CAEVES

Technical Overview

Reference Architecture

Version: 1.8

Date: June 20, 2025

Executive Summary

In today's digital landscape, organizations across industries are recognizing that their historical data archives represent untapped intellectual property that could drive innovation and competitive advantage. According to Gartner, by 2025, 80% of enterprises will shut down their traditional data centers, compared to 10% today, with data volumes growing at an unprecedented rate of 23% annually. This digital transformation imperative is driving organizations to modernize their data estates while simultaneously seeking ways to leverage AI to unlock value from decades of accumulated intellectual property.

CAEVES addresses this transformation through an innovative an AI-powered file platform built for Microsoft Azure. CAEVES revolutionizes how enterprises interact with their enterprise data and digital archives. Built natively for Microsoft Azure, it combines intelligent deep storage with Microsoft Copilot integration, enabling natural language interaction with archived content while optimizing storage costs and reducing infrastructure complexity. This transformation is particularly impactful across key industries:

Recent market analysis underscores the urgency of this transformation:

- IDC predicts that by 2025, global data creation will grow to 175 zettabytes, with enterprises managing 80% of that data
- Gartner forecasts that by 2025, 75% of enterprise-generated data will be created and processed outside a traditional centralized data center
- GigaOm research indicates that AI-powered data management can reduce storage costs by up to 70% while improving data accessibility by 85%
- The enterprise information archiving market is projected to reach USD 18.89 billion by 2030, with a CAGR of 14.1%

As organizations navigate this transformation, the integration of AI capabilities becomes crucial. According to McKinsey, companies that fully embrace AI-powered data management can expect to see a 50% increase in workforce productivity and a 20-30% reduction in total ownership costs.

CAEVES enables this transformation by bridging the gap between enterprise data and archives, introducing modern AI capabilities, turning historical data from a cost center into a strategic asset for future innovation and decision-making.

Solution Overview

CAEVES is offered through the Microsoft Azure Marketplace by CAEVES Technology, Inc. as a solution that seamlessly integrates with the Microsoft ecosystem. The solution consists of two primary components:

1. **CAEVES Instance**: A hardened virtual machine appliance that serves as a hybrid cloud gateway, presenting SMB/NFS* mapping through a caching/tiering mechanism, connecting to object storage container, used for data ingest and migration purposes.

2. **CAEVES Copilot Connector**: A containerized service using Azure Container Instance that enables indexing using various Microsoft AI solutions, including Graph API, Azure AI Foundry, Azure AI Search, Azure OpenAI studio, Azure ML enabling Copilot capabilities and future extensibility options.

Key Benefits

- Deploy CAEVES in your own hybrid or public Microsoft Azure cloud datacenter
- Connect to Microsoft AI & Machine Learning services, including Copilot, AI Search, OpenAI Studio and Graph
- Simplified data management with automated snapshots, versioning and backup
- Sustainable long-term data retention with immutable deep storage and archives
- Extensibility options within Microsoft Ecosystem (Purview, Sentinel, Monitor)

