HAMMERSPACE

Data-as-a-Service

David Flynn – CEO

Company Overview

HAMMERSPACE

- David Flynn, Founder and CEO (Fusion-io)
- Headquartered in Silicon Valley
- Cloud-native

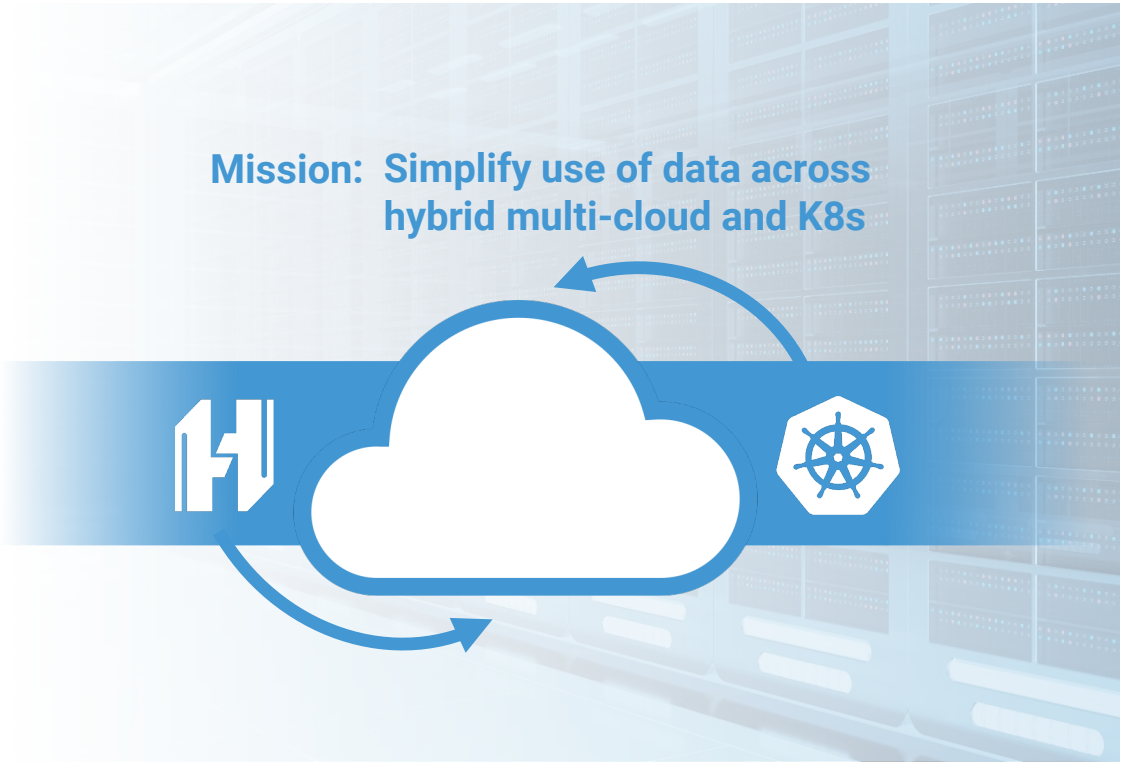

Today data is accessed and management through the infrastructure

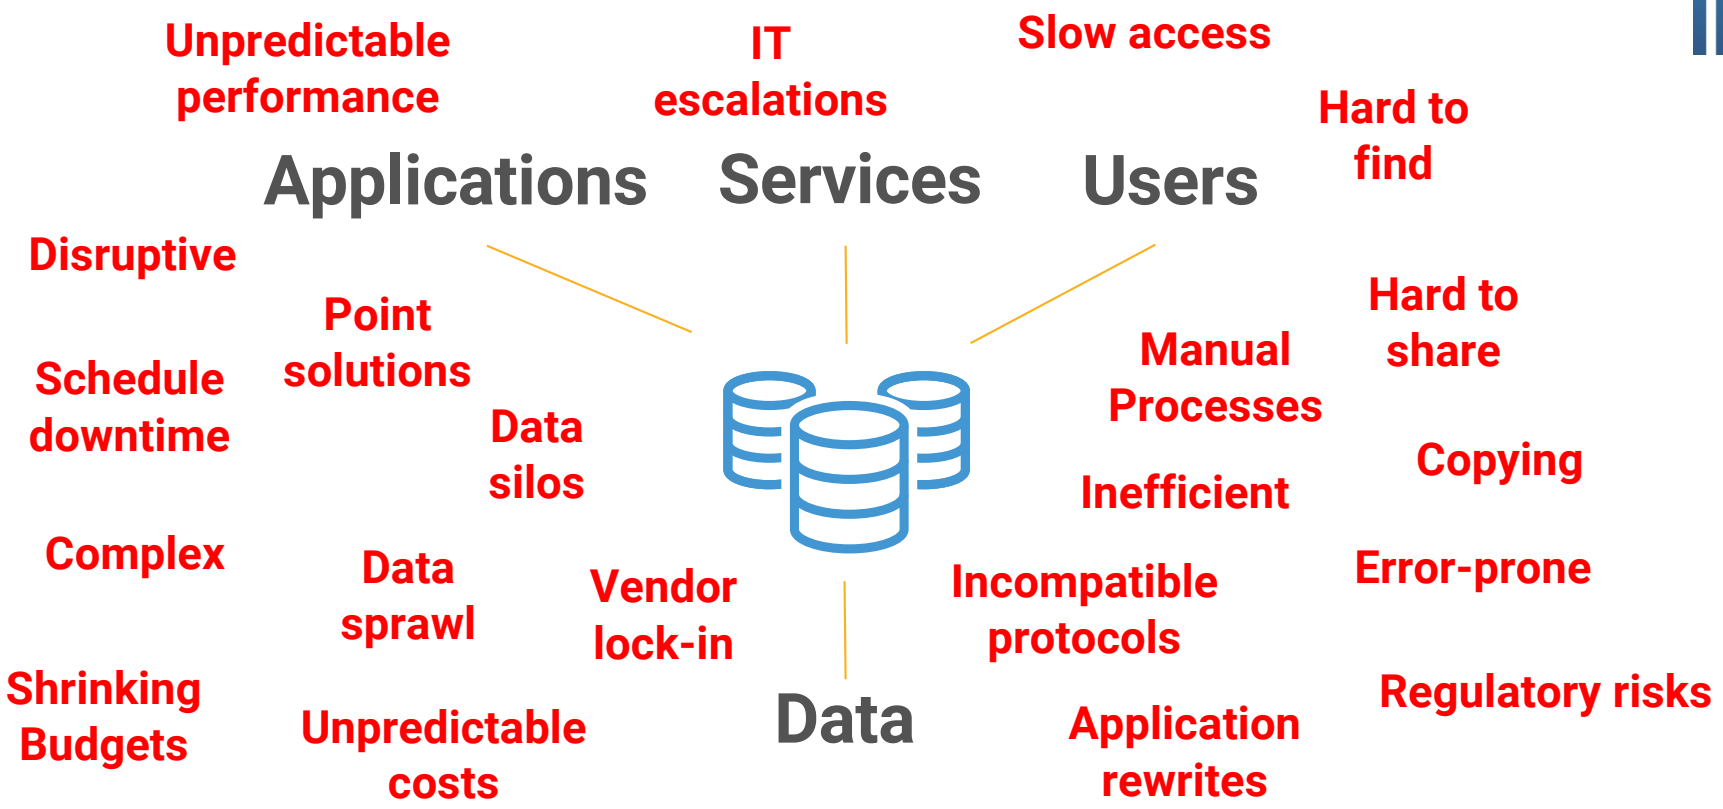

Infrastructure Is Getting More Complex Over Time

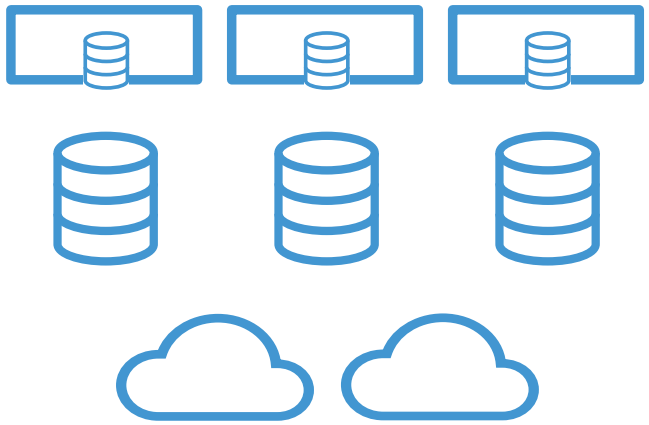

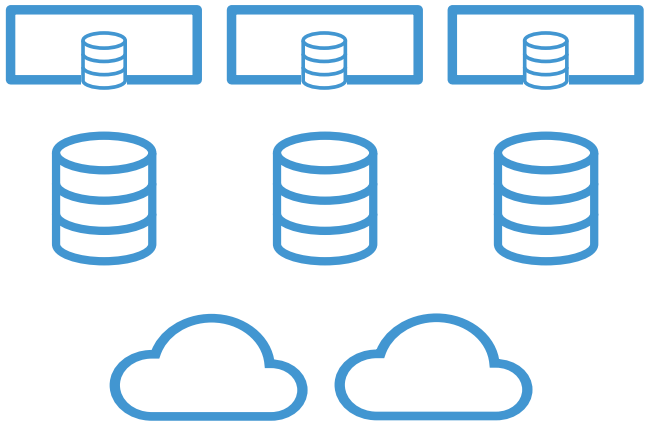

- Different design points
  - Performance, capacity cost and scale
- Different physical locations
  - across servers, across racks, across data centers

