HPE GreenLake

Improving the ROI for enterprise HPC and AI with HPE end-to-end clusters

Executive summary

As organizations in many industries struggle to overcome competitive and cost pressures, they are increasingly using high-performance computing (HPC) and artificial intelligence (AI) to drive greater innovation and productivity and deliver higher quality products and services faster at lower costs and risks. The use of HPC in business entities, enterprise HPC, is growing, is becoming mission critical, and has the potential to deliver a high return on investment (ROI).

However, deploying and scaling compute and storage capabilities is especially challenging as data volumes and variety explode. Scalable clusters that integrate compute, storage, networking, and software can address these challenges. But often, with rapid advances in technology components, many of these clusters get outdated. This inhibits clients from achieving maximum ROI.

The HPE end-to-end (E2E) cluster delivers technology innovation leadership with leading-edge components across compute, storage, software, networking, and services. It can be deployed on-premises, in the cloud, or as a service with worldwide service and support and a single point of contact to address and resolve the issues. This helps maximize the ROI for HPC and AI clients.

The promise of enterprise HPC

Organizations in every industry are under severe competitive pressure to innovate more, rapidly undergo a digital transformation, improve quality, enhance productivity, and reduce time to market, costs, and risks. HPC, analytics, AI, and machine learning (ML) are critical technologies to overcome these challenges.

With HPC and AI, businesses can deliver (Figure 1) higher quality products faster, optimize oil and gas exploration and production, improve patient outcomes, mitigate financial risks, and more. HPC also helps governments respond faster to emergencies, analyze terrorist threats better, and more accurately predict the weather. The return from HPC alone can be several times the initial investment.[1]

Consequently, the HPC market is expected to grow from $36.0 billion in 2022 to $49.9 billion by 2027, at a 6.7% annual growth rate.[2] As HPC is integrated with analytics and AI (Figure 1), the potential for additional growth and value is even greater.

Industry use case	Defense/ national security	Manufacturing	Electronic design automation	Energy	Healthcare & life sciences	Finance	Earth sciences
	• Combat systems • Video surveillance • Cybersecurity	• Computer aided engineering • Supply chain management • Predictive maintenance	• System/logic design • Simulation/layout • Testing/verification	• Seismic processing • Production optimization • Predictive maintenance	• Drug discovery • Clinical trials • Personalized medicine	• Risk analytics • Treasury analytics • Customer analytics	• Daily forecast • Seasonal forecast • Climate prediction

Solution domain	Simulation	Analytics	AI/ML

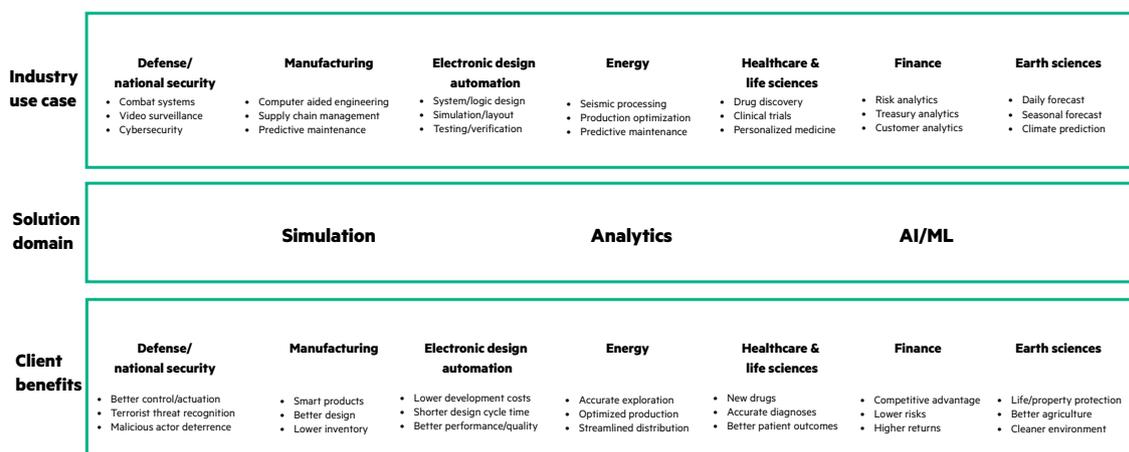| Client benefits	Defense/ national security	Manufacturing	Electronic design automation	Energy	Healthcare & life sciences	Finance	Earth sciences
	• Better control/actuation • Terrorist threat recognition • Malicious actor deterrence	• Smart products • Better design • Lower inventory	• Lower development costs • Shorter design cycle time • Better performance/quality	• Accurate exploration • Optimized production • Streamlined distribution	• New drugs • Accurate diagnoses • Better patient outcomes	• Competitive advantage • Lower risks • Higher returns	• Life/property protection • Better agriculture • Cleaner environment

