DATASHEET

KUMULUS DISTRIBUTED FILE SYSTEM

Scale-Out Performance for AI, HPC & Enterprise Workloads



Kumulus DFS is a high-performance, scale-out storage platform designed for organizations that require performance, scalability, resilience, and efficient data lifecycle management.

Built with a distributed architecture, Kumulus DFS provides a unified storage environment that can scale from a minimum cluster configuration to a large multi-node deployment as performance and capacity requirements grow.

By combining high-performance NVMe storage, scalable capacity expansion, global namespace architecture, and an integrated archive engine, Kumulus DFS helps organizations manage data efficiently across its entire lifecycle.

KEY BENEFITS

LINEAR SCALE-OUT
Expand performance and capacity by adding nodes as requirements grow.

GLOBAL NAMESPACE
Access distributed data through a single unified namespace, regardless of where the data physically resides.

HIGH-PERFORMANCE NVMe
Parallel data access across the distributed environment for demanding AI, HPC and enterprise workloads.

FLEXIBLE CAPACITY TIER
Combine high-performance NVMe with scalable JBOD expansion to place data on the right storage tier.

INTEGRATED ARCHIVE ENGINE
Built-in archive capability for policy-based migration, scheduling, automation, and lifecycle management.

MINIMUM 3-NODE EC CLUSTER
Distributed architecture with erasure coding for scalable performance and data resilience.

KUMULUS DFS vs TRADITIONAL SCALE-UP NAS

TRADITIONAL SCALE-UP NAS	KUMULUS DFS
Scale-up architecture	Distributed scale-out architecture
Performance dependent on controller	Parallel performance across nodes
Multiple file systems or islands	Global unified namespace
Bottlenecks as workload grows	Linear performance scaling with node expansion
Performance and capacity tightly coupled	Independent performance and capacity scaling
Separate archive infrastructure	Integrated Archive Engine / ARC
Complex licensing models	Perpetual licensing
Archive capacity licensing may increase	No capacity-based limitation for Archive license
Expansion may require redesign	Expand the distributed environment as required

PERFORMANCE
High throughput and parallel access for demanding workloads.

SCALABILITY
Linear scale-out architecture that grows with your business and data.

RESILIENCE
Erasure coding and distributed architecture ensure data protection and availability.

DATA LIFECYCLE
Complete lifecycle from active data to archive and long-term preservation.

KUMULUS DFS ARCHITECTURE

DISTRIBUTED BY DESIGN. BUILT FOR PERFORMANCE, SCALE AND RESILIENCE.

Kumulus DFS is a modern distributed file system that delivers high performance, scalable capacity, and intelligent data lifecycle management across your entire organization.

From high-performance workloads to long-term preservation, Kumulus DFS provides one unified platform for every stage of your data lifecycle.

1 GX-IO CLUSTER

GLOBAL NAMESPACE & INTELLIGENT ACCESS LAYER

- Global unified namespace
- Metadata management
- Intelligent I/O routing
- Centralized access to distributed resources
- Simplified data access for all workloads

2 UNF NODES (PERFORMANCE TIER)

HIGH-PERFORMANCE ACTIVE DATA LAYER

- Distributed scale-out architecture
- NVMe-based performance tier
- Parallel data processing
- Designed for high IOPS and throughput
- Scale performance linearly by adding nodes

3 JBOD EXPANSION (CAPACITY TIER)

CAPACITY EXPANSION LAYER

- Scalable JBOD expansion for massive capacity growth
- Online expansion without downtime
- Cost-optimizer tier for cold data
- Independent capacity scaling
- Ideal for large-scale data repositories

4 ARCHIVE ENGINE (ARC)

INTEGRATED DATA LIFECYCLE MANAGEMENT

- Policy-based data lifecycle management
- Automated data migration
- Archive scheduling & automation
- Data verification & integrity checks
- Move inactive data to long-term storage

