Hewlett Packard Enterprise

Business white paper

Hybrid HPC cloud

Bringing outstanding agility, simplicity, and economics to HPC using cloud technologies, operating methods, and business models

Table of contents

Executive summary

Across every industry, businesses are under severe competitive pressures to deliver and support higher quality and more custom products and services through digital platform business models. To address these challenges across the value chain, high-performance computing (HPC) unlocks several innovative data-driven computing workloads such as high-performance data analytics (HPDA), **artificial intelligence**/**deep learning** (AI/DL), and others. HPC is, therefore, growing at an unprecedented rate and many HPC customers who largely use on-premises data centers are exploring cloud computing to improve the end-user experience, agility, and economics.

Public clouds can be attractive for some customers by providing flexibility, scalability, pay-per-use, and lower costs for their HPC workloads. However, they are either low performing, expensive, or hard to customize for the increasing diversity and complexity of many industry-specific, data-intensive HPC and AI/DL workloads. On the other hand, for these workflows, private clouds (on-premises or off-premises delivered by a managed services provider) can be highly customized with legacy integration to deliver significant HPC capabilities with cloud end-user experience and agility.

A hybrid cloud that combines or augments a traditional HPC data center with a private or public cloud offers a better solution for a broad range of HPC workloads. As the market leader in HPC, Hewlett Packard Enterprise is delivering a purpose-built portfolio of Hybrid HPC solutions jointly with Intel®, several independent software vendors (ISVs), and cloud service providers (CSPs) to support customers with their cloud journey. This portfolio can be customized to specific industries, applications, and workflows. It is anchored on HPE Apollo servers based on innovative HPC foundation technologies from Intel.

HPC delivers high value across many industries

Businesses are investing in innovative HPC foundation technologies from Intel to deliver higher quality products faster, optimize oil/gas exploration and production (E&P), improve patient outcomes, mitigate financial risks, and more. **HPC** also helps governments respond faster to emergencies, analyze terrorist threats better, and even more accurately predict the weather.

The return on investment (ROI) from HPC is more than two times the capital invested.[1] As a fast-growing trend, simulation is being integrated more with HPDA and AI/DL with the potential to drive even greater value as shown in Figure 1.

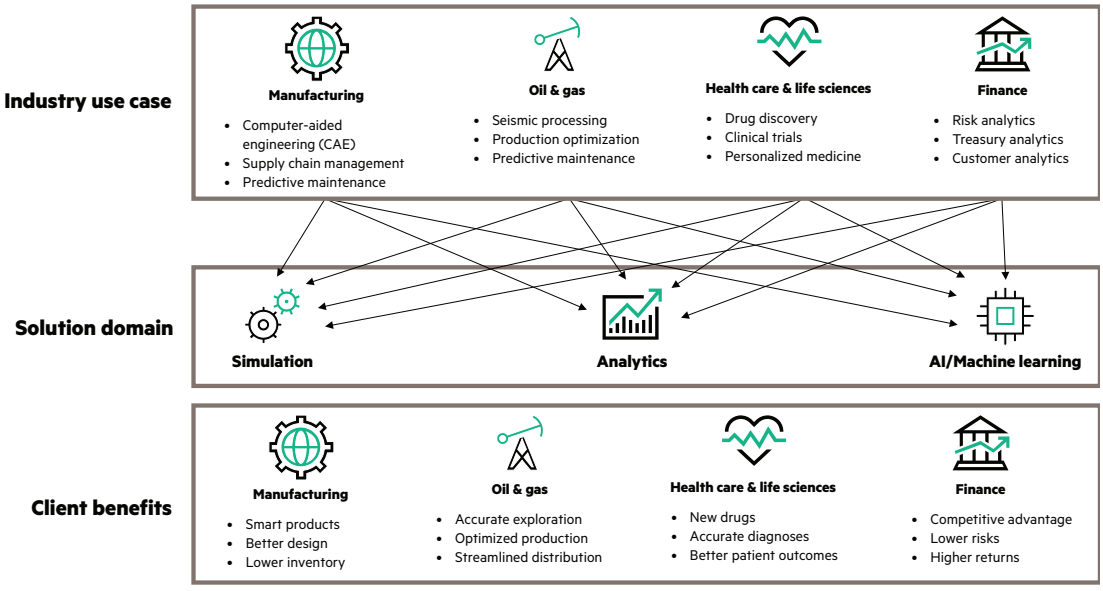

Figure 1. Integrating simulation, analytics, and AI/DL drives added customer value

[1] Hyperion (IDC) Paints a Bullish Picture of HPC Future—**hpcwire.com/2017/04/20/hyperion-idc-paints-bullish-picture-hpc-future/**

Designed for the convergence of HPC and AI, and supported by a broad software ecosystem, the Intel HPC platforms offered by HPE provide balanced HPC system performance, lower latency, greater capacity, and improved efficiency. Here are some noteworthy examples of customer use cases and benefits by industry that leverage a variety of innovative foundational technologies from Intel on-premises and on the cloud.

Manufacturing: By combining and complementing physics-based CAE models with Digital Twins[2] and **Internet of Things (IoT)** sensor data along with AI/DL models, it is possible to detect and predict product anomalies, as well as failures. This predictive maintenance capability helps take optimal corrective actions during operations, improve product maintenance, and further accelerates the new product development process. ANSYS (a leading CAE ISV) worked with PTC (a leading manufacturing ISV), Flowserve, National Instruments, and HPE to demonstrate how a system-level simulation model of a complex network of operating pumps can diagnose and solve operating problems faster than before.[3]

Oil and gas: In E&P, HPC workloads such as seismic processing and reservoir simulation are being integrated with analytics, AI/DL, and IoT. This improves operational efficiency and delivers optimized hydrocarbon production and distribution at the lowest cost per barrel. For example, Schlumberger's DELFI provides a novel, open, and extensible cloud-based integration environment for experts across the geophysics, geology, reservoir engineering, drilling, and production domains. This cognitive approach amplifies the capabilities of every expert—enabling them to automate tasks, use learning systems, interrogate richer data sources, and deliver better and faster E&P decisions.[4]

Life sciences and health care: Combine 3D modeling and simulation tools, as well as integrate and analyze information across patient health records and many data pools (genomics, imaging, clinical, behavioral, drug efficacy, claims, and more) for better health care. What's more, it is possible to plan therapies, practice surgical procedures, improve medical device design, deliver better patient outcomes, and help transition the health care industry to value-based care. To enhance personalized medicine, Stanford University, Dassault Systèmes (ISV), UberCloud (ISV), HPE, and Advania (CSP) are collaborating on the Living Heart Project, focused on personalized medicine for cardiovascular care.[5]

Financial services: Market activities require numerous complex financial simulations that involve several thousand calculation tasks, taking from a few seconds to several minutes each to finish. These computations typically run on large-scale HPC or grid infrastructures and require hundreds of megabytes of data such as the historical values of equity shares over several years. But each task usually uses only a small portion of that data and must be integrated with risk analytics so that financial firms can better adhere to an avalanche of stringent and complex regulatory requirements.[6]

Société Générale, a French multinational banking and financial services company partnered with Qarnot Computing and the Microsoft® Azure team to build a new financial simulation platform—Service Fabric. The new platform is flexible, scalable, environmentally responsible, and designed to support the growth of the bank's business in a rapidly changing economy.