Figure 1. Integrating simulation, analytics, and AI/ML drives add client value

But as the use of HPC and AI grows and becomes more mainstream at enterprises, meeting customer requirements in every industry is becoming more critical (Figure 2). Scalable clusters are designed to address these growing needs.

---

[1] The Business Value of Leading-Edge High Performance Computing: 2019 Update

[2] "High-performance Computing (HPC) Market by Component, Computation Type (Parallel Computing, Distributed computing and Exascale Computing)," MarketsandMarkets Research Private Ltd., 2020

Scalable clusters for enterprise HPC and AI

To achieve the full potential ROI from HPC and AI, clients across several industries are realizing that they must address similar business challenges. Moreover, their HPC and AI solution must deliver higher accuracy, faster time to results, higher performance, improved security and reliability, lower total cost of ownership (TCO), and a single point of contact for all support issues (Figure 2). Clients also need the flexibility to run their HPC and AI workloads on-premises and scale computing capabilities without large capital expenses by using cloud-like consumption-based pricing.

Business challenges	Industries	Customer requirements	HPC and AI solution
Lower costs	Defense/national security	High accuracy	Scalable cluster
Enhance productivity	Computer aided engineering	Faster time to result	High frequency CPU
Improve quality	Electronic design automation	Lower TCO	Large core count and memory capacity
Reduce time to market	Energy	Improved security	Fast memory bandwidth
Drive innovation	Healthcare & life sciences	Reliability	Greater I/O performance
Enhance collaboration	Financial services	Performance	Lower network latency and high bandwidth
Mitigate risk	Earth sciences	One vendor to resolve issues	Financing options
			Usage-based pricing

Figure 2. HPC solutions for business challenges and customer requirements across several industries

A scalable cluster system with HPC, interconnect, and storage provides these capabilities (Figure 2). Since these systems are expensive and need to be upgraded periodically, clients are also interested in financing options and usage-based pricing.

However, these scalable cluster solutions have been a mix of success and disappointment. While some organizations are relatively satisfied with their HPC investment, most are yet to fully maximize their ROI. This is because HPC clusters suffer many pitfalls in real-life operation.

Pitfalls of enterprise HPC clusters

HPC clusters typically have the same types of components: a collection of high-performance processors, interconnects, storage, and software to manage the cluster efficiently. Until recently, common wisdom was that customers needed a multivendor environment for HPC clusters if they wanted to get **best-of-breed** capabilities in every layer of the technology stack. However, this is hard to maintain because of the rapid evolution of technology and the associated high cost to upgrade to the most current technology level. So, at some time after initial deployment, many HPC clusters do not deliver the best ROI possible. Often this is because of circumstances beyond the customer's control.

These customers may not be getting the best level of service and support to operate their clusters since no single solution provider is fully responsible for delivering the best possible business outcomes from their HPC and AI solution. The root causes for this usually rest with the way these clusters are designed to operate. They are typically patched together with dated compute processors from various vendors, I/O and network speeds constrained by outdated technology, older storage systems from a storage vendor, and cluster management software from yet another vendor.

This **old way** of mixing and matching of components results in HPC systems that are complex to manage and have poor and inconsistent performance. Also, multiple vendors must be contacted when issues arise, so support costs can escalate quickly. HPC clusters are capital intensive and must be upgraded frequently as demand rises or fluctuates with many peaks and valleys. This drives the need for flexible financing options, usage-based pricing, and access to an external source of computing such as public cloud.

