



# Cost-Optimized Storage, Easy & AI Ready

Intelligent Deep Storage™

---

Jaap van Duijvenbode  
Head of Product & Customer Experience



The  
**IT Press Tour**

# IT Press Tour 2026, Sofia Bulgaria - Agenda

- Introductions
- Company Background
- Industry Trends + Mission & Vision
- Challenges & Pain Points
- Solution Overview & Technical Architecture Discussion
- Customer Example “NovaStruct Engineering” Live Demo
- Competitive Overview Discussion
- Roadmap
- Go-To-Market & Pricing Strategy
- Future

**The  
IT Press Tour**



**2012-2020**

## Co-founded Talon Storage Solutions, Inc.

Built and launched Talon FAST™ Edge Caching technology from zero to acquisition by NetApp, Inc. in March 2020.

**2020-2025**

## Acquired by NetApp, Inc.

Team helped to scale global revenue for NetApp Cloud Data Services team, aligning cloud storage, data management strategy with the major hyperscalers

**2024-2025**

## Founded CAEVES Technology, Inc.

Established leadership team, developing file system and cloud storage technology, researching AI extensibility options in “stealth mode”

**2025-2026**

## Launched CAEVES Intelligent Deep Storage™

Private / Public Preview in Microsoft Azure Marketplace, followed by global launch of CAEVES Intelligent Deep Storage™ and CAEVES Copilot Connector™

### Company Snapshot:

Headquarters: Morristown, NJ, USA

Global Offices: USA, The Netherlands, India

Funding: Privately held / funded

Microsoft Partner Status: ISV Success Partner

### Leadership Team Background:



# CAEVES Leadership Team



Shirish H. Phatak

CEO / CTO



Jaap Van Duijvenbode

VP of Product &  
Customer Experience



Andrew Mullen

SVP Sales & Alliances



Charles Foley

Advisor

## Industry Trends

# 291

## ZETTABYTES

291,000,000,000,000,000,000,000 BYTES

Global data creation will grow to 291 zettabytes by 2027, with 80-90% of that data being unstructured

## Industry Trends

85%

Unstructured data is **created once**, sits dark, untouched by AI or analytics

## Industry Trends

30%

Enterprise storage budgets now go toward maintaining cold or redundant data that delivers zero business insight

# The Future of Data Has Changed

The Era Of Passive Storage Is Over

Data Intelligence Is The New Infrastructure

Economics Are Shifting From Cost To Capability

Compliance & Innovation Now Go Hand In Hand

The Winners Will Make All Their Data Intelligent



## Our Beliefs

Your forgotten data holds tomorrow's breakthroughs  
it deserves to be alive, not buried

**CAEVES**

## Our Mission

Help organizations to optimize storage costs while bringing decades of “dark data” to life

**CAEVES**



## Our Vision

The future of enterprise storage is not a cheaper disk.  
It is a smarter data foundation



CAEVES

The logo for CAEVES features the word "CAEVES" in a white, sans-serif font. The letter "A" is stylized with a gradient from orange to yellow, resembling a flame or a mountain peak.



## Our Vision

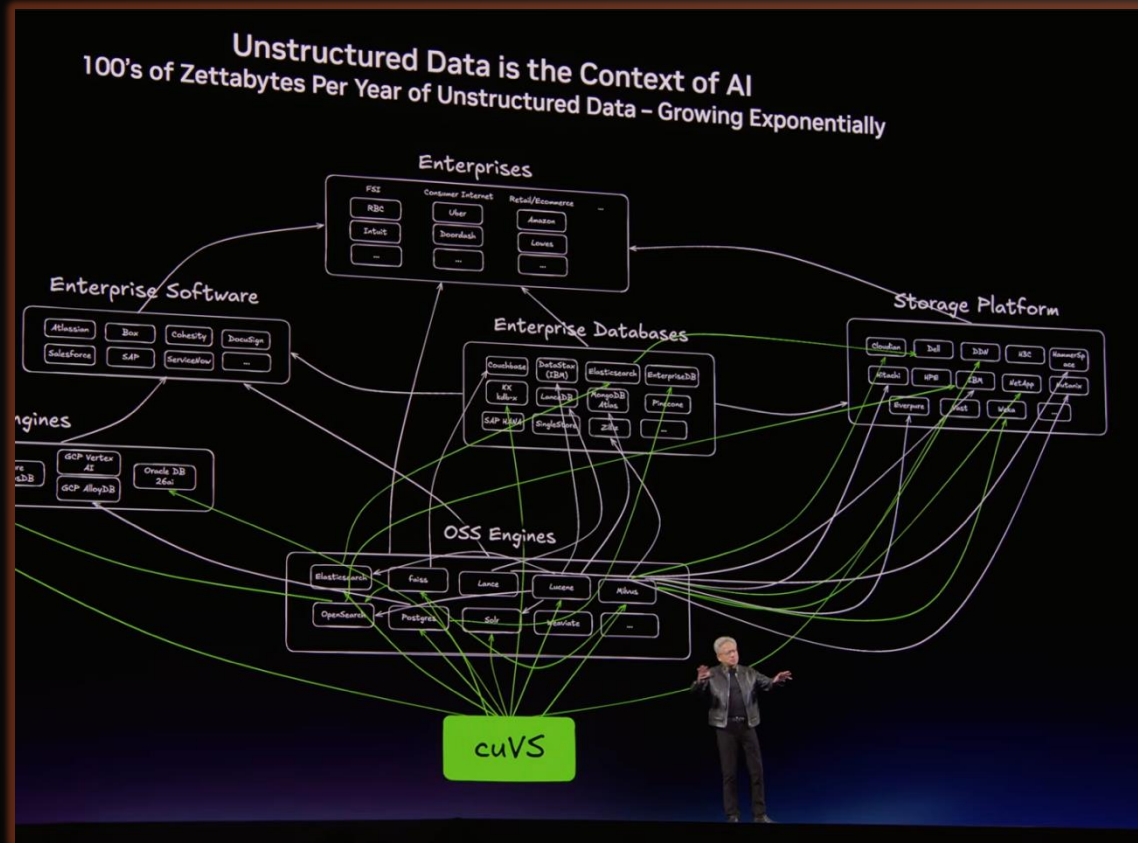
Cost-Optimized, Deep, Easy & AI Ready



CAEVES

The logo for CAEVES features the word in a white, sans-serif font. The letter 'A' is replaced by a stylized orange and yellow flame or mountain shape.

# Jensen Huang (NVIDIA CEO) on "Unstructured Data" in the AI era



NVIDIA GTC Keynote (March 16, 2026)

Until now, this data has been completely useless to the world.

Unfortunately, we can't query it, we can't search for it ...

There is no easy indexing of unstructured data ...

You'll have to understand it's meaning; it's purpose ...