Point Solutions Stretch Data Between Fragmented Silos

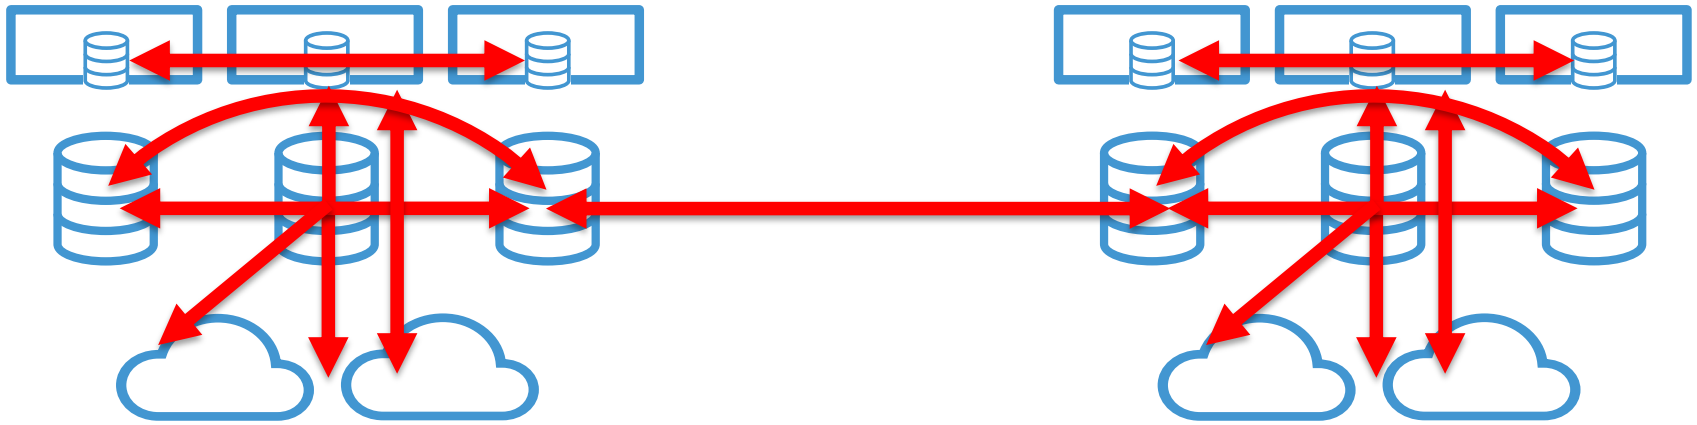

- Caching (Avere)
- Tiering (Komprise)
- Backup (Commvault, Rubrik, Cohesity)
- Gateway (Panzura, Nasuni)
- Scale-out (Isilon, Qumulo)
- Virtual SAN (VSAN, ScaleIO, Nutanix)
- Global namespace (Acopia, Alluxio, K-Mesh)
- Replication (SnapMirror, CloudSync, Aspera)

Tools for coping with the resulting mess

- Data catalogs (Collibra, IBM)
- Copy data management (Delphix, Actifio)
- DAMs MAMs (Bynder, MediaVault, Canto)

Like Managing A Library Through The Bookshelves

What's Needed Is Metadata

And Autonomic Data Services

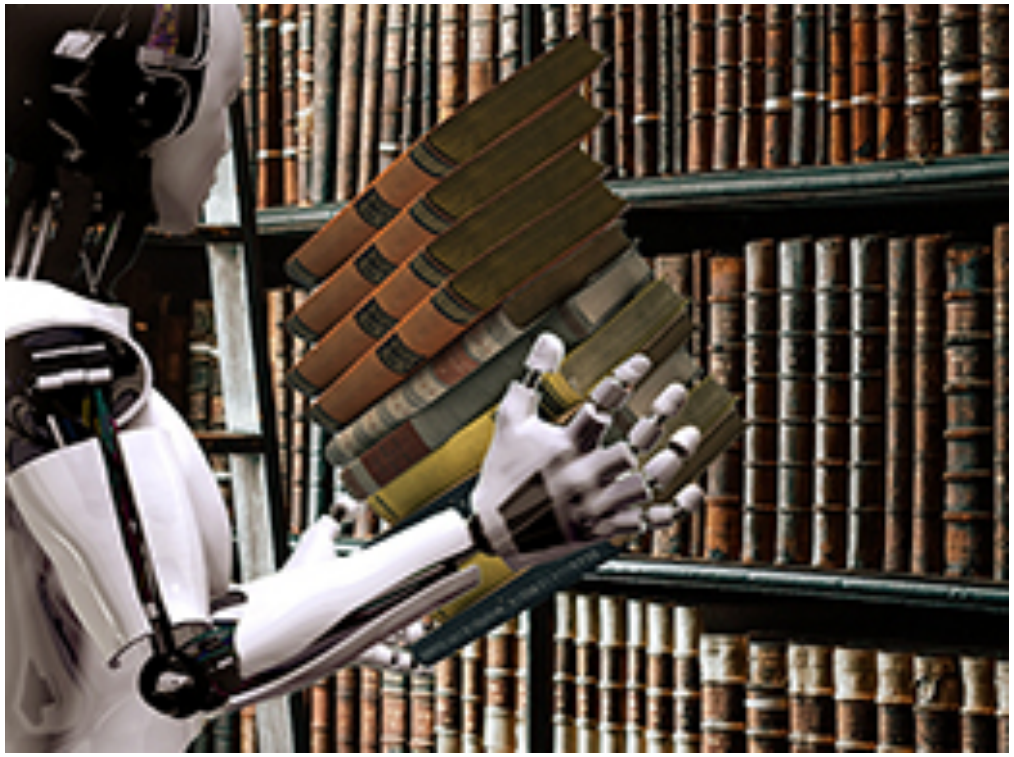

Block – File Services In The Client

- Geo shared
- Ultra-capacity
- Ultra-performance
- Metadata performance
- Declarative orchestration
- Heterogeneous environments

NAS – File Services In The Array

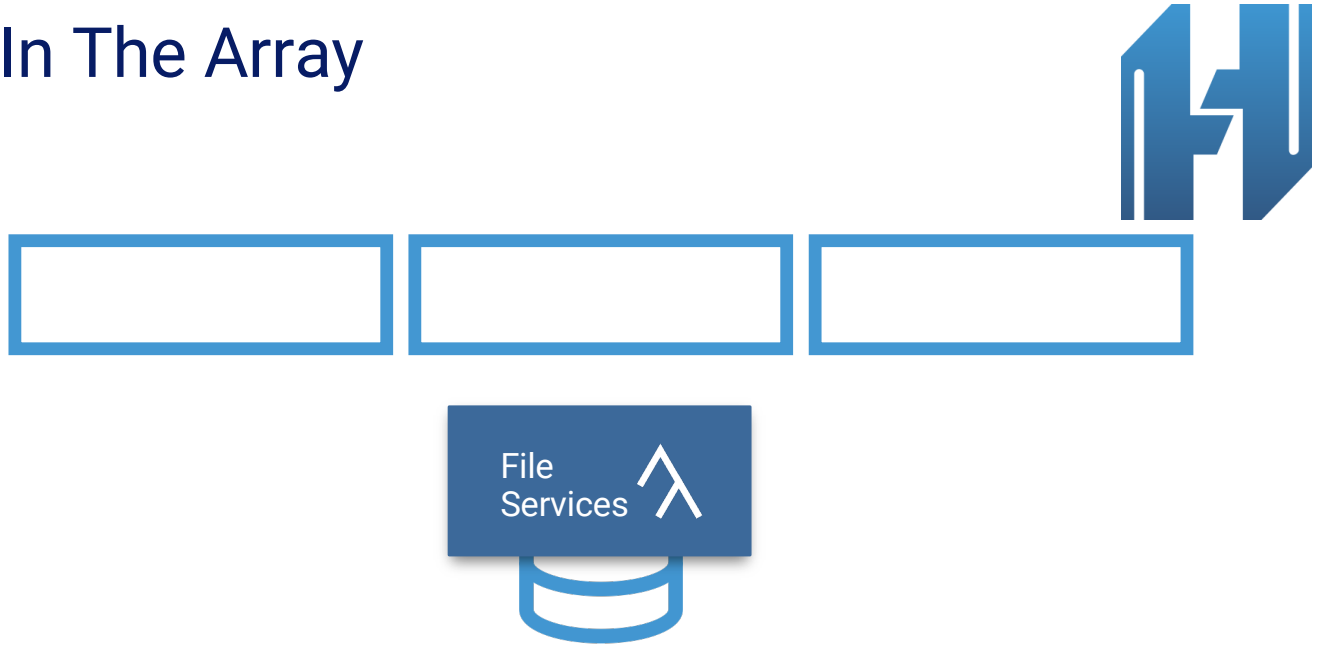

- Geo shared
- Ultra-capacity
- Ultra-performance
- Metadata performance
- Declarative orchestration
- Heterogeneous environments

Scale-out NAS – File Services Stretched

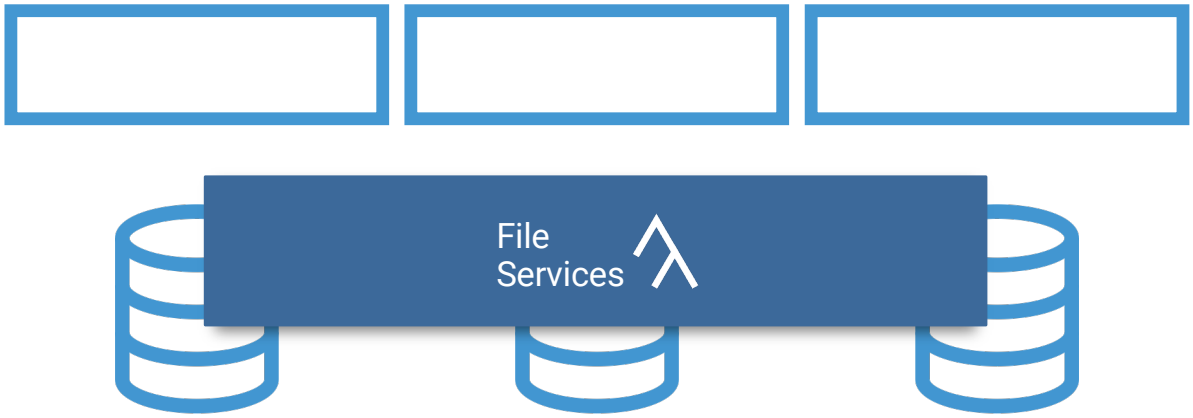

- Geo shared
- Ultra-capacity
- Ultra-performance
- Metadata performance
- Declarative orchestration
- Heterogeneous environments

