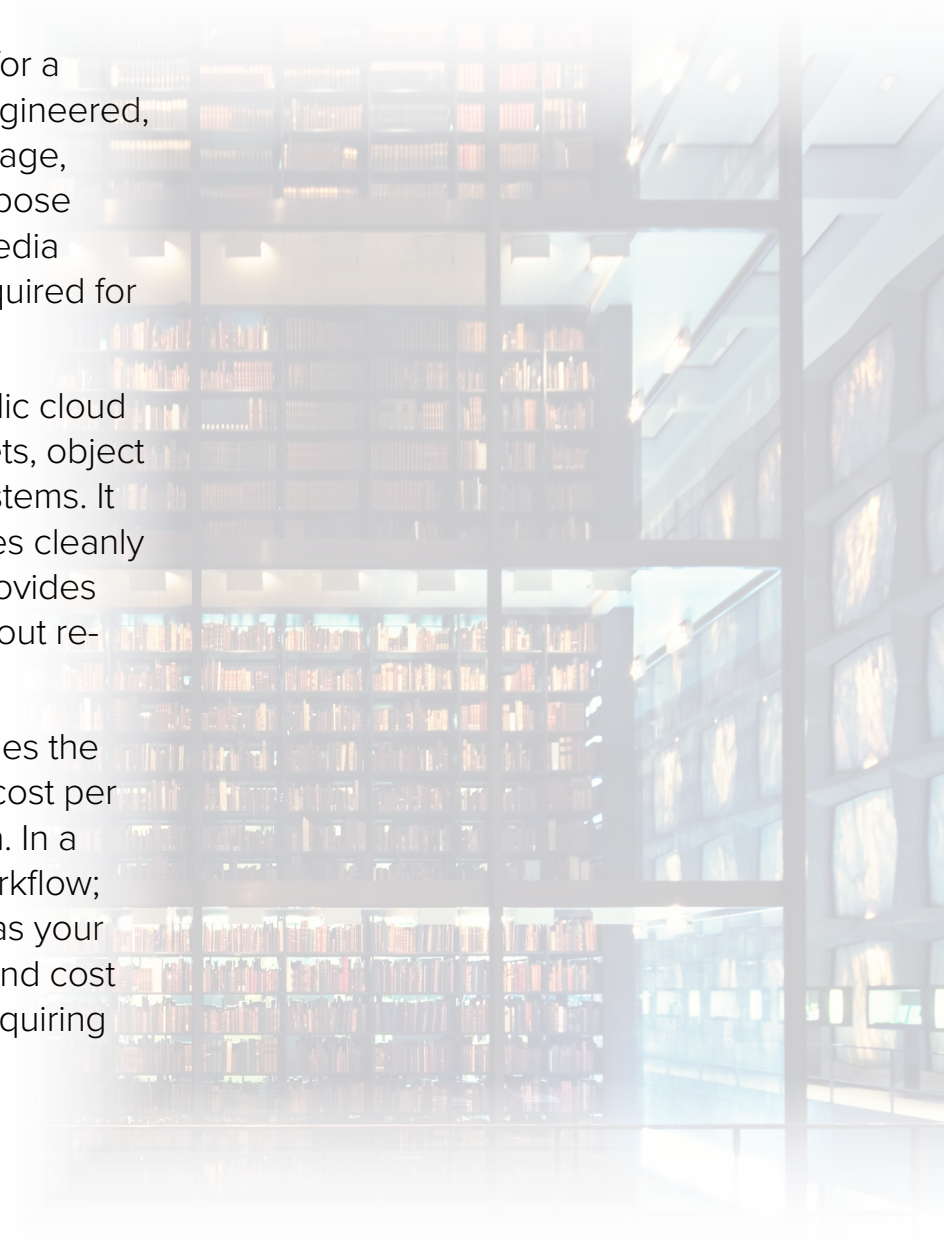

The result is a unified experience. Whether content is active, infrequently accessed, or deeply archived, it remains visible, referenceable, and governed through the same namespace. Files can be indexed, enriched, and analyzed without requiring teams to manage storage boundaries or understand where content physically resides.

A content lake provides a single home for everything—from the newest camera originals to decades-old archival masters.

