Hewlett Packard Enterprise | AMD

Business white paper

Check if the document is available in the language of your choice.

SUPERCOMPUTING TECHNOLOGIES AT ANY SCALE

Embracing a new era of compute

TABLE OF CONTENTS

We have entered an age marked by rapid growth, explosive amounts of data, and escalating competition. Organizations are working tirelessly to meet the demands of their industries, business requirements, and customer expectations, which are changing constantly. In addition, the emergence of disruptive technologies is causing continual shifts and driving progress across every sector, putting tremendous pressure on organizations to transform their operations from the ground up.

In this dynamic global economy, success is built on innovation. Organizations have an insatiable need for technology that delivers the right level of performance to accelerate critical insights, as well as boost their profits. The demand for more performance is mounting as workloads become increasingly more diverse and complex. Technological advances have made it possible to capture valuable information and convert it into actionable insights to create greater efficiencies and tap into new sources of revenue. These trends have redefined what it means to operate, leading organizations to reconsider how they think, learn, collaborate, and engage the world. They are now asking different and bigger questions, and they need new capabilities to answer them.

Organizations are examining their rising compute requirements to create strategic road maps for the future. Data center transformation is a major factor that affects planning, IT, and every area of operations. Developing a data center strategy that includes a future state of IT infrastructure lays the groundwork for enhanced capabilities, performance, and efficiency. The greater efficiencies that organizations can produce, the more customers they can serve and the better outcomes they can achieve. Those that invest in cutting-edge solutions will create an infrastructure that is designed to adapt and scale, reduce complexity, and unlock a new kind of compute.

Supercomputing completely rethinks how technology comes together to conquer the greatest challenges of today and prepare for the even bigger challenges of tomorrow. However, the term supercomputing is often misconstrued, causing organizations to shy away from high-performance computing (HPC) opportunities. Many believe they lack the financial resources, expertise, or even the need to adopt the latest technologies. As a result, the benefits of supercomputing are being widely underutilized in a number of industries. It's time to demystify supercomputing and its groundbreaking applications.

Data and workloads are becoming too complex for traditional compute to handle. Supercomputing technologies can deliver greater compute power and speed to overcome the limitations of inflexible technology, so organizations can optimize an extensive range of tasks, from small-scale activities to high-end workloads. These capabilities are used to make sense of events that are too difficult to see or measure in real life. The purpose of supercomputing is to provide the compute capabilities to accelerate data-driven workloads. The data-intensive nature of new application requirements will require ongoing innovation for the foreseeable future.

Advancements in supercomputing technologies are changing the face of the modern data center, as organizations implement capabilities today that are relevant for years to come. Addressing important scientific, engineering, and business challenges depends on continued investments in supercomputing. Strategic adoption of these technologies gives organizations the potential to evolve and lead their industries.

CAPITALIZING ON COMPUTING TRENDS

Big Data, edge computing, artificial intelligence (AI), machine learning, virtualization, and simulation are just some of the developments that are revolutionizing how organizations operate and make decisions. These capabilities are enabled by the creation and accessibility of vast amounts of data, backed by robust computational power. IT infrastructure must be able to support a variety of HPC uses, including efficiently processing and analyzing data from emerging applications and endpoints. Moving forward, supercomputing technologies will play a crucial role in every operating environment, to help ensure that organizations can execute data-intensive workloads quickly, efficiently, securely, and with optimal price-performance.

Supercomputing technologies are continuously evolving. No amount of scientific discovery, product enhancement, or engineering achievement is final—it is merely the next step in ongoing transformation. The latest advancements in compute provide extreme power and agility that were previously impossible. Broader use of supercomputing technologies is helping organizations capitalize on these trends and work at their maximum potential performance.

Organizations can use the components of a supercomputing system but on a smaller scale—at their scale. The same features and infrastructure can be delivered in an accessible way, enabling data-intensive compute across all industries. These technologies are designed to run different workloads and promote high utilization of assets, exceeding the limits of what industry-standard clusters can do. By doing so, organizations can harness the right amount of compute for their data center requirements, regardless of the types of workloads they run. This creates an era of new opportunities for organizations of every size and scope to realize better outcomes and secure their futures.

EMBRACING A NEW KIND OF COMPUTE

One of the greatest benefits of supercomputing capabilities is the ability to transform rapidly and gain immediate return on investment (ROI). Supercomputing technologies can be deployed into production and adapted as needed. These architectures are built for continuous evolution, which provides an upgrade path for nearly unlimited flexibility across multiple use cases and sizes of deployment. To reach a new level of performance, organizations are moving away from outdated systems to supercomputing technologies that can run high utilization for various workloads.

Organizations can implement next-generation compute for a broad range of uses. Applications in manufacturing, financial services, energy and gas, life sciences, earth sciences, climate, and academia are all data-intensive and, therefore, benefit greatly from supercomputing delivered at their ideal scale. Despite the value that supercomputing technologies bring, many are struggling to take advantage. Organizations face several pain points as they begin their journeys to innovation:

- Accommodating the data explosion/new business models
- Affording up-front CAPEX
- Managing increasingly sophisticated workloads, including HPC and AI
- Providing the flexibility to execute escalating workloads
- Meeting data center power and cooling demands
- Leveraging on-demand or elastic consumption
- Working with limited IT expertise

Advanced compute applications have a ravenous appetite for data, power, and other resources. To overcome these obstacles, organizations need systems that are engineered to meet all their requirements and prepare for future demands.