Technical Architecture

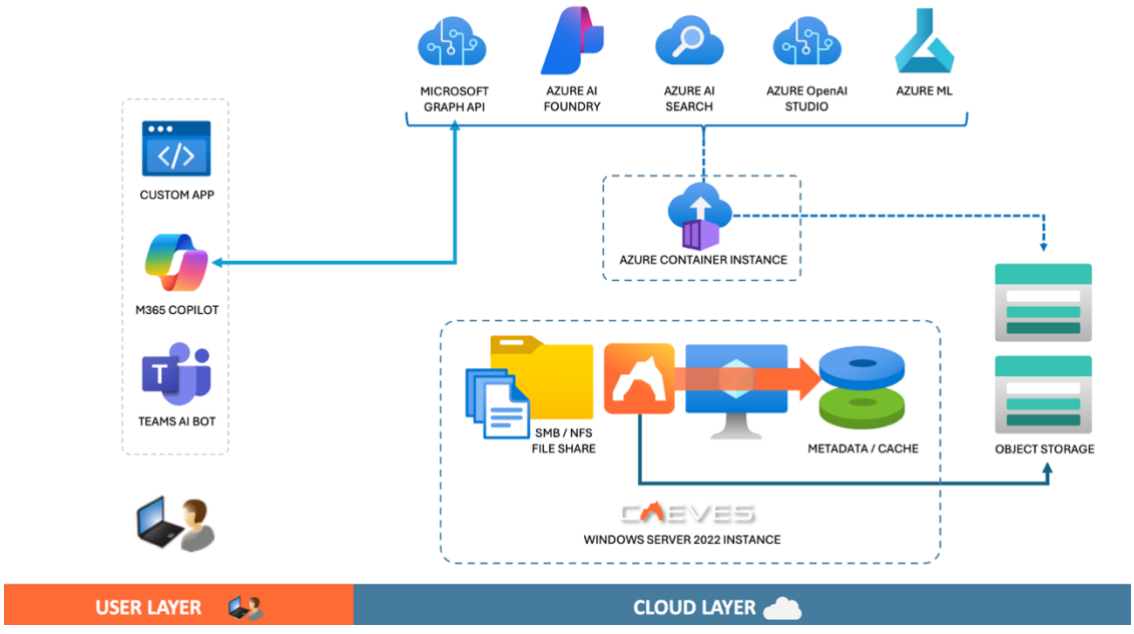

** coming soon*

Platform Features

- Deploy CAEVES in your own hybrid or public Microsoft Azure cloud datacenter
- Built on a Microsoft Azure certified platform with automated provisioning and configuration
- Intelligent Data Tiering with automated hot/cool/cold storage optimization
- Real-time Data Caching for 'Active' data, tiering to Cloud Object Storage
- Multi-protocol File Sharing Support (SMB/NFS*)
- Full support for ACLs & NTFS Permissions with Entra-ID / Active Directory authentication
- Simplified data management with automated snapshots, versioning and backup
- Multi-PB scaling options with predictable performance and cost profile
- Native Microsoft Copilot Connector for RAG-powered search capabilities*

Core Components

1. CAEVES Instance

- CAEVES hardened virtual appliance
- Supported VM sizes: Azure D4s v5 or D8s v5 or equivalent
- Two dedicated disks:
    - Operating System: Windows Server operating system
    - Metadata Disk: CAEVES data disk
        - Metadata Volume: Stores file system metadata
        - Cache Volume: Manages active data set
        - Snapshot Volume: Storage for VSS snapshots

2. CAEVES Copilot Connector

- Deployed as Azure Container App/Instance
- Connects Object Storage to Microsoft AI services
- Enables high-performance content indexing, correlate to file system
- Provides Copilot interface for archive interaction

** coming soon*

CAEVES

System Architecture Diagram (Customer Tenant)