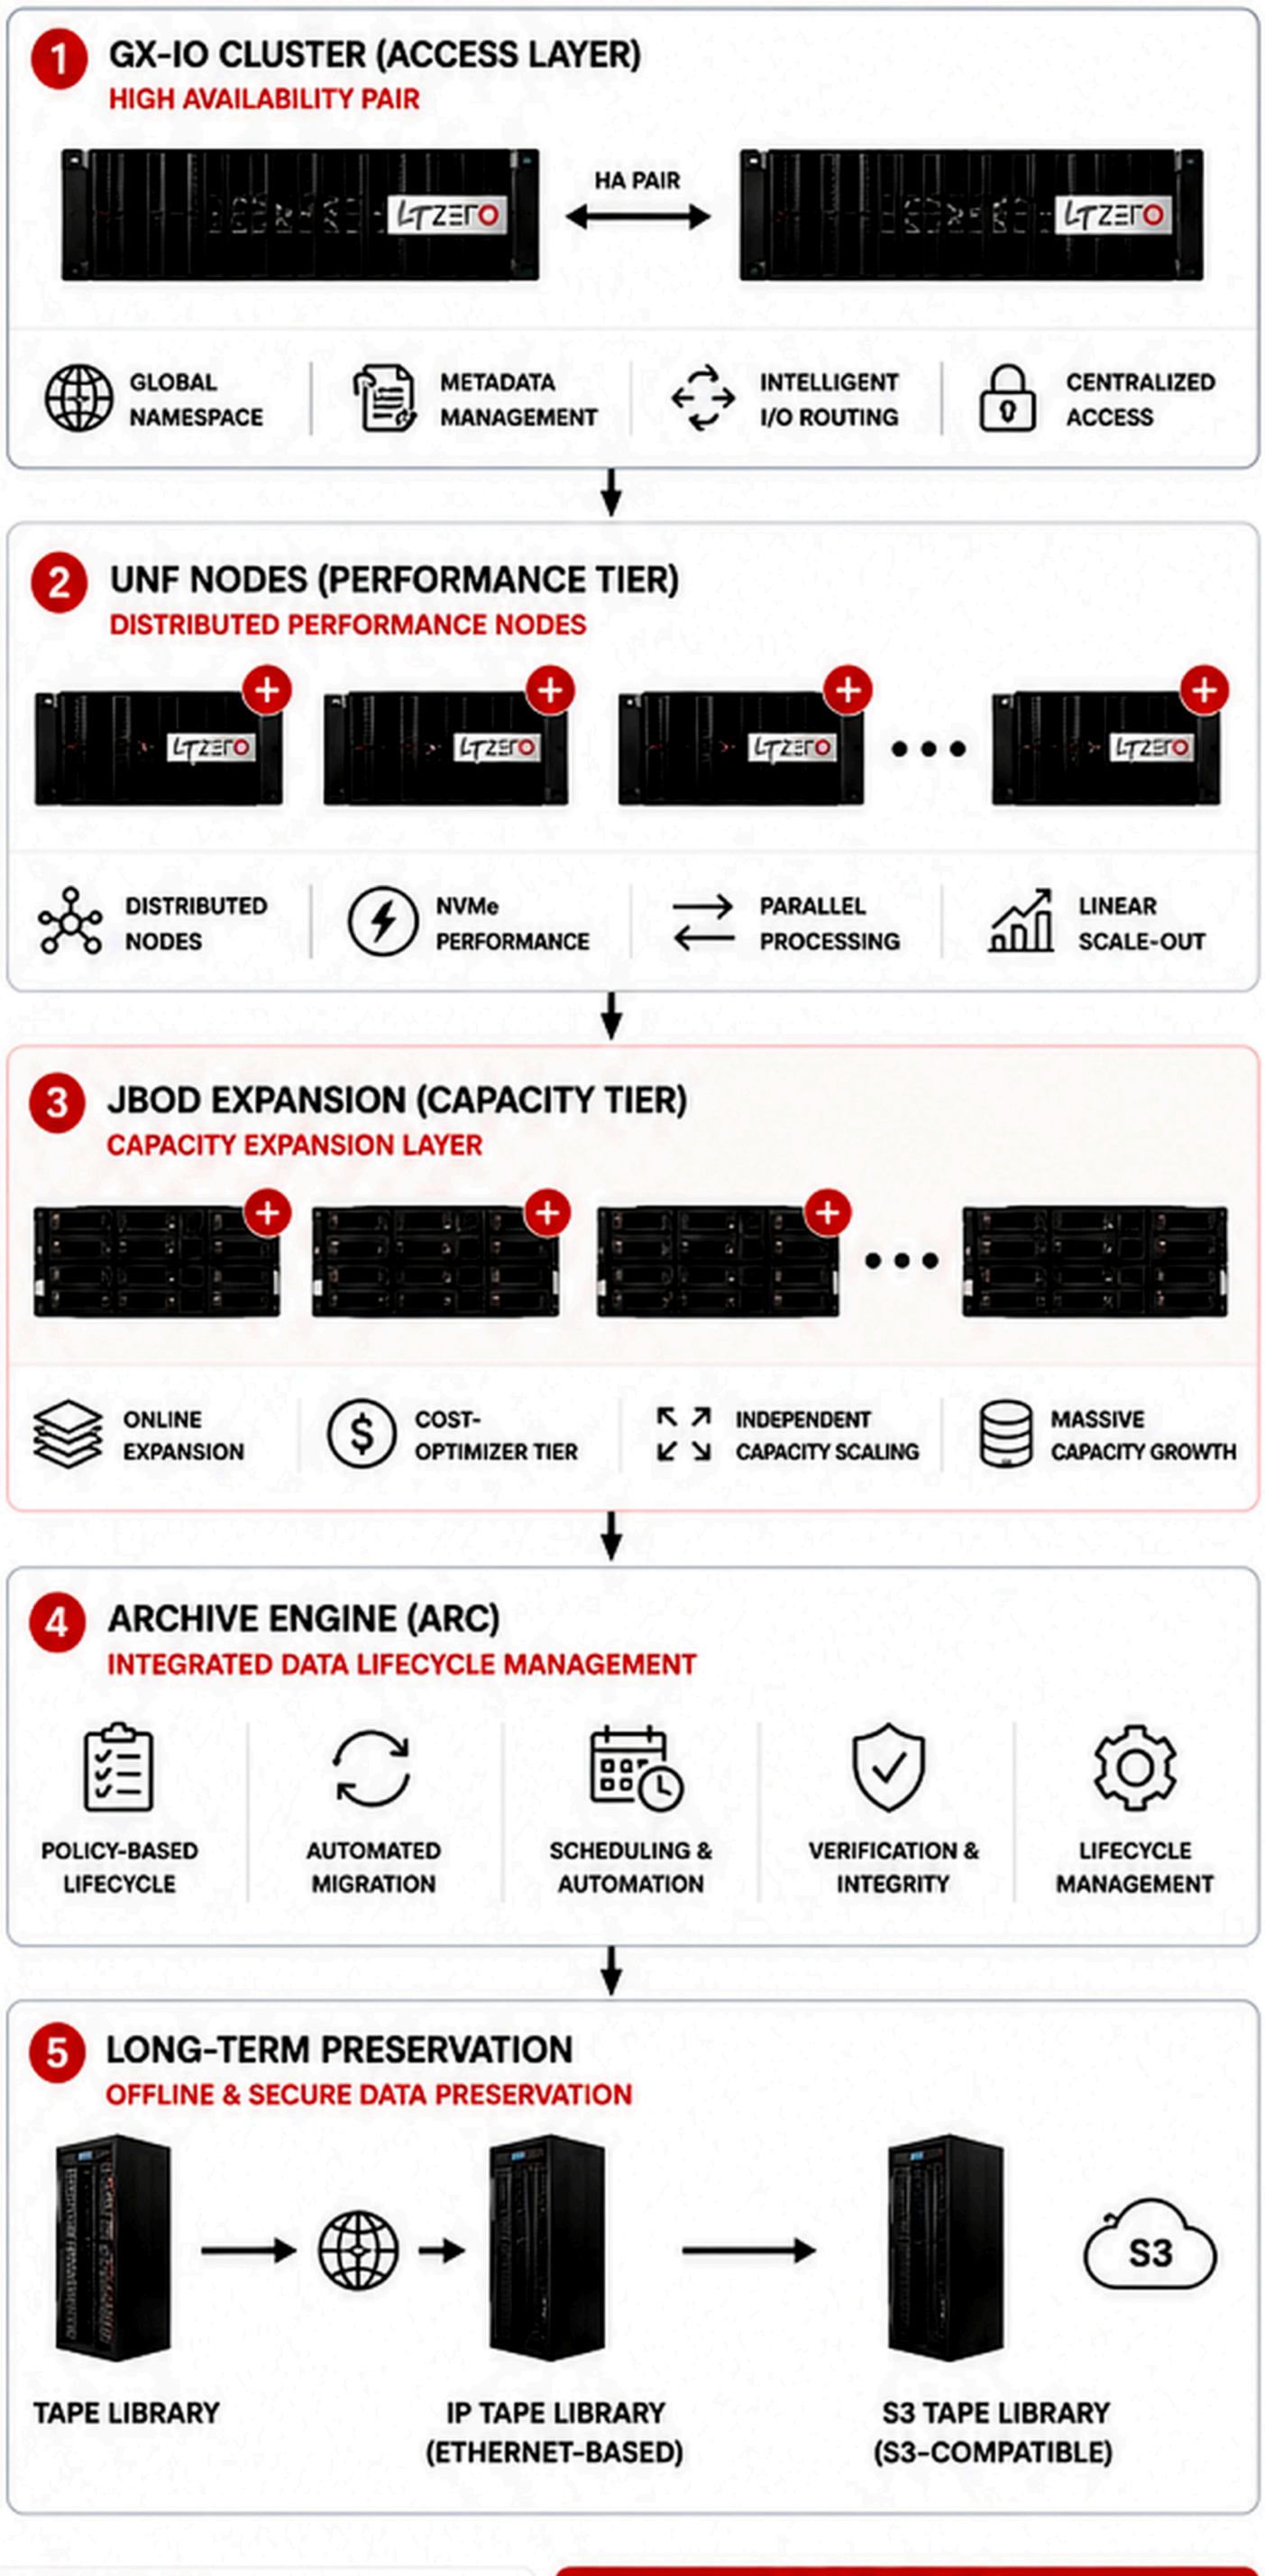
5 LONG-TERM PRESERVATION

SECURE & OFFLINE DATA PRESERVATION

- Tape Library
- IP Tape Library (Ethernet-based)
- S3 Tape Library (S3-compatible)
- Air-gap and offline protection options
- Built for long-term retention and compliance

ARCHITECTURE OVERVIEW

Kumulus DFS is built on a distributed architecture that separates access, performance, capacity, and preservation layers to deliver maximum scalability, resilience, and cost efficiency.



DISTRIBUTED BY DESIGN

- Minimum 3-Node Cluster
- Erasure Coding for data protection
- Linear scale-out architecture
- No single point of failure
- High availability and resilience
- Built for modern enterprise workloads

SCALE LIMITLESS.



KEY OUTCOMES

- High performance for demanding workloads
- Massive scalability
- Scale performance & capacity
- Cost efficiency