Hence, now customers are realizing that to help maximize ROI, it is more prudent to have a **new way**: have a single vendor with deep HPC and AI skills be responsible for the entire set of components of a HPC solution including financing, professional services, and support. It is the HPE E2E cluster solution (Figure 3). Hewlett Packard Enterprise also provides HPE Accelerated Migration to help customers upgrade their outdated clusters with PCIe Gen3 compute and storage nodes to high bandwidth PCIe Gen4-based solutions. This helps them keep up with increasing business demands.
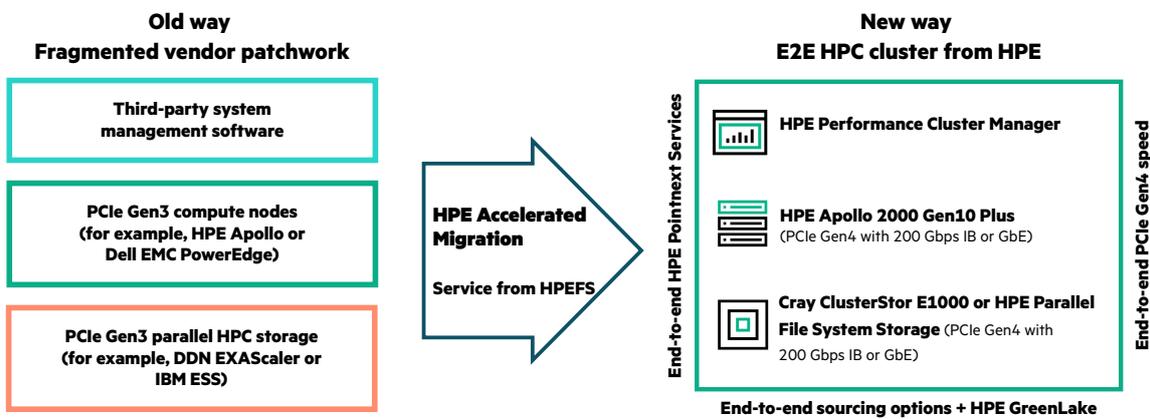
Figure 3. Migrate from the old way to a new way with an HPE E2E HPC cluster

Why upgrade now to the HPE E2E cluster?

With 37.2% market share of HPC installations, HPE is the market leader with a comprehensive end-to-end portfolio across compute, storage, software, networking, and services.[3] To keep up with growing business demands and get a much greater ROI and faster time to value, customers can replace their legacy, less performant, and harder to manage HPC cluster with the HPE E2E cluster built from the HPE Apollo 2000 Gen10 Plus system by (Figure 3):

- Swapping obsolete PCIe Gen3 compute and storage with PCIe Gen4 to remove compute, I/O, and network bottlenecks.

- Seamlessly moving from 56/100 Gbps to the cutting-edge 200 Gbps interconnects, to move data faster for better productivity.

- Quickly resolving challenging system-wide issues with a single point of contact and avoiding vendor finger-pointing to improve system uptime and productivity.

- Offering—unlike many other vendors—a subscription-based consumption model such as HPE GreenLake for HPC and other **as-a-service** options.

[3] "2019 Market Results, New Forecasts and HPC Trends," Hyperion Research, April 2020

The HPE Apollo 2000 Gen10 Plus system includes the following:

• AMD EPYC™ 7003 series processors—the latest generation of AMD processors delivering superior performance for HPC and AI

• Compute nodes with PCIe Gen4, which enables a 200 Gbps connection—effectively doubling transfer speeds between compute nodes and with storage

• Integration with a choice of high-performance HPE HPC storage solutions

• Modular software stack to accommodate variety of customer requirements

• Single point of contact for the support issues

• System upgrades that are easier with flexible financing options including pay-for-usage pricing options

These innovations are detailed in the following, starting with the AMD EPYC 7003 Series processor followed by the HPE E2E cluster solution stack.

High-performance AMD EPYC 7003 Series processor

HPE Apollo 2000 Gen10 Plus servers are powered by AMD EPYC 7003 Series processors. Built on 7 nm technology, the AMD EPYC 7003 Series processors bring together high core counts, large memory capacity, extremely high memory bandwidth, and massive I/O throughput (Figure 4) with the right balance to enable exceptional HPC and AI workload performance.