# Challenges & Pain Points

291 ZB

Global data projected  
by 2027

*IDC / Gartner*

80–90%

Of all enterprise data  
is unstructured

*IDC*

68%

Created once,  
never touched again

*IDC*

30%

Of storage budget  
wasted on cold data

*Gartner*

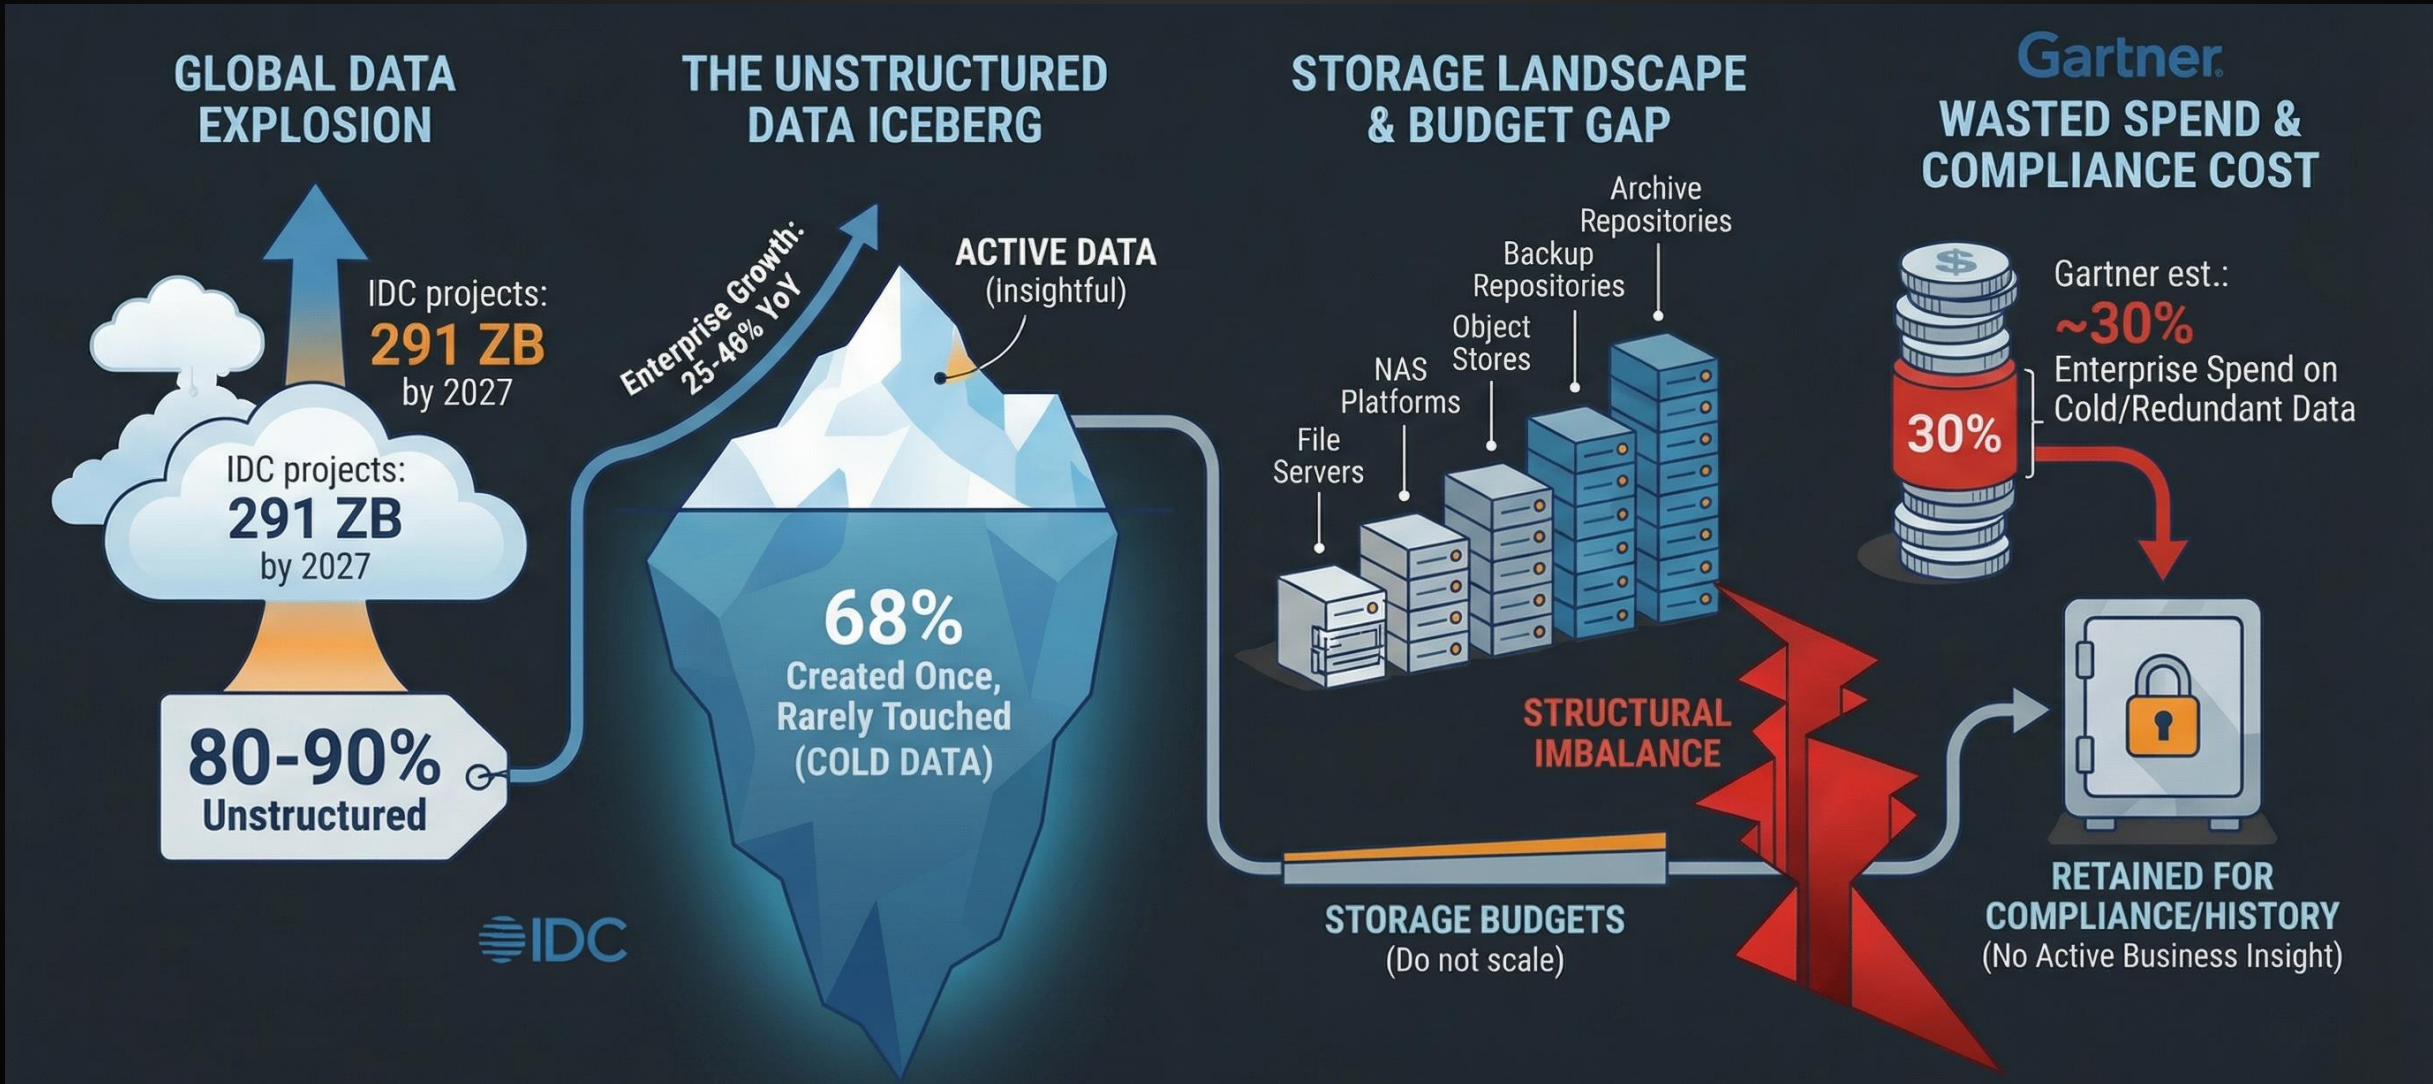
## The Economics Problem

Organizations continue to pay premium storage pricing, optimized for active data, for file shares where 60–80% of content has not been opened in months. The inefficiency is structural, persistent, and growing at 25–40% per year. Gartner estimates ~30% of enterprise storage spend delivers zero active business insight.

## The AI Readiness Problem

Modern Workplace, Search & AI tools, including Microsoft 365 Copilot can only surface knowledge from data that is indexed and accessible via Microsoft Graph. When data is archived to cold tiers or legacy systems, it disappears from Copilot, from Azure AI Search, and from your organization's institutional memory, permanently undermining AI investment ROI.

# Challenges & Pain Points



# The AI Readiness Gap | The Data Democratization Dilemma

Historical data provides context, trend validation, and institutional memory.

## LEGACY STORAGE ARCHITECTURES (Not AI Ready)



- Archives lack full indexing, search or visibility.
- Cold storage introduces latency and restore workflows.
- Permission models break when data moves.
- RAG implementations frequently duplicate data.

FRAGMENTATION

GOVERNANCE  
EXPOSURE

RISING COST  
LAYERS


## AI FAILURE

AI FAILS: Foundational data inaccessible, unstructured, or disconnected from security controls.



## AN UNSUSTAINABLE TRADE-OFF

× PREMIUM  
STORAGE  
COSTS



× REDUCE COST  
& accept dark  
archives

IT leaders forced into limiting decisions.  
CIOs exiting data centers, consolidating, preparing for AI.  
Compliance & Sovereignty are **non-negotiable**.

## THE STRUCTURAL RESET



**THE MARKET DOES NOT NEED FASTER STORAGE.  
IT NEEDS A STRUCTURAL RESET.**

A new model where cost, control, accessibility,  
and AI readiness operate together.

# Challenges & Pain Points

## Rising Storage Costs & Data Sprawl

Enterprise data volumes double every 2 years, driving uncontrolled infrastructure spend with no visibility into underutilized capacity.

## Expensive Cloud & File Storage for Archive Data

Premium cloud tiers and legacy NAS systems charge active-storage rates for rarely accessed archive and compliance data.

## Limited Growth Capabilities

On-premises storage arrays hit hard capacity ceilings, forcing costly forklift upgrades that delay business growth.