- Right data on the right tier
- Data lifecycle from active to archive & retention
- Data protection
- Resilient, reliable & secure

WHY KUMULUS DFS?

ENGINEERED FOR PERFORMANCE. BUILT FOR THE FUTURE.

Kumulus DFS is a modern distributed file system that delivers enterprise-grade performance, limitless scale, and integrated data lifecycle management in a single platform—without the complexity and cost of traditional storage.

KEY TECHNICAL ADVANTAGES

ENTERPRISE-GRADE PERFORMANCE

- NVMe-based performance tier for low latency and high throughput
- Parallel processing and distributed data access
- Optimized for I/O intensive applications

LINEAR SCALE-OUT ARCHITECTURE

- Add nodes and capacity independently
- No downtime, no data migration
- Scale performance and capacity seamlessly

INTEGRATED DATA LIFECYCLE MANAGEMENT

- Active, inactive to archive – all in one platform
- Policy-based automation across data tiers
- Built-in Archive Engine (ARC) for efficient data movement

DATA PROTECTION & RESILIENCE

- Erasure coding for maximum data durability
- High availability with HA pairs and distributed nodes
- No single point of failure

COST OPTIMIZATION

- Intelligent tiering from performance to capacity to archive
- Use cost-effective storage for long-term retention
- Reduce TCO without compromising performance

PERPETUAL LICENSING MODEL

- One-time purchase, perpetual use
- No subscription, no recurring fees
- Predictable cost, maximum value

NO CAPACITY-BASED LIMITATION FOR ARCHIVE LICENSE

- Archive Engine license included
- No limits on archive capacity or data movement
- True freedom to grow and retain

AT A GLANCE

HIGH PERFORMANCE
Built for demanding workloads with NVMe acceleration and parallel processing.

