



Unified analytics for hybrid deployments

Today's modern business needs to be rooted in data analytics; however, businesses are finding that data silos, disconnected data teams, and a lack of standards are making it difficult to achieve the insights required to spark that next innovation. Solving these problems requires rotating away from traditional data and analytics solutions to modern environments without impacting security, compliance, or adding more complexity.

The future isn't in one environment, and it isn't all cloud. The future is cloud-native but optimized for hybrid deployments, which means you need a secure, simple, and unified data platform experience that doesn't require moving all your data into one location. It should use automation to accelerate data engineers and analysts' processes. Lastly, it should deliver a cloud-native experience without requiring a jump from one proprietary platform to another or expose your business to the complexity of wide-open systems.

HPE Ezmeral Unified Analytics

Solve these challenges with a unified data experience that allows teams to securely connect to data where it exists today without disrupting existing data access patterns. From day 1, HPE Ezmeral Unified Analytics connects to datastores, data lakes, clouds, files, and objects exactly where they exist today in a controlled manner. Deploy on-premises, hybrid, or multicloud, or run with HPE GreenLake.

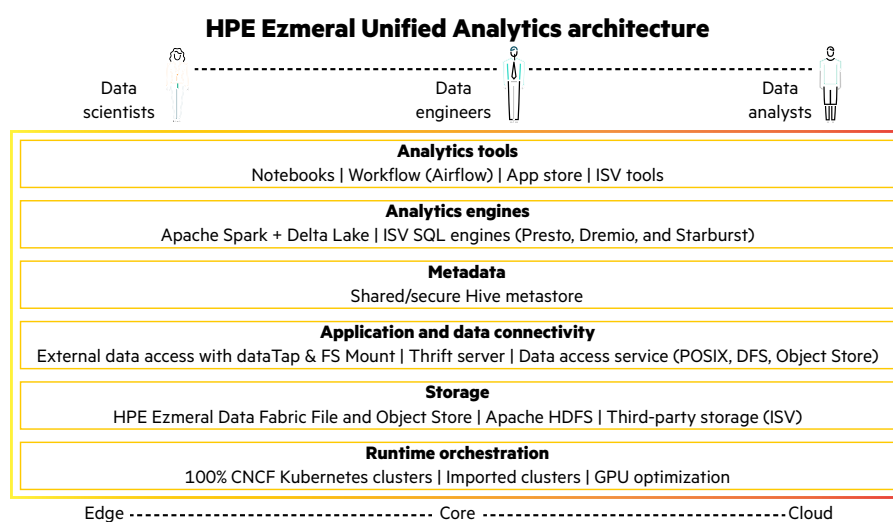


Figure 1. The modern platform connects analytics and data science teams to your data

The foundation of this solution is HPE Ezmeral Runtime, which delivers container orchestration of both cloud-native and non-cloud-native applications with persistent data. This 100% Kubernetes hybrid platform spans edge to cloud allowing enterprises to modernize their application deployments on bare-metal or virtual machines (VMs) on-premises, multiple clouds, and at the edge.

HPE Ezmeral Data Fabric File and Object Store provides a unified data experience across data lakes, files, objects, streams, and databases. It provides a persistent datastore without migrating or moving data, refactoring applications, or impacting service levels, compliance, and security. The result is increased elasticity, efficiency, and agility for your data-driven enterprise.

App-store accelerates development

Data engineers, data analysts, and data scientists each have a different point of view when it comes to analytics, which translates into them using different tools. Code-first data scientists are using Jupyter Notebooks while citizen data scientists prefer packaged solutions. All these tools are scattered across different systems limiting their productivity.

The consistent cloud-like user experience in HPE Ezmeral Unified Analytics offers single-click deployment of opinionated stacks or certified ISV solutions from a marketplace. The app workbench allows these personas to bring or build their own open-source stacks using the built-in app workbench. This unified approach to libraries, tools, and runtime engines helps to break down data silos without disrupting data access patterns or productivity.



