

Technology Spotlight

Linux Helps Drive HPC into New Markets

Sponsored by SUSE (<https://www.suse.com>)



Steve Conway, Earl Joseph, Bob Sorensen and Alex Norton
November 2019

HYPERION RESEARCH OPINION

The arrival of the standards-based Linux cluster has been heavily responsible for expanding the worldwide market for high performance computing (HPC) server systems nearly four-fold, from \$3.8 billion in 1996 to \$13.6 billion in 2018, en route to a Hyperion Research-forecast \$19.9 billion in 2023. In the most recent list of the world's TOP500 supercomputers (June 2019), every one of these powerful systems was running the Linux operating system, and nearly all of the systems are architectural clusters, even if not designated as such.

But those 500 supercomputers represent less than one percent of the HPC systems that are sold worldwide each year. The Linux cluster has democratized HPC by making it affordable for many more organizations. In 2018, 67% of the 90,980 HPC systems sold around the world, or 61,217 systems, were in the sub-\$100,000 Workgroup price band. The average price paid for one of these Workgroup HPC systems was just \$32,928 (granted, some of these purchases were simply additions to existing systems).

Beyond making HPC more affordable, the great advantage of Linux is its openness. By giving HPC users access to source code, Linux naturally evolves and adapts in lockstep with the global HPC community. That's a good thing, because during Linux's lifetime the global HPC community has undergone a great deal of evolution and adaptation to support new requirements and new environments.

At least as early as the 1990s, Linux clusters began migrating into the back offices of large investment banks to help quantitative analysts ("quants") price complex offerings, optimize portfolios and perform firmwide risk analysis. More recently, Linux has propelled HPC systems into the emerging era of digital transformation and AI, both on premises and in third-party cloud environments. A growing contingent of enterprise data centers have become first-time HPC adopters to accelerate business operations ranging from transaction processing to sales planning, supply chain and customer relationship management, affinity marketing, human resources, cyber security and fraud detection, to name just a few uses.

Competitive forces, especially the need to aim more complex questions at data structures and obtain answers in near-real time, are pushing enterprises up into the HPC competency space to run these challenging, data-intensive workloads. More enterprises are realizing that to out-compete rivals today, they have to out-compute them. Linux-based HPC is making that happen.

Vendor-supported Linux versions are especially important for businesses because companies typically must maintain high productivity and want to avoid hiring Linux experts themselves. This paper focuses on Linux as an important vehicle for driving HPC into new and challenging markets and applications, using SUSE to illustrate this major trend. SUSE Linux Enterprise is designed to meet the production computing requirements of organizations in government, industry and academia.

Note: this page is intentionally blank.

HETEROGENEOUS HPC DRIVES INTO NEW MARKETS

HPC arose in the 1960s to support government and academic research. HPC began moving into major industries in the 1970s, especially to accelerate the development of complex products (automotive, aerospace, oil and gas, financial services and pharmaceutical firms).

HPC has carved out a strong position in industrial markets. In 2018, nearly half (\$6.4 billion, or 47%) of \$13.6 billion in worldwide HPC server system revenues came from the private sector. Private sector spending on the larger HPC ecosystem (servers, storage, software and technical support) approximately doubled the figure for server systems alone.

FIGURE 1

2018 Worldwide HPC Systems Revenue by Application (\$000)

Bio-Sciences	1,245,865
CAE	1,521,850
Chemical Engineering	205,891
DCC & Distribution	780,184
Economics/Financial	746,418
EDA / IT / ISV	984,887
Geosciences	1,029,041
Mechanical Design	63,137
Defense	1,403,164
Government Lab	2,616,822
University/Academic	2,420,440
Weather	560,631
Other	127,757
Total Revenue	13,706,088

Source: Hyperion Research, 2019

In recent years, HPC has been moving into important new markets, new use cases and new environments.

- **New markets:** Enterprise IT. Historically, private-sector HPC has been used mainly for advanced research, especially upstream R&D to help design new products. Now, a growing contingent of large companies are integrating HPC systems into their enterprise data centers to accelerate live business operations ranging from human resources and supply chain management to sales analysis and revenue protection. Competitive forces are pushing enterprise data center requirements up into the HPC competency space, especially the need to aim more complex questions at data structures and get answers in near-real time. Firmwide analyses that used to

be done once a month may now need to be repeated multiple times a day, with more details and insight than ever before. Hyperion Research expects the number of enterprises exploiting transformational HPC technology to grow at a healthy rate.