## Lack of On-Demand Capacity

Traditional storage procurement takes months, leaving teams unable to respond to sudden data growth or project demands.

## Lack of Historical Access to Business Information

Archived files are siloed or offline, making retrieval slow and error-prone for audits, litigation, or analytics.

## Complex & Expensive Custom RAG Implementation

Building AI pipelines over fragmented data requires costly custom indexing, chunking, and embedding infrastructure per source.

## Search Stops at the Storage Boundary

Enterprise search tools index M365 content but go dark the moment files are moved to cold or archive storage tiers.

## "Dark Data" is Inaccessible to AI Services

Unstructured files outside active tiers, including decades of reports, emails, and contracts, are invisible to Copilot and AI models.

## Disparate Data Sources = Vendor Lock-In

Multi-vendor storage landscapes create dependency silos that inflate licensing costs and block migration or consolidation.



# Cost-Optimized Storage, Easy & AI Ready

Built for Microsoft Azure

CAEVES delivers cost-optimized, deep, and searchable storage for enterprise file data. Cut storage costs by up to 70% without losing access or future AI readiness.





## COST EFFICIENCY

Cut Archive Costs by 70%

Drive lower TCO through intelligent tiering to Azure's most cost-effective storage, keeping everything instantly searchable.

## SECURITY & SOVEREIGNTY



Your Cloud, Your Control

Deploys entirely within your Microsoft Azure tenant, hybrid or public cloud. All data stays under your governance.



## EASY TO DEPLOY & USE

Deploys in < 30 minutes

Deliver 50x faster search results using Microsoft 365 Search and Copilot. No IT tickets. No waiting for tape restores. Just instant answers from your archive data.

## AI-READY INNOVATION



Your Data 100% AI-Ready

Historical information becomes context for AI conversations, trend analysis, and strategic insights, without exposing sensitive data or retraining models.



## Our Promise

Your data, your rules — secured, governed and compliant, in your own Microsoft Azure environment

**CAEVES**

### 01 Archive Cost Reduction

Cut TCO up to 70% moving cold data off premium tiers to Azure Object Storage via intelligent lifecycle policies.

### 02 Legacy NAS Modernization

Retire aging NetApp/EMC/HP infrastructure while preserving NTFS permissions, SMB semantics, and user workflows.

### 03 M365 Copilot Data Expansion

Index historical archives into M365 Search & Copilot via CAEVES Copilot Connector™ — no custom RAG required.

### 04 Compliance & Retention Archiving

Enforce WORM immutability, versioning, and audit-ready governance entirely within your own Azure tenant.

### 05 Ransomware Protection & Recovery

Layered defense: hourly VSS snapshots, snapshot sealing, blob versioning, and rapid rollback — no tape needed.

### 06 Dark Data AI Activation

Make decades of file data AI-ready for Azure OpenAI Studio, AI Search, Microsoft Foundry, and ML workloads.

### 07 Hybrid Cloud Consolidation

Unify on-premises and remote environments into a single platform tiering to Azure Object Storage via SMB/NFS.

### 08 Edge-to-Cloud Deployments

Extend intelligent data management and AI readiness to field offices, construction sites, and distributed teams.

### 09 Infrastructure Simplification

Eliminate siloed backup, archive, and NAS gateway systems on native Azure storage with zero vendor lock-in.

### 10 M365 Search for Historical Data

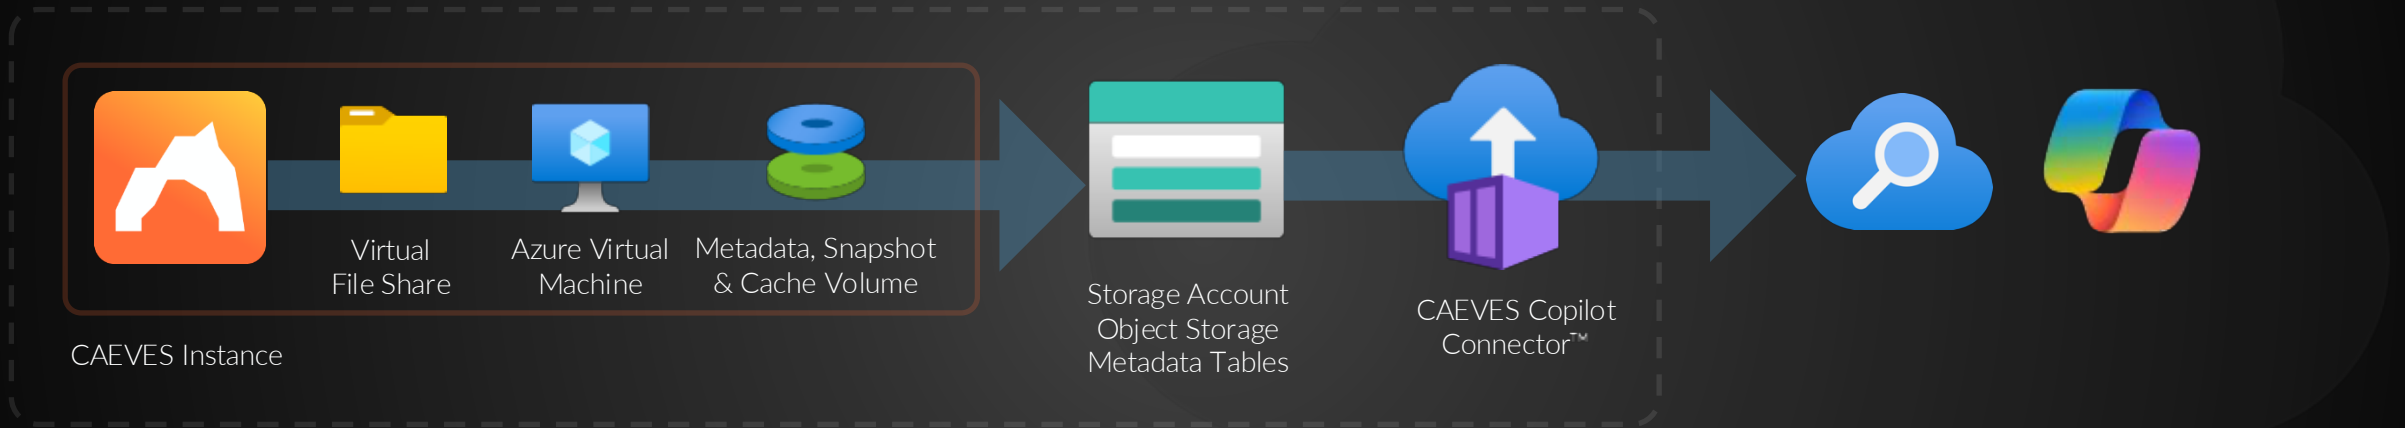
Deliver instant, permission-aware results across years of archived files through familiar Microsoft 365 tools.

# CAEVES Architecture / Data Flow Overview



## CAEVES

AZURE COMPUTE, STORAGE & CONTAINER INSTANCE



### Multi-tier Data Architecture

- 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

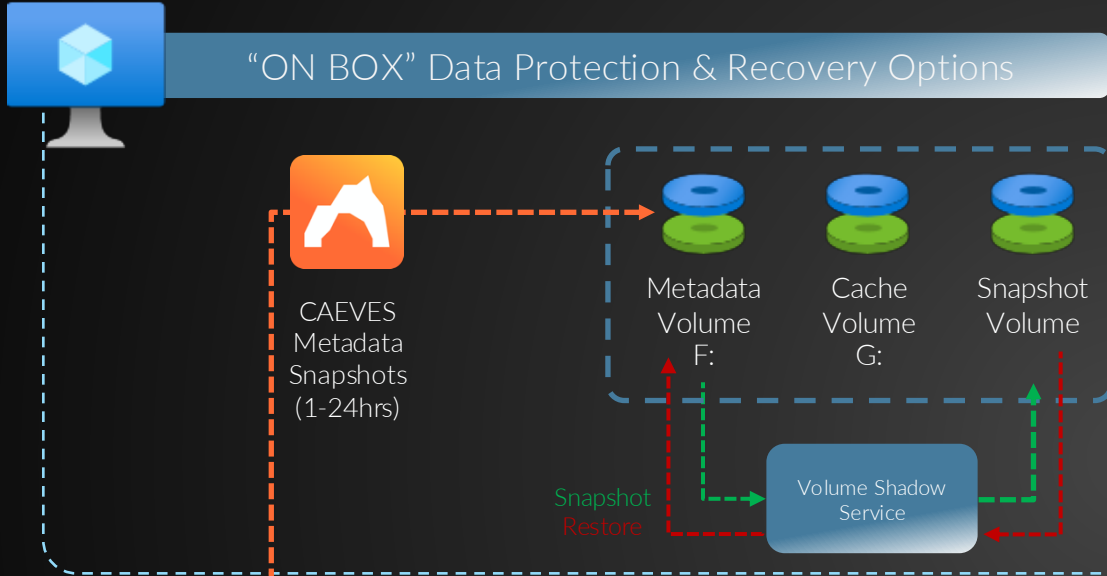
### Cloud Data & AI Integration

- Connects Azure Object Storage for long-term data storage and index-based retrieval
- CAEVES Copilot Connector™ for indexing
- Native Microsoft 365 Search & Copilot integration

