

White Paper

Sustainability and Energy Efficiency Found to be of Strategic Importance for HPC Datacenters

Sponsored by: Dell Technologies, Intel, and NVIDIA

Jaclyn Ludema and Mark Nossokoff
May 2023

HYPERION RESEARCH OPINION

There has been a recent shift in the HPC market priorities toward sustainability and energy efficiency in response to the current geopolitical climate and trends in HPC utilization. The cost of energy is on the rise worldwide following the COVID-19 pandemic, amidst financial recovery efforts, skyrocketing fuel source prices, and ongoing military conflicts.

The demand for HPC resources for AI/ML/HPDA workloads is increasing rapidly, and new CPUs/GPUs/xPUs with substantially larger power requirements than previous iterations gain popularity, while HPC sites and datacenters are experiencing the financial strain of high energy costs. In conjunction with additional factors, such as new government regulations on energy-efficient datacenters, and sustainable business practice priorities, many HPC center decision-makers have had to prioritize sustainability and energy efficiency in operations and procurement plans.

This white paper highlights key findings from recently conducted research to better understand the shift toward sustainability and energy efficiency in HPC datacenters, focusing on:

- Identifying motivations and strategies
- Uncovering impacts on future procurements
- Discovering current and future progress tracking methods for sustainability and energy efficiency goals

What is most interesting about the discussion of HPC sustainability is the numerous ways sites are addressing the challenge. Every site has different resources available to devote toward this goal, timelines for procurement and strategy implementation, and access to sustainable energy alternatives. Looking across the 93 sites surveyed in this study, trends were identified, providing insight into the decision-making process surrounding sustainability and the merit of different approaches.

STUDY DEMOGRAPHICS

Respondents that participated in this study have visibility into the operations of their organization's HPC computing environment either in the cloud or on-premises. Respondents were from APJ (16%), EMEA (25%), and North America (59%), and work in the academia (10%), government (12%), and industry (78%) sectors.

KEY FINDINGS

The following themes emerged from this research:

- Motivations for making sustainable choices
- Strategies of HPC sites being leveraged today and in the next 18 months
- Impacts of sustainability on procurement plans
- Requirements for tracking the progress of sustainability efforts

Motivations

As the number of AI/ML/HPDA workloads performed on HPC systems continues to increase, the drive to expand computing capabilities is coming up against the cost of required energy. 57% of all respondents were motivated by the idea of wanting to do more with less: expanding computing capabilities with lower power. This trend has been seen in previous studies on accelerated computing and the widespread adoption of more CPUs, GPUs, and xPUs, as these new components can be incredibly energy-consuming, and sites have had to reconsider energy usage in order to accommodate demand.

Within the global trend of greater awareness of environmental concerns, many companies and organizations have begun implementing business practices that intentionally consider environmental, social, and governance (ESG) standards. Companies and organizations with ESG goals recognize the widespread benefits of prioritizing sustainability, not just for the potential cost savings of energy efficiency. 48% of respondents cited ESG goals as a sustainability motivator.

In recent years, several countries in the EMEA and North America regions have ratified their sentiments toward sustainable technology, specifically identifying the need for sustainable datacenters in governmental regulations. HPC site decision-makers are obliged to take sustainability into consideration for compliance's sake. The impact of these new government regulations was observed in this study, being cited twice as frequently in North America and EMEA as a sustainability motivator than in APJ, where there is less government regulation surrounding sustainability and energy efficiency.

With the cost of energy skyrocketing, particularly in EMEA due to supply chain concerns of traditional energy sources, a growing motivation for sustainability prioritization is the immediate threat to operations of HPC sites. 33% of sites are encountering prohibitive energy costs.

The top motivations by region are:

- APJ: ESG Goals
- EMEA: Prohibitive Energy Costs
- North America: Wanting to do more with less

The impetus to make sustainable changes is felt to a varying extent by HPC sites depending on location and sector. When asked if sustainability goals will impact HPC budgets, respondents from APJ (80%) and EMEA (78%) were more likely to say "yes" than respondents in North America (47%). When divided into sectors, these respondents were more likely to say "yes" if in academia (78%), compared to government (55%) and industry (59%).

Impact of energy costs

When asked if energy costs impact operations within their HPC datacenters, 58% of respondents indicated multiple impacts. Of the respondents that indicated being affected by energy costs, 60% said energy costs will defer the purchase of new systems, 43% said energy costs will limit the capabilities of current systems, and 34% said energy costs will reduce operational hours.

Regional Variation

The impact of energy costs on the level of HPC investment and datacenter operations varies across the regions. Of the 55% of respondents that said energy costs were impacting their level of investment:

- 73% of APJ said energy costs are impacting their level of investment into new HPC systems, and 60% said energy costs are impacting operations at their datacenters.
- 65% of EMEA said energy costs are impacting their level of investment into new HPC systems, and 61% said energy costs are impacting operations at their datacenters.
- 55% of North America said energy costs are impacting their level of investment into new HPC systems, and 47% said energy costs are impacting operations at their datacenters.

