Technical white paper

Extending the viability of air cooling in high-performance data centers by upgrading to HPE Cray XD2000 Systems

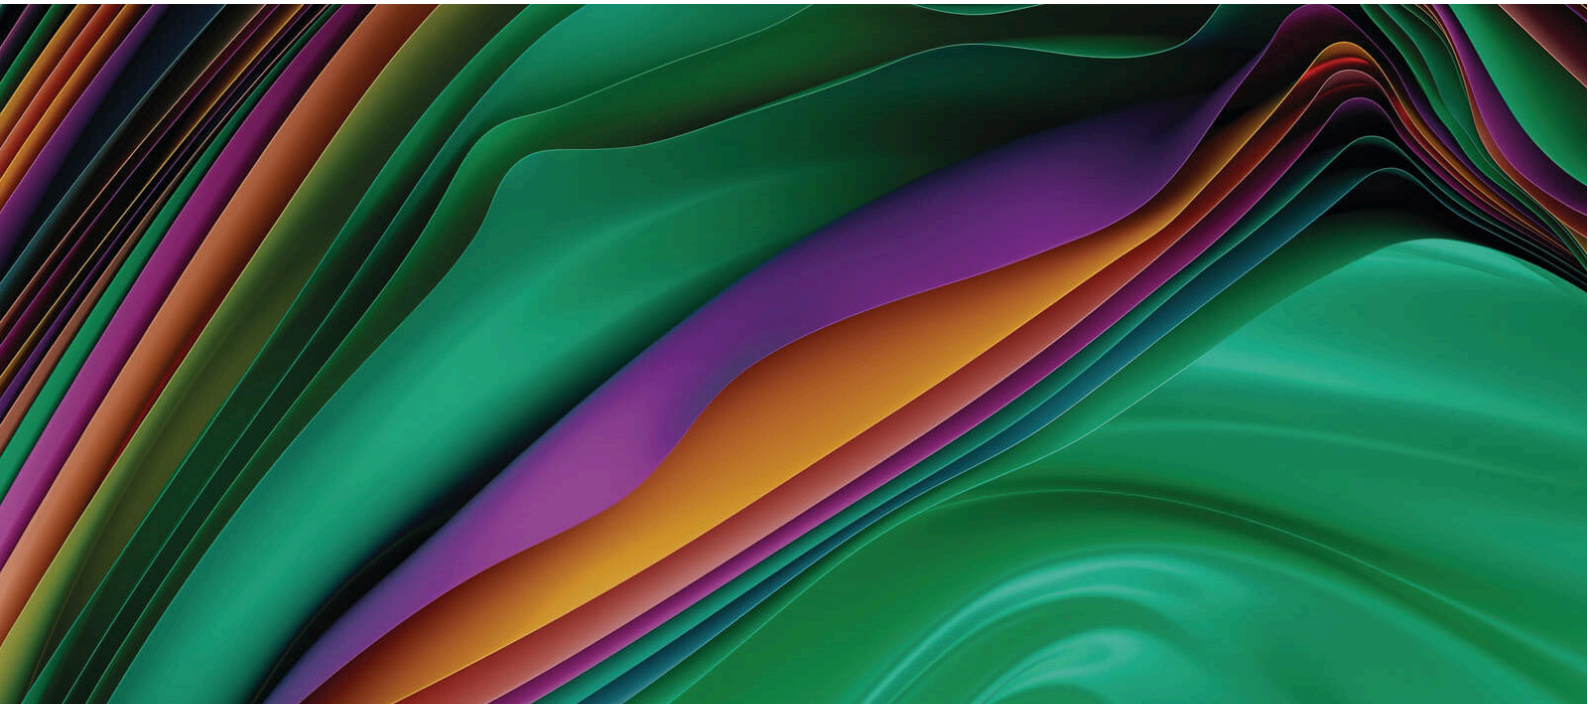

Contents

How new technologies from Intel® and HPE can help organizations maximize application performance while navigating the transition to liquid cooling.

Executive summary

For organizations operating high-performance data centers, cooling is a significant challenge. With the demand for compute ever-increasing, the performance required to deliver on these needs also needs to rise. Subsequently, high-end processor TDPs are also climbing with each generation, and we are rapidly approaching the point where cooling these processors with traditional air will become infeasible.[1] Liquid cooling is an obvious answer but brings new challenges, including cost, maintenance issues, and concerns about leakage and safety.

In this paper, we introduce the latest server solutions from Hewlett Packard Enterprise and Intel and explain how innovative new designs can help organizations deploy dense, high-performance servers while extending the viability of air cooling. With HPE Cray XD2000 Systems powered by 4th Gen Intel® Xeon® Scalable processors, organizations can:

- Realize dramatic improvements in application performance
- Help minimize disruptions to data center environments
- Avoid or delay the need to introduce liquid cooling
- Gradually adopt liquid cooling at their own pace

We also present a series of benchmarks conducted by HPE that explore the trade-offs between air and liquid cooling related to performance and power efficiency. With this information, data center managers can make informed decisions about how best to evolve their infrastructure based on their unique requirements.