AMD EPYC processors are the choice of next-gen exascale supercomputers and the EPYC 7003 series processors are the highest performing server CPUs[4] with unique features and are also highly affordable, often delivering superior performance to alternative processors while easily fitting within the budgets for HPC and AI environments of all sizes.

Compute
• AMD "Zen 3" x86 cores (64 core/128 threads)
• Up to 32 MB L3 cache/core, shared by each chiplet
• Flatter NUMA domain, reduced latency w/ smaller system diameter
• TDP range: 120W–280W

Memory
• 8 channel DDR4 with ECC up to 3200 MHz Option for 6 channel memory interleaving[5]
• RDIMM, LRDIMM, 3DS, NVDIMM-N
• 2 DIMMs/channel capacity of 4 TB/socket (256 GB DIMMs)

Performance
• Built to handle large scientific and engineering datasets with top performance—ideal for HPC workloads
• Infinity Fabric™ Gen2 (xGMI-2)

Integrated I/O—no chipset
• 128 lanes PCIe Gen3/4
  – Used for PCIe, SATA, and coherent interconnect
  – Up to 32 SATA or NVMe™ direct connect devices
  – 162 lane option (2P config.)
• Server controller hub (USB, UART, SPI, LPC, I2C, etc.)

Security
• Dedicated security subsystem
• Secure boot, hardware root of trust
• Secure Memory Encryption (SME)
• Secure Encrypted Virtualization-Encrypted State (SEV-ES)
• Secure Nested Paging (SNP)

Figure 4. AMD EPYC 7003 Series processor high-level processor design[6]

---

[4] As of 03/18/2021, result with 2 x AMD EPYC 7763 on the HPE ProLiant DL385 Gen10 Plus with a measured value of 821 for SPECrate®2017_int_base. This is higher than the previous best 2P server with an AMD EPYC 7H12 and a score of 717, spec.org/cpu2017/results/res2020q2/cpu2017-20200525-22554.pdf. SPEC® and the names SPECrate® and SPEC CPU® are registered trademarks of the Standard Performance Evaluation Corporation (SPEC). All rights reserved. See spec.org for more information.

[5] With certain DIMM population rules

[6] Image provided by AMD

AMD EPYC 7003 Series processor are industry leading particularly for HPC and AI with the following:

- **Compute**
  - High density—Up to 64-cores/128-threads per socket
  - Superior floating-point performance at single- and double-precision
  - Up to 32 MB L3 cache/core
  - All 32 MB allocated to a single core if needed
  - Significant application performance improvement where datasets fit more naturally in large cache

- **Superior memory bandwidth**
  - Provides eight channel DDR4 with ECC up to 3200 MHz
  - Helps improve memory performance
  - Is critical for all HPC workloads that are memory-bandwidth sensitive

As a market-leading HPC solution provider, HPE integrates the AMD EPYC 7003 Series processors in the HPE E2E cluster to deliver a unified compute and storage solution designed to simplify system and data management, reduce costs and complexity, and scale to deliver the exceptional performance needed for HPC and AI.

