CAEVES

SOLUTION DATASHEET

Cost-Optimized Storage, Easy & AI Ready

Built for Microsoft Azure

Enterprises today are burdened by exponential data growth. As file systems, NAS platforms, and archive repositories accumulate unstructured data, the economic and operational cost of storage has become a strategic bottleneck. Legacy storage solutions were built for retention, not for accessibility or intelligence, leaving vast volumes of enterprise information dark, underutilized, and expensive to maintain.

CAEVES Intelligent Deep Storage™ redefines intelligent deep storage for Microsoft Azure customers, Easy & AI-ready – Available in the Microsoft Marketplace

CAEVES combines the economics of cloud object storage, the usability of modern file systems, and the requirements of AI-ready data access into a single platform that operates within the customer's own Microsoft Azure tenant.

CAEVES software-defined storage solution delivers a high-performance file system interface to Azure Object Storage and allows organizations to dramatically reduce total cost of ownership (TCO) by up to 70% through data caching, tiering, and lifecycle management policies, while preserving continuous accessibility, security, and data governance.

As a result, organizations can unlock immediate strategic value from decades of historical data by integrating CAEVES Intelligent Deep Storage™ seamlessly with Microsoft 365 Search and Microsoft Copilot via the CAEVES Copilot Connector™ without data duplication, expensive RAG implementations, or disruptive changes to user and application workflows.

The Market Challenge | Dark Data. Explosive Growth. Structural Cost Pressure.

IDC projects global data creation will reach 291 zettabytes by 2027, with 80 to 90 percent unstructured. Enterprise environments are expanding at 25 to 40 percent year over year, and around 68 percent of unstructured data is created once and rarely touched again.

This data sits across file servers, NAS platforms, object stores, backup, and archive repositories, however, data storage budgets do not scale at the same rate. This imbalance is no longer incremental. It is structural.

Gartner estimates that approximately 30 percent of enterprise storage spend is consumed by cold or redundant data that delivers no active business insight. Large portions of enterprise capacity are allocated to information that is retained for compliance or history, but rarely accessed.

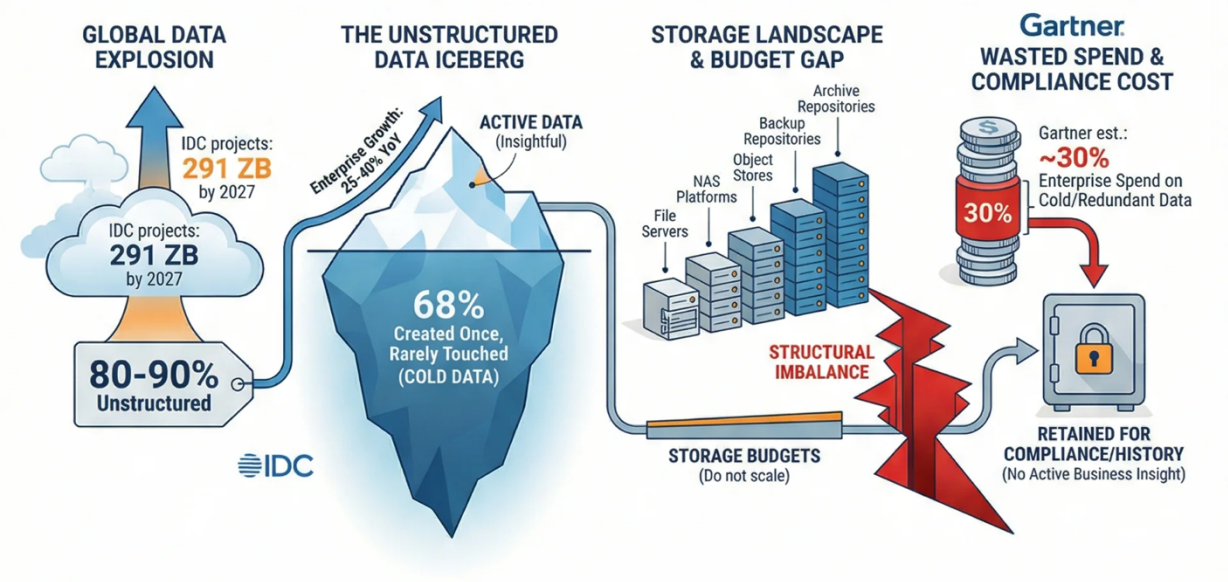

The consequences are measurable:

- ✕ Premium storage tiers carry inactive data for years
- ✕ Archives reduce cost but remove accessibility
- ✕ Intellectual property becomes operationally invisible
- ✕ Search and analytics initiatives operate on partial datasets

Enterprises Are Rich In Data

They are poor in accessible, governed intelligence

CAEVES Value & Promise | Transforming Storage Economics

CAEVES flips this equation by replacing expensive primary storage and legacy archive storage with intelligent, cloud-native deep storage built on Azure Object Storage. Rather than managing separate siloed systems, organizations now have a continuous storage layer that scales, protects, and democratizes unstructured data at long-term cost efficiencies, without disrupting user or application data access.

DRIVING COST EFFICIENCY
Cut Archive Costs by 70%

SECURE & COMPLIANT
Your Cloud, Your Control

EASY TO DEPLOY & USE
Deploys in < 30 minutes

Intelligent Deep Storage™

Built for Microsoft Azure

AI-READY INNOVATION
Your Data 100% AI-Ready

CAEVES delivers cost-optimized, deep and searchable storage for enterprise file and archive data, addressing the challenge of rising storage costs. Built natively for Microsoft Azure, our software-based solution slashes TCO by up to 70% through moving inactive "cold" data off premium tiers into deep object storage.

Unlike traditional archives that create "dark data", CAEVES keeps your files accessible and indexed. It provides AI-ready infrastructure, integrating seamlessly with Microsoft 365 Search & Copilot so your historical intellectual property remains a visible, strategic asset.

Key Benefits | Built for Microsoft Environments

- Offering fully available in the Microsoft Marketplace, provisioned in < 30 minutes
- Deploys & runs entirely inside your own Microsoft Azure (hybrid) cloud tenant
- Leverages native Azure Object Storage durability, resilience, and lifecycle capabilities without proprietary formats or vendor lock-in
- High-performance file access, data caching and tiering to Azure Storage (hot → archive)
- File and archive data remains indexed and discoverable using familiar interfaces (SMB) and Microsoft 365 Search & Copilot
- Provisioned, monitored, and managed like any other Microsoft Azure service

CAEVES

Unlock Your Enterprise Data & Archives | How CAEVES works

CAEVES delivers a structured architecture for consolidating and optimizing enterprise file and archive data within Microsoft environments.

Data is ingested through standard file interfaces and transparently tiered into durable object storage while preserving file semantics, metadata, and access controls.

Once stored, CAEVES integrates directly with Microsoft 365 Search and Copilot, allowing indexed, permission-aware access across historical and active data.

The outcome is a clean, extensible data foundation that aligns storage economics, security, and search with Microsoft-native services, without introducing proprietary dependencies or operational friction.