Driven by data, compute requirements are on the rise

According to the latest IDC Global DataSphere forecast, the new data created, captured, replicated, and consumed yearly is expected to double between 2022 and 2026.[2] Fueled by this growth in data, new predictive and analytic techniques, and competitive imperatives such as machine learning and AI, compute requirements are growing at a similar pace.

Today, high-performance systems are critical in applications ranging from manufacturing to financial services, life sciences, and data analytics. Manufacturers rely on HPC for structural design, computational fluid dynamics (CFD), and machine condition monitoring. Life sciences firms require large amounts of computing power for genomic analysis, surveillance, computational chemistry, and image analysis. Organizations are continually looking for ways to deliver additional computing capacity to keep pace with growing demands and stay a step ahead of the competition.

Cooling — a looming data center challenge

Meeting the high levels of performance required by today's data and compute-intensive applications is driving increased power and cooling requirements. Power, density, cooling, and concerns about sustainability are issues for almost all data center operators. Per-socket TDPs for top-bin CPUs range from 270 watts to 350 watts or more, and next-gen CPUs are likely to be even more power-hungry, reaching 400 to 500 watts per socket.[3]

The challenges are even more significant for GPUs critical to AI and machine learning workloads, where next-gen TDPs are expected to reach 700 watts.

A key issue is that silicon designs in modern processors are increasingly going 3D with components layered on top of one another. This presents new thermal challenges and requires that case temperatures be cooled to ever-lower levels to avoid component overheating and damage.[4] These conflicting trends are illustrated in Figure 1.

[1] Thermal Design Power (TDP) is defined as the theoretical maximum amount of heat generated by a CPU or GPU, usually expressed in watts, that a computer's cooling system must be designed to dissipate.

[2] "Worldwide IDC Global DataSphere Forecast, 2022–2026: Enterprise Organizations Driving Most of the Data Growth," IDC, 2022.

[3] Based on HPE internal estimates, 2022.

[4] T_{case} refers to the temperature at the interface between a CPU package and its heatsink.

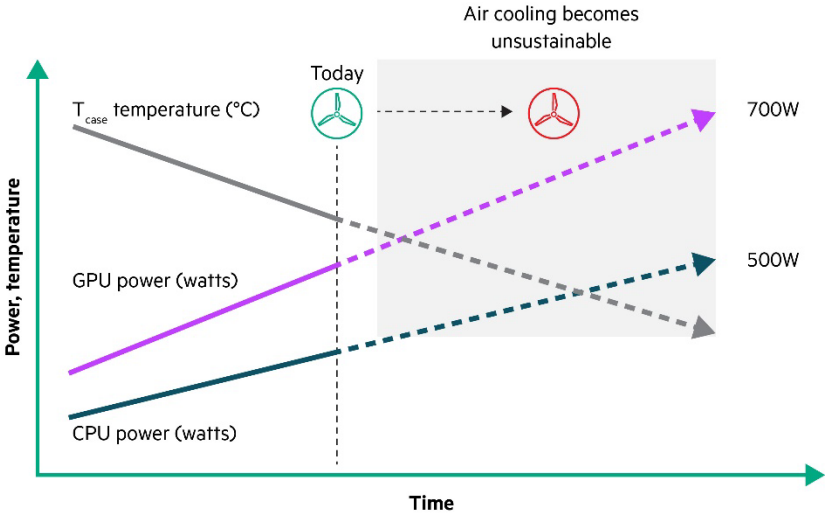

Figure 1. Air cooling is becoming unsustainable as component TDPs increase

Compounding this challenge, customers increasingly demand high-density racks that pack more computing power into a smaller data center footprint. Today, two-thirds of US data centers already have peak power demands of up to 16 to 20 kW per rack.[5] Per rack power consumption is rising quickly, with dense HPC racks already consuming 40–60 kW or more.[6]

Organizations will need to make trade-offs, either investing in new cooling technologies to accommodate next-generation processors and GPUs or settling for less powerful processors and more sparsely populated data center racks. It is no secret that liquid cooling provides better thermal transfer efficiencies than air cooling, so for many, liquid cooling is a logical path forward.

Liquid cooling

Liquid cooling spans a range of technologies, from rear-door chillers and heat exchangers to directly attached liquid cooling plates to immersion cooling. Liquid cooling can bring clear benefits:

Improved efficiency — In an analysis conducted by HPE, liquid cooling has been shown to reduce data center power usage effectiveness (PUE) and cooling-related power costs by up to 87.3%.[7]

Reduced environmental impact — Reducing power consumption with more efficient cooling can help organizations meet environmental, social, and governance (ESG) goals and reduce their data center's CO_2 equivalent (CO_2e) footprint.