HPE E2E cluster solutions for enterprise HPC

Figure 5 depicts the various components of the HPE E2E cluster solution stack for HPC and AI built with the HPE Apollo Gen10 Plus system, HPE high-performance storage solutions, and HPE GreenLake.
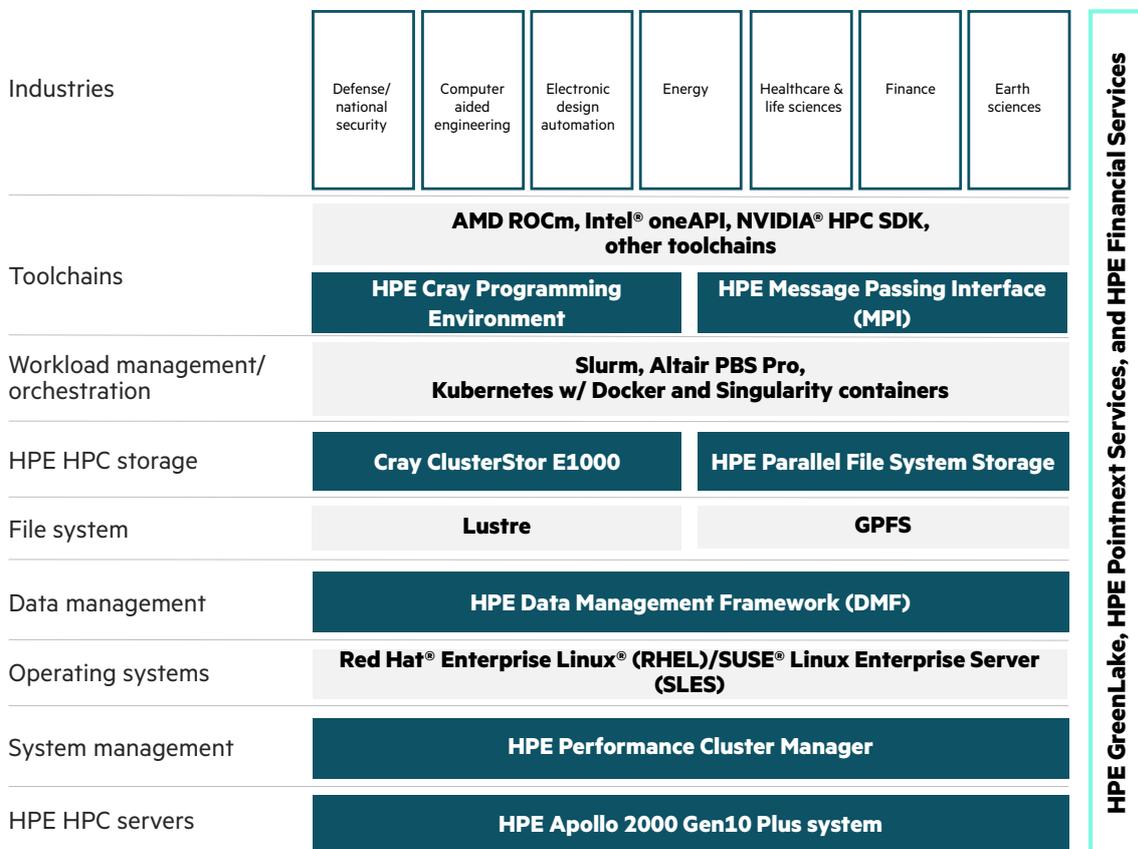
Industries	Defense/ national security	Computer aided engineering	Electronic design automation	Energy	Healthcare & life sciences	Finance	Earth sciences	**HPE GreenLake, HPE Pointnext Services, and HPE Financial Services**
Toolchains	**AMD ROCm, Intel® oneAPI, NVIDIA® HPC SDK, other toolchains**							
	HPE Cray Programming Environment			**HPE Message Passing Interface (MPI)**				
Workload management/ orchestration	**Slurm, Altair PBS Pro, Kubernetes w/ Docker and Singularity containers**							
HPE HPC storage	**Cray ClusterStor E1000**			**HPE Parallel File System Storage**				
File system	**Lustre**			**GPFS**				
Data management	**HPE Data Management Framework (DMF)**							
Operating systems	**Red Hat® Enterprise Linux® (RHEL)/SUSE® Linux Enterprise Server (SLES)**							
System management	**HPE Performance Cluster Manager**							
HPE HPC servers	**HPE Apollo 2000 Gen10 Plus system**							

Figure 5. HPE E2E cluster solution stack for HPC and AI

The HPE Apollo 2000 Gen10 Plus system (Figure 6) is a dense, multiserver platform that packs incredible performance and workload flexibility into a small data center space while delivering the efficiencies of a shared infrastructure. It is designed to provide a bridge to scale out architecture for traditional data centers, so enterprise customers can achieve the space-saving value of a dense infrastructure in a cost-effective and nondisruptive manner. Salient features include the following:
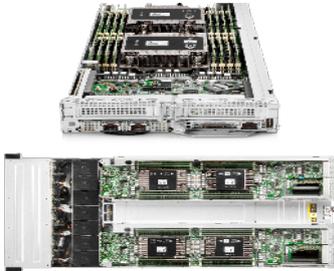
Figure 6. HPE Apollo 2000 Gen10 Plus system

