# Recovery First: Why Archive Protection Must Match the Medium

Why Replication — Not Erasure Coding — Ensures Long-Term Data Resilience on Tape

White Paper | Spectra Logic

# Table of Contents



## 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

# Executive Summary

Organizations today are retaining more data, for longer periods, under greater scrutiny than at any time in history. Regulatory mandates, cyber-resilience requirements, scientific reproducibility, intellectual property protection, and institutional memory have elevated archives from a back-office concern to a strategic asset. At the same time, the infrastructure used to store that data has become increasingly heterogeneous, spanning HDD, SSD, object storage platforms, cloud services, and tape.

In this environment, traditional measures of storage durability — often expressed as abstract probabilities or “nines” — are no longer sufficient. The more important question is no longer how rarely data might be lost in theory, but how reliably it can be recovered in practice.

Durability is necessary, but recoverability determines whether archived data is usable when it matters.

This distinction matters most for archives. Archive recovery rarely happens under ideal conditions. It often occurs years after data ingestion, under operational stress, with constrained resources, and with teams that may not have designed the original system. In these moments, theoretical durability gives way to a more practical concern: recoverability.

This white paper presents a recoverability-first framework for archive protection. It explains why erasure coding is highly effective in always-on disk environments, why applying erasure coding to tape is a category error, and why replication remains the media-appropriate protection model for tape-based archives. The objective is not to advocate for or against any single technology, but to reinforce a foundational architectural principle: data-protection mechanisms must align with the behavior of the storage medium they are protecting.

## From Durability to Recoverability

Durability has long been treated as the defining metric for storage systems. Expressed as a probability of data loss over time, it provides a useful abstraction for comparing systems under controlled assumptions. However, durability metrics often obscure the conditions under which recovery actually occurs.

For archives, success is not defined by how unlikely it is that data might be lost, but by whether complete, usable data can be recovered when required.

Recoverability is the practical ability to restore complete, usable data within an acceptable RTO (Recovery Time Objective) under real-world conditions — including degraded systems, operational stress, and long-term platform change.

This definition reflects several realities that matter deeply to archive operators:

- Recovery often occurs during volatile moments such as audits, legal discovery, cyber incidents, or investigations — not during calm, steady-state operation
- Systems may be partially constrained rather than fully failed
- Shared infrastructure resources may already be under load
- Personnel executing recovery may not be the original system designers
- Platforms, tools, and vendors may have evolved over time

### Implications for Archive Design

Durability metrics implicitly assume idealized recovery conditions: immediate access to all components, full system availability, and homogeneous failure domains. However, archives rarely enjoy such stability. Recovery often occurs while other workloads continue, while compliance timelines loom, or while infrastructure is intentionally segmented for security. Designs that optimize for durability alone may therefore perform well on paper but struggle when recovery must occur under pressure.

A recoverability-first perspective forces architects to ask not just "Is the data still there?" but "Can we get it back completely and predictably when it actually matters?"

Durability is a necessary input to archive design, but it is not the outcome. **Recoverability is the outcome** — and it is shaped as much by architecture and operations as by statistical reliability.

## Recoverability Over Time (A Quiet but Important Idea)

Over long retention periods, it becomes helpful to think not just about recoverability in the moment, but about how long recoverability can reasonably be sustained.

Architectures differ not only in how durable they are, but in how long they remain practically recoverable without extraordinary effort. This can be thought of informally as a system's recoverability horizon: the length of time an organization can expect to recover data with confidence without recreating the original storage system in full.

This horizon is influenced by factors such as:

- Dependency on proprietary layouts or metadata
- Configuration rigidity
- Operational complexity
- Platform and vendor continuity
- The degree of coupling between data and system interpretation

### Why This Matters for Archives

Archive strategies often succeed or fail years after initial deployment. During that time, organizations and personnel change, data volumes grow, and infrastructure evolves. Architectures that preserve independently recoverable units tend to maintain longer recoverability horizons because they reduce dependency on exact system reconstruction. Architectures that tightly bind data to a specific layout, configuration, or interpretation logic may shorten that horizon — even if the underlying media remains intact. Over decades, this distinction becomes more important than incremental gains in storage efficiency.