Hammerspace: File Services As a True Service

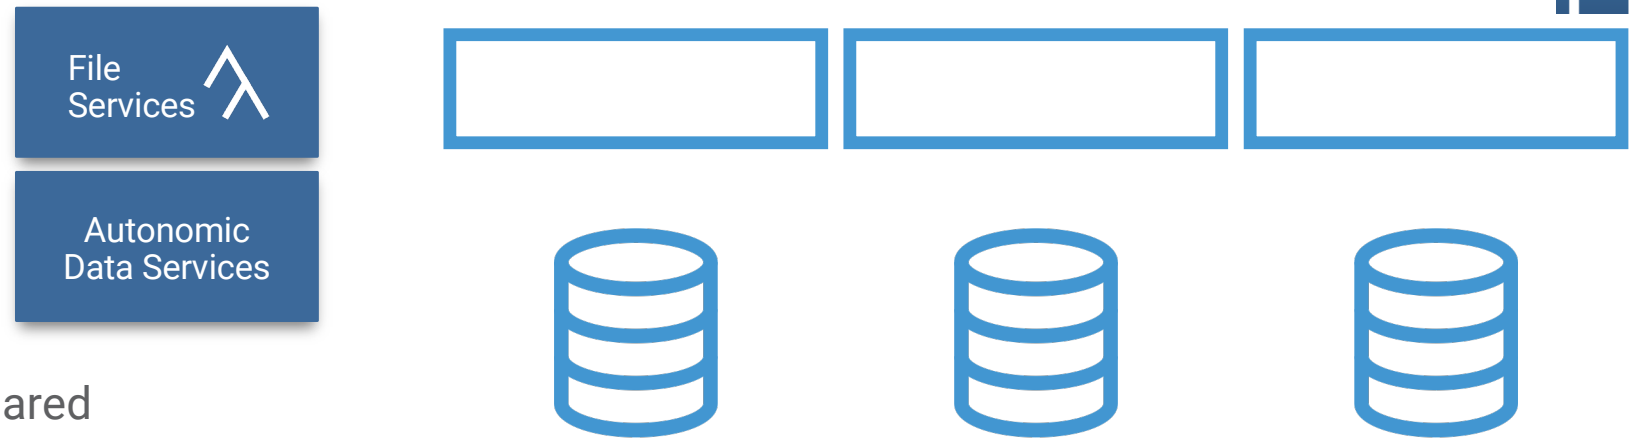

- Geo shared
- Ultra-capacity
- Ultra-performance
- Metadata performance
- Declarative orchestration
- Heterogeneous environments

Hammerspace – Ultra-Performance Local NVMe

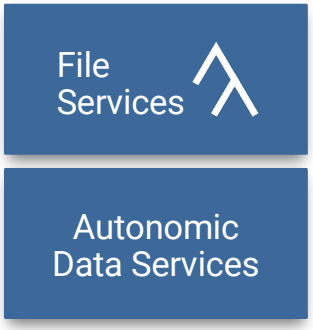

- Geo shared
- Ultra-capacity
- Ultra-performance
- Metadata performance
- Declarative orchestration
- Heterogeneous environments

Hammerspace – Ultra-Capacity Object Storage

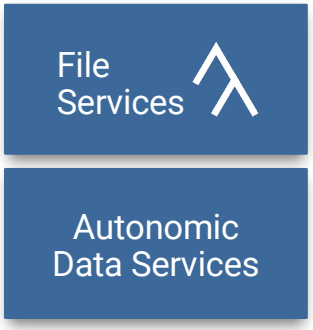

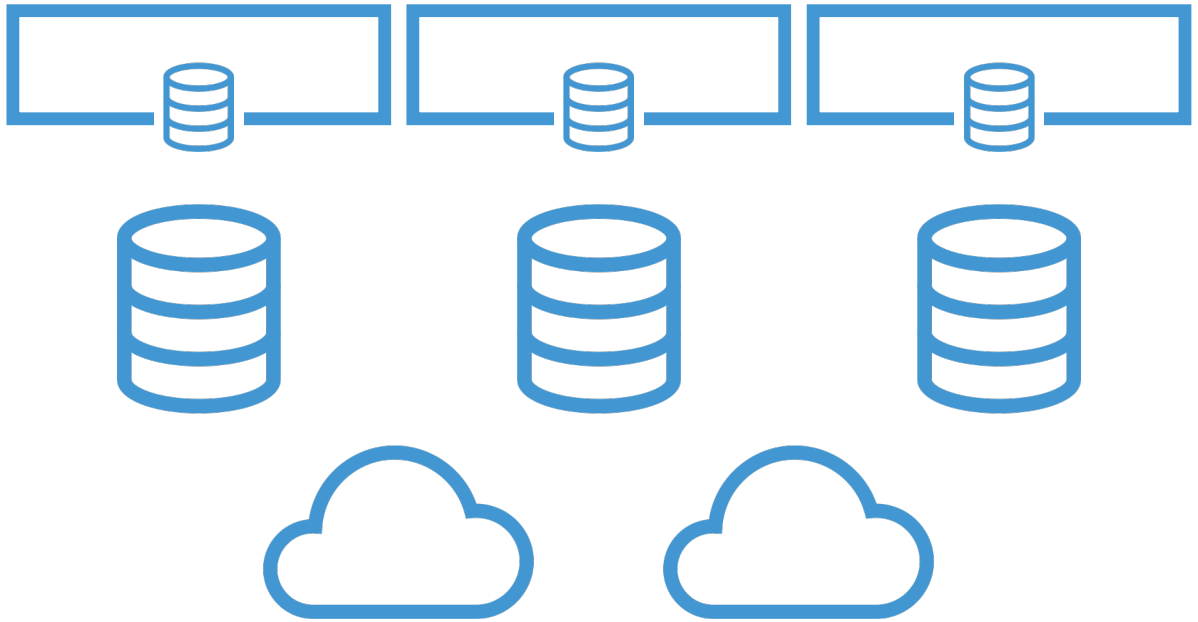

- Geo shared
- Ultra-capacity
- Ultra-performance
- Metadata performance
- Declarative orchestration
- Heterogeneous environments

Hammerspace – Multi-Site Geo Replication

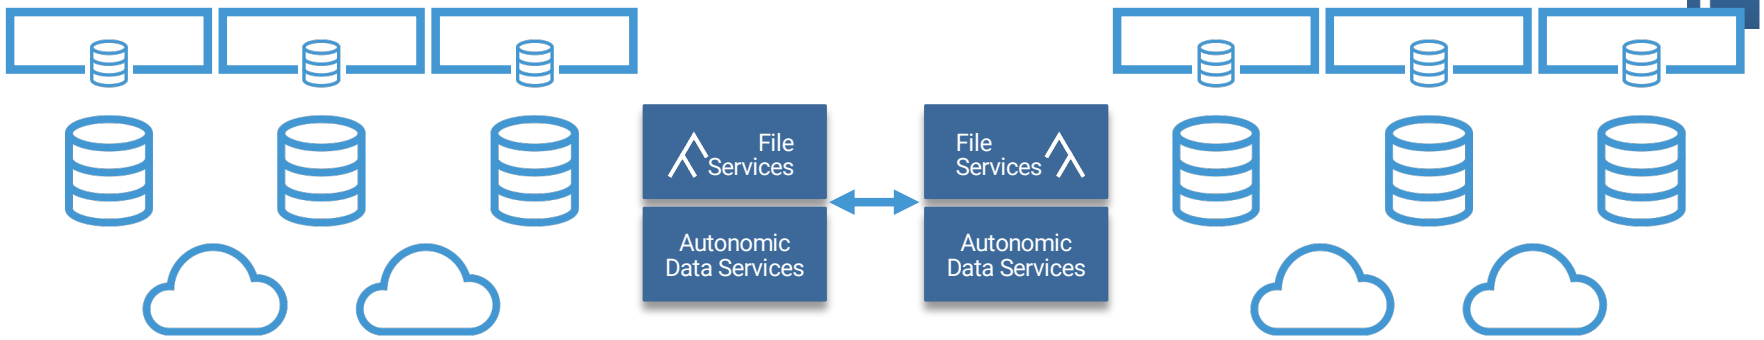

- Geo shared
- Ultra-capacity
- Ultra-performance
- Metadata performance
- Declarative orchestration
- Heterogeneous environments

Geo shared

- Mass (inertia, gravity)
- Space
- Time

Unifies untra-large and distant (General Relativity) with ultra-fast, close and small (Quantum Mechanics)

