HYPERION RESEARCH

Taxonomy

Hyperion Research Worldwide Taxonomy for HPC Server Tracking and Application Workload Segments, 2021

Melissa Riddle, Bob Sorensen, Mark Nossokoff, Earl Joseph, Alex Norton, and Steve Conway
March 2021

HYPERION RESEARCH OPINION

This taxonomy provides an overview of the market definitions of the technical computing server market industry segments and sub-segments. It is used by Hyperion Research analysts to generate market sizing, forecasts, and company models. Vendors and users in the technical computing space will find this taxonomy useful in understanding how Hyperion Research defines and measures the markets they serve or are considering entering. This document provides relevant definitions, tracking methodology, accounting rules, market segmentation, and workload categories.

Hyperion Research endeavors to track all HPC servers sold worldwide each quarter, along with relevant details pertaining to each system and product line. Areas of segmentation include:

- Competitive segments based on server price
- Broader market segments encompassing both servers and other HPC revenue sources
- Geographic segments by region
- Operating system segments based on the server's OS at time of shipment

This document only addresses the high-performance computing (HPC) or technical computing server portion of the overall server market, highlighting the following:

- Definitions and terminology used to track the technical computing market segments, as used in the HPC QView and HPC Forecast Database.
- Measurement methodologies.
- Definitions for HPC technical industry/application workload categories.

Note: This page is intentionally blank.

TABLE OF CONTENTS

IN THIS REPORT

This Hyperion Research report provides an overview of key technology and market definitions for the technical computing market (also referred to as the HPC market). Vendors and users in the technical computing space will find this taxonomy useful in understanding how Hyperion Research defines and measures the markets they serve or are considering entering. This taxonomy provides technology definitions, HPC QView tracking categories, and industry/application/workload categories as used in the technical computing market.

DEFINITIONS

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. An on-premise system primarily used for HPC workloads (at least 50%) can be referred to as an HPC system, or simply an HPC. Systems acquired by cloud service providers for the purpose of hosting cloud workloads are excluded, as Hyperion Research separately tracks spending for HPC usage in cloud environments.

Technical servers range from small servers costing less than $10,000 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 HPC users and major hardware original equipment manufacturers (OEMs) in the technical computing space to gather information on each vendor's quarterly sales. Specifically, Hyperion Research 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 processors per system, average number of nodes for each system, system revenue by geographical regions, and system revenue 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 Research refers to this data structure as the "HPC QView." In addition to the HPC QView, Hyperion Research 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

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") or major upgrades to existing systems. An ISS unit consists of processors, memory, embedded disk storage, cluster interconnect hardware/software, any bundled operating system(s), compiler(s), 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 multiple 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 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
- Any additional sales revenue from the channel. Hyperion Research does not count revenue coming from the channels and partners of the original equipment manufacturers (OEMs) or original system integrators (OSIs). All revenue Hyperion Research counts are based on direct sales from the OEMs and OSIs.

Accounting for Exascale and Pre-Exascale Systems

Exascale and pre-exascale systems typically have high costs and often involve a large portion of non-recurring engineering (NRE) in the server contract. As with other HPC servers, Hyperion Research recognizes the entire cost of the server contract as server revenue.

Hyperion Research's accounting rules record these server revenues as a lump sum at the time of system acceptance (as indicated by the purchaser), regardless of when actual payments may have been made. This large dollar amount accepted all at once may give the appearance of extraordinary growth or decline in a single quarter or year. Whenever an exceptionally costly system is accepted in a particular period, this will be clearly noted.

When necessary (such as when multiple exascale systems are accepted in the same period), these exceptionally large systems may be pulled out as a separate line item within the supercomputer segment to preserve forecasting of the general HPC market.

"Standard" Server Revenues and Processor Package Counts

"Standard" server revenues and "standard" processor package counts are provided at client request to facilitate direct comparisons with databases and reports from other sources. In particular, these figures have been adjusted to follow more closely IDC accounting rules for all servers instead of the usual Hyperion Research accounting rules. The "standard" numbers are always lower since IDC does not count certain HPC systems and some HPC revenue. For example, the Fujitsu Fugaku $1 billion system is not in the IDC numbers; NUDT in China builds and sells computers but IDC does not count them because they are a university; many supercomputers require additional NRE engineering, which is part of the sales price, and IDC does not include these revenues.

