HPE GreenLake

Empower engineers and scientists with on-premises, managed HPC

HPE GreenLake for HPC

Running high-performance computing (HPC) applications on cloud platforms that are not optimized for technical computing has been a challenge for most enterprise customers who were promised HPC at scale. HPC users already know that not all scientific and technical applications can be containerized for the cloud but have now learned that it's also difficult to manage and tune applications that they have chosen to run in the cloud. Bursting to the cloud from on-site HPC platforms has worked for some but can continue to be difficult to administrate and is expensive to operate at scale with a cloud service provider. The volume of HPC data makes ingress/egress expensive and slow, and HPC users are at increased risk for data security exposures beyond their corporate firewall. The difficulty of running HPC on public cloud infrastructure is exacerbated by the dwindling pool of HPC expertise caused by retirements, changing business priorities, and financial constraints.

Commercial HPC users need a solution that provides predictable costs to manage their operational budgets including tracking daily use of HPC resources and predicting the impact before they are invoiced. HPC applications are well known for their unique performance profiles and HPC customers require configuration flexibility in the underlying HPC platform to tune and optimize application performance and achieve maximum cost efficiency. HPC users want to focus on the speed and the ease of data ingress/egress and not the cost of data ingress/egress. HPC data stays most secure when it remains behind the corporate firewall and on-premises.

HPE GreenLake for HPC allows for valuable people and resources to be spent on core science and engineering, and not the management of complex HPC technology.

HPE GreenLake for HPC is an on-premises HPC solution that delivers the ease and efficiency of on-premises hosting with cloud-like data access. The solution includes a high-performance compute cluster, high-performance parallel file storage, and an HPC network supplemented with industry standard HPC tools. The entire solution is delivered, deployed, configured, and integrated into your network and is fully managed and supported by Hewlett Packard Enterprise.

HPE GreenLake for HPC is pre-assembled and configured by HPC specialists in the factory to help eliminate setup on-site and accelerate time to data insights. HPC implementation service from HPE stands up the HPC cluster and connects to your user and their data sources.

With HPE GreenLake for HPC, IT admins can avoid a long, steep learning curve to develop HPC skills. Instead, you leverage HPC expertise from HPE to manage the end-to-end HPC environment, including system operations, system maintenance, and problem reporting-escalation-resolution.

HPE GreenLake Central is an advanced software-as-a-service platform that provides HPC users with a consistent cloud experience for every application and data—on-premises or off-premises at a colocation. It provides you with insights and controls to manage your hybrid cloud estate, providing the benefits of public and private clouds.

HPE GreenLake for HPC is flexible and allows you to choose the instance types that are best fit for your application. You can mix and match the types and quantities of CPU nodes, GPU nodes, and storage nodes to personalize your configuration for HPC, artificial intelligence (AI), Big Data, price/performance, I/O bandwidth, or I/O operations.

HPE GreenLake for HPC is a consumption service. It has simple, effective metering via HPE GreenLake Central, so you can see who is using which resources and help eliminate surprise charges at the end of the month. In cases of unforeseen spikes in workload demands, the service enables your tasks to expand to on-premises resources instead of relying solely on the public cloud.

When the surge is over, you return to a normal operating state and only pay for what you use.* When capacity growth is planned, HPE GreenLake for HPC can scale in place at predictable rates for compute, GPU, and/or storage.

The new HPE Machine Learning Development Environment option for HPC users is a software solution that is designed to accelerate deep learning workloads in HPC environments and provide AI/ML model training at scale for the cloud. It is ideal for working in conjunction with HPE Cray HPC systems and leverages the power of GPUs to accelerate the training and inference phases of deep learning models. It does this by optimizing data movement between the CPU and GPU, and by providing a software environment that is optimized for deep learning workloads.

In an HPC environment, HPE Machine Learning Development Environment can be deployed on top of HPC schedulers such as Slurm or PBS Pro. This allows the platform to integrate seamlessly with other HPC workloads and to take advantage of the HPC cluster's resources such as those found in HPE GreenLake for HPC. HPE Machine Learning Development Environment can also leverage high-speed networking solutions from HPE, such as InfiniBand, to further accelerate data movement between nodes in the cluster.

HPE Machine Learning Development Environment provides a comprehensive solution for deploying machine learning (ML) workloads in an HPC environment, delivering high performance and scalability while leveraging the full power of HPC hardware and software.

* May be subject to minimums or reserve capacity may apply

Unlock the value of your data by deploying HPC workloads that provide a single unified data source for analytics, support data processing at scale, and enable centralized cloud access to your data, without the burden of managing any HPC infrastructure. Modernize your approach to data-intensive workloads and solve complex problems in less time at a lower cost by building complex AI models alongside HPC applications with machine learning development environment solutions.

Learn more at

HPE.com/us/en/GreenLake/high-performance-compute.html

Explore **HPE GreenLake**

Chat now (sales)