### Enterprise-grade Security

- 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

# CAEVES Snapshot Capabilities



## CAEVES Metadata Snapshot Mechanism (to Object Store)

- Full snapshot of metadata volume
- Full ‘Previous Versions’ support for folders, files
- Protecting individual file changes in Cloud Object Store
- In-place file restore, or external volume/folder/file restore
- Stored in Cloud Object Store along with corresponding data payload
- Disaster Recovery options, restore in-cloud, or hybrid datacenter
- Aligned with cloud versioning, WORM, and legal/compliance targets

## “OBJECT STORE” Data Protection & Recovery Options

CAEVES Metadata Snapshot [ID]

CAEVES Container Blobs

CAEVES Container

CAEVES Metadata Table

Name	Blob Type	Access Tier	Size	Content Type	Last Modified
7197004-700-482-8adc-07ab654cf93b_0_1000000000A0	Block Blob	Cool (inferred)	45.30 KB	application/octet-stream	07/06/2025, 00:01
7197004-700-482-8adc-07ab654cf93b_0_1000000000A1	Block Blob	Cool (inferred)	263,75 KB	application/octet-stream	07/06/2025, 00:01
7197004-700-482-8adc-07ab654cf93b_0_1000000000A2	Block Blob	Cool (inferred)	66,50 KB	application/octet-stream	07/06/2025, 00:01
7197004-700-482-8adc-07ab654cf93b_0_1000000000A3	Block Blob	Cool (inferred)	94,89 KB	application/octet-stream	07/06/2025, 00:01
7197004-700-482-8adc-07ab654cf93b_0_1000000000A4	Block Blob	Cool (inferred)	64,06 KB	application/octet-stream	07/06/2025, 00:01
7197004-700-482-8adc-07ab654cf93b_0_1000000000A5	Block Blob	Cool (inferred)	35,39 KB	application/octet-stream	07/06/2025, 00:01
7197004-700-482-8adc-07ab654cf93b_0_1000000000A6	Block Blob	Cool (inferred)	61,87 KB	application/octet-stream	07/06/2025, 00:01
7197004-700-482-8adc-07ab654cf93b_0_1000000000A7	Block Blob	Cool (inferred)	299,00 KB	application/octet-stream	07/06/2025, 00:01
7197004-700-482-8adc-07ab654cf93b_0_1000000000A8	Block Blob	Cool (inferred)	26,82 KB	application/octet-stream	07/06/2025, 00:01
7197004-700-482-8adc-07ab654cf93b_0_1000000000A9	Block Blob	Cool (inferred)	156,79 KB	application/octet-stream	07/06/2025, 00:01
7197004-700-482-8adc-07ab654cf93b_0_1000000000AA	Block Blob	Cool (inferred)	160,98 KB	application/octet-stream	07/06/2025, 00:01
7197004-700-482-8adc-07ab654cf93b_0_1000000000AB	Block Blob	Cool (inferred)	20,14 KB	application/octet-stream	07/06/2025, 00:01

	Purpose	Mechanism
<b>Cool Tier Blob Store</b>	Main storage for files & objects	Content-level storage from metadata refs
<b>Metadata Snapshots</b>	Uploaded via CAEVES snapshot mechanism	Retained using blob versioning + lifecycle
<b>Versioning</b>	Keeps past states of blobs & metadata	Blob versioning, blob snapshots
<b>Soft Delete Policy</b>	30 days for blobs & containers	Protection against accidental deletes
<b>WORM &amp; Immutable Storage</b>	For audit/compliance archive	Apply to monthly/yearly snapshots

## Before CAEVES: Traditional Storage



## After CAEVES: Cost-Optimized Storage

Rising archive/storage costs & data sprawl

Expensive primary tiers & legacy NAS systems

Different silos of data, multiple vendors

Complex & fragmented infrastructure

Vendor lock-in, different skills required

Finite growth capabilities, lack of on-demand capacity

Lack of historical access to business information

Costs scale linearly with data growth

Archive equals loss of access

Search stops at the storage boundary



Rising Costs & Data Sprawl

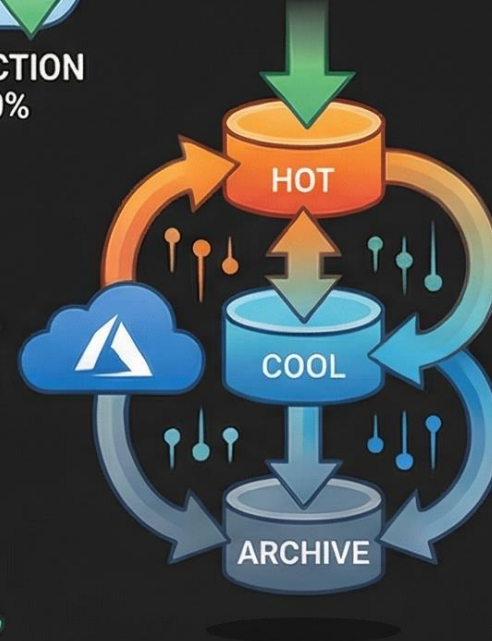


Inaccessible historical data for analytics or AI



COST REDUCTION UP TO 70%

Cost reduction up to 70% by tiering cold data



No rehydration required, no performance penalty

Unified management across on-premises, edge, & cloud

Preserved NTFS ACLs & Entra ID permissions

Infinite data growth & management options in Microsoft Azure

Dynamic intelligent tiering from hot to deep archive

Object storage absorbs exponential growth & cost

Data remain accessible, searchable, indexable

Archives ready for Microsoft 365 Search & Copilot

No rehydration required, no performance penalty

Data sovereignty & compliance within your Azure tenant

AI-ready infrastructure without redesign



# Key Takeaways

- **70% lower archive costs** through intelligent Azure tiering, no tape, no rehydration fees, no compromise
- **Search decades of data in seconds** using Microsoft 365 tools your team already knows
- Provision and secure in **your own Azure environment**
- Turn **dark data into actionable intelligence** with native Microsoft 365 Search & Copilot integration
- Scale predictably and sustainably with **pay-as-you-grow pricing** that evolves with your business needs



AVAILABLE IN THE  
MICROSOFT MARKETPLACE



# CAEVES

## “LIVE” DEMO

YOUR DATA, YOUR INTELLIGENCE,  
ALWAYS WITHIN REACH

---

# DEMO: NovaStruct Engineering



## Company Profile

- Company: NovaStruct Engineering
- Sector: Structural & Civil Engineering
- Locations: Americas, EMEA, APAC
- Employees: 20,000
- Core Expertise: Structures, bridges, transport infrastructure
- Typical Projects: Urban viaducts, tunnels, logistics hubs

## IT Profile / Challenges & Initiatives


- Datacenter exit, decommissioning DC
- Data increasing at 200% Y/O/Y
- Archive storage needs better visibility
- Microsoft Azure Cloud initiative

## Data Sources

Dell Technologies

NetApp IBM

Commvault



## Cloud Storage

aws wasabi

Google Cloud


# 2PB+



# Demo


CAEVES

# DEMO: NovaStruct Persona's



## Zach Ulrich (CIO)

- Responsible for defining NovaStruct's future data strategy
- Biggest headache: **Control data growth & costs, IT governance**
- Current project: DC exit strategy, moving to Microsoft Azure
- IT Steering committee is pushing for AI adoption



## Lars de la Rocha (Solutions Architect)



- Supporting IT demands of NovaStruct
- Biggest headache: **Keeping up with end user demands, industry trends, while maintaining a secure environment**
- Current project: Moving data from EOL storage platform to Microsoft Azure



## Lemmy Osbourne (Business)

- Project Designer / Architect
- Biggest headache: **Need to do more with less, hard to find historical information / reuse data**
- Current project: designing a large airport together with global team of architects, deadline end of 2026!

