Cubbit

Introduction to Cubbit

Mission

Enable The **Next Generation Cloud**

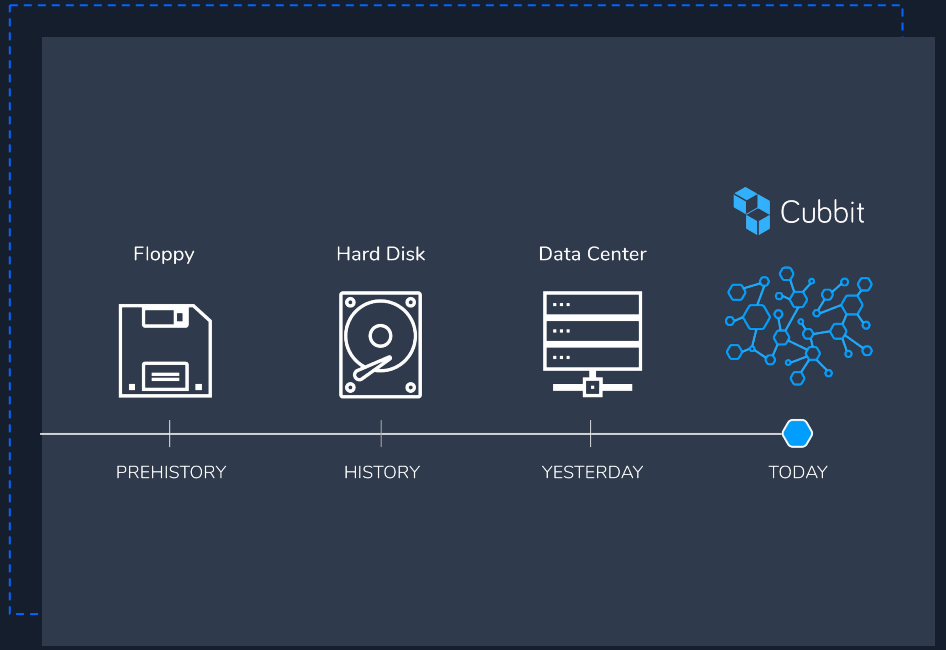

It Is All About Data

Data is the most important asset for every modern organization.
Metadata is key to understand the value of your data.
Machine-generated is much more than human-generated data.

Object storage is the solution but…
Traditional object stores are not optimized for edge

Edge is the where most data is generated
Data accessibility and **mobility** are the key to success

Cubbit DS3 is a decentralized Object Store designed for multi-cloud and edge use cases

How Does it Work?

Cubbit offers a SaaS control plane that collects new and reused resources across the network to exposes them as an Object Store.

Resources are organized in Pee-to-Peer networks (Swarms) and each node of the network can provide both capacity and access via S3 protocol.

The User has a full cloud experience while keeping control over data, infrastructure, and cost.

About Us

Total funding **$12M**
50 employees from 8 countries
Top 30 EU Deeptech Scaleups

120+ Enterprise Customers
5000+ Active Nodes
70 Countries Involved

Stakeholders meeting - March ‘23

Investors: AZIMUT, cdp, PRIMO VENTURES, BARCLAYS, GELLIFY, techstars_, Italian Angels for Growth

Partnerships: Commissione europea, gaia-x, CERN

Why Cubbit?

Digital Sovereignty

Cubbit DS3 is a zero-knowledge system: your data cannot be accessed, even by Cubbit. Data Sovereignty is by design. Cubbit is an European company enforcing ISO and GDPR compliance.

Hyper Resiliency

With Geo-distribution, Cubbit DS3 protects your data from human or natural disasters. With Object lock and versioning, it protects your data from ransomware, Data durability at the highest industry standards 99,999999999%.

Cloud Cost Reduction

Cubbit DS3 helps you save up to 80% compared to AWS or on-prem solutions. Our pricing is transparent, straightforward and predictable. You can set a budget knowing that you are going to stay within it - no hidden costs, no extra charge for multi-region.

Sustainability

Read our Green Paper here

Cubbit DS3 does not require additional, dedicated server farms to provide cloud services. The energy consumed to run and cool off such facilities is therefore saved. For each PB stored on Cubbit, you save up to 25.000 kg of CO_2 every year. It also helps to recycle underutilized hardware.

A Game Changing Paradigm

Secure, Resilient, Cost Effective

Instead of storing data in a single location, we encrypt data, split it in chunks, enforce redundancy and disperse objects in a geo-distributed network where customers have full control and ownership over their data.

Our software can run on any device connected to the internet.
Cubbit DS3 is a **zero-knowledge** system.

Today, **S3-compatible Object Storage**. Tomorrow, **Computing**.