Today data is accessed and management through the infrastructure

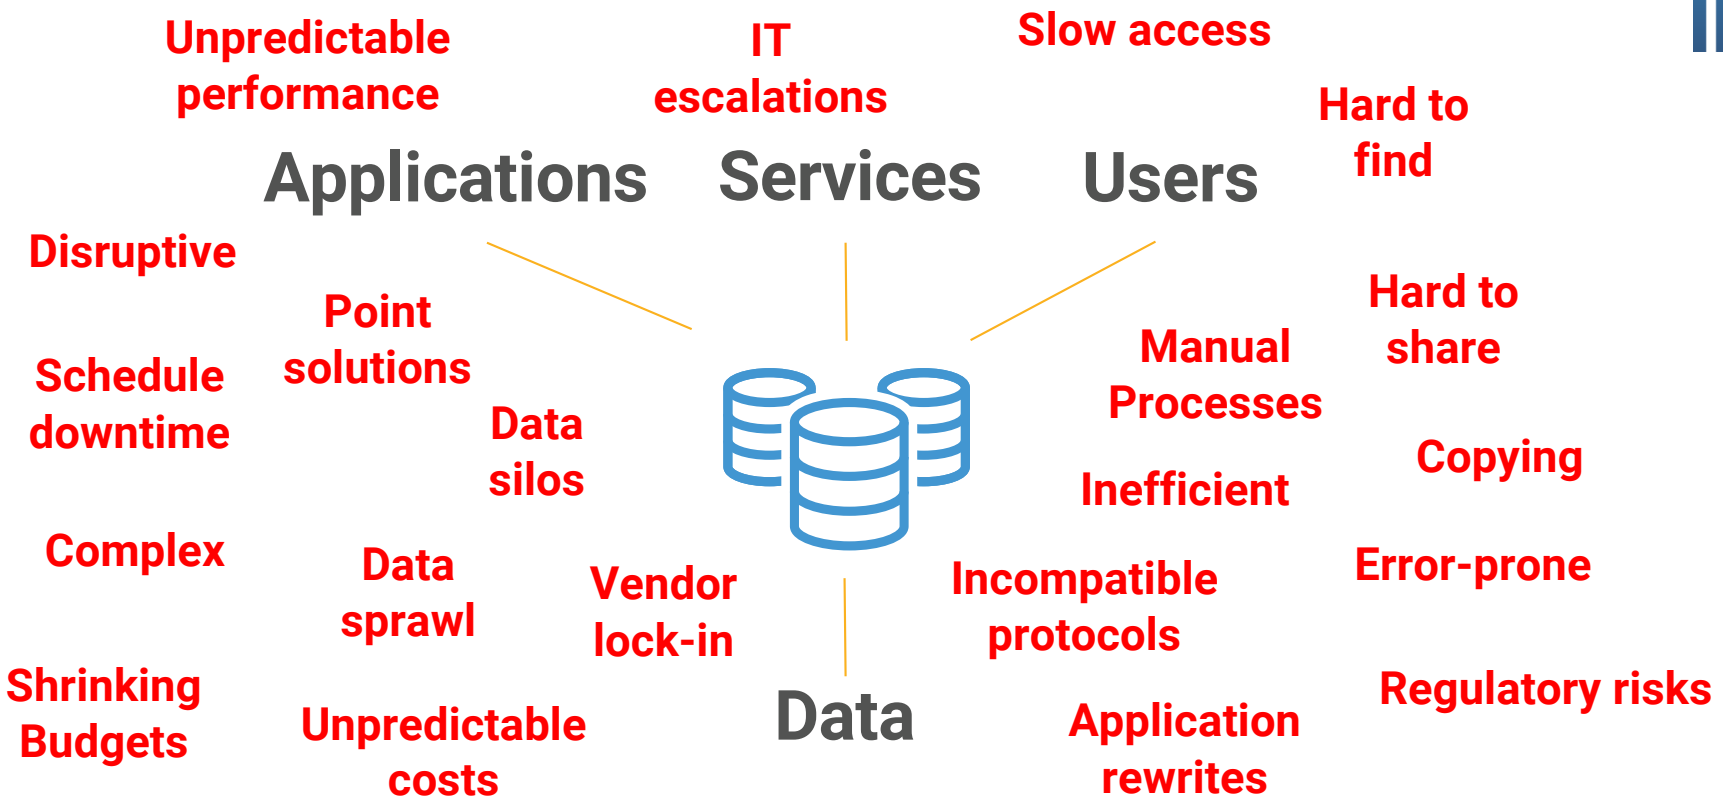

Data-as-a-Service data is accessed and managed through metadata

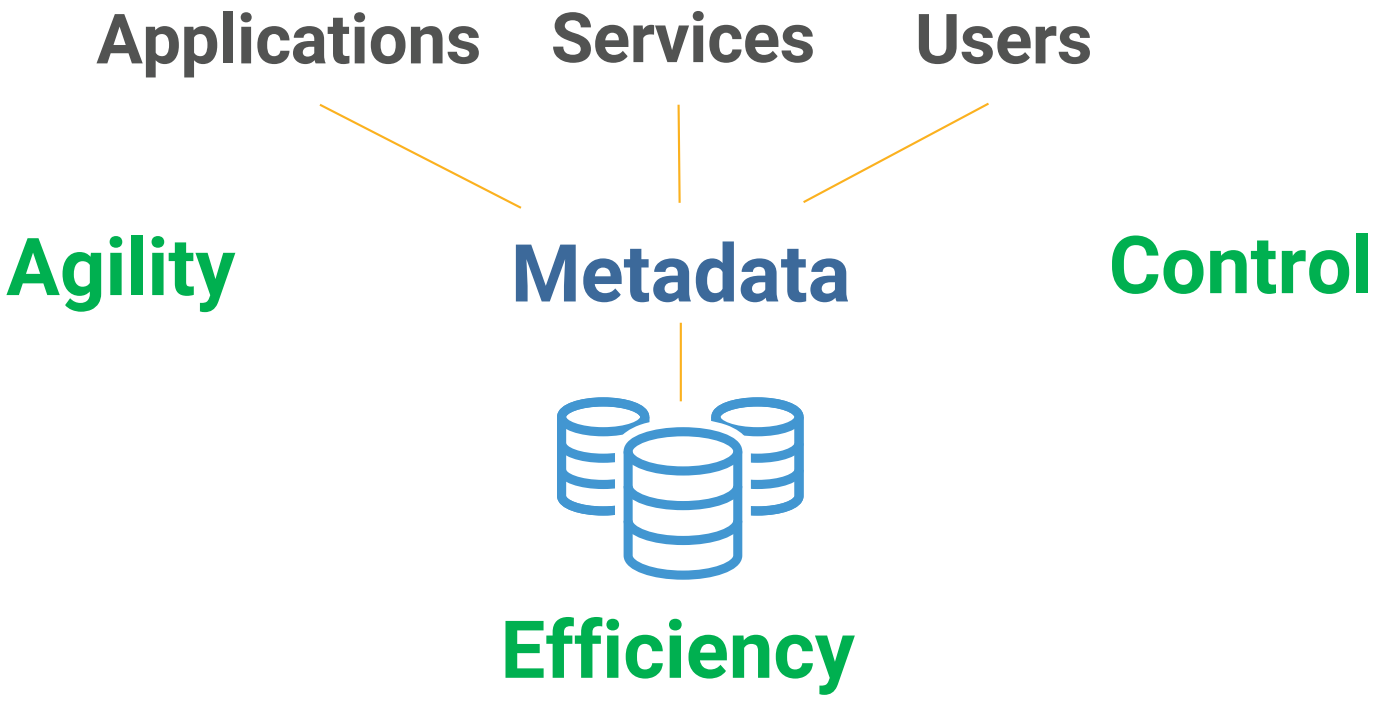

Data-as-a-Services meets an organizations conflicting needs

Agility

Data Consumers

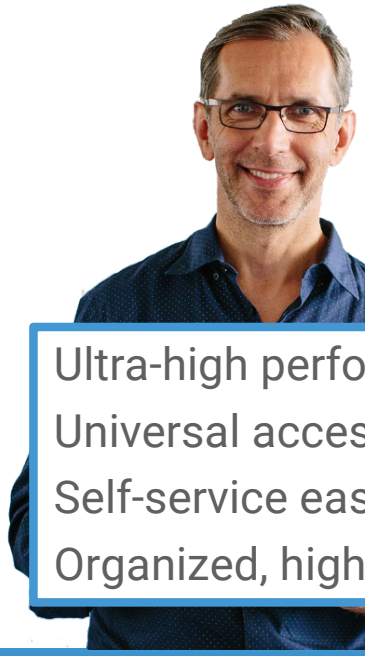

Accelerate time to value, maximize utility

- Ultra-high performance
- Universal access
- Self-service easy to use
- Organized, high quality

Control

Data Governance

Reduce risk

- Protection, recovery
- Security, governance
- Sovereignty, compliance
- Retention, provenance

