Have feedback on QuickSpecs? We're listening

HPE Data Fabric Software QuickSpecs

The rise of hybrid and distributed environments that span on-premises, cloud, and edge has redefined how enterprises manage and extract value from data.

Today, data is not just an asset; it is the foundation for intelligence, automation, and innovation. As organizations embrace AI and machine learning, the need for trusted, high-quality, and context-rich data has never been greater.

Market Overview

The rise of hybrid and distributed environments that span on-premises, cloud, and edge has redefined how enterprises manage and extract value from data. Today, data is not just an asset; it is the foundation for intelligence, automation, and innovation. As organizations embrace AI and machine learning, the need for trusted, high-quality, and context-rich data has never been greater.

Traditional data architecture struggle to keep pace with these demands. The proliferation of data silos, inconsistent governance, and fragmented security models hinder visibility and slow decision-making. To compete in an AI-driven world, enterprises require a unified, intelligent data fabric - one that ensures secure access, consistent governance, and continuous insight across all environments.

The modern Data Lakehouse addresses these challenges by providing a single platform to store, process, and analyze structured and unstructured data at scale. When extended with AI-powered governance, policy automation, and metadata intelligence, it evolves into a true data fabric, enabling trusted analytics, accelerating AI adoption, and transforming data into a strategic differentiator.

Challenges of the Modern Data Landscape

Enterprises today face mounting challenges as data environments grow in scale, distribution, and complexity. Data sprawls across hybrid and multi-cloud infrastructures, often trapped in incompatible formats and isolated silos that limit accessibility and reuse. Managing this sprawl introduces operational friction with rising storage costs, inconsistent performance, and increased administrative overhead. Governance and compliance add further pressure, with regulations demanding granular visibility and auditable control across every data source. At the same time, emerging AI workloads require low-latency, high-context access that legacy systems were never designed to deliver. The result is a landscape where agility, trust, and efficiency are difficult to achieve and mandates the need for a unified, intelligent data fabric that can seamlessly bridge these divides and enable AI-ready data operations at scale.

Product Overview

HPE Data Fabric Software is a next-generation data platform that unifies data management, governance, and intelligence across hybrid, multi-cloud, and edge environments. Designed for the AI-era, it transforms complex, distributed data ecosystems into a single, secure, and context-aware fabric enabling organizations to access, govern, and utilize data with unprecedented efficiency and trust.

The platform provides a unified, location-independent view of data, simplifying access to files, objects, streams, and databases while maintaining enterprise-grade security and policy control. Integrated metadata intelligence supports advanced lineage, auditability, and governance empowering teams to make informed, compliant, and timely decisions. It’s AI-powered control plane further automates data discovery, access, and compliance. This intelligent layer ensures data is consistently governed and readily available for both analytics and AI workloads. These capabilities form the backbone of a modern data platform that simplifies operations while driving greater business agility.

In essence, HPE Data Fabric Software delivers measurable cost savings and productivity gains by consolidating data from on-premises, cloud, and edge environments into a single logical store. By minimizing data movement and duplication, it optimizes resource utilization and operational efficiency allowing businesses to focus on innovation rather than administration, and to accelerate their journey toward data-driven and AI-enabled outcomes.

Overview

Business Outcomes

Drives Business Agility with Data Lakehouse: HPE Data Fabric Software enables organizations to create a robust data Lakehouse that supports structured and unstructured data, facilitating advanced analytics. This vendor-agnostic, location-independent solution drives business agility by simplifying the integration of new technologies and promoting data reuse across various projects, resulting in cost savings and increased efficiency.

Accelerates Data-Driven Insights: By unifying diverse data types such as files, objects, NoSQL databases, and real-time/batch streams into a unified data backbone, HPE Data Fabric Software provides a comprehensive view of data crucial for analytics, AI, and ML initiatives. This unified data environment empowers businesses to rapidly identify trends, anomalies, and relevant datasets for faster decision-making and strategic planning.