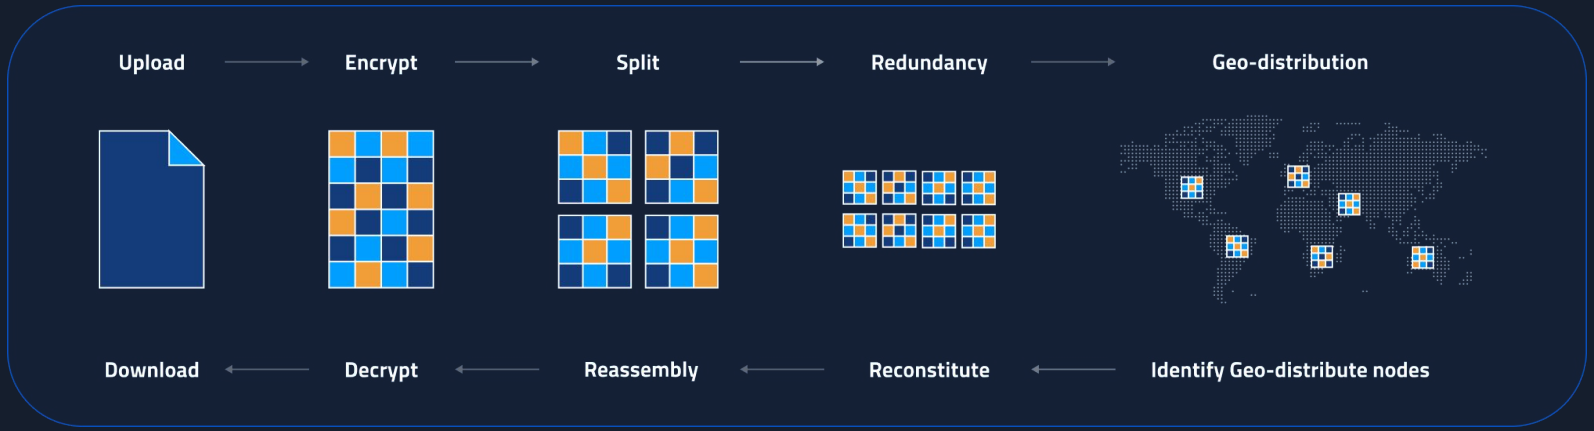

Innovative, yet mature

Our R&D Lab

B2C crowdsourced
5000+ Nodes

1st B2B region

120+ clients
Launched in 2022

Unmatched Security

Zero-Knowledge Approach

Cubbit DS3 is designed with privacy and data sovereignty in mind. To achieve this goal the user is the only one that has access to the data stored in the system and the keys to encrypt and decrypt them.

End-To-End Encryption

Every file is encrypted after ingestion, split in chunks with parity and dispersed in the network. Only the user, who have access to the encryption keys, can access the data.

GDPR and ISO 27001

Cubbit DS3 is ISO 27001 certified and responds to all the necessary requirements for GDPR compliance.

Ransomware Protection

Cubbit DS3 is compatible with Object Lock and Versioning S3 APIs, enabling the user to protect data against any tampering attempt.

Cloud With Data And Infrastructure Sovereignty

Hardware Flexibility

Cubbit DS3 nodes can run on any hardware or virtual machine. The user can decide to reuse existing resources or add new ones, data is distributed across all nodes to preserve the highest service availability and data integrity.

Local And Global S3 Access

Cubbit DS3 offers the possibility to access data through local or public access points depending on business needs and use cases.

Ease Of Use

Cubbit DS3 comprehensive web interface enables users to create and manage an object store with a few clicks.

Multitenancy

Cubbit DS3 is fully multi-tenant. Each tenant can be operated by a different user (or reseller), which has full control on all aspects of the storage resources, configuration, users and policies.

Cubbit Swarm Logical Layout

= GW Node

= St+GW Node

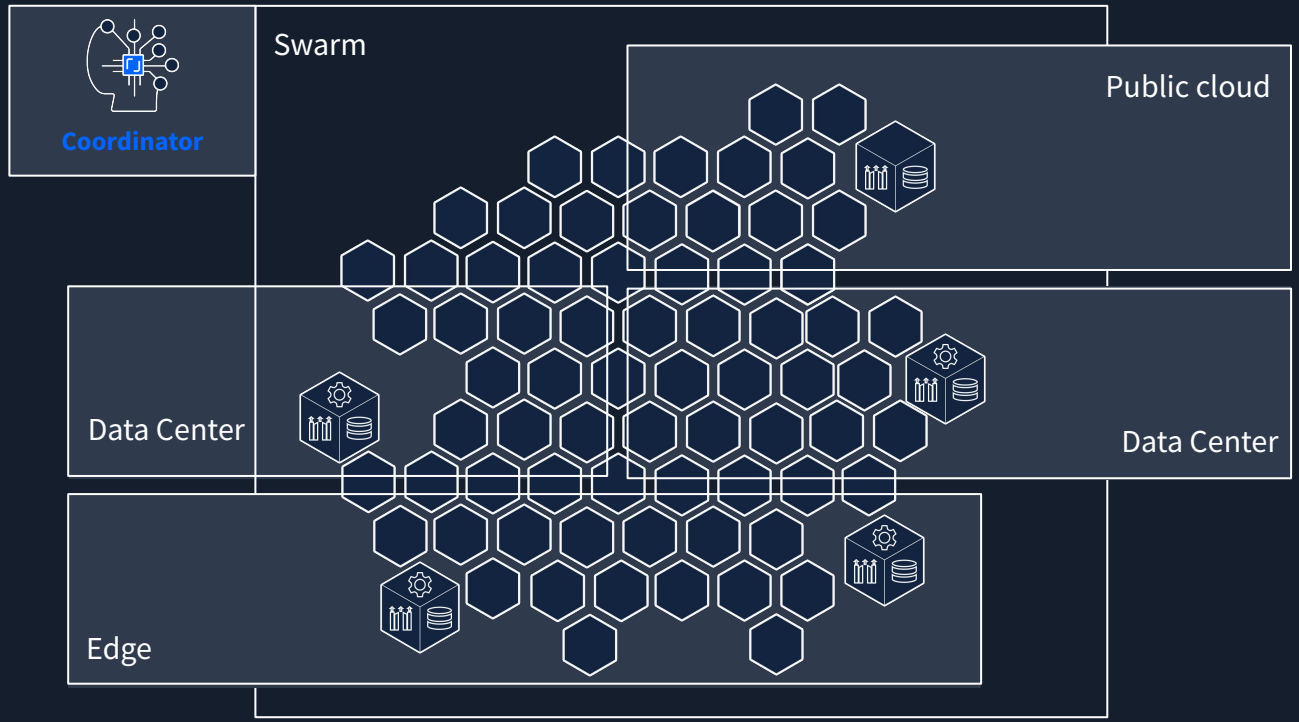

Use Cases

Backup and Recovery Local performance with global resiliency	**Archiving** Low cost and long term retention	**Log Management** Import logs from everywhere and consolidate them for analysis	**Content Delivery** Bandwidth optimization and better user experience
Internet of Things Scalability and single data domain at the edge	**AI/ML** Enable next-gen applications	**Surveillance** Safe data storage for your global security and compliance needs	**Data Analytics** Store at the edge, analyze in the cloud

