


Hewlett Packard
Enterprise



FUEL EDGE-TO-CLOUD DIGITAL TRANSFORMATION

Data means nothing without insights. See how global enterprises capitalize on their data with HPE Ezmeral—an open data and analytics platform purpose-built for the enterprise and all your data.

HPE EZMERAL


e-book

CONTENTS



MAKE DIGITAL TRANSFORMATION REAL

Modernize apps, unlock insights,
and automate ops



MAKE DIGITAL TRANSFORMATION REAL

HPE Ezmeral advances digital transformation by shifting time and resources from IT operations to business innovation. See how our customers use HPE Ezmeral to:

- **Modernize data-centric apps:** Accelerate time to value by deploying Kubernetes at scale with integrated persistent data storage for app modernization on bare metal or virtual machines (VMs)—in the data center, on any cloud, or at the edge.
- **Unlock insights:** Harness data and speed insights by operationalizing the end-to-end process of building data pipelines. Bring DevOps agility to the machine learning (ML) lifecycle and deliver a unified data fabric from edge to cloud.



TIER-1 BANK

Build a bridge to multicloud
for big data analytics



TIER-1 BANK

Build a bridge to multicloud for big data analytics

Situation

The bank needed to meet the demands of a growing number of analytics users. The existing situation wasn't scalable and presented a number of new problems:

- 9–18 month provisioning time.
- ~30% CPU utilization.
- Massive data duplication per data scientist.
- Couldn't keep up with data and had to erase 2 TB per day.
- General Data Protection Regulation (GDPR) compliance was in danger.
- "Cloud-first" corporate mandate didn't mesh with security requirements.

Solution design principles

- Architecture that scales independent of compute and storage.
- Data scientists can self-serve on demand via a request portal.
- The entire process is fully automated to reduce provisioning bottlenecks.
- Sandbox environments can get read-only access to the production data lake to protect production system integrity.
- Scratch space on the production system grants full visibility, trackability, and auditability.
- A single source of data provides GDPR compliance and eliminates redundant copies.
- A 100% cloud-ready solution.

Results

- Infrastructure utilization and scaling is significantly improved.
- Self-service and automated provisioning accelerates the entire data analytics lifecycle.
- Data copies are eliminated.
- Able to apply a common security model aiding in GDPR compliance.
- The solution bridges the data center and public cloud requiring only two lines of code to spin up environments in AWS.



>£6 million

infrastructure savings over three years

>1200x

faster time to value

2 lines of code

to change for AWS portability


~30 minutes

to provision, instead of 9–18 months



LARGE U.S. BANK

Eliminates data blind spots
with containerized Splunk



Large U.S. bank

Eliminates data blind spots with containerized Splunk

Situation

More than half of enterprise data can't be loaded and processed fast enough to be used in analytics.¹ As a result, there are huge blind spots in IT security and operations:

- Data is exceeding IT's ability to index, store, and analyze it.
- Low CPU utilization results in over-deployment to keep pace with data growth.
- Data centers are nearing maximum capacity as IT struggles to scale out data ingestion, processing, storage, and analysis.
- A growing backlog of data and security insights is forcing organizations to find ways to optimize their delivery and consumption of Splunk analytics.

Solution approach

- Turnkey cloud solution delivered by HPE GreenLake and powered by HPE Ezmeral Container Platform.
- Use HPE GreenLake to provide the platform-as-a-service (PaaS) solution managed by HPE up through the container and storage layer.
- Solve the blind spot with application, architecture, and consumption model modernization.
- Use a unique PaaS solution that allows for independently scaling indexers and search heads: up, down, and out.
- Get the power to leverage critical data for a full view of the IT security landscape.
- Deliver more value from existing Splunk investments with efficient, right-sized deployments.

