Intersect360 Research White Paper: HIGH-PERFORMANCE ANALYTICS DRIVES DELL CUSTOMER SUCCESS

MARKET DYNAMICS

Analytics and the New HPC

In supercomputing centers around the world, High Performance Computing (HPC) is used to unlock the secrets of the universe, from the dynamics of black holes to the formation of tornadoes, to understand how viruses bond to host cells and the long-term effects of carbon dioxide in the atmosphere. The science of these questions is complex, and HPC offers the massive computational power to understand them.

Considering the vast spheres of research, it would be easy to assume that government and academic research labs must be the dominant users of HPC, but this isn't the case. Commercial users account for more than half of HPC spending each year. This is due, in large part, to the many commercial applications for science and engineering, in industries such as manufacturing, oil exploration, or pharmaceuticals.

Beyond traditional scientific computing, HPC has seen adoption across a wider range of enterprise computing, due to the expanding role of analytics. Analytics isn't a brand-new science; it has existed as long as there has been data to analyze. However, the major data boom from successive trends in connectivity, cloud computing, and internet-connected devices has created troves of enterprise data beyond anything that existed previously. Analytics saw a boom in the "Big Data" era of the early 2010s, and it is only accelerating with the onslaught of machine learning and artificial intelligence (AI). Bolstered by these trends, worldwide spending on HPC is poised to exceed $60 billion worldwide in 2025.[1]

The machine learning algorithms behind AI can be deployed any time there is a wealth of data to draw on, coupled with a reward from making more intelligent inferences based on that data. Most domains within science and engineering fit this pattern, and beyond that, it is the perfect description of the world of analytics. HPC-using organizations in general are trending toward the incorporation of AI. In a 2021 survey, 81% of HPC users said they are already running machine learning as part of their environments or planning to implement machine learning within a year.[2]

[1] Intersect360 Research, *Worldwide HPC Market Model and Forecast*, 2021.

[2] Intersect360 Research, HPC User Budget Map survey data, 2021.

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

Powering Analytics with HPC and AI

HPC has been powering analytics for decades. "Data mining" was a catchphrase of the late 1990s, for emerging applications like fraud detection in insurance claims or credit card purchases, as the internet and digital transactions began to boom. Looking at patterns in existing data and marking known instances of fraud, data miners could use advanced HPC and visualization techniques to find other cases to investigate. Once patterns were determined, suspicious transactions could be flagged at initiation. These processes may not have used AI at the time, but the pattern of training and inference—using human guidance and intervention—is evident.

The field of analytics has continued to boom, both with the availability of data and with the power of the tools to search for insights. Visualization is still a critical tool for human examination of patterns in data. In a recent survey, 43% of HPC users cited the use of advanced visualization tools, and 31% of commercial HPC users were deploying HPC for business analytics applications.[3]

AI opens yet another door to the field of analytics by automating pattern recognition that used to require a human. Traditionally, most HPC applications have been *deterministic*; given a set of inputs, the computer program performs calculations to determine an answer. Machine learning represents another type of applications that is *experiential*; the application makes predictions about new or current data based on patterns seen in the past. We call this "artificial intelligence" because it approximates how we learn as humans. After all, a child does not become adept at catching a ball by learning advanced physics and mathematics, but through repetition and experience, learning how the ball behaves in most situations, and applying past knowledge to new situations. Both approaches have a future in HPC, and particularly in analytics.

New Dimensions of Scalability

The drive to incorporate new workloads has expanded the traditional notion of a fundamental HPC concept: scalability. Traditionally, scalability in HPC has essentially meant *more*. More cores, more flops, more memory, more bandwidth, more throughput: if you can make a bigger computer, that's scalability.

As high-performance workloads continue to expand, today's HPC environments are uncovering new dimensions of scalability. To meet the specific demands of various applications, heterogeneous computing is becoming commonplace. 91% of users incorporate accelerators into some portion of their HPC environments. HPC sites often find they need to incorporate a wide range of processing capabilities to serve expanding workloads.

Furthermore, no discussion of scalable computing would be complete without mention of cloud computing, which has rapidly become a viable alternative to on-premises

The field of analytics has continued to boom, both with the availability of data and with the power of the tools to search for insights.

[3] Intersect360 Research, HPC Software survey data, 2022.

infrastructure. Our most recent surveys indicate that over 60% of academic HPC sites make use of cloud computing, but typically for only a small portion of their overall HPC workloads. Most HPC sites find it advantageous to run on-prem when they can, for reasons of cost and data locality, bursting to cloud only for extra scalability, capability, or capacity when needed. These hybrid-cloud HPC environments bring new scalability challenges, such as managing data sovereignty, availability, and optimization across the total ecosystem. And in Intersect360 Research surveys, users said they used cloud resources for business analytics more than any other application domain.