# Erasure Coding in Its Proper Context

Erasure coding improves durability and storage efficiency by dividing data chunks into fragments, generating parity fragments, and distributing those fragments across multiple devices or nodes. If one or more components fail, missing fragments can be reconstructed from the remaining data and parity.

In this paper, **erasure coding refers specifically to wide-striped, distributed coding schemes** designed for disk- and object-based storage systems — not local parity protection within a single disk group.

Erasure coding is highly effective when the following foundational assumptions are met:

- **Always-on availability:** storage devices are continuously online and reachable
- **Random access and parallel I/O:** fragments can be read concurrently from many devices
- **Routine rebuild workflows:** component failures are expected, and rebuild is a normal background activity
- **Performance headroom:** systems are designed to absorb rebuild load without disrupting service levels

These characteristics are inherent to large-scale HDD and SSD environments, particularly distributed object storage systems. In those contexts, erasure coding delivers strong durability, improved capacity efficiency, and predictable operational behavior.

## Where Erasure Coding Excels — and Why That Matters

In disk environments, rebuild activity is an accepted cost of doing business. Systems are engineered to continuously rebalance data, absorb transient load spikes, and exploit parallelism across many devices. Recovery is typically automated, localized, and transparent to users. These conditions allow erasure coding to fulfill its promise: high durability at scale with manageable operational impact. Within its intended domain, erasure coding is a sound and proven engineering solution.

Understanding where erasure coding excels is critical because it clarifies why forcing it into environments that lack these conditions introduces risk rather than resilience.

## Why Tape Changes the Architectural Equation

Tape occupies a fundamentally different role in the storage hierarchy. It is not “slow disk.” Tape is **sequential, removable, and often offline by design** — qualities that make it uniquely valuable for long-term retention, deep archive, and cyber resilience.

Tape environments introduce characteristics that differ sharply from disk systems:

- Sequential access and mount/position workflows rather than random I/O
- Physical media handling, transport, and vaulting
- Finite shared resources inside the library (robots and drives)
- Retention horizons measured in decades
- Intentional offline survivability and air-gap security

### Operational Consequences of These Differences

The characteristics of tape-based infrastructure fundamentally shape how recovery unfolds. Tape systems are optimized for streaming large volumes of data predictably, not for assembling fragments from many sources simultaneously.

Offline storage, while valuable for security and preservation, introduces latency and coordination steps that must be accounted for in recovery planning. In archive contexts, success depends less on instantaneous access and more on predictable, repeatable workflows that can be executed correctly long after initial deployment.

These characteristics are strengths, not weaknesses because they prioritize recoverability over immediacy. But they also mean that protection models designed for always-on, random-access storage cannot be assumed to translate cleanly.

## Erasure Coding on Tape: A Category Error

Claims that erasure coding is “more durable” or “more efficient” than replication can be valid within disk-centric environments. When applied to tape, however, those comparisons are misleading.

The issue is not whether erasure coding works mathematically. It is whether it supports archive recoverability under real operating conditions.

### Scarce vs. Abundant Resources

In tape libraries, cartridges are abundant, while drives and robotic mechanisms are scarce, high-value resources. Protection schemes that increase mounts, seeks, and cross-media reads shift work away from the abundant resource and onto the constrained ones.

This shift negatively impacts:

- Throughput
- Concurrency
- Restore performance
- Operational headroom during recovery events

### Rebuild Behavior and Degraded RTO

In disk systems, rebuild is a background process executed against online devices. In tape environments, reconstruction may require mounting and scanning across multiple cartridges, consuming drives and robotics for extended periods.

During recovery events — when systems may already be constrained — this behavior impacts real-world RTO, not just theoretical availability. Archive recovery is rarely a single-threaded activity. It often overlaps with ongoing ingest, compliance, or investigation workflows.

### Fragment Coupling and Recovery Friction

Replication produces complete, independent copies. Erasure coding makes recoverability dependent on assembling a sufficient set of fragments, along with the correct metadata and layout rules to interpret them.

In real-world recovery scenarios, fragment-based approaches increase coordination requirements. More decisions must be made correctly, more dependencies must align, and more opportunities for delay or error exist. Over long retention periods, even small increases in procedural complexity compound. Recovery becomes less about reading data and more about orchestrating a precise sequence of operations — often under time pressure. This friction is largely invisible in durability math but very real during recovery.