Defer expensive data center upgrades — In space-constrained data centers, liquid cooling can enable denser rack configurations, helping maximize available space.

Improve reliability and predictability — Liquid cooling can prolong component life by providing stable operating temperatures, avoiding overheating conditions, and improving overall availability.

Despite these benefits, transitioning to liquid cooling is often easier said than done.

[5] How Power Density is Changing in Data Centers and What It Means for Liquid Cooling, JETCOOL Technologies Inc., March 2022.
[6] Based on HPE internal estimates, 2022.
[7] See the HPE white paper Addressing sustainability in the financial services industry. PUE refers to power usage efficiency, a measure of how much power is used to power data center servers vs. ancillary requirements such as lighting and air conditioning.

A challenging transition

Presently, air cooling is the predominant way of cooling high-performance servers. In a survey of 268 HPC sites from 252 organizations, Intersect360 Research found that 58% of respondents use air cooling exclusively.[8] 42% use liquid cooling in some systems, with the largest share using rear door chillers. Only 23% of commercial organizations operate fully plumbed racks with facility heat exchangers. For most, plumbing is extended to only a subset of their data center racks. In other words, there is still a long way to go before most facilities can fully embrace liquid cooling.

When deciding on a liquid cooling solution, customers must consider several factors: cost, sustainability, maintenance, and ease of installation.[9] Among commercial and industrial HPC users, the majority operate multiple clusters. According to the same Intersect360 Research study, 37% of organizations operate ten or more clusters ranging from entry-level HPC systems with 16 nodes or less to supercomputers that consist of more than 512 nodes. Upgrading these systems to liquid cooling poses technical, logistical, and financial challenges. These include:

- The added expense of operating two cooling systems instead of one
- A lack of standardization in cooling systems complicating adoption in multivendor environments
- Concerns about corrosion and safety hazards, such as risks of electrocution and arcing
- Increased operational complexity and the risk of cooling system failures

Organizations must consider multiple factors when introducing liquid cooling, including existing data center space, rack composition, power constraints, cooling capacity, utility costs, and projected growth requirements.

Extending the life of air cooling

Fortunately, new technologies from HPE and Intel provide data center managers with the flexibility to deploy the latest server hardware in air-cooled environments. Organizations can significantly extend the viability of air cooling by taking advantage of the latest HPE Cray XD2000 Systems powered by 4th Gen Intel Xeon Scalable processors. By deploying these Intel-based systems, organizations can:

- Help maximize performance while minimizing data center impact
- Avoid costly capital upgrades to data center facilities
- Protect existing investments in software and hardware

As illustrated in Figure 2, organizations can extend the life of air cooling without sacrificing performance, enabling them to gradually manage the transition to liquid cooling based on their schedule.

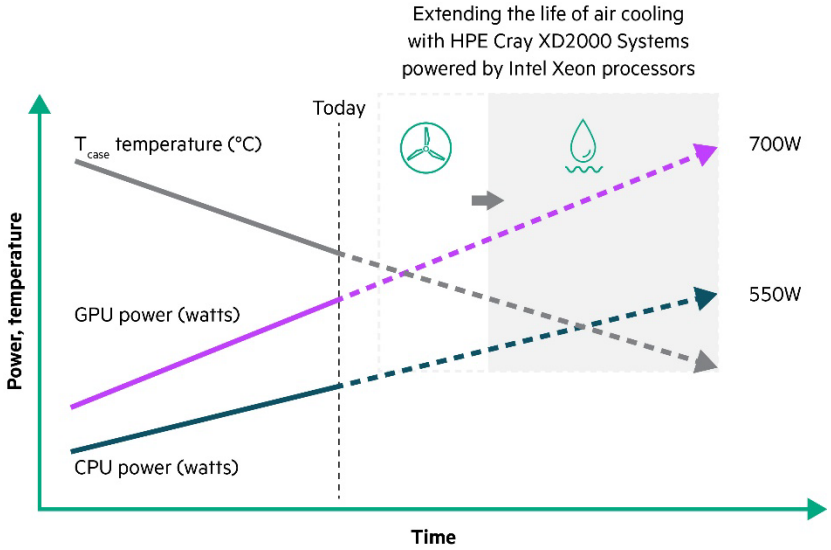

Figure 2. Extending the life of air cooling with HPE Cray XD2000 Systems

[8] "HPC Technology Survey 2021: Server Technologies and Configurations," Intersect360 Research, August 2021.
[9] 2022 Data Center Trends: Liquid Cooling Adoption Survey, JETCOOL Technologies Inc., 2022.

HPE Cray XD2000 Systems

With the HPE Cray family, HPE and Intel bring innovation from the world's most powerful supercomputers, making them available in commercial data center settings.[10] The HPE Cray XD2000 System is a dense, multiserver platform that packs exceptional performance and workload flexibility into a small data center space while delivering the efficiencies of a shared infrastructure.