Streamlines Data Access & Management: HPE Data Fabric Software simplifies data management and control with a single pane of glass view and a global namespace for effortless data discovery, for data across multiple locations and formats, including on-premise, public cloud, and edge devices. By providing a comprehensive 360-view of your data landscape, HPE Data Fabric Software ensures streamlined data governance, improves compliance with data security and privacy regulations.

Hybrid Security Simplified: HPE Data Fabric Software seamlessly integrates with existing IAM systems for authentication and authorization, offering a single security management system for hybrid environments. This integration, along with advanced features like automated policy management, minimizes risks associated with manual security management and reduces the cost.

Improved Intelligence with Advanced Analytics Support: HPE Data Fabric Software supports continuous intelligence and advanced analytics through seamless integration with certified open-source tools and offers native support for CSI drivers and other HPE containerization platforms. By leveraging NVIDIA GPU Direct technology, the software accelerates AI and machine learning workloads, ensuring optimal performance, scalability, and data-driven innovation for businesses seeking to enhance operational efficiency and decision-making.

Enables AI-Ready Operations: Through its AI-powered control plane and intelligent metadata governance, HPE Data Fabric establishes a trusted foundation for next-generation AI workloads. Data remains discoverable, governed, and context-rich - allowing enterprises to operationalize AI securely and at scale with confidence and agility.

Overview

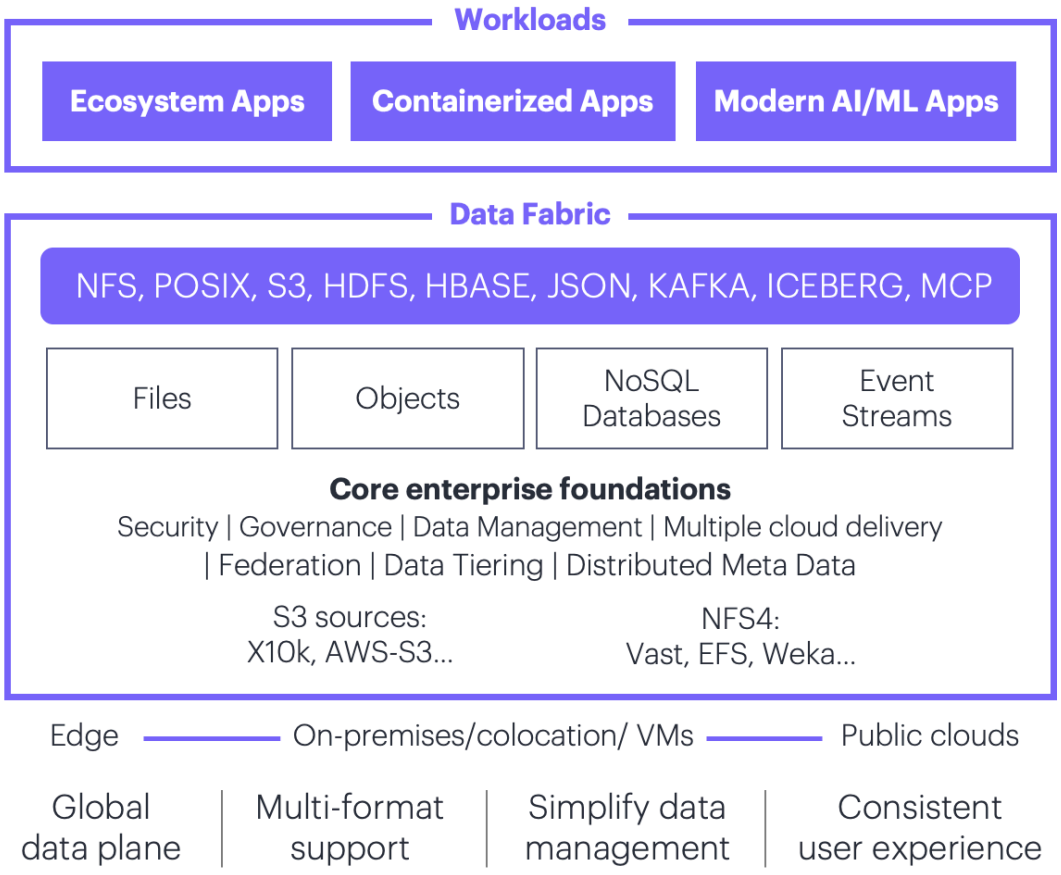

HPE Data Fabric Software Overview

What's New:

HPE Data Fabric 8.1 elevates next-generation capabilities that unify AI interoperability, strengthen metadata intelligence, and extend real-time analytics across the hybrid data landscape.

- Event-Driven Object Movement Framework: Introduces a new event-driven, object movement capability that enables near real-time data movement from any S3-compatible source such as Data Fabric, Alletra MP X10000 to any S3-compatible targets using native event notifications. This establishes the foundation for hybrid data plane, supporting use cases such as data tiering, real-time ingestion, cloud bursting, and cross-environment analytics.
- Federated Metadata Intelligence with Polaris and MCP Integration: Expands Apache Polaris with federated catalog capabilities across environments and integrates with the MCP server to enable governed, AI-driven access to both data and metadata through natural language. This enhances metadata visibility, control, interoperability, and strengthens governance and compliance across Data Fabric.
- Enhanced Tiering and Data Efficiency with Inline Erasure Coding: Introduces inline erasure coding for S3 object storage, improving storage efficiency by encoding data immediately on write, reducing overhead compared to replication and optimizing capacity utilization at scale.

Overview

Standard Features

Feature Set for HPE Data Fabric Software

Global Namespace

- Global view of all data sources (on-prem, cloud, edge)
- Federated sources, NFS, and S3
- Unified authentication policy & security
- Optimized cloud deployment (AWS, GCP, Azure)
- Apache Iceberg support
- NVIDIA Direct Support

Data Management:

- Support for files and objects
- Event Streams
- NoSQL Database
- Data Tiering: Hot (default), Warm (Erasure coded), and Cold (off-load to external stores)
- Support for NVMe, SSD, and HDD
- Mirroring and snapshots

Data Access:

- POSIX client (unlimited client instances)
- NFSv4 and NFSv3
- HDFS and HTTPFS APIs
- S3 API
- Kafka compliant streams API
- NoSQL databases for analytics and operational applications: a key-value database with HBase API and a JSON document database with OJAI
- Container Client including PACC, CSI (unlimited client instances)
- Model Context Protocol (MCP) Server

Core Data Capabilities:

- Resiliency and self-healing
- Auto-balancing
- Multi-site volume mirroring
- Volume Mirroring for Disaster Recovery

Performance and Optimization:

- Data Topologies
- Compression
- Multiple file server instances per node

Administration and Security:

- Administration using CLI, REST, GUI
- Rolling upgrades
- UI Installer
- Quotas
- Multi-tenancy (not for Object Store)
- Insights Service to track user behavior

Standard Features

- Unified security including authentication, authorization, wire and data-at-rest encryption
- Unified roles
- User-managed roles
- IPv6 support

HPE Data Fabric Software Supported Data Formats

Files: A distributed file system designed to handle structured and unstructured data with reliability and scalability. It features patented capabilities that ensure data protection and efficient storage management for diverse data types.

Object Store: A native object storage solution designed to efficiently manage objects and metadata for optimized access. As a fully S3-compliant storage system, it provides seamless integration with applications that require compatibility with S3 sources.

NoSQL Databases: Two databases for analytics and operational applications

- Columnar (Key-value) Database: Columnar Database is supported on the platform, and you can use HBase API to interact with this data.
- JSON Document Database: JSON document database is supported on the platform and one can use OJAI API to interact with the data.

Event Streams: A built-in publish-and-subscribe messaging system that supports Kafka APIs. It enables seamless integration without the need for additional process management, making it an ideal choice for real-time data processing and event-driven applications.