The Solution For Every IT Strategy

SaaS (managed) or On-prem (hosted) Control Plane

Federated cloud

Multi-tenant, secure and geo-fenced decentralized cloud object store with impressive TCO

Retail
Utilities
Manufacturing
Telco
Channel

Multi cloud

Next-gen data platform that offers a single data domain view across multiple clouds and on-prem

Banking
Finance

Edge

Zero-touch solution enabling performance at the edge while keeping data synchronized

Healthcare
Public Sector
MSPs
Retail
Telco

Private

End-To-End object storage solution for large enterprises

xSPs
Defence
Telco

120+ Customers

Insurance & Finance

Manufacturing

ICT

Food

Defence

Energy

Services

Transport

Public Administration

Cubbit

Technology

Cubbit Technology Components

	Definition
Node 	**A VM or a physical server in the Cubbit network**. Each node **belongs exclusively to a swarm**. The Node provides three services: • **Storage Agent:** responsible for managing disk space, • **S3 Gateway:** S3-compatible access point for the network of the referenced swarm, with a configurable **Cache**. Services can be independently enabled depending on the use case.
Swarm 	**A geo-distributed cluster of three or more Cubbit Nodes**. Each swarm identifies the set of nodes over which the system will distribute payloads related to the uploaded objects. **A swarm is always associated with only one tenant**.
Coordinator 	A resilient, scalable and multi-tenant service that **hosts objects metadata saved in one or more tenants**. Exposes API, CLI and UI dashboard to manage the configurations of the entire environment and its tenants. The coordinator can be managed (SaaS) or hosted (on-prem).