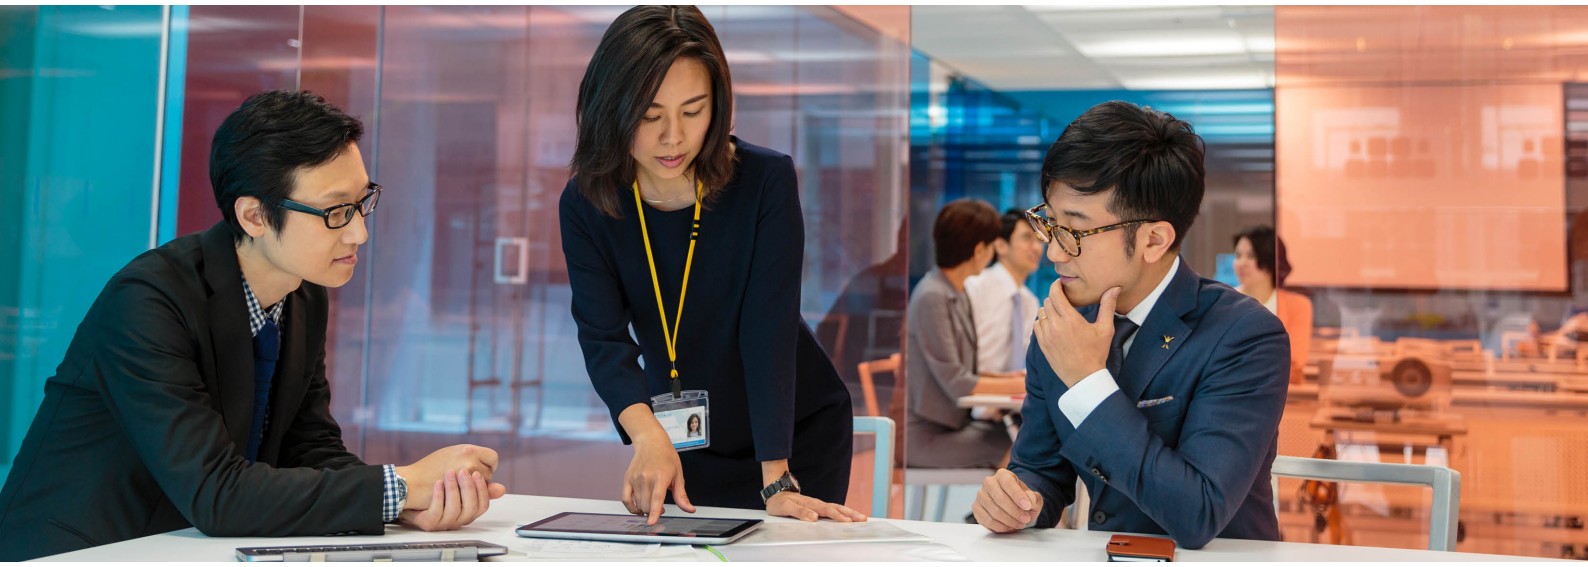

Supercomputing capabilities are helping organizations capitalize on the run-grow-transform (RGT) model. Many are focused on running—this means, completing the necessary tasks to keep their environment functioning. To simply run, organizations increase operational efficiencies to maintain their current IT capabilities. Expanding on these capabilities is essential to enhance how they work, enabling them to adapt to changing requirements as they grow. The best-run organizations are constantly searching for opportunities to realize efficiencies in how they run, so they can shift their budgets toward accelerating their digital transformation.

Investing in transformation is key to create new capabilities and drive value to new heights. By deploying a purpose-built system integrated with supercomputing technologies, organizations can achieve unmatched performance across all areas of operation. The next generation of compute is poised to accelerate innovation with a lower total cost of ownership (TCO) and increased ROI. Organizations that take advantage of these game-changing benefits will be positioned to work, compete, and excel in their industry's outcomes and secure their futures.

BRINGING SUPERCOMPUTING TECHNOLOGIES TO ALL

Hewlett Packard Enterprise and Advanced Micro Devices (AMD) are empowering organizations to enter the next frontier of compute. We are collaborating to jumpstart the journey to greater performance and groundbreaking results, introducing cutting-edge technologies that are engineered with organizations in mind. Our offerings include significant advancements in supercomputing to transform the data center and remain resilient for years.

Our solutions combine landmark HPE systems with the extreme processing capacity of AMD CPUs and GPUs. Comprehensive technologies, tools, and advisory services from HPE are designed to extend supercomputing capabilities from edge to cloud. We help ensure these solutions are integrated and tuned to create an environment that can be applied today and helps fuel progress as organizations continue to grow.

Together, we are breaking new ground in compute—with the capacity to build out, add features, and upgrade infrastructure when it is needed. Organizations that adopt these solutions are driving dramatic improvements across their operations, using the same feature sets as the largest supercomputing systems to accelerate technical, scientific, and business-related insights and secure better overall TCO.

