HYPERION RESEARCH

Special MCS Study

2017 HPC Multi-Client Study: HPDA (Big Data in HPC) Usage, Trends and Forecasts

Steve Conway, Earl Joseph, Bob Sorensen, and Kevin Monroe
September 2017

HYPERION RESEARCH OPINION

This report is derived from the fifth edition of Hyperion Research's end-user special study of the worldwide high-performance computing (HPC) market (Hyperion Research is the new name for the former IDC HPC analyst team). Hyperion Research coined the term “high-performance data analysis” (HPDA) to refer to data-intensive ("Big Data") workloads that require or benefit greatly from HPC resources, even though not all HPDA users consider themselves HPC users. These workloads include established data-intensive simulations and newer advanced analytics problems. HPDA problems are characterized by large data volumes, along with time-criticality and algorithmic complexity that are atypical for enterprise business workloads.

- Hyperion Research forecasts that worldwide HPDA server revenues will increase at a robust 17.0% CAGR between 2016 and 2021, creating projected server revenues of $4.0 billion in 2021. First-time HPC adopters from the commercial analytics side will represent $1.1 billion of that total.
- A subset of the HPDA server market, machine learning, deep learning and other AI methodologies (we treat them as horizontal methodologies, not vertical segments) will grow at an even faster pace (25.9% CAGR) to reach about $1.3 billion in 2021.

Despite these impressive growth projections, the worldwide HPDA market will remain in a largely exploratory stage through the end of our current forecast period in 2021. Soon after that year, we believe, the HPDA market will begin a considerably steeper rise, as mainstream, economically important HPC-dependent use cases such as precision medicine, autonomous vehicle design/networked operation, and the Internet of Things increasingly enter production usage and become norms. As a sign of things to come, HPC has already moved to the forefront of R&D for machine learning, deep learning, early AI and the IoT.

Note: this page is intentionally blank.

Executive Summary

HPDA Server and Storage Revenue Will Grow to Record Levels

Hyperion Research forecasts the overall HPDA server/storage ecosystem will be worth about $6.6 billion in 2021. Because HPDA use is present and growing in all major established HPC application/industry segments (primarily for data-intensive simulation, but increasingly also for data-intensive analytics), Hyperion Research treats most HPDA revenue as a workload view of the worldwide HPC market rather than as a market distinct from government, academia, aerospace, bio-life sciences, and other established HPC application/industry segments.

Hyperion Research anticipates total worldwide HPDA server revenues will increase at a robust 17.3% CAGR between 2016 and 2021, rising from $1.8 billion to $4.0 billion by 2021. We predict HPDA storage revenues will expand at an even-faster 19.0% CAGR, growing from $1.1 billion in 2016 to $2.6 billion in 2021.

Hyperion Research anticipates total worldwide HPDA server revenues will increase at a robust 17.3% CAGR

On average, about one-quarter of the aggregate capacity of the sites' HPC systems (26.3%) is devoted today to Big Data analytics, up from 22.7% of aggregate cycles in 2015. (Aggregate capacity is the total capacity of all HPC systems at a site; as Table 6 indicates, the average site in this study had 8.8 HPC systems.)

Big Data analytics usage is expected to increase substantially in the next 6-18 months, rising on average to 34.9% of the sites' aggregate HPC system capacity. As Table 13 shows, expectations for increased Big Data analytics use looking ahead 6-18 months have also increased strongly since the 2015 version of this study, growing by 8.9%.

New HPDA Segments

The smaller, faster-growing market for advanced commercial analytics using HPC requires different treatment. Hyperion Research identifies and tracks the new commercial HPDA market segments using these definitions:

- **Fraud and anomaly detection.** This "horizontal" workload segment centers around identifying harmful or potentially harmful patterns and causes using graph analysis, semantic analysis, or other high performance analytics techniques. The patterns may point to fraud, which is the deceptive exploitation or annotation of data for wrongful or illegal personal gain, or they may point to cyber security crime or insider threats, significant errors, or other anomalies that may deserve further investigation.
- **Affinity Marketing.** This segment covers the use of HPDA to promote products or services, typically using complex algorithms to discern potential customers' demographics, buying preferences and habits. From a mathematical standpoint, fraud/anomaly detection and affinity marketing are closely allied problems, one targeting exclusion and the other, inclusion.
- **Business intelligence.** The workload segment uses HPDA to identify opportunities to advance the market position and competitiveness of businesses, by better understanding themselves, their competitors, and the evolving dynamics of the markets they participate in.
- **Precision Medicine:** a medical model based on customizing healthcare, with medical decisions, practices, and products tailored to the individual patient rather than following a set of standardized procedures applicable to many patients. HPC is helping to lead the way toward precision medicine, especially through DNA sequencing and acting as a decision-support tool for physicians and other providers. (Alternative term: personalized medicine)

- **Commercial: other.** This catchall segment includes all commercial HPDA workloads other than the four just described. Over time, Hyperion Research expects some of these workloads to become significant enough to split out of this "other" category and command their own segments, as happened recently with precision medicine. An example of such a high-potential workload is the use of HPDA to manage large IT infrastructures, ranging from on premise data centers to public clouds and Internet-of-Things (IoT) infrastructures.

Key Factors Affecting the HPC Market

- **The market for HPDA systems is still formative, creating attractive opportunities for savvy vendors**. HPDA is a formative market that includes many promising use cases, only a few of which today are economically important enough to qualify as pursuable market segments. This fluid situation creates opportunities for savvy vendors to seize leadership positions with compelling horizontal (e.g., fraud detection) and vertical (e.g., precision medicine) solutions – and with server, storage, software, networking, and service offerings to underpin these solutions. With newer approaches, such as graph analysis, machine and deep learning, for example, vendors may uncover golden opportunities to differentiate in various ways, ranging from customized hardware platforms to smart models and algorithms for operating in the cloud, to vertical-specific software and services offerings.

Storage and data movement rank high as pain points.

- **Storage and data movement rank high as pain points.** In other recent Hyperion Research studies on the HPDA market, buyers listed attributes that would most command 10–15% premium pricing for HPC system purchases. Among the most important were, "higher-performance interconnects between nodes" and "higher-performance external I/O and storage." Storage and data movement have emerged as key pain points, and therefore also as key purchase criteria for many HPC buyers.
- **Machine/deep learning and cognitive computing.** Machine learning and its more profound relatives, deep learning and cognitive computing/AI, will play increasingly important roles in the HPDA market -- not only for advanced problem-solving analytics, but also for helping to relieve programmers by further automating software stack functions. HPC-supported machine learning differs from enterprise machine learning principally in the former's heavy use of parallelized algorithms. Hyperion Research expects HPC users to accelerate the evolution from today's mostly supervised machine learning to the emerging era of higher-value, unsupervised deep learning and AI/cognitive computing.

Machine learning and its more profound relatives, deep learning and cognitive computing, will play increasingly important roles in the HPC and HPDA markets.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF TABLES — Continued

LIST OF FIGURES

SITUATION OVERVIEW

In This Report

This Hyperion Research study presents our latest five-year forecast for the HPDA portion of the worldwide HPC market, covering the 2016-2021 period. To set the stage, the report begins with a forecast for the overall HPC market, then zeroes in on the HPDA subset.