Cubbit Swarm Logical Layout

= GW Node

= St+GW Node

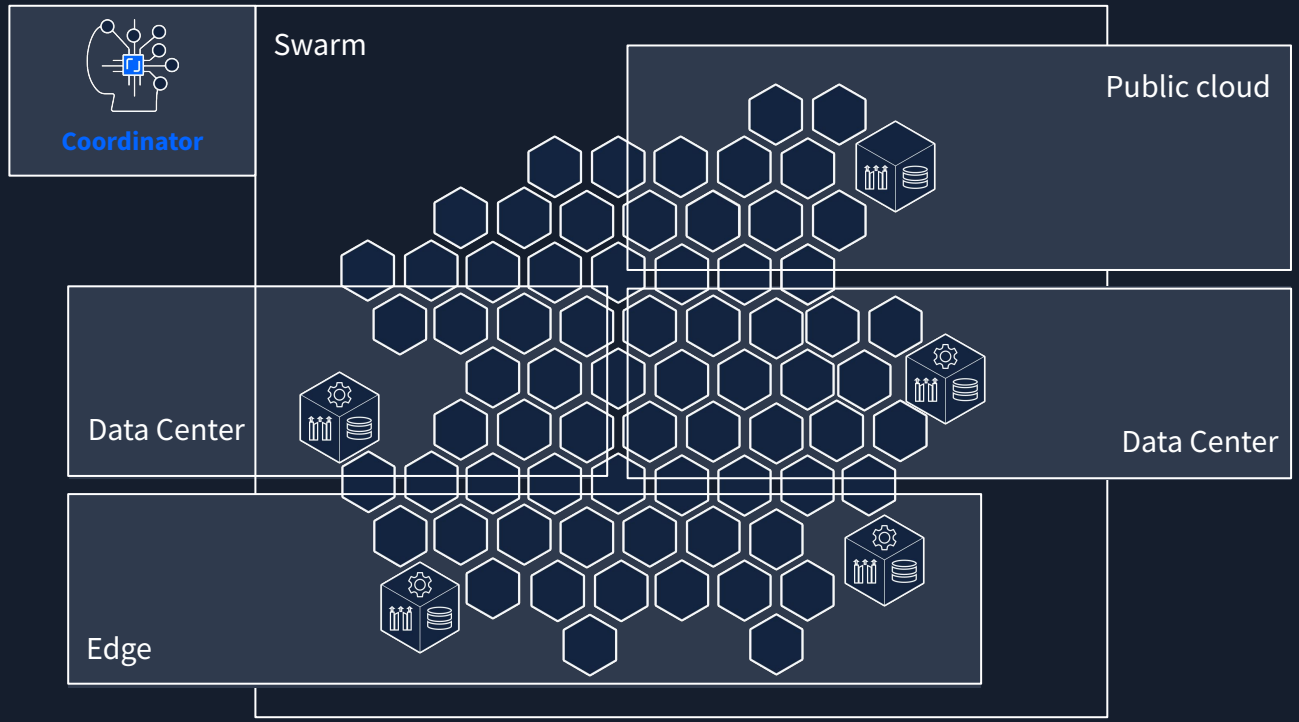

Swarm Data Reliability and Integrity

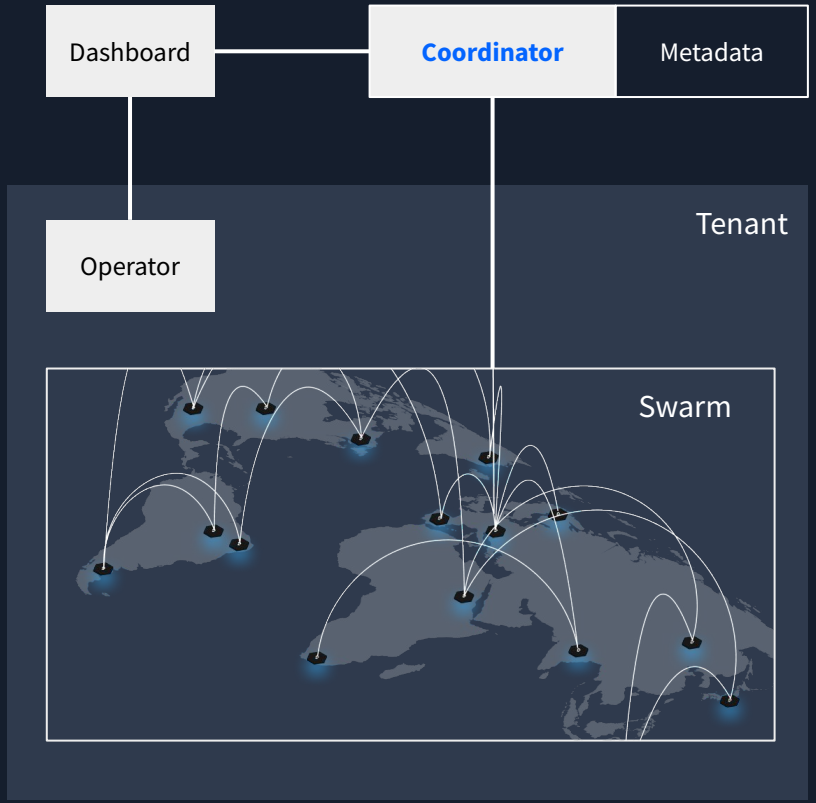

- **Geo-fenced** peer-to-peer network
- **Data is encrypted** in the local gateway, erasure coded and chunks are stored across the network
- Large data protection schemes and automated rebuilding operations ensure **data integrity and resiliency**
- **Smart data placement** algorithm
- Continuous **health monitoring**
- Dynamic **EC policies** allow for efficiency with small files

Swarm Security

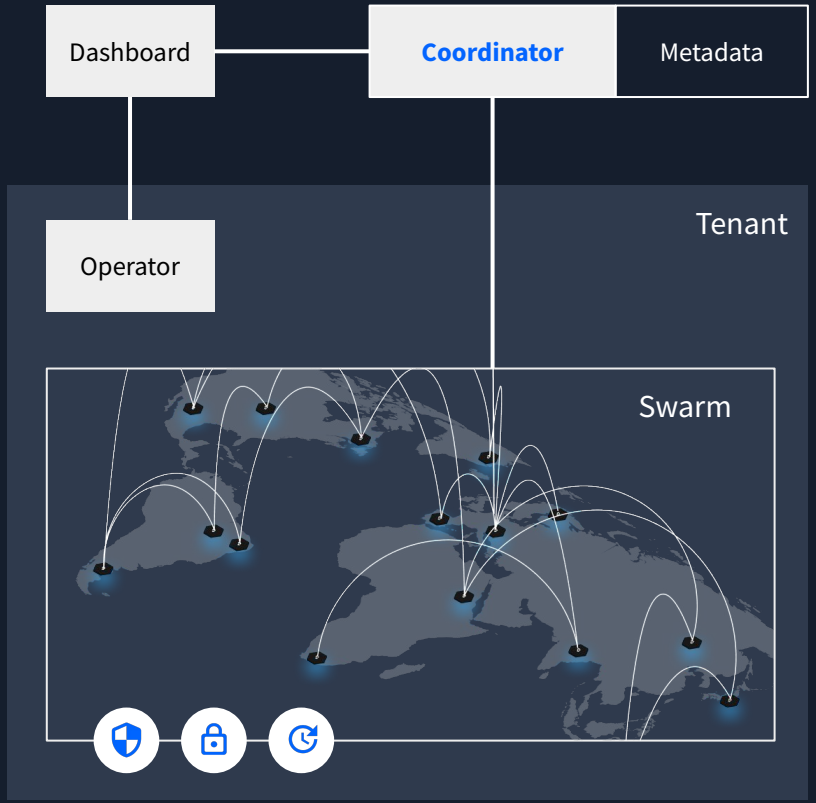

- **Zero-Knowledge** approach
- **End-to-End encryption**
- Support for **external KMS**
- Support for **S3 Versioning**
- Support for **S3 Object locking API**
- **ISO 27001** Certified