UNIFIED PLATFORM
One platform for all data—performance, capacity, and archive.

RESILIENT BY DESIGN
Distributed architecture with erasure coding ensures high availability and zero single point of failure.

COST-EFFICIENT
Intelligent tiering and integrated archive deliver the best TCO.

LIMITLESS SCALABILITY
Scale-out without limits. Add nodes, capacity, and performance on demand.

PERPETUAL LICENSING – NO HIDDEN COSTS

KUMULUS DFS	VS	TRADITIONAL STORAGE
• Perpetual license		• Subscription / recurring cost
• No capacity-based pricing		• Capacity-based license
• Archive Engine included		• Additional archive cost
• No hidden fees		• Hidden fees & charges
• Predictable TCO		• Unpredictable TCO

KUMULUS DFS vs TRADITIONAL STORAGE

FEATURE	KUMULUS DFS	TRADITIONAL STORAGE
Architecture	Distributed Scale-Out	Traditional Scale-Up
Scalability	Linear Scale-Out	Scale-Up (Limited)
Performance	High Performance (NVMe + Parallel)	Limited by Controller
Availability	High Availability (HA Pair + Erasure Coding)	Controller / Node Dependency
Data Protection	Erasure Coding (No Single Point of Failure)	RAID (Controller Dependency)
Data Lifecycle	Active – Inactive – Archive (Integrated)	Separate Solutions Required
Archive Engine	Included (No Capacity Limit)	Additional License / Extra Cost
Licensing Model	Perpetual License	Subscription / Recurring
Cost Efficiency	Optimized TCO	Higher TCO
Expansion	Online, Non-Disruptive	Complex, Downtime

IDEAL FOR

ENTERPRISE IT
Unify and protect enterprise data across all departments.

AI / MACHINE LEARNING
High-performance storage for AI training, inference, and data lakes.

HPC & RESEARCH
High throughput and low latency for compute-intensive workloads.

MEDIA & CONTENT
Store, manage, and archive massive media and content libraries.

OIL & GAS
Manage and retain large volumes of exploration and engineering data.

GOVERNMENT & DEFENSE
Secure, compliant, and air-gap capable data retention solutions.

DELIVERING MAXIMUM VALUE

MAXIMUM PERFORMANCE | UNLIMITED SCALABILITY | BUILT-IN RESILIENCE | LOWER TCO | COMPLETE DATA LIFECYCLE | FUTURE-PROOF PLATFORM

BOTTOM LINE

Kumulus DFS is the modern data platform that delivers high performance, limitless scale, and integrated data lifecycle management—without the complexity, hidden costs, or limitations of traditional storage.

ONE PLATFORM. COMPLETE DATA LIFECYCLE. UNLIMITED POSSIBILITIES.

FLEXIBLE SOLUTIONS. UNLIMITED POSSIBILITIES.

Kumulus DFS adapts to your business needs with flexible deployment options and seamless integration for backup, archive, analytics, and long-term data preservation.

IDEAL USE CASES

ENTERPRISE BACKUP
High-performance backup target with long-term retention.

DATA ARCHIVING
Cost-effective archive for inactive and historical data.

AI / MACHINE LEARNING
High throughput storage for training, inference and datasets.

HPC & RESEARCH
Parallel I/O and massive capacity for compute-intensive workloads.

MEDIA & CONTENT
Store, manage and protect large media and content libraries.

GOVERNMENT & COMPLIANCE
Secure, immutable storage for regulatory data and long-term retention.

DEPLOYMENT HIGHLIGHTS

MODULAR & FLEXIBLE
Start small and scale out linearly by adding nodes and JBOD enclosures without downtime.

HIGH AVAILABILITY
Built-in HA for GX-IO and UNF Nodes with erasure coding across all layers.

NON-DISRUPTIVE EXPANSION
Add capacity and performance on demand without interrupting services.

GLOBAL ACCESS & UNIFIED NAMESPACE
Access data from anywhere with a single global namespace across all sites and tiers.

SECURE & COMPLIANT
Integrated data protection, air-gap options, and long-term preservation for compliance and governance.