Efficiency

Infrastructure Operators

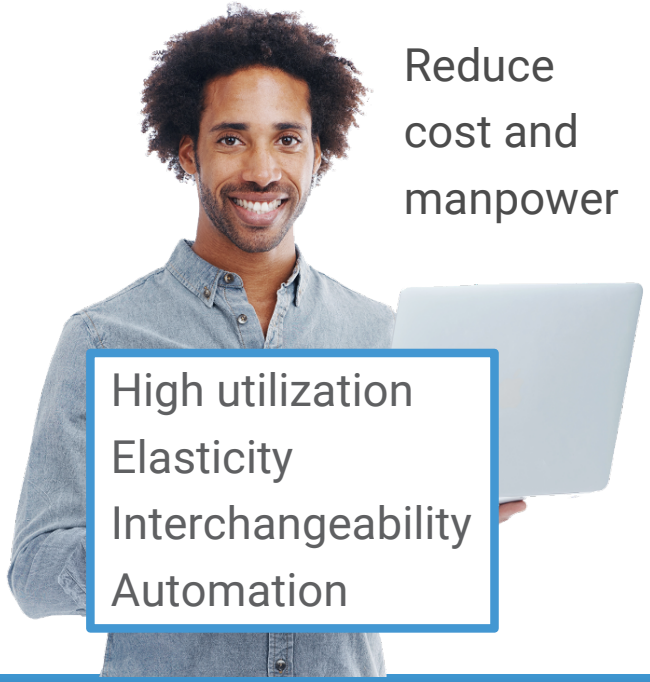

Reduce cost and manpower

- High utilization
- Elasticity
- Interchangeability
- Automation

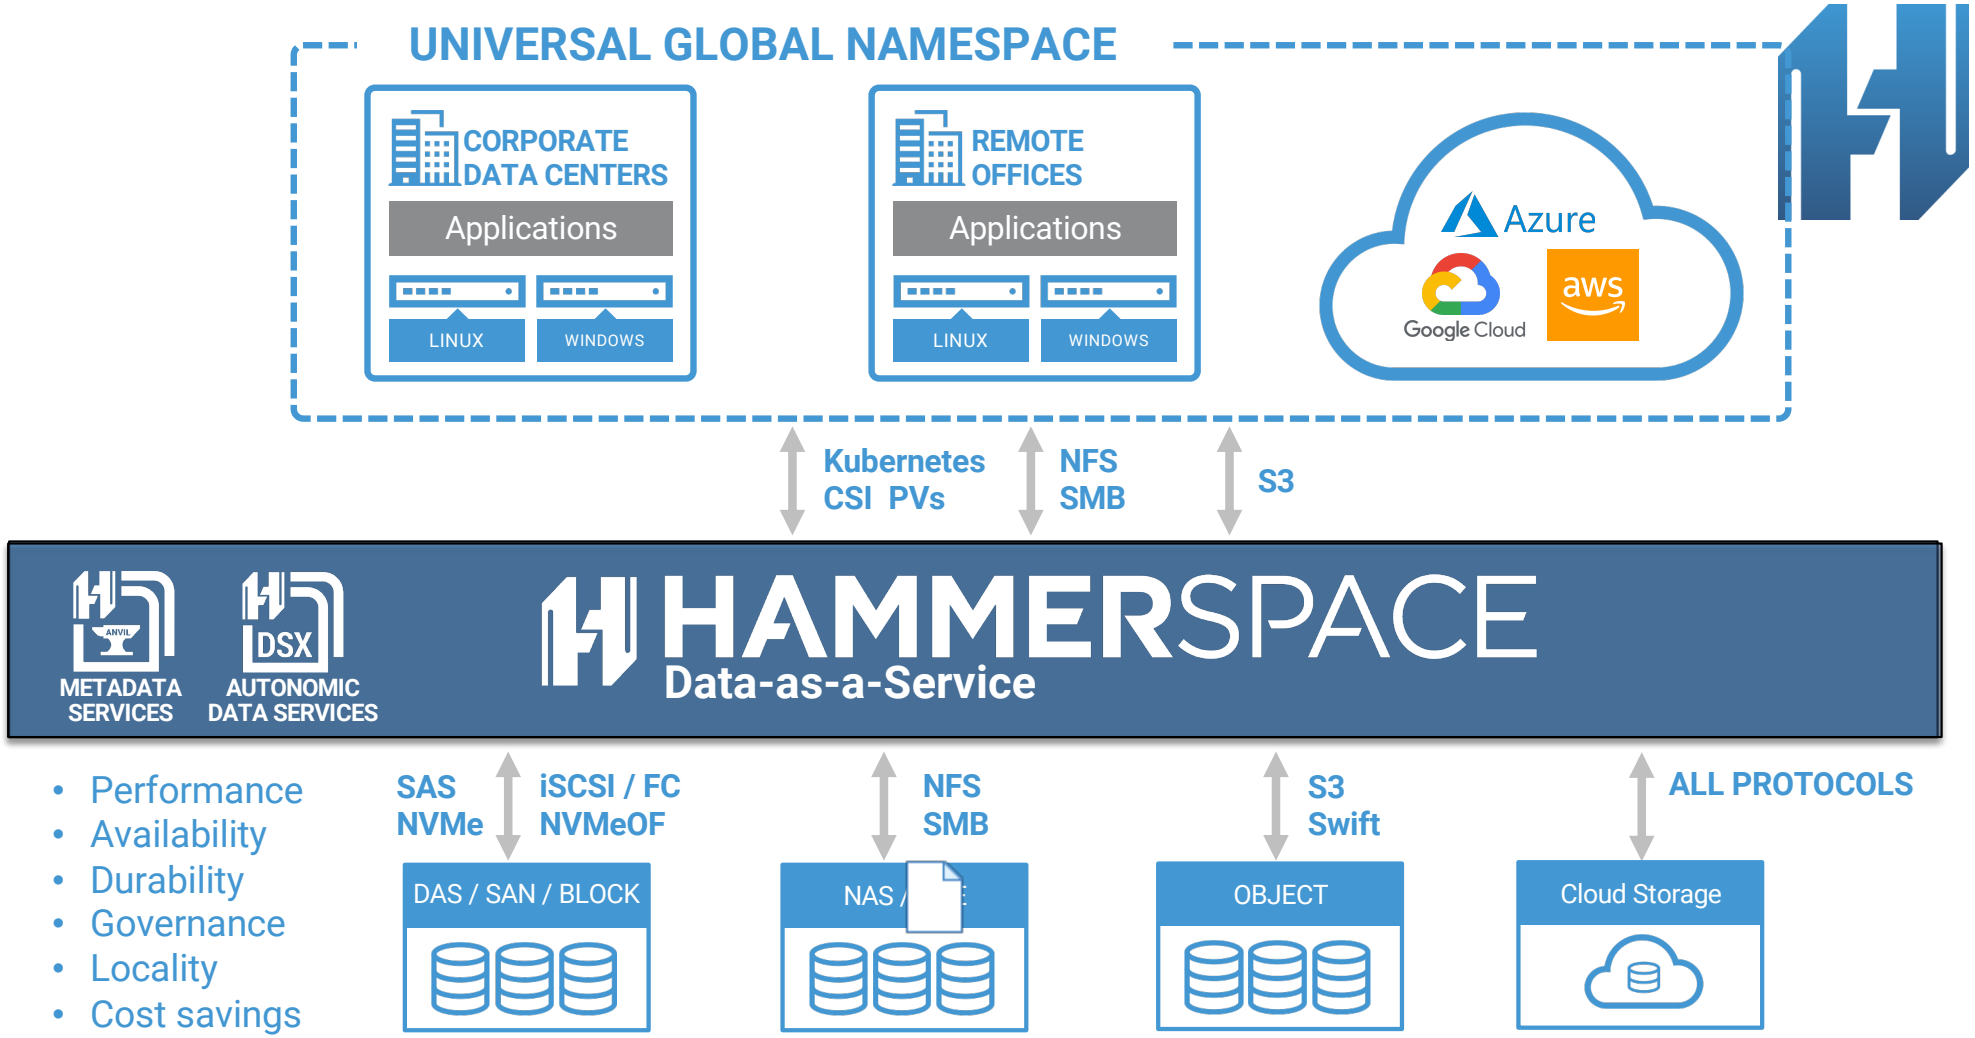
UNIVERSAL GLOBAL NAMESPACE
CORPORATE DATA CENTERS
Applications
LINUX
WINDOWS
REMOTE OFFICES
Applications
LINUX
WINDOWS
Azure
Google Cloud
aws
Kubernetes CSI PVs
NFS SMB
S3
METADATA SERVICES
ANVIL
AUTONOMIC DATA SERVICES
DSX
HAMMERSPACE
Data-as-a-Service
Performance
Availability
Durability
Governance
Locality
Cost savings
SAS NVMe
iSCSI / FC NVMeOF
DAS / SAN / BLOCK
NFS SMB
NAS /
S3 Swift
OBJECT
ALL PROTOCOLS
Cloud Storage

Hammerspace Key Capabilities

- Cloud and multi-cloud native
- Data is omni-present
  - Geo-spanning multi-site global namespace
  - Active-active eventually consistent with automatic collision disambiguation
  - On-demand and policy based granularly data replication
  - Consumable over file (pNFS, NFSv3, SMB), block and soon S3 object protocols
  - Kubernetes CSI driver with support for unified block and file
- Performance and scale
  - Multi-PB with linear performance scaling
  - Ultra-high performance across all I/O workloads
- Autonomic data placement
  - Utilize all storage types and vendors
  - File-level data management
  - Live migration / mobility
  - WAN optimization – deduplication, compression encryption to reduce bandwidth usage
- Full enterprise data services (snapshots, clones, undelete, auto-recovery)
- Enhanced metadata and metadata services
- Telemetry of access and performance