Swarm Performance

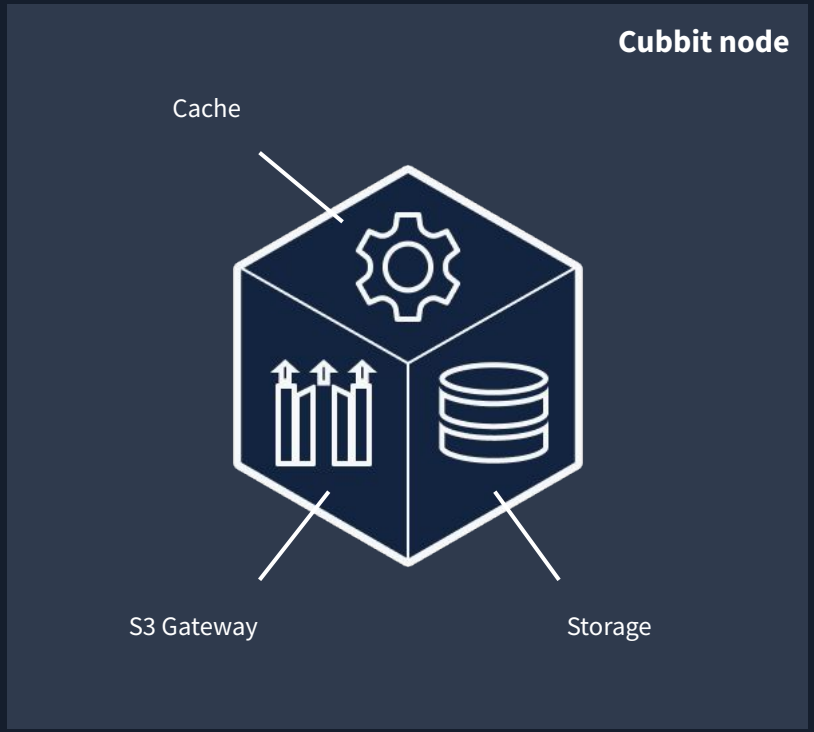

S3 Gateway

- S3 gateways can be local and global
- Data retrieved from most convenient nodes

Cache

- Improves read performance and enables CDN-like experience
- Minimizes write issues with poor network connections
- Provides persistency and performance at the edge
- Supports multiple media types