THE ECONOMICS THAT MAKE REPATRIATION IMPOSSIBLE TO IGNORE

Few transformations in media operations offer savings as dramatic and immediate as content repatriation. Once cloud-based content is moved into an on-premises content lake, the cost curve changes dramatically.

Eliminating egress fees alone can transform annual budgets. For organizations that frequently reuse catalog content—such as sports leagues, studios with evergreen titles, and broadcasters with decades-long archives—the financial relief is substantial. Storage bills that once fluctuated unpredictably stabilize into a fixed, manageable pattern.

The cold tier further amplifies the savings. Tape storage consumes a fraction of the power required by disk systems and offers unmatched long-term durability. Over time, many organizations see a 60–90% reduction in total storage spend, freeing budget for content creation, rights acquisition, infrastructure modernization, and AI development.

This is not theoretical. It's the direct result of aligning storage architecture with media's actual needs, rather than the general-purpose design of cloud platforms.

Quantum

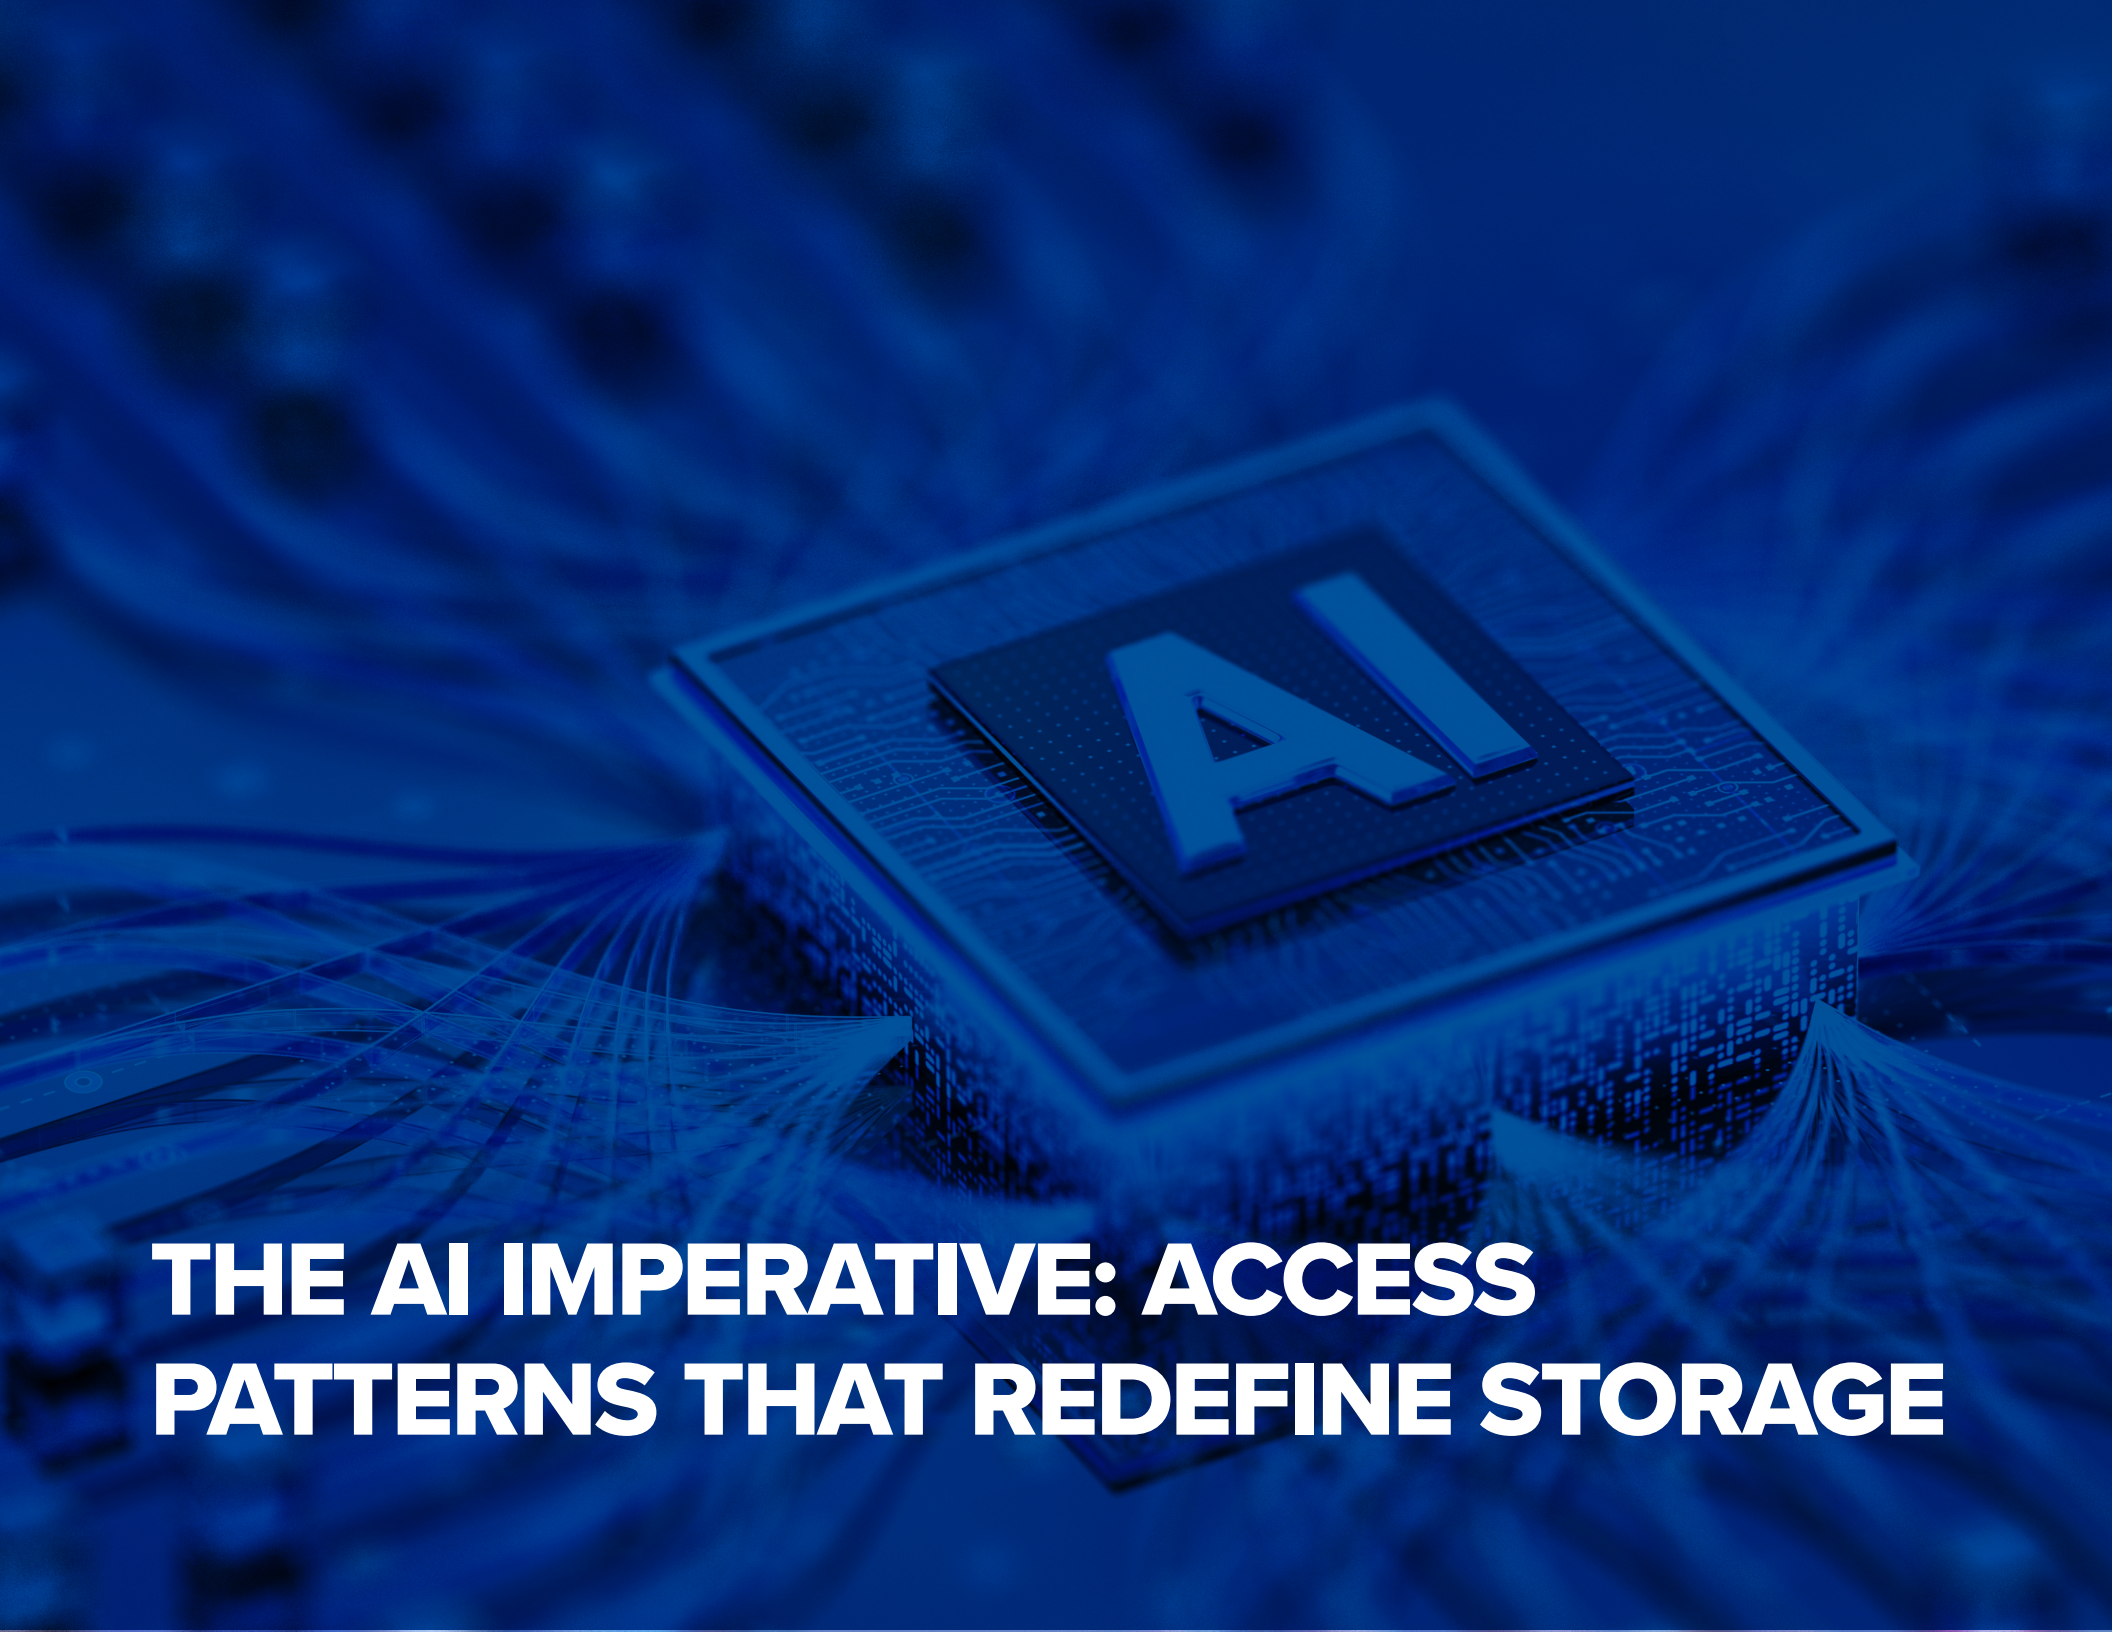

THE AI IMPERATIVE: ACCESS PATTERNS THAT REDEFINE STORAGE

AI is completely reshaping how media organizations interact with their content libraries. Where archives were once accessed intermittently and predictably, AI-driven workflows introduce continuous, high-throughput demand. Automated quality control, object and facial recognition, metadata enrichment, semantic search, and model training all depend on repeated, large-scale access to the same source material.