Compound Annual Growth Rate (CAGR)

In general, Hyperion Research reports predicted forecast growth rates as five-year compound annual growth rates (CAGR). The CAGR is an average annual growth rate, measured as a percentage. If each year of the forecast had the CAGR as its annual growth rate, the value predicted at the end of the forecast period would be reached exactly. The CAGR is calculated based on the first and last value of the forecast period only and is not directly affected by the values for intermediate years.

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

HPC Broader Market Segmentation

Hyperion Research's quarterly tracking and related data structures are focused on servers. The broader HPC market is also tracked and forecasted by the following categories:

- **Storage:** Additional storage purchased apart from the original procurement of a server

- **Middleware:** Software licensing fees for middleware such as compilers and schedulers
- **Applications:** Licensing fees paid for the use of proprietary application software
- **Service:** Repair or maintenance services for HPC systems. It does not include general professional services

Artificial Intelligence, Machine Learning, and Deep Learning

Artificial Intelligence (AI), Machine Learning (ML), and Deep Learning (DL) are specific terms describing the three main AI application types.

- **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, and other 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. The computer follows the base rules given to it.
- **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 additional explicit programming or human oversight. The computer develops its own rules.

AI, ML, and DL applications often use HPC servers. Hyperion Research tracks the subset of HPC server revenues generated by systems that are primarily used for these application types. These figures are also encompassed within the total HPC server data.

Technical Computing Taxonomy for HPC QView

Data Elements in the HPC QView

- **Manufacturer.** Identification of the technical computing server OEM.
- **Model.** The name of a technical computer. This is the standard model name used by the OEM for a product line. In many cases, Hyperion Research will combine multiple product models into one category in order to simplify the data collection process and to provide a common category that lasts over time, e.g., model 510/520/540 are combined into one set.
- **Competitive segment.** A category based on the average selling price of a technical server system (For detailed definitions refer back to the Technical Computing Market Segmentation section.)
- **Price band.** A price range to categorize the average selling price of a technical server system.

Hyperion Research uses the following 12 price bands to classify a technical server system:

- B01 – ≤ $2,999
- B02 – $3,000-$5,999
- B03 – $6,000-$9,999
- B04 – $10,000-$24,999
- B05 – $25,000-$49,999
- B06 – $50,000-$99,999
- B07 – $100,000-$249,999
- B08 – $250,000-$499,999

    - B09 – $500,000–$999,999
    - B10 – $1.0 million–$2.9 million
    - B11 – $3.0 million–$9.9 million
    - B12 – $10.0+ million
- **Processor package.** A processor package, or processor, is the actual item that is sold by companies like Intel and AMD. A processor package is not an individual core.
- **Average processor package count per system.** Average number of processor packages per system for a particular technical server model. Note that the average processor package count per system, times the number of systems sold, produces a metric closely equivalent to total "socket count shipped" in the Hyperion Research server tracker. Note that a HPC "system" refers to the full cluster complex and not just a node.
- **CPU type.** Hyperion Research categorizes the technical server market based on the design of the CPU and the instruction set. The main CPU types are X86, RISC, Vector, and Proprietary (Prop). Hyperion Research will add a new type of CPU as new designs come to the market. Note that CPU type is referring to the processor package and not the core.
- **CPU brand.** Brand of CPU within each CPU type. Hyperion Research uses the processor model name as the "brand" in this data field. For example, for CPUs of x86 type there are the Xeon brand and the Opteron brand today. Note that CPU brand is referring to the processor package and not the core.
- **Cores per processor.** Number of cores in each processor package for a technical server system.
- **Accelerators and co-processors.** Accelerators and co-processors, such as GPUs, coprocessors, and FPGAs, are specialized chips that improve system performance by supplementing the main system processors (CPUs).
- **Node.** In technical computing terms, a node is a sub-part of a server system that often contains a full copy of the OS, for example, a blade in a cluster. A node has one or more processor sockets that are attached to its own local memory, often on the same board. A node as defined in HPC terms is equivalent to a Server Unit in Server Tracker terms.
- **Average system node count.** Average number of nodes for a technical server system. The HPC node count is an estimated data field calculated based on the system architecture.
- **Unit.** Total number of technical server models sold within a given quarter. For a detailed definition, refer back to the Average Selling Price subsection under the section Technical Computing Server Revenue Accounting Rules section. Note that Hyperion Research defines a system "unit" as the full server (e.g., a cluster with 2,000 nodes is considered one unit). In the Hyperion Research overall server tracker, a system unit equals 1 node. In the HPC databases the terms "Units", "Systems", and "System Units" all refer to the same number, the full cluster complex or full computer system.
- **Revenue.** Total factory revenue of a technical server model sold within a given quarter.
- **Architecture.** Describes the configuration of the HPC server.
    - Cluster: A collection of independent nodes joined by an interconnect. The most popular architecture today.
    - SMP (symmetric multiprocessing): Multiple identical processors with a single shared memory subsystem and a single operating system. An "SMP-4 Way" has four processors per server. SMPs with different numbers of processors are named similarly.

