Zero-knowledge Cloud Storage Done Right

Privacy Without Performance Compromise

Many claim zero-knowledge = slow. Wrong. Discover how Everabyte® achieves performance 4x faster than encrypted competitors while maintaining complete client-side privacy — with zero trust required in any server or provider.

4x Faster
Than encrypted competitors

100Gbps
Max bandwidth scaling

25ms
US-East → EU latency

SOC 2 II
GDPR, HIPAA, SOC 2 Type II Compliant

KEY HIGHLIGHTS

- ✓ 4x faster than encrypted competitors
- ✓ US-East → EU: 25ms latency
- ✓ Bandwidth scaling: 1Gbps → 100Gbps
- ✓ GDPR Art 32, HIPAA BAA, SOC 2 Type II

CONTENTS

Engineering Playbook · 2026 · 48 pages · Everabyte® Privacy Engineering Unit

Table of Contents

Executive Summary

The prevailing assumption that zero-knowledge encryption inherently degrades performance is one of the most damaging myths in enterprise cloud storage. It has led countless organizations to choose between privacy and speed — accepting weaker encryption in exchange for acceptable throughput, or accepting poor performance in exchange for true privacy guarantees.

Everabyte® was purpose-built to destroy that false dilemma. Through a combination of client-side key management, hardware-accelerated AES-256-GCM encryption, a globally distributed 7-datacenter topology, and proprietary zero-copy transfer pipelines, we achieve what our competitors claim is impossible: zero-knowledge storage that is 4x faster than the industry average for encrypted cloud competitors.

This whitepaper presents the complete technical and commercial case. We examine the cryptographic foundations of true zero-knowledge architecture, dismantle server-side encryption myths, detail Everabyte®'s engineering decisions, present verified performance benchmarks, and show how any enterprise — from a 50-person startup to a Fortune 500 — can deploy zero-knowledge storage without sacrificing the SLAs their business demands.

4x Faster than encrypted competitors	**0** Bytes of plaintext ever leaving the client
25ms Cross-continental latency (US-East → EU)	**7** Global Tier III+ datacenters

The Privacy-Performance Paradox — Solved

For years, the storage industry offered a menu of bad choices. You could choose end-to-end encryption — true client-side key management where the provider never sees your plaintext — but you would pay for it in throughput. Encryption CPU overhead, larger ciphertext sizes, and the absence of server-side deduplication all conspired to make encrypted storage measurably slower than its plaintext counterpart.

The alternative was server-side encryption: the provider encrypts your data at rest, but holds the keys. Marketed as "encrypted," this approach provides almost no meaningful privacy guarantee. A government subpoena, a rogue employee, a breach of the provider's key management system — any of these events gives a third party full access to your data. Server-side encryption is, at best, encryption theater.

Everabyte® chose neither option. Instead, our engineering team asked a different question: what if we could make client-side encryption so efficient — at the hardware, network, and software levels — that the performance gap simply disappears?

Zero-knowledge Defined

In Everabyte®'s architecture, "zero-knowledge" means precisely this: at no point in the data lifecycle does any Everabyte® server, employee, subprocessor, or system have access to unencrypted client data or client encryption keys. This is enforced cryptographically, not by policy.

What This Whitepaper Covers

- Chapter 1 — The precise cryptographic definition of zero-knowledge storage and why it matters
- Chapter 2 — A myth-by-myth dismantling of server-side encryption marketing claims
- Chapter 3 — Everabyte®'s full technical architecture: key derivation, encryption pipeline, and network topology
- Chapter 4 — Verified performance benchmarks: throughput, latency, and IOPS compared to five leading competitors
- Chapter 5 — 7-datacenter redundancy architecture: how we achieve 99.9999% availability with zero-knowledge constraints
- Chapter 6 — Regulatory compliance: GDPR Art. 32, HIPAA BAA, SOC 2 Type II, and ISO 27001
- Chapter 7 — Enterprise deployment guide: migration, integration, and operational playbooks
- Chapter 8 — Post-quantum roadmap: CRYSTALS-Kyber integration timeline and forward-secrecy planning

Chapter 1: What Is Zero-knowledge Storage?

Beyond Marketing: A Precise Definition

The term "zero-knowledge" has been so thoroughly weaponized by marketing departments that it has nearly lost its technical meaning. To build on solid foundations, we must begin with a precise, engineering-grade definition.

In the context of cloud storage, a system is genuinely zero-knowledge if and only if the following conditions hold simultaneously:

1. The client generates and holds all encryption keys. The server receives ciphertext only.
2. The encryption and decryption processes occur exclusively on client-controlled hardware.
3. The key derivation function (KDF) used to generate keys from user credentials is non-exportable — the server receives a derived authentication token, never the master secret.
4. The provider has no technical capability — not just no contractual permission — to decrypt stored data.
5. This guarantee holds even in the event of a full server-side breach, a legal compulsion order served on the provider, or a malicious insider threat.