- **System performance and optimization**

  – 2x compute density of a traditional 1U server[7]

  – Expanded power capabilities

  – Software development and application acceleration tools for application performance at scale

- **Flexible scale-out building blocks**

  – Provide storage and I/O flexibility

  – Rightsize building blocks with future proof scalability

  – Offer comprehensive software portfolio to accommodate any workload

- **Comprehensive server security and management**

  – Secure from the start with HPE iLO 5 and silicon root of trust from HPE

  – Maintain system uptime and lower exposure to security risks with fully integrated cluster software

HPE HPC storage solutions span the whole storage hierarchy to accelerate time to insights while managing and protecting the valuable data in a customer's parallel file systems in a cost-effective way. Parallel file systems deliver aggregate speeds that exceed the architectural limitations of NAS, scale-out NAS, distributed file systems, or object storage. Two parallel file system products (Figure 7) with enterprise-grade support from HPE include the following:

- **Cray ClusterStor E1000 Storage System** comes with the open-source Lustre PFS. The engineered parallel storage system comes with purpose-built, high-performance storage controllers for extreme speed and scalability. It is for organizations that do not require enterprise-grade functionality but value extreme price/performance and scale. Key features include:

  – Up to 80 GB/s from just 24 SSDs in just two rack units

  – Up to 3.3 GB/s per SSD data transfer to the compute nodes

  – Connected with 200 Gbps HPE Slingshot or InfiniBand EDR/HDR or 100/200GbE

  – Benefits of open-source Lustre file system—No software license per TB or storage drive

  – Additional support with in-house Lustre R&D team

  – Built with AMD EPYC 7002 Series processors with PCIe Gen4

---

[7] HPE Apollo 2000 Gen10 Plus System

- **HPE Parallel File System Storage** embeds IBM Spectrum Scale a General Parallel File System (GPFS) for the enterprise. It is a software-defined storage solution built on cost-effective HPE ProLiant DL rack servers and offers a broad set of enterprise storage features: enterprise IT-grade data availability (backup and disaster recovery), data accessibility (NFS, SMB, HDFS, Object), and data compliance (audit log, industry certifications). Key features include the following:

  – Combination of a leading parallel file system in the enterprise with HPE ProLiant DL x86 rack servers

  – Starts as low as 27 TB in just four rack units and scales to more than 25 PB in a single file system (current testing limitation—not an architectural limitation)

  – Connected with InfiniBand EDR/HDR or 100/200GbE

  – No separate software license per TB or storage drive

  – Operational support services for the full product, both hardware and software, from HPE Pointnext Services

  – Built with AMD EPYC 7002 Series processors with PCIe Gen4

HPE Parallel File System Storage
(with embedded IBM Spectrum Scale parallel file system)

Cray ClusterStor E1000 Storage System
(with embedded Lustre parallel file system)

Figure 7. Two HPE Parallel File System Storage offerings with enterprise-grade support

System management with fully integrated cluster management software—The HPE Performance Cluster Manager offers customers the functionalities they need to manage their Linux-based HPC systems. The software provides system setup, hardware monitoring and management, cluster health management, image management, and software updates as well as advanced power and cooling management for clusters of any scale.

HPE Performance Cluster Manager turns even the most complex hardware into easily manageable systems capable of accommodating a growing variety of workloads. The software reduces the time and resources customers need to spend administering their systems—lowering total cost of ownership, increasing productivity, and providing a better return on hardware investments.

Operating systems are a choice of standard Linux operating systems with service-supported subscriptions for RHEL or SLES.

Data management with HPE Data Management Framework helps optimize storage resource utilization and data accessibility by introducing a hierarchical, tiered storage management architecture. Data is moved between tiers based on service-level requirements defined by the administrator. For example, frequently accessed data can be placed on a flash, high-performance tier; the data accessed less often can stay on hard drives in a capacity tier; and the data to be archived can be sent off to tape storage.

File systems that deliver fast I/O is critical for many HPC and AI applications. High-performance Lustre file systems such as the Cray ClusterStor E1000 Storage System or the HPE Parallel File System Storage based on IBM GPFS can be used.