Sector Variation

55% of academia, government, and industry sector respondents said energy costs are impacting their level of investment into new HPC systems. Sector responses vary regarding the impacts on datacenter operations:

- 78% of academia said energy costs are impacting operations at their datacenters.
- 27% of government said energy costs are impacting operations at their datacenters.
- 62% of industry said energy costs are impacting operations at their datacenters.

Moving to Cloud

The rise in popularity of HPC cloud services comes from a variety of factors, including the ability to outsource operational concerns such as energy costs and sourcing. With energy costs on the rise worldwide, the option to outsource energy costs to a cloud service provider (CSP) is expected to gain popularity. Paired with the global trend of new government regulations surrounding sustainability, sites are looking to CSPs to meet new energy efficiency requirements. This survey adds quantitative evidence of these cloud trends.

- 76% of respondents said energy costs influence their adoption of public cloud resources, to a great extent for 33% of respondents
- 62% said government regulation influences their adoption of public cloud resources

Strategy

A part of any good strategy is hiring the right leadership and the right team. 61% of respondents indicated that they have C-level or SVP level sustainability officers at their site. Industry and government drive this trend, with C-level or SVP level sustainability officers at 64% and 63% of sites, respectively. Sites with 6 or more HPC server systems are more likely to have a C-level or SVP level sustainability officer.


When discussing efforts to meet sustainability goals, approaches varied widely across HPC sites depending upon budget, access to renewable energy sources, procurement timelines, and time allotted for energy efficiency innovation. The dynamic pathways each HPC datacenter takes toward the shared goal of sustainability have yet to be explored in detailed comparison. To tee off this effort, this study asked respondents to select their top 2 priorities in their sustainability goals.

- 56% prioritize efficient hardware configuration (higher performance per watt)
- 42% prioritize optimized cooling and power distribution
- 34% prioritize optimized workloads (identifying efficient workloads on-premises and in the cloud)

Respondents were also asked to provide all priorities they are currently implementing and plan to implement in the next 18 months. The figure below compares these numbers side-by-side.

FIGURE 1

Priorities in Sustainability: Now vs. 18 Months



n=93

Source: Hyperion Research, 2023

All priorities that are implemented today are expected to continue 18 months from now, with smart power management/telemetry (+13%), renewable energy sourcing (+ 6%), extending equipment life (+3%), and effective equipment retirement (+3%) gaining popularity amongst priorities. All four of these priorities require significant planning, commitment, and resources to accomplish, suggesting that sites are interested in long-term solutions and plan to continue prioritizing sustainability into the future.

60% of respondents said sustainability initiatives will impact their on-premises HPC budgets. Interestingly, half of those respondents said sustainability initiatives will reduce their next on-premises purchase and the other half indicated that sustainability initiatives will increase their next on-premises purchase. Plans to reduce on-premises HPC budgets and save money for sustainability initiatives indicate that HPC sites are devoting resources and time to optimization efforts (i.e., reconfiguring hardware, streamlining code for energy efficiency, extending equipment life), or changing where HPC workloads are run. The other on-premises approach, increasing on-premises HPC budgets and using the increased budget towards sustainability efforts, suggests the purchases of new energy-efficient hardware or cooling systems (liquid cooling).

60% of respondents said sustainability initiatives will impact their on-premises HPC budgets

Hiring Practices

A quarter (25%) of respondents reported that sustainability practices impact their ability to recruit employees. That number jumps to 53% in APJ and 26% in EMEA, while North America trails at 16%. The impacts on recruiting vary widely. Some respondents indicate that a lack of awareness and knowledge base of sustainable coding practices makes it difficult to find the right person for the job. Sites that value sustainability are looking for people who can optimize code for energy efficiency, know how to extend the life of current hardware, and have an awareness of eco-friendly business practices. Other respondents remarked on the allure of showcasing environmental stewardship and energy efficiency, noting that employees are attracted to companies that focus on sustainability as part of their business practices.

Future Procurements

Whether addressing sustainability immediately or farther down the road, almost all (96%) of respondents to this study had insight into their datacenter's sustainability plans. For 9%, sustainability plans did not involve new equipment procurement, suggesting these sites are working with what they have now and identifying areas of efficiency. The majority (87%), however, said new equipment procurement will play a role in their sustainability plans, but timelines for equipment procurement timelines vary across sites.

87% said new equipment procurement will play a role in their sustainability plans

Procuring more sustainable equipment is expected to be a part of regularly scheduled procurement plans for 48% of sites, while 20% of sites will expedite procurement plans and 18% will delay procurement plans for sustainable solutions. The table below shows the breakout percentages of the different procurement timeline trends across the three different regions in this study.