Introduction

The HPC market is entering a perfect storm. For years, HPC architectures have tilted further and further from an optimal balance between processor speed, memory access, and I/O speed. As successive generations of HPC systems have upped peak processor performance without corresponding advances in per-core memory capacity and speed, systems have become increasingly compute-centric, and the well-known "memory wall" has gotten worse. Now comes the HPC Big Data era with its demands for superb memory and I/O capabilities, sometimes with little need for computing prowess or double precision. Data-intensive (high-performance data analysis) workloads are emerging relatively rapidly. As a result, Hyperion Research predicts that the global market for HPDA servers and storage will reach about $6.6 billion in 2021.

Hyperion Research predicts that the global market for HPDA servers and storage will reach about $4.9 billion in 2019.

Emerging data-intensive problems are exposing more limitations of established HPC architectural designs — not just in the memory wall itself, but also in how existing compute-centric architectures handle data movement throughout the system. It's important to make advances here, or data movement for emerging high-performance data analysis problems could become frustratingly slow and expensive.

Hyperion's studies shows the HPC community plans to deploy two main strategies for addressing data movement challenges. The first is to accelerate data movement via more capable interconnect networks. Think here not only of interconnect vendors such as Mellanox, Cisco and Cray, but also of Intel's Omni Path interconnect fabric architecture and newer competitors such as Bull (Atos), EXTOLL, and Numascale. The second approach is to reduce data movement at all levels. This approach, widely practiced in the enterprise (non-HPC) machine learning market, saves time and electricity use, but can limit problem resolution too severely for many HPC/HPDA users.

The shift toward better compute/data movement/storage balance will require a reorientation of planning and purchasing practices. Storage and data movement can no longer remain secondary considerations. The goal of using available budgets to maximize peak and LINPACK flops ("machoflops") will need to give way over time to a more singular focus on user requirements for sustained performance and time to solution. The TOP500 list will remain valuable for census-tracking large systems and trends over time, but the shift away from strong compute centrism will increase the importance of more balanced benchmarks for HPC buyers/users, along with HPC systems that have greater breadth-of-applicability.

Methodology

This HPDA report is a component of the fifth edition of Hyperion Research's demand-side tracking of the worldwide HPC market and augments Hyperion Research's established supply-side tracking. This study evaluates on-site HPC systems, requirements, purchasing criteria, and budgeting patterns of buyers. For this edition, Hyperion Research conducted in-depth interviews with purchasing decision-makers and influencers from HPC sites worldwide in the government, industry and academic sectors.

We collected data on 972 HPC systems. To promote candor, Hyperion Research promised not to associate survey responses with individual or organizational identities.

As in earlier editions, Hyperion Research strove to ensure that the industrial sector was especially well represented; in this edition, this segment represents 45.9% of the surveyed sites. This strong participation enables Hyperion Research to drill down meaningfully into individual sub-segments of the industrial sector.

Hyperion Research's demand-side study provides a deeper and more direct view into the living, evolving world of HPC data centers. Where supply-side tracking sees an HPC system's growth through multiple additions as multiple systems (i.e., separate shipments), the demand-side perspective recognizes the expanded resource as a *single* system and also identifies after-market additions such as interconnects, storage, value-added software, and accelerator technology. (For more on distinctions between supply-side and demand-side perspectives, see the following section, *The Demand Side Perspective*.)

Hyperion Research makes results of its comprehensive demand-side study available through a series of six reports that can be purchased as a discounted package or separately. Topics include:

- Processors, Clusters and System Attributes
- System Software and Middleware
- Storage and Interconnects
- Running HPC Workloads in External Clouds
- HPDA: Big Data in HPC
- Workloads by Industry/Applications
- New: HPC Budget Mix and Key Purchase Criteria

The Demand-Side Perspective: An Important Complementary View of the HPC Market

For more than 25 years, IDCs HPC team (now called Hyperion Research) has tracked worldwide HPC server revenues based on detailed quarterly supply-side reporting of vendors' factory shipments. Our aim, not perfectly realizable, has been to track every HPC-bound server system leaving a vendor's or subcontractor's manufacturing facility, and to characterize each system across numerous attributes related to actual ("street") pricing, component technologies, intended usage, and other characteristics of interest to Hyperion Research clients.

In 2008, we began augmenting vendor-centric, supply-side tracking and analysis with a new view of the HPC market: the buyer-centric, demand-side perspective. This complementary perspective is valuable for multiple reasons, as captured in the chart that follows:

Demand-side tracking	Supply-side tracking
"Sees" the reality of a single system growing in stages, expanded via multiple shipments and installments.	Treats each shipment leaving the factory as separate.
Captures information about "dark" clusters – sales vendors have not reported as part of HPC revenue, often because such clusters were purchased as components from computer-supply websites or outlets.	Does not capture specific information about "dark" clusters.
Follows the life cycle of a system configuration as it evolves at the user site. Identifies not only expanded processing power, but also on-site, after-purchase additions of storage, interconnects, software, alternative processors, and other components, as well as decommissioning of systems.	Focuses on initial system configuration, including internal storage bundled into a factory-shipped system.
Explicitly identifies the sites' largest and second-largest HPC systems, along with other installed HPC equipment. This helps identify emerging trends in system attributes and technology adoption.	Does not explicitly identify the largest systems.
Because reports are based on qualitative and quantitative user responses, reports capture valuable information about the intent of HPC buyers, including acquisition rationales and methods, along with budgetary and spending intent, including attributes for which buyers would pay a premium.	Does not provide qualitative information.

This 2017 fifth edition of Hyperion Research's demand-side tracking report captures information about sites representing 972 HPC systems. Demand-side tracking complements – but does not replace – our supply-side tracking, which keeps close tabs on more than 100,000 HPC systems per year. Hyperion Research conducts more deeply focused demand-side studies for private sponsors, and maintains close contact with demand-side requirements and trends through the Hyperion Research-operated HPC User Forum.

HPC Market Size and Overall Trends

The HPC Market in Revenues, Units and ASP

Worldwide factory revenue for the high-performance computing (HPC) technical server market grew 4.4% in full-year 2016 to a record $11.2 billion, up from $10.7 billion in 2015 and from the previous record of $11.1 billion in exceptionally strong 2012, according to the Hyperion Research Worldwide High Performance Technical Server QView.

HPC Market Forecasts

Worldwide HPC Server Market Forecast by Revenue, Units and ASP

As Table1 indicates, Hyperion Research predicts that the worldwide HPC server market will expand at a robust rate (5.8% CAGR) to reach $14.8 billion in 2021, up from $11.2 billion in 2016. With average selling prices (ASPs) essentially flat (1.4% CAGR), most of the predicted growth will be generated by increased unit sales (4.3% CAGR), with that figure predicted to top 103,000 server systems in 2021.

Table 1

Worldwide HPC Server Market Forecast by Revenue, Units and ASP

	2015	2016	2017	2018	2019	2020	2021	2016/2021 CAGR
Revenue $Millions	$10,727	$11,200	$12,019	$12,302	$12,988	$13,997	$14,819	5.8%
Units	99,817	83,519	87,130	90,897	94,827	98,926	103,203	4.3%
ASP $000	$107	$134	$138	$135	$137	$141	$144	1.4%

Source: Hyperion Research, 2017

HPC Server Revenues by Competitive Segments

Through our HPC QView product, Hyperion Research closely tracks quarterly revenues for HPC server systems at all price points. Table 2 displays our forecasts for worldwide HPC server revenue by Hyperion Research's competitive segments, which correspond to price bands.