This shift fundamentally changes storage requirements. AI systems are not optimized for occasional retrieval; they assume frequent scanning, parallel reads, and sustained throughput. The cost and performance implications of these access patterns are significant.

General-purpose cloud storage models are poorly aligned with this reality. Retrieval fees, cold-tier access delays, and usage-based pricing structures can quickly constrain AI initiatives, limiting how often content can be analyzed or forcing teams to narrow the scope of experimentation. In practice, this places artificial ceilings on innovation precisely where flexibility is most valuable.

When content is placed within a content lake architecture designed for persistent, high-frequency access, these constraints are removed. AI pipelines can scan large libraries repeatedly without triggering incremental penalties. GPU resources remain fully utilized rather than idle while waiting on data movement. Content discovery accelerates, and new workflows—such as automated localization, restoration, catalog analysis for monetization, or rights intelligence—become economically and operationally viable.

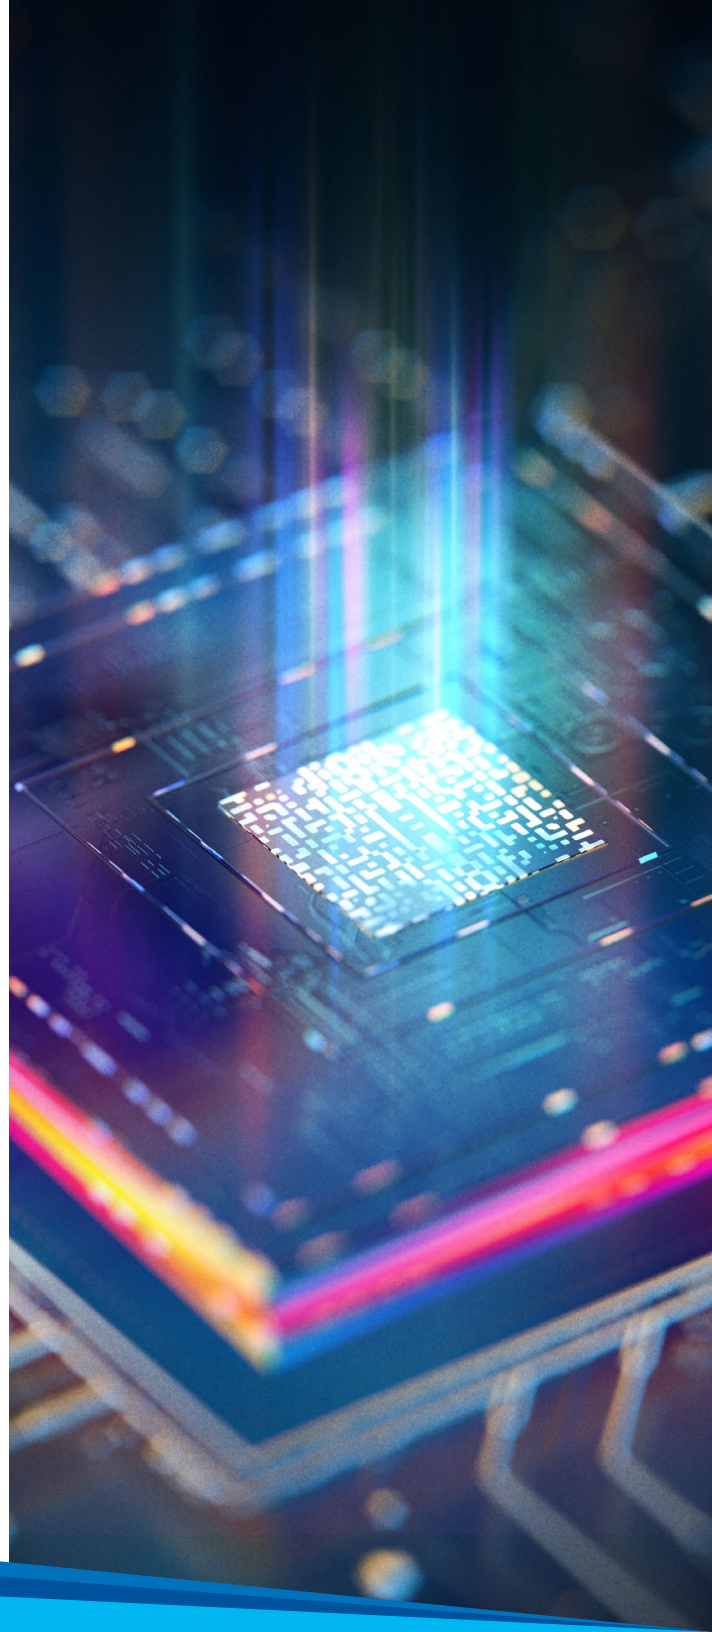

RESTORING INTENTIONALITY TO MEDIA PRESERVATION AND OPERATIONS

As media organizations operate at greater scale and under increasing pressure, one principle has emerged as decisive: infrastructure choices must be intentional. The complexity facing content operations today is not the result of poor decisions, but of years of growth layered onto platforms that were never designed for long-lived, frequently reused media assets.

A content lake approach restores that intentionality: by re-centering preservation and access within architectures designed for media realities, organizations regain control over where content lives, how it is accessed, and what it costs to use. Storage economics become predictable. Workflows stabilize. Preservation horizons extend confidently into the future. Decisions about content reuse, enrichment, and monetization are no longer constrained by uncertainty or platform-imposed penalties.

Just as importantly, this foundation prepares organizations for what comes next. AI-driven workflows, automated analysis, and new forms of content activation depend on fast, repeatable, and scalable access to large media libraries. A content lake enables these capabilities without forcing tradeoffs between performance, cost, and control. It allows experimentation without penalty and growth without architectural reinvention.

With a content lake at the core, media organizations can choose where cloud adds value, how to balance the mix of public and on-premises operations reside—while ensuring that preservation, access, and economics under their own governance.