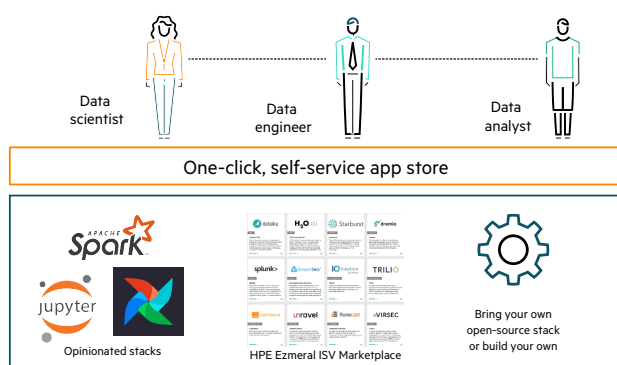






Figure 2. Unified approach to tools and runtime engines breaks down silos

Native integration of Apache Spark Kubernetes operator

Apache Spark (Spark) is an open-source data processing framework that speeds up processing tasks on very large data sets and can also distribute data processing tasks across multiple computers. It removes some of the programming burdens from the shoulders of developers and data science teams with an easy-to-use API that abstracts much of the mundane work of distributed computing and data analytics.

When using traditional data and analytic tools, data and scientific teams face a series of challenges some of which include the need to submit jobs through a choice of methods (REST, CLI), dynamic scaling using CPU/GPU acceleration, and the ability to deploy across on-premises, cloud, and edge.

Native-integration of the open-source Apache Spark Kubernetes operator into HPE Ezmeral Unified Analytics boosts data and scientific team productivity with:

-  Using multiple Spark versions
-  NVIDIA® GPU acceleration
-  Unified interface with Delta Lake to run batch and real-time workloads
-  Native integration with Airflow, MLflow, Kubeflow plus a vibrant certified HPE marketplace to enable the use of the most common analytics tools and engines from ISVs

Using industry-standard analytic and Big Data benchmarks, Hewlett Packard Enterprise completed testing using NVIDIA RAPIDS Accelerator for Apache Spark. The results of this test indicate that HPE Ezmeral accelerates Spark analytic workloads by 29x.¹

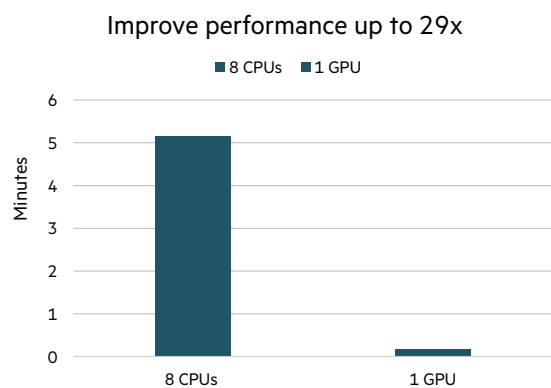


Figure 3. CPU/GPU testing for eight CPUs and one NVIDIA GPU

¹ We standardized our testing models based on **industry-standard** analytic and Big Data benchmarks. This benchmark was based on leveraging Kubernetes pods and storage managed by HPE Ezmeral, NVIDIA A100 Tensor Core 40GB GPUs and HPE ProLiant DL385 servers. All tests were conducted in a laboratory setting to simulate real-world conditions. User experiences may differ from those in our test scenarios.



HPE Ezmeral Unified Analytics goes beyond native Spark integration by adding additional functionality that delivers a true enterprise-ready data analytics solution:

- An elastic platform to scale Spark workloads across on-premises, cloud, and edge deployments
- Zero trust security
- Multitenant control plane
- Dynamic platform partitioning
- Full isolation and control
- Prebuilt ISV or bring/build your own data analytic tools
- 24x7 enterprise support

HPE Ezmeral Unified Analytics helps to increase data and scientific team productivity by working with the following open-source tools:



Run batch and real-time jobs from a unified interface
Offer full ACID transactions for consistent and reliable data



Shared back-end metadata store
Secured metadata access
Improved resource management (no separate database to manage)





Enables business intelligence apps to interface with Spark

Summary

Increase data and scientific teams' productivity with unified data access and data experience without changing applications or usage patterns. A self-service app store allows teams to deploy the popular open-source data science tools, libraries, and frameworks out of the box or integrated Jupyter Notebooks backed by 24x7 support from HPE.

HPE Ezmeral Unified Analytics unifies your data globally and makes it available to analytic teams—whether the data is at the edge, on an enterprise data warehouse on-premises, a cloud data lake, or on other cloud platforms. It unifies the analytic experience from SQL to the most demanding data engineering, analytics, and data science workflows for your teams. The cloud-native experience and built-in automation connect users and their tools to the right data, not the underlying data infrastructure. By using open-source tools and frameworks and enterprise-grade support, you get the flexibility to confidently mix and match open-source components.

Make the right purchase decision.
Contact our presales specialists.



 **Get updates**

Learn more at

hpe.com/us/en/software/ezmeral-unified-analytics.html

Visit **HPE GreenLake**