If any of these conditions is violated, the system is not zero-knowledge. It may be encrypted. It may be more secure than plaintext storage. But it does not offer the privacy guarantee that the zero-knowledge label implies.

The Cryptographic Primitives

A true zero-knowledge storage system is built on a stack of well-understood cryptographic primitives. Everabyte®'s implementation uses the following:

Key Derivation: Argon2id

User passphrases are never transmitted to any server. Instead, Argon2id — the winner of the Password Hashing Competition and the current gold standard for memory-hard key derivation — is applied client-side. The output is a 256-bit master key that never leaves the client device.

Argon2id is parameterized with a minimum memory cost of 64 MiB and 3 iterations, making brute-force attacks computationally prohibitive even on dedicated hardware.

File Encryption: AES-256-GCM

Each file is encrypted with AES-256 in Galois/Counter Mode (GCM), which provides both confidentiality and authenticated integrity. The 96-bit IV is generated using a CSPRNG on the client for each file and each chunk. The 128-bit authentication tag is stored alongside the ciphertext, enabling tamper detection without server-side verification.

Key Wrapping: HKDF-SHA-512

Per-file encryption keys are derived from the master key using HKDF-SHA-512, with a context string incorporating the file's unique identifier and the client's tenant ID. This ensures that compromise of any single file key does not expose the master key or any other file key.

Why AES-256-GCM Outperforms Legacy Modes

Unlike AES-CBC, which requires PKCS#7 padding and a separate HMAC for integrity, AES-GCM provides authenticated encryption with associated data (AEAD) in a single pass. Modern x86-64 and ARM processors include AES-NI and CLMUL hardware acceleration for GCM, enabling throughput exceeding 10 Gbps on a single CPU core — eliminating encryption as a performance bottleneck.

The Knowledge Gap: What the Server Knows

In Everabyte®'s architecture, the server maintains knowledge only of: ciphertext blobs (unintelligible without the client key), file metadata structure (size, creation timestamp, chunk count — but not filename, which is also encrypted), authentication tokens derived from a separate KDF branch, and billing and tenancy records.

The server does not know, and cannot derive: file contents, filenames or directory structure, the master encryption key or any per-file key, or the user's passphrase or any recoverable credential.

Historical Context: How We Got Here

Client-side encryption is not new. PGP, developed by Phil Zimmermann in 1991, established the principle of end-to-end encrypted communication. The S/MIME standard brought similar principles to enterprise email. BitLocker and FileVault demonstrated that device-level encryption could be transparent to users.

What is new — and what Everabyte® has pioneered — is the combination of genuine zero-knowledge architecture with enterprise-grade performance, reliability, and compliance at cloud scale. Previous zero-knowledge storage attempts (including Tresorit, SpiderOak, and early versions of various open-source solutions) demonstrated the security model but struggled to match the throughput and latency profiles that enterprise SLAs demand.

The breakthrough came from three simultaneous advances: ubiquitous AES-NI hardware acceleration in commodity server hardware, the maturation of QUIC as a transport protocol enabling faster connection establishment over lossy networks, and Everabyte®'s proprietary zero-copy encryption pipeline that eliminates memory allocation overhead in the encryption hot path.

Chapter 2: Client-Side vs Server-Side Myths

The Myth Taxonomy

Server-side encryption is a legitimate and useful tool in specific contexts. But the cloud storage industry has conflated it with zero-knowledge privacy, creating a taxonomy of myths that have materially harmed enterprise security posture. We examine each myth in turn.

Myth 1: "Our Data Is Encrypted At Rest"

"Encrypted at rest" has become the minimum-viable security claim for cloud vendors. The statement is technically accurate in almost every case. It is also almost completely meaningless from a privacy perspective.

When a vendor encrypts your data at rest using their own key management infrastructure, the protection provided is narrowly scoped. It protects against physical theft of storage media — an attacker who walks out of a datacenter with a hard drive gets encrypted blobs. It offers no protection against a lawful intercept order, a breach of the vendor's key management service (KMS), a malicious insider with KMS access, or an API vulnerability that allows unauthorized decryption.

The AWS S3 Default Encryption Case Study

AWS S3's default encryption (SSE-S3) uses AES-256 encryption managed by AWS. AWS holds the keys. A valid subpoena served on Amazon compels decryption. An AWS IAM misconfiguration can expose decryption capability to unauthorized principals. "Encrypted" does not mean "private." It means "the bits on disk are not plaintext." These are very different guarantees.

Myth 2: "Zero-knowledge Is Incompatible With Collaboration"

This myth has some historical basis — early zero-knowledge systems did struggle with multi-user scenarios, because key sharing was awkward and sharing required exposing the master key.

Modern zero-knowledge architectures, including Everabyte®'s, resolve this through public-key cryptography at the sharing layer. When User A shares a file with User B:

6. User A generates a per-share key by wrapping the file's symmetric key with User B's public key.
7. The wrapped key is stored on the Everabyte® server — but the server cannot unwrap it without User B's private key.
8. User B retrieves the wrapped key and decrypts it with their private key on their client device.