Workload management and orchestration tools that run workloads, either on bare metal or via containers or mix of both, using HPC-specialized workload managers such as Altair PBS Professional or Slurm for workloads, and/or Kubernetes for Singularity or Docker container for orchestration.

Software development libraries and tools are essential for development and acceleration of HPC codes. HPE Cray Programming Environment is a comprehensive set of tools for developing, porting, debugging, and tuning of applications designed to increase their productivity, application scalability, and performance. Alternatively, customers can also choose other leading open-source and commercial software development tools.

HPE GreenLake for HPC is a market-leading IT as-a-service offering for HPC and AI. It offers easy and affordable access to dedicated, powerful computing and analytics capabilities helping customers make faster decisions and reduce time to discovery. It avoids overprovisioning costs with elastic capacity ready for growth or unpredictable spikes. HPE GreenLake combines the simplicity, agility, and economics of public cloud with the security and performance benefits of on-premises IT.

This consumption-based IT model helps customers accelerate time to value, align IT economics with business priorities, simplify IT operations, and gain better control. HPE GreenLake for HPC brings a consumption-based HPC model on-premises—or in a colocation—that delivers superior flexibility, scalability, and control. With HPC and AI as a service, customers can design their own HPC and AI infrastructure solution using industry-leading HPE technologies or can standardize their service with presized configurations that are self-service and managed for them. A built-in technology refresh feature in HPE GreenLake engagement allows customers to benefit from the latest technology available in the market, so they can stay competitive. HPE can also buy out—and recycle—existing infrastructure to help meet sustainability targets.

Services: HPE Pointnext Services offers a spectrum of services to meet HPC and AI requirements—from services such as application tuning to more integrated advisory service offerings such as project management, on-site consulting, technical account management, and solution architecture consulting.

HPE HPC Cluster Management Solution and HPE Pointnext Services skilled consultants provide customers with assistance in installation, configuration, and understanding the management of the entire HPE cluster environment.

HPE financial support helps customers purchase HPC systems and upgrade them frequently, as this can be financially daunting for most organizations. HPE offers wide range of end-to-end sourcing options for the full HPC and AI infrastructure stack:

• Classic purchase: Customer owns it and runs it.

• Financing with HPE Financial Services (HPEFS): HPEFS finances it and customer runs it.

• HPE GreenLake for HPC: Customer subscribes and pays for what they use (pay for usage). Customer runs it.

• HPC as a service: Customer subscribes and pays for what they use. HPE runs it.

The HPE E2E cluster advantage for enterprise HPC and AI

Most HPC and AI systems have features that are cobbled together consisting of older computer nodes from various server vendors, I/O speeds constrained by dated technology, older storage solutions from a storage vendor, and a cluster management solution from yet another vendor. This mixing and matching of HPC components have resulted in HPC systems with poor and inconsistent performance. Moreover, multiple vendors must be contacted when issues arise with no one single vendor responsible for the performance of the system.

The answer to the challenge of a suboptimal multivendor HPC and AI system is the HPE E2E cluster solution, which offers technology innovation leadership and helps maximizes ROI for clients with the following:

- HPE E2E cluster provides a comprehensive solution across compute, storage, software, and services that is pretested for performance and compatibility.

- HPE Apollo 2000 Gen10 Plus compute nodes with PCIe Gen4 enable a 200 Gbps connection and double transfer speeds between compute nodes and with storage as compared to prior generation PCIe Gen3 solutions.[8]

- It includes integration with HPC storage and HPE Performance Cluster Manager.

- HPC as a service offers customers choice, flexibility, and speed to market with **pay for usage** options that offer financial and operational flexibility.

- It provides worldwide service and support with a single point of contact to address and resolve issues.

- It offers financing options through HPE Financial Services for easy upgrades and economical HPC and AI infrastructure modernization.

Learn more at
hpe.com/apollo

Visit **HPE GreenLake**

---

[8] Gen3 NIC cards can deliver 100 Gbps, Gen4 NIC cards can deliver 200 Gbps

**Make the right purchase decision.
Contact our presales specialists.**

Chat now (sales)

Call now

Get updates

**Hewlett Packard
Enterprise**