TABLE 1

Changes in Procurement Timelines Due to Sustainability Plans by Region

	APJ	EMEA	North America
A more sustainable solution will happen as a part of our regularly scheduled procurement plan	20%	52%	55%
Yes, a more sustainable solution is expediting our next procurement plan(s)	27%	17%	20%
Yes, a more sustainable solution procurement is delaying the timeline for our next procurement	33%	26%	11%
No, our sustainability goals do not involve new equipment procurement(s)	13%	4%	9%
No plans to pursue greater sustainability at our site	7%	0%	5%

n = 93

Q: Does sustainability affect the timeline of your future equipment procurement plans?

Source: Hyperion Research, 2023

While performance was selected as the most important factor in on-premises procurements, as is typical of most HPC procurements across the HPC market, the second most important factor selected by respondents suggests that environmental sustainability/energy efficiency is gaining importance in the procurement process. Industry leads this trend, with 37% of sites indicating that environmental sustainability/energy efficiency is their second most important factor in on-premises procurements.

The table below presents all factors in on-premises procurements considered in this survey question, and the percentage of respondents selecting these as the most important or second most important.

TABLE 2

Top Priorities for Next HPC On-premises Procurement

Priority	Most Important	Second Most Important
Performance	60%	18%
Price	16%	29%
Environmental sustainability/Energy efficiency	14%	35%
Positive brand reputation	2%	8%
Architecture (e.g., Accelerated compute nodes, CPU type)	8%	11%

n = 93

Q: Please select your top two most important priorities for your next HPC procurement

Source: Hyperion Research, 2023

Sustainability and Energy Efficiency Tracking

Smart power management (telemetry) usage is on the rise. Within the next 18 months, the number of sites implementing smart power management will increase to 40% from 27% using it today.

Nearly half (49%) of datacenters in this study said they are currently able to measure their carbon footprint, energy use or environmental impact. Sites that measure their carbon footprint, energy use or environmental impact are also able to track and report on progress towards sustainability progress, increasing the likelihood of success in sustainability efforts.


A Future Metric for Sustainability

We asked respondents to help build a metric that encompasses HPC sustainability in all its facets, based on factors that they view to be important. The most selected factors were energy use for the full datacenter, energy use efficiency ratios, and energy source (renewable energy, oil, gas). When considering that the most used metrics for HPC sustainability are energy use efficiency ratios, the requests to add to energy use of the full datacenter and energy sources to an algorithm suggests a need for a new approach. Moreover, 72% of respondents selected 3 or more factors they would want to be included in a sustainability metric, pointing to a need for a more holistic approach to tracking sustainability.

The figure below shows all variables considered by respondents to be worthwhile to include in a sustainability metric.

FIGURE 2

Factors that Should be Included in a Sustainability Metric



n = 93

Source: Hyperion Research, 2023

FUTURE OUTLOOK

With the expanding use of AI/ML/HPDA across the HPC market and the continued need for accelerated computing to enhance modeling/simulation with more accurate physics models and support training or large language models, the need to address sustainability and how energy-intensive innovation can continue into the future needs to be a part of HPC planning going forward.

Users should continue to choose HPC sites that prioritize sustainability. HPC sites are responsive to user preference and want to be recognized for their environmentally responsible brand reputation. Particularly with users in the weather & climate, and energy sectors, choosing environmentally responsible HPC sites for their workloads can be just as important as their environmentally beneficial innovations.

Vendors should include smart power management/telemetry products in their offerings to HPC sites. With 13% more sites in this study looking to implement improved telemetry in the next 18 months, the demand across the HPC market is likely to see an increase.

The world energy market continues to increase both renewable and nonrenewable energy sources in order to keep pace with demand. The shift to renewable energy sources will occur because of consumer choice, and many HPC sites are positioned to choose. Remarkably, this study reveals that 46% of sites have a say in their energy sourcing (16% off-grid/private energy supply and 30% on a local grid, but with a say in energy source prioritization), meaning the choice to use renewable energy sources lies with site management/administration.

The concepts explored in this study should be reexamined over the coming years, to identify trends throughout time and compare plans created now versus the realized outcomes. This shift in priority towards sustainability and energy efficiency has the potential to change the HPC landscape, and it will be interesting to see exactly how it evolves.

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). Hyperion Research provides thought leadership and practical guidance for users, vendors, and other members of the HPC community by focusing on key market and technology trends across government, industry, commerce, and academia.

Headquarters

365 Summit Avenue

St. Paul, MN 55102

USA

612.812.5798

www.HyperionResearch.com and www.hpcuserforum.com

Copyright Notice

Copyright 2023 Hyperion Research LLC. Reproduction is forbidden unless authorized. All rights reserved. Visit www.HyperionResearch.com to learn more. Please contact 612.812.5798 and/or email info@hyperionres.com for information on reprints, additional copies, web rights, or quoting permission.