- The Supercomputer segment for HPC systems priced at $500,000 and up will expand fastest during the forecast period (6.9% CAGR), Hyperion Research predicts, and will reach $5.7 billion, or 38.5% of all HPC server revenues, by 2021.
- The Divisional segment for HPC systems sold for $250,000 to $449,000 will grow at a healthy 4.6% CAGR to $2.8 billion in 2021, or 19.2% of all HPC server revenues.
- The Departmental segment for systems in the $100,000 to $249,000 price range will expand at 5.2% CAGR to $4.1 billion in 2021, representing 27.4% of all HPC server revenues.
- The Workgroup segment for sub-$100,000 HPC systems will grow at a 5.5% CAGR, producing revenues of $2.2 billion in 2021, or 14.9% of HPC server revenues worldwide.

Table 2

Worldwide HPC Server REVENUE Forecast, by Competitive Segment ($Millions)

	2015	2016	2017	2018	2019	2020	2021	2016/2021 CAGR
Supercomputer	$3,284	$4,091	$4,375	$4,324	$4,601	$5,249	$5,700	6.9%
Divisional	$2,212	$2,273	$2,404	$2,514	$2,655	$2,750	$2,848	4.6%
Departmental	$3,994	$3,147	$3,384	$3,538	$3,717	$3,885	$4,062	5.2%
Workgroup	$1,874	$1,689	$1,856	$1,925	$2,016	$2,113	$2,208	5.5%
Total	$10,727	$11,200	$12,019	$12,302	$12,988	$13,997	$14,819	5.8%

Source: Hyperion Research, 2017

HPC UNIT SHIPMENTS by Competitive Segments

As Table 3 shows, Hyperion Research forecasts that the number of HPC server systems (units) shipped will increase at a healthy clip (4.3% CAGR) during the 2016-2021 forecast period, rising to more than 103,000 units in 2021. That represents aggregate growth of 23.6% during this period. The CAGR figures for the competitive segments don't differ greatly, ranging only between 3.5% and 4.4%, but gains in the number of units shipped per year are heavily weighted toward the lower price bands. For instance, combined unit shipment gains in the Departmental and Workgroup segments (18,080 additional units) account for 92.0% of all unit shipment gains (19,654) during the forecast period.

Table 3

Worldwide HPC Server UNITS SHIPPED Forecast, by Competitive Segment

	2015	2016	2017	2018	2019	2020	2021	2016/2021 CAGR
Supercomputer	1,392	1,526	1,579	1,634	1,691	1,750	1,811	3.5%
Divisional	6,033	6,029	6,272	6,525	6,789	7,063	7,348	4.0%
Departmental	23,267	18,724	19,520	20,349	21,213	22,114	23,053	4.2%
Workgroup	69,125	57,240	59,759	62,389	65,134	68,000	70,991	4.4%
Total	99,817	83,519	87,130	90,897	94,827	98,926	103,203	4.3%

Source: Hyperion Research, 2017

HPC Servers AVERAGE SALES PRICE (ASP) by Competitive Segments

As Table 4 illustrates, Hyperion Research expects average selling prices (ASPs) to remain essentially flat during the forecast period for all segments except Supercomputers. The ongoing petascale/ exascale race will inject a bit more momentum into Supercomputer ASPs (3.3% CAGR).

Table 4

Worldwide HPC Average Sales Price Forecast, by Competitive Segment ($K)

	2015	2016	2017	2018	2019	2020	2021	2016/2021 CAGR
Supercomputer	2,359	2,681	2,771	2,646	2,721	3,000	3,148	3.3%
Divisional	367	377	383	385	391	389	388	0.6%
Departmental	172	168	173	174	175	176	176	0.9%
Workgroup	27	30	31	31	31	31	31	1.1%
Overall Average ASP	107	134	138	135	137	141	144	1.4%

Source: Hyperion Research, 2017

Forecast: The Broader HPC Market

Table 5 shows Hyperion Research's five-year revenue forecast for the broader HPC market that includes not just technical servers, but also storage, middleware and operating systems (OS), applications software, and technical support services.

To put this in historical perspective, worldwide revenues for the broader HPC market have grown from about $2 billion in 1990 to $22.4 billion in 2016. This 10-fold leap – owing much to the introduction of standards-based clusters in 2001-2002 – means HPC has been one of the fastest-growing IT markets in recent decades.

As Table 5 illustrates, in 2021 Hyperion Research expects broader HPC revenues to reach $30.3 billion, with nearly half of that total (48.8%) spent on computer servers, followed by storage (20.8%) and applications software (16.8%). Repair services (7.6%) and middleware/OS (5.9%) represent less significant spending categories. As it has been in recent years, storage will have the highest CAGR (7.8%) during the forecast period, owing in large part to the rise of data volumes for both simulation and analytics HPC workloads.

Spending for each category may vary greatly by scientific/engineering domain, however. For example, recent Hyperion research revealed that manufacturing firms may spend 4-5 times more per year on third-party applications software licenses than on computer servers.

Table 5

The Broader HPC Market

Worldwide HPC Compute, Storage, Middleware, Application and Service Revenues ($Billions)								
	2015	2016	2017	2018	2019	2020	2021	2016/2021 CAGR
Server	$10.7	$11.2	$12.0	$12.3	$13.0	$14.0	$14.8	5.8%
Storage	$4.3	$4.3	$4.5	$5.0	$5.3	$5.8	$6.3	7.8%
Middleware/ OS	$1.3	$1.3	$1.3	$1.4	$1.5	$1.7	$1.8	6.9%
Applications	$3.7	$3.7	$3.9	$4.2	$4.4	$4.8	$5.1	6.3%
Repair Services	$1.9	$1.9	$2.0	$2.0	$2.1	$2.2	$2.3	3.9%
Total	$21.9	$22.4	$23.7	$24.9	$26.3	$28.5	$30.3	6.2%

Source: Hyperion Research, 2017

Respondent Demographics

Sample in Total Servers, Processors, Nodes, and Peak PF

Hyperion Research conducted in-depth interviews with one or more representatives of 111 HPC sites around the world. These sites had a total of 972 technical server systems installed. The average in the present study was 8.8 systems per site. See Table 6.

As Table 6 shows, the total number of processors (physical socket parts) installed in the present study was 542,293. The aggregate core count was 6,287,210, yielding an average of 11.6 cores per CPU, a sharp rise from the 6.1 average figure in the 2015 version of the present study. There were 353,658 nodes among the installed systems, yielding an average of 1.53 processors per node and 17.8 cores per node.

Table 7 contains per-site information on the study's collection of 111 sites and 972 installed systems. On average, the sites had 4,886 processors (physical socket parts), 56.642 cores, and 3,186 nodes.