HPE is a global technology leader in supercomputing, and in collaboration with AMD, we are committed to delivering revolutionary capabilities to help answer the most complex questions. Now, you can leverage these technologies to boost supercomputing innovation.

Industry-leading supercomputing infrastructure

As your workloads rapidly evolve, the ability to choose your architecture becomes critical. HPE Cray supercomputing line powered by AMD processors enables you to address today's supercomputing challenges. These powerful systems are the firepower of computing environments, allowing you to streamline workflows from modeling and simulation to analytics and AI. HPE Cray supercomputers deliver reliable and exceptional performance for your applications by supporting a mix of AMD CPUs and GPUs, multisocket nodes, and other processing device options that may emerge—all integrated into a single infrastructure.

Some of our key differentiating features include:

- Air-cooling and liquid-cooling options to use the highest performing parts
- High-performance interconnects
- Highly agile and efficient storage
- Unique system architectures incorporating both CPUs and GPUs

HPE Cray supercomputers fulfill a host of new requirements for both research and business organizations, to deliver insights and innovation on a scale that has rarely been seen before. These systems are based on an array of HPE and legacy Cray technologies, including new capabilities from the HPE Cray EX architecture. HPE is working to make supercomputing technology ready out of the box. Very soon, organizations will be running workloads and transforming faster at an exceptionally low TCO.

The flagship system of this portfolio is HPE Cray EX—a supercomputer that is built from the ground up for massive compute and efficiency. EX stands for exascale. This system provides end-to-end direct liquid-cooling (DLC) to extract excellent performance from high-bin CPU and GPU options, making it a critical building block of HPE Cray pre-exascale and upcoming exascale deployments.

HPE Cray EX is highly integrated, including power and cooling and high-speed HPE Slingshot interconnects with a cableless connection. Because the ability to innovate is key, the system can be upgraded according to your business requirements. The architecture can run proficiently with multiple generations of CPUs, GPUs, and interconnects. You have the option to make these changes or updates in the field, to simplify and enhance your deployments while extending your supercomputing capabilities. The HPE Cray EX is designed to support up to 256 dual socket nodes in a single cabinet, as well as support high-wattage CPUs and GPUs, so you can execute diverse workloads and build the data center of the future.

For organizations that demand the utmost performance and sustainability, this supercomputer supports green initiatives with liquid-cooling and zero moving parts for the best possible energy efficiency. DLC cold plates rapidly remove heat from system components without exhausting hot air into the data center, which lets you operate at maximum capacity with confidence and at optimal price-performance.

The HPE Cray supercomputer is the second configuration of this powerhouse portfolio. The architecture is a standard 19-inch rack configuration supercomputer in a 2U compute server. This system provides the choice to implement a smaller system with the same feature set as the impressive HPE Cray EX systems. For organizations that don't need a full-scale deployment, or are taking evolutionary steps on their supercomputing journeys, the HPE Cray Supercomputer is available in standard data center racks with both air-cooling and liquid-cooling options. All HPE Cray supercomputers come DLC equipped, so you can transition from air to liquid to boost your performance.

Accelerated compute

To help organizations solve their pressing problems, HPE Cray supercomputers are powered by the world's leading processors. AMD EPYC™ CPUs and AMD Instinct™ GPUs are designed together to handle large datasets at superior performance to speed compute workloads, data-intensive models, and analysis techniques. Used by some of the world's fastest, most scalable data centers and supercomputers, AMD processors can enable rapid innovation and results.

Many of today's systems are heterogeneous—incorporating a blend of CPUs and GPUs—which gives organizations the flexibility of design to accelerate and optimize how they work. AMD processors offer high core count, more agility, and up to 8 channels of DDR4-3200 memory in the same feature set, whatever the size of your system. With this configuration, you can continue operating as you differentiate and scale, maintaining consistent performance across all stages of deployment.

AMD EPYC CPUs address use cases across verticals. These processors pack a punch with high-processing capacity and optimal throughput for the quickest time-to-solution. AMD EPYC CPUs have earned 200+ world records and counting by delivering best-in-class performance without compromise for applications such as Big Data, HPC, cloud, and AI. This instance is the best-performing x86 processor on the market,[1] providing up to 64 cores and up to 50% lower TCO[2] versus the competition. Third-generation EPYC processors offer up to 64 cores, enhanced memory performance, and significant changes throughout the Zen 3 microarchitecture yielding up to 19% IPC uplift[3] as compared to the previous generation. Organizations will now see twice the performance[4] versus the competition in HPC workloads. No matter which core count you choose, AMD provides the same features and functions at no additional cost.

With the unveiling of 3rd Gen AMD EPYC CPUs, organizations can expect unparalleled performance for all their supercomputing needs. The latest-generation processors are designed to elevate data center operations and bring transformational value rarely experienced before.

In conjunction with AMD CPUs, AMD Instinct MI100 GPUs are based on the new AMD CDNA architecture for the data center, which offers a new core design with double the computational efficiency of previous AMD data center GPUs.[5] It also includes Matrix Core Technology that targets acceleration for time-sensitive applications. The MI100 with AMD CDNA delivers over 1.2 terabytes per second of theoretical memory throughput[6] and up to a 37% speed-up of GPU-to-GPU theoretical bandwidth communication over the previous generation AMD Radeon Instinct compute accelerator,[7] so you can achieve extraordinary acceleration for your HPC and AI workloads.