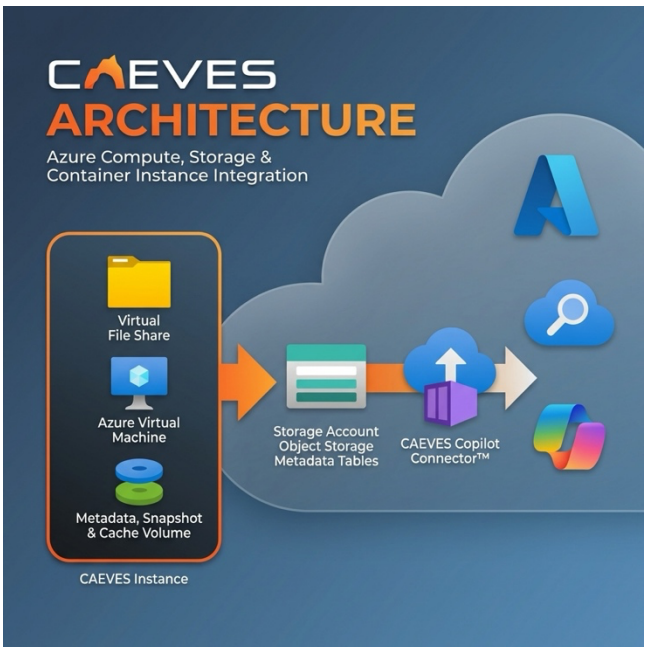

Two Phases. One Data Foundation

Lower cost, full access, and future intelligence built in

Phase 1: Transform & Optimize your "Dark Data" into CAEVES Intelligent Deep Storage™

- Deploy CAEVES with instant real-time file data access in less than 30 minutes
- Migrate your "dark data" to a secure CAEVES landing zone, backed by Microsoft Azure Object storage
- Preserve NTFS permissions and end-to-end security
- Automatically tier cold data to Azure Object Storage, from hot, to cool, cold and archive tiers
- Optimize costs up to 70% with automated lifecycle policies while maintaining seamless access

Phase 2: Index Intelligent Deep Storage™ to Microsoft Copilot, unlocking AI capabilities

- Make historical data instantly searchable and AI-ready by deploying the CAEVES Copilot Connector™
- Transparently integrate in Microsoft 365 Search & Copilot and Copilot Agents
- Eliminate complex RAG pipelines, expensive infrastructure and duplicate data silos
- Maintain security, compliance, and governance throughout the data pipeline
- Turn deep storage into an AI-ready strategic asset

CAEVES

CAEVES | Technical Overview

CAEVES presents an enterprise-grade real-time file system front-end, built on Windows Server technology, native NTFS and supporting SMB/NFS* protocols, while transparently connecting that namespace to Azure Object Storage as the authoritative data tier, subsequently indexing data into Microsoft 365 Search & Copilot using the CAEVES Copilot Connector™.

CAEVES unifies enterprise file services with Azure object storage in a single, future-ready architecture. Users continue to work with familiar interfaces, while data is intelligently tiered to cost-optimized Azure Object Storage under automated lifecycle policies. NTFS permissions, metadata, and snapshots are preserved throughout, ensuring governance and security remain consistent. The object storage layer becomes the authoritative foundation, delivering cloud-scale durability without changing user workflows.