Each HPE Cray XD2000 2U Chassis supports up to four HPE Cray XD220v 1U Servers powered by the latest 4th Gen Intel Xeon CPUs. Each server can be serviced without impacting the operation of other servers in the same chassis for maximum server availability. The HPE Cray XD2000 delivers up to 4 times the density of a traditional rackmount 2U server in standard racks and provides rear-aisle serviceability access.[11] Up to 20 HPE Cray XD2000 Chassis can be installed in either 42U or 48U HPE standard racks delivering up to 80 2P servers and 160 x 4th Gen Intel Xeon Scalable processors per data rack, subject to power and cooling considerations.

These systems offer a complete, scalable solution for customers requiring high-performance solutions. They feature flexible power and cooling options, including air cooling and direct liquid cooling (DLC), delivering superior performance while reducing TCO.

Figure 3. Density-optimized HPE Cray XD2000 Chassis supporting up to 4x HPE Cray XD220v 1U Servers

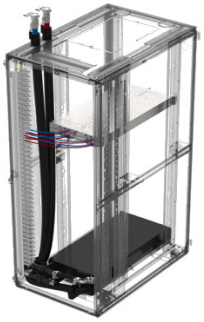

Figure 4. HPE data center rack with optional direct liquid cooling (DLC)

Table 1. HPE Cray XD2000 System with 4th Gen Intel Xeon and HPE Cray XD220v (1U) specifications

Component	Specifications
Processor	Up to two 4th Gen Intel Xeon Scalable CPUs — up to 56 cores / 112 threads per socket
OS support	Windows, RHEL, SLES, Ubuntu, HPE Cray OS, VMware®
Memory	8 DDR5 memory channels per CPU @4800 MT/s
Nodes	1U/2P half-wide servers; up to 4 per chassis
PCIe® lanes	65 PCIe Gen5 lanes with 4-link XGMI + 6 bonus PCIe Gen3 lanes per CPU
Storage	Supports up to 2x NVMe SSDs per node (4 for 2U) and/or M.2 2280/22110
Expansion slot	2x PCIe Gen5 x16 slot 1x PCIe Gen5 x16 MCIO slot 1x PCIe Gen5 x16 MCIO cable slot 1x PCIe Gen3 x4 M.2 Mezz
Power supply (hot plug)	CRPS PSU support 2400W 240V
Fans	4x 4056 fans per node
Cooling	Air-cooled or optional direct liquid cooling (DLC)

[10] Aurora Supercomputer and Argonne National Laboratory | HPE.
[11] HPE Cray XD2000 QuickSpecs.

Innovative engineering enables air cooling up and down the stack

Thanks to the design of the 1U HPE Cray XD220v Server residing in the HPE Cray XD2000 chassis; customers can benefit from the latest high-performance Intel Xeon Scalable processors in air-cooled environments. Customers can deploy fully populated HPE Cray XD2000 Racks using the latest processor technology without worrying about liquid cooling.

What makes this possible is the unique design of the Intel-powered HPE Cray XD220v Server illustrated in Figure 5. The HPE Cray XD220v is 20% wider than the previous generation HPE ProLiant DX220n designed for the HPE Apollo 2000 Chassis. However, this updated design still fits in industry-standard racks.

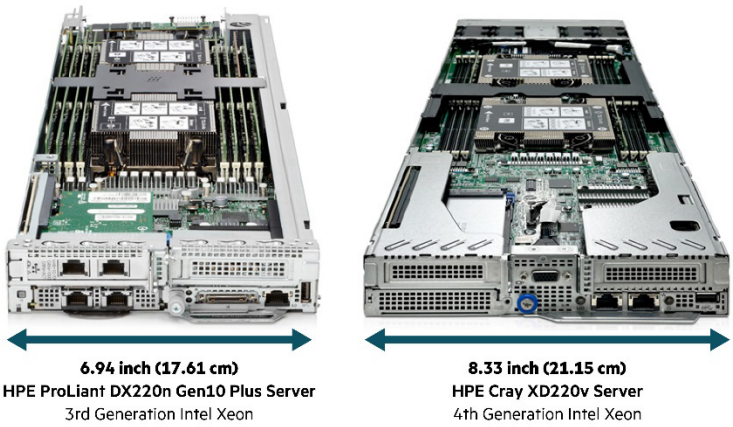

Figure 5. The HPE Cray XD220v supports efficient air cooling of the latest Intel Xeon processors

With larger heatsinks and additional cooling fans, this redesigned server supports the full range of 4th Gen Intel Xeon Scalable processors, from 12 to 56 cores per socket, to be efficiently air-cooled — including the most powerful 350 watt 56-core Intel® Xeon® Platinum 8480+ and Intel Xeon 9480 Max Series processors.