- **New use cases:** Studies by Hyperion Research and others show that HPC is nearly indispensable at the forefront of R&D for emerging AI use cases, including automated driving, precision medicine, cyber security, affinity marketing, smart city development and the Internet of Things. Figure 2 shows Hyperion Research's 2018 market sizing and robust five-year forecast for high performance data analysis (HPDA) servers and the AI subset of the HPDA server market.

FIGURE 2

HPC Systems Forecast for High Performance Data Analysis and HPC-Enabled AI

Forecast: Worldwide HPC-Based AI Revenues vs Total HPDA Revenues (\$ Millions)

	2018	2019	2020	2021	2022	2023	CAGR 18-23
WW HPC Server Revenues	13,706	14,495	15,780	17,376	18,983	19,947	7.8%
Total WW HPDA Server Revenues	3,153	3,598	3,932	4,737	5,467	6,450	15.4%
Total HPC-Based AI (ML, DL, and Other)	747	938	1,094	1,399	1,810	2,725	29.5%

Source: Hyperion Research 2019

Table 2

Forecast: Worldwide ML, DL & Other AI HPC-Based Revenues (\$ Millions)

	2018	2019	2020	2021	2022	2023	CAGR 18-23
ML in HPC	532	675	875	1130	1479	1940	29.5%
DL in HPC	177	216	301	392	510	665	30.3%
Other AI in HPC	38	47	66	80	95	120	25.9%
Total	747	938	1,242	1,602	2,084	2,725	29.5%

Source: Hyperion Research, 2019

- **New environments:** in addition to enterprise data centers, another new environment for HPC use is public, private and hybrid clouds. Recent Hyperion Research studies indicate that about 20% of all HPC workflows are now run in third-party clouds offered by cloud services providers (CSPs) or HPC system vendors. Hyperion Research forecasts that HPC spending for third-party cloud usage will grow beyond \$6 billion in 2023.

This trend toward HPC's increased heterogeneity is nicely described in the U.S. Department of Energy's January 2018 report, *Productive Computational Science in the Era of Extreme Heterogeneity*.

Hyperion Research has long said that software advances will be even more important than hardware progress in determining the future of HPC, especially system software. The HPC market has grown historically by attracting new waves of users, each of them less technically proficient than its predecessors. System software has consistently responded to this challenge by making HPC systems easier to manage and use, enabling HPC to drive into new market segments.

At the heart of system software is the operating system, and since its initial release in 1991 Linux has become the dominant operating system of the worldwide HPC community. Every one of the HPC systems on the June 2019 list of the world's 500 top supercomputers runs a version of Linux, and even the smallest entry-level HPC clusters typically are shipped with Linux. That's largely because HPC users have continuously advanced open-source Linux to address the community's increasingly challenging, heterogeneous requirements. Linux moved into industrial markets early on for modeling and simulation. As early as the 1980s, Linux clusters also began showing up in the back offices of large investment firms to support advanced financial analytics, an important precursor to today's AI/ML/DL workflows.

SUSE HPC SOLUTIONS

SUSE is priming its HPC solutions for adoption everywhere, including traditional supercomputing environments, small/medium enterprises, public/hybrid clouds (either all-in or bursting), and scalable GPU environments.

There is a strong connection between the rise of edge computing and having the high-performance computing infrastructure to support it. IoT and edge use cases include everything from autonomous vehicles crammed with sensors and compute devices, to farming and factory automation, smart cities and spaces, energy and utilities, robotics and artificial intelligence, retail and business intelligence, healthcare and pharmaceuticals, and much more. As more endpoints are connected to the IoT, enabling edge computing strategies with advanced networking technologies, an HPC platform available both on-premise and in the cloud with advanced tooling and management and software defined HPC storage, becomes vital. In many cases, data will need to be stored and processed nearer to where it's captured or generated. This will need to be balanced with the volume of data that needs to be transmitted back to centralized or consolidated data centers. Hybrid and multi-cloud strategies combined with robust security measures will become increasingly important.

