SC20

From Exascale to Everyscale™

Frequently Asked Questions

November 2020

Lenovo

NVIDIA

intel®

From Exascale to Everyscale™

Enterprises of all sizes get fast, efficient insights at Everyscale by adopting HPC innovation and exascale technology from Lenovo and Intel.

Question: What do you mean by From Exascale to Everyscale™?

Answer: If you're in an enterprise IT environment or someone that does not have access to the top research institutions in the world, you're probably wondering what exascale means for you. And, quite frankly, why should you care about a few discrete systems to which only a handful of researchers have access?

Here's why: The cascading of these exascale building blocks will allow you to compute faster, more efficiently, and at a lower cost than before. That may mean you can take on new workloads that weren't possible before or move that AI project from the lab to a large-scale deployment. **At Lenovo, we call this cascading of exascale era technology "From Exascale to Everyscale™."**

Question: What is Everyscale™?

Answer: Fundamentally, we take core building blocks of the new generation of technology and deliver it at "Everyscale," making it more accessible to everyone, from enterprises of all sizes to small research universities. This cascading approach includes customers just starting and exploring a proof of concept (POC), customers who want to democratize their IT investment by moving from POC to production, and customers who wish to optimize their existing data center performance and efficiency using the latest technologies.

Question: How does Lenovo and Intel help me get access to exascale era technologies?

Answer: Together, we deliver optimized architecture that drives increasing levels of performance, density, and efficiency within the HPC domain—while helping our customers gain actionable insights from their data faster. We provide proven expertise, simple access to HPC at Everyscale, and deliver faster insights and discovery from innovative technologies.

Question: How are you helping customers resolve resource gaps? Please provide some examples?

Answer: We offer proven expertise by leveraging Lenovo resources, tools, services, and support combined with Intel® technologies and open source contributions.

- **Unleash innovation**
  Consult with experts on projects, the right infrastructure to run a proof of concept, and proof of ROI before deployment all available in Lenovo's Customer Innovation Centers for HPC and AI.
- **Better manageability**
  Improve system manageability leveraging a variety of Lenovo system management software, including XClarity (simplifies and automates deployment and management of Lenovo infrastructure), xCAT (open-source distributed computing management software), and confluent software (provides access to various power and cooling data on the monitored hardware).
- **Identify opportunities**
  Discover opportunities in your data center by accessing our experts in Lenovo funded white-boarding workshops.

Question: How are you helping customers access HPC at Everyscale? Please provide some examples?

Answer: We offer Lenovo's flexible consumption models and industry-leading portfolio, combined with Intel's unmatched HPC portfolio, including Intel® Xeon Scalable processors and Intel® Select Solutions.

- **Flexible Consumption**
  Mitigate financial challenges when accessing HPC with Lenovo's choice in flexible consumption models (e.g., TruScale Infrastructure Services), including traditional HPC ownership, HPC Hardware-as-a-Service, and HPC Cloud.
- **Faster deployments**
  Minimize challenges of infrastructure evaluation and deployment with predefined, workload-optimized Intel® Select Solutions, validated by Lenovo, certified by ISVs, and verified by Intel.
- **Confidence at any scale**
  From a single rack installation to hundreds of racks, scale HPC with confidence using Lenovo Scalable Infrastructure (LeSI). This pretested solution is fully validated and built in a Lenovo factory, ready to deploy, all backed by solution level support.

Question: How are you helping customers get fast, efficient insights? Please provide some exampless?

Answer: We leverage maximum efficiency and performance using Lenovo's Everyscale innovation, fueled by Intel's latest technologies.

- **Optimized performance**
  Speed up overall compute performance with Lenovo ThinkSystem servers with Intel® Optane™ Persistent Memory, providing up to triple maximum capacity per node for more scalable memory closer to the CPU and lower data latency compared to solid-state drive technology.
- **Improved efficiency**
  Realize higher performance and energy savings with Lenovo Neptune™ liquid cooling technology using Intel® Xeon® Scalable processors, enabling critical server components to operate at lower temperatures.
- **Faster applications**
  Run AI applications faster using Intel® Deep Learning Boost, built into the Intel® Xeon® Scalable processor on Lenovo ThinkSystem Servers.