Unleashing supercomputing capabilities doesn't happen just by consolidating more hardware together—the individual components must also be faster. AMD Instinct MI100 GPUs are the first GPU with over 10 teraflops of performance, providing organizations with up to 11.5 teraflops of peak 64-bit floating point performance.[8]

By combining award-winning AMD processors, organizations can build systems that can propel discoveries today and prepare them for supercomputing requirements tomorrow. The AMD EPYC and AMD Instinct products will be integrated with the high-speed AMD Infinity architecture to boost the communication between CPU and GPU. Organizations will have a valuable opportunity to operate with unmatched agility, precision, and efficiency.

[1] Results based on the SPECrate®2017_int_base metric as of April 2021, see spec.org.

[2] Estimates based on AMD Server Virtualization TCO (total cost of ownership) Estimator tool, comparing the AMD EPYC and Intel® Xeon® server solutions required to deliver 320 total virtual machines (VM), requiring 1 core and 8 GB of memory per VM, with a minimum total solution memory requirement of 2.56 TB of memory. The analysis includes both hardware and virtualization software components. For 320 VMs and 1 core per VM, the Intel® Xeon® Gold 6246R Processor requires 10 2P servers, at a total hardware acquisition cost of $117078. The AMD 1P_EPYC_7742 solution requires 5 1P servers, at a total hardware acquisition cost of $61729.2. The AMD solution requires 50% fewer servers. (1 minus [5 divided by 10]) times 100 = 50% fewer servers. Virtualization software pricing as of 09/29/2020. This scenario contains many assumptions and estimates and, while based on AMD internal research and best approximations, should be considered an example for information purposes only, and not used as a basis for decision-making over actual testing. TCO results are estimates and are provided for informational purposes only. All pricing is in USD. TLC-11

[3] MLN-003: Based on AMD internal testing as of 02/1/2021, average performance improvement at ISO-frequency on an AMD EPYC 72F3 (8C/8T, 3.7 GHz) compared to an AMD EPYC 7F32 (8C/8T, 3.7 GHz), per-core, single thread, using a select set of workloads including SPECrate2017_int_base, SPECrate2017_fp_base, and representative server workloads. Learn more at spec.org.

[4] MLN-041A: Based on SPECrate2017_fp_base on 04/14/2021, a server powered by two 64C AMD EPYC 7763 CPUs has a score of 651, spec.org/cpu2017/results/res2021q1/cpu2017-20210219-24944.html. The current highest score on 2P Intel Xeon "Cascade Lake Refresh" Processor server is a score of 309 with a 2P Intel Xeon Gold 6258R based server, spec.org/cpu2017/results/res2020q3/cpu2017-20200915-23979.pdf. See spec.org for more information.

[5] AMD Instinct MI100 accelerators provide 120 compute units and 7,680 stream cores in a 300W accelerator card. Radeon Instinct MI50 accelerators provide 60 compute units (CUs) and 3,840 stream cores in a 300W accelerator card MI100-09.

[6] Calculations by AMD Performance Labs as of October 5, 2020 for the AMD Instinct MI100 accelerator designed with AMD CDNA 7nm FinFET process technology at 1,200 MHz peak memory clock resulted in 1.2288 TFLOPS peak theoretical memory bandwidth performance.

[7] Calculations as of September 18, 2020. AMD Instinct MI100 accelerators support PCIe Gen4 providing up to 64 GB/s peak theoretical transport data bandwidth from CPU to GPU per card.

[8] Calculations conducted by AMD Performance Labs as of September 18, 2020 for the AMD Instinct MI100 (32 GB HBM2 PCIe card) accelerator.

Networking technologies

HPE Cray supercomputers with AMD processors represent a major advancement in the speed and flexibility of computing. A key component of these industry-defining systems is HPE Slingshot, the latest version of scalable interconnect. HPE Slingshot is the new interconnect that allows HPE Cray systems to bridge the worlds of supercomputing, cloud, and the data center. It is raising the bar for supercomputing networking solutions and preparing organizations to transition from traditional cloud-based computing to supercomputing technologies.

Created to be the backbone of business operations, HPE Slingshot offers a host of features to manage demanding computational and data management workloads. HPE Slingshot networks are based on 64 port switches operating up to 200 gigabits per second per port switch for maximum network throughput. Its high-radix 64-port switch coupled with the Dragonfly topology scales to over 250,000 endpoints, with a maximum of three switch-to-switch hops between any endpoints. In addition to latency improvements in larger installations, this low-diameter network reduces network equipment, cabling, and power and cooling costs. It also facilitates the use of innovative adaptive routing algorithms that improve application performance. This high-speed interconnect is also cost-optimized to enable straightforward and cloud-like performance in a supercomputing environment.

To enhance interoperability with storage architectures and data centers, HPE Slingshot is Ethernet compatible. This means the HPE Slingshot switches can connect directly to third-party Ethernet-based storage devices, as well as data center Ethernet networks. Applications running on HPE Cray supercomputer nodes can exchange IP/Ethernet traffic with the outside world, making it easier and more efficient to ingest data from external sources—an increasingly important consideration in this highly networked and data-driven world.

