HPE's End-to-End Approach for Delivering Superior HPC Performance and Sustainability

Sponsored by HPE
Bob Sorensen
April 2021

EXECUTIVE SUMMARY

HPE is committed to the development of an end-to-end high performance computing (HPC) ecosystem spanning processors, servers, and data center, all effectively integrated to deliver industry-leading use-case capabilities. Perhaps as important, however, is HPE's larger pledge to reducing the overall customer cost and associated environmental impact of their HPC server and datacenter operations. Moreover, HPE is committed to ensuring that critical environmental and social responsibilities are tightly woven into their business strategy, with the goal of differentiating HPE in the marketplace by helping customers achieve their technical, financial, and sustainability goals.

HPE is the world's largest supplier of HPC systems, including some of the most powerful computers in the world. These systems are critical in the development of leading-edge scientific and engineering projects underway in the commercial, academic and government spaces globally. HPE's HPC solutions address a wide range of workloads, effectively scale up or out, work on-premises or in the cloud, all while offering a comprehensive software portfolio for the most demanding traditional HPC workloads and emerging analytics and artificial intelligence (AI) applications. HPE's technology partnership with leading-edge semiconductor provider, AMD, extends the vision of both organizations' commitment to offer performant, secure platforms optimized for a wide range of configurations. Likewise, AMD shares HPE's vision for embracing a broader scope of issues key to fostering global sustainability and environmental responsibility.

Addressing growing global awareness that efficient energy use and sustainability can improve the way people live and work, HPE has wholly committed to developing transformative IT solutions to help users reap the benefits of growing edge-to-data center connectivity, while reducing the environmental footprint of their overall IT infrastructure. HPE sustainable products help customers energize existing sectors while addressing new ones. Additionally, HPE strives to expand critical IT access to people around the world, reflecting the technology industry's unique potential to deliver sustainability at scale. HPE has a number of programs that are already effectively addressing global concerns with climate and related energy issues not only for their products, but for their internal operations. HPE is serving as a strong example of what an end-to-end sustainability policy can achieve while consistently delivering leading-edge high-end compute and storage capability.

HPC OPERATIONS: AN EXPANDING COST, POWER, AND ENVIRONMENTAL FOOTPRINT

The hallmark of HPC has been one of continuously improved performance. Over the last two decades, the computational capability of the highest-end HPC systems have on average doubled every 15 months, leading to a more than a 90,000X performance improvement during that span. Without innovation, performance gains increase energy consumption, and such an increase in performance would not have been sustainable without the substantial improvement in energy efficiency technology that occurred during this same period. The unprecedented amount of performance that can now be obtained per watt, has enabled important breakthroughs across a broad range of science and engineering research areas, enabled continual and significant design and manufacturing productivity enhancements, and helped open new important use cases in the areas of modeling/simulation, big data analysis, and artificial intelligence (AI).

These inexorable gains, however, have come at a price. Each new generation of HPC systems is more complex, more expensive, and consumes more power than its predecessor. Whereas high-end HPC systems could once be purchased for $25 million or less, the newest levels of HPC, those in the exascale range, have price tags near $600 million apiece. The current fastest HPC system in the world, Fugaku at Japan's RIKEN, a national research facility, had a 2021 price tag that exceeded $1.1 billion.

Likewise, data centers, both serving both enterprise and HPC applications, are struggling to adequately address increasing computer rack-level power and cooling demands, driven primarily by the growing use of high performance processors and related co-processors or accelerators. Currently, some high performance processors draw power that can exceed 200W, and some graphics processing units (GPUs) are operating at over 300W. As a result, high-density rack configurations in the HPC space are moving from average power ratings of 20 kW to 40 kW, with estimates of reaching 70 kW and beyond in the next few years. Indeed, some of the highest-end custom liquid cooled HPC system racks consume 70 to 120 kW, with power densities tending to double every seven years.

The increasing levels of power needed to operate complete HPC systems are equally notable. Fifteen years ago, the fastest HPC system in the world required 1.4 megawatts (MW), while today HPC systems in development have power requirements that can exceed 25 MW. In addition, that power, in the form of heat, has to be managed and removed, requiring increasingly complex, efficient, and often costly counterpart cooling schemes.

The utility costs associated with operating these machines has quickly becomes an issue. An HPC industry rule of thumb is that one megawatt of power is roughly equivalent to a utility cost of $1 million per year, and those costs can be significantly higher based on local energy rates. As a result, power costs for a 25 MW HPC system could be over $125 million over its usable lifecycle, a significant fraction of the initial purchase price of the machine.