Question: Why are Lenovo and Intel collaborating on this effort?

Answer: Last year, Lenovo and Intel announced a collaborative effort to use exascale era technologies and innovation to help customers of all sizes access and leverage HPC at Everyscale. This "From Exascale to Everyscale™ campaign brings it to life.

Question: Where can I get more information on From Exascale to Everyscale™?

Answer: Learn More at the Exascale to Everyscale website, located at lenovo.com/everyscale.

Question: What is Lenovo's leadership position in HPC, and have you received any awards or recognition to substantiate this position?

Answer: Yes, Lenovo is the #1 TOP500 supercomputer provider1, reflecting our commitment to groundbreaking innovation in HPC, Analytics and AI for businesses of any size. Also, the annual HPCWire awards2 recognize the best and brightest innovators within the global HPC community. This year, HPCWire recognized Lenovo in three (3) categories:

1. Best Use of HPC in Industry (Automotive, Aerospace, Manufacturing, Chemical, etc.)
2. Top 5 Vendors to Watch
3. Best HPC Storage Product or Technology

Further, our customer, the University of Birmingham (UK), won the HPCWire award for Best HPC Collaboration (Academia/Government/Industry).

[1] Source: https://top500.org/ Nov 2020 List

[2] Source: https://www.hpcwire.com/

Lenovo ThinkSystem SD650-N V2

Provides unprecedented compute power and density while simultaneously reduces energy consumption and lowers overall data center costs.

Question: What is Lenovo announcing?

Answer: Lenovo is announcing the ThinkSystem SD650-N V2, a new member of our ultra-dense family of systems featuring Neptune™ direct-to-node liquid cooling. The system features two next-generation Intel® Xeon® processors, and four board mounted NVIDIA A100 GPUs in a 1U tray.

Question: What's unique about this system?

Answer: Through our partnership with NVIDIA, we have expanded direct-to-node, warm-water cooling to GPUs and accelerators. This application of liquid cooling to GPUs will allow Lenovo to deliver nearly 3 petaFLOPS of compute power per rack[1]. Lenovo is the first vendor to bring this type of direct liquid cooling for GPUs to market.

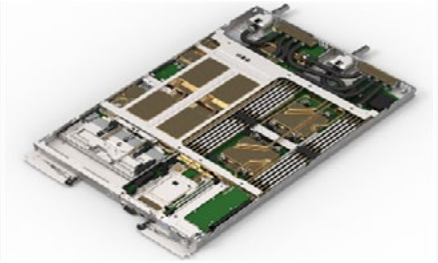

[1] Substantiation: Lenovo internal with NVIDIA. A single node = 84TF (78TF for the GPU, 5.5TF for the CPU**); therefore, a rack = 3PF (84TF per node x 36 nodes/rack). October 2020.

Question: How is the new system different from the existing ThinkSystem SD650?

Answer: Going forward, the new chassis will accommodate the new ThinkSystem SD650-N V2 GPU nodes, as well as future generation compute platforms. Customers will be able to mix both the SD650-N V2 and future compute systems in the new chassis. The new chassis is not backward-compatible to the current generation SD650 platform.

Question: Who is the target market or customer for the ThinkSystem SD650-N V2?

Answer: The ThinkSystem SD650-N V2 is an ultra-dense, high performance system. Customers trying to pack as much compute power in the smallest data center footprint will want to see this system. HPC customers who are considering or have already made the jump to energy-efficient water cooling are prime candidates as well.

Question: When will these systems begin to ship?

Answer: We anticipate shipping the system in the first half of 2021.