CAEVES | Architecture & Data Flow

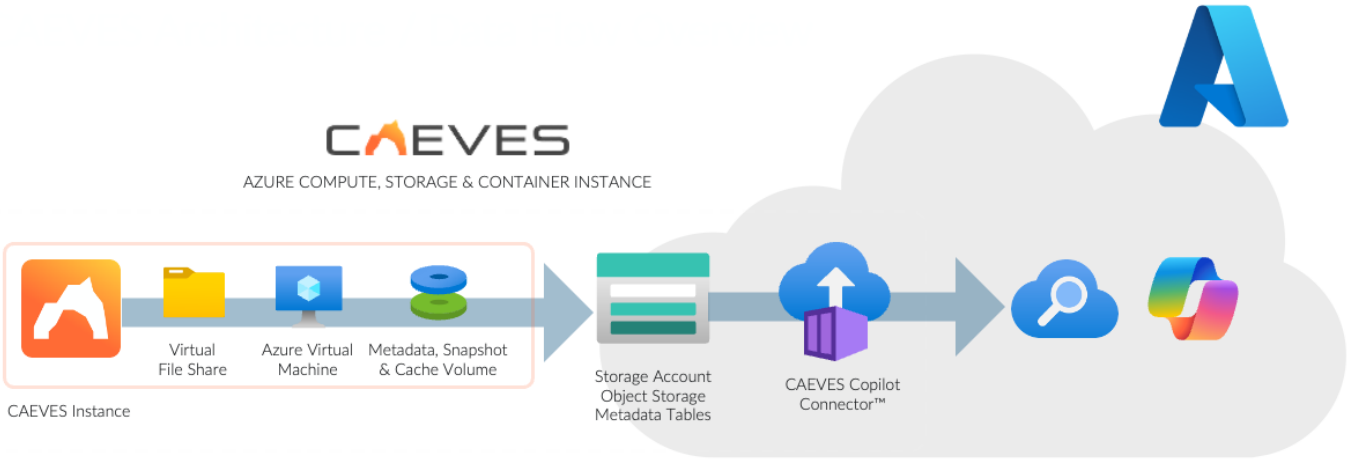

Multi-tier Data Architecture	Cloud Data & AI Integration	Enterprise-grade Security
▪ CAEVES VM Instance for data ingest ▪ Intelligent caching & data tiering (Hot/Cool/Cold & Archive storage) ▪ Public or Hybrid Cloud deployment options ▪ Multi-protocol support (SMB/NFS*) ▪ Integrated VSS snapshot functionality	▪ Connects Azure Object Storage for long-term data storage and index-based retrieval ▪ CAEVES Copilot Connector™ for indexing ▪ Native Microsoft 365 Search & Copilot integration	▪ Secure Entra ID / AD authentication ▪ End-to-end ACL / NTFS permissions support ▪ Item-level permissions in Microsoft 365 ▪ Filtering options for file types, source locations ▪ Staging Options to departments, limited staff

For CIOs, this means dark data is no longer stranded or economically misaligned. Historical file data becomes searchable, permission-aware, and AI-ready through native integration with Microsoft 365 Search and Copilot, without duplicating repositories or introducing complex pipelines. The result is lower storage costs, built-in ransomware resilience, and a scalable data platform that supports compliance, analytics, and AI initiatives within the organization's own Microsoft Azure environment.

CAEVES | Intelligent Deep Storage™ Fundamentals

When a file is created or modified, CAEVES file system commits metadata and full NTFS security descriptors locally, caches and tiers the data and transfers the payload to object storage using an optimized data pipeline. The local file is cached/stubbed depending on policy, enabling automatic tiering in Azure Object Storage across Hot, Cool, Cold and Archive tiers, without disrupting user access.

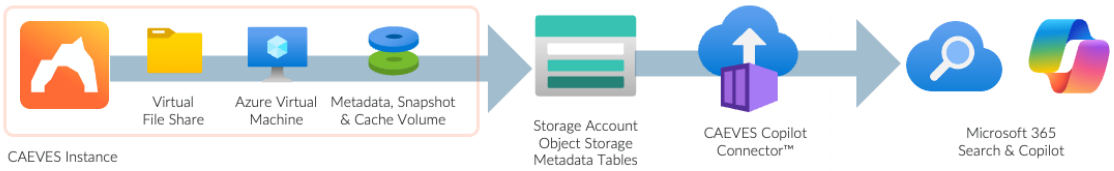

Metadata and file-to-object mappings are persisted on the Metadata Volume and stored in an Azure Table separately, creating a lightweight but durable control plane and processing interface for external services, i.e. AI, search and data analytics.