With Service Fabric, Société Générale built a robust, stateful microservices architecture in no time, giving them more time to focus on business needs and rely on the platform for resiliency, load balancing, and scalability. Now that the simulation platform is in production, the team is focusing on integrating new types of simulations and scaling the platform to handle them. The goal is to enable more customer applications to move away from legacy systems.

With HPC systems running into a more complex environment as illustrated with the use cases here, deploying, managing, and using HPC is becoming more challenging. Hence, many corporations are considering Hybrid HPC, which leverages cloud-based HPC options to augment their on-premises HPC capabilities.

[2] Market-leading HPC solutions for manufacturing—**h20195.www2.hpe.com/v2/Getdocument.aspx?docname=a00021803enw**

[3] Creating a Digital Twin for a Pump—**ansys.com/About-ANSYS/advantage-magazine/Volume-XI-Issue-1-2017/creating-a-digital-twin-for-a-pump**

[4] DELFI Cognitive E&P Environment—**software.slb.com/delfi**

[5] The Living Heart Project Wins Three Prestigious Awards for HPC Simulation—**hpcwire.com/solution_content/hpe/government-academia/living-heart-project-wins-three-prestigious-awards-hpc-simulation/**

[6] Société Générale's complex financial simulation platform expands on Azure Service Fabric architecture—**customers.microsoft.com/en-us/story/societe-generale-complex-financial-simulation-platform-expands-on-azure-service-fabric-architecture**

Why HPC cloud and what is Hybrid HPC?

The key challenges faced by enterprises with HPC include:

- To rethink the data center strategy essentially by going off-premises and handing over the data center operations to co-location providers and cloud providers. According to Gartner, by 2025, 80% of enterprises will have shut down their traditional data center versus 10% today.[7]
- To be more agile by leveraging built-to-purpose IaaS and PaaS technologies inherited from the cloud—the use cases depicted in the previous section require rethinking what is expected from an HPC cluster. More than ever before, HPC clusters have to deliver simulation, analytics, and AI insights to multiple information system domains such as CRM manufacturing, logistics, field operations, and more.
- To be less exposed to growing up-front investments—with the expansion of HPC across the product lifecycle, companies need models that allow a reduction of up-front investments while helping to speed-up deployment. Consequently, consumption-based payment models are becoming essential for on-premises deployments in the data center and at the edge. In this scenario, companies can leverage an on-site buffer to scale up or down on demand, paying only for the capacity they use.

The bottom line: Handing over data center operations to a third party, deploying IaaS and PaaS technologies with HPC clusters, and moving to asset-light operations with consumption-base models are definitely cloud attributes that can be summarized by the term Hybrid HPC.

Introducing Hybrid HPC

Hybrid HPC leverages cloud-based options to augment on-premises HPC capabilities. A Hybrid HPC cloud (Figure 2) provides a customer with a choice of deployment options and can overcome limitations of public clouds while retaining all the benefits. It extends very familiar on-premises HPC systems with a flexible private cloud infrastructure (on-premises or off-premises) including IaaS and PaaS features.

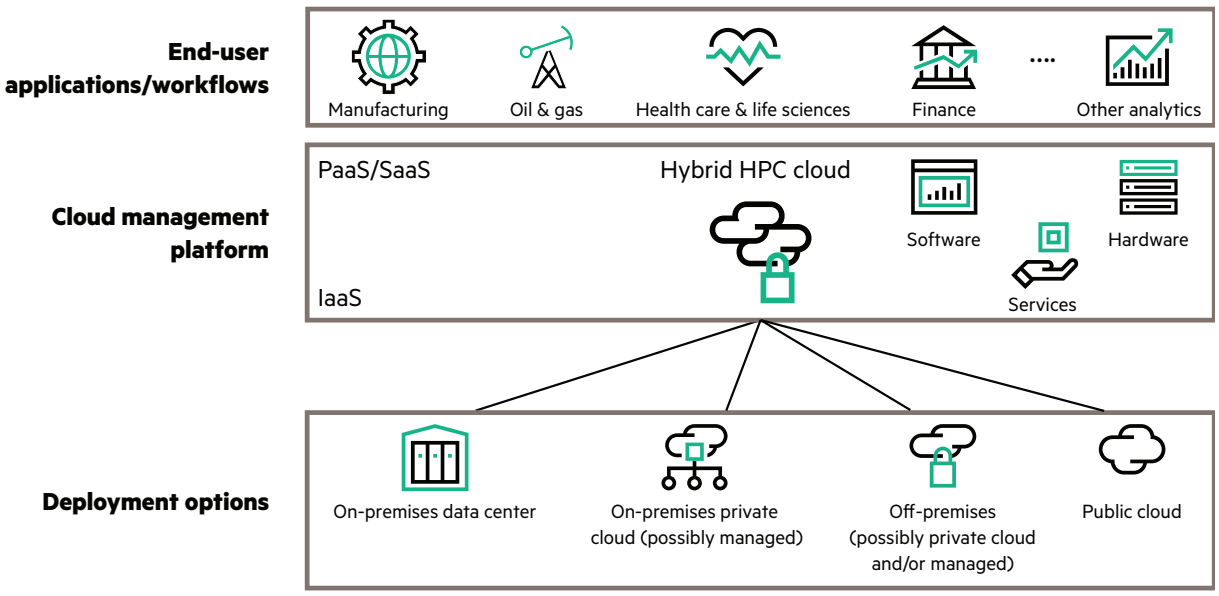

Figure 2. Hybrid HPC deployment models across industries

[7] The Data Center Is Dead, and Digital Infrastructures Emerge—**gartner.com/doc/reprints?id=1-4WDVFD0&ct=180417&st=sb**

With Hybrid HPC, customers—across many industries—benefit from an unprecedented level of flexibility, automation, and economics without compromising reliability, security, compliance, and performance for the entire HPC workflow. In short, customers can get the benefits of both private and public clouds.

Compared with a public cloud-only option, Hybrid HPC can be customized to specific workloads/applications to further:

- Reduce the total cost of ownership (TCO) for HPC workflows because private cloud systems can be precisely configured to maximize performance. The processor, RAM, storage, and even other system software can be highly customized. This is not always possible in a classic public cloud environment.