However, standard Ethernet can be poorly suited for HPC workloads, with high latency and overhead expenses. To overcome these issues, HPE developed an customized, high-performance, and low-latency Ethernet protocol that has smaller headers, support for smaller packet sizes, credit-based flow control, reliable hardware delivery, and a full suite of HPC synchronization primitives. HPE Slingshot not only uses this customized protocol for internal communication but can also intermix standard Ethernet traffic on all ports at packet-level granularity. This global approach allows HPE Cray supercomputers to comfortably straddle the supercomputing and data center worlds.

HPE has worked closely with AMD to couple HPE Slingshot with AMD EPYC and AMD Instinct processors. You can choose how to establish your network—in a standard rack for HPE Cray supercomputers or in a custom package for HPE Cray EX supercomputers. We deliver the same performance, connectivity, and cost savings for any scale.

Integrated storage

The amount of data stored and processed by compute systems has increased exponentially, especially in an age where traditional modeling and simulation have converged with methods like AI and machine learning. Often, these processes are running simultaneously on the same system in mission-critical and business-critical applications.

HPE Cray supercomputers integrated with Cray ClusterStor E1000 provide high-performance storage that is well-suited for the unique demands of supercomputing. This new breed of parallel storage is engineered to overcome data center challenges with these transformative capabilities:

- Rapidly feed data to CPUs and GPUs to power different workloads at scale
- Help eliminate input/output bottlenecks that slow time-to-innovation and force processors to sit idle waiting for data
- Reallocate your supercomputing budget from addressing costly storage issues to objectives such as upgrades and new deployments
- Unify your architecture with supercomputing technologies from HPE that are carefully integrated for simplicity, control, and faster problem resolution

Cray ClusterStor E1000 is a balanced, zero bottleneck design built on the AMD EPYC processor's PCIe Gen4 system architecture that supports modern networks like HPE Slingshot to rip through any roadblocks. To ensure high efficiency, the open-source parallel file system Lustre is embedded in the storage system. HPE provides customer services in-house for Lustre, which scales out near-linearly, without the need for software licensing per terabyte capacity or per storage drive. This allows organizations to reap the benefits of open-source solutions while getting high-level support.

This degree of data performance is not exclusive to large supercomputing sites. Organizations can start with as little as one-fifth of a data center rack and expand as needed. Now, you can gain unprecedented throughput delivered to your CPUs and GPUs and focus your budget on areas that accelerate time-to-insight.

All-inclusive software stack

The HPE Cray software stack has evolved from decades of supercomputing expertise and is now coupled with a container-based model to raise the bar for availability, reliability, resiliency, and interoperability. The software stack has a proven track record of helping organizations get results faster.

HPE offers an extensive software portfolio, so you can choose the right mix of solutions to fit your exact needs. This includes system development, system management, and data management tools to turn even the most complex hardware into your ideal solution. These offerings are designed to meet your compute requirements:

- **Scalability:** HPE has developed its own operating system—HPE Cray Operating System—that is tuned for scalability and built for scale, so organizations can adapt and grow their technology infrastructures with ease.
- **Performance:** Organizations need an architecture that can run different applications efficiently. HPE Cray software helps optimize your underlying architecture to execute workload execution for workloads with utmost performance.
- **Productivity:** Increased visibility and control into your supercomputing needs enable you to automate actions and effectively utilize resources to complete jobs as quickly as possible. The HPE Cray Programming Environment is a fully integrated software suite designed to improve productivity, application scalability, and performance and it is the only environment that can be used to program both CPUs and GPUs.
- **Cloud:** The HPE Cray software stack prepares your environment to perform like a supercomputer and run like a cloud. Cloud gives you the agility and elasticity to harness the right compute capabilities for your requirements, so your resources are available any place, anytime.

Now, organizations can speed up their results with supercomputing technologies and get into production quickly. You can install HPE Cray software and deploy your system over bare metal in minutes, rather than hours or days. During the setup, hardware elements are automatically detected and configured, and compute nodes are provisioned in parallel for top system performance in one session. For organizations looking to expand, new nodes can be added to your system without requiring the system to shut down, so you can grow your cluster without disrupting service.

Besides tools created by HPE Cray, we also offer leading third-party software along with popular open source software to provide 360-degree solutions that are long term. We integrate the end-to-end hardware and software features you need to help maximize the value of your system.

Professional, Advisory, and Operational Services

Your success depends on overcoming IT challenges and adapting to new technologies and means of operating. HPE Pointnext Services offers modern solutions to help you transform, all while training your workforce and identifying tools that allows you to excel in supercomputing.

HPE Pointnext Services enables you to design and implement your digital journey map by defining your goals and requirements, planning for potential roadblocks, and selecting the right technologies and services to deliver your ideal solutions. We work with you to underpin your vision to automate, help optimize, and scale your operations. Bringing your IT together under one team gives you one place to call for your supercomputing needs.

For HPE Cray EX deployments, we provide premium HPE Cray Advanced Support. This offering is designed to be one of the most advanced, comprehensive support experiences for supercomputing technologies. Unlike other industry offerings, Advanced Support is based on the performance-based service-level agreement under one price. With hardware and software services and the option of on-site technical support, we can resolve any system problems with HPE and eligible third-party solutions.