Lenovo ThinkSystem SR670 V2

The most versatile Lenovo GPU platform for AI and Enterprise workloads.

Question: What is Lenovo announcing?

Answer: A new addition to our ThinkSystem series, the SR670 V2, supports up to eight full-size (NVIDIA® A100 Tensor Core GPUs / NVIDIA® V100 Tensor Core GPU) or half-sized GPUs in a new 3U frame. It has models that utilize both traditional air and Lenovo Neptune™ liquid-to-air cooling technology. The SR670 V2 supports high GPU density for massive computational workloads, AI and Deep Learning (training and inference), and VDI optimized for the enterprise.

Question: What's unique about this system?

Answer: The SR670 V2 offers up to 160 TFLPS DP peak performance[1] per 2CPU 3U node. Multiple form-factor GPUs work for AI Training and inference workloads as well as enterprise workloads such as VDI.

Additionally, like the SR670, the SR670 V2 offers "integrated modularity," a design that includes no hard-wired PCI slots, instead of using modular flexible PCI lanes are used, so customers have maximum functionality. Integrated modularity separates peripherals from the CPU and allows customers to reconfigure PCI peripherals on the fly by moving cables. An example of this is the ability to assign GPUs to either CPU as the workload demands dynamically.

[1] Substantiation: Lenovo internal with NVIDIA. A single node with 8 x A100 NVIDIA GPUs = 161.5TF (8 GPUs x 19.5 TF Rpeak = 156TF for GPU, plus 5.5TF for the CPU) and a rack = 2.9PF (161.5TF/node x 18 nodes/rack). October 2020.

Question: How is the SR670 V2's cooling different from some of the other Lenovo Neptune™ cooled systems?

Answer: Rather than using direct-to-node cooling (as with the SD650-N V2), the SR670 V2 leverages Neptune™ liquid-to-air heat exchanger, which delivers the cooling properties of liquid in an air-cooled system. This capability allows users to get the performance of a liquid-cooled system without installing any plumbing in the data center. The system also allows for flexible hot or cold aisle configurations.

Question: What are the ThinkSystem SR670 V2 configurations?

Answer: In addition to the choice of cooling architecture, the system offers configurations that support either traditional PCI or NVIDIA's NVLink interconnect. This choice allows for the use of flexible GPU interconnect fabric and lower-cost PCI or higher performance NVLink.

Lenovo ThinkSystem SR860 V2 and SR850 V2

New 4S scalable Mission Critical servers featuring 3rd Generation Intel Xeon Scalable processors for next-gen AI and analytics applications.

Question: What is Lenovo announcing?

Answer: At Tech World 2020, Lenovo announced two new additions to our mission-critical compute portfolio; the ThinkSystem SR860 V2 (4S 4U) and the ThinkSystem SR850 V2. Both servers leverage 3rd Generation Intel® Xeon® Scalable processors, increased memory speeds, and expanded storage and GPU support to address emerging workloads and use cases.

Question: What's unique about these systems?

Answer: Lenovo is one of very few OEMs to introduce platforms that support Intel's Cooper Lake-6 platform and Optane Persistent Memory 200 series. These unique offerings give Lenovo customers access to the next generation of workload and application support, especially for transactional database, analytics, and AI use cases.

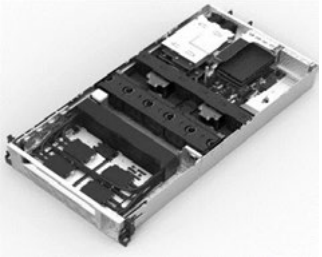

Question: How are the new systems different from their predecessors?

Answer: Lenovo has completely revamped the platforms to support the latest in data center technologies. The SR860 V2 supports copious amounts of ultra-fast memory, considerable onboard storage capabilities with support for up to 48 2.5-inch drives, up to 24 NVMe drive support for latency-sensitive applications, and even up to 4 double-wide 300W GPUs. The SR850 V2 offers highly dense compute power that expands CPU, memory, and onboard storage capabilities too, next-gen networking cards, and hot-swap rear 7mm boot drives. Both systems get updated physical security and onboard monitoring capabilities too.