Table 6

HYPERION RESEARCH HPC End User MCS Sample Profile

Attribute	Sample Total
Number of Servers in the Sample	972
Number of Sites in the Sample	111
Average HPC Servers per Site	8.8
Total Processor Count (sockets)	542,293
* Average CPUs per system	557.9
Total Processor Core Count	6,287,210
* Average Cores per system	6,468.32
* Average Cores per CPU	11.6
Total System Node Count	353,658
* Average Nodes per system	363.8
Total Peak PF Installed	565
* Average PFlops per system	0.58
N = 111	

Source: Hyperion Research, 2017

Table 7

Sample Characteristics: Changes in Per System Overall Averages

	2015 Value	2017 Value
Average number of processor sockets/site	2,966	4,886
Average number of cores/site	52,617	56,642
Average number of nodes/site	1,539	3,186
Average length of time that a system is installed	4.41	3.62
N = 111		

Source: Hyperion Research, 2017

Sample by Sector

Figure 1 and Table 8. Hyperion Research deliberately weighted the sample of 111 sites hosting 972 HPC systems in favor of industry participants (45.9% of the sites), to allow meaningful explorations of individual industry segments (see Figure 1). Academic sites were also heavily represented in the study (36.0%). Fewer government sites were profiled (18.0%), but these sites were carefully chosen to reflect a wide range of government sites within the global HPC community.

FIGURE 1

Sample by Major Sector

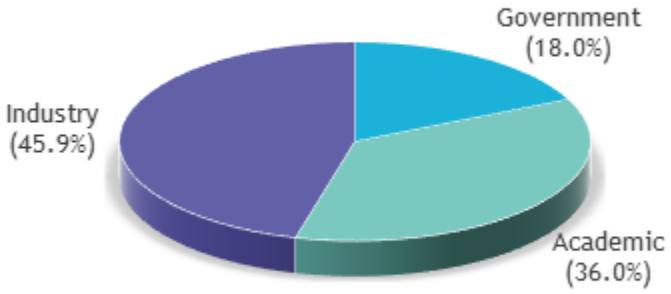

Source: Hyperion Research, 2017

Table 8

Sample by Major Sector

Q1.a. Is your organization in government, commerce/industry, or academia?		
Sector	Number of Responses	Percentage of Sample
Government	20	18.0%
Academic	40	36.0%
Industry	51	45.9%
Total	111	
N = 111		

Source: Hyperion Research, 2017

Sample by Industries

As noted previously (Figure 1), Hyperion Research deliberately weighted the sample in favor of industry. Figure 2 and Table 9 display the breakdown of industrial sites by industry segment. The financial services and manufacturing segments were most heavily represented within the population of

industrial sites. We were pleasantly surprised at the number of financial institutions willing to provide information for this study, given that industry's historical reluctance to participate in third-party surveys.

FIGURE 2

Industrial Mix in the Sample

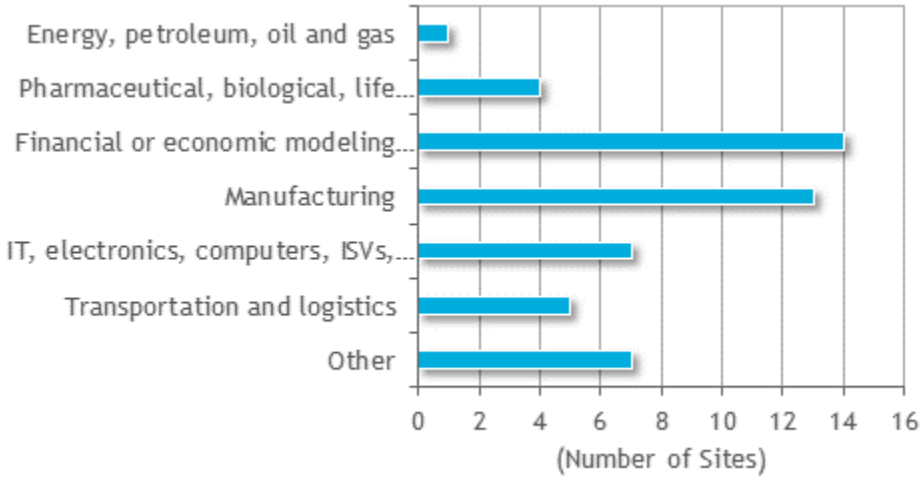

Source: Hyperion Research, 2017

Table 9

Sample By Industry

Q. If commercial or industry, which one best fits your organization?

Sector	Number of Responses	Percentage of Sample
Energy, petroleum, oil and gas	1	2.0%
Pharmaceutical, biological, life sciences, healthcare	4	7.8%
Financial or economic modeling and BI	14	27.5%
Manufacturing	13	25.5%
IT, electronics, computers, ISVs, ITSP, cloud provider, telco	7	13.7%
Transportation and logistics	5	9.8%
Other	7	13.7%
Total Industry	51	100.0%
N = 111		

Source: Hyperion Research, 2017

Number of HPC Servers Per Site

We next asked respondents, "How many HPC server systems of all types (clusters, SMP systems, et al.) does your organization have?" As Table 10 shows, 41.4% of the surveyed sites said they have 2-5 HPC systems, and about one in five (21.6%) operated more than 15 systems. The average figure was 8.8 HPC server systems per site (see Table 6).

Table 10

Number of Servers Per Site

Q. How many HPC server systems of all types (clusters, SMP systems, etc.) does your organization have?

Attribute	Number of Responses	Percentage of Sample
a) 1	12	10.8%
b) 2 to 5	46	41.4%
c) 6 to 10	24	21.6%
d) 11 to 15	5	4.5%
e) More than 15	24	21.6%
Total	111	100.0%
N = 111		

Source: Hyperion Research, 2017

Growth Driver: Pent-Up Demand

Three-quarters (78.1%) of the surveyed sites reported that they have pent-up demand for HPC systems, defined as work that users would like to run, but can't run due to limited system resources (Table 11). Nearly half of the sites (46.9%) said they have pent-up demand equivalent to no more than 50% above the aggregate capacity of their HPC systems, while 31.1% of the sites reported pent-up demand larger than that, in some cases (13.5%) 100% or more beyond the aggregate capacity of their systems.

Table 11

Estimated HPC Growth Potential: Pent-up Demand at End User Sites

Estimated amount of "pent-up demand" that exists for HPC systems (work that users would like to run, but can't due to limited resources)		
Response	Number of Responses	Percentage of Sample
None, we are still in the mode of filling up the system	21	21.9%
≤ 24%	31	32.3%
25-50%	14	14.6%
51-100%	17	17.7%
101-200%	11	11.5%
201-300%	1	1.0%
> 300%	1	1.0%
N = 96		

Source: Hyperion Research, 2017

BIG DATA IN HPC (HPDA): WHERE BIG DATA MEETS HPC

Drivers for the HPDA Market

The factors driving businesses to adopt HPC for big data analytics (i.e., HPDA) fall into a few main categories:

- **High complexity.** HPC technology allows companies to aim more complex, intelligent questions at their data infrastructures. This ability can provide important advantages in today's increasingly competitive markets. HPC technology is especially useful when there is a need to go beyond query-driven searches to discover unknown patterns and relationships in data – such as for fraud detection, to reveal hidden commonalities within millions of archived medical records, or to track buying behaviors through wide networks of relatives and acquaintances. Hyperion Research believes that HPC technology will play a crucial role in the transition from today's static searches to the emerging era of higher-value, dynamic pattern discovery and cognitive computing.
- **High time criticality.** Information that is not available quickly enough may be of little value. The weather report for tomorrow is useless if it's unavailable until the day after tomorrow. At PayPal, enterprise technology was unable to detect fraudulent transactions until after the charges had hit consumers' credit cards. The move to high-performance data analysis using HPC technology corrected this problem. For financial services companies engaged in high

frequency trading, HPC technology enables proprietary algorithms to exploit market movements in minute fractions of a second, before the opportunities disappear.