  Typically, public cloud providers only offer predefined configuration options. The options are a specific number of cores, GBs of RAM, and storage. The user has to choose a configuration that meets their minimum needs in one of these three options. For example, if the user needs 8 cores but less RAM and storage, the user has to compromise. They may have to choose a predefined configuration that has 8 cores, but the RAM and storage in the chosen configuration may be much more than what is needed. So, despite efforts to minimize unnecessary resources, predefined increments could result in configurations with superfluous IT infrastructure resources. Excess resources, in turn, incur an unnecessary additional cost for running the workload.

  In addition, as HPC becomes more data intensive and requires low-latency access to large data volumes, many cloud managed service providers with expertise in secure network connectivity and management can offer better prices compared to public cloud providers.

- Enhance security and compliance to protect highly sensitive information and workloads using a purpose-built HPC multi-tenancy as a core part of the HPC infrastructure. Multi-tenancy allows multiple users (tenants) to share the same physical environment while ensuring complete separation between them. Each tenant's data is isolated and remains invisible to other tenants. This helps address a customer's unique security, industry, and compliance requirements. In a public cloud environment, it is typically a challenge to address these highly unique customer requirements.
- Deliver higher value PaaS/SaaS capabilities to build, customize, and deploy HPC workflows. HPE and many managed service providers have deep HPC, industry, and application expertise to help customers in their migration, deployment, and access to these higher value cloud services.

By endowing or complementing on-premises systems with cloud-like technologies, Hybrid HPC can improve user experience, agility, and economics for all HPC customers. Its use is growing.

There are several options to deploy HPC. Pros and cons of each major option are discussed here starting with the on-premises data center option at one end of the spectrum followed by public cloud at the other end.

Comparing Hybrid HPC options

On-premises data center
The entire HPC infrastructure is all on-premises. Usually, most customers often optimize their workflows and typically get better predictable performance. Unlike on many "classic" public cloud providers, costs associated with data transfer, assured deletion of data, and runaway jobs are not a concern. But time delays to deliver the facilities, labor, and energy often inhibit HPC growth.

Using Tier-1 public cloud providers including HPC-specific technologies
Public clouds are designed to work across a broad range of general-purpose workloads including email, messaging, collaboration, calendaring, and so on.

In some cases, using the public cloud for HPC or compute-intensive applications can be good enough. Three types of HPC workloads with no legal constraints to run in the public cloud can easily leverage it with minimal adaptations—massively parallel applications, small-scale parallel applications (usually less than one socket), and parallel applications with small data sets.

Altogether, these three classes of applications represent a significant portion of HPC workloads. In addition, public clouds also provide the following benefits:

- Unified, location-independent platform for data and computation, pay-per-use, and availability for small organizations and individual end users
- Linear scaling with parallel execution for some workloads
- Ability to quickly scale to a large number of compute resources when needed
- Basic infrastructure deployment and management complexities handled by a cloud provider
- Better collaboration between engineers and scientists with a centralized computing environment

Despite the many benefits of public clouds—even with specialized public cloud with HPC technologies such as accelerators and high-performance interconnects—there are many shortcomings.[8]

Here are some reasons why it is hard or expensive to customize public clouds to handle the diversity and complexity of many industry-specific, data-intensive HPC and AI/ML workloads:

- **Application performance** is critical for competitiveness and financial results and is bottlenecked because of latency, memory, I/O, or restrictive service-level agreements (SLAs).
- **ISV licensing costs** are often the largest cost component of HPC solutions. Many leading ISVs and HPC-specialized ISVs have agreements with public cloud providers. However, very few have highly flexible pay-per-use licensing models.
- **Data mobility** in and out of the cloud is costly and cumbersome, as many HPC applications require low-latency access to major data feeds. This promotes adverse data lock-in and increases the time to deployment for new environments.
- **Mapping and tuning HPC workloads** on a specific public cloud require additional special skills. This could be hard and expensive for current HPC customers without compromising on security, supervision, administration, networking topology, and more.

[8] Optimal Workload Placement for Public, Hybrid, and Private Clouds, Intel white paper, 2017

HPC private cloud on-premises
This option converts an on-premises data center, retaining all its advantages, to a private cloud. It provides benefits of using IaaS and PaaS for more flexible deployment and operating models.

IaaS is hardware (compute, storage, and network) delivered as a programmatic service. IaaS gives developers more control over the entire application stack, or when the application requires isolation from other applications within the stack. In addition, an IaaS implementation exposes infrastructure services as application programming interfaces (APIs) and provides a user interface that makes it easy for developers to consume compute, network, and storage.

PaaS provides a pre-provisioned environment with an operating system, abstracted middleware, and infrastructure that allows developers to rapidly deploy applications without having to procure or provision servers. PaaS enables rapid application deployment through self-service, on-demand tools, resources, automation, and a hosted platform runtime container in private or public clouds.

It frees users from having to install and manage all the hardware and software to develop or run new applications and provides flexible deployment and operating models.

HPC off-premises (possibly managed services and/or private cloud)
As described in the previous section, off-premises operations have a lot of benefits by providing the data center capacity needed that meets workload placement and sustainability imperatives.

Off-site co-location with a trusted managed service provider can reduce HPC management costs with an operational expenditure (OPEX) model. The model scales to meet new business requirements and provides better energy capacity, redundancy, security and data protection, reliability, uptime, as well as disaster recovery and business continuity. It also helps customers with 24x7 support, maintain service provider diversity, and take advantage of volume pricing without having to negotiate and manage multiple contracts and service agreements.

Hybrid HPC cloud is growing

IDC estimates that worldwide whole cloud revenues will reach $554 billion in 2021, more than double those of 2016[9] and CSPs will account for 76% of cloud-related infrastructure hardware and software spending. SaaS is the largest cloud category as enterprise software providers increasingly adopt cloud-based models. Gartner predicts that the high-profit margin SaaS market[10] could hit $117.1 billion in 2021.[11]

The HPC market is also growing. Hyperion Research predicts a compound annual growth rate (CAGR) of 5.8% for HPC, from $11.2 billion in 2016 to $14.8 billion in 2021. Additionally, cloud-based HPC could grow to over $10 billion by 2021.[12]

[9] Worldwide "Whole Cloud" Revenues Will Reach $554 Billion in 2021, According to a New Forecast from IDC—**idc.com/getdoc.jsp?containerId=prUS43285217**

[10] Taking software-as-a-service at more than face value—**morganstanley.com/ideas/saas-valuations**

[11] Plenty of growth in store for cloud computing, according to Gartner, will hit $186B in 2018—**geekwire.com/2018/plenty-growth-store-cloud-computing-according-gartner-will-hit-186b-2018/**

[12] Market Forecast—Worldwide Regional HPC Market Forecast and Country Snapshot, 2016–2021, Hyperion Research, August 2017