# Conclusion: Where is NovaStruct today?



**2PB+ consolidated, data center exited:** All scattered data unified in one intelligent platform

**30 years searchable in seconds:** Engineers find designs, audit compliance, and learn from history instantly

**AI-ready and scalable:** 200% growth managed automatically, critical archives now answering questions through Copilot

# Customer Propensity Example 2000TB (Growth + Monthly ISV / ACR)

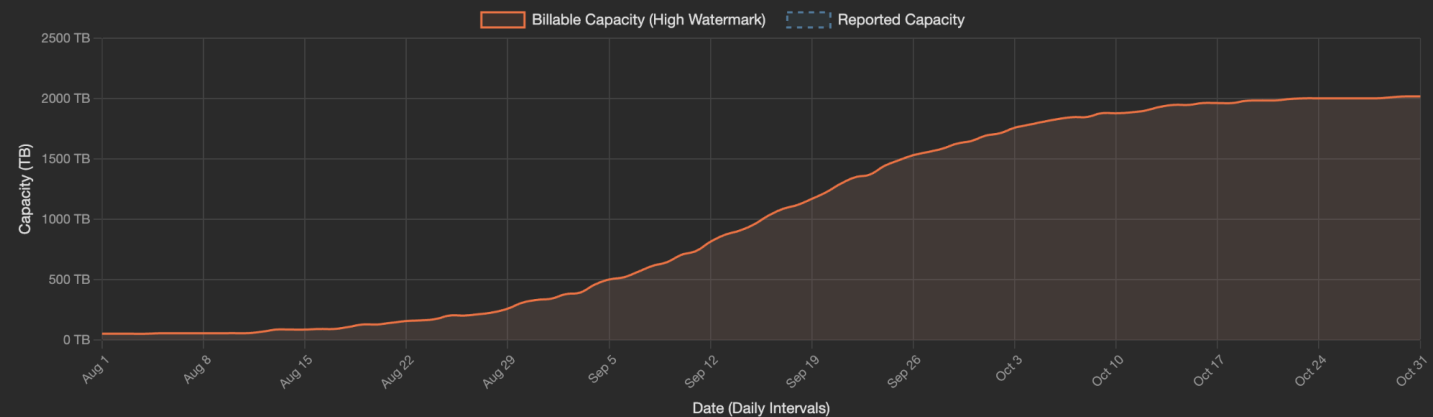
## CAEVES Revenue & Utilization Report

3-Month Analysis: August 2025 - October 2025

CUSTOMER <b>Jaap@shpresearch.onmicrosoft.com</b>	SUBSCRIPTION ID <b>30110926-03e1-41cf-d9e9-83db3ba558c2</b>	PLAN <b>caeves-basic-plan-3</b>	INSTANCE <b>shpr-jvctest4</b>
STATUS <b>Active (Subscribed)</b>	AZURE TENANT ID <b>30c492df-c7ea-4ddd-9016-b7706724302b</b>		

Total Net Revenue (3 Months) <b>\$40,730</b> After 3% Microsoft Commission	Current Capacity <b>2,000 TB</b> 1,953.13 TB Growth	Current Tier <b>Tier 5</b> \$0.01/GB/Month	Avg Monthly Net Revenue <b>\$13,577</b> +2,310% Growth
--	---	--	--

### Capacity Growth Trend



## ISV / Azure Consumption Monthly Cost (USD)

CAEVES Licensing \$31,100.00

Azure Storage \$20,480.00

Azure Compute \$275.00

Azure Disk Cost \$34.56

Azure Infrastructure Cost \$20,789.56

Total Monthly Cost **\$51,889.56**

CAEVES Licensing Cost/GB/mo \$0.016

Total Monthly Cost per GB **\$0.026**

# Competitive Intelligence (Discussion)

CAEVES transforms cold, legacy storage into secure, governed, AI-ready knowledge – data stays in your own Azure tenant with full compliance and AI access. Unlike most competitors, CAEVES stores files in **open, standard formats (Azure Object Storage)** – no proprietary backend lock-in, direct object store access always preserved.

Product	Performance Profile	Protocols	Tiering	Search / Copilot	Open Backend / No Vendor Lock-in	Service \$/GB-mo
<b>CAEVES</b>	Premium SSD + Object tiers	SMB / NFS / Azure Object Storage / SDK	Caching, Tiering (Hot → Archive)	✅ Integrated (native)	✅ <b>Open – Azure Blob. Direct access. Zero proprietary format.</b>	0.01 – 0.03 *
Egnyte	Cold Storage	SMB / Web / REST	Tiered	✅ Internal	❌ Proprietary (EnSync)	~0.22 *
Nasuni	Edge cache + object	SMB / NFS	Policy-based	⚠️ Metadata only	❌ Proprietary (UniFS)	~0.085 *
Panzura	Global FS	SMB / NFS	Hot / Archive	✅ Add-on	❌ Proprietary (CloudFS)	0.07 (0.023 archive) *
NetApp (ONTAP + FabricPool)	Edge + SAN/NAS, Cloud tiering	SMB / NFS / S3	FabricPool (Hot → Cold object)	⚠️ Limited (partner RAG)	⚠️ WAFL/ONTAP format (cold tier via native S3)	0.06 – 0.15 *
Dell EMC (ECS / PowerScale)	Object-native + Scale-out NAS	S3 / NFS / SMB	Policy-based (Hot → Archive)	❌ None (no Copilot)	⚠️ Proprietary namespace (ECS metadata layer)	0.01 – 0.12 *
CTERA	Edge + object	SMB / NFS	Policy-based	❌ None	❌ Proprietary (CTERA FS)	0.02 – 0.08 *
<b>Hyperscaler Offerings / Platform Costs</b>						
Azure Files	HDD / SSD	SMB / NFS	Hot / Cool	❌ None	✅ Open (NTFS/NFS native)	0.01 – 0.16
Azure Blob	Object only	REST / SDK	Hot → Archive	⚠️ Add-on	✅ Open (native object)	0.002 – 0.02+
Azure NetApp Files	SSD / NVMe	SMB / NFS	Standard, Premium, Ultra w/ Cool	❌ None	⚠️ ONTAP format	0.15 – 0.39 (0.08 cool)
AWS S3 (+Gateway)	Object + EC2	S3 / NFS	Standard → Glacier	⚠️ Add-on	✅ Open (native S3)	0.023+

**CAEVES Unique:** Your data is stored directly in your own Azure Object Storage environment in native object format. No CAEVES-proprietary encoding, no vendor-specific file system layer – your files remain fully accessible through multiprotocol file access (SMB/NFS) standard cloud APIs at all times. Migrate away or access data directly, any time, with zero extraction cost.

The logo for CAEVES, featuring the word "CAEVES" in a white, sans-serif font. The letter "A" is stylized with a gradient from orange to yellow, resembling a mountain range or a stylized letterform.

