Intersect360 Research White Paper: VERNE: COOLER HPC

EXECUTIVE SUMMARY

High Performance Computing (HPC) continues to expand, fueled by never-ending demands of scientific computing. Because of their applicability to both technical computing and machine learning, GPUs are growing increasingly common in HPC. At the same time, CPUs themselves are getting more powerful, packing increasing numbers of cores and transistors into each socket. The net result is that systems are getting denser, both physically and computationally. Today's HPC components and systems can draw substantially more power than their predecessors, taxing the datacenter. Even if a facility has the electrical capacity to run and to cool today's HPC systems, the cost to do so may be prohibitive.

Cloud computing itself is not a mechanism for eliminating power, cooling, and carbon footprint, because the same workloads are still being run, only elsewhere. Very few HPC users have become fully immersed in cloud computing, and most workloads across the industry remain on-premises. The primary reason is economics. Because HPC resources tend to be both specialized and fully utilized, there is an economic advantage to dedicated infrastructure.

At the confluence of these trends—the desire to pursue greener HPC and to reduce facilities costs, while still maintaining control of data and consistency of environments—an increasing number of HPC users are pursuing an intermediate option: facilities outsourcing. Based in Iceland, Verne is able to capitalize on these trends. With high-bandwidth connectivity to both mainland Europe and North America and year-round access to cool temperatures and free, renewable power, Verne addresses the problem of running large-scale HPC/AI installations with lower power. The entire island country operates on green energy, a combination of abundant geothermal and hydroelectric power.

With its carbon-free, high-density datacenters and high-bandwidth, low-latency connections to the U.S. and Europe, Verne has delivered consistent customer successes for HPC-using organizations. These examples highlight the role Verne plays at the confluence of major trends in HPC. Organizations seeking facilities cost reduction and greener IT can leverage low-latency connectivity to Verne's high-density Dell Technologies HPC and AI infrastructure, without sacrificing performance or control.

With its unique location and its strategic engineering investments, Verne is uniquely positioned to serve green HPC to a growing number of customers.

White paper sponsored by Verne, in conjunction with Dell Technologies.
Full neutrality statement at https://www.intersect360.com/reports/neutrality-statement/.
P.O. Box 60296 | Sunnyvale, CA 94088 | Tel. [888] 256-0124 | www.intersect360.com |
info@Intersect360.com

MARKET DYNAMICS

More Computing Power; Less Electrical Power

Driven by the inexhaustible demands of scientific computing, the High Performance Computing (HPC) market continues to expand. Today, the need for scalable, high-performance architectures is driven by not only the traditional technical computing of science and engineering, but also expanding use cases in enterprise business, including rafts of new investment in high-performance analytics and artificial intelligence (AI).

As use cases have expanded, the types of computing architectures have expanded with them, including specialized processing elements, such as GPUs, that may excel at specific tasks, like machine learning. Because of their applicability to both technical computing and machine learning, GPUs are growing increasingly common in HPC.

In the most recent Intersect360 Research surveys, 74% of respondents said they have accelerators as part of their HPC/AI environments; 95% of these are GPUs. And while deployments with single or dual GPUs per node are still common, the most prevalent configurations now have four GPUs per node, and 15% of those with GPUs now have "broad usage" of nodes with more than four GPUs. [See chart.]

GPUs per HPC/AI server node: "Broad usage"
Intersect360 Research, HPC/AI Technology Survey, 2022

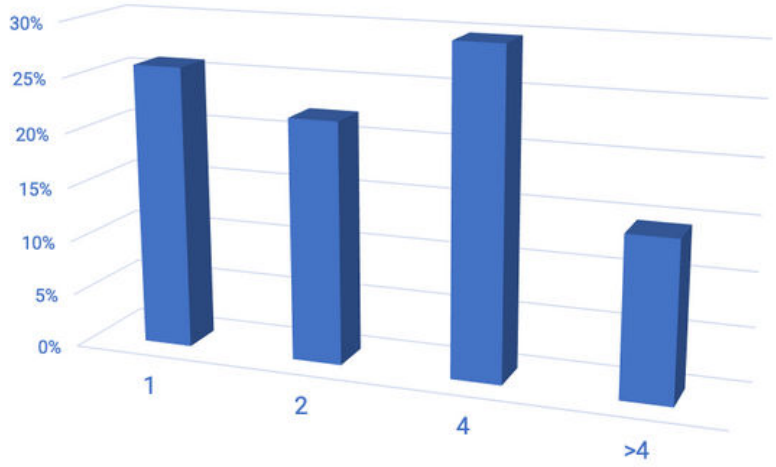

Given these trends, it is natural that we should see the influence of analytics and AI in HPC configurations, and we do. According to Intersect360 Research surveys, 78% of HPC users have engaged with machine learning, and 37% have integrated machine learning into their HPC applications for steering or simulation optimization.[1]

Today's HPC components and systems can draw substantially more power than their predecessors, taxing the datacenter. Even if a facility has the electrical capacity to run and to cool today's HPC systems, the cost to do so may be prohibitive.