Question: Who is the target market or customer for the ThinkSystem SR860 V2 and SR850 V2?

Answer: The new SR860 V2 and SR850 V2 systems are ideal for customers running mission-critical applications that cannot afford downtime. Multiple redundant systems combined with next-gen compute and storage capabilities maintain availability for key applications like transactional, in-memory databases, artificial intelligence-enabled business analytics, and enterprise resource management operations.

Question: Are the ThinkSystem SR860 V2, and SR850 V2 supported for SAP HANA workloads?

Answer: Yes, the new SR860 V2 and SR850 V2 systems are scalable to support 2 and 4 socket SAP HANA workloads. SAP HANA information can be found on the SAP HANA certification web site and at www.lenovo.com/sap.

Question: Does the ThinkSystem SR860 V2 hold any world record SAP application benchmarks?

Answer: Yes, the SR860 V2 holds multiple SAP BW Edition for SAP HANA application benchmark records across three workload sizes of 1.3B, 5.2B, and 7.8B records. SAP application benchmark results are posted on www.sap.com/benchmark and Lenovo Press at https://lenovopress.com/servers/benchmarks/sap-bw.

Question: When will these systems begin to ship?

Answer: Both of these systems are currently shipping as of November 12, 2020.

SR860

SR850 V2

GOAST (Genomics Optimization and Scalability Tool)

Accelerate genomics analytics and increase sample throughput at costs much lower than solutions relying on expensive accelerators.

Question: What are the two (2) GOAST configurations, and what do they deliver?

Answer:

1. **GOAST Base** delivers faster execution times, higher sample throughput, and 3X as many more genomes per day than other competing 2-socket solutions.
2. **GOAST Plus** reaches GPU-level speeds but at CPU-level costs—this means accelerated performance at 50% the cost of solutions relying on expensive accelerators and proprietary software.

Lenovo GOAST Base (cost-optimized) Built on a standard 1U 2-Socket server, Lenovo ThinkSystem SR630

Server Type	1x ThinkSystem SR630
Processor	2 x Intel Xeon Gold 624R (24 cores, 205W, 3.0 GHz
Memory	384GB RAM, 12 x ThinkSystem 32 GB, TruDDR4 2933 MHz (2R x 4 1.2V RDIMM)
Storage	ThinkSystem 2.5” PM1643a 1.92TB Entry SAS 12Gb Hot Swap SSD or SSD-based DSS-G GPFA Network Storage

Lenovo GOAST Plus (performance-optimized) Built on a single 8-Socket server – Lenovo ThinkSystem SR950

Server Type	1x ThinkSystem SR950
Processor	8 x Intel® Xeon® Platinum 8280 (28 cores, 205W, 2.7GHz)
Memory	1.5TB RAM, 48 x ThinkSystem 32GB TruDDR4 2933MHz (2R x 4 1.2V) RDIMM
Local Storage	4 x 3.2TB U.2 NVMe PCIe 3.0 x 4 Hot Swap SSD or SSD-based DSS-G GPFS Network Storage

Question: What is the value of the GOAST solution for scientists?

Answer: We help scientists understand data faster and make discoveries sooner with GOAST - Lenovo's easy-to-use genomics-optimized technology solution.

GOAST leverages specially-tuned hardware to accelerate the open-source GATK software suite from the Broad Institute, thus enabling genomics analytics at speed-ups up to 167x compared to typical GATK environments. We make no code-level changes to GATK. Instead, our recipe optimizes parameters fed to the tools while sprinkling key system tunings to Lenovo hardware so it can run at its fullest capacity when performing genomics analytics.