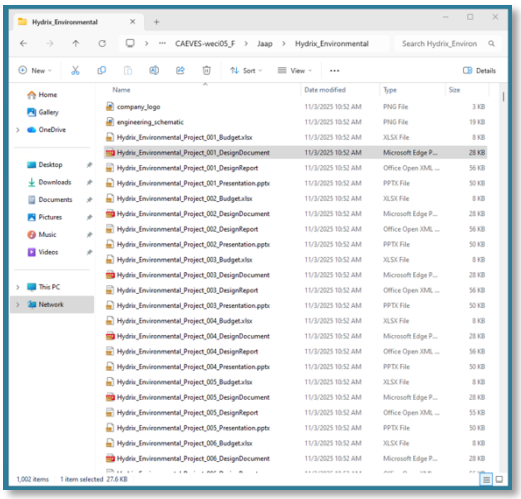

Throughout the process, identity and access semantics remain consistent: NTFS ACLs are preserved, federated with Active Directory and Entra ID, and enforced for every read and recall operation. Users continue to interact with a live file share, while the underlying storage operates with object-scale economics.

CAEVES | Data Protection, Ransomware & Disaster Recovery

CAEVES creates frequent, point-in-time snapshots at the file system level, including hourly file snapshots and daily, weekly and monthly full-volume snapshots with tiered retention. These snapshots preserve both file data and NTFS security descriptors, enabling users to restore previous versions directly from the live file system. In parallel, CAEVES captures and persists metadata snapshots to cold object storage on a scheduled basis, creating an additional recovery anchor independent of the active file system.

By combining snapshot-based rollback, optional object immutability (WORM), soft delete, and anomaly detection on file operations, CAEVES provides a strong defense against ransomware and accidental data loss. Repositories can be sealed after snapshot creation, preventing unauthorized modification of protected states. In the event of corruption or malicious encryption, organizations can rapidly revert to a known-good snapshot without requiring traditional archive restore workflows, supporting both operational resilience and disaster recovery objectives.

CAEVES | Copilot Connector™

In parallel, the CAEVES Copilot Connector™ indexes content directly from Azure Object Storage and Azure Tables into Microsoft 365 Search and Copilot. The connector processes files at scale, extracting text and enriching content where needed, embedding inherited ACL information into the search index, thereby grounding security permissions, with the option to implement filtering rules, helping to provide data guardrails before indexing to Microsoft Copilot.

In the following example, IT administrators and end users can immediately discover and retrieve customer records, file content, and associated metadata through Microsoft 365 Search. In this example, our fictional organization (NovaStruct) is querying project-related information for their customer HYDRIX.

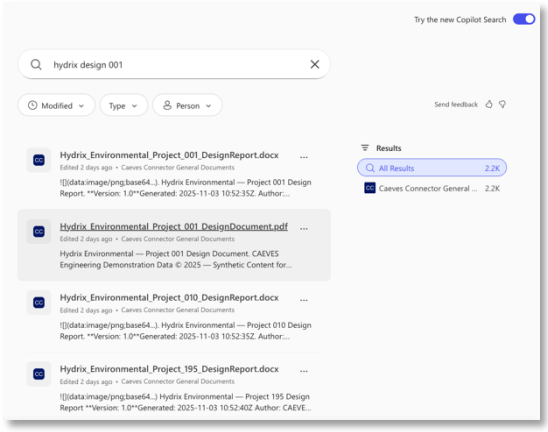

Microsoft 365 Search surfaces results indexed through CAEVES and enables direct file access using a secure, time-bound SAS URL generated for the corresponding Azure Object Storage object. This ensures controlled, audited access to the underlying content without exposing the storage account directly.

The search experience is powered by the CAEVES Copilot Connector™, which is registered as an external data source within Microsoft 365. The connector is deployed and managed as part of the CAEVES Copilot Connector™ instance, maintaining synchronized indexing, metadata integrity, and permission-aware access enforcement across the environment.

Within Microsoft 365 Copilot, our company NovaStruct has deployed a dedicated “NovaStruct Demo” agent configured to leverage CAEVES as a secure external data source. This agent enables retrieval-augmented generation (RAG) workflows grounded in the organization’s deep storage repository. Both IT administrators and end users can submit natural language prompts that trigger contextual retrieval from CAEVES-indexed content, while preserving inherited NTFS and Entra ID-based permissions.