Results

- Workload-optimized infrastructure helps eliminate bottlenecks at scale.
- Containerized software brings agility and scale to the containerized Splunk operator.
- Shrinks the infrastructure footprint and significantly lowers TCO.
- Adds new use cases deploying new indexers and search heads in minutes.
- HPE GreenLake enables flexible pay-for-use consumption with no upfront outlay and the ability to scale up or down quickly and take advantage of burst options.
- Fully managed, on-premises cloud experience with no patching, performance tuning, or maintenance required and no need to hire hard-to-find Kubernetes skills.
- Containerized solution is multicloud / Hybrid Cloud ready.

¹ Splunk, [The State of Dark Data](#), 2019.



e-book

Up to

17x

indexer throughput improvement

Up to

10x

reduction in infrastructure

Up to

12 indexers

per host

Up to

8.7 TB

ingest per host per day




Read the [Splunk white paper](#).

e-book

MERCEDES BENZ AG

Intelligent, autonomous
cars make driving safer



MERCEDES BENZ AG

Intelligent, autonomous cars make driving safer

Situation

European automotive company, Mercedes Benz AG, was seeking to accelerate development of autonomous driving functions. To do this, it needed a data platform that could:

- Collect massive amounts of data from car cameras and sensors.
- Make this data available to developers and data scientists for analysis.
- Train algorithms that enable cars to make smart decisions in real time.
- Use data to identify anomalies and refine autonomous driving systems.
- Access and share data globally with high performance on a massive scale.
- Minimize replication and avoid data duplication to optimize hardware resources.

Solution approach

The company deployed an HPE Ezmeral Data Fabric platform that can:

- Efficiently manage and analyze millions of miles of vehicle test data.
- Extract data from test vehicles quickly and efficiently.
- Make data available to developers and data scientists around the world.
- Enable direct access to data, including from legacy systems.
- Leverage deep learning techniques to train algorithms.
- Scale storage easily and effectively.
- Deliver high availability and data resiliency.
- Enable developers to run containerized applications.

Results

- Accelerates the development of autonomous driving systems.
- Enhances vehicle and road safety through improved driving systems.
- Delivers economic advantage with a platform that doesn't require extra investment in tools.
- Optimizes compute resources by accommodating multi-tenancy on its data platform.
- Lowers barriers to entry for additional projects.



5-20 TB


of data generated each day

Multi-PB

system is ready to scale

10 million miles

of driving data in a training set



From starting as a small project involving about 100 users, it now has several hundred data scientists and developers using its clusters without experiencing system breakdowns. Not only does this optimize resources and increase compute power for each data scientist, it also lowers barriers to entry for additional projects.


—HPE Case Study, [“Accelerating autonomous car development with ready access to global data”](#)

Read the [case study](#).



FORTUNE 50 TELCO

Security solution detects anomalies and attacks



FORTUNE 50 TELCO

Security solution detects anomalies and attacks

Situation

One of the key tasks of the chief security officer at this Fortune 50 telco is to detect anomalies and attacks. However, the company's security platform was no longer meeting their requirements in the following areas:

- **Vendor disruption:** Sensitive data needed to stay on premises, but one of the core technologies would soon be offered only through the cloud.
- **Scalability:** The multi-vendor solution was cumbersome to scale.
- **Performance:** Deleting data, ad hoc queries, and major compactions all caused significant application delays.

Solution approach

- Enterprise storage with promotable mirrors, replication, and snapshots.
- Performance to support business initiatives and continued data growth.
- Disaster recovery and high availability, including the ability to seamlessly fail over container location database (CLDB) services.
- Highest performance per server scale tested versus all other distributed NoSQL vendors.