Our genomics optimizations are based on extensive systematic testing of the factors affecting performance at our in-house Genomics R&D lab. As a result of our R&D work, we have developed two configurations to accelerate variant calling workflows: GOAST Base and GOAST Plus, optimized for cost vs. extreme performance, respectively. Compared to standard solutions, GOAST Base delivers 18x-45x improvement, while GOAST Plus reaches accelerations up to 167x.

GOAST improves lab productivity by allowing you to process more genomes concurrently—this means higher throughputs and shorter execution times. Processing more genomes means you get answers faster, and your discoveries save more lives at a fraction of the cost of alternatives. The tech revolution of accessible innovation is here, be it for basic research, infectious disease, or precision medicine.

Question: Is the GOAST solution for Genomics Analytics only?

Answer: With the GOAST architecture, we've made it possible to run other life sciences workloads on a single, optimized platform. From Genomics to Transcriptomics, Molecular Dynamics, and CryoEM. Moreover, the system tunings and infrastructure in GOAST will benefit any workloads requiring high core, high memory, fast storage, and/or high I/O (e.g., most bioinformatics software pipelines).

Configuration Size	Workload		
	Molecular Dynamics[3]	Cryo-EM[4]	Genomics Analytics[5]
Small 20 compute nodes ~2 PB storage	2,509.74 ns/day	One (1) microscope	25K whole genomes per year
Medium 40 compute nodes ~4 PB storage	5,019.48 ns/day	Two (2) microscopes	50K whole genomes per year
Large 80 compute nodes ~8 PB storage	10,038.96 ns/day	Four (4) microscopes	100K whole genomes per year

Question: Where can I get more information on Lenovo GOAST?

Answer:

- Ted Talk (Oct 2020) Combining high performance computing & AI to solve humanity's greatest challenges
- GOAST Technical White Paper (Oct 2020) Genomics Optimized Servers Accelerate COVID-19 Discovery with High-Throughput Analytics

Contact our HPC Experts:

- North America: Francois Corradino <fcorradino@lenovo.com>
- Europe and the Middle East: Noam Rosen <nrosen1@Lenovo.com>
- Asia Pacific: Kumar Raghavan <kthirugokarn@lenovo.com>
- Australia and New Zealand: Joao Almeida <joaoalmeida@lenovo.com>
- Latin America: Roberto Brandao <rbrandao@Lenovo.com>

DAOS Storage on ThinkSystem SR630

Question: What is the DAOS Exascale Storage Stack?

Answer: The Distributed Asynchronous Object Storage (DAOS) is an open-source software-defined object store, designed from the ground up for massively distributed Non-Volatile Memory (NVM). DAOS takes advantage of next-generation NVM technology like Storage Class Memory (SCM) and NVM express (NVMe). It presents a key-value storage interface and provides features such as transactional non-blocking I/O, advanced data protection with self-healing on top of commodity hardware, end-to-end data integrity, fine-grained data control, and elastic storage to optimize performance and cost.

Question: What are some of the challenges with existing distributed storage options?

Answer: The emergence of data-intensive applications in business, government, and academia stretches the existing I/O models beyond limits. Modern I/O workloads feature an increasing proportion of metadata, combined with misaligned and fragmented data. Conventional storage stacks deliver poor performance for these workloads by adding a lot of latency and introducing alignment constraints.

Question: How can these challenges be addressed?

Answer: The advent of affordable large-capacity persistent memory combined with an integrated HPC fabric offers a unique opportunity to redefine the storage paradigm and support modern I/O workloads efficiently. This revolution requires a radical rethinking of the complete storage stack. To unleash the full potential of this new technology, the new stack must embrace a byte-granular shared-nothing interface from the ground up. It must be able to support massively distributed storage for which failure will be the norm, while preserving low latency and high bandwidth access over the fabric. DAOS is a complete I/O architecture that aggregates SCM and NVMe storage distributed across the fabric. It provides globally-accessible object address spaces, consistency, availability, and resiliency guarantees without compromising performance.