9. User B can now decrypt the file. Everabyte® servers at no point had access to the file's plaintext key.

This construction — known as hybrid encryption — is the same mechanism used by Signal, WhatsApp's end-to-end encryption, and Apple's iMessage. It is battle-tested and provides full collaborative capability without compromising the zero-knowledge property.

Myth 3: "Zero-knowledge Prevents Search"

Partially true, but overstated. Full-text search of encrypted content is genuinely challenging. However, Everabyte® supports client-side search indexing — the search index is built and maintained on the client, encrypted with the file encryption key, and stored on the server. Search queries never leave the client device; they are executed against the locally cached or locally decrypted index.

For large enterprises with petabyte-scale datasets, distributed client-side indexing is provided through Everabyte®'s local sync daemon, which maintains an encrypted SQLite index on each authorized device. Index synchronization between devices is encrypted under the master key.

Myth 4: "Zero-knowledge Means Slower Performance"

This is the central myth this whitepaper exists to destroy. The performance argument against zero-knowledge storage typically invokes three sources of overhead: CPU cycles consumed by encryption, larger ciphertext (AES-GCM adds a 16-byte auth tag per chunk), and the absence of server-side deduplication.

We address each:

CPU Overhead

On any x86-64 CPU manufactured after 2010 with AES-NI support, AES-256-GCM throughput exceeds the network bandwidth of any commercially available internet connection. The benchmark: an AWS c6i.xlarge (4 vCPUs, 2021 era) achieves 25+ GB/s AES-GCM throughput in a single thread. Your network uplink is the bottleneck, not the cipher.

Ciphertext Overhead

AES-GCM adds 12 bytes (IV) + 16 bytes (auth tag) = 28 bytes per chunk. With a 4 MB chunk size, this overhead is 0.0007% — completely negligible.

Deduplication

Server-side deduplication is incompatible with client-side encryption — two identical plaintext files produce different ciphertext under different keys, so the server cannot detect them as duplicates. Client-side deduplication, performed before encryption, resolves this. Everabyte®'s client performs SHA-256 content addressing before encryption, achieving deduplication rates comparable to server-side approaches.

Feature	Everabyte® ZK	Provider-Managed
Key Control	Client-held always	Provider-held
Breach Exposure	Zero (ciphertext only)	Full plaintext risk
Legal Compulsion	Cannot comply (no key)	Must comply
Throughput	4x faster (benchmarked)	Baseline
Collaboration	Full (hybrid encryption)	Full
Search	Client-side index	Server-side (full)
Dedup	Client-side SHA-256	Server-side (full)

Chapter 3: Everabyte® Zero-knowledge Architecture

System Overview

Everabyte®'s zero-knowledge architecture is organized into four logical planes: the Client Plane (where encryption, decryption, key management, and indexing occur), the Transfer Plane (zero-copy encrypted data movement over QUIC), the Storage Plane (distributed object storage across 7 Tier III+ datacenters), and the Metadata Plane (encrypted metadata management, ACL enforcement, audit logging).

The critical invariant enforced across all planes: no plaintext data and no client key material ever exists outside the Client Plane.

Client Plane Architecture

Key Management Subsystem

The Everabyte® client maintains a local key store, protected by the device's secure enclave (Apple Secure Enclave, Intel SGX, or Android StrongBox, depending on platform). The master key is derived from the user's passphrase using Argon2id and is wrapped under the secure enclave's hardware-bound key before any at-rest storage on the device.

For enterprise deployments, a centralized key management server (KMS) can be deployed on-premise or in the customer's own cloud account. The Everabyte® service never has access to this enterprise KMS — key retrieval occurs directly between the client and the customer's KMS over mTLS.

Encryption Pipeline

Files are chunked into 4 MB segments before encryption. Each chunk is independently encrypted with AES-256-GCM using a per-chunk IV derived from HKDF(master_key, file_id || chunk_index). This construction enables:

- Parallel chunk encryption across CPU cores (typically 8-32x parallelism on enterprise hardware)
- Random-access decryption of arbitrary file regions without full-file decryption
- Partial upload resumability — interrupted uploads can be resumed from the last committed chunk boundary
- Efficient range requests for streaming media files

Zero-Copy Pipeline

The most significant performance innovation in Everabyte®'s client is the zero-copy encryption pipeline. Traditional encryption implementations involve multiple memory copies: read plaintext from disk into buffer A, copy to buffer B for encryption, copy ciphertext to socket buffer for network transmission. Each copy consumes CPU cycles and memory bandwidth.

Everabyte®'s pipeline uses Linux kernel splice() and sendfile() system calls combined with io_uring for asynchronous I/O, eliminating intermediate copies. Plaintext is read from disk directly into AES-NI-encrypted output, which is written directly to the QUIC socket. On benchmarked hardware, this reduces CPU utilization per GB transferred by 60% compared to naive implementations.

io_uring: The Linux I/O Revolution

io_uring, introduced in Linux kernel 5.1 (2019), provides an asynchronous I/O interface that eliminates syscall overhead for high-throughput I/O workloads. Everabyte®'s client leverages io_uring for all disk and network I/O, enabling throughput that approaches the theoretical limits of the underlying hardware. On NVMe-equipped clients, we observe sustained read throughput of 6+ GB/s feeding the encryption pipeline.

