


## CASE STUDY

High Performance Computing (HPC)  
2nd Generation Intel® Xeon® Scalable Processors  
Intel® Server System S9200WK Family



# Lawrence Livermore National Laboratory Readies Deployment of 5.4 petaFLOP Cluster from Penguin Computing

## Latest Intel HPC components plus innovative liquid cooling from CoolIT Systems power Magma

### Magma Highlights:

- Largest Intel® Xeon® Platinum 9200 processor series-based supercomputer in the United States
- 760 dual-socket compute nodes (Intel® Xeon® Platinum 9242 processors—72,960 cores)
- Built on Intel® Server System S9200WK family
- Built by Penguin Computing with CoolIT Systems advanced liquid cooling



*"We're continually tracking advancements in technologies and looking for capable and economic HPC solutions for LLNL scientists. Our workloads—although they are intensive on the network—they are most intensive on memory bandwidth. So, we looked at the Intel solutions based on the Intel® Xeon® 9200 series processors."*

– Dr. Matt Leininger, Senior Principal HPC Strategist at LLNL

### Executive Summary

[Lawrence Livermore National Laboratory \(LLNL\)](#) is one of the Tri-Labs of the [National Nuclear Security Administration](#), where High Performance Computing (HPC) clusters in the Commodity Technology Systems (CTS-1) program provide over 25 peta-FLOPS of computing capacity across the three labs. For its latest HPC acquisition, LLNL turned to Penguin Computing to build a new system using the latest Intel HPC technologies and a unique liquid cooling solution from [CoolIT Systems](#).

### Challenge

The National Nuclear Security Administration's (NNSA) core mission is the responsible stewardship of America's nuclear stockpile through the application of unparalleled science, technology, engineering, and manufacturing. Computer simulations are critical to the work done by the scientists in the NNSA. They use high performance computing (HPC) resources at three national laboratories abbreviated as the Tri-Labs—Sandia National Laboratories (SNL), Los Alamos National Laboratory (LANL), and Lawrence Livermore National Laboratory (LLNL). In 2016, the Tri-Labs began acquiring a fleet of new HPC systems built by [Penguin Computing](#) under the Commodity Technology Systems program (CTS-1), which has delivered more than 25 petaFLOPS of computing across the three institutions through 2019.

CTS-1 systems are used as the everyday workhorses for Tri-Lab scientists and engineers researching a range of problems in hydrodynamics, materials science, molecular dynamics, and particle transport. Some of the CTS-1 systems are also dedicated to institutional computing and collaborations with industry and academia. The NNSA needed additional capacity at LLNL dedicated to the simulation of 2D and 3D physical systems for parametric studies. The new system being built is called Magma.

### Solution

Magma, with 760 compute and user nodes plus 12 infrastructure nodes, is a liquid-cooled supercomputer built by Penguin Computing using Intel® Server System S9200WK family chassis with Intel® Xeon® Platinum 9242 processors (compute nodes), Intel® Xeon® Platinum 8200 processors (management and file system access nodes), Intel® Omni-Path Architecture (Intel® OPA) fabric, and a unique liquid cooling system designed by CoolIT Systems. Delivered in the first quarter of 2020, it offers LLNL scientists 5.4 petaFLOPS of computational capacity (theoretical).

## First Large-scale Intel® S9200WK Server-based Supercomputer in the U.S.

When NNSA needed to add compute cycles to their fleet of CTS-1 systems, the Intel Server System S9200WK family and Intel Xeon Platinum 9200 processors were in early launch.

“We’re continually tracking advancements in technologies and looking for capable and economic HPC solutions for LLNL scientists,” explained Dr. Matt Leininger, Senior Principal HPC Strategist at LLNL. “Our workloads—although they are intensive on the network—they are most intensive on memory bandwidth. So, we looked at the Intel solutions based on the Intel Xeon 9200 series processors.”

“One of the things we like about the Intel Xeon Platinum processors,” continued Leininger, “is that they have a tremendous amount of memory bandwidth per node, and therefore we can remove that bottleneck from our application and deliver both capable and economical cycles to our mission critical applications.”


The advanced processors offered LLNL’s HPC architects a new level of compute performance compared to their existing CTS-1 resources built on Intel Xeon E5-2600 v4 processors in 2016.<sup>1</sup>

“We also required the system to be liquid cooled,” added Leininger. “Liquid cooling allows LLNL to utilize the higher performance processors in a high-density solution while also easing the air-cooling requirements within our data centers.”