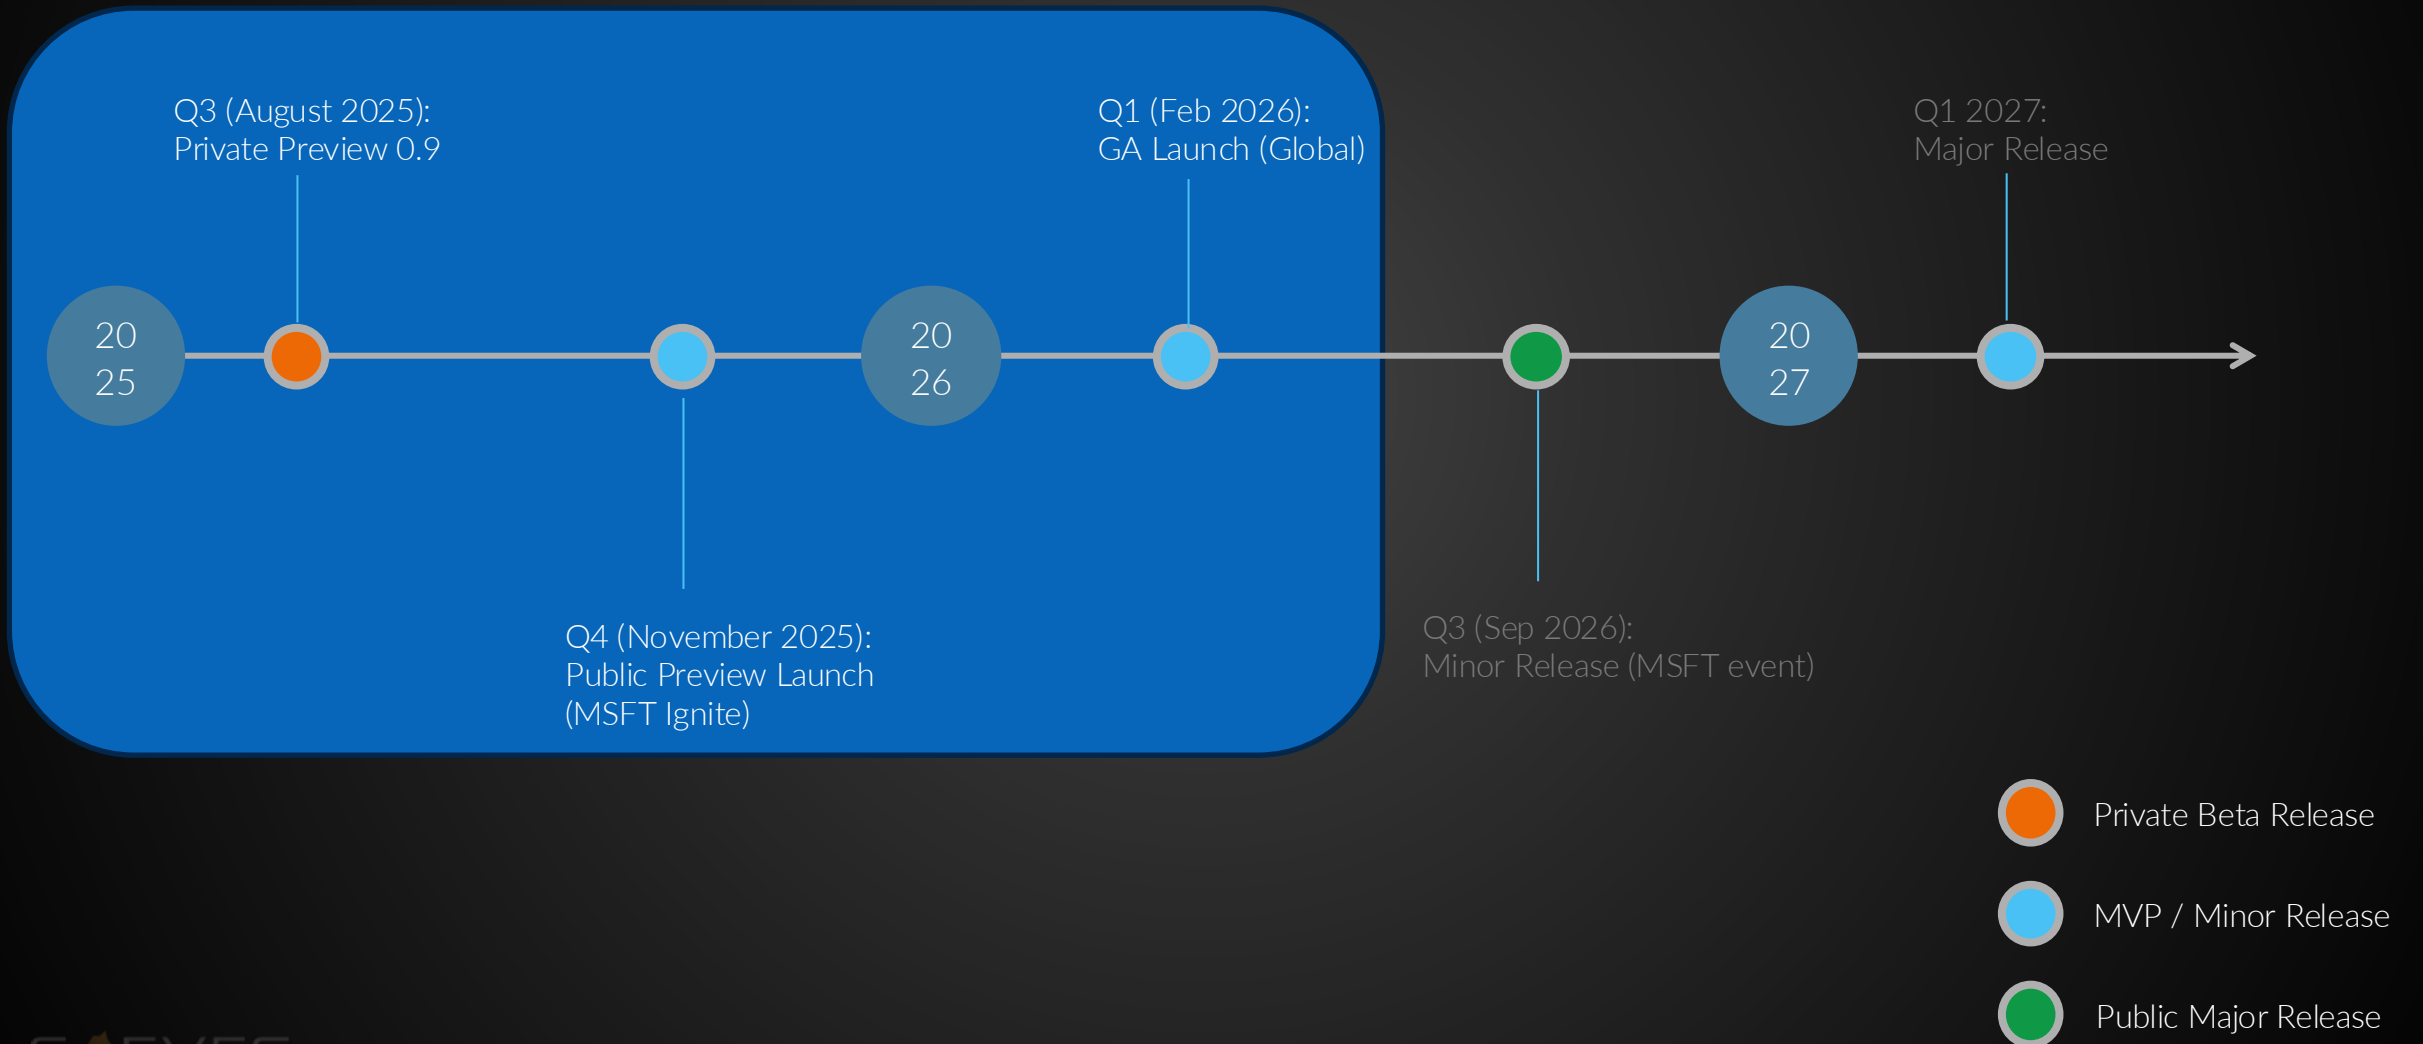
# CAEVES

# ROADMAP

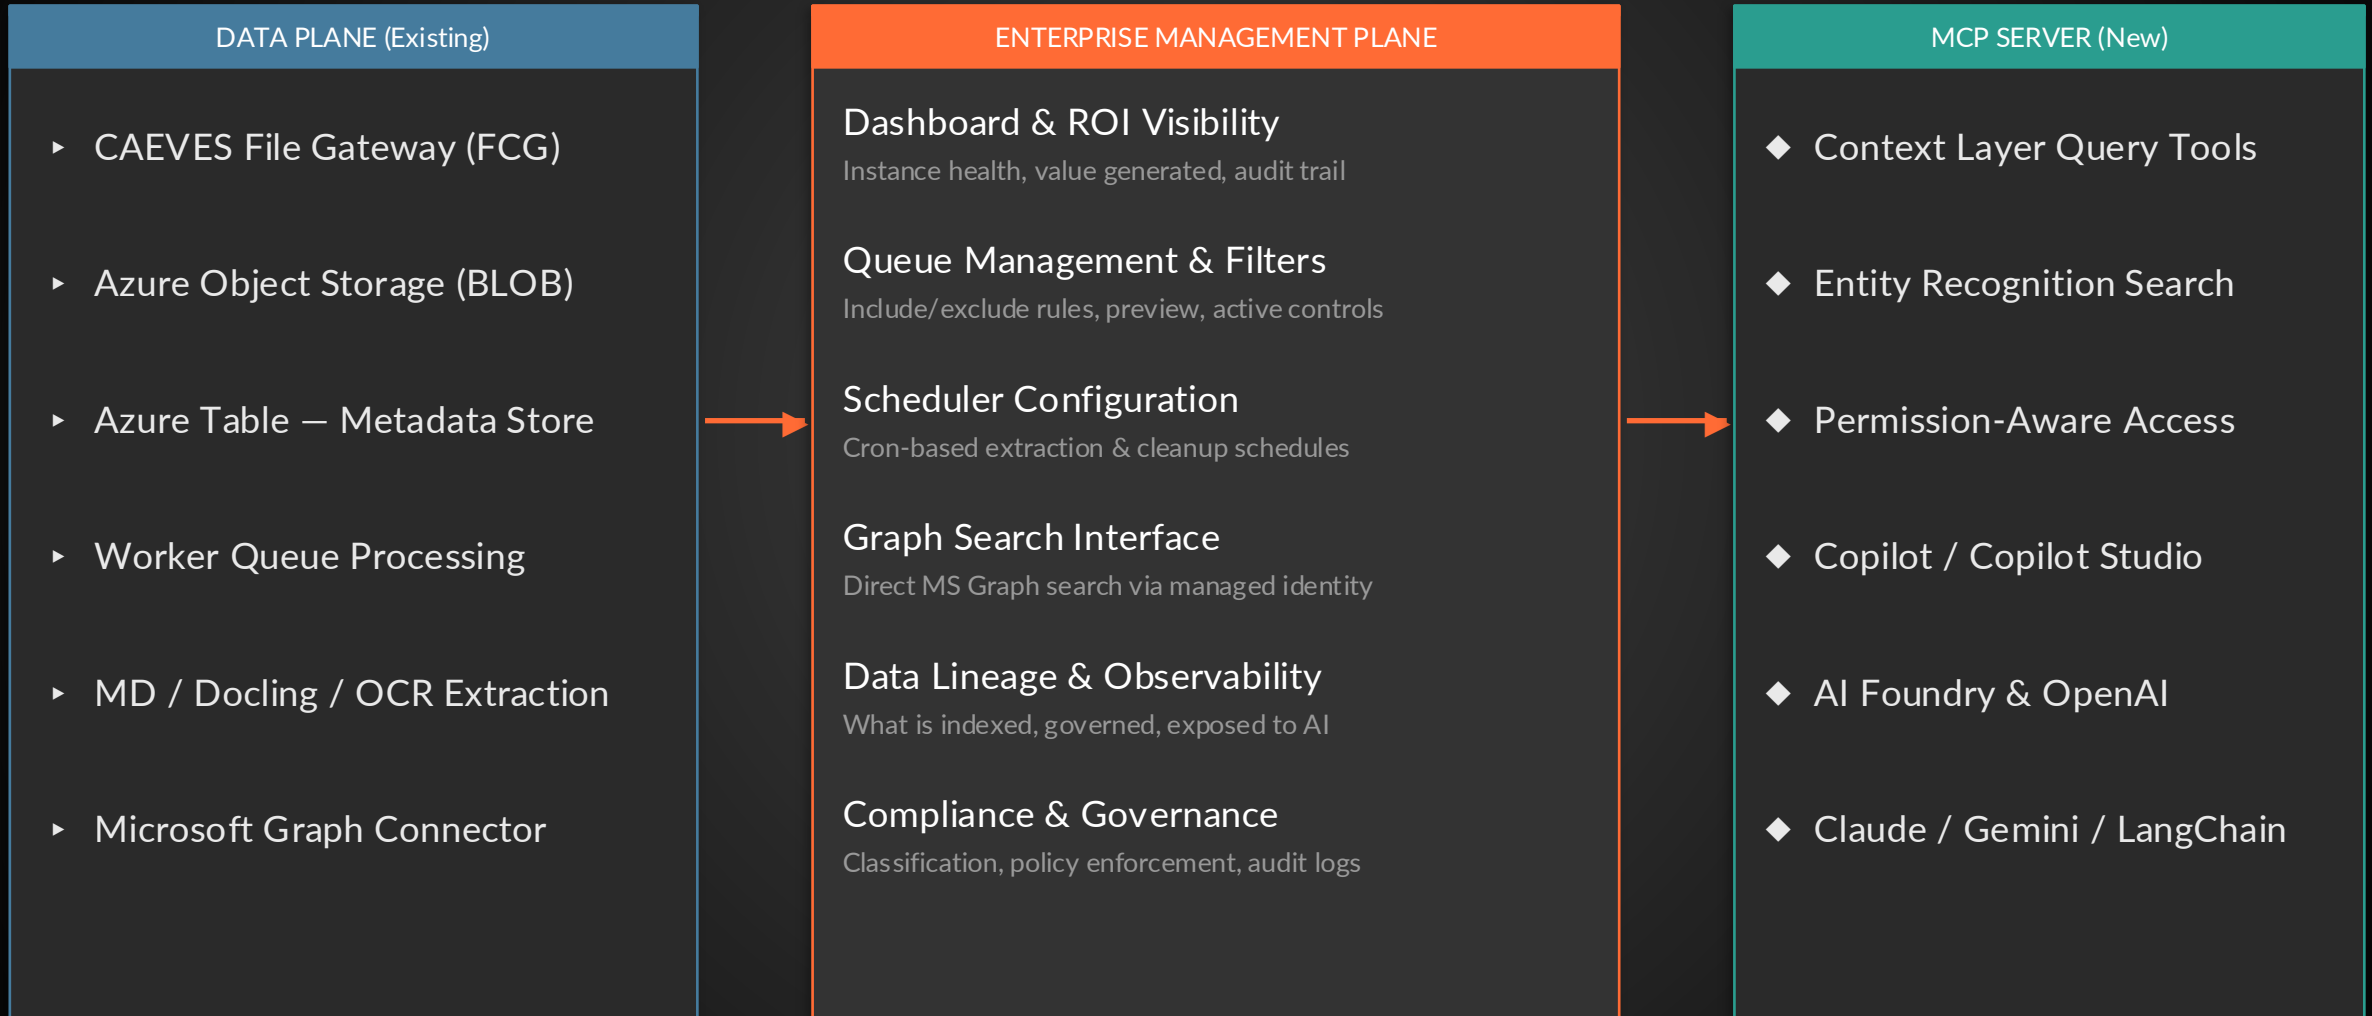
YOUR DATA, YOUR INTELLIGENCE,  
ALWAYS WITHIN REACH

---

# PRODUCT - Roadmap



# Enterprise Management Plane + MCP Server



CAEVES

# GO-TO-MARKET & AVAILABILITY

YOUR DATA, YOUR INTELLIGENCE,  
ALWAYS WITHIN REACH

---

# Go-To-Market

## GO-TO-MARKET CHANNELS



### Azure Marketplace

MACC-eligible; one-click deploy; global reach from day one.



### Microsoft Co-Sell

Azure Sales, STU/ATS/CSA and Product Strategy teams aligned to drive more Cloud Data & AI.



### SI / VAR Channel

Microsoft Gold/Solutions Partner SIs; deal registration + co-marketing.



### MSP & OEM

Multi-tenant MSP mode (future).

## PRICING MODEL

### Capacity-based Subscription via Microsoft Marketplace

Per TB under management/month

No upfront commitment. Pure outcome alignment.

### Pay-as-you-Grow -> Annual Commitment

Pre-committed TB volume; discounted

1-year and 3-year tiers. MACC eligible.

CAEVES cost is typically recovered within the first month of tiering savings for organizations with >200 TB under management.

# Generally Available in Microsoft Marketplace

PR Newswire | Send a Release

## Intelligent Deep Storage™ from CAEVES Now Available in the Microsoft Marketplace

NEWS PROVIDED BY CAEVES Technology →  
Feb 09, 2026, 08:45 ET

Microsoft customers worldwide can now discover and deploy CAEVES Intelligent Deep Storage™ through Microsoft Marketplace, accessing trusted solutions that accelerate innovation and business transformation with unified integration across Microsoft products




NEW YORK, Feb. 9, 2026 /PRNewswire/ -- CAEVES Technology, Inc. today announced the availability of CAEVES Intelligent Deep Storage™ in the Microsoft Marketplace, the unified online destination for customers to buy trusted cloud solutions, AI apps, and agents to meet their business needs. CAEVES customers can now discover and deploy trusted solutions through Microsoft Marketplace, with smooth integration and streamlined management across Microsoft Azure and other Microsoft products.

CAEVES Intelligent Deep Storage™ is a software-defined deep storage solution designed to help enterprises modernize archive and unstructured data within their own Microsoft Azure environments. The solution enables policy-driven data tiering to Azure object storage while maintaining secure, transparent access for enterprise workloads, compliance requirements, and AI-enabled discovery. Through integration with Microsoft 365 search and Microsoft 365 Copilot, organizations can securely extend historical data into modern search and AI workflows while maintaining existing identity, security, and governance controls.

“We're pleased to welcome CAEVES Intelligent Deep Storage™ to the growing Microsoft Marketplace ecosystem, making it easier for customers to discover and deploy solutions that work seamlessly with Microsoft products.” said Cyril Belikoff, vice president, Microsoft Azure Product Marketing.