Transfer Plane: QUIC Protocol Adoption

Everabyte® adopted QUIC (RFC 9000) as its transport protocol in 2024, replacing the legacy TLS-over-TCP stack. The performance implications are substantial:

Connection Establishment

QUIC combines transport and cryptographic handshakes, achieving 0-RTT connection establishment for returning clients (where the server's transport parameters are cached from a previous session). For new connections, QUIC requires 1-RTT — identical to TLS 1.3 over TCP, but without TCP's additional SYN/SYN-ACK round trip.

The practical impact: for workloads involving many small file uploads (common in photo backup and document collaboration scenarios), Everabyte® clients achieve connection setup in under 5ms to our nearest datacenter, versus 45-80ms for TLS-over-TCP connections to competing services.

Multiplexing Without Head-of-Line Blocking

TCP suffers from head-of-line blocking: a lost packet stalls the entire connection until retransmission is complete. QUIC multiplexes independent streams over a single UDP connection; a lost packet stalls only the stream it belongs to, while other streams continue unimpeded.

For parallel chunk uploads (Everabyte® uploads 8 chunks simultaneously by default), QUIC's multiplexing ensures that a packet loss event affecting one chunk stream does not degrade the other seven. In high-latency, moderate-loss network conditions (common in mobile and international scenarios), this produces throughput improvements of 2-3x over TCP-based competitors.

Storage Plane: Distributed Object Storage

Everabyte®'s storage layer is a custom distributed object store built on top of NVMe-backed hardware in 7 Tier III+ datacenters. Unlike commodity object stores (S3-compatible APIs built on spinning disk or hybrid storage), Everabyte® uses all-NVMe storage arrays with a custom replication protocol optimized for encrypted object storage.

Erasure Coding

Objects are stored using a Reed-Solomon erasure coding scheme with a 6+3 configuration: each object is split into 6 data shards and 3 parity shards. The 9 shards are distributed across 9 distinct failure domains (combination of datacenter, rack, and drive array). The system tolerates the simultaneous loss of any 3 shards — equivalent to losing 3 complete failure domains — without data loss.

Replication Topology

Primary shards are placed in the datacenter nearest to the uploading client. Secondary shards are placed in datacenters in different geographic regions, ensuring that a regional infrastructure event does not result in data unavailability. Shard placement follows a deterministic consistent hashing scheme, enabling any node in the cluster to serve any shard without centralized coordination.

Datacenter	Location	Tier
DC-1	Ashburn, Virginia (US-East)	Tier III+
DC-2	Amsterdam, Netherlands (EU-West)	Tier III+
DC-3	Singapore (APAC)	Tier III+
DC-4	São Paulo, Brazil (LATAM)	Tier III+
DC-5	Sydney, Australia (APAC-South)	Tier III+
DC-6	Johannesburg, South Africa (EMEA-South)	Tier III+
DC-7	Mumbai, India (APAC-West)	Tier III+

Chapter 4: Performance Reality (4x Faster)

Benchmark Methodology

All benchmarks presented in this chapter were conducted using a standardized test harness deployed on identical hardware in each competitor's nearest datacenter to our primary test client location (AWS us-east-1 region). The test client was a c6i.8xlarge instance (32 vCPUs, 64 GB RAM, 25 Gbps network) running Ubuntu 22.04 LTS.

Tests were conducted across three workload profiles: sequential large file upload (10 GB synthetic dataset, 1 file), random small file upload (10 GB total, 10,000 files averaging 1 MB), and mixed workload (70% sequential, 30% random, representative of enterprise backup workloads).

All tests were run with encryption enabled on all platforms. For platforms offering zero-knowledge (client-side) encryption, it was enabled. For platforms offering only server-side encryption, the default encryption setting was used (which represents their best-case scenario — our worst-case scenario from an apples-to-apples standpoint, since we are measuring encrypted throughput).

Sequential Upload Throughput

Sequential upload throughput is the most commonly cited benchmark in cloud storage marketing. It represents the peak throughput achievable when uploading a single large file with sufficient network bandwidth.

680 MB/s Everabyte® — sequential upload (zero-knowledge)	**155 MB/s** Competitor B — encrypted (client-side, legacy stack)
170 MB/s Competitor A — encrypted (provider-managed keys)	**190 MB/s** Competitor C — server-side encryption

The 4x advantage is consistent with our architectural analysis. Competitor B's client-side encryption stack — a 2019-era implementation using OpenSSL's AES-CBC mode with a single-threaded encryption pipeline — leaves substantial hardware capacity on the table. Our parallel chunk encryption across 32 cores, combined with the zero-copy pipeline, saturates the available 25 Gbps network link.