PERPETUAL LICENSING – NO HIDDEN COSTS

KUMULUS DFS	VS	TRADITIONAL STORAGE
• Perpetual license		• Subscription / recurring cost
• No capacity-based pricing		• Capacity-based license
• Archive Engine included		• Additional archive cost
• No hidden fees		• Hidden fees & charges
• Predictable TCO		• Unpredictable TCO

DATA LIFECYCLE WORKFLOW

1. **GX-IO ACCESS & METADATA LAYER** — GX-IO: Metadata Management, Intelligent & Centralized Access provides a single global namespace across the distributed DFS environment.
2. **PERFORMANCE NVMe HOT TIER** — High-performance NVMe tier for active applications and demanding workloads.
3. **SCALE-OUT CAPACITY JBOD CAPACITY TIER** — Cost-effective capacity expansion using JBOD with Erasure Coding for warm/cold data.
4. **INTEGRATED ARCHIVE ARCHIVE ENGINE (ARC)** — Archive Engine (ARC) manages lifecycle, migration, verification and retention policies.
5. **LONG-TERM PRESERVATION TAPE / IP TAPE / S3 TAPE LIBRARY** — Secure long-term retention with offline, air-gapped, or S3-compatible preservation options.

EXAMPLE CONFIGURATIONS

USE CASE	CAPACITY RANGE (USABLE)	SUGGESTED CONFIGURATION (EXAMPLE)		
		GX-IO HA PAIR (NODES)	UNF NODES (MIN.)	JBOD ENCLOSURES (MIN.)
SMALL ENTERPRISE	Up to 500 TB	2 Nodes	3 Nodes	3 Enclosures (Min.)
MEDIUM ENTERPRISE	500 TB – 5 PB	2 Nodes	5 – 9 Nodes	5 – 9 Enclosures (Min.)
LARGE ENTERPRISE	5 PB – 50+ PB	2 Nodes	9+ Nodes	9+ Enclosures (Min.)
SERVICE PROVIDER / CLOUD ARCHIVE	50 PB+	2 Nodes	9+ Nodes (per site)	9+ Enclosures (per site, scalable)

Note: The number of JBOD enclosures must be equal to or greater than the number of UNF Nodes.

LONG-TERM PRESERVATION OPTIONS

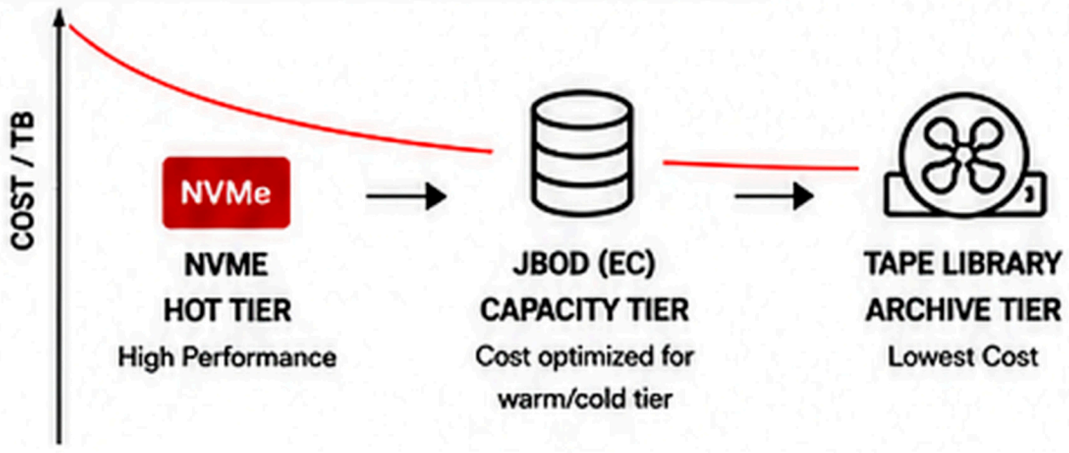


TAPE LIBRARY	IP TAPE LIBRARY (ETHERNET-BASED)	S3 TAPE LIBRARY (S3-COMPATIBLE)
• Fibre Channel (FC)	• IP-based connectivity	• S3 compatible interface
• High capacity	• Easier deployment	• Cloud integration ready
• Air-gap capable	• (No SAN infrastructure)	• Cost-effective archive over distance
• Enterprise grade	• Air-gap capable	• Air-gap capable
• LTO 7 – LTO 10	• LTO 7 – LTO 10	• LTO 7 – LTO 10