HPE has proven track record of helping customer increase the use of their solutions, so they can get results faster. Our 23,000 experts and 30+ ecosystem partners worldwide are available to help you accelerate innovation and reach your desired outcomes.

Partnering with HPE gives you the expertise to accelerate your data center transformation. It's time to free your resources from day-to-day heavy lifting, so you can stay ahead of what's next.

Cloud and financial services

HPE is delivering market-leading compute technologies as fully managed, pay-per-use services from HPE GreenLake. We aim to democratize the use of supercomputing by helping eliminate the concerns of system and operational costs and limited technology expertise through consumption-based options. Organizations can deploy these services in any data center, whether on-premises or in a colocation facility, to leverage data growth faster and accelerate time-to-insight.

HPE GreenLake, powered by AMD-based processors, is dramatically simplifying this experience by speeding up deployments as much as 75% and reducing capital expenditures by up to 40%. Our services bring the cloud experience to every environment, which enables you to free up capital, boost operational and financial flexibility, and refocus on running your projects in as little as 14 days.

HPE GreenLake services for HPC remove the complexity and cost associated with supercomputing deployments by delivering fully managed, pre-bundled services based on purpose-built systems, software, storage, and networking technologies that come in small, medium, or large options. Organizations can order these through a self-service portal to choose the right configuration for your workload needs and receive services to get started quickly. As part of the offering, organizations gain several features to easily manage and control your deployments. HPE GreenLake Central offers an advanced software platform for you to manage and help optimize your services. The self-service dashboard allows you to run and manage clusters on your own—without disrupting workloads—through a point-and-click function, while HPE consumption analytics provides at-a-glance analytics of usage and cost based on metering through HPE GreenLake.

Now any organization, from the major corporations to startups, can easily adopt supercomputing technologies. According to HPC and mission-critical solutions experts at HPE, we are transforming the market by delivering industry-leading solutions through HPE GreenLake. These cloud services enable any business to access powerful HPC and AI capabilities that will power their ability to advance critical research and achieve bold customer outcomes.

HPE is one of the few system vendors that delivers products and services spanning from edge to cloud. Very soon, HPE GreenLake will support supercomputers of any size for every data center, so you have the capacity and scalability you need to meet unexpected spikes in resource demand, all in line with your budget. Together with our experts, you can make technology that seemed unattainable the foundation for your future.

FUELING GROWTH AND INNOVATION IN MULTIPLE INDUSTRIES

Every industry can benefit from supercomputing technologies, and HPE is working to bring these capabilities into the data center. Through a wide variety of use cases, numerous organizations are realizing better outcomes, whether that means increased profits and productivity, delivering better services, or discovering a breakthrough. HPE supercomputing technologies have impacted operations in sectors including manufacturing, energy and gas, financial services, healthcare and life sciences, earth sciences and climate, and academia.

Manufacturing

The manufacturing sector has grown increasingly more competitive as organizations work to accelerate the development of products and services. Organizations are applying computer-aided engineering (CAE), computer-aided design (CAD), and simulation techniques to hone and diversify their offerings to remain competitive, while revealing operational efficiencies to increase their bottom lines. Powering high-fidelity models is a top priority for manufacturers to gain real-time insight. The benefits of CAE/CAD technology are limitless, but only with the right infrastructure in place. The ideal production environment is flexible and cost-effective, enabling organizations to keep pace with industry trends and customer expectations. To accomplish these goals, manufacturers are exploring the potential of supercomputing to advance modeling and simulation.

The use of HPC technologies like supercomputing is responsible for significant progress in the manufacturing sector, driving improvements across product design, testing, and development and meeting the rising demand for sustainable practices. HPE compute is extending the use of CAE/CAD capabilities, so today's manufacturers can produce high-quality products and services quickly and accurately.

HPE and AMD have teamed up with Ansys, a pioneer in simulation software, to help manufacturers uncover engineering insight with the power of supercomputing. These solutions enable organizations to implement technologies that can support the extreme scaling of demanding CAE applications to facilitate high-resolution models and more simulations. These capabilities provide immediate and predictive insight into how products like autonomous vehicles, aircrafts, and medical devices behave and perform under real-world conditions. Organizations that use HPE supercomputing technologies have the competitive edge to deliver safer, more effective products at low TCO.

Energy and gas

Advancements in supercomputing are accelerating energy and gas exploration, allowing organizations to execute compute-intensive workloads at high speeds. These insights are crucial to pinpoint the location of energy reservoirs that exist underground or beneath the seafloor, as well as determine the ideal materials and methods to extract these resources for transportation and refinement. Traditional compute environments lack the performance to rapidly capture data and convert it into actionable intelligence, which can result in unplanned downtime, catastrophic failure, and lost revenue. Energy and gas organizations need to modernize their IT to improve their operations.

A leading producer of the energy and chemicals has built its legacy by investing in cutting-edge supercomputing technologies to enhance their upstream capabilities and support a global downstream network. The business has made supercomputing innovation a top priority to solve issues related to exploration and discovery, such as worker safety and resource optimization, all the way to production. Their hope was to develop a new system that would allow them to dramatically reduce risks and promote efficiency with the use of sophisticated imaging and deep learning algorithms.