The redesigned system features a special baffling to optimize airflow and 16 fans (40 mm each) per HPE Cray XD2000 Chassis for reliable cooling of dense server configurations in the most demanding HPC environments. Better still, these servers are designed to support air cooling of future Intel Xeon processors. This translates into exceptional investment protection and flexibility. Organizations can deploy HPE Cray XD2000 Systems with air cooling today and easily add liquid cooling in the future.

This advantage is unique to the Intel-powered HPE Cray XD220v Server. Servers powered by competing processor technologies with similar TDPs require liquid cooling, adding cost and complexity to server deployments.

Direct liquid cooling (DLC)

While Intel-powered HPE Cray XD2000 Systems are at home in air-cooled environments, for customers with suitably equipped data centers, these systems also support plug-and-play support for DLC. Engineered and supported by HPE, DLC offers clear advantages compared to third-party or immersion-based cooling solutions. These include:

- Efficient thermal transfer for more efficient cooling
- No need for expensive hazardous chemicals or specialized fluids
- Server equipment remains easily accessible for serviceability
- HPE server racks connect directly to facility water supplies without secondary plumbing

By taking advantage of DLC, organizations can substantially reduce data center PUEs and cooling-related costs and improve energy efficiency. Options are available for CPU only or CPU plus memory cooling.

Advanced power and thermal management

HPE offers a complete software portfolio for customers deploying high-performance Intel-based clustered systems. This portfolio is validated, integrated, and performance-enhanced by HPE so that customers can select the right software for their mix of applications — all from one source. HPE offers operating systems, cluster managers, job schedulers, remote visualization solutions, HPC tools and libraries, and more.

HPE Performance Cluster Manager (HPCM) is a fully integrated system management solution offering all the functionality required to manage Intel-based HPE Cray XD2000 Systems. HPCM aggregates system management and remote management information from the standard DMTF Redfish® API supported on the HPE Cray XD2000.[12] The software also provides system setup, hardware monitoring and management, image management and updates, power management, and integrations with ISV and open-source software solutions.

HPCM helps optimize power and cooling in air- and liquid-cooled environments, giving administrators visibility to power faults, various overheating conditions, fan failures, and alerts about insufficient power or cooling. Management of liquid-cooled systems is further simplified with flow and pressure alarms, leak sensors, and flood tray alarms for trouble-free operation.

4th Gen Intel Xeon Scalable processors

The latest generation of Intel Xeon processors brings scalable, balanced architectures that integrate CPU and GPU functionality with the Intel oneAPI open software ecosystem, delivering a leap in data center performance, efficiency, and security. These processors deliver industry-leading performance and power efficiency for a wide range of applications, from CFD to quantitative analysis calculations to analytic workloads.[13]

The latest Intel Xeon Platinum processors offer substantial advantages over the previous generation, including:

- Up to 56 cores and 112 threads per socket[14]
- 105 MB of L3 cache per CPU with dynamic sharing with Intel Smart Cache technology
- PCIe 5.0 support, offering twice the I/O throughput of the previous generation[15]
- Compute Express Link (CXL) technology
- 8 x DDR5 4800 MT/s memory channels
- Optional high bandwidth memory (HBM) support in the Intel Xeon Max Series processors

A key advantage of 4th Gen Intel Xeon Scalable processors is support for built-in accelerators. These set of instruction extensions offer dramatic performance and efficiency advantages for a wide range of applications. Among the built-in accelerators are:

- Intel Advanced Matrix Extensions (AMX) for deep learning and inference
- Intel QuickAssist Technology (QAT) for offloading encryption, decryption, and compression
- Intel Data Streaming Accelerator (DSA) for accelerating data-intensive workloads
- Intel Advanced Vector Extensions 512 (AVX-512) for providing vector instruction support
- Intel Data Direct I/O Technology (DDIO) for facilitating direct communication between controls, adapters, and the host CPU memory and cache

With these accelerators, application developers and ISVs can boost application performance by leveraging various Intel oneAPI toolkits to take advantage of on-chip hardware features.[16] They can also substantially improve power efficiency. For example, in a real-time object detection AI workload, using Intel AMX vs. FP32 boosted performance per watt by over 14x.[17] The Intel oneAPI enables developers to write once and deploy on any device from CPU to GPUs to DPUs to FPGAs.

[12] Redfish is an industry standard protocol specified by the Distributed Management Task Force (DMTF) that provides a standard RESTful interface for managing servers, storage, networking, and converged infrastructure.
[13] See published Intel benchmarks showing Intel Xeon Scalable processor leadership at Supercomputing 22 | Performance Index.
[14] These and other specs presented here are based on the Intel Xeon Platinum 8480+ processor. Details of specifications are available at Intel Xeon Platinum 8480+ Processor 105M Cache 2.00 GHz Product Specifications.
[15] PCIe Gen5 employs a 32-bit bus clocked at 32 GHz for an effective bandwidth of 128 GB/s. PCIe Gen4 by contrast is clocked at 16 GHz for an effective bandwidth of 64 GB/s.
[16] See details on the various Intel oneAPI Toolkits.
[17] Accelerating HPC and AI with Intel, slide 23.