Example of Cubbit Region

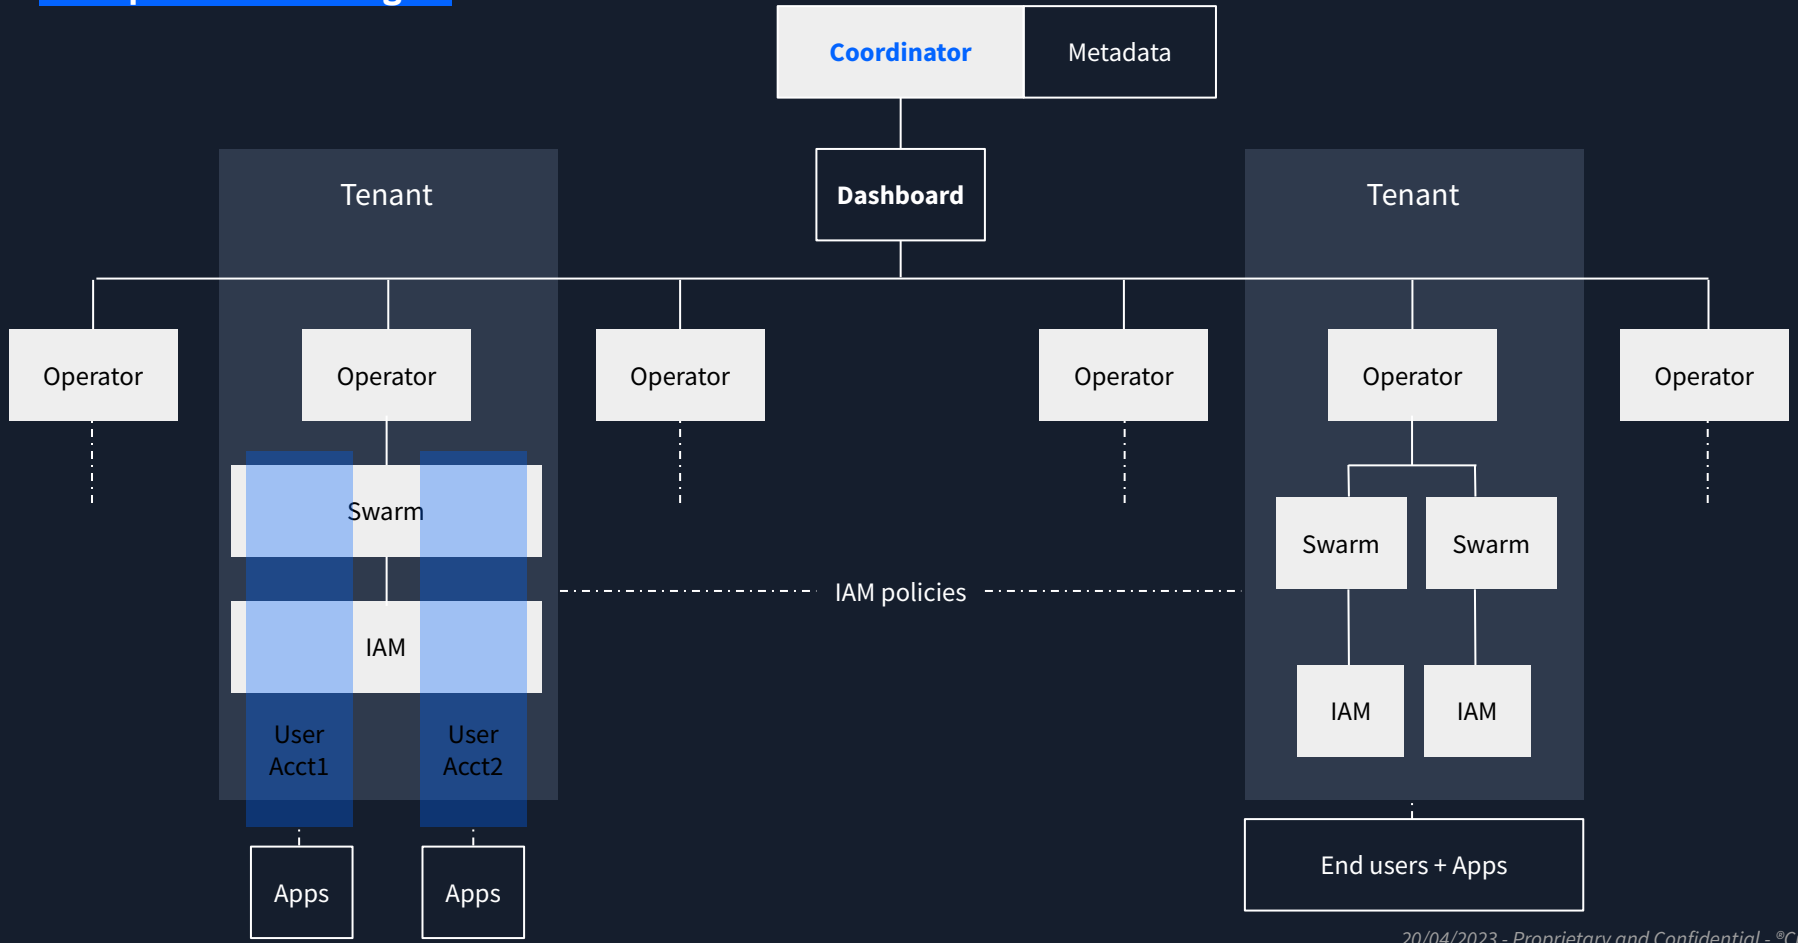

One Object Store, four deployment models

SaaS (managed) or On-prem (hosted) Coordinator

Federated cloud

Multi-tenant, secure and geo-fenced decentralized cloud object store with impressive TCO

Virtual and physical nodes are installed in users' data centers. S3 service can be local or public.

Multi cloud

Next-gen data platform that offers a single data domain view across multiple clouds and on-prem

Virtual Nodes are installed in public clouds and on-prem to provide shared capacity and access.

Edge

Zero-touch solution enabling performance at the edge while keeping data synchronized

Nodes are installed on edge servers and appliances providing access and shared storage capacity.

Private

End-To-End object storage solution for large enterprises

Nodes are installed in one or more data centers. The user keeps control of both nodes and coordinator

Federated cloud

Industry: Telcos, System integrators, regional ISPs, MSPs or ISVs

Need: Want to **provide competitive cloud services, but they can’t or don’t want to afford initial capex investment**.

Current solution: Big and risky CAPEX investment and low ROI, use hyperscalers’ services with low margin and unpredictable costs.

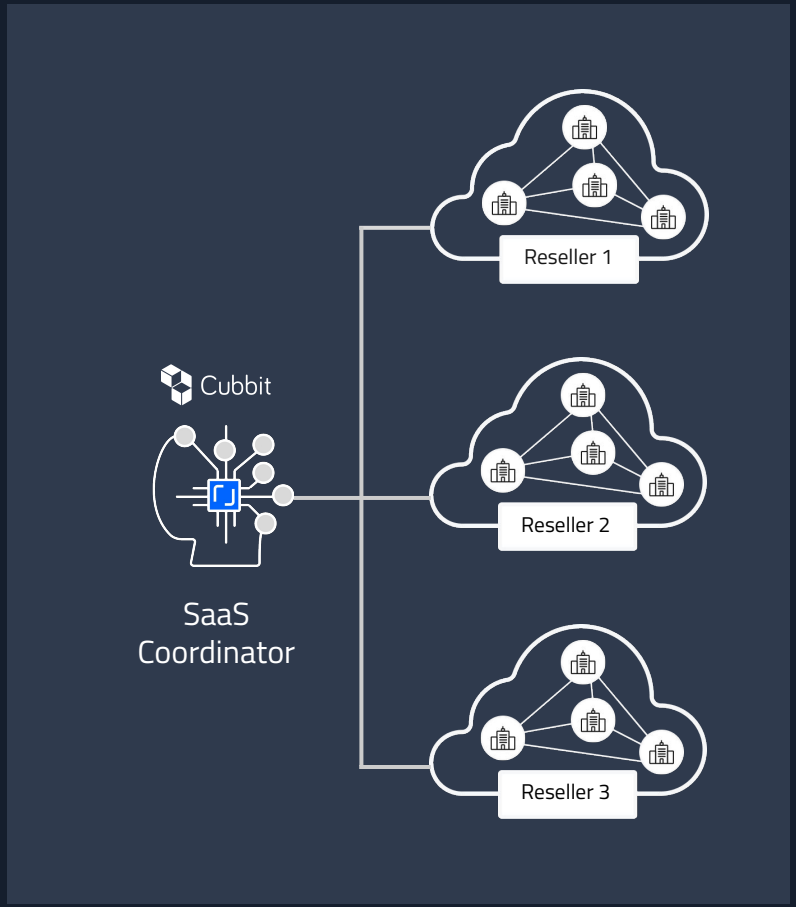

Federated cloud

→ Every reseller can become a cloud storage provider in 3 clicks.

→ One coordinator managed by Cubbit can handle hundreds of swarms at the same time.

Benefits:

- Hyper-resiliency
- Data sovereignty
- Minimal capex investment
- Easy to manage
- Price competitive against hyperscalers
- Re-use of wasted storage resources

Use-case: Backup, disaster recovery, log ingestion.

Private

Industry: **Banking, defence, healthcare, public sector**

Need: Existing on-premise solutions are difficult to manage and not flexible.