- **High variability.** People generally assume that big data is "deep," meaning that it involves large amounts of data. They recognize less often that it may also be "wide," meaning that it can include many variables. Think of "deep" as corresponding to lots of spreadsheet rows and "wide" as referring to lots of columns (although a growing number of high-performance data analysis problems don't fit neatly into traditional row-and-column spreadsheets). A "deep" query might request a prioritized listing of last quarter's 500 top customers in Europe. A "wide" query might go on to analyze their buying preferences and behaviors in relation to dozens of criteria (affinity marketing). An even "wider" analysis might employ graph analytics to identify any fraudulent behavior within the customer base (fraud/anomaly detection).

Volume of Big Data Workloads in 2017

Surveyed sites were next asked about their use of Big Data analytics on HPC servers, a category that includes MapReduce/Hadoop, graph analytics, semantic analysis, machine learning, deep learning, early AI and other Big Data methods. As Table 12 indicates, on average about one-quarter of the aggregate capacity of the sites' HPC systems (26.3%) is devoted today to Big Data analytics – up from 22.7% of aggregate cycles in 2015. Big Data analytics usage is expected to increase substantially in the next 6-18 months, rising on average to 34.9% of the sites' aggregate HPC system capacity. As Table 12 shows, expectations for increased Big Data analytics use looking ahead 6-18 months have also increased strongly since the 2015 version of this study, growing by 8.9%.

Percentage of Aggregate System Capacity Running Big Data Jobs

Table 12

Percentage of System Running Big Data Jobs

Q. Of all the workloads on your systems, what percentage of your total system utilization is devoted to Big Data analytics?

	Percentage of Systems (2015)	Percentage of Systems (2017)	Change
Currently used for Big Data analytics - Percent	22.7%	26.3%	3.6%
Expected 6 to 18 months from now - Percent	26.1%	34.9%	8.9%
N = 87			

Source: Hyperion Research, 2017

Percentage of Aggregate System Capacity Running Big Data Jobs, by Sector

As Table 13 shows, Big Data analytics workloads already utilize a substantial minority of aggregate HPC system capacity, on average, at the surveyed government sites (24.6% of capacity) and industry sites (30.3%), with lower average utilization (12.8%) in academia. These utilization figures have risen

markedly since our 2015 version of this study, as the table shows, except for the relatively flat growth reported by the industrial sites. In the next 6-18 months, all the segments expect substantial growth in utilization for Big Data analytics workloads.

Table 13

Percentage of System Running Big Data Jobs, By Sector

Of all the workloads on your systems, what percentage of your total system utilization is devoted to Big Data analytics (MapReduce/Hadoop, graph analytics, semantic analysis, etc.)?				
In 2017:	Percentage of System in Government	Percentage of System in Academia	Percentage of System in Industry	Overall
Currently used for Big Data analytics - Percent	24.6%	12.8%	30.3%	26.3%
Expected 6 to 18 months from now - Percent	33.1%	22.9%	39.2%	34.9%
In 2015:				
Currently used for Big Data analytics - Percent	18.2%	8.4%	32.6%	22.7%
Expected 6 to 18 months from now - Percent	19.8%	15.6%	33.4%	26.1%
N = 87				

Source: Hyperion Research, 2017

Top HPDA Big Data Application Areas

The largest percentage of the sites (45.8%) identified knowledge discovery as their top data-intensive analytics application, followed by graph analysis (26.2%) and more distantly by semantic analysis (17.8%). See Table 14. Because "knowledge discovery" is a relatively vague term, it is fair to assume that it was seen as including much of what did not fit into the more precise terms "graph analysis" and "semantic analysis."

Table 14

Top HPDA Applications in 2017

Q. What application categories do your top 3 data-intensive analytics applications fall under?

	Number of Responses	Percentage of Responses
App #1 - Graph Analysis	28	26.2%
App #1 - Semantic Analysis	19	17.8%
App #1 - Knowledge Discovery	49	45.8%
App #1 - Other	11	10.3%
N = 87		

Source: Hyperion Research, 2017

Table 15 splits out the findings shown in Table 14 in two ways: first by competitive segment and second by the most important data-intensive application, second-most-important data-intensive application, and third-most-important data-intensive code.

- Among the sites' number one (most important) data-intensive applications, graph analytics was most popular among government sites (33.3% of government sites), while semantic analysis (19.0%) and knowledge discovery (47.6%) scored highest within industry.
- When asked about their second-most-important data-intensive applications, industry led in graph analysis (23.5% of industry sites), academia led in semantic analysis (27.8%), and government scored highest in knowledge discovery (50.0%).
- Table 15 also shows results for the sites' third-most-important data-intensive codes.

Table 15

Top HPDA Applications in 2017

What application categories do your top 3 data-intensive analytics applications fall under?						
	Government Number of Responses	Government Percentage of Responses	Academia Number of Responses	Academia Percentage of Responses	Industry Number of Responses	Industry Percentage of Responses
App #1 - Graph Analysis	3	33.3%	9	25.7%	16	25.4%
App #1 - Semantic Analysis	1	11.1%	6	17.1%	12	19.0%
App #1 - Knowledge Discovery	4	44.4%	15	42.9%	30	47.6%
App #1 - Other	1	11.1%	5	14.3%	5	7.9%
App #2 - Graph Analysis	1	16.7%	6	16.7%	12	23.5%
App #2 - Semantic Analysis	1	16.7%	10	27.8%	10	19.6%
App #2 - Knowledge Discovery	3	50.0%	14	38.9%	26	51.0%
App #2 - Other	1	16.7%	6	16.7%	3	5.9%
App #3 - Graph Analysis	3	37.5%	8	25.8%	16	34.0%
App #3 - Semantic Analysis	1	12.5%	8	25.8%	8	17.0%
App #3 - Knowledge Discovery	3	37.5%	10	32.3%	21	44.7%
App #3 - Other	1	12.5%	5	16.1%	2	4.3%
Total Respondents						
N = 87						

Source: Hyperion Research, 2017

Middleware Used for HPDA Applications

When asked which middleware software they use for data-intensive workloads, the sites named only three software packages with any real frequency: Hadoop (54.0% of the sites), SPARC (33.3%), and LexisNexis' HPCC software (26.4%). See Table 16.