[1] Intersect360 Research, HPC/AI Software survey, 2022. 78% engagement with machine learning omits "not sure" responses.

P.O. Box 60296 | Sunnyvale, CA 94088 | Tel. [888] 256-0124 | www.Intersect360.com |
info@Intersect360.com

At the same time, CPUs themselves are getting more powerful, packing increasing numbers of cores and transistors into each socket. The net result is that systems are getting denser, both physically and computationally. Today's HPC components and systems can draw substantially more power than their predecessors, taxing the datacenter.

Even if a facility has the electrical capacity to run and to cool today's HPC systems, the cost to do so may be prohibitive. The cost of the data center facilities—power, cooling, building costs, etc.—represented 16% of users' overall HPC budgets in 2021, when these costs are considered part of the budget. These costs are often hidden, as more than half of users say that facilities costs are not part of the HPC budget.[2] This doesn't mean they are free. At the organizational level, these are still significant expenses.

Cloud to the Rescue?

Cloud computing has been one of the dominant topics in enterprise computing for the past two decades, allowing companies to outsource datacenter capabilities to specialists that can operate at economies of scale. And HPC has been subject to this trend: 71% of HPC users make at least some use of cloud computing,[3] usually for occasional peak needs beyond what their own in-house datacenters can provide.

However, very few HPC users have become fully immersed in cloud computing, and most workloads across the industry remain on-premises. At the individual application level, users report the application is run entirely on-premises 84% of the time, a mix of on-prem and cloud 15% of the time, and exclusively in the cloud only 1% of the time. [See chart.] The primary reason is economics. Because HPC resources tend to be both specialized and fully utilized, there is an economic advantage to dedicated infrastructure.

It is also dubious to cite cloud computing itself as a mechanism for eliminating power, cooling, and carbon footprint, because the same workloads are still being run, only elsewhere. Imagine your friends who allege they are saving water in their drought-affected area by showering at the gym. This makes little sense; they have only moved the water usage, not eliminated it. The argument has merit only in the corner cases in which either the gym is far enough away not to be affected by the same drought, or it is somehow using water more efficiently, such as by purifying reclaimed water. So it is with remote computing: the overall power consumption problem is best alleviated if it is not subject to the same local restrictions and if it generates power or cools the computing infrastructure more efficiently.

[2] Intersect360 Research, HPC/AI Budget Map survey, 2022.
[3] Intersect360 Research, HPC/AI Technology survey, 2022.

HPC Application Affinity to Cloud Computing
Proportion of individual applications run on-premises to in cloud, on seven-point scale
Intersect360 Research, HPC/AI Software Survey, 2022

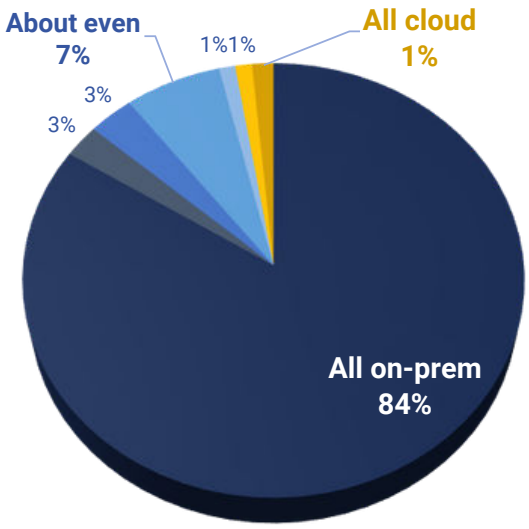

Because HPC resources tend to be both specialized and fully utilized, there is an economic advantage to dedicated infrastructure.

Economics aren't the only argument against cloud computing for HPC, particularly in commercial markets. As stated in Intersect360 Research reports, "HPC can frequently be distinguished from general business computing in that companies generally will use HPC applications to gain advantage in their core endeavors—e.g., finding oil, designing automobile parts, or protecting clients' investments—as opposed to non-core endeavors such as payroll management or resource planning." As such, organizations can have concerns, or even corporate restrictions, preventing the exportation of data to system's not owned and managed by the organization. Data sovereignty—including ownership, possession, monitoring, control, and stewardship—is a topic of growing scrutiny for many organizations, and it prevents the adoption of cloud for many HPC applications.