Open source solutions will continue to take a leading role in almost everything from edge to core to cloud. This includes a long list of things such as sensors, smart devices, data storage, HPC and analytics, public/private/hybrid and multi-cloud implementations, automation and tooling, and more. As enterprise customers move beyond the software-defined data center to also embrace edge and cloud computing, SUSE is evolving with the market, delivering new hybrid and multi-cloud capabilities and application delivery innovations that aim to help customers transform their digital infrastructures in their own way, as quickly as they need to.

SUSE Linux Enterprise Server for ARM has enhanced support for edge workloads, supporting many different system-on-a-chip (SoC) processor options. This broadens support for storage and industrial automation applications on 64-bit ARM server and IoT devices. For 64-bit Raspberry Pi devices, it now supports full HDMI audio and video and provides an ISO image for faster installation. SUSE Linux Enterprise HPC includes many popular tools for managing and monitoring data-intensive workloads and supporting HPC applications such as AI and machine learning, complex simulations, and advanced analytics at the edge, in the core and in the cloud.

HPC is emerging into the mainstream, with commercial organizations trying to harness the power of HPC to solve their business issues. SUSE continues to invest in HPC solutions in support of those initiatives around edge computing, AI/ML and advanced analytics. SUSE HPC solutions encompass SLE HPC, with HPC tools such as Slurm for workload management and Ganglia for performance monitoring; SUSE Enterprise Storage for Ceph-based HPC storage in both primary (ultra-fast HPC) and secondary (archival) tiers; and cloud-ready HPC images via Azure and AWS marketplaces.

In summary, SUSE investments range from expanding the company's HPC OS platform and tools to make it easier for enterprises to adopt HPC in-house in support of a new generation of applications, to software-defined storage based on Ceph that is designed to handle the volumes of data being generated in high performance environments, to creating HPC environments in the cloud that aim to provide ultra-high scalability and additional resources on-demand. To address the rise of AI/ML and the containerization of AI. SUSE's HPC software infrastructure and tools are designed to provide an in-house HPC environment in a wide range of industries:

- SUSE Linux Enterprise High Performance Computing is SUSE's OS that includes the HPC module, a bundle of popular HPC tools like Slurm and Ganglia that are supported by SUSE
- SUSE Enterprise Storage is Ceph-based and now is ranked in the IO500 HPC storage list, qualifying it as a viable option for primary and secondary HPC storage.
- SLE HPC is now available on the public cloud marketplaces such as Microsoft Azure and AWS.
- SUSE Services for HPC are designed to provide outstanding technical and professional support.

SUSE also boasts a strong partner ecosystem, including independent hardware vendors such as HPE/Cray, Intel, Dell and Lenovo; ISVs such as Bright Computing, Sylabs/Singularity, Univa and ANSYS; and leading CSPs, including Azure, AWS and Google. Cray relies on SUSE HPC within their Cray Linux Environment, helping to propel SUSE to market share leadership among paid-for HPC OS systems in the latest TOP500 report.

SUSE SUCCESS STORIES

Day & Zimmerman (Construction: Philadelphia, Pennsylvania)

As a major construction and engineering company with projects all around the world, Day & Zimmermann relies on SAP applications to ensure its 46,000 employees and contractors in more than 150 locations around the world get paid on time. To keep its core SAP HANA databases online 24/7, the company built a highly available cluster using SUSE Linux Enterprise Server for SAP Applications, reducing unplanned downtime to less than 30 minutes per year.

SAP recommended running the new SAP HANA system on SUSE Linux Enterprise Server for SAP Applications. SAP reports that installation of the software is fast and easy, because 90 percent of the parameters are set automatically. Day & Zimmermann's production SAP HANA environment runs on a multi-node cluster, with the SUSE High Availability Extension providing automatic, near-instant failover between primary and secondary nodes if a fault occurs. With its clustered architecture, Day & Zimmermann has seen availability rise to around 99.995 percent. Over the past four years, it has experienced just two hours of unplanned downtime.

As the company's IT landscape continues to grow, SUSE Linux Enterprise Server has become a corporate standard for all new applications – not just for SAP. For example, the company recently set up an AMD EPYC 3-node Veeam scale-out backup repository that runs on the SUSE operating system. Day & Zimmerman's use of SUSE Linux Enterprise Server promises to produce significant cost savings by reducing reliance on proprietary operating systems and licensing.