Costs aside, there are growing concerns about the impact of such power requirements on the environment, from the perspective of both the power needed to run the system as well as removing and then dispersing the resultant heat. Both HPC system suppliers and users are focusing on ways to eliminate waste and address inefficient approaches to HPC energy consumption without sacrificing performance.

HPE: A CORPORATE VISION FOR SUSTAINABLE TECHNOLOGY

HPE is taking a holistic approach to producing new generations of computing systems that are more intelligently designed and efficient than ever before but that still adhere to HPE's goal of delivering the highest performant systems currently available. In addition, HPE consistently works to meet those goals while operating for good in the communities in which they design, build, and deploy such systems.

HPE's commitment to sustainability is evident as every level in its overall hardware, software, and services solutions. HPE offers high performance nodes and servers that reflect a fresh approach to energy conscious technology using innovative liquid cooling techniques that can effectively lower energy needs and total cost of operations. From a software perspective, HPE offers customers the ability to optimize power usage of their HPC systems. Their cluster management software collects, analyzes, and optimizes power for all system hardware -- CPU, GPU, rack, chassis, nodes, rack AC, bulk DC -- and orchestrates overall cooling systems operations, while maximizing workload efficiency.

Other HPE sustainability exemplars include:

- HPEs GreenLake, an IT as a service offering that targets right-sizing end user IT configurations and retiring assets responsibly while funding future innovation.
- HPE's Pointnext Data Center Technology Services that comprises a complete lifecycle of consulting services from strategy and planning to the complete commission of a data center.
- HPE Financial Services that offers financial and asset lifecycle solutions to enable worldwide access through the purchase or rental of HPE hardware and software for current needs as well as to support a long-term strategy with options for pre-owned equipment acquisitions, retire and recover plans, and aggressive refresh cycles.

THE HPE AND AMD TECHNOLOGY PARTNERSHIP: PERFORMANCE AND EFFICIENCY--INSEPARABLE BY DESIGN

HPE's HPC solutions can address a wide variety of workloads with the ability to effectively scale up or out, work on-premises or in the cloud, offer a comprehensive software portfolio and options for density-optimized compute for the most demanding traditional HPC workloads as well as emerging big data and AI applications. HPE currently offers high performance solutions in a broad range of architectures and configurations well suited to almost any data center and its specific workload composition: its line of HPE Cray supercomputers and HPE SGI 8600 systems lead the way in world-class performance. Its HPE Apollo 2000 Gen10 Plus and Apollo 6500 Gen10 Plus are engineered with density and agility to support the most demanding traditional HPC and AI applications.

- HPE's HPC systems offer additional flexibility by extending on-premises capabilities with on-premises, off-premises, hybrid cloud options or through a consumption-based model, allowing users to design a system that best meets their performance, provisioning, and budgetary considerations.

HPE's commitment to delivering high compute capability with a dense and energy efficient profile is clearly evident in one of their recent offerings, the HPE Apollo 2000 Gen10 Plus. The system is specifically targeted to support rack-scale efficiency in configurations optimized with the right compute, flexible I/O, and storage options.

- Designed as an enterprise-level, density optimized, 2U shared infrastructure chassis for up to four HPE ProLiant XL225n Gen10 Plus compute nodes, the system can seamlessly be used as a building block for small, mid-sized, or large HPC clusters.
- A 42U rack fits up to 20 HPE Apollo 2000 system chassis, accommodating up to 80 servers per rack driving reductions in overall data center space, power, and cooling demands.

In order to deliver the highest performing and secure systems and solutions for their customers, HPE has established a strong technology partnership with AMD, one of the industry's leading innovators for CPUs and GPUs that provides record-breaking technologies for a wide range of HPC, cloud, and enterprise applications. For example, AMD recently announced the EPYC™ 7003 series processors that extended the total AMD EPYC family of processors world record count to over 200 on industry-standard benchmarks, including a number of energy efficiency records (see endnote 1). The design goal of "Zen3" was to increase performance without unduly increasing power and HPE delivered a SPECpower_ssj® 2008 world record on the HPE Apollo 2000 Gen10 Plus using four of the HPE ProLiant XL225n Gen10 Plus servers running 64-core AMD EPYC 7762 processors[2].

The EPYC™ 7003 Series processors deliver up to 64 "Zen 3" cores per socket for large-scale virtualization, HPC, dense computing, and leadership per core performance for single-threaded applications. In addition, the 128 PCIe® Gen4 lanes and up to 256MB of L3 cache (or up to 32MB per core) provide highly efficient and powerful I/O configurations, as well as highly predictable scaling of HPC application performance while reducing cache misses.

