

## Quick Take

# An Accelerator View of the 2018 HPC Server Market

Alex Norton, HPC Research Analyst  
May 2019

### HYPERION RESEARCH OPINION

---

Hyperion Research recently closed the books on the 2018 year for the high performance computing market in our market tracking database, the QView. 2018 was another high growth year, with the server market reaching \$13.7 billion, and the total market (servers, storage, software and technical support services) totaling more than \$27.6 billion. The market has continued to grow at a steady rate of around 6% for the past 5 years.

Although there is a plethora of data built into the QView, which contains all of the HPC sales for the past 25 years, this series of reports highlights some of the key trends and numbers. This Quick Take shows the accelerator market from a five-year historical view. NVIDIA has always been the leader in accelerators with their GPU for the high performance computing market, especially for the most expensive machines in the world. Intel's Phi accelerator had some success in the HPC market, never quite to the level of NVIDIA, but was generally well used until 2017 when sales dropped, and subsequently the Phi product line was discontinued in 2018. Contrastingly, NVIDIA continued to prosper, and although they had a slight dip at the high end of the market in 2017, the NVIDIA GPU has continued to be the leading supercomputer accelerator. In 2018, NVIDIA had an enormous supercomputer board shipment number, mainly due to the Summit and Sierra systems being accepted at Oak Ridge National Lab and Lawrence Livermore National Lab, respectively. The other category includes FPGAs, ASICs, and any other accelerators in the market.

The "other" category is about to become far more diverse, as a slew of new accelerators are gaining traction in the market, especially accelerators targeted for AI workloads. Many of the other accelerators here are FPGAs, but ASICs are also starting to gain some traction among HPC users. Table 1 below shows the breakdown of the different segments of the HPC market, and the accelerators sold into each of those segments in 2018, as well as their five-year CAGR.

**TABLE 1****Accelerator Market (By segment)**

	Segment	2018	5 Year CAGR
Sum of NVIDIA Board Count	cs1-Supercomputer	100,389	30.4%
	cs2-Divisional	27,010	27.9%
	cs3-Departmental	31,875	36.2%
	cs4-Workgroup	13,917	22.7%
	Total	173,191	31.2%
Sum of Other Acce/CoProc Board Count	cs1-Supercomputer	3,573	14.0%
	cs2-Divisional	1,048	1.4%
	cs3-Departmental	1,311	16.6%
	cs4-Workgroup	1,476	11.8%
	Total	7,408	11.5%
<b>Total</b>		<b>180,599</b>	<b>29.0%</b>

Source: Hyperion Research, 2019

**FUTURE OUTLOOK**

The HPC market for accelerators promises to expand at a healthy rate, but the longer-term leaders in this market are hard to predict. NVIDIA dominates the GPU market today, although other processor makers are starting to invest in GPU technology as alternatives to NVIDIA. The FPGA is proving to be an interesting alternative to the GPU as well, yet their momentum has not taken off quite yet. The exascale race is shining some light on accelerators, mainly for diverse workloads, and GPUs have proven to be great boosters for Linpack performance. This is not to say that they don't do anything else well, but it's clear they play an integral role in increasing the performance for Linpack runs.

The AI market is a major change agent in the accelerator sector. Right now, there are several startups emerging with AI-centric processors, intended to act as AI accelerators. Furthermore, several large companies are investing in companies or technologies that could enhance their accelerator platforms, such as Intel's purchase of Nervana Systems. Nervana Systems is an AI software company, and their

technology was combined with Intel's processor technological expertise to create the Nervana chip, a neural network chip.

- It is still too early for many of these technologies to accurately assess their potential technological impact or eventual market position.
- With the increase of diversity in workloads and the convergence of HPC and AI workloads, the accelerator market will be right at the center of change.

## APPENDIX: DEFINITIONS USED IN THIS REPORT

---

### Definition of Technical Computing (HPC)

Hyperion Research uses the terms *technical computing* and *high-performance computing (HPC)* to encompass the entire market for computer servers used by scientists, engineers, analysts, and other groups using computationally and/or data-intensive modeling and simulation applications.