Universal Global Namespace

Fully managed, Active-Active, Multi-site

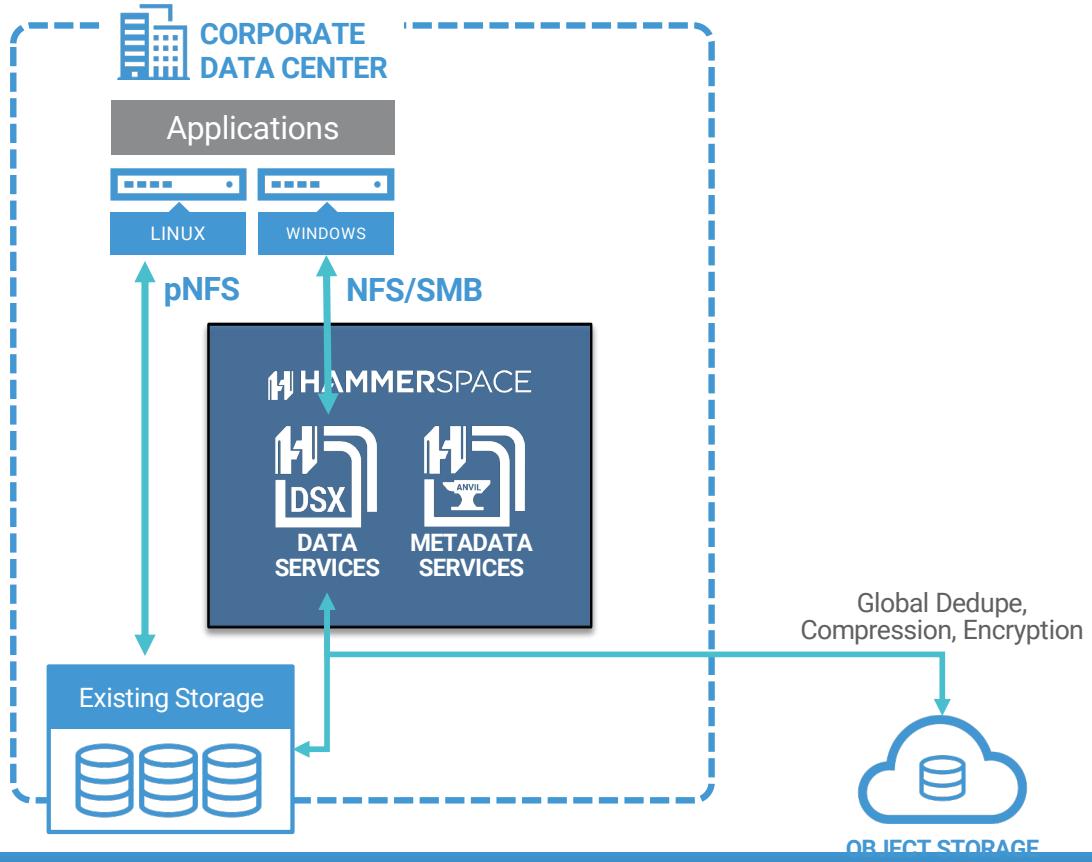

Universal Global Namespace

Fully managed, Active-Active, Multi-site

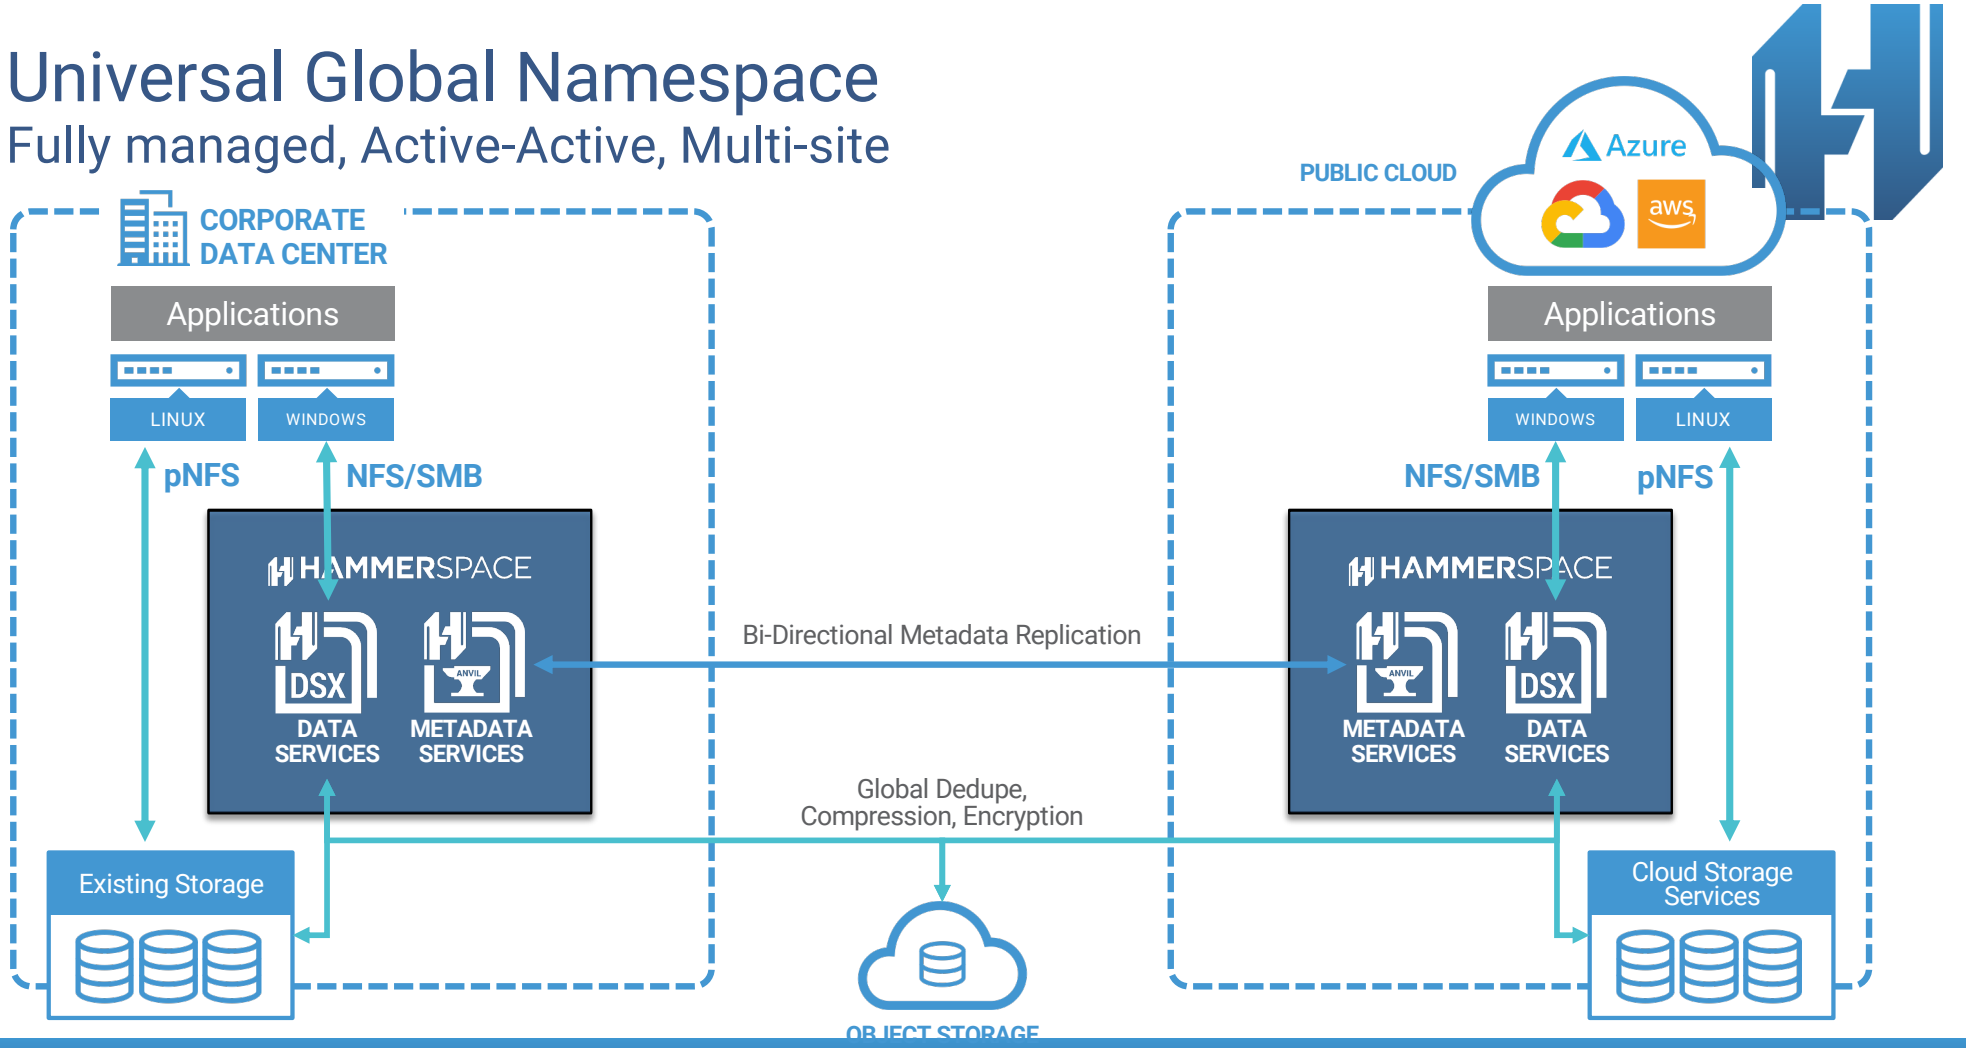

Defy Data Gravity – Rise above the infrastructure

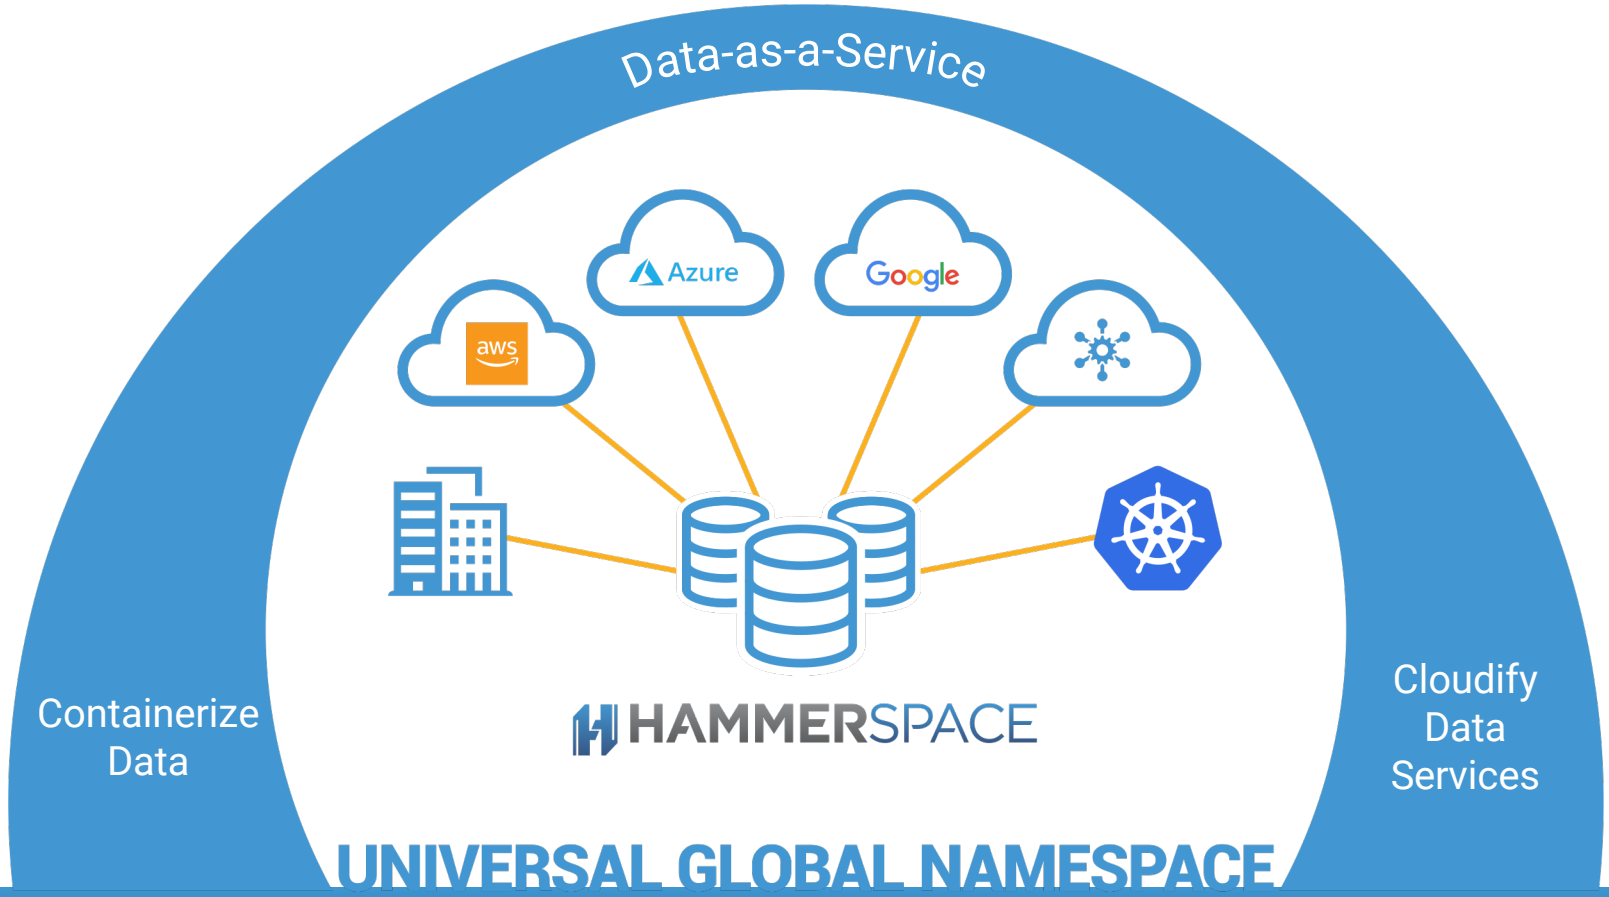

Broad applicability across Use Cases

Data Mobility Without Disruption

- Data Migration, Performance load-balancing, Tiering, Transparent Recovery

Multi Site Data Management

- Collaboration, Datacenter Migration & Migration, Active DR, File Services in the cloud

Burst

- Leverage cloud capacity when needed, Expand workloads into the cloud, Migrate to the Cloud

Unified File Data Services

- Single namespace across all vendors, Enterprise data services, Tiering to Object Storage, Kubernetes CSI

On-Ramp to Hybrid Multi-Cloud

Simplify Burst-to-Cloud and Multi-site workflows

- Turn on Hybrid Cloud in 10 minutes
- Use existing data sources in-place
- Make any storage cloud-native
- Break vendor lock-in
- Cost vs. performance optimization

Deploy Hammerspace → Connect Storage → Run Apps Anywhere

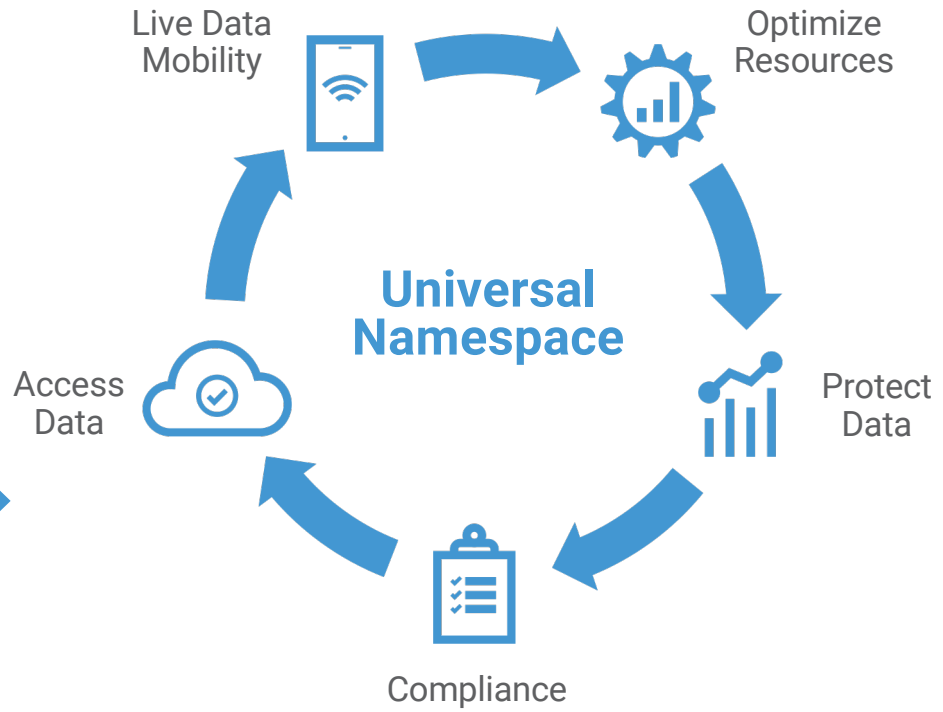

Bring Enterprise Data Services to Kubernetes

DevOps and Database workflows

- Make any storage K8s-native
- Multi-cluster support on any platform, local and over distance
- Instantly start stateful apps anywhere
- Native integration using container storage interface (CSI)
- Data protection using snapshots
- Active-active Disaster Recovery