Cross-Continental Latency: US-East to EU

For organizations operating across geographies, cross-continental latency is a critical SLA metric. We measured latency for a 4 KB metadata operation (file stat, equivalent to "does this file exist and what is its size") from our US-East test client to each provider's EU endpoint.

25ms Everabyte® US-East → EU (p50)	**85ms** Competitor A US-East → EU (p50)
28ms Everabyte® US-East → EU (p99)	**210ms** Competitor B US-East → EU (p50 — TLS over TCP)

Our 25ms p50 latency is achieved through a combination of QUIC 0-RTT connection resumption (eliminating handshake overhead for returning clients), anycast routing to the nearest datacenter PoP, and a metadata caching layer that serves file stat operations from an in-memory cache on the first hop rather than querying the underlying object store.

Bandwidth Scaling: 1 Gbps to 100 Gbps

Enterprise workloads are rarely static. A backup job that runs at 1 Gbps overnight may need to scale to 100 Gbps during disaster recovery scenarios when an organization needs to restore petabytes of data in the shortest possible window. Everabyte®'s architecture supports dynamic bandwidth scaling without pre-provisioning.

Scaling from 1 Gbps to 100 Gbps is achieved through: automatic client-side parallelism scaling (the client automatically increases chunk concurrency as bandwidth allows), BGP anycast routing distributing load across multiple datacenter ingress points, and Everabyte®'s custom congestion control algorithm — a modification of BBR v3 tuned for encrypted large-object transfers.

Disaster Recovery Scenario: 100 TB Restore

In a tested disaster recovery scenario (100 TB restoration from Everabyte® backups), a client with 100 Gbps datacenter connectivity restored the full dataset in 2.2 hours. The same dataset restoration from a traditional encrypted backup service (TLS over TCP, server-side encryption) required 11.4 hours — 5x slower. In a real ransomware recovery scenario, the difference between 2.2 hours and 11.4 hours of downtime is existential for many businesses.

Small File Performance: The Hidden Bottleneck

Large file throughput benchmarks dominate vendor marketing materials because they favor all vendors equally. Small file performance — the handling of millions of files averaging under 1 MB — is where architectural differences become decisive.

Small file workloads stress the connection establishment path, metadata operations, and chunking overhead. Our QUIC-based architecture, with 0-RTT connection resumption and a

metadata service optimized for high-IOPS workloads, maintains throughput efficiency at small file scales where TCP-based competitors suffer severe degradation.

File Size	Everabyte® (ops/sec)	Competitor Avg (ops/sec)
4 KB	45,000	8,200
64 KB	38,000	11,500
1 MB	12,000	4,800
10 MB	3,200	1,600
1 GB	85	22

Chapter 5: 7 Datacenter Global Redundancy

The Availability Equation

True enterprise storage requires three availability guarantees that are in tension with each other: geographic diversity (data survives regional disasters), low latency (data is close to the clients that access it), and consistency (all clients see the same data, always). Traditional distributed storage systems force tradeoffs between these three properties — the CAP theorem being the most famous formalization of this tension.

Everabyte®'s zero-knowledge architecture introduces a fourth constraint: all replication and consistency mechanisms must operate on ciphertext, without any server-side decryption. This rules out certain consistency approaches that require content inspection (such as server-side deduplication-based consistency checks) but does not fundamentally alter the availability calculus.

Datacenter Topology Design

The 7-datacenter topology was designed to satisfy three geographic requirements: no two adjacent datacenters should share a common failure domain (defined as a natural disaster zone, a national power grid, or a regional internet exchange point), every major population center should have a datacenter within 100ms round-trip latency, and the erasure coding scheme (6+3) should be distributable such that any single datacenter failure results in at most 1.5 lost shards (well within the 3-shard tolerance).

Failure Domain Analysis

The 7 datacenter locations were selected after analysis of historical infrastructure failure events between 2015 and 2025. The analysis covered: major cloud provider regional outages (AWS, Azure, GCP), national-level BGP routing incidents, submarine cable cuts affecting internet connectivity, and natural disaster impacts on datacenter availability.

No two Everabyte® datacenters share: the same Tier-1 internet backbone provider (AS), the same national regulatory jurisdiction, the same tectonic fault zone, or the same power grid interconnect.

99.9999% Designed availability (six nines)	**6+3** Reed-Solomon erasure coding ratio
31 sec Maximum annual downtime at six nines	**3** Simultaneous datacenter failures tolerated

Replication Without Decryption

Standard distributed storage replication protocols often rely on content inspection for consistency verification. For example, read-repair protocols in Cassandra-style systems compare the actual content of replicas to detect divergence. Since Everabyte® servers never have plaintext access, such content-inspection-based protocols must be replaced with ciphertext-safe alternatives.

Everabyte® uses HMAC-SHA-256 content fingerprints of ciphertext chunks for consistency verification. The HMAC key is derived from the file's metadata (object ID and chunk index), ensuring that fingerprint comparison reveals whether two ciphertext chunks are identical without revealing anything about the plaintext. This allows our replication system to detect and repair divergent replicas without any plaintext exposure.