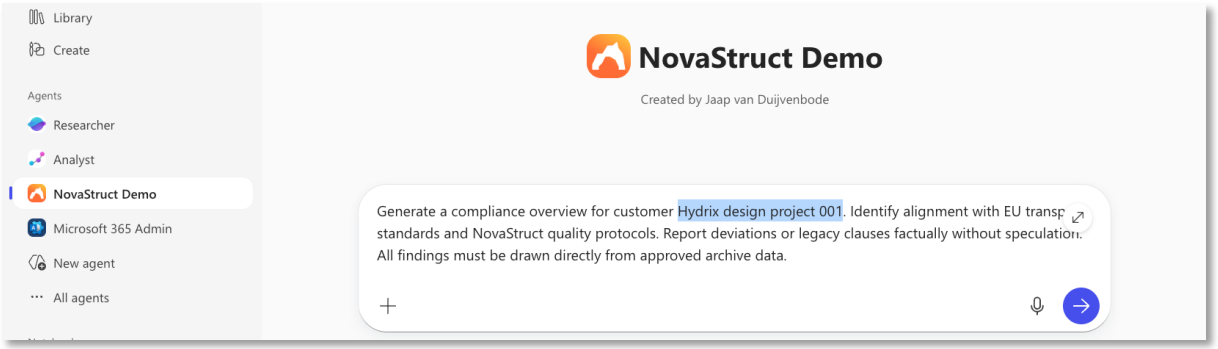

In this scenario, NovaStruct generates a compliance overview for project data associated with the customer HYDRIX. The agent, configured within Microsoft 365 Search and Copilot with predefined governance and compliance instructions, retrieves only authorized documents and metadata from CAEVES.

Once the user confirms the Copilot prompt, the agent initiates the report generation process, synthesizing insights directly from securely indexed deep storage data without duplicating or relocating the underlying content.

Microsoft 365 Copilot has generated a compliance overview for project information related to HYDRIX by grounding its response exclusively in data stored within CAEVES and accessible to the requesting user.