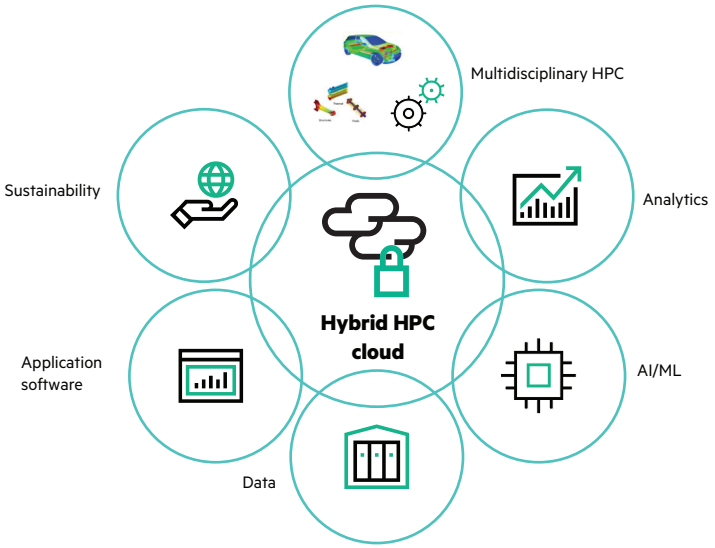

Figure 3. Growth drivers of Hybrid HPC clouds

Some key growth drivers of the cloud technologies and deployment models in HPC (Figure 3) across many industries include:

1. The rise of digital platforms. Hybrid HPC models with off-premises operations and IaaS are suitable to provide compute and storage horsepower to digital platforms.
2. The growing use cases discussed in previous sections combining simulation, analytics, and AI are captured within the HPDA segment; HPDA is growing between 17% and 25%[13] and includes the even more rapidly growing areas of AI, ML, and DL. IDC predicts that global spending on cognitive and AI solutions will increase at a CAGR of 55.1% over the next few years, exceeding $47 billion by 2020.[14] HPDA applications are a strong driver for the adoption of PaaS in HPC.
3. Pressure to be environmentally conscious and promote sustainability through green HPC.[15]
4. High growth and profit of leading HPC-related ISVs such as Dassault Systèmes, Siemens PLM, PTC, GE Digital, Autodesk, ANSYS, and their rapid adoption of SaaS models.

High-level architecture of the Hybrid HPC cloud from HPE

As the market leader in HPC with a 36.8% market share,[16] HPE is delivering a portfolio of Hybrid HPC components (based on Intel's innovative HPC foundation). The portfolio consists of data center options delivered by PRSP and co-location partners, improved agility with IaaS and PaaS options, and a consumption-based model.

Intel's innovative HPC foundation, based on the Intel® Xeon® Scalable processors, includes critical platform innovations in memory, storage, and acceleration technologies to address the complex spectrum of diverse HPC workload requirements. Given the criticality of HPC fabric in scale-out deployments, Intel Omni-Path Architecture (Intel OPA) provides low-latency interconnect for a scalable performance of multinode environments, such as AI training applications.

While large public clouds are an option for some organizations, other users have unique requirements sometimes better served by specialized providers with domain expertise. HPC users want easy access to high-performance bare-metal clusters and the convenience of on-demand provisioning, but they want the environments to be tailored to their unique workloads. They need an environment that supports a variety of applications as well as multiple deployment models including on-premises, off-premises, and hybrid clouds. They also need to deliver a variety of HPC-as-a-service offerings ranging from simple IaaS to PaaS to complete SaaS solutions.

[13] Market Forecast—Worldwide Regional HPC Market Forecast and Country Snapshot, 2016–2021, Hyperion Research, August 2017

[14] Worldwide Cognitive Systems and Artificial Intelligence Revenues Forecast to Surge Past $47 Billion in 2020, IDC, 2016

[15] Why you'll be using liquid cooling in five years—**hpe.com/us/en/insights/articles/why-youll-be-using-liquid-cooling-in-five-years-1710.html**

[16] Hyperion Market Update: "Decent" Growth Led by HPE; AI Transparency a Risk Issue—**hpcwire.com/2017/11/15/hyperion-hpc-market-update-decent-growth-led-hpe-ai-transparency-risk-issue/**

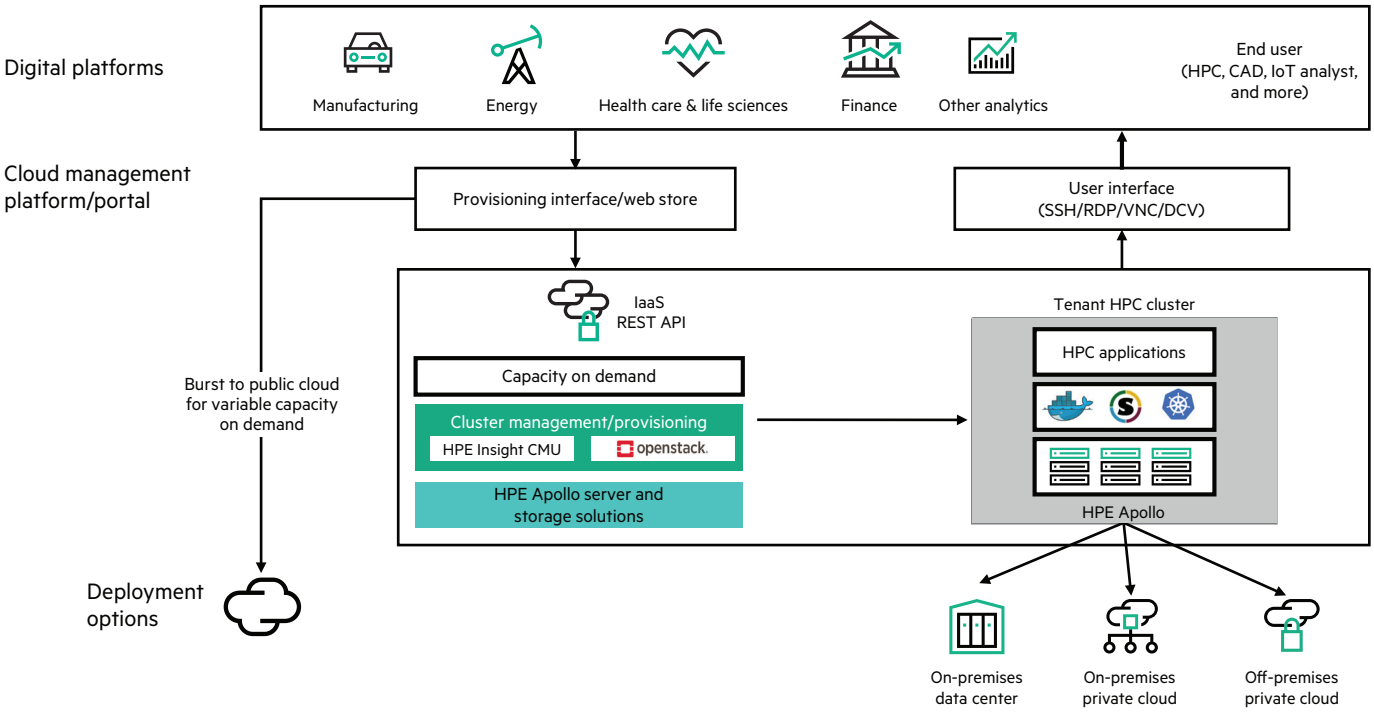

Figure 4. High-level view of Hybrid HPC cloud solution

To address these challenges, HPE has developed a Hybrid HPC IaaS solution and reference architecture, which is purpose-built to deliver a cloud experience to IT administrators and end users. The solution is based on various innovative technologies from Intel.