Guerbet (Healthcare/Medical: Roissy, France)

With more than 2,800 employees in the five continents, Guerbet is a global leader in medical imaging. Guerbet is headquartered in France and generated annual revenues of €790 million (USD\$872 million)

in 2018. Guerbet's employees are committed to delivering solutions that empower medical professionals around the world to diagnose and treat their patients quickly.

With operations in over 80 countries, and long and complex supply chains, the company knew that rationalizing its business processes would be vital to operate efficiently and cost-effectively as its growth gathered pace. To achieve its goals, Guerbet consolidated 21 ERP systems to create a future-ready platform based on SAP Business Suite powered by SAP HANA, running on SUSE Linux Enterprise Server for SAP Applications.

Guerbet uses SUSE Manager to automate Linux server provisioning, patching and configuration. Guerbet reports that SUSE Manager automatically generates detailed, timely and granular reports that help demonstrate to regulators that Guerbet complies with industry best practices. Since going live with SUSE Linux Enterprise Server for SAP Applications, Guerbet has never experienced unplanned downtime.

Leibniz Supercomputing Center (High Performance Computing: Munich, Germany)

A world-class academic IT service provider and supercomputing pioneer, the Leibniz Supercomputing Center (Leibniz-Rechenzentrum, or LRZ) hosts SuperMUC-NG, one of the most powerful computers in the world. The supercomputer is built with a combination of technologies from Intel and Lenovo and runs on SUSE Linux Enterprise High Performance Computing.

LRZ, located in Garching, which is near Munich, Germany, is the IT service provider for all universities in the Munich area, as well as a growing number of research organizations throughout Bavaria. LRZ also plays an important role as a member of the Gauss Centre for Supercomputing (GCS), delivering top tier HPC services on the national and European level. LRZ is an institute of the Bavarian Academy of Sciences and Humanities.

Big data could hold the key to understanding the origin of the universe, the composition of matter itself and many more huge questions that have intrigued academics for centuries. To advance the frontiers of knowledge with ground-breaking research, scientists require access to high-performance computing environments that enable them to process vast amounts of complex data quickly and efficiently.

LRZ has used SUSE solutions for over two decades and considers SUSE to be one of the best operating systems for performing standard HPC workflows and for compatibility with LRZ's HPC software stack. LRZ says that the SUSE operating system is scalable enough to accommodate future hardware expansions.

FUTURE OUTLOOK

Since its introduction in the 1998-2000 era, the standards-based Linux cluster has come to dominate the worldwide HPC market. All 500 systems on the June 2019 list of the world's TOP500 supercomputers run Linux, as do most of the more than 90,000 other HPC systems sold around the world each year. Linux's salient attribute is openness that allows this operating system to evolve and adapt in lockstep with the global HPC community. This openness has helped Linux to democratize HPC and move it into new areas, including AI, cloud computing and enterprise data centers. HPC is indispensable today at the forefront of R&D and practice for the most important AI use cases, including automated driving systems, business intelligence, precision medicine, affinity marketing, fraud and anomaly detection,

cyber security, smart cities and the Internet of Things. Major edge computing initiatives around the world rely heavily on HPC.

Linux predominates in all of these situations that help define the new era of AI and digital transformation. Hyperion Research believes that Linux is here to stay and will remain the operating system of choice as HPC conquers new markets and use cases.

Enterprises that adopt HPC typically prefer vendor-supported Linux versions to meet the requirements of production computing environments. The same holds true for an important subset of HPC sites in government and academia. SUSE has already proven its value in the historical HPC market and has extended its reach into new segments including agriculture, construction, healthcare, publishing and the retail industry. The company's close links to SAP users mean that SUSE is well grounded in the rising enterprise market for digital transformation and AI, on premises and in third-party clouds such as Azure, AWS and Google Cloud (SUSE is available on all three).

Hyperion Research therefore believes that SUSE is well positioned to benefit from, and help drive, the robust growth we forecast for the global HPC market.

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 multiuser 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 and www.hpcuserforum.com

Copyright Notice

Copyright 2019 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.