HPE Data Fabric Software

Data challenges

Diverse, complex data challenges quality assurance

67% of data scientists' time is spent on **data preparation** tasks

48% of data scientists struggle with **access rights** to diverse data sources

Diverse data types distributed across multiple sources

Increased complexity

- Difficulty in analyzing large volumes of data from distributed sources
- Data mapping and reconciliation
- Monitoring and maintenance

- Multiple data types / formats
- Lots of diverse workloads often requiring the same data

Higher costs

- Licensing for diverse tools
- Infrastructure
- Personnel
- Unpredictable cost – egress/networking

Reduced flexibility

- Data quality
- Security risks
- Reduced agility

Top concerns we hear from global leaders regarding their future data, AI & analytics roadmaps

Managing distributed data is complex

Productivity is stifled by siloed development

Costs can be high and unpredictable

To achieve the full potential for AI and analytics **across** hybrid multi-cloud environments, **today's workloads demand modern software that is...**

Hybrid by design

Make the technology invisible to abstract the developer from the complexities of deployment by solving the 'across' challenges of hybrid cloud

Open

Teams have the freedom to use the tools, tech, and data format best suited for the task

Committed to open-source providing evergreen tooling without vendor lock-in

Extensible

Flexibility to adapt as technology and needs evolve

Easily add new types of data, sources and bring in tools, including open-source, ISVs, and custom apps

Consistent experience

Delivers a simple, fully managed, end to end hybrid cloud solution with a consistent experience and predictable consumption-based costs