Control of HPC resources has another benefit for HPC users: reproducibility of results. Cloud computing can create inconsistency in HPC environments, and subtle variations in hardware and software configurations have the potential to change the results of simulations. When comparing runs over time, organizations want to know that answers vary depending on the model and its inputs, rather than the underlying configuration. Owning and controlling an HPC environment creates consistency and confidence.

Furthermore, there is efficiency in architecting solutions specifically for the high-density computation of HPC and AI. Although certain high-performance instances are increasingly available, most cloud resources are not built around HPC. When a datacenter is built for HPC and AI workloads, it can incorporate more efficient power and cooling strategies for high-density configurations.

Facilities Outsourcing: Off-Premises Ownership

At the confluence of these trends—the desire to pursue greener HPC and to reduce facilities costs, while still maintaining control of data and consistency of environments—an increasing number of HPC users are pursuing an intermediate option: facilities outsourcing. By renting datacenter space from a company that specializes in hosting high-density computing, organizations can reduce their facilities costs, particularly when moving away from locations where space or power are costly. If the hosting company is optimally located, carbon footprint can be reduced, while power bills are lowered.

Intersect360 Research identified the trend toward facilities outsourcing for HPC in recent years and officially began tracking it as a spending category in 2021. As more organizations emphasize green IT, we expect facilities outsourcing to grow as a strategy for reducing both carbon and cost while maintaining control of HPC.

INTERSECT360 RESEARCH ANALYSIS

Verne: Cooler HPC

Based in Iceland, Verne is able to capitalize on these trends. With low-latency connectivity to both mainland Europe and North America and year-round access to cool temperatures and free, renewable power, Verne aims to solve the problem of running large-scale HPC/AI installations with more efficient power consumption.

Iceland offers unique advantages for hosting high-power datacenters. The entire island country operates on green energy, a combination of inexhaustible geothermal and hydroelectric power. The same geology that provides tourist-friendly hot springs is an endless source of zero-carbon, renewable energy for datacenters. Furthermore, Iceland offers a year-round climate that is consistent and cool without becoming hostilely frigid. Environments that are too cold provide challenges for facility operations, increasing the overall cost to support important systems that are exposed to the elements, such as cooling equipment. The prevailing conditions in Iceland therefore create an opportunity to host HPC datacenters that are truly green—not just "greenwashing" by moving the problem, but actually reducing the carbon footprint of HPC resources.

Of course, not everything comes for free. HPC datacenters need to be engineered to take advantage of power and cooling to provide green, high-performance resources. Verne offers state-of-the-art datacenter hosting services beginning at a single rack.

Racks offer a range of cooling options and compute densities. Racks using CFD-optimized free air cooling go up to 50 kilowatts; liquid-cooled racks can support over 100 kilowatts of power in a single cabinet. Individual racks can be combined into "data halls" and "ultra high density

With low-latency connectivity to both mainland Europe and North America and year-round access to cool temperatures and free, renewable power, Verne aims to solve the problem of running large-scale HPC/AI installations with lower power.

PODs" for multiple megawatt requirements,[4] supporting highly scalable applications. Verne facilities have full fault tolerance and resilience, with 100% uptime guaranteed.

Extending In-House to Off-Premises

The decision to-cloud-or-not-to-cloud actually contains two distinct decisions that can be teased apart: ownership and location. With public cloud services, both are outsourced; with on-premises, in-house IT, both are in-sourced. Facilities outsourcing services, like those provided by Verne, allow for the economic and data sovereignty benefits of HPC ownership while incorporating the cost and carbon benefits of colocation. As such, Verne–hosted resources act as an extension of in-house operations. It is still a corporate datacenter, just a remote, carbon-free one.

For this to be effective, efficient data movement is critical, and here Verne benefits from Iceland's location and its strategic investments in data connectivity. With major fiber trunks in place, Verne data centers offer high-bandwidth connections to both the U.S. and Europe. Average one-way latency to New York is 40.6 milliseconds, to Paris 21.8 milliseconds, and to London 18.9 milliseconds. [See chart.]

