

## Special Report

# Perspectives from SC25

Mark Nossokoff, Bob Sorensen, and Earl Joseph  
November 2025

### HYPERION RESEARCH OPINION

---

While attendance declined year-over-year in the post-covid era (approximately 16,500 at SC25 vs. approximately 18,000 at SC24), there was no shortage of energy and engagement for the users and vendors who participated in St. Louis. This was especially evident on the show floor, with a more than 10% increase in the number of vendors represented (a record 560 at SC25 vs. 500 at SC24). As is our custom, the Hyperion Research team of analysts has compiled its primary takeaways and perspectives from more than 70 customer meetings, workshops, session attendance, and other interactions at the event:

- The Top 500 list continues its slow decline as an HPC sector bellwether.
- The US government shutdown had a visible negative impact on user attendance and exhibitors.
- The nine new US Department of Energy (DOE) machines announced the week before SC25 created much buzz, especially given the dearth of details surrounding them.
- AI factories remain a central topic of discussion, although a universally accepted definition has yet to be established.
- Interest in quantum computing is on the rise within the HPC user base.
- Innovations in liquid cooling and related energy and power management is becoming mainstream within the HPC and AI communities.
- The future of mixed precision is really a two-part discussion.
- Skepticism about AI's potential to produce quality HPC software is widespread.
- There was a noticeable increase in niche HPC supply chain exhibitors, including those that sell connectors, pumps, valves, power supplies, and even cooling fans.
- Despite the recent developments in national-level policies on tariffs, export controls, and sovereignty dynamics happening around the world, there was little official SC coverage on the impact of these policies on the HPC sector overall.

In addition to this SC25 report, Hyperion Research's breakfast briefing market update slides are available at <https://tinyurl.com/5e6dkufy>. The HPC market is in a strong growth mode, heavily driven by AI, cloud usage, and increased spending on traditional HPC applications.

## PERSPECTIVES AND TAKEAWAYS FROM SC25

---

### **The Top 500 list continues its slow decline as an HPC sector bellwether**

Only 45 new machines were added to the list, a turnover of just 9%, and there was no movement in the top 10 except for two performance increases for US DOE's Lawrence Livermore National Laboratory's El Capitan and the EuroHPC Jupiter Booster at Germany's Julich Supercomputer Center.

- Indeed, the largest new system, the CHIE-4 system installed by Japan's SoftBank, only ranks at number 17 on the list. The Shaheen-III GPU partition at KAUST, one of only five Nvidia CPU-GPU superchip systems added to this most recent list, weighed in at number 18.
- China continues its boycott of the Top500 list, adding no new machines once again.

Perhaps more illustrative, the Top 500 list is no longer a single list but has diversified to better reflect the increasingly diverse range of HPC workloads. Statistics are now being gathered and presented along with the traditional Top 500 list for the Green Top500 (quantifies HPC energy efficiency), the HPCG Top 500 (stresses more typical HPC workloads), and the HPL-MxP benchmark (captures the emerging convergence of HPC and AI workloads).

### **The US government shutdown had a visible negative impact on user attendance and exhibitors**

While the US government shutdown technically ended prior to the start of the show, many government organizations scaled back the number of attendees from their areas, and many researchers and scientists were not able to secure funding and approval to attend. Additionally, government agencies who typically have a large and popular presence on the show floor (e.g., NASA) had no exhibits this year. Hopefully this won't turn into a recurring trend.

### **The nine new US DOE machines announced the week before SC25 created much buzz, especially given the dearth of details surrounding them**

With the announcement of nine new machines prior to the start of SC25, the US DOE displayed a strong commitment to increased investment in provisioning the leadership computing infrastructure for both HPC and AI capabilities. Targeted use cases included several machines aimed at extreme-scale traditional HPC workloads and others aimed at large-scale AI workflows. Unique to some of the machines is the employment of a cloud service provider to manage several of the AI machines, seemingly in an "on-prem cloud" operating environment.

That said, there were scant details in several key areas, including the cost of each machine, expected power consumption, and anticipated utilization of the AI machines between DOE mission and external, non-DOE users.

### **AI factories remain a central topic of discussion, although a universally accepted definition has yet to be established**

Nations around the world are investing heavily in AI "factories", although the approach, system designs, and use cases vary greatly. Europe was early in announcing and setting up AI factories, while the US more recently launched a number of new AI systems. The concerns about AI's power requirements seem to have taken a back seat to the need to install large systems that are currently required to make AI applications successful.

## **Interest in quantum computing is on the rise within the HPC user base**

As many believe that it may be only 2-3 years until quantum computing can clearly demonstrate performance advantage over classical counterparts, HPC end users were looking more attentively than previous SC gatherings at what quantum computing can add to the classical HPC ecosystem.

Many of the major QC vendors were on the exhibition floor, and instead of fielding questions about what a quantum computer is and how it works, they were addressing attendees' interests in exploring QC use case opportunities, QC/HPC software requirements, and details on integration requirements of QC into a classical HPC environment. Not only were the panel sessions and BOFs that addressed these issues typically standing room only, but that so many were scheduled by SC25 planners equally validated HPC interest in this emerging technology.