Current solution: Traditional object storage solutions

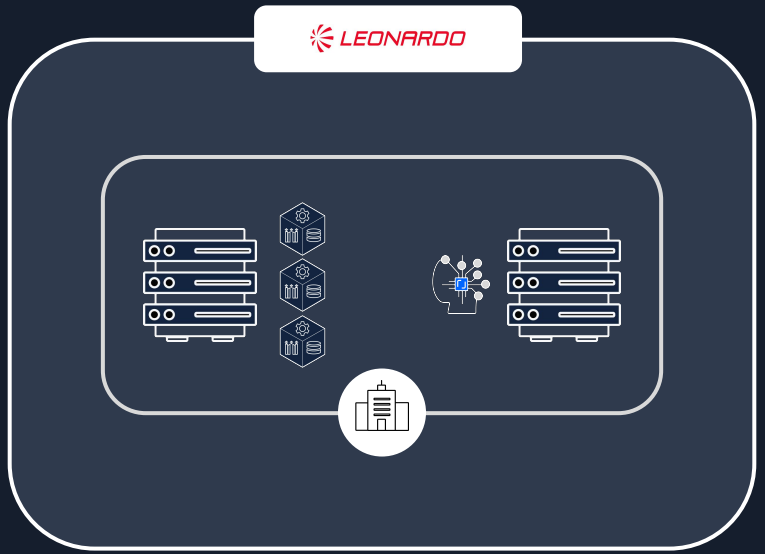

Private

→ **Deploy Cubbit's full architecture in a datacenter** to achieve better infrastructure efficiency while improving reliability and control.

Benefits:

- Low cost
- Improved efficiency
- Local performance
- Disaster recovery
- End-to-end solution
- Re-use of wasted resources

Use-case: Backup, secondary storage, archiving, log ingestion

Multi cloud

Industry: media and entertainment, manufacturing, retail, finance and banking.

Need: Users and applications **need to access the same data from multiple cloud providers and on-prem** at competitive prices.

Current solution: Expensive data replication and data management (compliance/governance) across clouds.

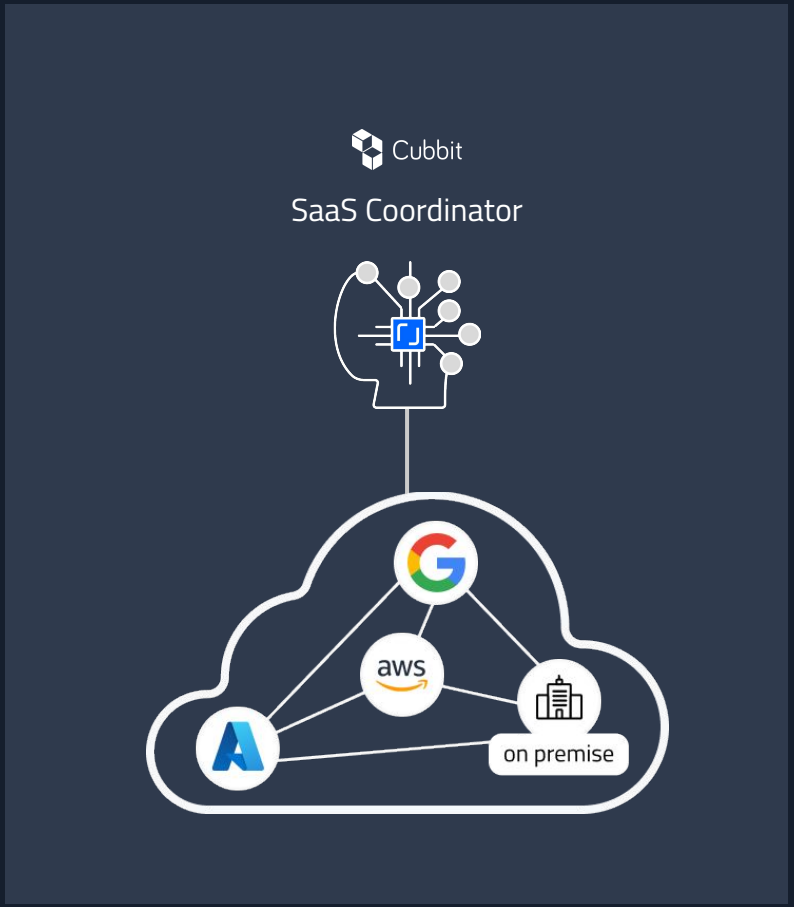

Multi cloud

→ **Install a Cubbit node with the gateway enabled in every cloud** and **access the same data simultaneously**.

→ Swarm can spread across multiple environments.

Benefits:

- **No lock-in**
- Single data domain
- Performance
- Cost savings (due to reduced data movement and no replication)
- Agility

Use-case: data analysis, data repatriation and migrations, disaster recovery, cloud cost optimizations, next-gen applications.

Edge

Industry: manufacturing, healthcare, retail, utilities, public sector.

Need: **Data is generated in decentralized locations, might be used locally and might need replicated on the cloud** for analysis. Decentralized locations need persistence, single data domain, and reliable storage.

Current solution: Local storage in remote nodes and traditional data replication methods, complex application stack.