Table 16

Middleware to Support HPDA Workloads

What types of middleware system software are you using for managing your data intensive workloads (e.g., Hadoop, HPCC, etc.)?		
Type Software Used	Number of Responses	Percentage of Respondents
Hadoop	47	54.0%
SPARC	29	33.3%
HPCC	23	26.4%
IBM IIS/Watson/WebSphere/ MQSeries/InfoSphere/DataStage	6	6.9%
In-house /home grown	4	4.6%
Oracle WebLogic/Coherence/ERP native	4	4.6%
Red Hat/Linux	2	2.3%
HPC Manager	2	2.3%
Cylc	1	1.1%
Python framework	1	1.1%
Lustre	1	1.1%
NUMAlink	1	1.1%
BigInsight	1	1.1%
Exadata	1	1.1%
Open MPI	1	1.1%
Intel HPC	1	1.1%
Informatica	1	1.1%
CEP	1	1.1%
TIBCO Spotfire	1	1.1%

Table 16

Middleware to Support HPDA Workloads

What types of middleware system software are you using for managing your data intensive workloads (e.g., Hadoop, HPCC, etc.)?		
Type Software Used	Number of Responses	Percentage of Respondents
Ignite	1	1.1%
Regred	1	1.1%
Techila	1	1.1%
LoadLeveller	1	1.1%
FMW	1	1.1%
Clarisse	1	1.1%
Sun HPC Cluster Tools	1	1.1%
Arcwieve	1	1.1%
Hbase	1	1.1%
Splunk	1	1.1%
PVM	1	1.1%
R	1	1.1%
Mesos	1	1.1%
Vertica	1	1.1%
Experimental	1	1.1%
Number of Responses =	143	
N = 87		

Source: Hyperion Research, 2017

Types of Systems Used for HPDA Applications

Hyperion Research asked respondents whether they run their Big Data analytics jobs today on the same HPC systems where they run other types of workloads (e.g., simulation). As Table 17 shows, three-quarters of the sites (78.2%) said yes to this question, while four in 10 (39.1%) said they employ separate systems. Respondents were allowed to give more than one answer in cases where they

sometimes run Big Data analytics and other applications on the same system, sometimes on separate systems.

Table 17

For your CURRENT Big Data analytics, do you run these on?

	Number of Responses	Percentage of Respondents
The same HPC system used for other types of workloads	68	78.2%
A separate HPC system or Big Data appliance	34	39.1%
N = 87		

Source: Hyperion Research, 2017

We next asked whether the sites plan to run Big Data analytics and other types of applications on the same or separate systems 6-18 months from today. The findings in Table 18 confirm that there is a trend toward using separate systems, compared with what the sites do today (Table 17).

Table 18

IN THE NEXT 6-18 MONTHS, do you plan to run your Big Data on?

	Number of Responses	Percentage of Respondents
The same HPC system used for other types of workloads	59	67.8%
A separate HPC system or Big Data appliance	43	49.4%
N = 87		

Source: Hyperion Research, 2017

Plans for Using Machine Learning, Deep Learning and/or Cognitive Applications

Tables 19 and 20 show survey findings for the HPC sites' current and planned use of machine learning, deep learning and cognitive computing/artificial intelligence (AI). Hyperion Research treats these terms as *methodologies* rather than distinct market segments such as bio-life sciences, defense, or weather/climate. ML, DL and AI are more analogous to methods used in the long-standing HPC modeling and simulation market, such as computational fluid dynamics (CFD) or finite element analysis (FEA). Each of these methodologies: ML, DL, CFD, FEA -- is applicable across multiple market segments. In addition, more than one of these methodologies (e.g., simulation and advanced analytics) may be used in combination to solve an HPC problem.

Our definitions are as follows:

- **Artificial Intelligence (AI):** a broad, general term for the ability of computers to do things human thinking does (but NOT to think in the same way humans think). AI includes machine learning, deep learning (a.k.a. cognitive computing) and more minor methodologies.
- **Machine learning (ML)**: a process where examples are used to train computers to recognize specified patterns, such as human blue eyes or numerical patterns indicating fraud. The computers are unable to learn beyond their training and human oversight is needed in the recognition process.
- **Deep Learning (DL):** an advanced form of machine learning that uses digital neural networks to enable a computer to go beyond its training and learn on its own, without explicit programming or human oversight.

Current Use of ML, DL, AI

As Table 19 shows, a sizeable majority of the respondents (70.1%) said they run machine learning applications today on HPC systems. Substantial minorities run deep learning workloads (42.5%) and AI jobs (37.9%).

Table 19

For your CURRENT Big Data analytics, do you run these types of applications?

	Number of Responses	Percentage of Respondents
Machine learning	61	70.1%
Deep learning	37	42.5%
Cognitive computing/artificial intelligence	33	37.9%
N = 87		

Source: Hyperion Research, 2017

Future Use of ML, DL, AI

Table 20 reveals that the use of ML, DL and AI will increase among the surveyed sites in the next 6-18 months, with each of these methodologies begin exploited by a moderately larger percentage of the sites than is the case today (see Table 19).

Table 20

IN THE NEXT 6 TO 18 MONTHS, do you plan to run these types of applications?

	Number of Responses	Percentage of Respondents
Machine learning	67	77.0%
Deep learning	55	63.2%
Cognitive computing/artificial intelligence	42	48.3%
N = 102		

Source: Hyperion Research, 2017

Plans for Using HPC for IOT Applications

Only about one in seven of the sites (15.6%) uses HPC for Internet of Things (IoT) applications today (Table 21), but that numbers should more than double (36.7%) in the next 6-18 months (Table 22). HPC is already beginning to assume important roles in IoT where dense nodes with high data rates are needed, as well as for network wellness management and security.

Table 21

Are you CURRENTLY using HPC for anything related to the Internet of Things (IoT)?

	Number of Responses	Percentage of Respondents
No	81	84.4%
Yes	15	15.6%
N = 96		

Source: Hyperion Research, 2017

Table 22

IN THE NEXT 6 TO 18 MONTHS, do you expect to use HPC for anything related to the IoT?

	Number of Responses	Percentage of Respondents
No	50	63.3%
Yes	29	36.7%
N = 79		

Source: Hyperion Research, 2017

Plans for Using External Clouds for Big Data/HPDA Applications

Table 23 shows that today only a minority of the sites (41.1%) run any Big Data/HPDA applications in external cloud environments, although about one in seven sites (14.7%) assigns a substantial percentage (25% or more) of their Big Data/HPDA workloads to external clouds.

Table 23

Approximately what percent of all your Big Data application workload is run on external cloud computing TODAY?

	Number of Responses	Percentage of Respondents
0%	56	58.9%
1-9%	11	11.6%
10-24%	14	14.7%
25-50%	6	6.3%
> 50%	8	8.4%
N = 95		

Source: Hyperion Research, 2017

In the next 6-18 months (Table 24), the proportion of the sites running some Big Data/HPDA workloads in external clouds will rise from a minority of 41.1% (Table 22) to a majority of 54.7%. Sites assigning a substantial proportion of their Big Data/HPDA workloads (25% or more) to external clouds will jump from 14.7% today to 23.3%.

Table 24

IN THE NEXT 6 to 18 MONTHS -- Approximately what percent of all your Big Data application workload will be run on external cloud computing?

	Number of Responses	Percentage of Respondents
0%	39	45.3%
1-9%	8	9.3%
10-24%	19	22.1%
25-50%	8	9.3%
> 50%	12	14.0%
N = 86		

Source: Hyperion Research, 2017

Worldwide HPC-based Artificial (AI) Market Revenue Snapshot

Hyperion Research forecasts that the worldwide HPC server-based AI market will expand at a 29.5% CAGR to reach more than $1.26 billion in 2021, up more than three-fold from $346 million in 2016. We define the HPC AI market as a subset of the high-performance data analysis (HPDA) market that includes machine learning (ML), deep learning (DL) and other AI workloads running on HPC servers.

