

ACCELERATE ANALYTICS WITH HPE EZMERAL RUNTIME ENTERPRISE AND STARBURST ENTERPRISE

Make better decisions with fast access to your data

Data analytics means understanding trends, extracting patterns, and creating information that can be used to drive smarter business practices. Data helps companies reduce costs, increase competitive advantages, and make better, faster decisions. Data analytics improves business outcomes and can drive performance throughout an organization.

Companies face many hurdles in achieving this outcome because their data is scattered and siloed. To access and analyze data, many companies rely on traditional data warehouse products with outdated, monolithic solutions that breed inefficiency and ultimately cannot help business analysts run fast analytics on all their data. This prevents the business from making better and more timely decisions to improve their company's performance.

HPE Ezmeral Runtime Enterprise makes it easy to build industry-scale data analytics and artificial intelligence (AI)/machine learning (ML) applications since it contains significant analytics advances with Apache Spark and MLflow integration. HPE Ezmeral Runtime Enterprise also includes a multitude of other easy-to-use features to simplify app modernization, data and model collaboration, policy management, and runtime security. A variety of data access options are also supported including native access within the HPE Ezmeral Runtime Enterprise across tenants, external data in data lakes or object stores, and from edge to cloud with the embedded HPE Ezmeral Data Fabric.

Hewlett Packard Enterprise is partnering with Starburst Enterprise to help companies build distributed cloud applications that can easily and quickly access and analyze both traditional data sources, as well as new data sources located everywhere—in data centers, cloud, and even vendor environments. Using Starburst Enterprise on HPE Ezmeral

Runtime Enterprise, companies can develop modern cloud applications that can access and analyze data wherever they reside, with any business intelligence (BI) tool, to run interactive analytical queries.

Starburst Enterprise is a fully supported, production-tested, and enterprise-grade distribution of the open-source Trino (formerly known as Presto SQL) MPP SQL query engine. When combined with enterprise BI, Starburst Enterprise serves as a data consumption layer to decrease time to value by providing HPE Ezmeral Runtime Enterprise data analytics with speed at scale. Starburst supplements BI dashboards and reports by providing a secure single point of access to a wealth of new data without the need to centralize it first. This allows users to do more without extra help from IT or data engineering.

KEY CAPABILITIES OF STARBURST ENTERPRISE

Analyze anything: Allow your data scientists to interactively query any data source, from traditional warehouses to cloud data lakes, through their favorite BI tools (for example, Tableau, Qlik, Superset, and such). Trino's flexible architecture allows you to perform analytics federated across multiple data sources at the same time.

Separation of compute and storage: Starburst Enterprise can query data wherever it lies, allowing companies to separate storage and compute during their HPE Ezmeral Runtime Enterprise deployment. There is no need for companies to move their data, as companies can provision HPE Ezmeral Runtime Enterprise compute resources to their exact needs, scaling up or down based on analytics demand, which results in significant cost savings.

Solution brief

ETL workloads: Starburst Enterprise for Trino is ANSI SQL compliant for supporting commands such as create table and insert. It functions as the SQL engine for extract, transform, and load (ETL) jobs, providing a single platform for both query and migration needs. For example, archive data from an Apache Hadoop cluster could be moved to HPE Ezmeral Data Fabric, allowing federated Trino queries against that data, as well as data from other sources that are not ready for migration.

Rapid access: Starburst Enterprise is a high-performance distributed SQL query engine. Companies can query data in cloud data lakes as quickly as they can query on other storage platforms, allowing them to save on storage without sacrificing insights. Data from disparate sources or stored in different formats can be joined and operated on as if it were in the same database.

Global security and access controls: The Starburst Enterprise software supports authentication and authorization of data sources, along with tables, rows, and columns within those sources, that each user can access. Starburst also supports encryption for intra-cluster data streaming and data at rest.

KEY CAPABILITIES OF HPE EZMERAL RUNTIME ENTERPRISE

On-premises and/or on the public cloud: Can be deployed on-premises, in the public cloud, or in a hybrid environment that includes both public cloud and on-premises resources.

Make the right purchase decision.
Contact our presales specialists.



Chat



Email



Call



Get updates

Hewlett Packard Enterprise

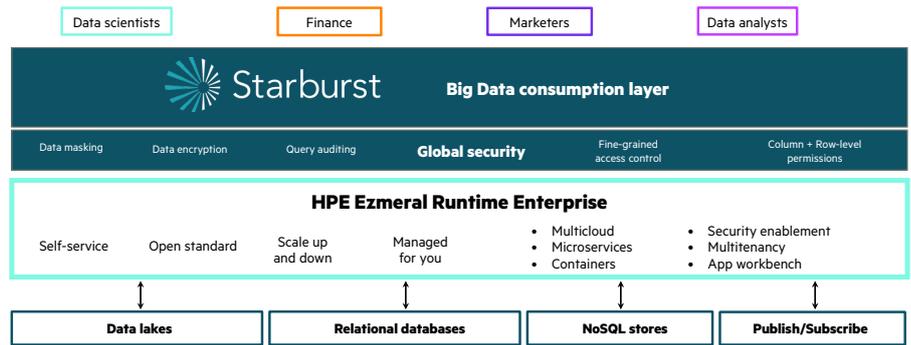


FIGURE 1. Starburst Enterprise data consumption layer on HPE Ezmeral Runtime Enterprise

Integrated platform for AI, Big Data analytics, and data science:

Is purpose-built for Big Data (large scale) and/or AI applications—including data science, analytics, ML, and deep learning (DL)—using enterprise-grade security, networking, and support for a variety of local and remote storage options.

Multitenancy and enterprise-grade security model:

Integrates with enterprise LDAP and Active Directory authentication systems. Administrators can create groupings of users and resources that restrict access to jobs, data, or clusters based on department membership and/or roles. The result is an integrated, safe, and multitenant infrastructure.

Leverage existing storage and compute resources:

Repurposes existing large-scale data deployments. Multiple storage protocols are supported (NFS, HDFS, HDFS with Kerberos, and others). The platform itself can be deployed on physical and virtualized infrastructure, along with major public clouds (Amazon Web Services, Google™ Cloud Platform, and Microsoft Azure).

Superior performance: Provides storage I/O optimizations to deliver data to Big Data and AI applications without the penalties commonly associated with virtualization and containerization or the separation of compute and storage resources. The CPU cores and RAM in each host are pooled and then partitioned into virtual resource groups based on tenant requirements.

WHY STARBURST ENTERPRISE ON HPE EZMERAL RUNTIME ENTERPRISE?

With Starburst Enterprise, HPE Ezmeral Runtime Enterprise users can:

- Modernize, update, migrate, and move data at their own pace—without forced data migrations
- Have a single point of access for distributed and disparate data sources
- Reduce time to insights without having to build complex and time-consuming ETL pipelines

With Starburst Enterprise on HPE Ezmeral Runtime Enterprise, data architects, BI analysts, and AI/ML teams have fast, reliable, and stable access to the data they need to do their best work—no matter where that data resides. Starburst Enterprise delivers data to these HPE Ezmeral Runtime Enterprise users quickly and securely, without the need for ETL or expensive data warehouses.

LEARN MORE AT
hpe.com/us/en/software/marketplace/starburst.html

© Copyright 2021 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Google is a trademark of Google LLC. Active Directory, Azure, and Microsoft are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. All third-party marks are property of their respective owners.

a50004198ENW, Rev. 1