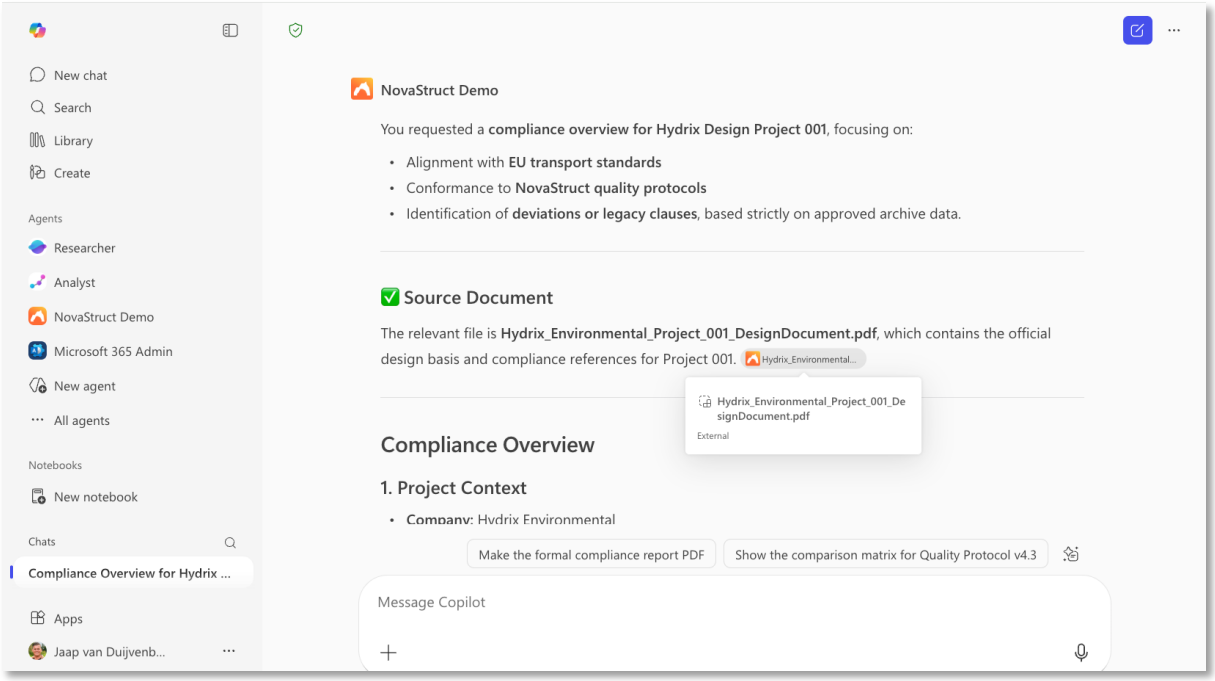

The NovaStruct Demo Agent processed the prompt using retrieval-augmented generation, ensuring that only content indexed from CAEVES deep storage was searched and contextualized. Permission enforcement remains fully intact, as NTFS ACLs are preserved and applied within the Copilot indexing framework, meaning users can only see and reason over documents they are entitled to access. Copilot references the exact source documents used in generating the report, provides full citation transparency, and enables direct download of those documents via secure SAS URLs. The CAEVES identifier within the Microsoft Copilot experience confirms that all referenced information originates from the CAEVES deep storage platform provisioned in the customer's Microsoft Azure environment.

CAEVES | Bring Decades Of Data To Life.

CAEVES Intelligent Deep Storage™ and the Copilot Connector™ ensure that discovery and AI-driven reasoning provides end-to-end security and authorization. As a result, data flows seamlessly from file system to object storage to search and AI, without duplicating repositories or introducing separate archive workflows.

The outcome is a unified architecture where cost-optimized deep storage, enterprise-grade security, and AI-ready search coexist within the customer's Azure tenant, maintaining consistent governance across the entire pipeline.

CAEVES | Pricing Model, Tiered, Built for Growth

Our pricing model combines the flexibility of monthly billing with the strategic commitment of an annual agreement, so you can scale on your terms without compromising budget efficiency.

Customers subscribe to a 12-month contract, but are billed monthly based on the maximum capacity used in that billing cycle. This means you can ramp up or down as needed, with billing that adapts to your real-world usage, perfect for unpredictable workloads and dynamic scaling.

We've designed a tiered structure that rewards scale, ensuring your total cost of ownership (TCO) improves as your data footprint grows:

Storage Capacity Tier	Price (USD/GB/Month)	Notes
0–5 TB	$0.00	Free tier — great for pilots & testing
6–50 TB	$0.03	Ideal for smaller teams or departments
51–200 TB	$0.025	Scales with efficiency
201–1000 TB	$0.02	Designed for mid-size enterprise scale
1001 TB and beyond	$0.01	Enterprise-grade, hyperscale economics

By combining dynamic usage billing with volume-based discounts, our model ensures that you're never locked into a static plan, and always positioned to optimize costs as you grow.

CAEVES | Feature Matrix

1. Solution Overview

Architecture Model	Hybrid File System + Azure Object Storage, indexed to Microsoft 365 Search & Copilot
Deployment	Azure Marketplace, customer tenant deployment using software installation package or virtual machine image
Target Use Cases	Archival file storage, deep storage cost optimization, distributed file consolidation, AI-ready deep storage
Primary Value	Up to 70% cost reduction while maintaining live file access, and indexing to Microsoft 365 Search & Copilot

2. Solution Architecture

CAEVES Instance	Windows Server Virtual Appliance
CAEVES Instance Disk Storage	Premium SSD
CAEVES Instance Volumes	Metadata, Cache, Snapshots
Disk-based Caching	File-aware Data Caching
Stub-Based Data Tiering	Write-back file operations support Transparent recall from object storage
Azure Object Storage	Authoritative payload layer for tiered file storage and snapshots
Azure Table Storage	Metadata persistence, AI processing
Lifecycle Policies	Configurable Hot/Cool/Cold/Archive tiering
Geo-Redundancy	Supported via reference architecture
Cross-Region Failover	Supported via reference architecture