HPE Data Fabric Software

Federated data layer
The backbone of AI initiatives

Need a new approach that federates data across all formats and locations

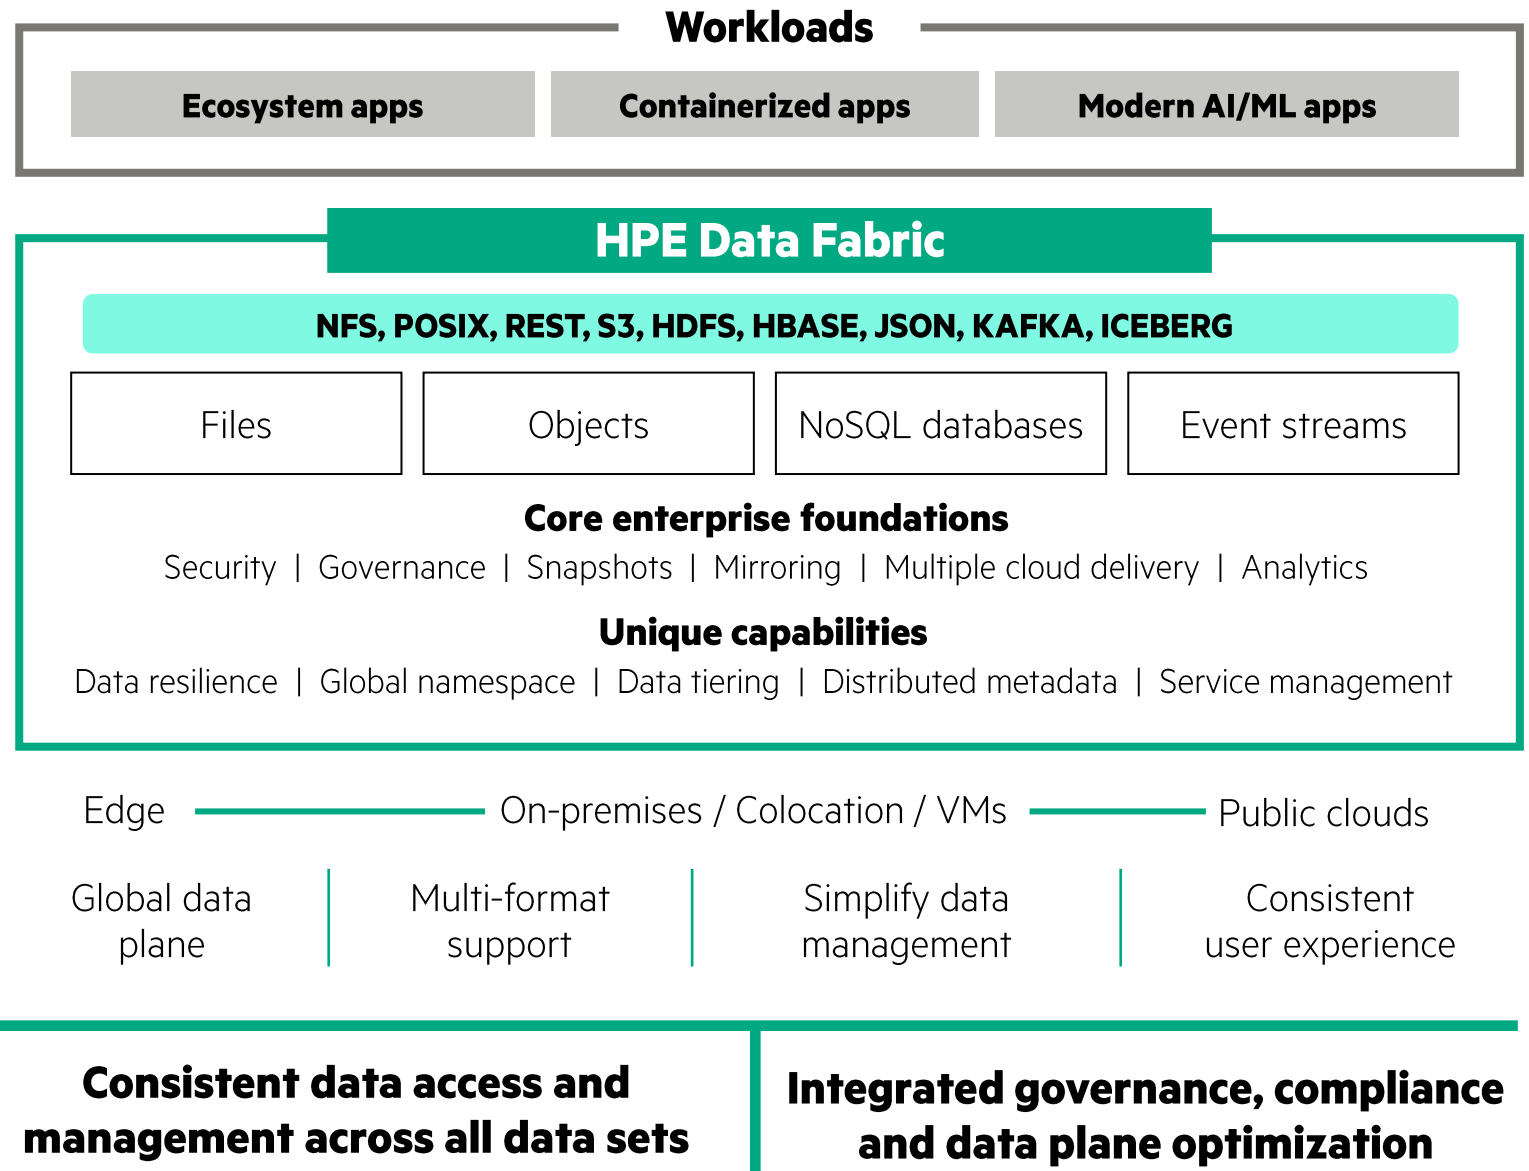Data accessibility and consistency

Scalability and flexibility

Enhanced data security and privacy

Improved AI model development & deployment

HPE Data Fabric Software

An **intelligent data plane** with radically simplified management, control and governance for analytic and AI applications across all enterprise storage assets

Workloads

Ecosystem apps	Containerized apps	Modern AI/ML apps

HPE Data Fabric

NFS, POSIX, REST, S3, HDFS, HBASE, JSON, KAFKA, ICEBERG

Files	Objects	NoSQL databases	Event streams

Core enterprise foundations

Security | Governance | Snapshots | Mirroring | Multiple cloud delivery | Analytics

Unique capabilities

Data resilience | Global namespace | Data tiering | Distributed metadata | Service management

Edge ———————— On-premises / Colocation / VMs ———— Public clouds

Global data plane	Multi-format support	Simplify data management	Consistent user experience