Question: What are Intel and Lenovo doing together with DAOS?

Answer: DAOS is developed primarily by Intel, and the DAOS storage stack is an essential part of the multi-year strategic partnership of Intel and Lenovo in High Performance Computing. Under this partnership, Lenovo is actively contributing to the development, documentation, testing, and hardening of the DAOS software stack. We are also developing optimized hardware configurations for DAOS servers on current and future generations of Intel Xeon platforms.

Question: What are the typical workloads and use cases for DAOS?

Answer: Applications that perform many small and/or unaligned I/O requests will benefit most from DAOS. While DAOS provides a traditional POSIX interface, the best performance will be achieved when directly interfacing with the DAOS Storage API. Frameworks like MPI-IO and HDF5 have already been enabled to work directly over DAOS, providing immediate value to many HPC and data analytics applications. Apache Spark is supported through a DAOS backend to its Hadoop file system interface.

Question: Which benchmarks demonstrate the performance that DAOS can achieve?

Answer: The IO500 benchmark suite contains a representative set of synthetic HPC storage benchmarks. It measures bandwidth (IOR), metadata rates (mdtest), and filesystem traversals (find) and contains both "easy" and "hard" use cases. DAOS excels on the IO500 benchmarks, and in particular, on the "hard" workloads. See https://www.vi4io.org/io500/start.

Full application workloads are currently benchmarked on DAOS storage, and results will be shared in the near future. Please contact us at hpcstorage@lenovo.com if you are interested in a DAOS Proof-of-Concept with your specific storage workload.

Question: What is the best Lenovo Server to implement a DAOS Storage Solution?

Answer: The Lenovo ThinkSystem SR630 is a 1U, 2-socket Intel Xeon server that is perfectly suited for DAOS. Each of its two sockets runs one instance of the DAOS data plane and manages half of its hardware resources (CPU cores, Optane PMem, NVMe, NICs).

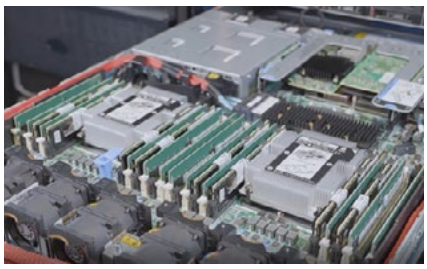

Question: How is Lenovo supporting DAOS in its HPC Innovation Center?

Answer: Lenovo's worldwide HPC Benchmarking Center is located in Stuttgart, Germany. Here we operate a cluster of 12x DAOS Servers (ThinkSystem SR630) that is used for internal benchmarking and customer proof-of-concept activities. Each DAOS server is equipped with 12x Optane DC Persistent Memory Modules, 8x NVMe disks, and 2x EDR InfiniBand. The total raw capacity is 150 TB NVMe + 18 TB DCPMM, with a peak bandwidth of ~240 GB/s (read). This DAOS cluster can be accessed from Slurm jobs that execute on EDR compute nodes.

Question: Where can I get more information on DAOS?

Answer:

- DAOS User Group (DUG20) meeting at SC 2020 LINK
- Intel/Lenovo DAOS Paper in the SC-Asia 2020 proceedings (LNCS 12082) LINK
- Lenovo SR630 Product Guide posted on Lenovo Press (LP1049) LINK
- Lenovo Optane PMem Product Guide on Lenovo Press (LP1066) LINK
- Future of HPC Storage and DAOS Video from ISC 2020 LINK

Lenovo Services: Workshops, Power & Cooling Offerings, Neptune Focus

Question: What is a Services-Led Transformation, and how does it benefit our customers?

Answer: With more than 3,000 implementations and their related project experiences, Lenovo Professional Services has a team of experts with a vast knowledge of industry best practices to break down customer problems in their digital transformation journey. Services can also integrate hardware and software into a well-planned solution to achieve desired business outcomes for customers.

