



## HPE Discover Highlights

**By Jean S. Bozman**

HPE's Discover conference unveiled key additions to the company's HPE GreenLake as-a-service offering, which provides a consistent cloud environment across a customer's enterprise.

Originally announced in 2017, HPE GreenLake is evolving as it becomes a fast-growing platform that generates nearly \$900 million of annual recurring revenue (ARR) for HPE, which has [more than \\$27 billion in annual revenues](#).

HPE CEO Antonio Neri confirmed at the [HPE Discover conference](#) in Las Vegas that HPE is currently supporting 70 services on the GreenLake platform – and that it will add even more by year's end. The GreenLake services are accessed on a pay-per-use basis, providing subscription-based, end-to-end cloud services across SMBs, enterprises and governmental organizations. Examples include services for Microsoft SQL Server and Microsoft Azure; SAP/HANA; and high-performance computing (HPC)— currently included among the GreenLake-supported services.

Neri told the HPE Discover audience – with 8,000 attending in-person in Las Vegas and many more online – that HPE GreenLake clearly shows the trajectory of HPE's current and future investments in distributed data services for Core, Cloud and Edge computing.

As announced, the new additions include the [HPE GreenLake for Data Fabric](#) software for hybrid data discovery and metadata management; and HPE GreenLake for Private Cloud Enterprise, supporting on-premises clouds and ensuring customer control of protected data. To drive home the impact of deploying HPE GreenLake across customers' enterprises, HPE made a point of listing several "light-house" global customers, including BMW, Disney, and Home Depot. IT executives and business executives from these companies spoke at HPE Discover.

### What to Expect

HPE is moving quickly in the as-a-service space, but it will not be alone. Systems vendors, software vendors and cloud services providers (CSPs) see the opportunity, as well. The as-a-service space is becoming increasingly competitive, given customer requirements to migrate applications and data into the Cloud and the Edge, while retaining control and security of private-cloud data. HPE plans to differentiate its HPE GreenLake portfolio by adding more services; by growing its ecosystem of services partners – and by delivering those services on-demand to its global customer base.

Examples of this competition include Dell Technologies' Apex console for managing distributed data services. IBM has a Data Fabric capability for data discovery and metadata management,



storage-as-a-service and provides Red Hat Cloud Satellite for remote data management. Competition is growing, as well, in an increasingly multi-cloud world, with new partnering agreements: Cloud service providers and co-location (co-lo) firms are expected to provide enterprise-wide access to distributed data services, along with AI/ML software tools for data discovery and automation of service management.

## **Worldwide Drivers for as-a-service (aaS) offerings**

Why is the as-a-service space such an active area for growth?

Several factors are contributing to this, including:

- the pace of enterprise cloud migrations accelerated in 2021, due to the pandemic -- and the need to modernize and transform aging data-center applications.
- the rapid growth of data in the Cloud and at Edge locations. IDC expects that most of customers' net-new data will be generated in Edge locations by 2025, generating double-digit Edge growth for systems, storage and services from [2022 through 2025](#). Edge locations include remote offices, expanded retail distribution, manufacturing sites, energy-processing locations).
- customers are demanding control of their enterprise data, both for on-premises private clouds and for data repositories in the public clouds. These demands are coming first from regulated industries (finance, healthcare, oil/gas/energy), but are expected to widen globally as legacy applications and cloud-native applications reside in distributed data resources, spanning Core, Cloud and Edge locations worldwide.
- regulatory environments require customer compliance with data-privacy rules that vary across countries and regions (e.g., the series of European Union and California regulations for personally identifiable data [PII] ).
- As-a-service offers are based on vendor-based management and support, reducing the need for enterprise and SMB customers to have all the skill-sets in-house that would otherwise be needed to develop, deploy, and maintain the as-a-service capabilities.

Now, customer's mission-critical data, which used to be housed "inside" the corporate data center, is now flowing across the entire enterprise. This is happening, even though access to customers' most-protected data must be restricted to authorized users – and managed according to governmental compliance regulations.



Case-in-point: On-premises data delivery to private clouds is increasingly important to the world's largest companies, many of which operate in regulated industries, like finance, healthcare, and telecoms. HPE's GreenLake for Data Fabric capabilities, as announced at the HPE Discover conference on June 28, now in limited availability, will be generally available by year's end (2022).

## The June Announcements at HPE Discover

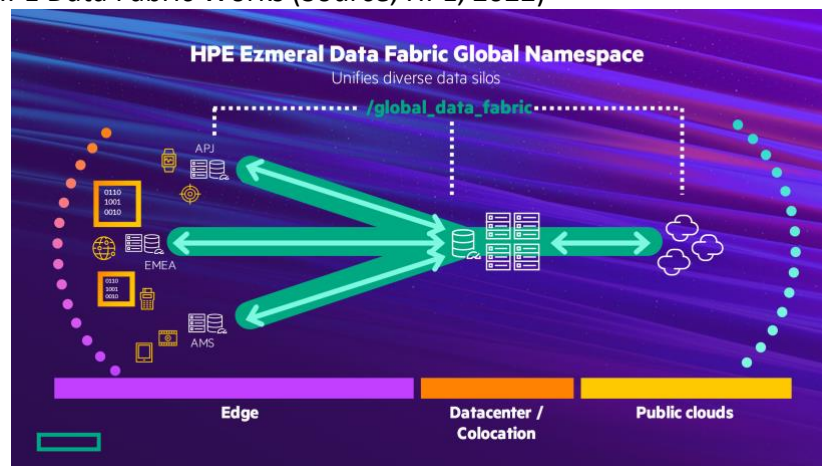
### HPE Data Fabric

HPE's June announcements unveiled more layers of the HPE GreenLake platform, including a new Data Fabric capability linking data resources in core, cloud, and edge locations. The HPE GreenLake for Data Fabric delivers a hybrid data fabric that is wrapped by HPE GreenLake hardware, software, and services.