Geo-Distributed Consistency: CRDT-Based Metadata

File metadata (name, directory structure, modification timestamps — all stored encrypted) uses a CRDT (Conflict-free Replicated Data Type) approach for eventual consistency across datacenters. CRDTs allow concurrent metadata updates from different geographic locations to be merged without conflict resolution requiring a central authority.

The practical benefit: if a user uploads a file from a US-East client while simultaneously a collaborator modifies a file attribute from an EU client, both operations can proceed independently and will converge to a consistent state without coordination latency. For most enterprise use cases, this is equivalent to strong consistency, while delivering the availability benefits of eventual consistency.

Recovery Time and Point Objectives

Enterprise SLAs are ultimately expressed as RTOs (Recovery Time Objectives) and RPOs (Recovery Point Objectives). Everabyte®'s architecture delivers:

- RTO: < 15 minutes for full datacenter failure (failover to secondary datacenter, automated DNS rerouting)
- RPO: Near-zero (replication lag < 5 seconds under normal operating conditions)
- RTO for single object restoration: < 2 seconds (served from nearest available shard)
- RTO for full namespace restoration: scales with data volume — 100 TB in 2.2 hours at 100 Gbps

Chapter 6: Compliance & Regulatory Framework

The Regulatory Landscape for Cloud Storage in 2026

The regulatory environment for cloud storage has become substantially more complex since GDPR came into force in 2018. Organizations operating internationally now navigate a matrix of requirements including GDPR Article 32 (EU), HIPAA Security Rule and BAA requirements (US healthcare), SOC 2 Type II (US enterprise audit standard), ISO 27001 (international information security management), CCPA/CPRA (California), PDPA (Thailand), LGPD (Brazil), PIPL (China), and sector-specific frameworks including PCI DSS (payments) and FedRAMP (US federal government).

A genuine zero-knowledge architecture provides significant compliance advantages across this matrix. When the processor (Everabyte®) genuinely cannot access personal data, the risk and obligation landscape changes fundamentally.

GDPR Article 32: Technical and Organizational Measures

GDPR Article 32 requires controllers and processors to implement "appropriate technical and organizational measures" to ensure data security "appropriate to the risk." Recital 83 explicitly mentions pseudonymization and encryption as appropriate measures.

Everabyte®'s zero-knowledge architecture addresses Article 32 at multiple levels:

- Encryption in transit: QUIC with TLS 1.3 for all data movement
- Encryption at rest: AES-256-GCM, client-side keys
- Ongoing confidentiality: cryptographically enforced — not policy-dependent
- Integrity and availability: 6+3 erasure coding across 7 Tier III+ datacenters
- Resilience: automatic failover, < 15 minute RTO
- Restoration: deterministic, tested quarterly

For GDPR purposes, Everabyte® is a Data Processor when handling EU personal data on behalf of customers. Our DPA (Data Processing Agreement) is available for review and signing via the enterprise portal.

HIPAA: BAA and the Encryption Safe Harbor

HIPAA's Breach Notification Rule includes an important provision: breaches of Protected Health Information (PHI) that has been encrypted in accordance with NIST guidelines are exempt from breach notification requirements — because encrypted data without the decryption key is unusable to an attacker.

Everabyte® qualifies for this safe harbor by design. Our AES-256-GCM encryption using client-held keys meets the NIST SP 800-111 encryption specifications referenced in the HIPAA guidance. A breach of Everabyte®'s infrastructure — exposure of ciphertext — does not trigger

HIPAA breach notification obligations for customers storing PHI on Everabyte®, because the ciphertext is cryptographically unusable without the client's keys.

We offer a signed Business Associate Agreement (BAA) to all healthcare customers under our Enterprise and Business plan tiers.

The HIPAA Encryption Safe Harbor in Practice

Under 45 CFR § 164.402, a breach of unsecured PHI requires notification to affected individuals, HHS, and potentially media. "Unsecured" means not rendered unusable, unreadable, or indecipherable. AES-256 encrypted PHI where the keys are not compromised is considered secured. Everabyte® customers storing PHI are protected by this safe harbor for any breach of Everabyte®'s infrastructure, since the keys are held client-side and are not part of Everabyte®'s infrastructure.

SOC 2 Type II

Everabyte® maintains a SOC 2 Type II report covering the Trust Services Criteria for Security, Availability, and Confidentiality. The report covers a 12-month observation period and is audited by an independent PCAOB-registered CPA firm.

Key findings from our most recent SOC 2 Type II report (covering January–December 2025): zero exceptions noted across all tested controls, 99.9998% actual availability achieved during the observation period, all penetration testing engagements resolved with no critical or high findings carried forward, and full encryption-at-rest and in-transit controls verified through technical testing (not solely through management representation).

ISO 27001 Certification

Everabyte®'s Information Security Management System (ISMS) is certified to ISO/IEC 27001:2022. The certification covers our product development, infrastructure operations, customer support, and key management processes. Annual surveillance audits and triennial recertification audits are conducted by an accredited certification body.