This ensures the customer is getting the best possible solution featuring the latest in market trends and anticipating their future needs according to their own unique strategic goals.

Question: What are Lenovo-funded workshops, and what's the customer value?

Answer: Principal Consultant-led workshops are typically valued at $4,000 but funded by Lenovo as an investment to help customers drive business decisions as they kickstart their Analytics to AI journey. With Lenovo-funded workshops, customers can connect to a subject matter expert (our Principal Consultants) to white-board various options to accomplish their objectives.

Question: Is this open for anyone?

Answer: Lenovo-funded workshops are designed for customers with propelling business problems that need clarity from subject matter experts.

Question: Can we skip the workshops and just start with an Assessment?

Answer: It's not recommended. We highly encourage customers to start with a white-boarding workshop session. This not only allows us to pinpoint exactly what the customer needs instead of offering a cookie-cutter solution, but it also allows for Lenovo's direct investment in the customers' long-term success as true business partners.

Question: What are the Lenovo Professional Services capabilities supporting Power & Cooling Technology for the Data Center?

Answer: Lenovo Professional Services help you determine whether your Data Center Power and Cooling Infrastructure meet your business expectations with technical assessments such as Power and Cooling Resiliency Assessment, Data Center Baseline Cooling, and Dense/HPC System Thermal Modeling.

Power and Cooling Best Practices

- Identify opportunities for IT system and infrastructure improvements with Data Center Best Practices Workshop, a one-day, onsite interactive workshop to review power, cooling and energy efficiency associated with IT systems and data center infrastructure.
- Ensure your data center has the required power and cooling capacity for present and future IT plans with Data Center Power and Cooling Assessment.

Lenovo Neptune Services

- Get help with your Data Center design, planning, installation, and implementation with DTN Cooling for Neptune Portfolio Systems.
- Talk to the experts when planning, designing, installing, or implementing a Rear Door Heat eXchanger with Data Center Rack Cooling with Rear Door Heat eXchanger.

Question: What's the process for engagement?

Answer: To see if your business issue qualifies for a no-cost workshop, reach out to your services representatives in each GEO. The workshops are scheduled on a first-come, first-served, limited basis. Alternatively, please reach out to our Lenovo Professional Services HPC and Power & Cooling Experts:

- **Jerrod Buterbaugh**, jbuterbaugh@lenovo.com
- **Milind Salokhe**, msalokhe@lenovo.com

Email Contact	Geo / Coverage
dcg_ps_na@lenovo.com	North America (USA, Canada)
dcg_ps_ap@lenovo.com	Asia Pacific
dcg_ps_la@lenovo.com	Latin America (Brazil, Mexico, SSA)
dcg_ps_emea@lenovo.com	Europe, Middle East, Africa

Question: What are the different types of power and cooling assessments available to customers?

Answer: There are several available specifically for customers with those needs, including:

- TCO analysis
- Thermal modeling
- Optimization
- Thermal validation
- Problem resolution
- Thermodynamic data analysis
- Power and cooling resiliency
- Power distribution analysis

SAP HANA Enterprise Cloud, customer edition

New SAP HANA Enterprise Cloud customer edition Offering® Leveraging Lenovo TruScale, Enabling New Digital Transformation Capability.

Question: What are Lenovo and SAP announcing?

Answer: SAP delivers an OPEX private-cloud service with SAP HANA Enterprise Cloud, customer edition, leveraging Lenovo TruScale Infrastructure Services, and Lenovo ThinkSystem and ThinkAgile servers and ThinkSystem storage that are SAP HANA certified and monitored and managed by Lenovo and SAP white-glove services. It is a turn-key ACV cloud offering fully compliant with SAP HEC white-glove standard of service, architecture, & security – delivered on customer premises – and backed by an SLA only SAP can assure. The offering enables customers running SAP applications on-premises with traditional perpetual software licenses and capital equipment to move to a cloud subscription model.

Question: How does the SAP's HEC ce service compare to SAP's hosted HEC?