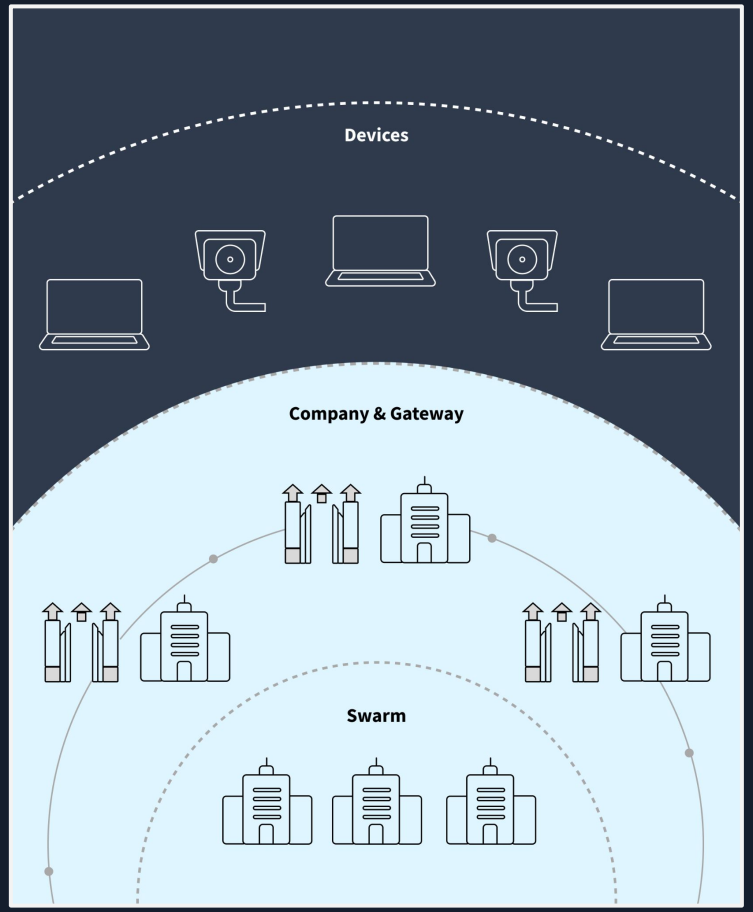

Edge

→ **Install Cubbit Nodes in remote locations**.

→ The S3 gateway provides persistent storage, cache provides performance. Swarm provides reliability.

Benefits:

- Single data domain
- Local performance
- Disaster recovery
- Automatic data sync
- Better infrastructure efficiency

Use-case: IoT, data collection, 4.0 industries, big data analytics

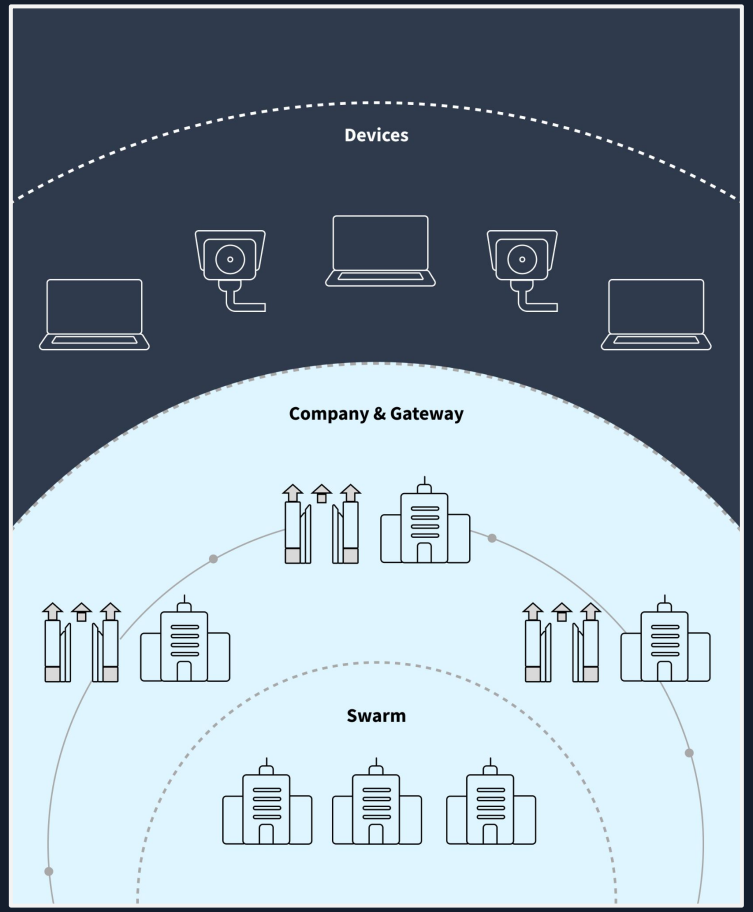

Example

→ **Challenge:** Retail chain with 800+ stores. 2 servers in each store (HA configuration with shared storage). 2.4PB of unused space.

→ **Solution:** SaaS coordinator with a single Swarm made out of VMs running on the existing infrastructure. The S3 gateway on the nodes provides resiliency and global data access. Cache provides CDN-like experience.

→ **Results:** The S3 gateway receives backup data and stores it automatically on the Swarm for DR. Promotional and training videos are cached on the local nodes for faster access and limited bandwidth consumption. No additional hardware required.

Benefits

	On-Prem	Cloud	Cubbit
Secure			
Scalable			
Easy to Adopt and Grow			
Low $/GB and Performance			
Simple and Robust			
Flexible			
No Lock-in			
Covers Edge-Core-Cloud			

What's next?

Serverless computing

→ Provide developers with **tools to develop Lambda Functions at the edge**

→ Small code snippets triggered by events process data locally before synchronizing it to a swarm

Benefits:

- Improved ETL
- AI inferencing
- Data and metadata augmentation
- Data classification at the edge
- Improved edge-cloud efficiency

Use-case: IoT, data collection, 4.0 industries, big data analytics