Technical servers range from small servers costing less than \$5,000 to the large-capability machines valued in hundreds of millions of dollars. In addition to scientific and engineering applications, technical computing includes related markets/applications areas such as economic analysis, financial analysis, animation, server-based gaming, digital content creation and management, business intelligence modeling, and homeland security database applications. These areas are included in the technical computing market based on a combination of historical development, applications type, computational intensity, and associations with traditional technical markets.

### Tracking Methodology

Each quarter, Hyperion Research analysts conduct interviews with major hardware original equipment manufacturers (OEMs) in the technical computing space to gather information on each vendor's quarterly sales. Specifically, Hyperion collects data on the number of HPC systems sold, system revenue, system average selling price (ASP), the competitive segment that a system falls into, architecture of the system, average number of processor package per system, average number of nodes for each system sold, system revenue distribution by geographical regions, and system revenue distribution by operating systems.

Hyperion Research records all of the above-mentioned information and merges it into a master database, which contains over 50 data fields; some of these fields contain actual data gathered from the OEMs as described previously, some are calculated based on the actual data, and some are only used for special data cuts.

Hyperion Research then creates a pivot table based on this master database. Data tables with different views of the technical computing market can then be created from this pivot table. Hyperion refers to this data structure as the "HPC QView." In addition to the HPC QView, Hyperion maintains other HPC technical computing data structures, for example:

- HPC end-user demand-side data structure
- HPC application/industry segmentation data structure
- HPC application software data structure
- HPC server and processor sales by country database

### Hyperion Research Technical Computing Server Revenue Accounting Rules

#### *Initial System Shipment*

Initial system shipment (ISS) characterizes the first sale of a system (previously referred to as a "new footprint"), and also includes major upgrades to existing systems. An ISS unit consists of processors, memory, embedded disk storage, cluster interconnect hardware/software, any bundled operating system, compiler, math/statistical library, parallel computing, database, and networking software that would typically be configured when it leaves the OEM's factory floor. Note that separately acquired

software is not included, e.g. often the database software is purchased separately as is most ISV application software. It is recognized as a shipment only when the complete system or cluster is installed and accepted. In addition, major upgrades that include processors are treated as an ISS in the quarter that it is accepted. External user storage and all paid services are excluded from the ISS revenue value. If a system is paid for over a number of quarters, for example, via service or R&D contracts, Hyperion Research determines a value for the whole system when it is finally accepted by the buyer.

### ***Average Selling Price***

Average selling price (ASP) is the value of an initial system shipment (ISS) unit configured as it is typically sold. A unit is the whole computer system or cluster complex. The ASP includes the base configuration plus any add-ons or upgrades typically sold when the system or server is first delivered to a customer. For upgrade only sales, the ASP is based on the value received by the OEM for the upgrade. This includes the primary system interconnect used for inter-processor communication and any system disks and system software necessary for operation, but not any additional networked-attached storage or additional software packages (e.g., all user application software is excluded). Typically, an HPC system is shipped with large amounts of memory relative to systems sold for nontechnical workloads, which can add significantly to the average selling price. Hyperion Research assumes that all servers are shipped with an operating system. The portion of the operating system license fee that is shipped with the server is included in the factory-revenue figures, if paid to the OEM. Note that often in HPC the OS is free, so that its inclusion doesn't change the ASP.

The value of the system and the units associated with the sale are only counted when the system is fully installed and accepted. At this point the full value of the contract is recognized.

Neither ISS nor ASP should include the following:

- Any external storage purchased separately from the initial server system
- Any extra services purchased for the initial server system shipped
- Any application software, regardless of whether the application software is part of the contract or not
- Any additional sales revenue from the channel (Hyperion does not count revenue coming from the channels and partners of the original equipment manufacturers or original system integrators [OSIs]. All revenue Hyperion counts should be based on direct sales from the OEMs and OSIs.)

## **Hyperion Research Technical Computing Market Segmentation**

Based on input from HPC vendors and end users, Hyperion Research created four competitive segments to reflect the trends in the HPC technical server market. These competitive segments are based on average selling prices and defined as follows:

- **Supercomputers:** Technical servers that sell for \$500,000 or more
- **Divisional servers:** Technical servers that sell for \$250,000-499,999
- **Departmental servers:** Technical servers that sell for \$100,000-249,999
- **Workgroup servers:** Technical servers that sell for less than \$100,000

## 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 2019 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.