With Intel Xeon Platinum 9200 processors with liquid cooling, Penguin was able to provide a competitive high-density system offering with outstanding performance per core.<sup>2</sup>

Penguin Computing has deployed over 1,000 CTS-1 racks based on their Tundra Systems (built on the OCP—Open Compute Project) architecture, which enables very high-density in a DC-powered rack. The Penguin Computing Tundra solution included a DC-powered version of Intel OPA switch and several options for both liquid and air cooling. The OCP design gave the Tri-Labs high capacity in a smaller space than standard rack designs.

“The increased memory bandwidth of the Intel Xeon Platinum 9242 processors was compelling, along with the quick availability allowing for a quick deployment, were the two major driving factors in selecting the configuration,” stated Ken Gudenrath, DOE Director at Penguin Computing. “To ensure a quick and complete cluster solution, we partnered with Intel using our Relion XE2142eAP 2U4N server in a standard EIA rack.”



**Innovative cooling technology offers high serviceability and provides very stable liquid cooling across the DIMMs.**



**Magma, with 760 compute and user nodes plus 12 management and file access nodes, is a liquid cooled supercomputer built by Penguin Computing featuring Intel® Xeon® Platinum processors and Intel® Omni-Path Architecture fabric. (Photo courtesy Lawrence Livermore National Laboratory/Meg Epperly)**

“One of the things we like about the Intel Xeon Platinum processors is that they have a tremendous amount of memory bandwidth per node, and therefore we can remove that bottleneck from our application and deliver both capable and economical cycles to our mission critical applications.”

—Dr. Matt Leininger, Senior Principal HPC Strategist at LLNL

Penguin Computing purchased fully integrated Intel Server System components—with liquid cooling support—and worked with Intel and CoolIT Systems to design the remaining liquid direct-chip cooling and Cooling Distribution Units (CDU) for the datacenter.

“The quick collaboration amongst all the stakeholders allowed for fast design, contract execution, delivery, and ultimate acceptance of Magma,” added Gudenrath. “Our final goal was achieved when we completed several initial high-performance Linpack runs and submitted these for qualifying on the November 2019 Top500 list.”

Magma comprises 760 dual-socket compute nodes built on Intel S9200WK servers with Intel Xeon Platinum 9242 processors (total of 72,960 cores). Twelve more nodes provide management and file system access, using 2nd Generation Intel Xeon Gold Scalable processors. Like other CTS-1 systems, the fabric is based on Intel OPA. Due to the higher performance node requirements, LLNL doubled the on-node network performance by adding a second Intel OPA host adapter for each node.

Magma performance previews ranked #69 on the November 2019 Top500 list with 3.4 petaFLOPS from 650 compute nodes.<sup>1</sup> The system’s theoretical peak is 5.4 petaFLOPS using all 760 compute nodes.

### Innovative Cooling Enables Maximum Performance

HPC architects have integrated liquid cooling into systems for several years. Numerous CTS-1 purchases included direct-to-chip liquid cooling. According to Leininger, LLNL’s experience with liquid cooling shows that it makes a significant difference with modern processors. Prior to more advanced CPU designs, air cooling offered adequate thermal protection to run the systems at full performance. Today’s

modern processors require liquid cooling to reach their maximum compute capability.

Direct-to-chip liquid cooling in a system the size of an HPC cluster has always introduced a level of complexity into the design. Since direct-to-chip cooling brings coolant right to the component(s), the cooling solution adds difficulty—and thus service cost—to replacing parts, sometimes bringing an entire server down to move the cooling structure out of the way, replace the component, and bring it back into service. System leaks or failures amplify the difficulties and costs.

CoolIT considers itself a trusted solution provider to the OEMs, working closely to deliver cooling solutions that optimize performance while optimizing serviceability.

“We’re used to designing innovative, custom cooling systems for large clusters,” commented Jason Zeiler from CoolIT Systems. “Our own CDU interface between the facility liquid, subfloor piping, and the secondary side technology in the rack. Our offering to the market is as a technology leader, integration collaborator, and solution provider.”