A high-level view of this solution is shown in Figure 4. Today, the solution is deployed using HPE Insight Cluster Management Utility (CMU) along with various open-source components.

As the solution evolves, it leverages additional OpenStack® and Kubernetes components to provide cloud operators with greater flexibility. The Hybrid HPC IaaS solution can:

- Auto-provision state-of-the-art infrastructure for modern HPC workloads on demand
- Present a variety of service entry points including application-level (SaaS), platform-level (PaaS), and infrastructure-level (IaaS) offerings
- Provide performance and security comparable to best-in-class on-premises clusters
- Support a variable cost consumption-based model attractive to HPC users

The following HPC cloud operators can benefit from this solution:

- Cloud-service providers offering hosted infrastructure or application services
- Systems integrators delivering HPC-managed services
- Private or public enterprises employing a shared-services model to deliver private cloud to groups within their organizations
- HPC, analytics, or AI ISVs delivering SaaS models either themselves or in partnership with one or more cloud providers

Hybrid HPC IaaS solution from HPE puts service provider partners and customers firmly in control of their environment. Cloud operators can choose multiple entry points, relying on HPE for infrastructure only, taking advantage of various aspects of the IaaS solution to deliver self-service IaaS, PaaS, or SaaS environments. Partners can use the solution to deliver their own software solutions or work with HPE partners to deliver full-stack SaaS solutions including third-party ISV application services.

Getting started on the Hybrid HPC cloud journey with HPE

Through any of the cloud deployment options, HPC customers can begin their value-accretive cloud journey now. They can either transform their on-premises HPC infrastructure into a private cloud or expand their operations through an off-premises private cloud as an initial first step on their cloud journey. This enables them to get many public cloud attributes (Table 1), without the drawback of the public cloud.

Table 1. Key benefits of starting the cloud journey using Hybrid HPC cloud offerings from HPE

HPE offerings	Related public cloud attributes
HPE Datacenter Care for Hyperscale	Single point of operations **HPE Datacenter Care** offering provides proactive support and an assigned team to simplify the operations lifecycle of the on-premises, hosted, or managed HPC system, tailored to fit the way one specific company operates IT.
Hosting with one of the **HPE Partner Ready Program** partners	**No data center to manage** Data center upgrades (power and cooling density, floor space, and more) on schedule for HPC systems are often the main challenge for most IT departments. HPE can work with IT departments to define the best strategy for either an on-premises data center upgrade or using a third-party data center provider registered within the **HPE Partner Ready Program**.
HPE GreenLake Flex Capacity	**Pay as you grow** Consume HPC in the data center, saving on the cost of overprovisioning, with reserved capacity ready for growth or unpredictable spikes in use. **HPE GreenLake Flex Capacity** has also already been adopted by many service providers.
IaaS from HPE	**Flexibility** On-premises HPC clusters have typically been designed in a monolithic way so that nominal performance can be provided. Nevertheless, customers are looking for ways to match new and fast-growing business needs, including a growing number of HPC users with no HPC skills and new types of simulations that require operation environments other than bare-metal compute. These new and fast-growing businesses require the capability of dialing up and down or repurposing system partitions on-premises or off-premises.
Multicloud enabled by cloud management platform (CMP) solutions	**Managing multicloud environment** All companies expect public cloud bursting to handle peak loads. To meet this goal, the public cloud partitions must be integrated into the existing end-user environment so that the legacy application workflows remain unchanged when bursting. Middleware such as **Micro Focus Hybrid Cloud Management (HCM)**, Cycle Computing CycleCloud, or Rescale ScaleX can provide flows and connectors for bursting from on-premises HPC to a public cloud. CMPs provide a means for a cloud service customer to bring together disparate cloud environments and to manage the deployment and operation of applications and associated data sets across multiple cloud service infrastructures.

Market-leading HPC systems and software solutions

HPE provides unified compute and storage solutions that simplify system and data management, reduce costs and complexity, and scale to deliver excellent performance for next-generation HPC, HPDA, and AI/DL workloads. These solutions[17] have three pillars—partnerships, platforms, and services.

Platforms: HPC systems from HPE provide a choice of servers (for example, HPE Apollo family) that incorporate Intel HPC foundation components ensuring a balanced, high-performance, and scalable HPC environment. Among the key technologies components are:

- **Intel Xeon Scalable processors** offer up to 28 cores and highly enhanced per core performance, significant increases in memory with up to six memory channels, and I/O bandwidth expanding up to 48 PCIe lanes. These innovations combined with HPC advances such as Intel Advanced Vector Extensions (Intel AVX 512) deliver up to 1.73X HPC performance compared to prior generations. When using the latest Intel optimized AI libraries these processors offer up to 127X improvement in training throughput over previous generation processors.
- **Intel Xeon Phi** processors provide ultra-wide vector capabilities in highly parallel computing with up to 72 powerful and efficient cores.
- **Intel OPA switches** and PCIe adapters are ideal for large-scale HPC applications or neural network training requirements delivering 100 Gb node-to-node connectivity, low latency, and the ability to scale to very large numbers of nodes. The use of Intel OPA fabrics provides HPC cloud operators a distinct advantage over competitors that typically use slower 10GbE or FDR InfiniBand interconnects.
- **Intel Optane SSD** is a new technology in the data center that is extremely responsive under any load and delivers the low latency, high performance, and high endurance demanded by HPC clusters.
- **Intel 3D NAND SSDs** provide greater storage performance with higher density and reliability than hard disk drives. Both SSD technologies deliver the high-performance and reliability that users expect with increased density and reliability, making them excellent choices for cloud operators.

The high-value cluster software from HPE/partners drives additional performance.

HPE Apollo Gen10 servers are high-density, energy-efficient, and memory-intensive systems that scale from small (**HPE Apollo 2000**) to medium (**HPE Apollo 6000**) to large supercomputers (**HPE SGI 8600**). HPE Apollo systems combine a modular design with power distribution and air/liquid cooling techniques, providing more outstanding performance per square foot.