Transparent and secure data access to AI, analytics and data lakes

Consistent data access and management across all data sets

Integrated governance, compliance and data plane optimization

HPE Data Fabric Software — Vision

HPE Data fabric Software helps users access, manage, organize and govern enterprise data, across a variety of formats, in a **single**, consistent, easy to use, edge to cloud **data plane** optimized for analytics workloads
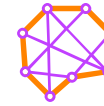
Access and manage data universally across data sources, and locations

Simplifies management, access-control and governance of data across on-prem, public cloud, and hybrid deployments

Single enterprise-grade, battle tested platform that supports most AI and analytical data formats

Best-in-class security features including geofencing, encryption at rest, extensible security policies

First to autonomously optimize cost, performance, compliance, etc and radically simplify admin experience for complex storage deployments

# Delivering a unified approach to hybrid data lakehouse architectures

Transparent, secure access for AI, analytics and all data lakehouse data

Consistent data access and management across all data assets and formats

Integrated governance, compliance and data plane optimization

View all data assets in a **global namespace**

Simplify many data sources into **unified views**

Data fabric
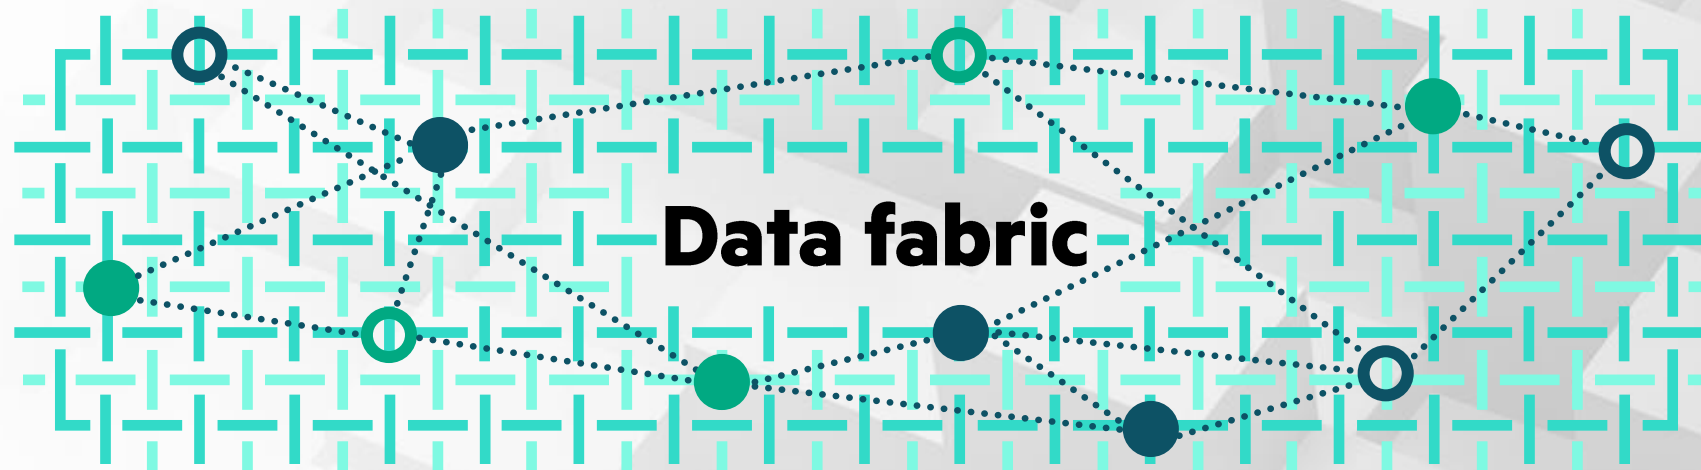
Enable all lakehouse data formats into **one data plane**

Manage all your data assets though one hybrid data management plane

Consolidate and automate governance, compliance, and security alerting and reporting

Optimize data placement to improve efficiency and reduce costs

Edge ···· **Cloud** ···· **VM** ···· **Data center**

HPE Data Fabric Software

Key differentiators

Unified view and management

Provides unified view of data across multiple storage systems, simplifying data management—reducing complexity, latency and cost

Multi-format support

Out-of-box support for the most popular analytics data formats (files, objects, tables, & streams) provides maximum flexibility of use