The EPYC family of 7001, 7002, and 7003 Series processors are built on a "chiplet" design that is both flexible and power efficient. The chiplet approach allows AMD to build a single processor package using several different Compute Complex (CCX) and Compute Die (CCD) combinations and connect them using a die-to-die interconnect to flexibly assemble I/O, memory, and processor cores. Using the cutting edge 7nm process node, this chiplet design offers efficient server CPUs without sacrificing performance and feature, all while helping significantly reduce TCO and operating expenses (OPEX) for HPC, enterprise customers, and cloud providers alike.

Sustainability-Enabling Software

In addition to their energy-efficient hardware, HPE also offers energy-efficient options designed to enable the highest level of performance while supporting the most rigorous energy conservation opportunities.

HPCM Power and Cooling Management Features

The HPE Performance Cluster Manager (HPCM) offers an end-to-end power management for hardware including CPUs, GPUs, racks, chassis, and nodes. HPCM can also set limits to trigger a power cap based on environmental failures (power or thermal), data center power capacity or varying workloads, planned brownouts or time of day guidelines. HPCM can also manage power consumption per user, per group and per job for all consumed energy. In conjunction, HPE has a comprehensive data collection and analysis scheme targeted to optimize the operation of the rack, server, and data center cooling environment, support more effective decision-making on cooling requirements and related equipment procurements, reduce overall systems downtime, and increase end-to-end reliability.

AIOps: Automating Performance, Reliability, and Efficient Operation for HPC

HPC system architectures are becoming increasingly complex, necessitated by the technical requirements of the growing diversity of heterogenous workloads that span traditional modeling and simulation, big data analytics, and AI-based applications. The ability to monitor, analyze, and optimize such operations can rapidly expand beyond the ability of any traditional performance monitoring schemes.

To help address this growing problem, HPE is using AI/machine learning (ML) to develop an advanced, non-threshold-based, real-time analytics tool designed to reduce data center downtime via rapid and early anomaly detection that performs at scale, speed, and automatically.

In addition, HPE is developing predictive capabilities to improve data center energy efficiency and sustainability with an initial focus on power usage effectiveness, predictive scheduling of cooling for large jobs, water usage effectiveness, and carbon usage effectiveness.

HPE'S ADVANCED COOLING TAGETED FOR BOTH SERVERS AND THE DATA CENTER

In today's world, many data centers are struggling to adequately address increasing computer rack-level power and cooling demands, driven by the growing use of high performance processors and related co-processors or accelerators as well as requirements for space-saving high-density computer capabilities. These necessities are driving new solutions for adequate power and cooling capabilities at both the rack/server level and for the data center at large. Alternate cooling strategies in use today can span from 5-10kW/rack raised floor hot and cold aisle configurations to 100kW/rack fully closed loop fanless liquid cooling schemes.

Seeking to ensure power and cooling capabilities that can operate effectively beyond the 50kW barrier where liquid cooling techniques are usually required, HPE offers a blend of capabilities that address cooling at both the IT equipment level as well as at the data center level. Building on the cooling concepts developed for the rack-level DLC used in the HPE Apollo 2000 Gen10 Plus, HPE has developed a companion system-level cooling scheme to address large data center cooling requirements as well.

To ensure that the best possible design choices are made for both current and planned data center configurations with an emphasis on energy concerns, HPE offers its Pointnext Data Center Technology Services (DCTS). DCTS comprises a complete lifecycle of consulting services from strategy and planning to the complete commission of a data center. Key services include road map planning, current state versus future state scenarios, IT sourcing options, and total cost of ownership analysis.

In a recent example that highlights the impact HPE data center design skills can achieve, HPE assisted a major electronics company in designing a state-of the-art HPC data center that minimized energy and water use and had a reduced carbon footprint, without compromising the performance required for their intensive research and development data-simulation needs. Compared with a traditional data center, this design could result in an annual site reduction of 25,000 metric tons of carbon emissions, 400 million gallons of water at the source, and $10 million in site energy costs.

HPE'S HOLISTIC SUSTAINABILITY PROJECTS AND PLANS—MORE THAN JUST A SLOGAN

HPE is dedicated to developing transformative solutions to help users reap the benefits of growing edge-to-data center connectivity, while reducing the environmental footprint of their IT infrastructure. Spearheading the effort is HPE's Living Progress program, HPE's companywide and multifaceted initiative to apply the innovation engine of HPE to create sustainable solutions that meet the technology demands of the future. Highlights include:

Climate-related Initiatives

HPE was the first major IT company to set science-based targets to reduce greenhouse gas emissions across the value chain, including its internal operations and related supply chain functions. HPE's climate goals align with the recommendations of the internationally recognized Paris Climate Agreement to limit global average temperature rise to well below 2°C from pre-industrial levels to substantially reduce the risks and effects of climate change.