Hyperion Research treats machine learning, deep learning and AI as *methodologies* rather than distinct market segments such as bio-life sciences, defense, or weather/climate. ML, DL and AI are more analogous to methods used in the long-standing HPC modeling and simulation market, such as computational fluid dynamics (CFD) or finite element analysis (FEA). Each of these methodologies: ML, DL, CFD, FEA -- is applicable across multiple market segments. In addition, more than one of these methodologies (e.g., simulation and advanced analytics) may be used in combination to solve an HPC problem.

Interest in ML, DL and other AI work within the HPC community has grown large enough for us to size HPC server spending related to these methodologies and provide a five-year forecast in this document.

Definitions Used in This Forecast

- **Artificial Intelligence (AI):** a broad, general term for the ability of computers to do things human thinking does (but NOT to think in the same way humans think). AI includes machine learning, deep learning (a.k.a. cognitive computing) and more minor methodologies.
- **Machine learning (ML)**: a process where examples are used to train computers to recognize specified patterns, such as human blue eyes or numerical patterns indicating fraud. The computers are unable to learn beyond their training and human oversight is needed in the recognition process.
- **Deep Learning (DL):** an advanced form of machine learning that uses digital neural networks to enable a computer to go beyond its training and learn on its own, without explicit programming or human oversight.

Worldwide HPC AI Server Revenues vs. All HPDA Server Revenues

Figure 3 displays Hyperion Research's five-years forecasts for high performance data analysis (HPDA) server revenues, along with our five-year projection for the subset of revenues for HPDA servers used primarily (>50% of cycles) for ML, DL or other AI workloads. We predict that during the period shown in Figure 1 (2015-2021), HPC server revenue for the whole AI category (ML, DL, et al.) will expand at a 29.5% CAGR to reach $1.2 billion, or about 31% of the $4.0 billion total for all HPDA server revenue.

FIGURE 3

Worldwide Server-Based AI Revenues

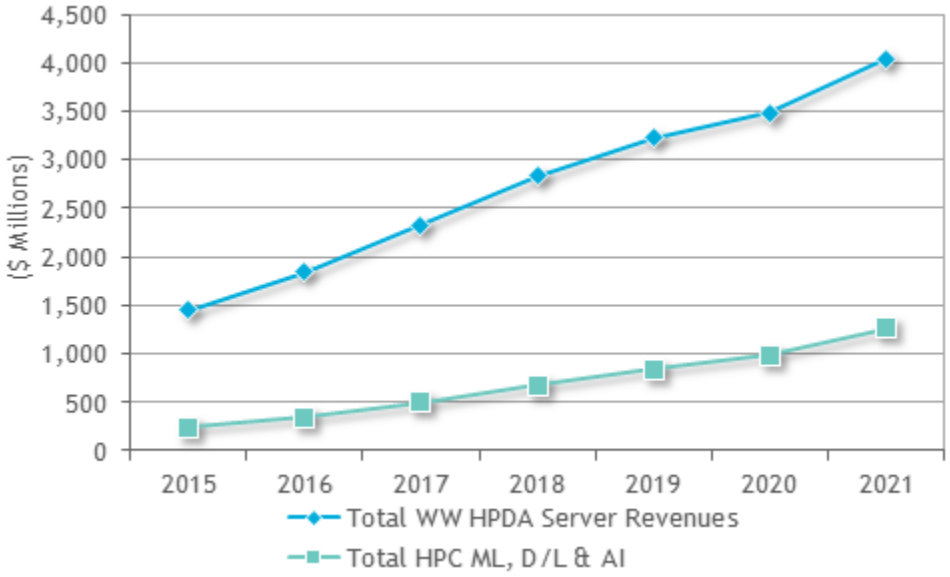

Source: Hyperion Research, 2017

Table 25 shows the revenue figures associated with the Figure 1 graph. As Table 1 indicates, worldwide revenue for the HPC AI server market was $346 million in historical year 2016, representing 18.8% of worldwide HPDA server revenues. We predict that the CAGR for the AI portion of the HPDA server market (29.5%) will be substantially higher than the CAGR for the overall HPDA server market (17.0%) during the forecast period.

Table 25

Worldwide HPC-Based AI Revenues vs Total HPDA Revenues (Millions)

	2015	2016	2017	2018	2019	2020	2021	CAGR 16-21
Total WW HPDA Server Revenues	$1,455	$1,845	$2,333	$2,830	$3,224	$3,488	$4,040	17.0%
Total HPC-Based AI (DL, ML, and Other)	$246	$346	$501	$673	$845	$986	$1,260	29.5%

Source: Hyperion Research, 2017

Worldwide HPC AI Server Market Forecast by Methodology

Figure 4 and Table 26 show Hyperion Research's five-year forecast for worldwide HPC AI server revenues, split out by AI methodology areas.

- We expect DL to grow exceptionally fast from its tiny 2015 starting point ($15 million) to exceed $600 million in 2021 (81.5% CAGR), as the prerequisites for production-grade DL, especially access to adequate data volumes, begin to exist in more HPC-supported markets.
- ML growth should be impressive though less stellar (16.0% CAGR). As DL becomes practical for more use cases in more HPC market segments, some of the most challenging ML workloads will migrate to DL methods while some of the least challenging ML job types trickle down to desktop computers and over time to portable electronic devices.

FIGURE 4

Worldwide M/L, D/L & AI HPC-Based Revenues

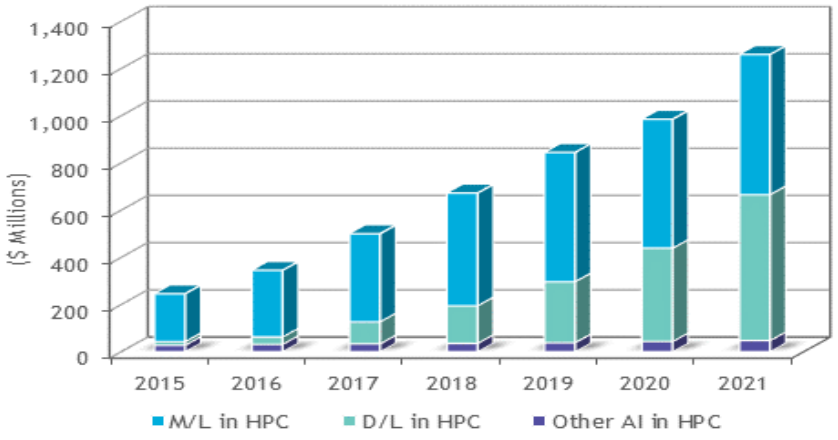

Source: Hyperion Research 2017

Table 26

Worldwide M/L, D/L & AI HPC-Based Revenues ($ Millions)

	2015	2016	2017	2018	2019	2020	2021	CAGR 16-21
M/L in HPC	$203	$282	$373	$478	$548	$546	$594	16.0%
D/L in HPC	$15	$31	$93	$159	$258	$395	$618	81.5%
Other AI In HPC	$28	$32	$34	$36	$39	$44	$48	8.4%
Total	$246	$346	$501	$673	$845	$986	$1,260	29.5%

Source: Hyperion Research, 2017

High Growth is Projected for HPDA Software

Table 27 compares the current budgets for different types of HPC software, along with the expected growth in budgets for the next year. HPDA software currently represents 7.7% of HPC center budgets (for centers that do HPDA work), and is expected to grow to 14.0% over the next year.