Chapter 7: Enterprise Deployment Playbook

Migration Strategy: From Existing Storage to Everabyte®

Migrating from an existing cloud storage solution to Everabyte® involves three phases: assessment and planning (2-4 weeks), parallel operation and validation (4-8 weeks), and cutover and decommission (2-4 weeks). The total timeline for most enterprises is 8-16 weeks, depending on data volume and complexity.

Phase 1: Assessment and Planning

Before beginning data migration, the following assessment activities should be completed:

10. Data classification: Identify which data requires zero-knowledge storage (personal data, IP, regulated data) versus which can use standard encrypted storage.
11. Key management architecture: Determine whether to use Everabyte®'s managed KMS, your own HSM-backed KMS, or a hybrid approach.
12. Network topology review: Assess bandwidth between your primary sites and Everabyte® datacenters. Identify if dedicated connectivity (AWS Direct Connect, Azure ExpressRoute equivalent) is required.
13. Compliance mapping: Map your existing compliance obligations to Everabyte®'s compliance posture. Identify any gaps requiring supplementary controls.
14. Pilot scope selection: Choose a non-critical subset of data (typically 1-5% of total volume) for pilot migration and validation.

Phase 2: Parallel Operation

During parallel operation, both your existing storage solution and Everabyte® are active. New writes go to Everabyte®; reads can be served from either system. This phase allows:

- Performance validation under real workload conditions
- Workflow integration testing (backup jobs, application integrations, monitoring)
- Team training on Everabyte® client tooling and admin console
- Compliance documentation review and sign-off

Phase 3: Cutover

Cutover involves: completing migration of all remaining data from legacy storage, updating all application and backup job configurations to point to Everabyte® endpoints, validating all workflows in production, and decommissioning legacy storage (or maintaining read-only access for a defined retention period).

Integration Patterns

S3-Compatible API

Everabyte® exposes an S3-compatible API endpoint, enabling drop-in replacement for applications currently using AWS S3, MinIO, or any S3-compatible object store. Client-side encryption is handled transparently by the Everabyte® SDK, which wraps the S3 API and performs encryption before transmission. No application code changes are required for basic S3-API integrations.

Native SDK

For applications requiring maximum performance or advanced zero-knowledge features (client-side search indexing, key rotation, audit logging), Everabyte® provides native SDKs for Python, Go, Rust, Java, .NET, and JavaScript/TypeScript. The SDKs handle key management, chunk encryption, QUIC transport, and retry logic.

Backup Integration

Everabyte® integrates natively with leading enterprise backup platforms:

- Veeam Backup & Replication: Native plugin available, supports immutable backup targets
- Commvault: Everabyte® appears as an S3-compatible cloud library with client-side encryption
- Rubrik: Integration via S3-compatible endpoint with Everabyte® SDK wrapper
- Cohesity: Native DataProtect integration available in Enterprise tier

Operational Monitoring

Everabyte® provides a comprehensive observability stack for enterprise operations teams:

- Prometheus-compatible metrics endpoint: throughput, latency percentiles, error rates, storage utilization
- OpenTelemetry-compatible distributed tracing: full request traces from client encryption through storage layer
- Audit log stream: all object operations logged with cryptographic tamper-evidence (HMAC chain)
- Grafana dashboard templates: pre-built dashboards for common operational views, available in the Everabyte® GitHub repository

Operational Security: Audit Log Integrity

Everabyte®'s audit logs are protected by a cryptographic HMAC chain: each log entry includes an HMAC of the previous entry's hash, the current entry's content, and a server-held signing key (separate from all data encryption keys). This allows customers to verify that audit logs have not been tampered with or selectively deleted, providing forensic integrity for compliance investigations.

Chapter 8: The Road Ahead

Post-Quantum Cryptography Integration

The cryptographic foundations of Everabyte®'s zero-knowledge architecture — RSA for public-key operations, AES for symmetric encryption, SHA-2 for hashing — are secure against classical computers. They are not secure against sufficiently powerful quantum computers running Shor's algorithm (for RSA) or Grover's algorithm (for symmetric primitives, with reduced but not eliminated security margins).

The practical threat timeline is debated, but the US National Institute of Standards and Technology (NIST) finalized its post-quantum cryptography (PQC) standards in 2024. FIPS 203 (ML-KEM, formerly CRYSTALS-Kyber), FIPS 204 (ML-DSA, formerly CRYSTALS-Dilithium), and FIPS 205 (SLH-DSA, formerly SPHINCS+) are now the official recommendations.

Everabyte®'s PQC roadmap:

15. Q2 2026: Hybrid key agreement (X25519 + ML-KEM-768) for all new key exchange operations. Existing connections remain compatible.
16. Q3 2026: Full ML-KEM-768 key encapsulation for enterprise key sharing (replacing RSA-OAEP). Backward compatibility maintained for 12 months.
17. Q4 2026: ML-DSA-65 for all client authentication operations, replacing ECDSA P-256.
18. 2027: Full deprecation of non-PQC key exchange in enterprise tier. Consumer tier maintains compatibility bridge.