Each toolkit provides components such as compilers, optimized libraries, and trace/debugging tools for different use cases. Use cases range from general compute to high-performance computing to end-to-end AI and machine learning acceleration.

For data center managers, this is an important innovation. By taking advantage of these onboard accelerators, applications that would have previously required expensive power-hungry GPUs can run directly on the latest Intel Xeon Scalable processors. Also, as ISVs take advantage of these libraries and accelerators, application performance will improve as new software application versions are released.

Superior application throughput

Based on the industry standard SPEC CPU® 2017 benchmark, the latest Intel Xeon 8480+ CPUs deliver an impressive 100% throughput increase for the SPECrate® 2017_fp_base metric compared to 3rd Gen Intel Xeon 8380 processors.[18] Based on the SPECrate® 2017_int_base metric, the same comparison yields a 66% throughput increase.[19] Both of these results are illustrated in Figure 6 and were obtained on different generations of HPE ProLiant DL360 Servers.[20] This throughput increase provides a solid incentive to upgrade to servers running the latest Intel Xeon processors. With 4th Gen Intel Xeon-based servers, the same throughput can be achieved with half the data center footprint compared to 3rd Gen processors.[21] By doubling throughput per server, organizations have the potential to reduce data center space, the number of racks required, network drops, and associated management costs.

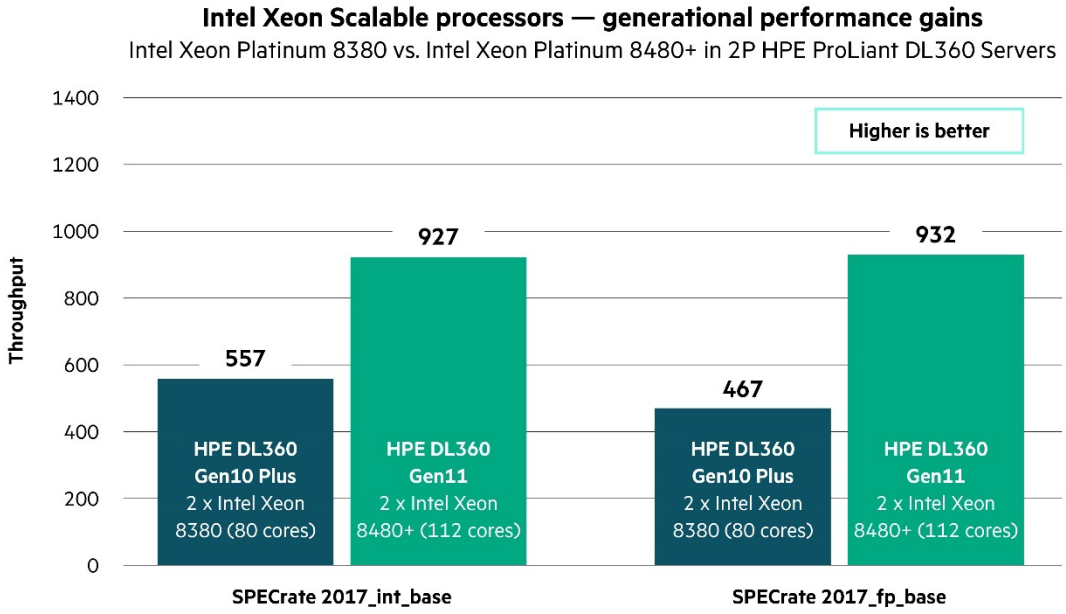

Figure 6. Intel Xeon Scalable processors — generational performance gains based on SPEC CPU 2017 benchmark results

Internal testing at HPE has shown that similar generational improvements can be obtained with HPE Cray XD2000 Systems powered by 4th Gen Intel Xeon 8480+ processors.

[18] The 2P Intel Xeon Platinum 8480+ (112C) scored 932 SPECrate 2017_fp_base: spec.org/cpu2017/results/res2023q1/cpu2017-20221204-32903.html. The 2P Intel Xeon Platinum 8380 (80C) scored 467 SPECrate 2017_fp_base: spec.org/cpu2017/results/res2021q2/cpu2017-20210524-26430.html. 932/467 represents an ~100% throughput gain.

[19] The 2P Intel Xeon Platinum 8480+ (112C) scored 927 SPECrate 2017_int_base: spec.org/cpu2017/results/res2023q1/cpu2017-20221204-32905.html. By comparison, the 2P Intel Xeon Platinum 8380 (80C) scored 557 SPECrate 2017_int_base: spec.org/cpu2017/results/res2021q2/cpu2017-20210524-26427.html. 927/557 represents a 66.4% throughput gain.