Results

- Ingests almost 30 billion records per day.
- Gets fast, consistent response times from thousands of queries per hour while ingesting data.
- Integrates with current homegrown applications.
- Reduces the overall size of the scale-out database cluster.
- Enables real-time queries and analytics.



e-book

Thousands

of queries per hour

Real-time

queries

30 billion

records ingested per day

5.7 PB

initial footprint



HPE ON HPE

Analytics as a service empowers
a global business team



HPE ON HPE

Analytics as a service empowers a global business team

Situation

As a multinational enterprise, HPE had thousands of scattered point solutions for analytics and reporting, with little data governance and certification, and limited end-to-end visibility. Data scientists and team leaders faced a tedious, manual process to pull together data to drive the business. HPE was looking to:

- Improve the speed, agility, and economics of analytics services.
- Break down data silos that were impeding insights.
- Reduce the costs and uncertainty of bringing data in from multiple sources.
- Build a true analytics platform—a single source connecting data from edge to core that every group in the enterprise can draw from to make informed business decisions.
- Provide a single version of the truth via validated, quality data.

Solution approach

- Build an agile data storage solution to support a hybrid data and analytics-as-a-service platform.
- Scale compute and storage building blocks independently.
- Deliver a cloud experience, enabling users to spin up customized analytics projects in minutes.
- Allow data to be ingested from any source and fed into a data lake where it is consumable by business groups.
- Gain the ability to easily add technologies such as GPU nodes, memory-centric nodes, and tiered storage nodes in the future.

Results

- Scales storage and compute independently, improving agility to respond to new workloads, and drastically reduces the cost of delivering global analytics services.
- Achieves near-real-time ingestion and processing of data from any source.
- Delivers a self-service cloud experience, allowing users to spin up clusters in a matter of minutes across a hybrid environment.
- Empowers global teams to make data-driven business decisions.
- Provides one version of the truth, from edge to core, for more consistent reporting.
- Open ecosystem allows the business to bring their own stack and innovate at the speed of open source to get the most value from their data.



One

version of the truth, from edge to core

10 million

orders per day in real time

Minutes

to spin up clusters

Open

ecosystem



“Data scientists are distributed across the enterprise. We can empower them with access to analytics as a service wherever they are, and do that in a manner that’s fast, cost-effective, and architecturally sound—that all matters.”


—Dave Carlisle, Chief Technology Officer, HPE IT

Read the [case study](#).



OROCK TECHNOLOGIES

Deploy FedRamp-certified
data science as a service



OROCK TECHNOLOGIES

Deploy FedRamp-certified data science as a service

Situation

ORock is a Hybrid Cloud service provider that was seeking to introduce capabilities for data science, Kubernetes orchestration, and big data. The solution needed to:

- Deliver high-performance processing for demanding workloads such as artificial intelligence (AI) and machine learning (ML).
- Meet the stringent requirements of the U.S. federal government.
- Enable customers to run workloads as a service.
- Provide resources that match customer needs without extensive customization.
- Offer the flexibility of an open source platform.
- Reduce costs while enhancing performance.
- Optimize operational efficiency.

Solution approach

- Data-science-as-a-service platform providing optimized operational efficiency, business agility, and support for sensitive data and workloads
- Comprehensive cloud infrastructure and storage
- End-to-end hybrid service via HPE Ezmeral Container Platform and HPE GreenLake

Results

- Brings new capabilities to market faster with deeper data insights
- Provides multicloud capabilities with superior security, performance, predictability, and control
- Delivers end-to-end hybrid service via HPE Ezmeral Container Platform and HPE GreenLake
- Meets strict FedRamp regulations for security, compliance, and performance



e-book

Self-service

data access

FedRamp

certified

Low cost, zero egress

for technical and operations teams

End-to-end

security



“Our new data-science-as-a-service solution is fundamental to how organizations can consume AI and ML technologies at speed and scale. This is a much easier way for IT teams to experience popular AI and ML tools that they love, with the most secure cloud in the market utilizing HPE ProLiant Gen10 servers.”