To reach this goal, the business unveiled the latest and most powerful addition to its supercomputing environment. The new system presents new opportunities in both exploration and development, with the ability to make more informed decisions on operations in the field and investments. The supercomputing system was created by HPE to deliver 55.4 petaflops of peak computing power, making it a top ten-ranked supercomputer

on the TOP500®. This groundbreaking capability can process and image some of world's largest geophysical datasets to generate detailed 3D earth models. As a result, the business can run simulations with dedicated seismic analytics designed to assist in the discovery of oil reservoirs, while reducing overall risks and costs by identifying when and where to drill.

These capabilities are the next step in the energy giant's transformation, complementing a suite of advanced technologies that are reshaping core operations and reinforcing their leadership in geoscience. The system created by HPE will be used to guide future investments in innovative technology for production and resource allocation, helping the business expand their influence in the energy and gas sector.

Financial services

Financial services institutions (FSI) make millisecond-decisions that affect every area of operations—from identifying and resolving cybersecurity threats to developing better financial products and accelerating intelligence to boost profit margins.

Utilizing data from sources such as transactions, payments, and operational workflows, financial institutions can predict emergent trends, determine favorable stock trades, improve risk management, and enable fraud detection for real-time events. However, organizations operating on outdated technologies often struggle to efficiently capture and derive intelligence from the available data. To compete in the fast-paced FSI market, decisions must be made instantly and with precision. This presents a tremendous hurdle for organizations operating on legacy infrastructure that struggle to execute data-heavy workloads at scale.

Building resiliency, adapting to changing regulatory requirements, and developing more effective ways to engage customers will be key competencies moving forward. With the use of supercomputing, financial institutions can overcome the congestion and latency issues caused by traditional compute. HPE has already helped organizations evolve and reach new heights of financial performance with solutions that are designed to streamline core operations and fuel compute across any environment. We have delivered cutting-edge technologies and tools for FSI for over 35 years. And the HPE supercomputing technologies powered by AMD processors are the next stage of this important legacy, making supercomputing more accessible and affordable for organizations of all sizes.

Healthcare and life sciences

Data growth is impacting healthcare and life sciences in profound ways. Many organizations in these sectors are unable to manage the flood of data being generated by an increasing number of sources. The problem is compounded for those that lack the capacity to channel this information into deep insights. Compute technologies that have been used for the past decade are no longer powerful enough to analyze this crucial data, causing scientists, researchers, and caregivers to fall behind while organizations struggle to work and innovate.

To enter the next phase of medical and scientific advancement, healthcare and life sciences are transitioning to a new type of infrastructure to accelerate vital workloads and scale as data volumes continue to rise. Supercomputing is aiding discovery across these industries and revolutionizing what we know about ourselves and the world around us. Adopting an HPE Cray supercomputing system at the right scale with the right mix of AMD CPUs and GPUs allows organizations to select the right amount of compute for their specific use cases. Four primary applications include personalized genomics, biomedical imaging, and computational chemistry.

Personalized genomics is the basis of disease research, drug discovery, and expanding the scope of precision medicine. Scientists run complex simulations to understand biological and chemical interactions that can lead to new drug therapies and curing diseases. Finding markers in DNA and correlating patterns requires organizations to sequence thousands of genomes. To support their research, laboratories, pharmaceutical groups, and medical departments around the world are using supercomputing to process highly complex, granular, and unstructured data. These technologies provide the essential power and flexibility required for high-speed analytics.

Supercomputing technologies are used to enhance biomedical imaging. Visualizing complex phenomena in the human body is one of the greatest challenges for caregivers and healthcare researchers. Vital applications of biomedical imaging include cardiac imaging, functional brain imaging, and interventional imaging for guided surgery. Creating a high-resolution image in the 3-dimensional space of the human body is an incredibly precise task. Supercomputing technologies are capable of delivering the speed and accuracy to power these images in real-time.

Computational chemistry is used at every level, in universities to labs of all sizes. Computational chemists and researchers develop sophisticated programs to model and study the structures, functions, and interactions of biologically important molecules. To do this effectively, they rely on supercomputing performance to apply AI and machine learning techniques at scale. Supercomputers can derive insights from vast amounts of data to answer pressing questions in biochemistry. These capabilities enable scientists to accelerate research, make predictions, and solve great medical mysteries.

Earth sciences and climate

Supercomputing technologies from HPE are redefining how organizations approach weather and climate studies. HPE Cray systems are being used to power some of the most critical projects across the industry, helping scientists and researchers achieve better fidelity and faster predictability by accelerating key workloads.

CSC, the Finnish IT Center for Science, undertakes massive research projects across numerous scientific disciplines. Climate change is a top priority as rising sea levels and shifting weather patterns raise crucial questions about the safety and sustainability of our world. To take on these initiatives, scientists require top-notch computing resources to execute data-heavy and compute-intensive studies.

CSC enlisted the help of HPE to build a leading-edge system to fast-track mission-critical simulations. LUMI, based on the HPE Cray EX, is one of the world's fastest supercomputers with a peak performance of 552 petaflops per second. With its vast computational capacity, LUMI will execute workloads faster than previous systems, combining AI, deep learning, and sophisticated simulations with data analytics to fuel weather and climate research. The number-crunching capability of LUMI is accelerated by a partition of AMD Instinct GPUs and a planned future partition of 64-core AMD EPYC CPUs, making it one of the most advanced platforms for AI.