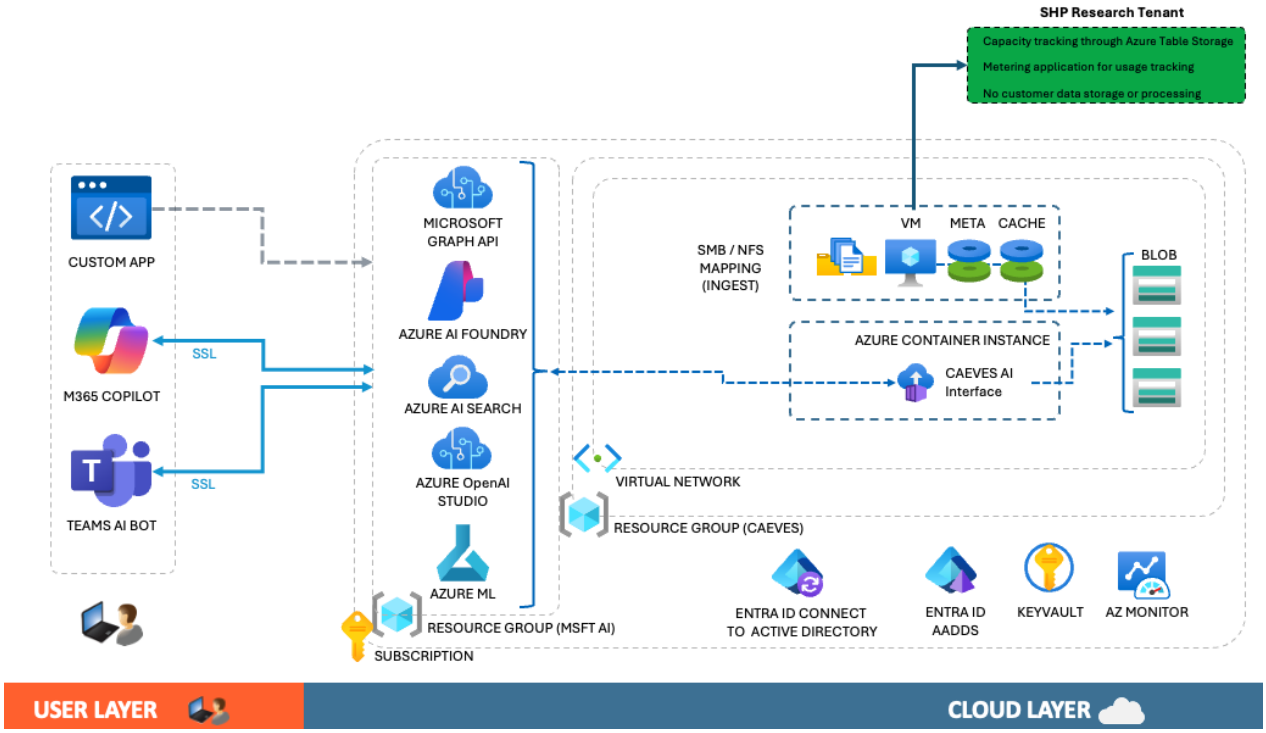

Deployment Model

Prerequisites

- Azure Subscription with Marketplace access
- Sufficient quota for required VM sizes
- Microsoft Entra ID Application Registration for ACI / Object Store / AI services
- Network connectivity requirements:
    - Private Virtual Network (recommended)
    - DNS configuration for Active Directory or Azure AD DS integration
    - Azure Private Link (recommended)
    - Azure Bastion Host (recommended for management access)

Provisioning Process

1. Activate the CAEVES Subscription

Head to the Azure Marketplace and search for the CAEVES platform offer. Choose the appropriate Azure subscription and review the pricing tiers, plans, and EULA.

The pricing is usage-based, measured in terabytes of storage, and your first 5TB are completely free—making it easy to get started without any commitment. Once you accept the offer, you'll be directed to a short onboarding form to configure your account and contact preferences.

2. Provision your CAEVES Instance in your datacenter or in Microsoft Azure

After activation, you'll be guided to automatically deploy the CAEVES Instance using a Virtual Machine template. This process provisions all necessary Azure resources—such as the VM, Azure Table for meta information, and object storage. If you prefer to deploy on-premises, the same CAEVES instance can be provisioned with pre-built deployment scripts and documentation. Either way, setup is fully automated, securely deploys in your datacenter or Microsoft Azure tenant and typically takes less than 30 minutes from start to finish.

3. Migrate your data into CAEVES

Once your CAEVES instance is active, begin migrating data directly into the CAEVES data platform. The system captures and registers metadata into the Azure Table, while the actual content is stored securely in the object store. Whether you're migrating enterprise data or archiving cold data, backups, or entire project folders, CAEVES ensures high availability and WORM-compliant storage for long-term retention and audit-ready compliance.