3. CAEVES Instance Enterprise File Services

SMB Version	2.x / 3.x
SMB Multichannel	Yes
SMB Encryption in Transit	Yes
NFS Protocol	Yes*
NTFS ACL Preservation	Yes
Azure AD Kerberos	Yes
On-Premises AD Domain Join	Yes
Entra-ID Federation	Yes, via Entra-ID AD Connect
Cross-Domain Trusts	Yes

4. Azure Object Storage Data Lifecycle & Cost Optimization

Automatic Object Storage Tiering	Hot, Cool, Cold, Archive
Lifecycle Policies	Configurable
Object Storage Primary Tier	Yes
Transparent File Recall	Yes
Restore Workflows Required	No
File Size Limits	NTFS limits only
Cost Savings	Up to 70% vs hot-tier storage

5. Snapshot & Ransomware Protection

VSS File Snapshots	Hourly
Previous Versions	User-accessible “Previous Versions”
Full Volume Snapshots	Daily, Weekly, Monthly (USN-based)
Retention Tiers	Hourly to Annual
Immutable Storage (WORM)	Yes
Soft Delete	Yes, 30 days default
Snapshot Sealing	Yes, using Volume Shadow Copies
Rapid Rollback	Yes

6. Observability & Logging

Application Logging	Yes
Azure LogAnalytics Integration	Yes, for CAEVES Copilot Connector™
3rd Party Logging options	Yes, using syslog / event forwarding

7. Microsoft 365 Search & Copilot Integration

CAEVES Copilot Connector™	Azure Container Instance, supporting indexing to Microsoft Copilot
Direct Object Storage Indexing	Yes, using metadata (process) table
Processing Scale	Unlimited, horizontal scaling options, claims-based processing via table
Supported File Types	Documents: PDF (.pdf), Word (.docx, .doc), PowerPoint (.pptx, .ppt), Rich Text Format (.rtf), OpenDocument (.odt, .odp). Spreadsheets & Data: Excel (.xlsx, .xls, .xlsm, .xlsb), CSV (.csv). Text & Markup: Plain text (.txt), Markdown (.md), JSON (.json), XML. Images: PNG, JPEG, JPG, PJP, JFIF, GIF, WEBP. Archives & Others: Zip archives, Email files, Loop components
OCR Support	Yes
Docling Indexing	Yes
Image & Audio / Video Indexing (GPU)	Yes*
Markdown Normalization	Yes
ACL-Aware Indexing	Yes
Microsoft Graph Enforcement	Yes
Custom RAG Required	No
Pricing Model	Free of charge

8. Deployment & Commercial Model

Availability	Microsoft Marketplace
Deployment Time	~30 minutes
Provisioning	Automated ARM templates, available in the Microsoft Marketplace
Environment	Customer-controlled Azure tenant
Hybrid Support	Yes, can deploy CAEVES instance in a hybrid or on-premises datacenter, with object store residing in Microsoft Azure
Microsoft Licensing	SaaS Subscription
Free Tier	5 TB
Pricing Model	Tiered, usage-based, see https://caeves.com/products

** Coming soon*

Get started in the Microsoft Marketplace at https://www.caeves.com/signup

For more information, check https://www.caeves.com

Disclaimer

This document is provided for informational purposes only and reflects CAEVES Technology's current product capabilities and direction as of the date of publication. Features, functionality, and roadmap items are subject to change without notice. Nothing in this document constitutes a binding commitment, warranty, or guarantee of performance. Deployment results, cost savings, and performance outcomes may vary depending on configuration, environment, and usage.

CAEVES Technology, Inc. assumes no responsibility for errors or omissions in this document.