According to Zeiler, memory failure rates in large systems are high across the industry, so easily serviceable memory cooling is a high priority going forward in HPC clusters. According to Leininger, DIMMs are the single most-replaced component in the CTS-1 clusters. So, Magma uniquely brings liquid cooling right to memory.

“Our design allows for high-density memory heat capture,” explains Zeiler. “It provides very stable liquid cooling across the DIMMs. A key objective for our design was to provide both a cost-effective, high heat capture solution for memory while also maintaining very high serviceability. Our design allows for a high number of insertion cycles per DIMM, allowing them to be removed and replaced without any significant impact to the liquid cooling design.”

With the size of Magma, and especially the amount of DIMMs per server, adding traditional liquid cooling directly to the DIMMs had the potential of significantly magnifying service complexity. But, CoolIT used innovative blind-mate, dry-break quick disconnect connectors to mate the component piping to the server board in each server and between the server and coolant manifold in the back of the rack.

Blind-mate, dry-break connectors automatically mate with a chassis manifold at the component level without having to manually disconnect the plumbing. Unplugging a server automatically unplugs the coolant lines without leaking.

“The Intel server design is very user friendly for liquid cooling with blind-mate connectors,” explained Zeiler. “Liquid cooling often adds an element of mild complexity with manual-mate style quick disconnects. But a blind mate design automatically engages with the system, so users are not really interacting with the connector at all. It’s just as simple as electronic blind mate connections.”

“Our admins really like the serviceability around the memory and that its being liquid cooled as well,” added Leininger. “LLNL was worried that the liquid cooling serviceability would be complex and potentially messy. However, CoolIT designed a clean and non-invasive solution. So that was one thing that definitely impressed our folks as we were making decisions.”

## RESULT

Magma was deployed in the first quarter of 2020. Of the 11 CTS-1 systems on the latest Top500 list, Magma holds the highest ranking at 69. Compared to Jade, a CTS-1 system at LLNL, Magma has 35 percent fewer cores than Jade, but it delivers more than 1.2X higher Rmax.<sup>1</sup> Jade is built on Intel Xeon E5-2695 v4 processors, illustrating the performance benefit of latest generation Intel Xeon processors.

Magma will deliver an additional 5.4 petaFLOPS of theoretical peak performance to NNSA resources, offering more than 25 petaFLOPS total computing capacity to the Tri-Labs for Stockpile Stewardship and scientific discovery.<sup>2</sup>

## Solution Summary


Needing more computing capacity for its national security mission, NNSA funded the building of Magma, a theoretical 5.4 petaFLOPS supercomputer built on the latest Intel Xeon Server Systems S9200WK family using Intel Xeon Platinum 9242 processors and Intel OPA fabric. Liquid cooling across the chassis includes direct cooling to the memory with a blind-mate, dry-break system that simplifies serviceability. Working with Penguin Computing, who designed Magma, CoolIT provided the unique and innovative cooling solution.

### Where to Get More Information

- Learn more about [Penguin Computing’s solutions](#).
- Find out more about [CoolIT Systems solutions](#).
- Learn more about [Intel Omni-Path Architecture](#).
- Learn more about [Intel Server System S9200WK](#).
- Explore the capabilities of the [2nd Generation Intel Xeon Scalable processors](#) with integrated Intel® Deep Learning Boost capabilities for accelerated AI inferencing.

### Solution Ingredients

- Intel Xeon Platinum processor 9200 series
- 772 nodes (760 compute nodes with 72,960 cores)
- Built on Intel Server System S9200WK product family
- Built by Penguin Computing with CoolIT Systems advanced liquid cooling



<sup>1</sup> Data provided by Lawrence Livermore National Laboratory

<sup>2</sup> <https://www.enterpriseai.news/2019/11/18/ai-ready-2nd-generation-intel-xeon-platinum-9200-processors-demonstrate-leadership-performance/>

For details, visit <http://www.intel.com/2019xeonconfigs/> (Intel Xeon Scalable processors – claim #31). For additional detail visit <https://www.intel.com/content/www/us/en/high-performance-computing/performance-for-hpc-platforms.html>

<sup>3</sup> <https://top500.org/lists/2019/11/>

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit [www.intel.com/benchmarks](http://www.intel.com/benchmarks).

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available security updates. No product or component can be absolutely secure.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy. Results have been estimated or simulated. Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.