A cloud experience without risks

HPE GreenLake for File Storage helps eliminate trade-offs

These days, there's much talk about going to the public cloud for IT infrastructure. Enterprises want the compelling benefits of the cloud experience — simplicity, agility, and as-a-service consumption — but migrating to public cloud involves persistent concerns, especially performance at AI (artificial intelligence) scale, data sovereignty and security, and cost control. These concerns are deterring IT managers from moving to the public cloud.

Into this mix, Hewlett Packard Enterprise delivers data services via the HPE GreenLake edge-to-cloud platform, including HPE GreenLake for File Storage, which combines an intuitive cloud experience with the security, control, performance, and cost advantages of on-prem infrastructure.

To understand the groundbreaking benefits of HPE GreenLake for File Storage, we need to know the importance of public cloud trade-offs. First, a little data storage history.

HPE GreenLake

The move from centralized to distributed

In the early days of computing and storage, IT infrastructure was centralized and built on proprietary mainframe computers in the data center. Users used dumb monitors, storing all data in the data center. With the introduction of the personal computer and its 160 KB floppy disks, UNIX® open systems servers, network-attached storage (NAS), and the rapid growth of the internet, we would eventually generate huge amounts of distributed data, particularly file data. At this point, modern data storage looks very different from the old, centralized data center approach.

Today, handling all that distributed data is a significant challenge.

- How do you manage it, keep track of it, and back it up?
- What strategies can you use to make storage more easily available to users and application owners?
- How do you make it shareable and eliminate data silos?
- What can be done to protect it and keep it secure?
- And now, with the proliferation of data-intensive workloads such as AI and machine learning (ML), how can your data storage infrastructure deliver the required performance to extract insights and derive value from your data in aggregate?

Cloud trade-offs drive the focus back to on-prem enterprise storage

In a distributed data world, storing, managing, and protecting all your data is more daunting than ever. Organizations are now focused on data-first transformation initiatives — many of which rely on migrating to the public cloud to simplify and re-centralize data management.

Farming out IT data storage infrastructure can simplify operations, reduce demand on IT staff, and trim hardware purchases, but storing data in the public cloud also involves unavoidable trade-offs. The drawbacks include security, performance and latency, cost at scale, data sovereignty and management, and flexibility. The weight of these trade-offs has pushed many businesses to keep data storage infrastructure on-premises — reversing the earlier focus on the public cloud and shifting it again toward the data center.

In file data storage infrastructure, the convenience of using the public cloud comes at the cost of losing control of your security protocols, including the critical ability to put physical security measures in place for hardware. Data security is increasingly important in the face of data breach proliferation. Physical and software protection measures to control and restrict access to data are primary factors in formulating an IT strategy.

Now consider performance and latency. Public cloud storage places a premium on performance. There is a higher cost for meeting IOPS and throughput performance requirements beyond baseline levels. Therefore, we're not only concerned about performance but also about the cost at scale of public cloud storage as data volume increases. It is difficult to achieve enterprise performance at scale for large data volumes for your most demanding workloads while keeping costs down.

Data sovereignty, control, and management are also eroded when file data storage infrastructure is external to an organization. Data is your most important asset. In fact, it's the life force of your enterprise, so preserving your ownership of it is of paramount importance.

When you use the public cloud for your data storage infrastructure, you're paying for storing your data. However, putting your data in the public cloud is much easier than taking it back. It is designed like a one-way street, as egress fees to take back your data can be highly expensive. This kind of lock-in reduces your agility and cost control.

The last major trade-off has to do with flexibility. As technology companies are all shifting toward subscription-based models, they want subscribers to commit for the long run to generate a recurring revenue stream. It makes practical business sense as well. Which company wouldn't want to secure ongoing income by being entrenched as part of a customer's business infrastructure?

The key question is whether you, as the customer, have options and flexibility in your subscription. Can you choose between purchasing your hardware or going purely with a subscription? Most importantly, can your hardware live on-prem, and can you handle the data management yourself? As we'll see, the HPE approach addresses this customer demand for flexibility, along with simplicity, security, performance, cost-effectiveness, and data sovereignty.

A new paradigm in HPE GreenLake

We created the HPE GreenLake platform to provide an intuitive hybrid cloud experience — and the benefits of public cloud — without the drawbacks and trade-offs. With over 50 cloud services spanning storage, networking, compute, and workloads, the HPE GreenLake platform empowers you to accelerate your data-first modernization flexibly with a unified experience, full control, and broad visibility into your data, bringing the cloud to wherever your apps and data live.

Now, there's HPE GreenLake for File Storage

HPE GreenLake for File Storage delivers an on-premises cloud management experience for file storage without the security risks and trade-offs of the public cloud.

With an intuitive cloud experience and radically simple file data management, HPE GreenLake for File Storage reduces operational and management overhead for your IT staff, line-of-business application owners, and data scientists. Powered by the HPE GreenLake platform, you get streamlined deployment, easy file share creation, unified storage management with a single cloud console, and automated, nondisruptive upgrades.

In addition, you can modernize your data management with a comprehensive suite of cloud data services and a unified operating experience on the HPE GreenLake platform. This simple, end-to-end, self-service cloud experience across the file storage management lifecycle — accessible from anywhere, on any device — empowers your IT staff to work on strategic initiatives rather than be weighed down by day-to-day operations.

You can easily maintain your storage infrastructure and make nondisruptive upgrades while benefitting from investment protection and high ROI with data-in-place upgrades. Efficiently scale up and out by adding controller and storage nodes as needed. Keep all your data on a single tier of fast, scalable, and affordable all-NVMe flash to help eliminate silos, complex data movement workflows, or application wait times. New insights become possible when exabytes of data are available in real time.

HPE GreenLake for File Storage offers enterprise performance at scale to accelerate the most data-intensive applications, including AI, ML, deep learning, Big Data, high-performance computing, life sciences, financial analytics, media and entertainment, large data lakes, and high-speed access to data archives. Built on VAST Data software and the HPE Alletra Storage MP modular hardware platform for both file and block storage, HPE GreenLake for File Storage provides a highly differentiated file storage offering.

By avoiding trade-offs, the solution makes storage simple while delivering performance at AI scale via a unique and modular hardware architecture, helping eliminate hardware silos, islands of data, and tiered data pipelines. Consider the high degree of flexibility in deploying the storage resources you need with common hardware. End users can quickly set up file storage, access and share data, and complete their work quickly, while IT staff will no longer have to worry about trade-offs in security, performance and latency, cost at scale, data sovereignty and management, and flexibility.

Next, you can leverage HPE GreenLake for File Storage to unlock more value from your aggregated data. Simple setup and fast job completion increase the productivity of your data scientists and line-of-business application owners. You'll get results faster as your team is not burdened by cumbersome, legacy file systems that require technical expertise and intricate setup processes. And your highly compensated subject-matter experts no longer need to spend their valuable time and effort performing routine IT tasks and instead focus on innovation.

As organizations bring data back on-premises, HPE GreenLake for File Storage presents a compelling solution. You might even say that HPE has introduced an ideal file storage offering at the right time and place to meet industry needs, particularly in AI, as enterprise storage requirements and the cloud operational model converge to establish an entirely new operating paradigm.

Learn more at

HPE.com/us/en/HPE-GreenLake-file-storage.html

Explore **HPE GreenLake**

Chat now (sales)

Hewlett Packard Enterprise

a50009537ENW, Rev. 1