The U.S. Department of Energy's Oak Ridge National Laboratory (ORNL) is also home to an HPE Cray EX supercomputing deployment. Comprising two HPE Cray EX supercomputers, the system delivers a combined 7.2 petaflops of numerical weather prediction capability to support the U.S. Air Force and Army missions worldwide. The system features 800 nodes, each equipped with two AMD EPYC processors, with future plans to enable heterogeneous compute. Each four-cabinet system can be expanded to 1,024 nodes, accommodating AMD Instinct GPU-based nodes that would boost throughput by 10x.

With extraordinary performance, the ORNL is breaking new ground in weather forecasting for a number of targeted scenarios—including forecasting stream flow, flooding, land submersion, and water depths. The end-to-end supercomputing technologies made possible by the HPE Cray EX can enable greater speed and dedicated performance to advance simulations in weather forecasting that were not possible before. The ORNL plans to accomplish this by implementing comprehensive cloud physics and creating a global hydrology model that involves simulating hundreds of watershed and drainage basins to increase the accuracy in predicting future events.

In another case study, the National Oceanic and Atmospheric Administration (NOAA) wanted to implement a significant upgrade to the computing capacity, storage space, and interconnect speed of its Weather and Climate Operational Supercomputing System. The upgrade would keep the agency's supercomputing capacity on par with other leading weather forecast centers around the world.

About AMD

For 50 years AMD has driven innovation in high-performance computing, graphics, and visualization technologies—the building blocks for gaming, immersive platforms, and the data center. Hundreds of millions of consumers, leading Fortune 500 businesses and cutting-edge scientific research facilities around the world rely on AMD technology daily to improve how they live, work, and play. AMD employees around the world are focused on building great products that push the boundaries of what is possible.

NOAA selected two new HPE Cray supercomputers, an operational primary as well as a backup, to dramatically increase their performance. The new systems effectively triple their previous capacity and double the storage and interconnect speeds, allowing the agency to develop better forecasting models with higher resolution, as well as detailed earth-system models using larger ensembles, advanced physics, and improved data assimilation.

Academia

Supercomputing technologies help universities gain a competitive edge in their technical, engineering, and scientific applications. Academic use cases range from simulation tasks to data-intensive research, a diversity of workloads, which require universities to expand their processing capabilities or miss out on the next wave of discoveries.

The Indiana University (IU) is a well-known institute that relies on the latest technologies and tools to facilitate their research. IU utilizes highly dense and parallel computing to expedite their activities and open up new performance capabilities that were once considered impossible. To achieve this, the university set out to build an AI-enabled system that would give researchers the power to explore new frontiers in medicine, cybersecurity, climate change, and more.

IU partnered with HPE to deliver Big Red 200, a high-density, water-cooled cluster system that contains nearly 100,000 cores and leverages the HPE Slingshot technology to deliver unmatched performance and energy efficiency. The system is the first of the HPE Cray EX supercomputers to be deployed and is specifically engineered for IU's state-of-the-art data center. Big Red 200 features 672 compute nodes, each equipped with 256 gigabytes of memory and two 64-core AMD EPYC processors. IU is making continuous upgrades to the data center, which include warm water cooling and the addition of 64 AMD GPU-accelerated nodes, each with 512 gigabytes of memory. After this expansion, Big Red 200 is expected to operate at a peak rate of 6 petaflops, more than 6x faster than its predecessor.

Moving forward, IU has continued plans to expand their supercomputing capabilities with HPE to support more scientists, meet their research objectives, and foster economic development.

CONCLUSION

HPE and AMD are helping organizations use supercomputing technologies at any scale, with a revolutionary approach that makes it possible to overcome the barriers to transformation. Rethought and re-engineered, we have introduced a new form of compute to address today's diversifying needs, with the same groundbreaking performance of the largest supercomputing systems. By bringing together the latest technologies and new workloads, HPE and AMD provide a compelling vision based on unrivaled performance, increased ROI, and low TCO.

HPE offers a truly comprehensive portfolio of supercomputing technologies, tools, and support services to organizations across the board, regardless of the size, scope, specifications, or budget. Powered by industry-leading AMD CPUs and GPUs, HPE Cray supercomputers are pioneering the next era of science, discovery, and achievement. We give you the control to develop your ideal solution, and as your workloads rapidly evolve, you can adapt your architecture. By deploying or upgrading individual components or implementing a full-scale supercomputing environment, you can meet any compute requirements that may emerge.

Supercomputing technologies are key to discovering unparalleled insights and ensuring your success, now and far into the future. It's time to start planning today. Let HPE empower your journey with the next advancements in this powerful supercomputing legacy.

LEARN MORE AT

hpe.com/us/en/compute/hpc/supercomputing/cray-exascale-supercomputer.html

amd.com/en/campaigns/high-performance-computing

Make the right purchase decision.
Contact our presales specialists.

Chat **Email** **Call**

Get updates

Hewlett Packard Enterprise

a50003570ENW, May 2021