- In 2016, HPE established the goal of reducing their manufacturing-related greenhouse gases by 15% by 2025. In 2018, HPE had already reduced emissions by 2% over their 2016 number and the organization is on track to meet their 2025 goal.
- In 2019, HPE set a new target to reduce absolute emissions from their transportation logistics by 35% by 2025, relative to 2016; this goal builds on successful reductions of those emissions of 29% compared with 2016 levels.

Energy Conservation

HPE has an aggressive multiprong approach to driving down energy requirements for the design and marketing of energy efficient IT products as well as in-house operations.

HPE technology is setting records for products with leading-edge performance while delivering an ever-lower carbon footprint. In 2019, efficient IT products and services represented nearly $7.7 billion in revenue at HPE. The approach combines energy efficiency that delivers optimum levels of power, storage and connectivity and resource efficiency. Companion efforts target data center designs that require the least amount of support equipment and staff for power conversion, cooling, and resiliency.

Complementing these customer-oriented efforts, HPE has a broad range of inward-facing activities designed to reduce its overall environmental footprint, spanning programs for product return, reuse and recycling, material and packaging schemes, energy and water consumption, ozone depletion and air pollutants, and both hazardous and nonhazardous waste generation.

- HPE deployed a range of actions to reduce on-site electricity consumption that resulted in an operational energy use reduction of almost one-third between 2016 and 2019.
- Between 2016 and 2019, HPE increased their energy use sourced from renewables from 20 percent to 41 percent of their overall operations energy consumption.

HPE/AMD: Realizing One Plus One Equals Three

HPE's partnership with AMD complements the vision of both organizations to deliver a performant, efficient, secure platform that can be optimized for a wide range of configurations in enterprise applications, virtualized and cloud computing environments, high-performance computing jobs, or data

analytic applications. Perhaps more important, AMD shares HPE's vision for embracing a broader scope of issues key to foster global sustainability and environmental responsibility.

AMD has a track record for responsibly developing and delivering cutting-edge technologies that help solve some of the world's most complex challenges and doing so with environmental stewardship in mind[3]. AMD was the first to semiconductor company to receive approved science-based climate targets by the Science-based Targets initiative with goals that spanned supply chain, operations, and products. One of the goals – the 25x20 Energy Efficiency Initiative – exceeded a 25x improvement in energy efficiency for processors powering mobile devices by achieving a 31.7x gain from 2014-2020[4].

- The gains achieved in 2020 by the 3rd Gen AMD Ryzen processors exceeded historical trends by 2x and resulted in 84% less typical use energy consumption and 80% less average compute time for a given task.
- In addition, AMD EPYC processors have a sustained history of performance-per-watt leadership represented by the twelve world records in HPC energy efficiency.

AMD is delivering generation over generation performance improvements in areas that are important to our customers such as Performance per Watt (PPW). With the 2nd Gen of EPYC processors, the 7002 Series, AMD delivered up to 1.5x the PPW of the 1st Gen EPYC. Now the 3rd generation, EPYC 7003 CPUs more than doubles the performance of the competition, allowing researchers to reduce the time needed to find the next scientific breakthrough by either running more simulations in a given period of time or using bigger data sets and more complex models[5,6].

The forthcoming Frontier and El Capitan exascale systems – expected to be among the world's most powerful HPC systems – will be powered by AMD CPUs and GPUs in partnership with HPE. These systems will further push the boundaries of breakthrough simulations and modeling to advance the understanding of the science of weather, sub-atomic structures, genomics, physics and more.

SUMMARY

HPE is committed to the development of a robust and rapidly expanding HPC system ecosystem spanning power efficient leading-edge processors to data center power and cooling technologies, all with the goal of delivering industry-leading performance capabilities. Perhaps as important is HPE's larger pledge to reducing the overall cost and environmental impact of their HPC systems and associated data center energy demands. Moreover, HPE is committed to ensuring that environmental, social, and governance responsibilities are tightly integrated into their server business strategy, with the goal of differentiating HPE in the marketplace by helping customers achieve their technical, financial, and sustainability goals. HPE is committed to partnering with leaders like AMD that share a sustainability vision to deliver the best technology-driven solutions for their customers.

As one of the world's leading HPC system suppliers, HPE has a long history in the development of performant technology solutions that offer flexibility in design, installation, and operation. HPE is well positioned to meet any future emerging challenges that are the hallmark of the HPC sector. The ever increasing base of users and use cases for such high-end systems can only benefit from HPE's mandate to offer flexibility as a fundamental product objective, offering users the opportunity and assistance to select the system that best matches their computational, organizational, and deployment needs.