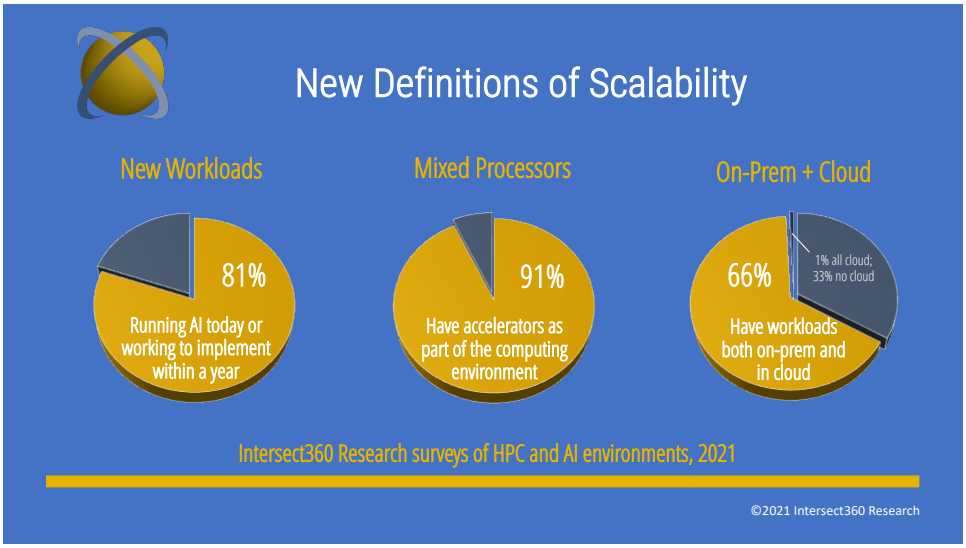

INTERSECT360 RESEARCH ANALYSIS

Today's HPC environments need to scale in new ways, combining 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. Successful solution providers will be those that have technologies that can span this expanded environment, combined with domain-specific expertise in helping organizations achieve meaningful breakthroughs.

This is the role of Dell Technologies. With trusted products across both computation and data management, Dell Technologies is the industry leader in total HPC solution revenue[4] (See chart below), and Dell was the most-cited vendor in a survey of HPC users in 2021.[5] Dell Technologies leverages this breadth of offerings with converged solutions that incorporate HPC, data analytics, and AI[6] and offers tailored solutions incorporating the latest in AI for a wide range of scientific domains.[7]

Dell Technologies is the industry leader in total HPC solution revenue, and Dell was the most-cited vendor in a survey of HPC users in 2021.

[4] 2020 HPC market share by revenue. Intersect360 Research HPC market model and forecast data, 2021.

[5] Intersect360 Research, HPC Technology Survey, http://www.intersect360.com/LiteratureRetrieve.aspx?ID=159114

[6] https://www.dellemc.com/en-us/collaterals/unauth/brochures/solutions/hpc_ai_convergence_brochure.pdf.

[7] https://www.delltechnologies.com/en-us/solutions/data-analytics/machine-learning/ready-solutions-for-ai.htm.

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

Dell Technologies provides systems, solutions, and support for AI-augmented HPC, including servers incorporating the newest processing technologies from AMD. With its third-generation AMD EPYC™ processors and a steady drumbeat of HPC application benchmarks, AMD has raced back to the forefront of many HPC users' minds. AMD EPYC processors are now in use at 70% of surveyed HPC sites, and in broad usage at 23%, representing a tremendous gain from previous years. (See chart below.)

AMD processors are now in use at 70% of surveyed HPC sites, and in broad usage at 23%, representing a tremendous gain from previous years.

Combined HPC Server and Storage Revenue Share, 2020

Intersect360 Research, 2021

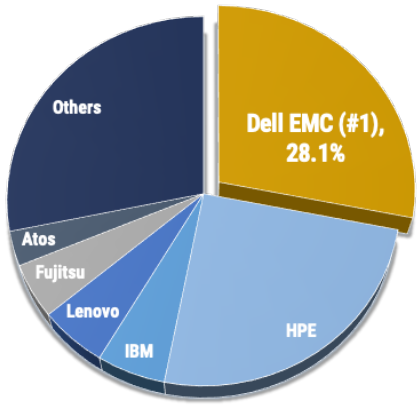

Percentage of Users Using AMD CPUs as Part of HPC Environment [8]

Intersect360 Research, 2021

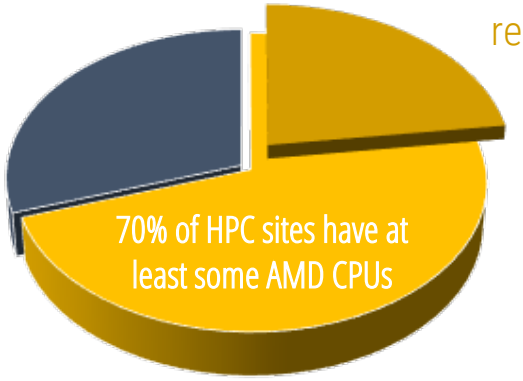

23% of HPC sites report "broad usage" of AMD CPUs

[8] Intersect360 Research HPC Technology Survey, 2021.