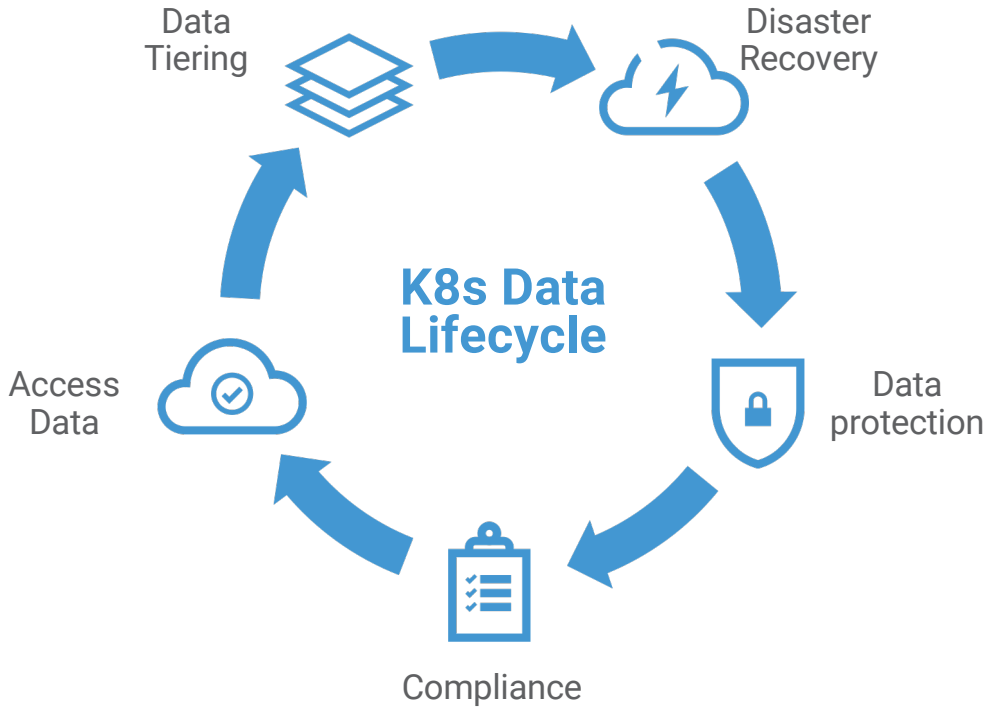

The Customer Journey

Streamline Hybrid Multi-cloud

- Cloud Performance Management
- Cloud Cost Management
- Reduce Legacy Infrastructure

DevOps DataOps

- Multi-Cluster Kubernetes Data Management
- DevTest Workflows
- Stateful Apps in Production
- Data Protection for Containers

On-Ramp to Hybrid Multi-Cloud

- Burst-to-Cloud
- Multi-site Collaboration
- Autonomic Data Mgmt
- Active-DR

Demo

When you need to do this…

Hybrid Multi-Cloud
• On-ramp • File data services • Data management • Data protection • Security and governance • Burst to cloud • Lift-and-shift to cloud • Analytics & ML • On-prem white-box

Kubernetes
• On-ramp • Multi-cluster • Databases • Systems-of-record • Production to TestDev • Analytics • CI/CD • Content management • High performance

Multi-Cluster Kubernetes in Production

DevOps and Database workflows

- Instantly start Kubernetes workloads anywhere
- Use any storage
- Reduce storage and transfer costs
- Data protection and Active-active Disaster Recovery

HAMMERSPACE.com

Data-as-a-Service

Hammerspace Architecture

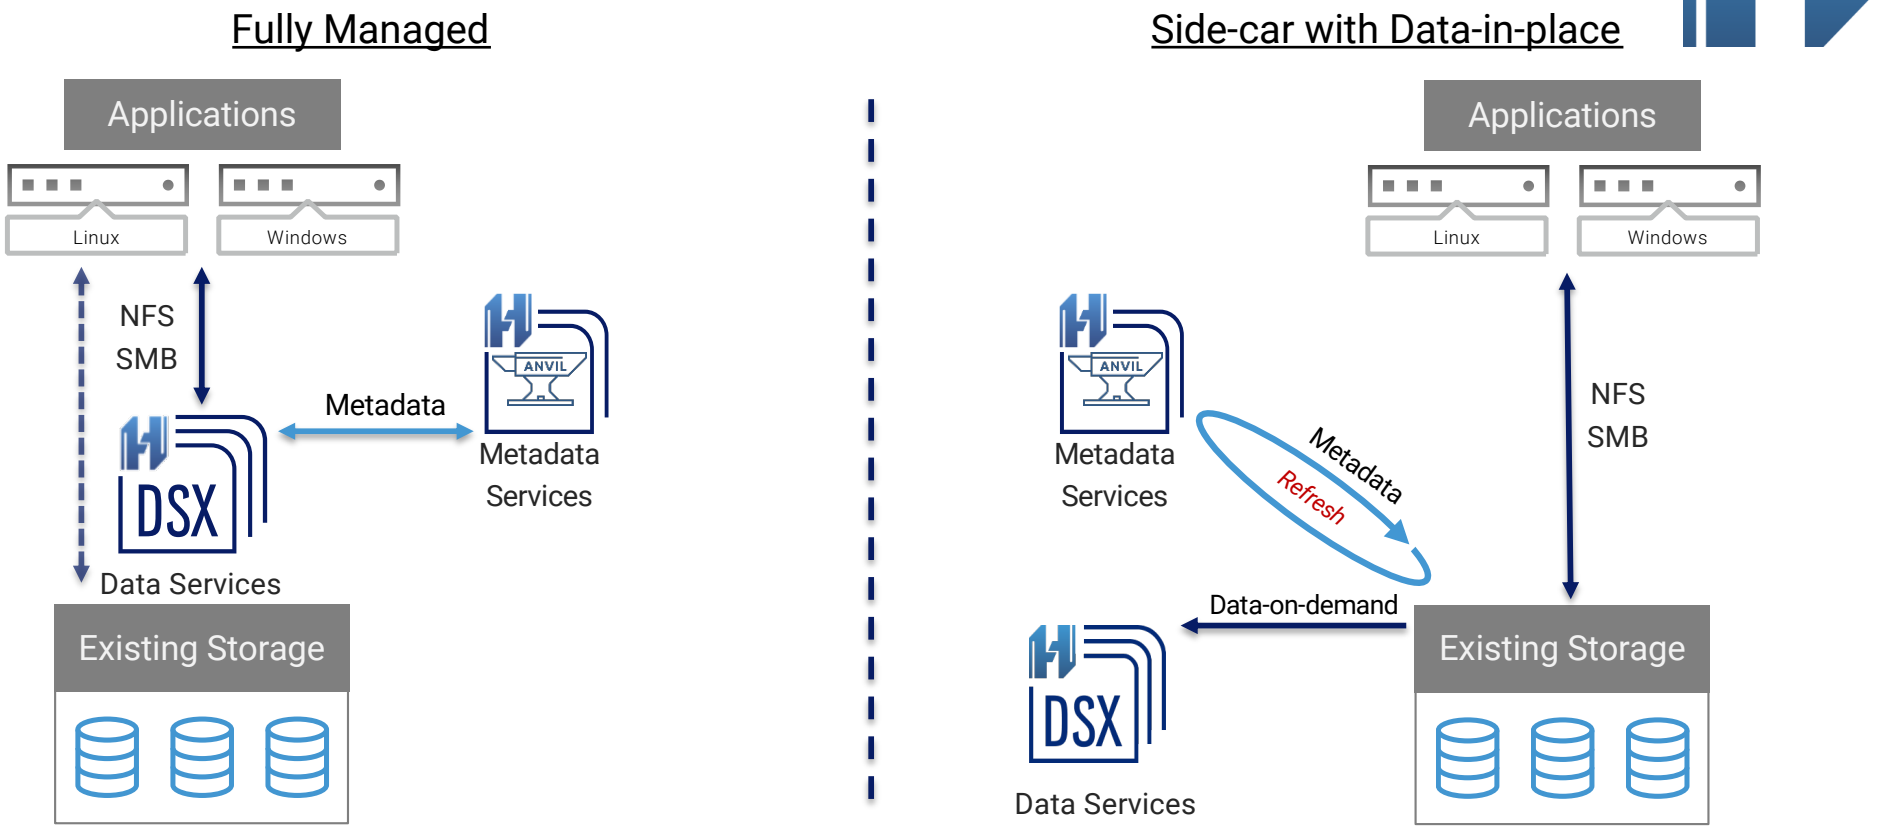

APPENDIX

Company Overview

- Headquartered in Silicon Valley
- Industry's top talent
- Uniquely technology
- Cloud native
- 6 years $100M in the making
- General availability Q1'19

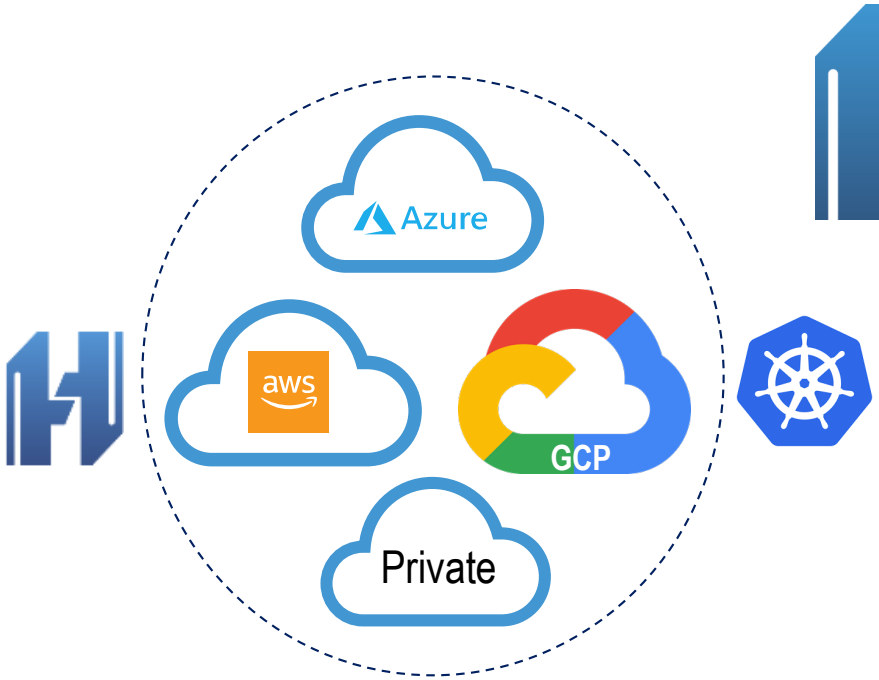

The Mission:

Simplify data use and management across hybrid multi-cloud and container environments

What is Hammerspace Data-as-a-Service?