Based on HPE Ezmeral Data Fabric, this HPE solution ingests, centralizes, and stores files, objects, NoSQL databases, along with real-time and batch streams, into a single logical data store. That is providing enterprise-wide data access to users and applications without having to copy or migrate the data. All data types share policy-driven security and data management across multiple physical clusters and locations, a technology that increases analytic team productivity and data reuse.

Importantly, the HPE Data Fabric solution leverages a unified global namespace that connects enterprise-wide data across Core, Edge, and Cloud, enabling access to hybrid data across multiple data resources. A single query point is used to access data -- no matter where the data is located (see HPE diagram in Figure 1, below) This reduces the need for IT organizations to provide unique access procedures for each user.

Figure 1: How HPE Data Fabric Works (Source, HPE, 2022)





HPE GreenLake for Data Fabric is a managed data solution that includes hardware, software, and services, eliminating installation and configuration tasks. This allows a customer's teams to focus on business outcomes, and driving new insights from enterprise data. This approach supports enterprises that have an analytic mandate to leverage insights from their rapidly growing data resources. It allows customers to manage their IT infrastructure, while avoiding the technical debt associated with analyzing multiple petabytes (PB) or exabytes (EB) of data.

## **HPE GreenLake for Private Cloud Enterprise**

HPE GreenLake for Private Cloud Enterprise is a new service that supports many styles of distributed data resources -- bare-metal deployment, VMs and container workloads – all of them managed within an extensible, shared environment. Roles and personas see “customized” views of the data resources. The Private Cloud Enterprise software works with metadata across an enterprise environment. Data types supporting this Private Cloud solution span a wide range of data resources and data silos, including customers’ block data, file data and object data.

The GreenLake private-cloud services include SAP HANA, Microsoft SQL Server, AI/ML software tools, and high-performance computing (HPC). As for cloud service provider (CSP) data, GreenLake supports access to data stored in Amazon Web Services (AWS), Microsoft Azure Cloud and the Google Cloud Platform. Multi-cloud access to multiple CSP services is supported.

Following authentication of user IDs, the roles (e.g., data analysts, business analyst, programmers, business-unit managers) gain a tailored view to the applications and data they use in their day-to-day work. This provides access to services (including SaaS services) that span many customer sites, encompassing local, remote and a variety of cross-cloud services.

Customers can use HPE GreenLake for Private Cloud Enterprise for traditional enterprise applications and for cloud-native applications that were built for use with containers and Kubernetes orchestration. All applications are accessed using cloud attributes, including self-service, scalability for applications, pay-per-use subscription models and managed services.

From an analyst perspective, the key business advantage of this private-cloud capability will be that large data repositories will remain where they are (e.g., in centralized data-center sites or in decentralized data hubs), while access is being extended to wider groups of end-users throughout the entire organization.

That makes data processing more efficient – while reducing the slow-down effects of network latency that would occur if large amounts of data – in the multi-PB range – were transferred from site to site. With broader access to distributed data, large data repositories will stay where



they are, while broader data access is extended to business-unit users and IT users who are currently working with “data hubs” across the organization

## Summary

Customers are already migrating their applications and data to storage resources in the Cloud and Edge. Additionally, they are adding cloud-native application and data. That’s why they are looking for efficient methods to pull together geographically distributed data.

As announced in June at the HPE Discover conference, HPE GreenLake uses AI/ML to find “patterns in the data” that identify usable data and support end-to-end access to data resources. That ability can be applied to view the current status of batch, transactional and real-time data across the business, for the purpose of improving business outcomes.

The new HPE Greenlake for Data Fabric unifies different data types across hybrid environments to increase data integrity, leading to more complete business insights based on cross-enterprise data. Support for a broad range of APIs, protocols, and languages supports for existing analytic infrastructure, such as data lakes and warehouses -- and Spark-based workloads and cloud-native applications.

The new HPE GreenLake for Private Cloud Enterprise supports bare-metal deployment, VMs and container workloads in a shared environment. For efficiency purposes, large data repositories stay where they are, while access is extended to business and IT users across the organization.

Customers’ enterprises are unique: they reflect the use of data-center data and legacy applications that reflect decades of past IT decisions. The types of data resources available to the enterprise often reflect the results of earlier merger & acquisitions (M&A) decisions. Now, old and new data resources must live, or co-reside, in a distributed mix of Core, Cloud and Edge data “hubs.” For enterprises, the world that is taking shape in 2022 combines on-premises resources with a broad mix of off-premises cloud services and co-lo provider services.

In effect, using a car-tire analogy, this is where the “rubber meets the road,” leveraging a wide variety of data resources that are increasingly being leveraged in New Normal infrastructure. If the current move to transform all customer computing is to work, then customers will demand that all providers offer new software frameworks that can access and manage the wide “universe” of their applications and data, across their enterprises and organizations.