Furthermore, AMD is poised to become the first processor vendor to enter the HPC market with both CPUs and GPUs in an integrated package, connected by AMD Infinity Fabric™ technology. This combination promises to improve CPU-GPU communications, improving performance for machine learning and other accelerated applications. In addition, direct memory addressing can simplify programming and porting of applications.

Internally, Dell Technologies pursues advancements in HPC and AI through its HPC & AI Innovation Lab, where the company's engineers test and optimize new generations of technologies for processing, networking, and storage, including solutions based on AMD processors.[9] Globally, Dell Technologies also operates its HPC & AI Centers of Excellence, which showcase the latest solutions and provide community collaboration opportunities with the wider HPC community[10], and its Worldwide Customer Solution Centers, which offer remote access capabilities for testing and optimizing customer-specific workloads in collaboration with solution specialists.[11]

Most importantly, Dell Technologies is helping its customers derive real value from their investments in computation and machine learning, including the advanced analytics powered by HPC and AI. Dell Technologies is well-positioned to do this, based on its breadth of offerings and its actions to broaden HPC access, such as through its Dell HPC Community[12] and direct communications from its HPC engineering teams.[13] This direct-access, domain-specific support is critical in empowering innovation for those users investing in their next breakthroughs.

High-Performance Analytics with Dell and AMD

As one example of using HPC for advanced analytics, the National Health Service of Wales is using Dell HPC solutions based on AMD processors as part of the Cloud Infrastructure for Microbial Bioinformatics (CLIMB) project. According to a Dell customer case study:

> *… a 160-core HPC cluster sits next to the sequencing system to do some initial results generation. … At the other end of this first analysis run, close to the sequencing instrument, the data is shuttled to an OpenStack-based cluster that can provide other testing environments, and has backup and resilience systems in place. This way, if a lab goes down, it is still possible to get patient data without interruption, and do fast analysis.*[14]

According to Dr. Tom Connor, Bioinformatics Lead for the Pathogen Genomics Unit at Public Health Wales, "We do our genome sequencing, we put our fragments back together, and

---

[9] http://delltechnologies.com/innovationlab.

[10] http://delltechnologies.com/coe.

[11] http://delltechnologies.com/csc.

[12] http://www.dellhpc.org.

[13] http://www.hpcatdell.com.

[14] Dell Technologies, "A New Sequence for Bioinformatics HPC," https://www.delltechnologies.com/asset/en-us/products/ready-solutions/customer-stories-case-studies/dell-ukhealth-cardiff-case-study.pdf.

then we run through an analysis process which searches that blueprint to identify which drugs are going to work."[15]

Dell and AMD are also behind the new "Anvil" supercomputer at Purdue University, "a powerful new supercomputer that will provide advanced computing capabilities to support a wide range of computational and data-intensive research spanning from traditional high-performance computing to modern artificial intelligence applications," according to Purdue.[16] Funded by a grant from the National Science Foundation (NSF), Anvil will be part of the Extreme Science and Engineering Discovery Environment (XSEDE), furthering scientific discover at Purdue and across the wider research community.

In one more example, the High Performance Computing Center at Texas Tech University recently presented on its work on advanced data visualization techniques to improve data management and analytics, as part of the NSF-funded Cloud and Autonomic Computing (CAC) industry-university cooperative research center. This project uses analytics and scientific visualization to understand the efficiency and behavior of the systems themselves, presenting "cutting-edge visualization, monitoring and management solutions for HPC systems to understand the status of high-performance computing platforms." TTU's newest, most powerful HPC system, "Nocona," is central to this initiative. Nocona has 240 Dell PowerEdge C6525 nodes powered by AMD EPYC 7702 processors.[17]

The common question with analytics is, "What can we do with all this data?" The answer sometimes isn't obvious. The new HPC is fueled by analytics and AI. Across domains, there is a deep wealth of data. If harnessed, it can unlock new discoveries, which themselves will spawn new areas of research. While the fundamental drivers of research remain unchanged, these converged, high-performance solutions are critical to enabling new generations of insight. With its domain-specific knowledge and technology solutions across data management and computation for AI-augmented HPC, Dell Technologies is well-positioned to help is customers achieve new insights for the next generation.

To learn more, please visit delltechnologies.com/hpc.

[15] Dell Technologies customer video, "Unleashing the power of genomics for national health," https://www.youtube.com/watch?v=jWaVtJ1aIIk.

[16] https://www.rcac.purdue.edu/anvil.

[17] Texas Tech University SC21 presentation, "Advanced Visualization and Data Analysis of HPC Cluster and User Application Behavior," https://www.youtube.com/watch?v=sdoTGiD9n-U.

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

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