—Gregory Hrcir, President and CEO, ORock Technologies, Inc


Watch the [video](#).



e-book

DXC.TECHNOLOGY

Accelerate advanced
driver-assisted systems



DXC.TECHNOLOGY

Accelerate advanced driver-assisted systems

Situation

DXC operates one of the largest data fabrics, with a robotic solution that allows manufacturers to accelerate their autonomous vehicle programs. DXC was seeking a partner that could help the company:

- Stretch their data fabric cost-effectively with speed and scalability.
- Deploy a robust file system that could handle hundreds of petabytes of data daily.
- Integrate diverse data sources.
- Simplify data management, security, and compliance.
- Distribute large amounts of data to multiple global sites to give data scientists direct data access.
- Reduce research and development (R&D) time and costs.

Solution approach

- HPE Apollo Systems, HPE Edgeline Converged Edge Systems, and HPE Ezmeral Data Fabric providing speed, scalability, and the ability to geo-distribute data
- A single optimized platform that spans from the car to the data lake to enable developers
- Open architecture that supports multiple protocols and allows integration with whatever application deployment or virtualization models are used to plug into the data fabric
- Data fabric that can stretch to hundreds of petabytes and more than a million data points
- Secure global access to data worldwide

Results

- Eliminates additional point solutions
- Benefits from a robust file system that scales
- Achieves multiprotocol support
- Reduces time and cost for R&D
- Gains direct access to distributed global data



100s of PBs

and more than a million data points

More than

1000 cores

and 250 CPUs

75% reduction

in R&D time

1.5 TB

of data generated per car, per day



“RoboDrive is a game-changer in developing autonomous driving vehicles. It’s an end-to-end hyperscale machine learning and AI platform based on the HPE Ezmeral Data Fabric, which is used by some of the largest manufacturers in the world for development of their autonomous driving algorithms.”

—Dragan Rakovich, Director, DXC Luxoft Analytics & Engineering, Data to Insights GP

Watch the [video](#).



BIDTELLECT

Crunch big data for smarter digital advertising



BIDTELLECT

Crunch big data for smarter digital advertising

Situation

The Bidtellect mission requires a solution that can support context-driven optimization, engagement measurement, and creative services. To do this, the company needed a data platform that could help them address multiple issues:

- Major performance issues with their Cloudera environment, including application compatibility issues with Kafka.
- Support issues that left Bidtellect to fix problems on their own.
- High infrastructure and licensing costs.
- Phenomenal growth requiring a partner and solution that can scale with support.
- Space and budget constraints require the ability to process more data in their existing petabyte scale footprint.

Solution approach

- Use a single platform to plan and execute digital advertising without using cookies.
- Modernize the application stack and reduce the cost of hardware required by Cloudera.
- Build a foundational layer for large-scale custom applications with a global file system, data lake for analytic applications, and software-defined storage for large-scale container environments.
- Provide a trusted data repository for data scientists, developers, and IT with access controls.
- Scale transparently as the business grows, without downtime or IT intervention.
- Provide visibility and governance across all applications and data through an easy and flexible user experience.

Results

- Achieves higher optimization for compute and storage
- Saves cost of Kafka licenses using HPE Ezmeral Data Streams
- Enables moving from Cloudera Impala onto Apache Drill, creating a fully secure data fabric environment
- Gets the support they need as the company grows, through HPE Financial and Professional Services
- Gains a trusted partner for digital transformation and reliable, excellent customer support



Real-time

data analytics

Milliseconds

to insights

Excellent

customer support

Global access

to data



“HPE Ezmeral Data Fabric helps our customers receive more value out of their media, allowing us to blow the competition out of the water. The combination of HPE Ezmeral Data Fabric and our optimization technology allows us to review past advertisement placement behavior and performance, which was not possible with our previous vendor.”

—Mike Conway, Chief Technology Officer, Bidtellect

Watch the [video](#).