HPE Data Fabric 8.0 Add-on options:

HPE Data Fabric Analytics for Apache Spark

Spark 3.5.5.500: Upgraded Spark version in Ecosystem Pack to enhance large-scale data processing and analytics capabilities within Data Fabric, supporting advanced machine learning, graph processing, and real-time streaming applications.

HPE Data Fabric Analytics Premium Package

Includes Spark, Hadoop, Drill, Apache Ranger, and Apache NiFi, providing a comprehensive suite of tools for data management, processing, analytics, and security within the Data Fabric ecosystem.

Apache Hadoop 3.4.1.400

Offers distributed storage and processing capabilities for large datasets across clustered systems within the Data Fabric environment.

Apache Drill 1.21.2.350

Provides a powerful SQL query engine for Data Fabric users to explore and analyze large-scale datasets, enabling interactive analytics and data exploration on Hadoop-based systems.

Apache Hive 3.1.3.950

Simplifies data analysis and manipulation operations within Data Fabric, allowing users to read, write, and manage large datasets using SQL-like queries.

Standard Features

Apache Airflow 3.1.7.0*

Facilitates the creation, management, and scheduling of complex data pipelines within Data Fabric by enabling users to define and execute Directed Acyclic Graphs (DAGs) of tasks using Python.

Apache Ranger 2.4.0.200

Ensures data protection and compliance across diverse data sources and applications within Data Fabric through centralized security administration and access control.

Apache NiFi 1.28.0.300

Offers efficient data management for Data Fabric users by providing a robust data integration and processing platform that supports dataflow automation, stream processing, and real-time data ingestion.

Apache Flink 2.0.0.100

Delivers real-time stream processing and event-driven analytics through a distributed processing engine that enables low-latency insights and continuous data computation at scale.

Notes:

- *Airflow is not supported on SLES 15 SP2. In addition, Airflow is supported only with Ubuntu 20.04 and 22.04; it is not supported with Ubuntu 18.04.
- The list above highlights key components of the Analytics Premium package and is not exhaustive. For a complete list, please refer to the official documentation: https://docs.ezmeral.hpe.com/datafabric-customer-managed/78/c_ecosystem_intro.html

Standard Features

HPE Data Fabric File and Object Store Capability Matrix	
Capabilities	**HPE Data Fabric File and Object Store**
Distributed File and Object store	Yes
Unified security including authentication, authorization, policy-based security, external KMIPkeystore, and encryption (on-wire & at-rest)	Yes
Support compression, auto-balancing, resilient through replication	Yes
Highly scalable, global namespace, quotas, auditing, rolling upgrades	Yes
Data topologies, data tier support: 'hot' for fast data; 'warm' for Erasure coded and 'cold' to offload data to 3rd party S3-based external stores.	Yes
Container Client including PACC, Flex Volume Plugin, CSI (unlimited client instances)	Yes
Snapshots and mirroring of data across clusters (used for data safety, access from multiple geo locations and disaster recovery)	Yes
Operating systems supported: RHEL, SUSE, Ubuntu and Oracle Enterprise Linux	Yes
Management using CLI, REST and GUI	Yes
Support for NVMe, SSD & HDD	Yes
Support for NFSv4 & NFSv3	Yes
Clients - Hadoop for MapR-FS, HTTP-FS	Yes
Multiple file server instances per node	Yes
Supports the following for File store only: – Multi-tenancy – Container client including PACC, Flex Volume Plugin, CSI – Data Topologies – Platinum POSIX client	Yes
Supports the following for Object store only: – S3 API, Buckets, Objects, IAM accounts – Policies for accounts, buckets, objects – Write Once Read Many (WORM) times support – S3 Select, S3 Gateway	Yes
HPE Data Fabric Installer	Yes
HPE Data Fabric Monitoring	Yes
Hive Metastore	Yes
Apache Ranger – Table replication that can be used for disaster recovery and multi-site access – Unified Security including authentication, authorization, encryption (wire and data-at-rest) and auditing – Asynchbase (Client for MapR-DB Binary Tables) – HBase (Full Service and as the client interface to MapR-DB Binary Tables)	Yes