Modern media preservation is no longer about where content is stored. It is about how confidently it can be preserved, accessed, and activated over time. Organizations that adopt a content lake strategy are not simply modernizing storage. They are establishing a durable operating model—one that delivers predictability today and flexibility for whatever challenges and opportunities come next.

ACTIVE SCALE COLD STORAGE: PRESERVATION BUILT FOR THE LONG TERM AND FOR AI

Long-term preservation places unique demands on infrastructure. Media content must remain durable, accessible, and economically viable for decades—often far beyond the lifespan of any single technology generation. At the same time, AI-driven workflows require the ability to scan, analyze, and reprocess vast libraries repeatedly without incurring punitive access costs or excessive power consumption.

Quantum ActiveScale® Object and ActiveScale Cold Storage are designed to support **content repatriation** by enabling organizations to bring large media libraries back on-premises and under their control—without sacrificing the flexibility of cloud-based workflows. By combining cloud-compatible object storage with an integrated cold tier, ActiveScale allows organizations to build modern on-premises content lakes that retain massive volumes of content at extremely low cost, while significantly reducing power and cooling requirements. Content remains visible and governed within a single namespace, even as it moves seamlessly between performance and cost tiers.

This architecture enables practical repatriation at scale. Storage economics become predictable, energy consumption is minimized, and content can be retained indefinitely while remaining readily accessible for AI-driven analysis, reuse, and monetization. For organizations seeking to preserve content for the long term while reducing cost and operational friction, ActiveScale provides a durable foundation for both today's workflows and future innovation.

Visit www.quantum.com/activescale to learn more about how ActiveScale Object Storage and ActiveScale Cold Storage support content repatriation and long-term, low-cost, AI-ready media preservation.

Quantum

Quantum delivers end-to-end data management solutions designed for the AI era. With over four decades of experience, our data platform has allowed customers to extract the maximum value from their unique, unstructured data. From high-performance ingest that powers AI applications and demanding data-intensive workloads, to massive, durable data lakes to fuel AI models, Quantum delivers the most comprehensive and cost-efficient solutions. Leading organizations in life sciences, government, media and entertainment, research, and industrial technology trust Quantum with their most valuable asset – their data. Quantum is listed on Nasdaq (QMCO). For more information visit www.quantum.com.