Average One-Way Latencies to Selected Cities from Iceland
Source: Verne

Du b lin	10.0 ms
Copenhagen	14.9 ms
Frankfurt	17.5 ms
Amst er d a m	17.8 ms
London	18.9 ms
Paris	21.8 ms
Halifax	33.7 ms
New York	40.6 ms

This level of connectivity allows organizations to use their Verne –hosted resources either completely independently or as hybrid extensions of in-house IT. This mirrors how most companies use cloud computing resources today, bursting to accommodate peak workloads or to access specialized resources, such as for machine learning training. And for those who do leverage public cloud for its bursting capabilities, Verne offers connectivity as well, with high-bandwidth links to Amazon Web Services, Google Cloud Platform, IBM Cloud, Microsoft Azure, and Oracle Cloud Infrastructure.

[4] https://verneglobal.com/

The connectivity to both in-house and public cloud lets organizations keep access scalable computing wherever their respective data sets are. When adopting any technology that is dependent on getting more value out of data at scale, companies need to think about architecture options, including on-premises, public cloud, and facilities outsourcing providers like Verne.

The net result is that Verne offers many of the advantages of cloud computing without the liabilities. Organizations maintain data sovereignty, consistent performance, and the economic benefits of dedicated systems, while also reducing facilities burdens and reducing carbon emissions to zero.

Dell Technologies for HPC

Today's HPC environments need to combine multiple technology elements to merge scientific computing with large-scale data analytics and artificial intelligence. The wide range of high- performance workloads requires HPC solutions that are flexibly capable of multiple approaches.

With trusted products across both computation and data management, Dell Technologies is the industry leader in total HPC solution revenue[5] [See chart], and Dell was once again the most-cited server vendor in an Intersect360 Research survey of HPC users in 2022.[6] Verne is a Dell Technologies Titanium Partner, delivering best-of-class Dell HPC and AI as-a-Service offerings.

Combined HPC Server and Storage Revenue Share, 2020
Intersect360 Research, 2021

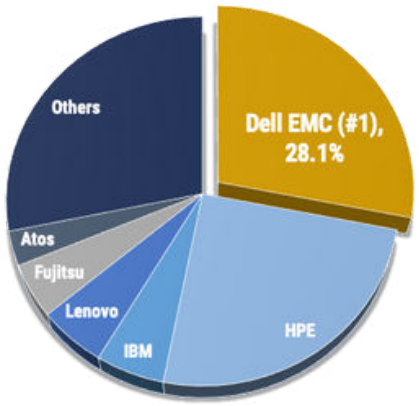

Verne is a Dell Technologies Titanium Partner, delivering best-of-class Dell HPC and AI as-a-Service offerings.

[5] 2020 HPC market share by revenue. Intersect360 Research HPC market model and forecast data,
[6] 2021. Intersect360 Research, HPC Technology Survey, 2022.

P.O. Box 60296 | Sunnyvale, CA 94088 | Tel. [888] 256-0124 | www.Intersect360.com |
info@Intersect360.com

Customer Successes for HPC

With its carbon-free, high-density datacenters and high-bandwidth, low-latency connections to the U.S. and Europe, Verne has delivered consistent customer successes for HPC- using organizations. BMW, for example, has leveraged Verne for crash testing and computational fluid dynamics (CFD) simulations for its i-series cars, with an 82% savings in energy costs.[7]

Wirth Research, a CFD consultancy, uses Verne to deliver innovations for its customers, moving both CFD workloads and its own internal virtual desktop infrastructure (VDI) to Verne . According to President and Technical Director Nick Wirth, "The cost of annual service, power, and connectivity from Verne is equivalent to about three months of our previous electricity bill alone. The fact that our research is now powered totally by renewable energy is the cherry on top." [8]

These examples highlight the role Verne plays at the confluence of major trends in HPC. Organizations seeking facilities cost reduction and greener IT can leverage low-latency connectivity to Verne's high-density Dell Technologies HPC and AI infrastructure, without sacrificing performance or control. With its unique location and its strategic engineering investments, Verne is uniquely positioned to serve green HPC to a growing number of customers.

For more information about Verne, visit https://verneglobal.com/.

"The cost of annual service, power, and connectivity from Verne is equivalent to about three months of our previous electricity bill alone. The fact that our research is now powered totally by renewable energy is the cherry on top."

–Nick Wirth, Wirth Research

VERNE

DELL Technologies
TITANIUM PARTNER

[7] https://www.forbes.com/sites/peterdetwiler/2014/12/02/verne-global-iceland-will-become-a-mecca-for-data-processing

[8] https://verneglobal.com/app/uploads/2021/10/Verne-Global_WirthResearchCaseStudy.pdf

P.O. Box 60296 | Sunnyvale, CA 94088 | Tel. [888] 256-0124 | www.Intersect360.com |
info@Intersect360.com