Table 27

Software Budgets and Growth Plans

	Current Year	Next Year
System software: OS - Current HPC Budget Allocation %	6.7%	
System software: OS - Next Year % Up or Flat or % Down		3.5%
System SW: all MW, compilers, file system etc. - Current HPC Budget Allocation %	8.8%	
System SW: all MW, compilers, file system, etc. - Next Year % Up or Flat or % Down		6.6%
Applications software - Current HPC Budget Allocation %	14.6%	
Applications software - Next Year % Up or Flat or % Down		4.6%
Big data analysis and/or analytics - Current HPC Budget Allocation %	**7.7%**	
Big data analysis and/or analytics - Next Year % Up or Flat or % Down		**14.0%**
N=99		

Source: Hyperion Research, 2017

IN SUMMARY

HPDA – the market for data-intensive computing ("Big Data") needing HPC resources – dates back to the Cold War era: think of data-intensive simulation for weather forecasting, and data-intensive analytics as used by classified government agencies and by large investment banks. Even as these established uses continue to grow organically, the HPDA market is poised for unprecedented growth, principally on the analytics side, as HPC has moved to the forefront of R&D in machine learning, deep learning, and other early AI methodologies. Because of the technical challenges that must be mastered before HPC servers can be put to production use in this brave new world, HPDA analytics is still mostly a high-growth, exploratory market today. We don't expect the truly sharp ramp-up of this market to begin until after our current forecast period ends in 2021, driven by mainstream, economically important uses such as precision medicine and autonomous vehicle design and networked operation.

To put this in perspective, Hyperion Research forecasts that the worldwide HPC server system market will grow at a healthy 5.8% CAGR (considerably faster than the enterprise server market) to reach a record $14.8 billion in 2021. The broader HPC market that includes servers, storage, middleware, application software and technical support service will amount to about $30.3 billion in that year.

- We project that HPDA subset of the HPC market will expand at a speedier 17.0% CAGR to reach $4.0 billion in 2021, and first-time HPC adopters from the commercial analytics side will represent $1.1 billion of that total.
- A subset of the HPDA server market–machine learning, deep learning and other AI methodologies (we treat them as horizontal methodologies, not vertical segments)–will grow at an even faster pace (25.9% CAGR) to reach about $1.3 billion in 2021.

We project that HPDA subset of the HPC market will expand at a speedier 17.0% CAGR to reach $4.0 billion in 2021

FORECAST METHODOLOGY

The forecasts in this study are based on multiple information sources, including Hyperion Research's technical computing systems quarterly census database, vendor results for the historical years, discussions with vendors and users on future business directions and expectations, end-user studies, and in-depth interviews with users. The multi-client approach used for this study spreads the cost of the research equitably across all clients who purchase one or more of the six reports generated from the study, making the price of the reports affordable to more organizations.

The forecasts were developed based on Hyperion Research's technical computing systems forecast model, which targets compute servers. This model initially considers competitive segments (supercomputers, technical divisional servers, technical departmental servers, and technical workgroup servers), forecasting system unit shipments, revenue, and average sales price by industry/application segment. The forecasts include estimates for second-tier and new-entrant vendors selling into the HPC server market space.

The forecasts provided in this study include only server systems used in technical computing applications. Systems sold into commercial (nontechnical) applications and desktop technical computers are not included in this study.

Note: Numbers in this document may be inexact due to rounding.

Synopsis

This Hyperion Research report is based on the fifth edition of Hyperion Research 's HPC end user-based tracking of end user trends, buying behaviors and future plans. This report focuses on the newer use of big data combined with high computation requires, called HPDA. Hyperion Research forecasts that revenue for HPDA servers will grow robustly (17.3% CAGR), increasing from $1.8 billion in 2016 to $4.0 billion in 20219. HPDA storage revenue will near $2.6 billion by 2021.

According to Steve Conway, Hyperion Research senior vice president of research, "The organic growth of data-intensive simulation and the proliferation of newer analytics methods with real-time requirements, including machine learning, deep learning and cognitive methods, will propel growth in the HPDA market. The analytics side of HPDA is already growing quickly and its growth rate promises to accelerate beyond our current forecast period ending in 2021."

APPENDIX: SITES SURVEYED IN THIS STUDY

Table 28

Appendix: List of Organizations in The Study

Organization	Organization	Organization
Abbott Laboratories	Iowa State University	Peer Technologies GmbH & Co KG
Air Force Research Laboratory	J.P. Morgan	Pfizer
ANV Syndicates	Jefferson Lab	Philly Shipyard
Bank of America	JetBlue	Princeton University
Bank of America Merrill Lynch	John Wiley & Sons	Procter & Gamble
Berkeley Lab	JPL/NASA	Rothschild
Boston College	Kiln Group	RSC Group
Boston University	Kinder Morgan	RWTH Aachen University
Briggs & Stratton	L&L Products	Sandia National Laboratories
Brunel University	Leibniz Supercomputing Center LRZ	Scottish Borders Council
Carrefour Retail	London School of Economics	Shanghai Jiao Tong University
CERFACS	Manchester City Council	Shared Services Canada
CINECA	Marks and Spencer	Singapore Airlines
Citi	Mayo Clinic	Singtel
City of Chicago	Mellanox	Skanska
City of Westminster	Miami University	SUNY
City University of New York	Michigan State University	Tesco
CNA	MIT / CSAIL	TUI Europe
DataSwing	MITRE	Turner Construction
Dept. Of Defense	NASA Goddard Space Flight Center	Universität Hamburg, Scientific Computing (Informatics)
DFJ	NASA Langley Research Center	Universiti Putra Malaysia
Diagenode	Nationwide	University of Coimbra

Table 28

Appendix: List of Organizations in The Study

Organization	Organization	Organization
DKRZ GmbH (German Climate Computing Centre)	Navy Research Laboratory	University of Edinburgh
Dubai Trade (Government)	NCAR/UCAR	University of Liverpool
Emaar	Neoanyang Technological University	University of Maryland
EverBank	Netflix	University of Minnesota
eXact-lab srl	New York Life Insurance Company	University of Tennessee
Frobozz Magic	NIIT Technologies	University of Tennessee
GigaBack	NOAA	University of Wales
Grayhill	Northwestern University	University of Warwick
Hammersmith and Fulham Council	Nous Infosystems	UPS
Hewlett Packard Enterprises	NYC Department Of Transportation	Vanderbilt University - ACCRE
HLRS	Oak Ridge National Laboratory	Wells Fargo
HSBC	Onx	Western Michigan University
Hughes Network Systems LLC	Pacific Northwest National Laboratory	Western Oregon University
Institute Process Engineering, Chinese Academy of Sciences	Partnership	WWU Münster
International Academy, Research, and Industry Association - Advanced Computing	PDC	XPO Logistics
N = 111		

Source: Hyperion Research, 2017

About Hyperion Research, LLC

Hyperion Research, consisting of the former IDC high performance computing (HPC) analyst team, provides HPC information, analysis, and recommendations based on technology and market trends. 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.hpcuserforum.com and www.HPCatHyperion Research.com