KEY SPECIFICATIONS

Architecture	Distributed Scale-Out File System
Minimum Cluster	3 UNF Nodes (Erasure Coded Cluster)
Access Layer	GX-IO Cluster (HA Pair)
Performance Tier	NVMe-based (Hot / Warm tier)
Capacity Expansion	JBOD Expansion (Erasure Coded)
Archive Engine	Integrated Archive Engine (ARC)
Protocols	NFS, SMB, S3 via Gateway, HDFS Compatible
Erasure Coding	End-to-End EC across all tiers
High Availability	HA Pair for GX-IO, EC Protection for Data
Expansion	Linear Scale-Out (Add Nodes & JBOD)
License Model	Perpetual License
Archive License	No Capacity Limitation
Supported Tape	LTO 7 – LTO 10
Deployment	On-Premises / Hybrid / Multi-Site
Management	Centralized Monitoring & Lifecycle Management

COST EFFICIENCY WITH TIERED ARCHITECTURE



SOLUTION SUMMARY

PERFORMANCE WITHOUT COMPROMISE
High-performance NVMe tier and parallel processing deliver consistent throughput for demanding workloads.

SCALE WITHOUT LIMITS
Add nodes and capacity independently to grow seamlessly as your data grows.

PROTECT & PRESERVE
Integrated lifecycle management and multiple archive options ensure data is secure, compliant and preserved for the long term.

LTZERO KUMULUS DFS

ONE PLATFORM. FROM PERFORMANCE TO PRESERVATION.

Kumulus DFS unifies high-performance storage, massive scale-out capacity, integrated archive, and long-term preservation into a single platform. Designed to grow with your business and built to protect what matters most.



THE KUMULUS DFS PLATFORM

GX-IO

ACCESS & METADATA LAYER

Metadata Management, Intelligent & Centralized Access provides a single global namespace across the distributed DFS environment.

PERFORMANCE

NVMe HOT TIER

High-performance NVMe tier for active applications and demanding workloads.

SCALE-OUT CAPACITY

JBOD CAPACITY TIER

Cost-effective capacity expansion using JBOD with Erasure Coding for warm/cold data.

INTEGRATED ARCHIVE

ARCHIVE ENGINE (ARC)

Manages lifecycle, migration, verification and retention policies.

LONG-TERM PRESERVATION

TAPE / IP TAPE / S3 TAPE LIBRARY

Secure long-term retention with offline, air-gapped, or S3-compatible preservation options.

WHY CHOOSE KUMULUS DFS?

DISTRIBUTED SCALE-OUT ARCHITECTURE

Add nodes and capacity linearly without forklift upgrade. Grow as your data grows.

BUILT-IN DATA PROTECTION

Erasure Coding across all tiers ensures data resiliency and high availability without performance compromise.

INTEGRATED ARCHIVE ENGINE

Lifecycle management, data migration, verification and retention — all included, no additional software required.

PERPETUAL LICENSING

Pay once, own forever. No recurring software license cost.

NO CAPACITY LIMITATION FOR ARCHIVE LICENSE

Archive as much as you need without license based on TB or slot. True freedom to scale.

OPTIMIZED TCO

Cost-optimized architecture from performance to preservation delivers best TCO for the entire lifecycle.

TRUSTED TO PROTECT WHAT MATTERS MOST

Kumulus DFS is engineered for organizations that demand performance, scalability and long-term data protection without compromise.

HIGH PERFORMANCE

Built for demanding workloads with consistent throughput and low latency.

HIGH AVAILABILITY

HA architecture and Erasure Coding ensure maximum uptime and data resiliency.

LIMITLESS SCALABILITY

Scale-out nodes and JBOD capacity on demand without disruption.

SECURE & OFFLINE

Air-gap capable archive with Tape Library options for ultimate data preservation.

YOUR DATA. OUR PLATFORM.
POWERING YOUR FUTURE.

EXPERT TEAM

Experienced professionals dedicated to your success.

END-TO-END SUPPORT

From deployment to operations, we've got you covered.

PROVEN SOLUTIONS

Battle-tested architecture trusted by enterprise customers.

FUTURE READY

Continuously innovating to meet the demands of tomorrow.

YOUR DATA. OUR PLATFORM.

One platform for entire data lifecycle — from performance to preservation.