CARESTREAM

Enhance patient care
with AI as a service



CARESTREAM

Enhance patient care with AI as a service

Situation

Carestream, a leading global provider of medical imaging systems, needed to accelerate the execution of their strategy to differentiate services by using mobile and AI technologies. However, the company:

- Needed to improve accuracy, automate radiology workflow processes, and reduce cycle times
- Lacked data science and MLOps best practices required to use AI algorithms
- Needed a scalable platform and operational model capable of spanning from the core to the edge

Solution approach

- HPE Ezmeral MLOps with HPE GreenLake provides a scalable, future-proof platform that leverages powerful AI technology to streamline, transform, and improve the medical imaging process, adding value across the entire imaging chain for both patients and providers.
- The end-to-end AlaaS platform uses HPE storage, servers, HPE Ezmeral software, and partner solutions.
- HPE will handle the total solution execution, including ongoing care, AI operations, and infrastructure management.
- NVIDIA® GPUs enable each HPE server to tackle multiple imaging and analysis tasks simultaneously.
- Time-consuming local software updates will be replaced by cloud hardware and service upgrades.
- HPE experts will help Carestream transfer to an as-a-service model.

Results

- Accelerates AI research, innovation, and adoption
- Reduces development and operational costs
- Enhances X-ray solutions and workflows for medical professionals and patients, facilitating the next generation of AI-enabled imaging devices for medical diagnosis
- Accelerates development, compliance, and operational processes such as deploying updates of new machine learning and deep learning models
- Transforms Carestream to an as-a-service business model
- Delivers a new revenue stream and business model with AI-enabled mobile diagnosis
- Enables Carestream to reduce its AI-related R&D costs, while accelerating the company's AI-powered imaging, workflow automation, and analytics innovations



e-book

Simultaneous

imaging and analytics processing

End-to-end

AlaaS platform

Edge

enabled

100,000

pieces of equipment integrated



“Carestream was already ahead of the curve in using AI for medical imaging, and this close collaboration with HPE helps us scale to support the evolving needs of our global customer base. HPE GreenLake Cloud Services enables Carestream to embrace the Hybrid Cloud so we can focus on supplying customers with superior on-premises imaging hardware, while delivering AI services and software updates through the cloud.”

—Dharmendu Damany, CTO, Carestream Health, Inc.


Watch the [video](#).



e-book

NOW IT'S YOUR TURN

Edge-to-cloud digital transformation



NOW IT'S YOUR TURN

Edge-to-cloud digital transformation

Business leaders can gain visibility and control across all clouds. HPE Managed Cloud Controls ensures compliance and cost control across multiple public cloud environments to simplify operations and optimize costs.

App developers and DevOps can deploy containers at scale with Kubernetes. HPE Ezmeral Container Platform provides an enterprise-grade solution to accelerate your app modernization efforts and run data-driven apps with persistent container storage—in your data centers, on any cloud, or at the edge.

Data science and MLOps can accelerate time to insights. HPE Ezmeral MLOps addresses the entire ML pipeline, operationalizing the end-to-end process so you can focus on business outcomes.

Data analytics and DataOps can enable data-driven apps with no movement of data. HPE Ezmeral data fabric is a globally distributed data store that simplifies management and provisioning for data-driven applications and eases movement between on-premises and cloud environments.

ITOps and admins can drive consistent service delivery and business agility. HPE automation software empowers you to focus on innovation and improving execution reliability and efficiency that results in increased agility, reduced risk, and lower costs.

Security and SecOps can speed app authentication in distributed clouds. HPE security software and open source contributions increase productivity and agility by enabling developers to build inter-app authentication into containerized apps to straddle legacy data centers and cloud-based infrastructure.



RESOURCES

Read the blog

- [HPE Ezmeral: Uncut](#)

Learn more at

- hpe.com/ezmeral
- [Interactive demos](#)

Data means nothing without insights. Don't wait to make digital transformation real by modernizing your apps, automating operations, and unlocking insights from edge to cloud.



Check if the document is available in the language of your choice.



Make the right purchase decision. Contact our presales specialists.



Chat



Email



Call



Get updates