Standard Features

Capabilities	HPE Data Fabric File and Object Store
HPE Data Fabric Event Data Streams: – Automatic partition balancing – Distributed publish-subscribe messaging infrastructure – Kafka API support, Kafka C Client, Kafka Java Client, Connect REST Proxy, Kafka Connect HDFS & JDBC, Schema Registry, Streams – KSQL, multi-site stream replication, multi-tenancy – Unified Security including authentication, authorization, encryption (wire and data-at-rest) and auditing – Kafka wire-level protocol support (POC only)	Yes
HPE Database: – Key-Value Database using HBase API – JSON document database using OJAI API – Automatic compactions, Change Data Capture (CDC) – HPE Data Fabric DB Data Access Gateway – Multi-master table replication – Multiple file server instances per node – Multi-tenancy, Policy-based security, Resiliency and self-healing, secondary indexes, strong consistency – Table replication that can be used for disaster recovery and multi-site access – Unified Security including authentication, authorization, encryption (wire and data-at-rest) and auditing – A synch base (Client for MapR-DB Binary Tables) HBase (Full Service and as the client interface to MapR-DB Binary Tables)	Yes
HPE Data Fabric Event Data Streams: – Automatic partition balancing – Distributed publish-subscribe messaging infrastructure – Kafka API support, Kafka C Client, Kafka Java Client, Connect REST Proxy, Kafka Connect HDFS & JDBC, Schema Registry, Streams – KSQL, multi-site stream replication, multi-tenancy – Unified Security including authentication, authorization, encryption (wire and data-at-rest) and auditing – Kafka wire-level protocol support (POC only)	Yes
Apache Iceberg	Yes

Standard Features

HPE Data Fabric for Analytics Capabilities Matrix		
Capabilities	**HPE Data Fabric Analytics for Apache Spark**	**HPE Data Fabric Analytics Premium Package**
Apache Spark Ecosystem: – Apache Spark Core, Mllib, Astreaming and YARN – GraphX – HPE Data Fabric DB Binary and OJAI Connectors and for Apache Spark – HPE Data Fabric Streams Integration – SparkR, Spark Standalone and Spark on YARN support – Helm Charts for Spark	Yes	Yes
GPU based acceleration	Yes	Yes
SQL ecosystem with Drill: – Drill Explorer, Monitoring, query and admin UI, Drill standalone and Drill-on-YARN support – File formats (Text, JSON, Parquet), Impersonation – JDBC/ODBC drivers, Multiple data type support – Queries on File, on Hive tables and views – Queries on HPE Database tables and secondary indexes – Schema-less ANSI-compliant distributed SQL query engine, SQLLine		Yes
Hadoop Ecosystem – Apache HBase, Hive, Hive Query Engine, Hue, Yarn – Apache MapReduce v2 (MapReduce v1 is not supported) – HPE Data Fabric DB OJAI Connector for Apache Hive		Yes
Apache Airflow		Yes
Apache NiFi		Yes
Apache Zeppelin		Yes
Apache Livy		Yes
Apache Tez		Yes
Apache Flink		Yes
Model Context Protocol (MCP) Server		Yes
Polaris Catalog		Yes

Consume IT on your terms

GreenLake brings the cloud experience directly to your apps and data wherever they are -the edge, colocations, or your data center. It delivers cloud services for on-premises IT infrastructure specifically tailored to your most demanding workloads. With a pay-per-use, scalable, point-and-click self-service experience that is managed for you, GreenLake accelerates digital transformation in a distributed, edge-to-cloud world.

- Get faster time to market
- Save on TCO, align costs to business
- Scale quickly, meet unpredictable demand
- Simplify IT operations across your data centers and clouds

Managed services to run your IT operations

GreenLake Management Services provides services that monitor, operate, and optimize your infrastructure and applications, delivered consistently and globally to give you unified control and let you focus on innovation.