Services: Purpose-built HPC solutions and advisory services deliver faster time to value. HPE Centers of Excellence (COE) and 24x7 support ensure efficient operations with minimal disruptions.

Implementation scenarios with Hybrid HPC cloud solutions from HPE

There are several options to implement an HPC cloud. Each scenario has its benefits and challenges. HPE and partners are well positioned to create the Right Mix of Hybrid HPC clouds that is tailored to each company's needs. The different scenarios are listed as follows along with their salient characteristics to be compared with the first one.

[17] Market-leading HPC solutions for manufacturing—**h20195.www2.hpe.com/v2/Getdocument.aspx?docname=a00021803enw**

Scenario 1: Bursting to Tier-1 public cloud provider (with or without HPC-specialized hardware)

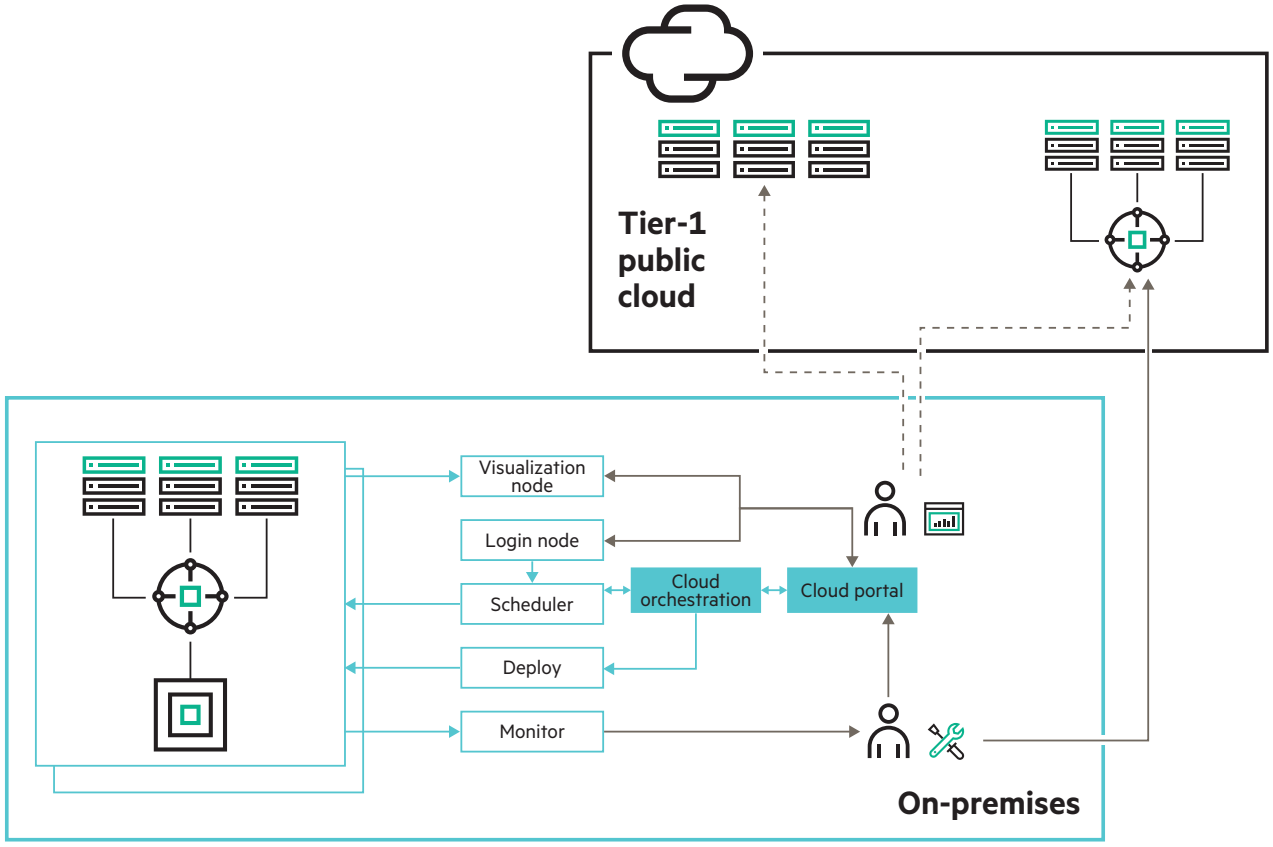

Figure 5. Bursting to the public cloud

Bursting to a Tier-1 public cloud provider offers flexibility and the capability to meet new and unexpected demands, but it is only suitable for a limited number of massively parallel HPC workloads (even with HPC specialized hardware).

While the public cloud offers a great amount of resources on demand, some factors can be limiting for those who need repeatability, predictability, and reliability through SLAs. Cost predictability can also be a challenge when using public clouds for bursting because hidden costs around bandwidth, data transfer, and storage often appear afterward.

Data security is often a key concern because it must be handled by the end-user customer, it can breach local regulation, and it can be very time-consuming and expensive. Another drawback of bursting to the public cloud is that the customer still has HPC assets on-premises. Other common issues that arise could be the latency when using remote visualization, as well as the time to transfer data to and from the cloud.

This scenario is recommended only for certain workloads and where a part of the HPC pipeline lends itself to this sort of environment due to the factors listed previously. For example, a bank could use cloud bursting for some embarrassingly parallel Monte Carlo simulations on data that isn't sensitive. Some steps of risk analytics typically require short bursts of compute capabilities to meet compliance reporting requirements. Using on-premises capacity for these types of intermittent calculations could be overdoing it.

Scenario 2: Bursting to regional HPC cloud service provider

Alongside public cloud Tier-1 providers, there are regional service providers, whose core business relies on providing built-for-purpose bare-metal HPC solutions.