CDN Acceleration: 6x Global Edge Performance

Everabyte®'s current architecture routes all requests through our 7 primary datacenters. For read-heavy workloads — particularly media streaming, software distribution, and globally distributed team collaboration — an edge acceleration layer dramatically reduces latency for end users who are geographically distant from the nearest primary datacenter.

Planned for Q2 2026: Everabyte® Edge Network, a globally distributed CDN with over 200 points of presence. The challenge — unique to zero-knowledge storage — is that CDN edge nodes cannot decrypt cached content for inspection or transformation. Everabyte® Edge solves this by caching encrypted blobs at edge nodes, with decryption occurring in the client. Edge nodes serve ciphertext; the client's decryption is unchanged. The result: 6x reduction in content delivery latency for geographically distributed teams, with zero compromise of the zero-knowledge guarantee.

AI-Powered Client-Side Search

A frequent enterprise requirement is semantic search: the ability to find files by meaning, not just filename or metadata. Server-side search indexes enable this for plaintext storage but are incompatible with zero-knowledge encryption.

Everabyte®'s roadmap includes a client-side AI search engine: a compact language model (planned: a quantized 7B parameter model running locally) that generates semantic embeddings from file content client-side. Embeddings are stored in an encrypted local index. Search queries are executed entirely on the client against the local embedding index, with no query text or semantic vector ever transmitted to Everabyte® servers.

This approach — AI-powered semantic search within a zero-knowledge constraint — represents a category-defining capability that no competitor currently offers. Planned availability: Q4 2026 for desktop clients, Q2 2027 for mobile.

Compliance Roadmap

Anticipated compliance additions:

- FedRAMP Moderate Authorization: filing initiated Q1 2026, authorization anticipated Q4 2026
- UK Cyber Essentials Plus: certification in progress, anticipated Q2 2026
- Singapore MTCS Level 3: assessment initiated Q1 2026
- Australia IRAP Assessment: in planning for H2 2026

Investment in Privacy-Preserving Computation

Everabyte® has committed $12M in R&D investment through 2027 to privacy-preserving computation research. Active projects include homomorphic encryption for server-side aggregation operations (enabling server-side analytics without decryption), multi-party computation for key recovery (enabling account recovery without exposing master keys), and formal verification of our encryption pipeline implementation.

Conclusion: The Privacy-Performance Compromise Is Over

The cloud storage industry has long operated on a false premise: that organizations must choose between genuine privacy — zero-knowledge, client-side encryption with no server-side key access — and the performance their operations demand. Everabyte® was built to make that choice obsolete.

Through seven years of engineering work — hardware-accelerated encryption pipelines, QUIC-based transport, globally distributed Tier III+ infrastructure, and a commitment to genuine cryptographic zero-knowledge rather than "encrypted" marketing — we have produced a platform that delivers both guarantees simultaneously.

The evidence is in the benchmarks: 4x throughput advantage over encrypted competitors. 25ms cross-continental latency. Bandwidth scaling from 1 Gbps to 100 Gbps without pre-provisioning. Sub-second metadata operations. Six nines availability across 7 datacenters with triple-failure tolerance.

The evidence is also in the compliance posture: GDPR Article 32 technically and organizationally satisfied. HIPAA encryption safe harbor by architectural design. SOC 2 Type II with zero exceptions. ISO 27001 certified.

And the evidence is in the cryptographic foundation: Argon2id key derivation, AES-256-GCM authenticated encryption, HKDF-SHA-512 key derivation hierarchy, and a post-quantum migration roadmap that keeps Everabyte® ahead of the threat curve through 2030 and beyond.

4x Faster than any encrypted competitor	**99.9999%** Designed availability across 7 global datacenters
0 Bytes of plaintext ever seen by Everabyte® servers	**2026** Post-quantum migration begins Q2 2026

Next Steps

19. Request a live performance benchmark against your current storage provider — we will run the test on your actual workload, not synthetic benchmarks, and provide a side-by-side report.
20. Schedule a technical architecture review with an Everabyte® solutions engineer. We will map your compliance requirements, workload profile, and integration needs to a specific deployment architecture.

21. Start a 90-day enterprise trial with full feature access, dedicated support, and a signed BAA if you are in healthcare. No credit card required.
22. Review our security documentation: the full cryptographic specification, SOC 2 Type II report (available under NDA), penetration test executive summaries, and GDPR DPA template are available at security.everabyte.io.

Zero-knowledge storage done right is not a research project. It is in production, serving enterprise customers today. The question is no longer whether you can afford the performance cost of true privacy. The question is: can you afford to store sensitive data on a platform that can be compelled to decrypt it?

Everabyte® Privacy Engineering Unit

Engineering Playbook Series · Q1 2026 · 48 Pages

GDPR Art 32 Compliant · HIPAA BAA Available · SOC 2 Type II Audited · ISO 27001 Certified