[20] The Intel Xeon 8380 performance results were obtained on a 2P HPE DL360 Gen10 Plus Server. The Intel Xeon 8480 results were obtained on a 2P HPE ProLiant DL360 Gen11 Server.

[21] Based on the SPECrate 2017_fp_base benchmark results presented here. Since each HPE ProLiant DL360 Gen11 Server running Intel Xeon 8480+ processors delivers twice the throughput compared to the previous generation, only half as many servers are required to achieve the same aggregate throughput.

Measuring the impact of air vs. liquid cooling

In March 2023, HPE ran a series of six internal benchmarks to evaluate the latest Intel Xeon processors' performance and power efficiency in air- and liquid-cooled environments. The tests involved an HPE Cray XD2000 Chassis fully populated with 4 x HPE Cray XD220v compute nodes, each with two Intel Xeon 8480+ processors. The benchmarks included SPEC CPU 2017 (SPECrate 2017_int_base and SPECrate 2017_fp_base), three separate SPEChpc™ 2021 benchmarks, and a High-Performance Linpack (HPL) benchmark.

For each benchmark, results were obtained in both air- and liquid-cooled configurations. Details were tabulated, including performance, power consumed by the HPE Cray XD2000 Chassis while the benchmarks ran, and performance per kW. The results of these tests are summarized in Table 2.

As shown, the latest top-bin Xeon processors in air-cooled configurations delivered performance on par with that achievable with liquid cooling. The variance between the air and liquid-cooled result for all six benchmarks was less than ~3%.[22]

Liquid-cooled configurations delivered slightly better performance because the higher temperatures in the air-cooled configurations led to higher leakage current in silicon. This resulted in a higher power draw, leaving less power available for boosting clock frequencies within the processor's fixed TDP budget.[23]

Table 2. Comparing performance and power requirements in air- and liquid-cooled HPE Cray XD2000 Systems

	Performance			Chassis power (watts)			Performance per kW		
Benchmark	**Air cooled**	**DLC**	**Delta**	**Air cooled**	**DLC**	**Delta**	**Air cooled**	**DLC**	**Delta**
SPECrate 2017_int_base (est.)* (per HPE Cray XD220v Server)	862	887	2.9%	4397	3768	–14.3%	196.04	235.40	20.1%
SPECrate 2017_fp_base (est.)* (per HPE Cray XD220v Server)	903	913	1.1%	4580	3940	–14.0%	197.16	231.73	17.5%
SPEChpc 2021, small, MPI only, 448 ranks (est.)* (per HPE Cray XD2000 Chassis)	2.98	3.03	1.7%	4685	3973	–15.2%	0.64	0.76	19.9%
SPEChpc 2021, tiny, MPI + OpenMP, 64 ranks, 14 threads (est.)* (per HPE Cray XD2000 Chassis)	30.00	30.80	2. 7%	4539	3862	–14.9%	6.61	7.98	20.7%
SPEChpc 2021, small, MPI + OpenMP, 64 ranks, 14 threads (est.)* (per HPE Cray XD2000 Chassis)	3.52	3.55	0.9%	4701	3982	–15.3%	0.75	0.89	19.1%
Linpack — Socket HPL Rmax (GF/s) (per socket)	3121	3167	1.5%	4565	3923	–14.1%	683.68	807.29	18.1%
Averages			**1.8%**			**–14.6%**			**19.2%**

* Results shown in Table 2 are estimates based on internal performance testing conducted by HPE in April 2023. SPECrate 2017_int_base and SPECrate 2017_fp_base estimated results are per server. SPEChpc 2021 estimated results are per chassis and were run across four servers. Linpack results are per socket.

The average impact of air vs. liquid cooling across all six benchmarks in Table 2 is shown in Figure 7. On average, the liquid-cooled configurations delivered 1.8% better performance and consumed 14.6% less power. The liquid-cooled HPE Cray XD2000 delivered a 19.2% boost in power efficiency measured in terms of throughput per kW.

[22] On average, the liquid-cooled configurations ran 1.8% faster across the six benchmarks. The largest impact was seen with the estimated SPECrate 2017_int_base result, where the liquid-cooled configuration ran 2.9% faster.

[23] Intel Turbo boost algorithm takes operating temperature into consideration, explaining the slight difference in performance.