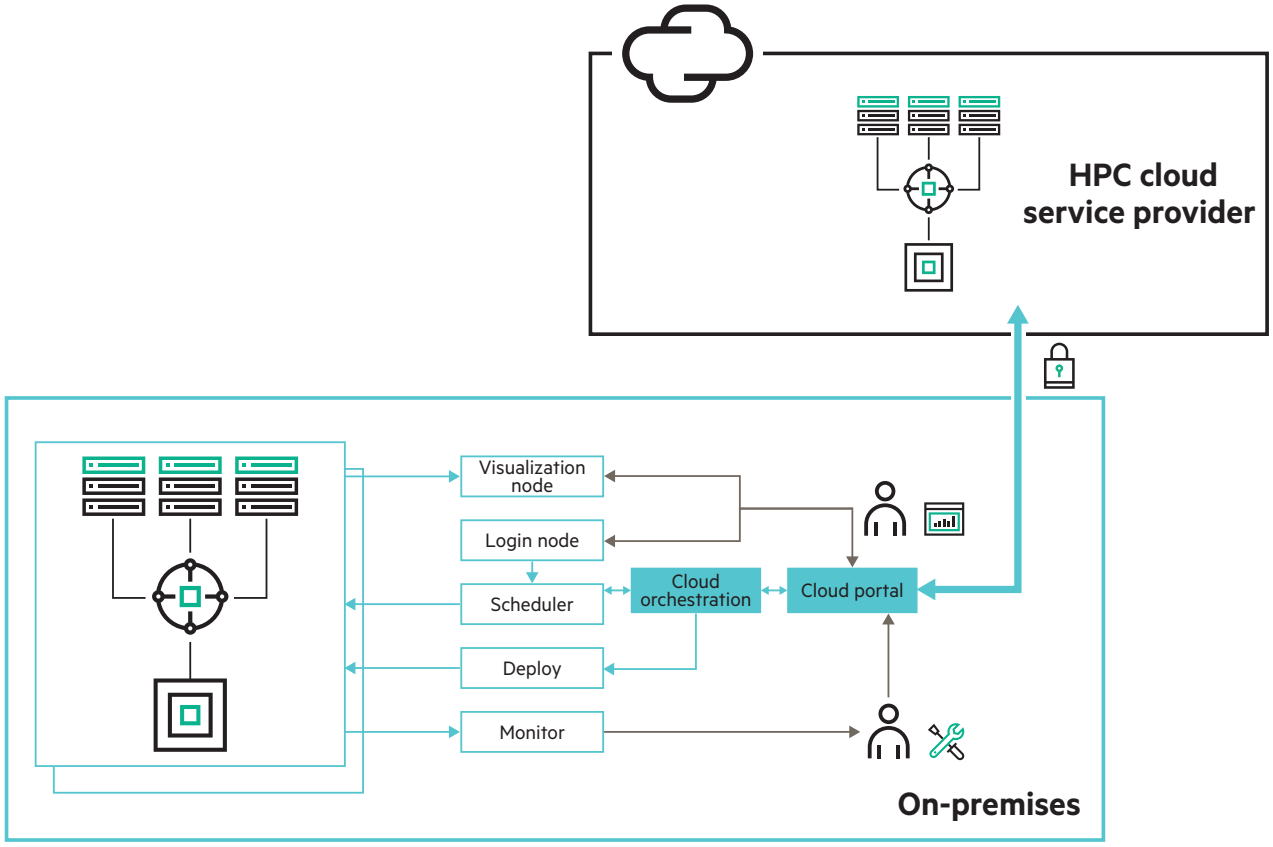

Figure 6. Bursting with a regional HPC cloud service provider

Bursting into a regional HPC cloud service provider offers the same advantages as bursting to a Tier-1 public cloud provider, but it:

- Allows for a more secure environment
- Enables the pre-sharing of data for speed
- Provides a bare-metal option for HPC hardware; offering predictable performance for HPC jobs that can be repeated in a reliable manner

Such providers offer SLAs where the time to execute the compute job and the price are predictable and controlled by the user, with support from an HPC support team, operated by HPC experts.

Regional HPC cloud providers offer best-in-class security ensuring job and data isolation, usually at the hardware level. This allows network segregation and dedicated resource for the lifespan of the jobs. A critical component when leveraging sensitive data is when the handling is controlled by local legislation such as in the highly regulated health care and financial services industries.

Moreover, HPC tenants are accessed through their own login node that runs on physical hardware, which is the single-entry point to the tenant's environment.

HPE recommends service providers that are ISO 9001 and ISO 27001 certified and have implemented ITIL® processes in their IT and hosting operation workflows. This ensures that the provider has an auditable and proven environment and has practices in place to detect and mitigate breaches or incidents that impact the security of the environment.

Service providers also have the advantage of having diverse and resilient backbone connections either onto the internet or with the possibility of connecting directly onto a customer's corporate or research network of choice. This reduces the latency and cost of data movement of potentially large data sets. For example, for next-generation sequencing (NGS) in the life sciences industry—where petabytes of data must be processed—many analytics algorithms benefit from the fast and scalable compute on the cloud with co-located storage. A genomics cloud service provider can not only provide better agility and economics but can also provide better security for sensitive patient data.

Scenario 3: Hosting and bursting to regional HPC cloud service provider

In most cases, enterprises choose an HPC cloud hosting provider to reduce the overhead of hosting, management, and support of the HPC environment. This HPC cloud provider should provide a flexible solution that is suited to each customer's needs. While some only need floor space and power, others want to have management further up the HPC stack, and some want it all the way up to the HPC application layer and licensing support.