- **Hammerspace is *not* storage**
- **Hammerspace is *not* a point solution to data silos**

Hammerspace is the extradimensional, instantly accessible storage area which is used to explain how magicians, animated, comic, and game characters can produce objects out of thin air.

https://en.wikipedia.org/wiki/Hammerspace

- **Hammerspace is the extra-dimensional existence of data**
- **Hammerspace is a universal namespace**
- **Hammerspace is the containerization of data**
- **Hammerspace unlocks data from infrastructure silos**

Data-as-a-Service Reduces the complexity of hybrid multi-cloud

Easy to get started

- Cloud native
- Use apps as-is
- Use any storage
- Non-disruptive
- Auto deployed
- Avoids point-solutions

Manages Itself

- Objective driven policy
- Continuous optimization
- Avoid IT intervention
- Dynamic scaling
- Dynamic load balancing
- Live data mobility
- Automated recovery

Unify Kubernetes Data Management with Universal Namespace

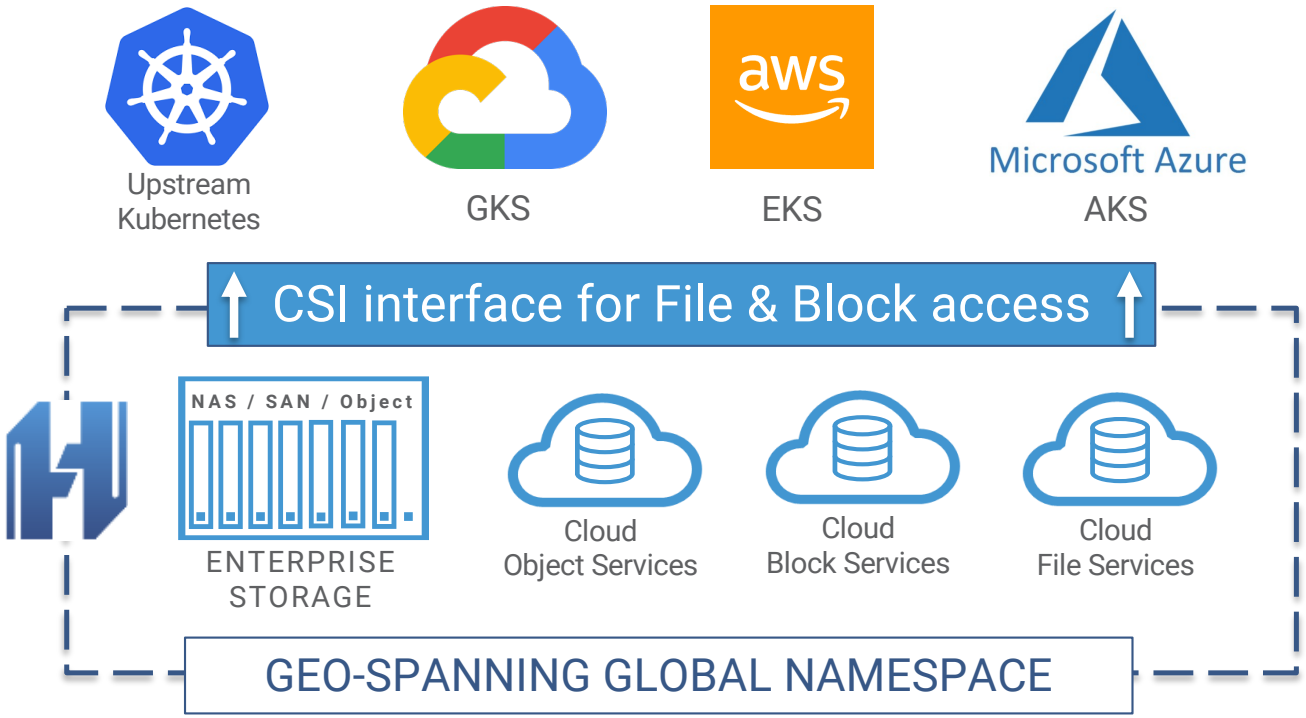

Enterprise Cloud Data Services

Hybrid Multi-Cloud Data Services


Metadata Management


Automation


Data Mover


Security


Global Namespace


Data Governance


Data Protection

Pay-as-you-Go


Capacity Managed


Software-only


Cloud Marketplace

Deliver Data-as-a-Service Anywhere


Microsoft Azure

Legacy Data Center Adoption Progression

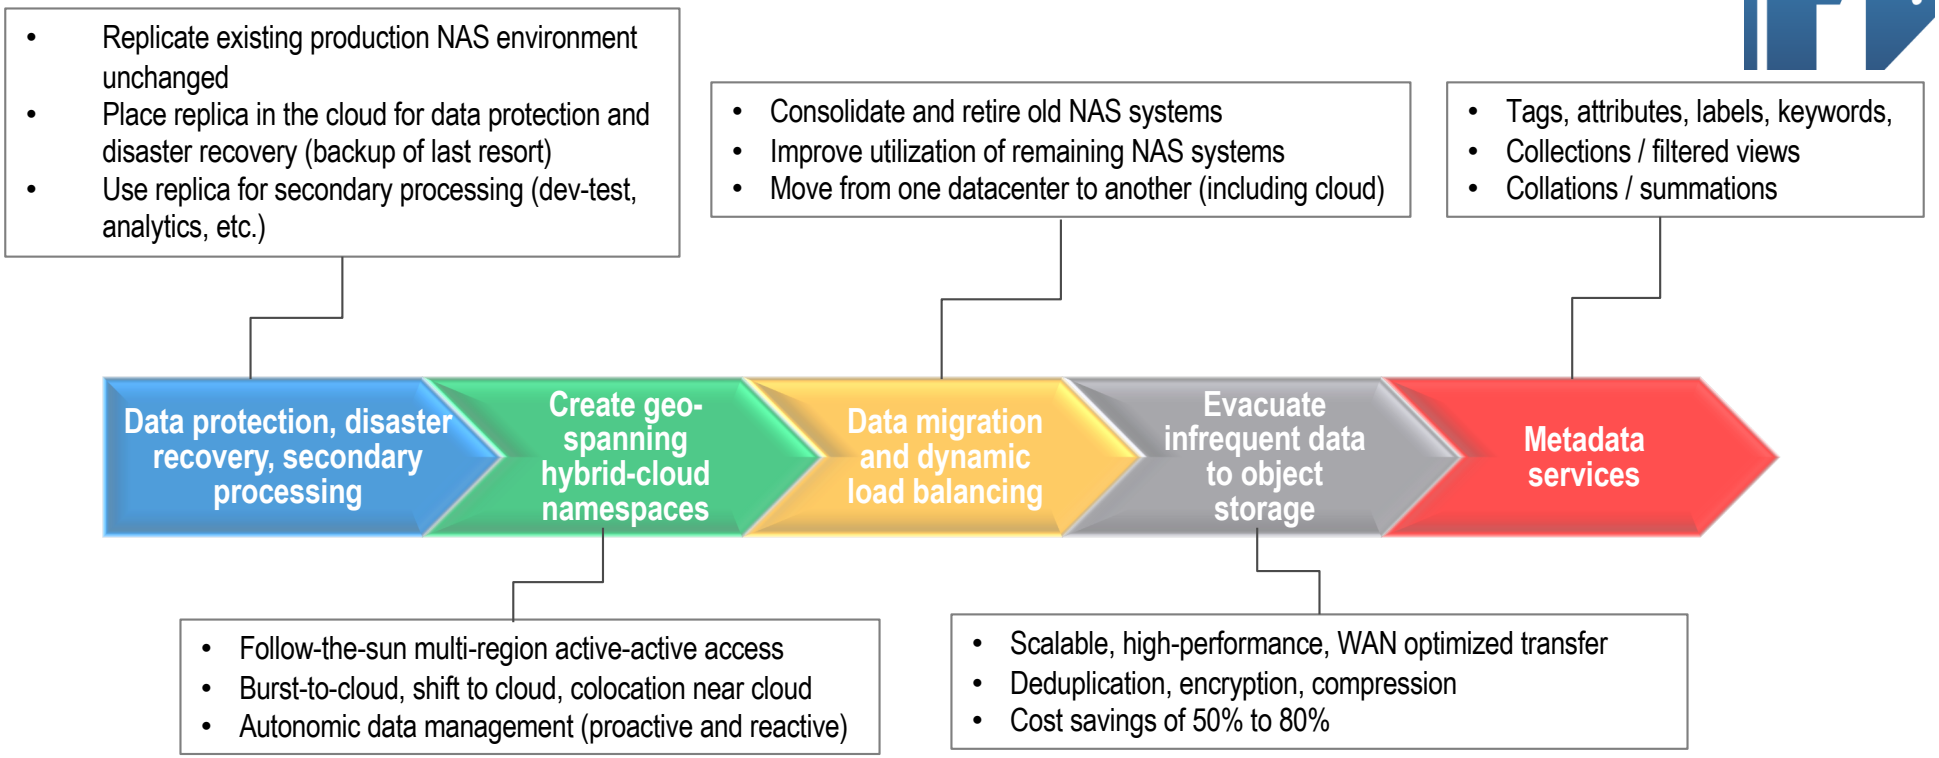

Broad Platform Support

Cloud & Kubernetes

Upstream Kubernetes

amazon web services
EKS

GKS

Azure
AKS

redhat
TECHNOLOGY PARTNER
OpenShift

Ecosystem

NetApp

CLOUD NATIVE COMPUTING FOUNDATION

EMC² ISILON

SCALITY

CLOUDIAN

IBM Cloud
Object Storage

HGST
a Western Digital brand

DELL EMC ECS

Open Telekom Cloud

HITACHI
Inspire the Next
Hitachi Vantara