## **Innovations in liquid cooling and related energy and power management is becoming mainstream within the HPC and AI communities**

Another holdover from past shows is the continued growth of real estate on the exhibition floor consumed by cooling, energy, and power management solution vendors. This reflects the increasing demand created by virtually all new machines integrating the latest accelerator technologies requiring liquid cooling.

Liquid cooling, both direct-to-chip and immersion implementations, are now recognized as required for leading edge HPC and AI infrastructure solutions and vendors are scrambling to establish technology leadership and market share.

## **The future of mixed precision is really a two-part discussion**

There are two distinct issues surrounding the growing interest in mixed precision for the HPC sector:

1. The first is the somewhat logical movement for HPC end users to look for opportunities to realign legacy science and engineering codes away from FP64 towards lower precision formats, such as FP32 and FP16, while still preserving code accuracy. The accepted reality is that such lower precision implementations can reduce power demands, ease memory constraints, and overall provide better time to solution. A number of efforts were discussed in panel sessions to detail the software tools that can help explore conversion processes. Such efforts will likely become more prevalent in the next few years.
2. Conversely, discussions surrounding the growing processor design emphasis on low precision data formats, currently bottoming at FP4, was met with little enthusiasm by HPC end users. Outside the well documented performance benefits of low precision formats targeted for AI application, there was a general consensus that these AI-centric formats offer little, if any, potential advantage for overall science and engineering community.

## **Skepticism about AI's potential to produce quality HPC software is widespread**

A number of researchers exploring AI schemes to generate advanced scientific or engineering code highlighted a number of pernicious issues. Some stated that the ability of AI to effectively manage large chunks of code conversion was limited.

- For example, one researcher cited the requirement to break down large legacy codes into small code sections in order to get correct, validated AI-generated output code. In those cases, the onus for taking complex and lengthy legacy codes and dividing them into small

independent code sections fell to the programmers and code developers, and the process was seen as overly complex with little current automated support for such conversions.

- Likewise, another researcher noted that AI generated code was often convoluted and difficult to understand by code developers. The researcher noted that AI generated code was not generally human understandable and did not follow traditional coding formalisms, complicating efforts to optimize, update, debug, or even understand the flow and intent of AI generated software.
- Another expert noted that AI code generation was generally acceptable for Python codes but was much less effective for typical science and engineering codes centered on traditional HPC languages such as C or Fortran.

### **There was a noticeable increase in niche HPC supply chain exhibitors including those that sell connectors, pumps, valves, power supplies and even cooling fans**

In addition to many of the perennial favorites, this year sported a noticeable increase in the number of HPC supply chain vendors, many centered on cooling components, both liquid and air, as well as racks, boards, and power supplies. It was clear that many of these vendors were not targeting the highest end of computing but instead were looking to gain a presence in the more diverse mid and low end HPC sector. This may be a harbinger that the SC conference, and indeed the HPC sector in general, is no longer all about the highest performant computers but is expanding to embrace the technical requirements of general-purpose, and more numerous, HPC suppliers that target more conventional HPC markets.

### **Despite the recent developments in national-level policies on tariffs, export controls, and sovereignty dynamics happening around the world, there was little official SC coverage on the impact of these policies on the HPC sector**

Besides the obvious impact of the US government shutdown limiting the presence of a number of USG participants, there was little attention paid to the current spate of national policies directed at either promoting or protecting national-level advanced computing interests. Granted, the SC conferences have always strived to put technology first, but as a clear dual-use technology HPCs have always been subject to government strategies, often with mixed results. Perhaps HPC researchers, suppliers, end users, and relevant policy makers could benefit from SC-enabled opportunities to have open discussions on government policies and practices that can either boost or hinder the trajectory of the sector.

## **FUTURE OUTLOOK**

---

Despite lower attendance, SC25 appeared to be a clear success from many perspectives:

- A record number of exhibitors provided a broad range of current HPC- and AI-related products and solutions and created expectations of continued innovations.
- A wide breadth, depth, and expertise of topics was shared at conference sessions.
- Vibrant and productive conversations, both formally and informally, took place across a diverse set of venues and networking events.

The global HPC community took full advantage of being able to congregate and collaborate to provide meaningful opportunities for continued advancement in technological innovations. Scientists,

engineers, and researchers across government, industry, and academia are directly involved in these conversations and are looking forward to advanced tools and capabilities to accelerate their research and discovery for solutions to the world's most challenging issues.

Attention now turns to SC26 in Chicago, the first time it has been selected to host an SC show.

## 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 & vendor analysis for multi-user technical server technology used for HPC and HPDA (high performance data analysis). We provide 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.HyperionResearch.com](http://www.HyperionResearch.com) and [www.hpcuserforum.com](http://www.hpcuserforum.com)

---

## Copyright Notice

Copyright 2025 Hyperion Research LLC. Reproduction is forbidden unless authorized. All rights reserved. Visit [www.HyperionResearch.com](http://www.HyperionResearch.com) to learn more. Please contact 612.812.5798 and/or email [info@hyperionres.com](mailto:info@hyperionres.com) for information on reprints, additional copies, web rights, or quoting permission.