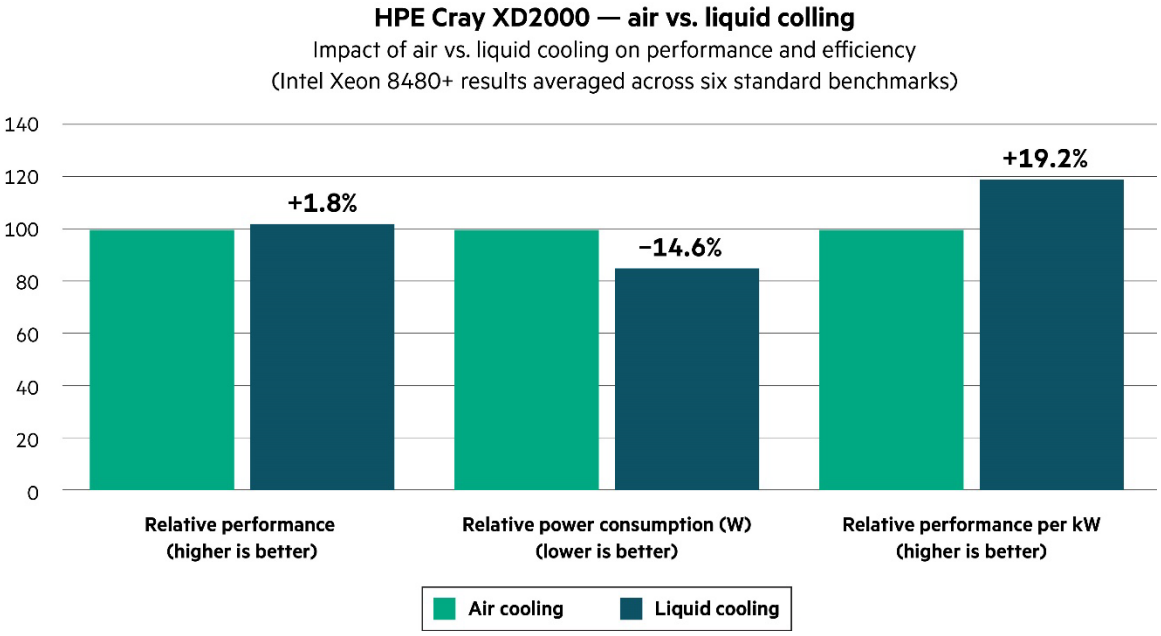

Figure 7. HPE Cray XD2000 System powered by Intel Xeon 8480+ processors — the average impact of air vs. liquid cooling on performance, power consumption, and power efficiency across six standard benchmarks

The results in Figure 7 show that while liquid cooling delivers superior power efficiency, air-cooled servers running the latest Intel Xeon 8480+ processors deliver excellent performance. Air-cooled HPE Cray XD2000 Systems are an excellent solution for customers that are either unable or not yet ready to make the transition to liquid cooling.

Protecting existing investments

Today, many customers run applications that leverage the powerful features of the Intel Xeon Scalable processor family. Whether applications come from commercial ISVs or are developed in-house, developers frequently use Intel instruction set extensions such as Intel AVX-512. They may also rely on the Intel oneAPI Math Kernel Library (MKL) or other high-performance libraries to achieve the highest performance levels.[24]

Deploying HPE Cray XD2000 Systems based on the latest 4th Gen Intel Xeon Scalable processors provides customers with a seamless, risk-free migration path that avoids the need for application recompilation, protecting investments in software tools and application software.

Help maximize performance, flexibility, and value

With power requirements for high-end CPUs rising, many organizations are considering liquid cooling to increase density and improve cooling efficiency. However, this transition can be expensive and disruptive, and not all organizations are ready to take this step.

Fortunately, the latest HPE Cray XD2000 powered by 4th Gen Intel Xeon Scalable processors provide customers the flexibility to navigate this transition at their own pace. With Intel-powered HPE Cray XD2000 Systems, customer can:

- Avoid or delay expensive data center upgrades and refits by extending the viability of air cooling
- Experience up to twice the throughput of previous-generation servers[25]
- Deploy dense, energy-efficient servers to help maximize data center space
- Protect existing investments in software and tools
- Gradually adopt energy-efficient direct liquid cooling at their own pace

[24] Intel oneAPI Math Kernel Library.

[25] 2P Intel Xeon Platinum 8480+ (112C) scored 932 SPECrate 2017_fp_base — spec.org/cpu2017/results/res2023q1/cpu2017-20221204-32903.html. 2P Intel Xeon Platinum 8380 (80C) scoring 467 SPECrate 2017_fp_base — spec.org/cpu2017/results/res2021q2/cpu2017-20210524-26430.html. 932/467 represents an ~2x performance improvement. SPEC, SPEC CPU, SPECfp, and SPECrate are trademarks of the Standard Performance Evaluation Corporation. All rights reserved. All stated results are as of April 15, 2023. See spec.org for more information.

Learn more at

HPE.com/servers/CrayXD2000
Intel.com/Xeon

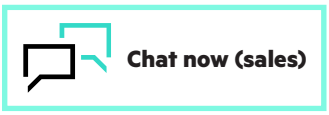

a50008911ENW

Hewlett Packard Enterprise