Global namespace

Enables seamless data federation allowing users to query and access data across different storage locations without data movement

Hybrid and multi-cloud deployment

Works across on-premises, cloud and edge environments offering vendor agnostic flexibility

Optimized for AI and analytics

Designed to accelerate ML workflows with Apache Iceberg support, improving query performance and scalability

Governance and compliance

Integrated governance, and policy driven encryption, access control and regulatory compliance to meet industry standards (SOC2, HIPPA, GDPR)

What sets HPE Data Fabric Software apart?

|	HPE Data Fabric	Competition
Data federation	Unified Global Namespace across storage platforms	Requires data ingestion into platform-native format
Hybrid and multi-cloud deployment	Deploy across on-premises, cloud, or at the edge like a true data mesh	Primarily cloud-dependent
Multi format support for modern AI and analytics workloads	Apache Iceberg support, real time query acceleration	Limited support for multi-format AI workloads
Security and compliance	End to end encryption, RBAC and audit trails	Requires third-party integrations for compliance
Storage optimization	Tiered storage and intelligent data placement for cost efficiency	Standard data storage models

Global namespace

- **Unified naming system** that organizes data under a consistent naming scheme
- **Abstracts** data's physical location
- Provides **multiple entry points**, but structure remains the same for all users
- **Easily scale** as data volumes/users grow and add new data sources

Instant access to your entire data ecosystem
Global namespace

- Reduces time consuming data transfers between systems
- Enables real-time analysis and decision-making
- Maintains data integrity and consistency
- Reduces costs, improves governance
- Protect investments in existing big data estates
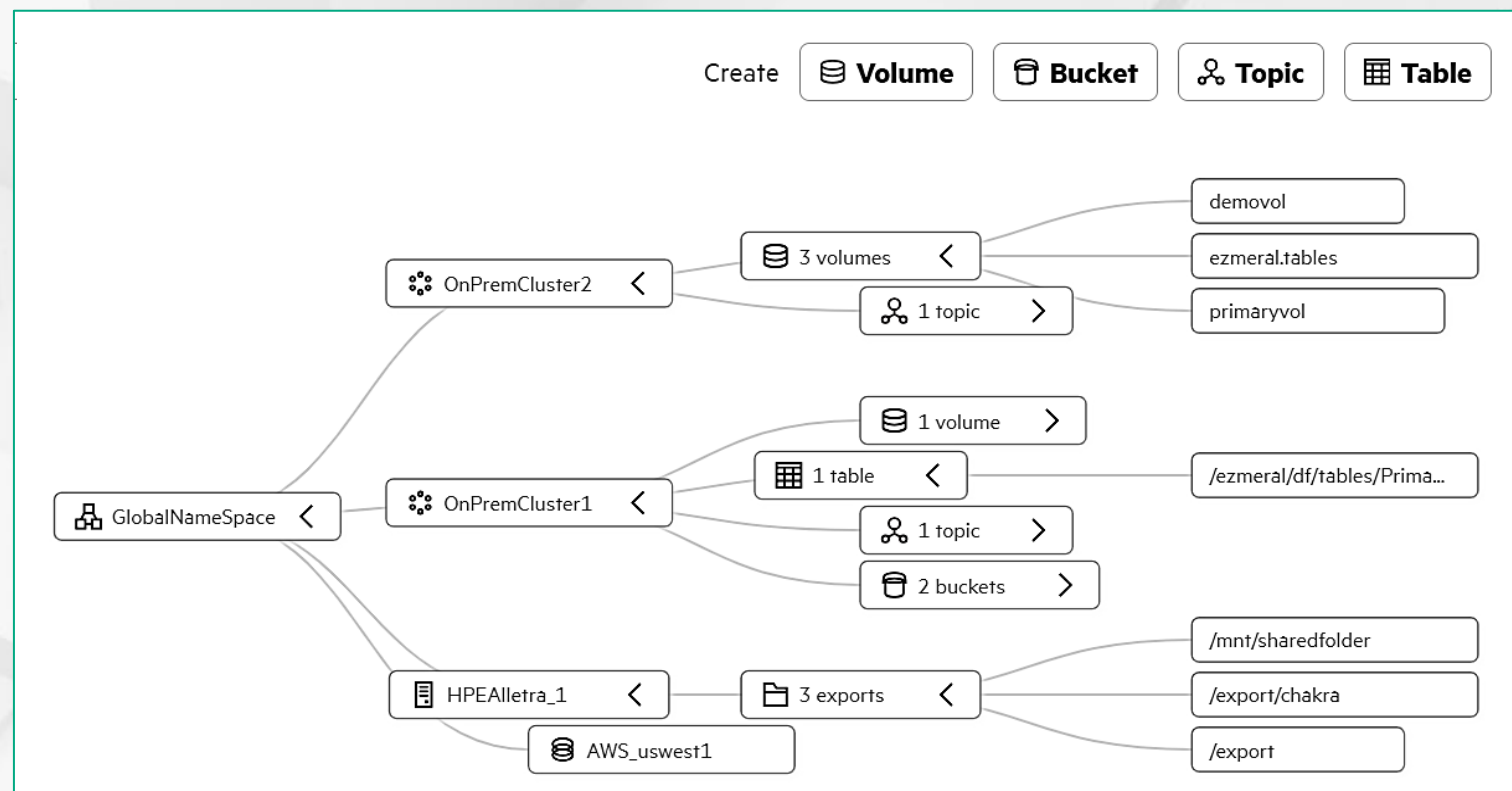
HPE Data Fabric Software
Customer benefits