<https://www.prnewswire.com/news-releases/intelligent-deep-storage-from-caeves-now-available-in-the-microsoft-marketplace-302681584.html>



 <b>CAEVES - 1. Intelligent Deep Storage - Subscription</b> Caeves Technology SaaS 1. Get started by activating the CAEVES Subscription in your own Azure environment. Price varies Subscribe	 <b>CAEVES - 2. Instance Quickstart</b> Caeves Technology Azure Application 2. Deploy your CAEVES Instance in Azure. This is the "landing zone" for your historical data. Price varies Create	 <b>CAEVES - 3. Copilot Connector Quickstart</b> Caeves Technology Azure Application 3. Securely connect your data to Microsoft 365 Search & Copilot with our Copilot Connector. Price varies Create
--	--	---

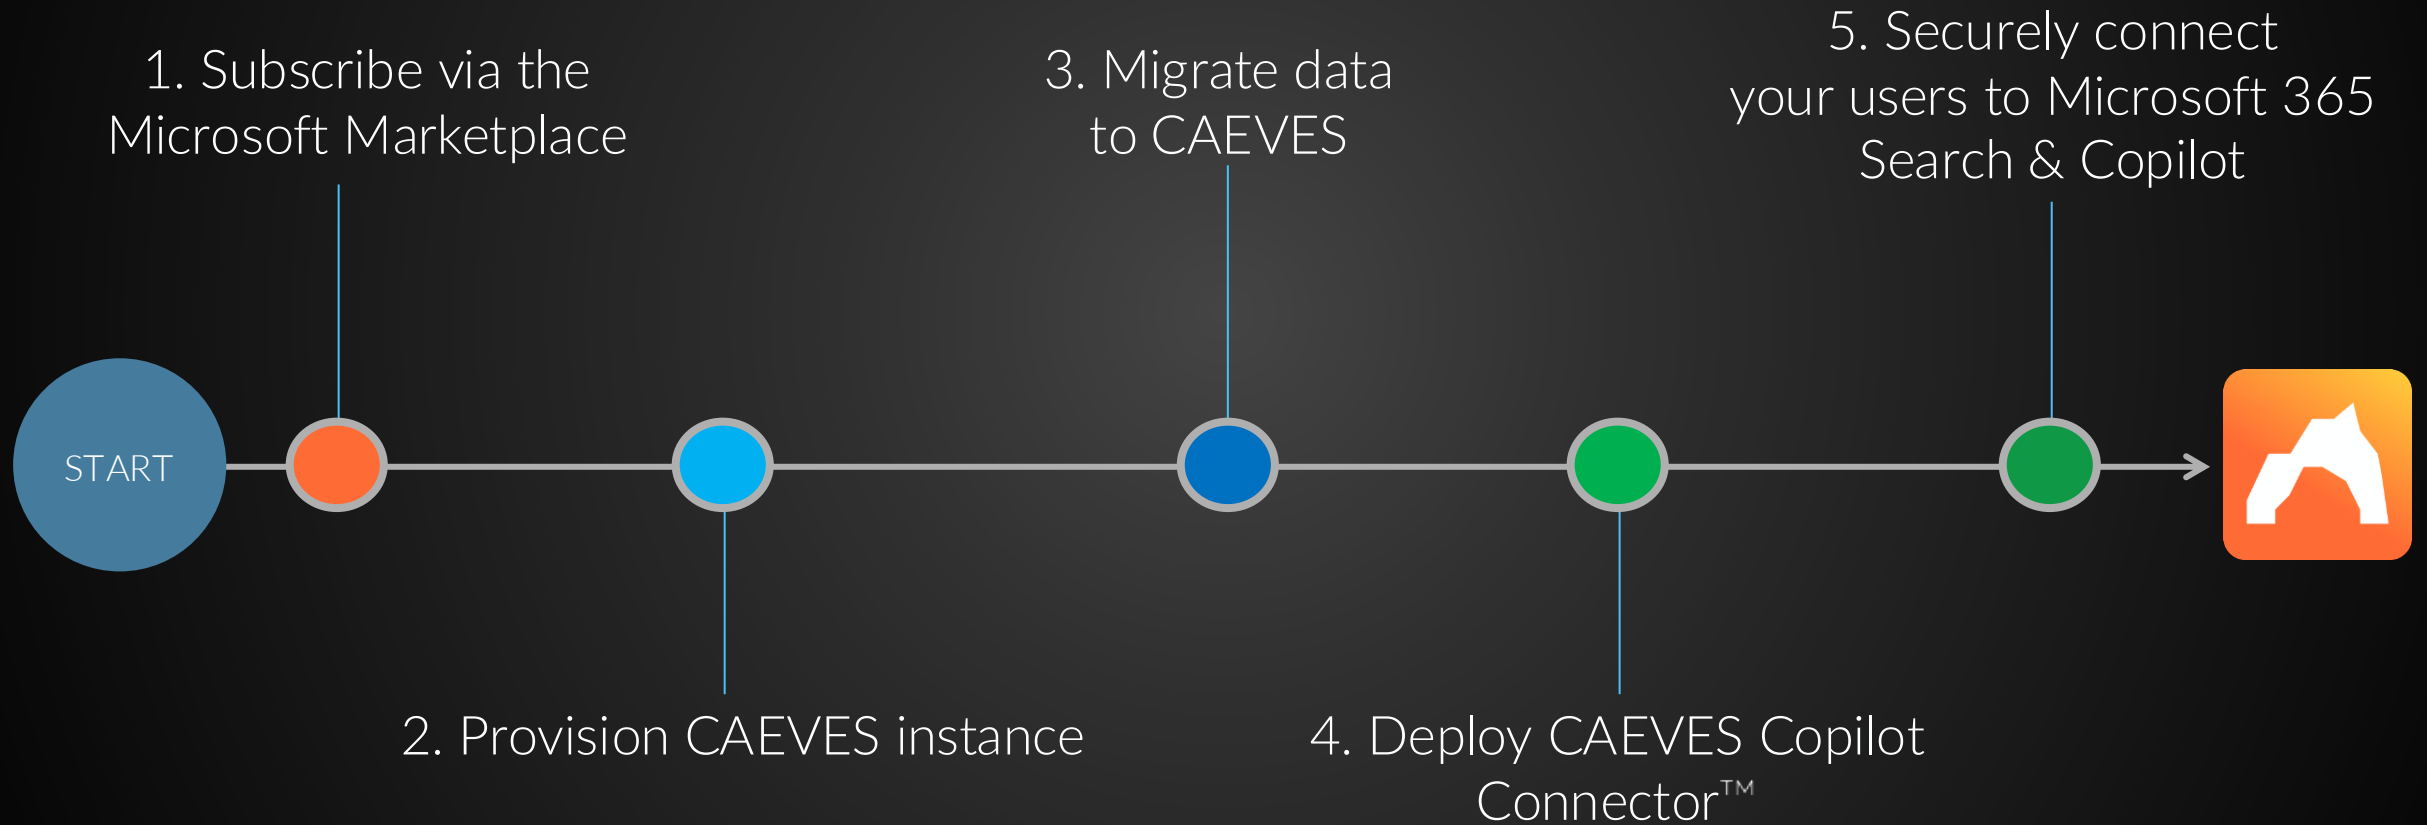
# CAEVES Pricing Structure

Storage 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

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 sign 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.

# CAEVES Activation, Migration & Onboarding



The logo for CAEVES, featuring the word "CAEVES" in a white, sans-serif font. The letter "A" is stylized with a gradient from orange to yellow, resembling a mountain peak or a stylized letterform.

# CAEVES

# FUTURE

YOUR DATA, YOUR INTELLIGENCE,  
ALWAYS WITHIN REACH

---

# Future

CAEVES is building toward becoming the intelligence, data & context layer for enterprise data estates within the Microsoft Azure ecosystem.

## From Storage Management to Data Intelligence



The analytics and metadata enrichment capabilities in CAEVES form the foundation for content classification, data sensitivity scoring, and data lineage tracking. Storage optimizer today; data intelligence platform tomorrow.

## Autonomous Storage Operations & Cost Optimization



The logical endpoint of policy-driven tiering is policy that writes itself. CAEVES is building toward a model where the platform recommends, simulates, and executes optimization decisions autonomously.

## Copilot, Copilot Studio & MCP as the Enterprise Interface



M365 Copilot is becoming the primary interface through which knowledge workers interact with enterprise data. CAEVES intends to be the infrastructure that makes Copilot comprehensive, and extending to Model Context Protocol, leaving no institutional knowledge invisible.

## Strategic Company Trajectory



Azure Marketplace traction, SI channel buildout, team expansion. The founding team has navigated the full cycle from creation to exit. CAEVES is being built to create significant enterprise value.

## Cost Efficiency & Sustainability

What if you could find a sustainable, cost-effective spot for your data, without losing access or compliance?



CAEVES



→ Cut costs by up to 70% with intelligent tiering to Azure object storage.

## Data Governance & Consolidation

What if you could keep decades of history in a single, intelligent platform that's AI-ready?



CAEVES



→ Consolidate legacy file shares & archives into one secure, searchable environment.

## Intelligent Search & AI

What if your entire archive was instantly searchable, indexable, and ready for AI conversations?



CAEVES



→ Turn dark data into AI-powered assets for Microsoft Search, Copilot & analytics.

## Infrastructure Simplification

What if you could retire traditional file services & NAS gateways, but keep file-aware access and security?



CAEVES



→ Simplify infrastructure while preserving authentication, permissions & audit trails.

## Key Takeaways: Three Things to Remember

70% of your storage cost is recoverable, without losing any data or any AI capability.

**CAEVES**

## Key Takeaways: Three Things to Remember

CAEVES is the only platform that keeps Microsoft 365 Search & Copilot connected without vendor lock-in.

**CAEVES**

## Key Takeaways: Three Things to Remember

This team has built and exited a company in this exact space. We know how to **build, scale and execute**.

**CAEVES**

# CAEVES

## Q&A

YOUR DATA, YOUR INTELLIGENCE,  
ALWAYS WITHIN REACH

---



The  
**IT Press Tour**