HPE software services

HPE Hybrid Cloud Professional Services provide expert consulting, implementation, and support to help organizations harness the power of HPE's Data Fabric platform. These services are designed to simplify data management, accelerate analytics, and enable seamless data integration across hybrid, edge, and multi-cloud environments. The goal is to help businesses unlock the value of their data for improved decision-making and operational efficiency. HPE Hybrid Cloud Professional Services are tailored to help businesses modernize their data ecosystems, improve operational efficiency, and unlock the full potential of their data assets. Whether for analytics, AI, IoT, or hybrid cloud management, these services provide the expertise and support needed to succeed in today's data-driven economy.

Key Offerings include:

Assessment and Planning

- Data Architecture Review: Evaluate current data systems and recommend an optimal architecture for implementing HPE Data Fabric.
- Migration Planning: Create a roadmap for transitioning from legacy systems to HPE Data Fabric with minimal downtime.

Design and Implementation

- Data Fabric Deployment: Design and implement an HPE Data Fabric solution to meet specific business needs.
- Data Integration: Enable seamless integration of disparate data sources, including cloud, on-premises, and edge systems.
- Hybrid and Multi-Cloud Enablement: Configure the HPE Data Fabric to support hybrid and multi-cloud architectures for enhanced flexibility.

Optimization, Modernization & Migration

- Performance Tuning: Optimize HPE Data Fabric configurations for better performance and scalability.
- Modernization Services: Upgrade existing data systems to an HPE Data Fabric solution to improve efficiency and reduce costs.
- Advanced Analytics Support: Design, develop and deploy pipelines and tools for real-time and batch analytics to empower data-driven insights.
- Migration: Assist with migrating from legacy systems or other platforms to HPE Data Fabric.

Service and Support

HPE Hybrid Cloud Platform Services

- Proactive Monitoring and Support: 24X7 monitoring and proactive management to ensure high availability and performance.
- Capacity Planning: Support for scaling storage and processing capabilities as business needs grow.

HPE training

HPE data and analytics training delivers the skills and expertise for success with Big Data, containers, AI/ML and deep learning—including HPE Data Fabric. Live virtual instructor-led and eLearning courses are available.

- Data and Analytics Courses
- HPE Learn On-Demand

HPE Support Services are sold by Hewlett Packard Enterprise and Hewlett Packard Enterprise Authorized Service Partners:

- Services for customers purchasing from HPE or an enterprise reseller are quoted using HPE order configuration tools.
- Customers purchasing from a commercial reseller can find HPE Support Services at https://ssc.hpe.com/portal/site/ssc/

Ordering Information

HPE Data Fabric File and Object Store subscriptions include an electronic license to use on one terabyte of storage, Customer's license capacity under management is defined in the Additional Licensing Authorizations (ALA) for

Notes:

- Each subscription SKU includes a term license to use on one core plus the associated 24x7 HPE Tech Care support.
- HPE Data Fabric is licensed by the number of unique cores available to the kernel in the OS on which the HPE software is directly installed, regardless of the number of threads in each core.

All HPE Data Fabric Software subscriptions include support with 24x7 coverage.

HPE Services HPE Tech Care Service

HPE Services HPE Tech Care Service is the operational service experience for HPE products. Tech Care goes beyond traditional support by providing access to product specific experts, an AI driven digital experience, and general technical guidance to not only reduce risk but constantly search for ways to do things better. HPE Services HPE Tech Care Service has been reimagined from the ground up to support a customer-centric, AI driven, and digitally enabled customer experience to move your business forward.

HPE Data Fabric Software products are covered by HPE Services HPE Tech Care Service Essentials, providing 24x7 global support coverage.

Notes:

- The following end user information is required at the time of order to receive the electronic license:
  - End-user organization name
  - End-user organization address
  - End-user contact name
  - End-user email address
  - Reseller and distributor organization name and address

HPE Data Fabric Software Subscription Options

SKUs are available in HPE's new BRIM subscription ordering system as well as the traditional S4 ordering system. Please work with the Sales Team to determine which is the best system for placing your order.