Transparent data access to AI, analytics and data lakes

- Unified data view across multiple storage systems
- Real-time insights and decision making

Consistent data management solution across all data sets

- Unified data management
- Integrated with wide set of tools and systems
- Software defined platform

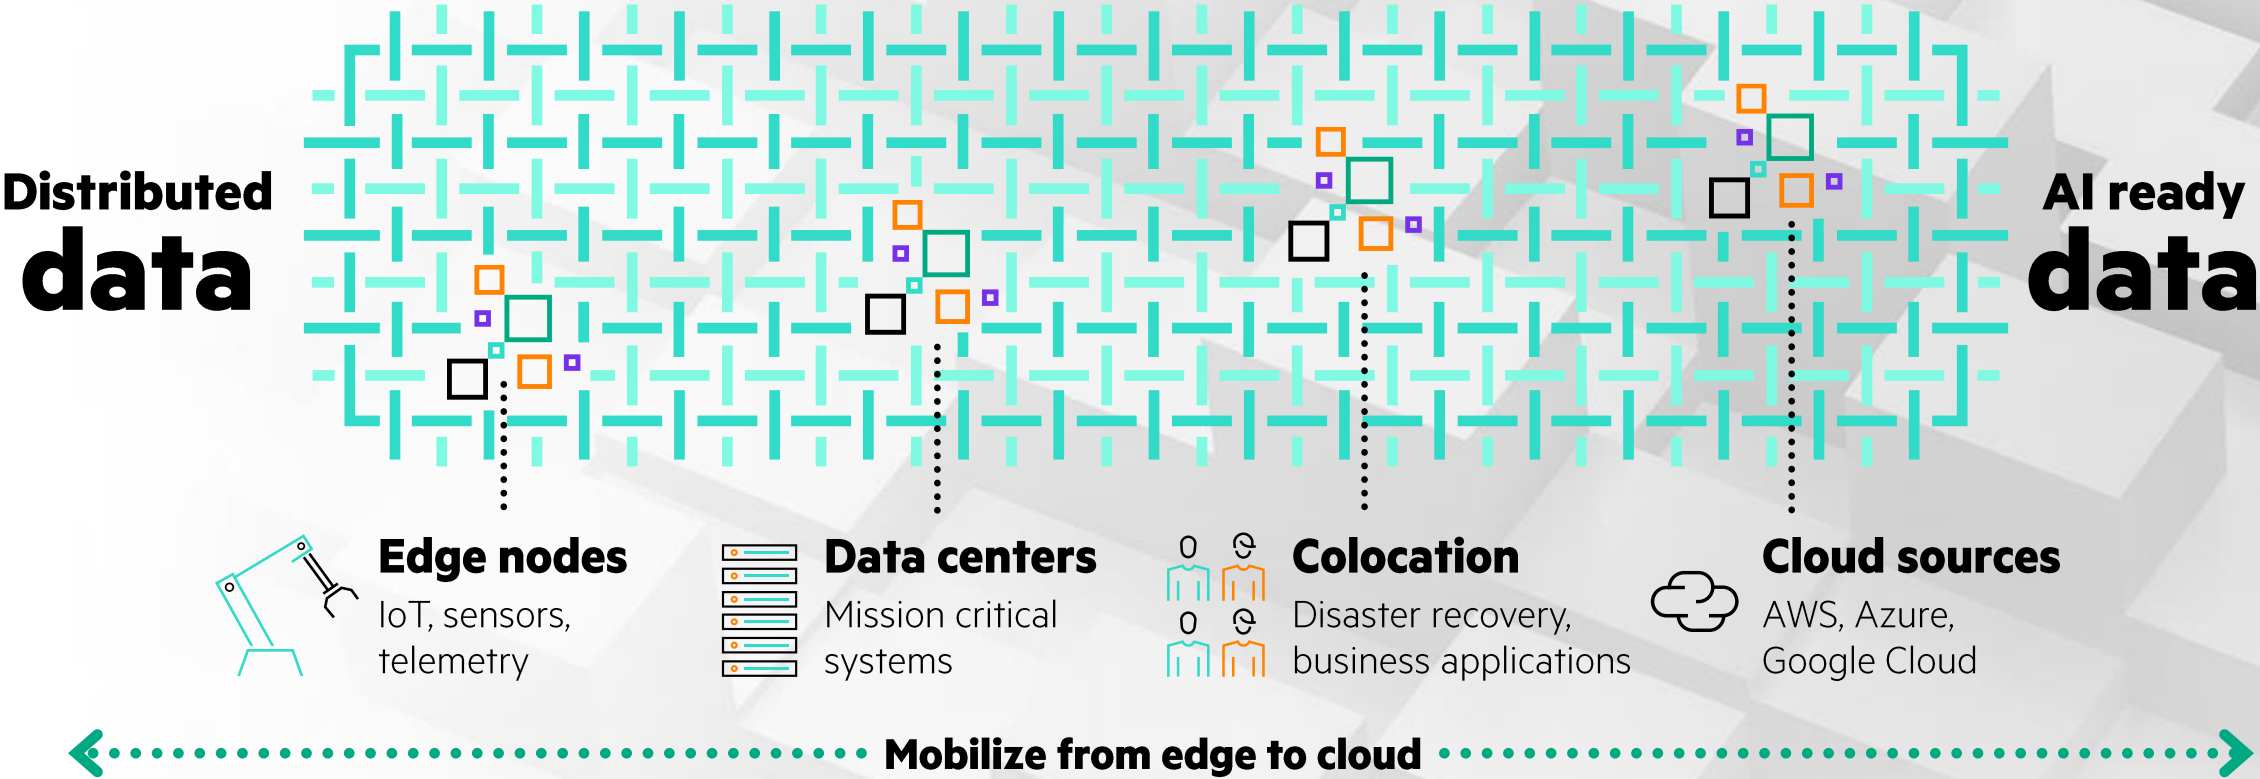**Secure governance, compliance and data plane optimization**

- Single point for collecting, analyzing and enforcing compliance and governance
- Consolidated data in a single platform
- Robust security and compliance

Unifying distributed data to accelerate AI insight

Universal access | Global namespace | Automated data tiering | Multi-protocol

HPE Data Fabric Software

Distributed data

AI ready data

Edge nodes
IoT, sensors, telemetry

Data centers
Mission critical systems

Colocation
Disaster recovery, business applications

Cloud sources
AWS, Azure, Google Cloud

Mobilize from edge to cloud

Leading the way to AI data management

Summary

Unified data management

Unified view of data across multiple sources simplifying management and reducing latency and costs

Multi-format

Supports multiple protocols in a single deployment for AI and analytics use cases

Software-defined

100% software-defined platform integrating with existing storage infrastructure

Data lakehouse for AI/analytics

Real-time insights for better decision making and integrates with data engineering and analytic tooling

Robust security and compliance

Meet stringent compliance and security requirements with built-in access control and encryption

Seamless hybrid integration

Users can grow their data lakehouse without compromising on performance and scalability

Thank you