HPE, however, recognizes that the HPC world must operate within a growing list of additional considerations imposed by global attention to issues such as energy efficiency, sustainability and changing environmental conditions. As such, HPE is taking a leadership position not only as a world-class supplier of HPC technology, but as one that seeks to address and provide real solutions for some of these most pressing societal issues.

ENDNOTES

1 - For a complete list of EPYC-22 world records see http://amd.com/worldrecords

2 - The AMD EPYC 7763 processor scored an overall CPU Blade world record result on SPECpower_ssj 2008 overall ssj_ops/watt as of 3/11/2021 as published at https://spec.org/power_ssj2008/results/res2021q1/power_ssj2008-20210223-01073.html. The score of 17696 overall ssj_ops/watt is the overall CPU, 4-node, Blade best in class compared to all other SPEC POWER® 2008 server results published on the SPEC® website.

3 – For more information on AMD's environmental stewardship, visit https://www.amd.com/en/corporate-responsibility/planet

4 – For more information on the 25x20 Energy Efficiency Initiative, visit www.amd.com/25x20

5 - Comparison based on best-performing systems published at www.spec.org as of 04/13/2020. 2x EPYC™ 7742 powered server has set the 2-socket (1U) world record on Linux® using the SPECpower_ssj® 2008 benchmark with a score of 20,214 overall ssj_ops/watt (1U on SUSE® Linux Enterprise Server 12 SP4). https://www.spec.org/power_ssj2008/results/res2020q2/power_ssj2008-20200313-01019.html. The next highest published competitive score is 13,465 overall ssj_ops/watt 1U on Red Hat® Enterprise Linux Server 7.6 (Maipo) on a 2-socket (1U) Xeon® 8280 powered server http://www.spec.org/power_ssj2008/results/res2019q2/power_ssj2008-20190313-00937.html as of 4/13/2020. SPEC®, SPEC POWER® and SPECpower_ssj® are registered trademarks of the Standard Performance Evaluation Corporation. For more information about SPEC POWER® 2008, see www.spec.org/power_ssj/

6 - Results as of 04/14/2021 using SPECrate®2017_int_base. The AMD EPYC 7763 scored 839, http://www.spec.org/cpu2017/results/res2021q1/cpu2017-20210219-24936.html which is higher than all other 2P scores published on the SPEC® website. SPEC®, SPECrate® and SPEC CPU® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information.

SOURCES AND ADDITIONAL INFORMATION

Publications and reference sources used in this document:

- The Top500 List (www.top500.org)
- *Cooling Strategies for the Modern Data Center*, A Technical Whitepaper, Hewlett Packard Enterprise, a00104981ENW, September 2020, Rev. 1
- *Apollo 2000 Power and Cooling Challenges and Opportunities for Differentiation*, TekTalk, Hewlett Packard Enterprise, May 2020
- *HPE Apollo 2000 Gen10 Plus: Are your computing needs exceeding you budget & space*? White Paper, Hewlett Packard Enterprise, a00109717enw, January 2021
- *HPE Performance Cluster Manager, Power and Cooling Capabilities*, Hewlett Packard Enterprise a0004744ENW, February 2021
- *Artificial Intelligence for IT Operations; a study on high-performance compute*, Technical white paper, Hewlett Packard Enterprise, a50003274ENW, December 2020

- *People, Plane, Purpose, 2020 Corporate Citizenship Summary*, Advanced Micro Devices, PID# 20408343-A
- *Improving the energy efficiency of modern supercomputers: Novel approaches for driving optimal performance per watt*, Technical White paper, Hewlett Packard Enterprise, a00028651enw, October 2017
- *Living Progress Report 2019*, Hewlett Packard Enterprise, a00097537enw, May 2020

About Hyperion Research, LLC

Hyperion Research provides data-driven research, analysis and recommendations for technologies, applications, and markets in high performance computing and emerging technology areas to help organizations worldwide make effective decisions and seize growth opportunities. Research includes market sizing and forecasting, share tracking, segmentation, technology, and related trend analysis, and both user and vendor analysis for multi-user technical server technology used for HPC and HPDA (high performance data analysis). Hyperion Research provides thought leadership and practical guidance for users, vendors, and other members of the HPC community by focusing on key market and technology trends across government, industry, commerce, and academia.

Headquarters

365 Summit Avenue
St. Paul, MN 55102
USA
612.812.5798
www.hpcuserforum.com and www.HyperionResearch.com