BRIM

Description	SKU
HPE Data Fabric Software File and Object Store E-LTU	R9J44AAE
HPE Data Fabric Software Analytics for Apache Spark E-LTU	R9J45AAE
HPE Data Fabric Software Analytics Premium Package E-LTU	R9J47AAE

S4

Description	SKU
HPE Data Fabric Software File and Object Store 1-year 24x7 E-LTU	S0M33AAE
HPE Data Fabric Software File and Object Store 2-year 24x7 E-LTU	S0M52AAE

Configuration Information

Description	SKU
HPE Data Fabric Software File and Object Store 3-year 24x7 E-LTU	S0M34AAE
HPE Data Fabric Software File and Object Store 4-year 24x7 E-LTU	S0M35AAE
HPE Data Fabric Software File and Object Store 5-year 24x7 E-LTU	S0M36AAE
HPE Data Fabric Software Analytics for Apache Spark 1-year 24x7 E-LTU	S0M37AAE
HPE Data Fabric Software Analytics for Apache Spark 2-year 24x7 E-LTU	S0M38AAE
HPE Data Fabric Software Analytics for Apache Spark 3-year 24x7 E-LTU	S0M39AAE
HPE Data Fabric Software Analytics for Apache Spark 4-year 24x7 E-LTU	S0M40AAE
HPE Data Fabric Software Analytics for Apache Spark 5-year 24x7 E-LTU	S0M41AAE
HPE Data Fabric Software Analytics Premium Package 1-year 24x7 E-LTU	S0M47AAE
HPE Data Fabric Software Analytics Premium Package 2-year 24x7 E-LTU	S0M48AAE
HPE Data Fabric Software Analytics Premium Package 3-year 24x7 E-LTU	S0M49AAE
HPE Data Fabric Software Analytics Premium Package 4-year 24x7 E-LTU	S0M50AAE
HPE Data Fabric Software Analytics Premium Package 5-year 24x7 E-LTU	S0M51AAE

Learn More:

- HPE Learn On Demand: https://learn.software.hpe.com/
- HPE Data Fabric Software: hpe.com/DataFabric
- Support & documentation: hpe.com/support/hpesc
- User documentation:
  https://docs.ezmeral.hpe.com/
  https://docs.ezmeral.hpe.com/datafabric-customer-managed/79/

Summary of Changes

Date	Version History	Action	Description of Change
04-May-2026	Version 14	Changed	This version represents a major content refresh aligned to HPE Data Fabric Software v8.1, with emphasis on AI-ready data fabric positioning, metadata intelligence, and event-driven architectures. Key Components information was updated.
03-Nov-2025	Version 13	Changed	Overall updates have been made to reflect brand updates across all aspects of the product.
28-Jul-2025	Version 12	Changed	Update survey link.
07-Apr-2025	Version 11	Changed	Overview section was updated
17-Mar-2025	Version 10	Changed	QS name updated
02-Dec-2024	Version 9	Changed	Overview, Standard Features, Service and Support and Configuration Information sections were updated.
07-Nov-2022	Version 8	Changed	Overview, Standard Features, Service and Support and Configuration Information sections were updated.
15-Aug-2022	Version 7	Changed	Service and Support and Configuration Information sections were updated.
01-Aug-2022	Version 6	Changed	Service and Support and Configuration Information sections were updated.
02-May-2022	Version 5	Changed	Overview and Configuration Information sections were updated.
15-Nov-2021	Version 4	Changed	Service and Support section was updated
06-Apr-2021	Version 3	Changed	Overview and Configuration Information sections were updated.
01-Mar-2021	Version 2	Changed	Overview, Standard Features, Service and Support and Configuration Information sections were updated.
01-Feb-2021	Version 1	New	New QuickSpecs

Copyright

Have feedback on QuickSpecs? We're listening

Chat now

a50002556enw - 16711 - Worldwide - V14 - 04-May-2026
HEWLETT PACKARD ENTERPRISE
Hpe.com