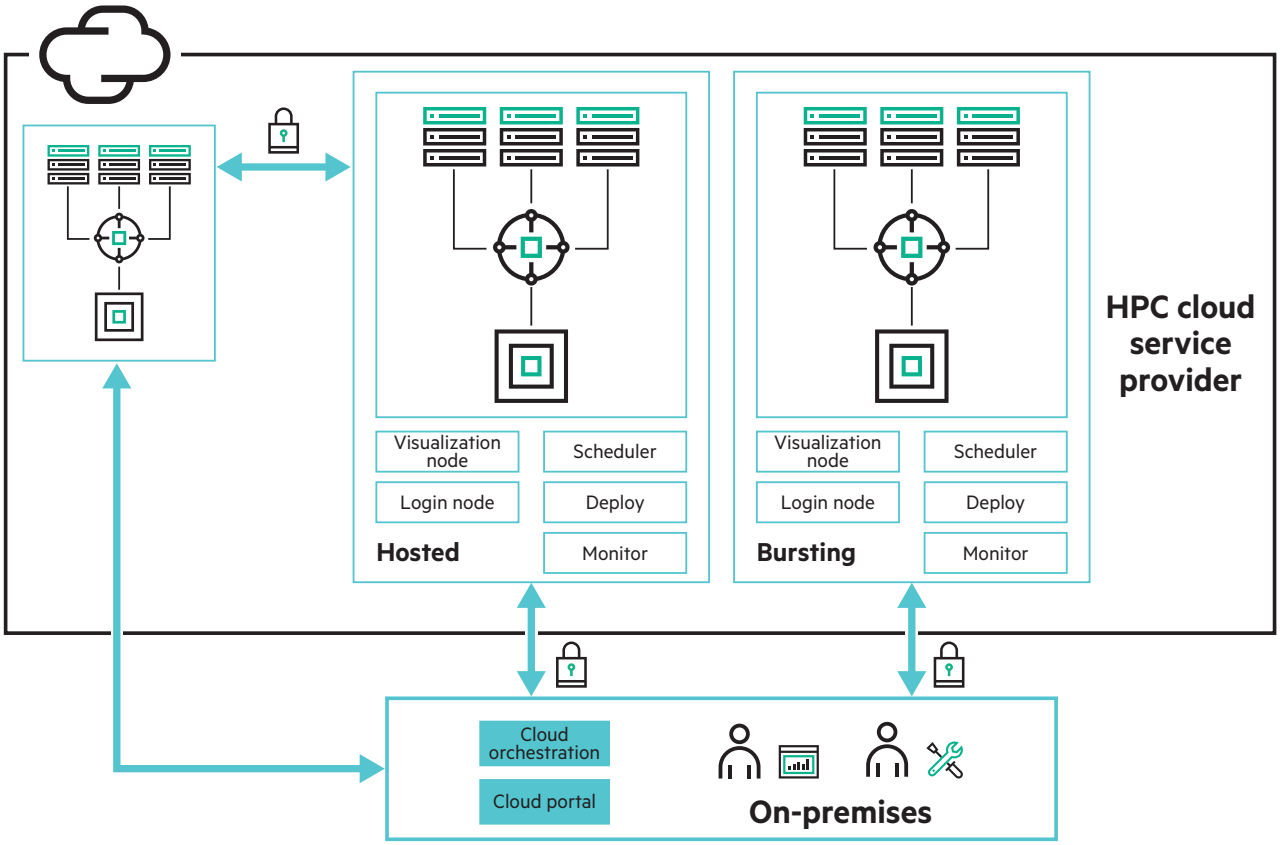

Figure 7. Hosting and bursting to regional HPC cloud service provider

Some regional HPC cloud service providers can offer the best of both worlds, namely the cost-effectiveness and control of an on-premises HPC cluster, along with the scalability of an HPC cloud available for bursting. In this case, the cost, performance, and availability become predictable.

Another benefit is perfect consistency between the HPC operating environment whether hosted or with bursting. The same software images can be used in both environments.

This is the preferred scenario especially for CAE where there are many diverse applications (with expensive licensing costs) spanning structures, fluids, electromagnetics, and other disciplines, each with its own unique infrastructure needs.

Hybrid HPC cloud solutions implementation from HPE—Customer examples

By providing a range of innovative Hybrid HPC options and high-value services, HPE and its partners are driving and accelerating the digital transformation of customers worldwide. Here are some examples:

Siemens AG: Using HPE GreenLake Flex Capacity on their HPC/PLM private cloud

HPE GreenLake Flex Capacity: The **Siemens Energy example** shows how Siemens Energy uses HPE GreenLake on an internal cloud, which supplies workstation and HPC capabilities to a large pool of engineers across the U.S., Europe, and Asia (18 locations). The cluster is actually split in two locations—the U.S. and Germany.

HPE GreenLake gives Siemens the agility, pay-per-use billing, and the rapid scalability of a cloud-like model in the privacy of Siemens controlled data centers. It allows them to quickly access more compute power when required with the very low latency of being on-premises. This supports Siemens' cloud strategy while providing the advantages of on-premises IT and allowing Siemens to remain in control of all technical configurations and specifications.

> "HPE GreenLake Flex Capacity is a unique way of consuming IT-Infrastructure as an OPEX model and provides Siemens a flexible experience with the benefits of on-premises IT. We only pay for what we consume."
>
> – Dr. Henner Tilgner, Head of Global Engineering Computing, Siemens AG

HudsonAlpha: Accelerating the search to cure humankind's most insidious diseases

The **HudsonAlpha case study** highlights how three billion DNA base pairs are processed rapidly with a private cloud, powered by HPE Synergy solutions to help deliver on the tremendous promise of genomics research to cure cancer and other menacing diseases.

HudsonAlpha needed a solution that would meet its demands of intensive research, collaboration, and analysis, as well as make it a pioneer in scientific breakthroughs. The organization required a robust infrastructure that was powerful enough to meet researchers' needs and accommodate the immense rate of growth, scale of data, and be manageable enough to run cost-effectively.

HPE Synergy has helped HudsonAlpha address these scalability challenges and spiraling costs, through understanding the benefits of a composable infrastructure. It has distinct advantages over older technologies as it's easier, faster, and more efficient. With HPE Synergy, overprovisioning and stranded capacity can be decreased because flexible resource pools of compute, storage, and fabric are shared and can be quickly configured for any workload.

Resources can be managed through templates and dynamically allocated to meet the needs of any application. Increased operational efficiency and rapid deployment of IT resources are achieved with a unified API providing full programmability to provision physical infrastructure as easily and quickly as virtual machines.

HudsonAlpha initially considered the public cloud, however, HPE Synergy proved superior to the public cloud option by providing the benefits of both on-premises and the cloud with its Hybrid IT approach. This helped to reduce the amount of time it would take to migrate the company's existing data to the cloud. It would also improve application performance and help eliminate the transfers that some cloud providers charge for data uploads, which can completely negate any benefits of **moving data to the cloud**.

By retaining the best of the cloud and traditional IT, the architecture was able to offer additional benefits—importantly, the ability to scale with increased workload demands.

The key benefits included:

- An annual cost benefit of $2.23 million, estimated by IDC, was achieved by adopting HPE Synergy and 48% lower cost of operations projection than that of an alternative legacy environment.[18] The bulk of this benefit was tied to improved business productivity.
- Support cost savings—33% fewer staff hours for management and maintenance tasks, as HPE Synergy was able to be managed and supported more efficiently than its legacy infrastructure.[19]
- Ability to scale up its compute capacity without increasing the number of hours that IT staff spend managing the platform.
- Significant reduction in the additional costs that would have been incurred, related to storage and licensing, if an alternative approach to solving its scalability challenges had been opted for.

"Immediately it provided an increase in speed. We're able to bring all the needs of our 34 companies and 200 researchers together and use our resources in the most efficient manner possible."

– Jim Hudson, Co-founder and Chairman, HudsonAlpha Institute for Biotechnology

The HPE advantage

As more HPC customers embark on their cloud journey, they need a reliable partner with deep HPC, Big Data analytics, and AI/DL expertise. Working with Intel, HPE provides a comprehensive, purpose-built **Hybrid HPC solutions** portfolio jointly with ISVs and CSPs.

Unlike "classic" public cloud providers, HPE customers across many industries can **customize** their specific applications and workflows to drive outstanding performance, scale, consumption-based operations and economics from cloud technologies, and still benefit from built-to-purpose HPC infrastructure.

Learn more at
hpe.com/us/en/solutions/hpc-high-performance-computing.html

[18], [19] HudsonAlpha Maximizing the Impact of Its Genomics Research with HPE Synergy, IDC ExpertROI® SPOTLIGHT (Sponsored by Hewlett Packard Enterprise and Intel), June 2017 **idcdocserv.com/US42599317**

Sign up for updates

a00029377ENW, October 2018, Rev. 1