Answer: SAP HEC CE follows a standard HEC reference architecture and security standards designed and pre-built for SAP S/4HANA deployment and is deployed in the customer's data center or co-lo facility. The reference architecture is globally consistent and adopts best practices across the multi-cloud ecosystem.

Question: What services does Lenovo deliver, and what services does SAP provide?

Answer: SAP delivers the cloud solution to the customer and provides SAP application services, and SAP leverages Lenovo TruScale infrastructure as a service with white-glove services to deploy, manage, maintain, and support the infrastructure.

Question: What Lenovo servers and storage are used for HEC ce reference architecture?

Answer: The Lenovo systems include: ThinkSystem SR950, ThinkSystem SR850P, ThinkAgile VX, and ThinkSystem DM7100 and can scale to meet the most demanding SAP application workloads.

Question: Who will sell this solution?

Answer: SAP sells the solution in collaboration with Lenovo.

Question: Is VMware software also included in the SAP HEC ce offering?

Answer: Yes, VMware is deeply integrated with the SAP HANA Enterprise Cloud management, monitoring, and reporting service delivery toolchain with Lenovo systems. With these partnerships, the SAP HANA Enterprise Cloud, customer edition, offers SAP customers the merits of an operation expense (OPEX) business model, elastic computing, and concierge managed services.

Question: Where can I find more information?

Answer: Information is posted on the following web sites.

SAP HEC customer edition:

- Go to: www.sap.com/hec
- Contact: HEC_Global_GTM@sap.com

Lenovo Systems, TruScale Infrastructure Services, and SAP alliance:

- Go to: www.lenovo.com/truscale and www.lenovo.com/sap
- Contact:
  - North America (USA, Canada)
  - Asia Pacific (GCG, ANZ, ASEAN, Japan, Korea, ISA)
  - Latin America (Brazil, Mexico, SSA)
  - Europe, Middle East, Africa

For more information on VMware Solutions:

- Go to: https://www.vmware.com/solutions/business-critical-apps/sap-virtualization.html

Question: Why are SAP customers interested in having the HEC delivered on-premises?

Answer: Key reasons SAP customers may want to keep their SAP applications on-premises include:

- SAP system entanglement with non-SAP systems
- Customer data sovereignty, privacy, and residency requirements
- Regulatory & governance
- Public cloud scarcity / no POPs
- Bilateral contract requirements (single throat to choke)
- Difficult to migrate, multiple entangled non-SAP legacy landscapes
- Need for low-latency, high-throughput communications
- Current infrastructure due for refresh (high HW maintenance costs)
- Low tolerance for risk

Question: What makes the SAP / Lenovo relationship so valued?

Answer: SAP and Lenovo have a long and trusted history together, delivering technology innovation. We have a truly 360-degree relationship as a partner, supplier, and customer. Our trusted partnership is based on both Lenovo and SAP, using what we sell and selling what we use. Lenovo has also been a long-standing Global Technology Partner, holding benchmark records for SAP HANA as well as providing a wide array of SAP HANA certified solutions for SAP customers. Lenovo, as a customer, is planning to deploy SAP S4/HANA and embarking on the journey to become an intelligent enterprise. SAP runs its SAP HANA platform on Lenovo systems and now leverages Lenovo TruScale Infrastructure Services for SAP HEC CE. Most importantly, both companies are committed to integrating digital technology into all areas of a business to change how companies operate and deliver value to customers fundamentally.

Lenovo and SAP enjoy more than 25 years of co-innovation. Some of the most significant milestones and innovations in 2020 include:

- Lenovo's lifetime shipments of SAP HANA systems reached over 15,000 worldwide, ranking it among the top global leaders in this segment.
- Lenovo is the 2019 SAP Pinnacle Award winner for the infrastructure category.
- World record SAP application performance with a long track record of world record SAP application benchmarks (13 current world records as of 9/1/2020)