This mismatch of assumptions is why applying erasure coding to tape is  a category error. It is a disk-centric protection model forced onto a medium with fundamentally different behavior.

## Encapsulation and Technical Obsolescence

For long-term archives, one of the most underestimated risks is technical obsolescence — the inability to recreate the platform, software, or interpretation required to retrieve data decades later.

Replication minimizes this risk because the durable unit is a complete copy. Complete copies are easier to move, audit, preserve, and recover using deterministic processes.

Erasure coding changes the durable unit to a coded set: fragments, parity, metadata, and layout rules that must be interpreted together. This creates encapsulation, increasing dependency on the originating system’s ability to analyze and reassemble those pieces.

### Media Survival vs. Data Usability

Archives fail not only when media degrade, but also when data remain intact but become difficult to interpret or recover. Long-term preservation requires three things to remain viable: the media, the interpretation, and the recovery process. Encapsulation increases the likelihood that one of these elements will drift out of alignment over time.

When erasure-coded archives require static configuration after initialization, the risk deepens. Over a 10–20+ year horizon, hardware generations, library configurations, sites, and operational models will change. A static encoding scheme can effectively lock recoverability to a specific platform, shortening the recoverability horizon even if the media itself remains intact.

Preservation architectures should therefore minimize tightly coupled, system-specific dependencies that are difficult to unwind over time.

## Human and Operational Realities

Archive recovery is rarely executed by the same team that designed the original system. Over long retention periods, staff turnover, organizational change, and evolving priorities are inevitable.

This reality amplifies the value of:

- Simplicity
- Deterministic workflows
- Clear failure domains
- Minimal system coupling

Systems often remain technically durable long after they become operationally fragile.

### Procedural Drift Over Time

Even well-documented systems experience procedural drift. Processes evolve, shortcuts emerge, and institutional knowledge fades. Architectures that rely on implicit assumptions or complex coordination are more vulnerable to this drift. Designs that emphasize clarity and independence age more gracefully because they reduce the cognitive load placed on future operators during high-pressure recovery situations.

Architectures that require deep system knowledge or precise reconstruction of historical configurations increase the likelihood of recovery delays or errors — especially under stress. This is not a question of competence; it is a recognition that human and operational factors are integral to recoverability.

## Why Replication Aligns with Tape

Replication aligns naturally with tape's strengths because it preserves complete, independent copies.

For tape-based archives, replication provides:

- Deterministic recovery without fragment assembly
- Predictable workflows that support RTO planning
- Reduced operational coupling
- Longer recoverability horizons
- Strong alignment with offline and air-gapped strategies

In this context, replication is efficient and intentional. It prioritizes recoverability, simplicity, and independence over metrics that do not translate to archive outcomes.

## Practical Guidance: Matching Protection to the Medium

No single durability model is optimal across all storage tiers. The most resilient architectures apply media-appropriate protection.

**Use erasure coding when:**

- Storage is always-on and online
- Random access and parallel I/O are available
- Rapid rebuild workflows are expected
- Capacity efficiency at scale is a priority
- The medium is HDD or SSD

**Use replication when:**

- Media is removable or offline by design
- Recoverability and deterministic restore matter most
- RTO realism under degraded conditions is critical
- Long-term retention or cyber resilience is the goal
- The medium is tape

The architectural principle is straightforward: match the protection model to the behavior of the medium and evaluate success by recoverability — not abstract durability mathematics.

## Conclusion: Recoverability as the Defining Archive Metric

Erasure coding is a powerful, proven technology in the environments for which it was designed: always-on, random-access disk systems. Tape archives are different by design.

For long-term retention, preservation, and cyber resilience, recoverability over time — not theoretical efficiency — defines success. Applying erasure coding to tape introduces complexity, coupling, and obsolescence risk that can undermine long-term recovery, even when the media itself remains intact.

Replication remains the media-appropriate durability strategy for tape because it prioritizes deterministic recovery, realistic RTO, and long-term independence from specific platforms.

**When archive value matters most, measure success by recoverability — and match the protection model to the medium.**


Toll Free: 800-833-1132 • 303-449-6400 • 6285 Lookout Road • Boulder, CO 80301 USA
V2-012026