4. Connect your data with our AI connector (coming soon)

Soon, you'll be able to unlock powerful insights and search across your enterprise data with our CAEVES AI Connector. This connector enables intelligent data discovery, natural language queries, and seamless integration with tools like Microsoft Copilot. Stay tuned—this capability will make your data and archives not just storage, but an intelligent, searchable, AI-ready knowledge base.

Tenants and Security

Customer Tenant Components:

- CAEVES VM instance(s)
- Storage accounts for data persistence
- Container instances for AI integration

CAEVES Technology, Inc. Tenant Components:

- Capacity tracking through Azure Table Storage
- Metering application for usage tracking
- No customer data storage or processing of PII data

Network Architecture

Network Requirements

1. **Internal Network Access**
    - SMB and NFS protocols for file access / data ingest or migration
    - SMB configuration optimized for:
        - Latest protocol versions and dialects
        - RDMA support
        - Multi-channel capabilities
    - Internal communication between VM and Container Apps
2. **Management Access**
    - Azure Bastion Host (preferred)
    - RDP access through private IP (alternative)
    - Future HTTPS access to control plane
3. **External Communication**
    - HTTPS/SSL to SHP Research, Inc. Azure Tables for metering / billing
    - End-user Copilot access via HTTPS to Graph API
    - No direct public internet access required

Security Considerations

- Private Virtual Network deployment
- Azure Private Link for secure service access
- Network segmentation for management and data traffic
- Restricted management ports (RDP/HTTPS)

Integration Points

Microsoft 365 Integration

- Graph API, Azure AI Foundry, Azure AI Search, Azure OpenAI studio, Azure ML
- Copilot integration for natural language interaction
- Content indexing and search capabilities

Azure Services Integration

- Azure Blob Storage for data persistence
- Azure Container Apps for AI services
- Azure Monitor for operational insights

Monitoring and Management

Operational Monitoring

- Azure Monitor integration options
- Log Analytics workspace integration (future)
- Custom metrics for capacity tracking

Capacity Management

- Automated reporting to CAEVES Technology, Inc. tenant (non PII-data)
- Usage metering for billing and reporting purposes
    - Capacity
    - Number of files analyzed
- Capacity optimization recommendations

Future Extensibility

The solution architecture supports future integration with:

- Azure AI Foundation Models
- Azure Fabric
- Additional Microsoft Copilot capabilities

Security and Compliance

Data Security

- All data remains in customer's Azure tenant
- End-to-end encryption for data at rest and in transit
- Integrated Entra-Id / AADDS / Active Directory authentication

Compliance

- Azure Marketplace certification (TBC)
- Microsoft security baseline adherence
- Regular security updates and patches

Sizing and Capacity Planning

Virtual Machine Specifications

Minimum Requirements:

- 4 vCPU / 16GB RAM (e.g., D4s v5)
- Suitable for smaller enterprise data & archive environments
- IO-intensive workload optimized

Recommended Specifications:

- 8 vCPU / 32GB RAM (e.g., D8s v5)
- Ideal for medium to large enterprise data & archive environments
- Enhanced performance for concurrent operations

Note: While these are the recommended VM types, the solution can be adapted to other compute-optimized or general-purpose VM sizes that meet the minimum CPU and memory requirements.

Storage Configuration

Azure Blob Storage:

- Standard General Purpose v2 storage accounts
- Supported access tiers:
    - Hot access tier: For frequently accessed archives
    - Cool access tier: For less frequently accessed data (Default)
    - Cold tier: For inactive data

Local Storage:

- OS Disk: Windows Server operating system
- Data Disk: CAEVES data disk
    - Metadata Volume: Size based on number of objects
        - Files and Folders
    - Cache Volume: Size based on active dataset requirements
        - Number and size of files in local cache
    - Snapshot Volume: Size based on VSS retention policies