- MPP (massively parallel processing): Similar in structure to a cluster with multiple nodes, each with 1 processor.

Technical Computing Geographic Segmentation

Hyperion Research segments the technical server market by the following geographic regions:

- North America (including the United States, Canada, and Mexico)
- EMEA (including Western, Central and Eastern Europe, the Middle East, and Africa)
- APAC (including Eastern Asia, Southern and Southeastern Asia, and Oceania)
  - In addition, APAC without Japan, and Japan separated
- Rest of the World (mainly Latin America)

In addition, as a separate data structure, Hyperion Research tracks 26 different countries and regions by server revenues, processors shipped, application workloads, and broader market categories.

Technical Server Operating System Segmentation

Hyperion Research assumes that all servers are shipped with an operating system. A portion of the license fee for the operating system shipped with the server is included within the factory-revenue figures. The fee is only associated with new operating systems licenses and not with installed-base licenses.

Hyperion Research tracks the following operating systems by model:

- Linux
- Unix (excluding Linux variants)
- Windows (all types)
- Other

TECHNICAL COMPUTING APPLICATION WORKLOAD DATA STRUCTURE

HPC Application/Industry Workload Categories

Hyperion Research identifies and tracks the following technical computing workload segments (note that this is a separate data structure from the HPC QView):

- **Biological sciences.** This workload centers on applications such as genomics, proteomics, pharmaceutical research, bioinformatics, drug discovery, bioanalytic portals, and agricultural research. Computational techniques include database searching and management, molecular modeling, and computational chemistry. These workloads appear in commercial, academic, and institutional research environments. Systems that are specifically targeted for these workloads should be included; systems purchased for more general scientific and R&D environments should be counted in the university and academic, national laboratories and research centers, or national defense segments.
- **Chemical engineering.** This workload centers on applications such as molecular modeling, computational chemistry, process design, and chemical analysis. It includes all chemistry applications that are not directly related to biosciences research and development. These workloads appear in commercial, academic, and institutional research environments.

- **Computer-aided engineering (CAE) and product design.** This workload centers on applications such as finite element modeling and analysis, mechanical computer-aided engineering, civil engineering, structural analysis, computation fluid dynamics (CFD), crash testing, NVH, and solid modeling. Like CAD applications, these CAE tasks are used to design automobiles, aircraft, running shoes, ski equipment, sailboards, beer bottles, and other everyday items. Workloads include those tasks generally accomplished by engineers, not drafters.
- **Mechanical design**. Focuses on computer-aided design (CAD) and drafting. This workload centers on applications such as mechanical computer-aided design; 2D, 2.5D, and 3D design and drafting; and 3D wire frame. Design and drafting applications require graphics capability but are less compute intensive than design engineering and analysis applications (see "CAE: Product Design" above). CAD tasks are typically performed by designers and drafters and are often used in conjunction with CNC machines. Users are found primarily in manufacturing industries such as automotive, aerospace, heavy machinery, and consumer goods.
- **Digital content creation and distribution (DCC&D).** This workload category centers on applications such as 2D and 3D animation, film and video editing and production, and multimedia authoring for both CD and Web pages that utilize sophisticated graphics content. This category also includes servers used for image rendering, content management, and distribution of finished products for areas such as film, TV, commercial animation, advertising, product styling, and industrial design as well as servers used for large-scale games. These workloads are developed in large part in concert with scientific visualization research and technologies. In addition, the creation of special effects and animation for motion pictures requires significant amounts of computational capacity.
- **Economic and financial modeling**. This workload centers on applications such as econometric modeling, portfolio management, stock market and economic forecasting, and financial analysis. The segment includes both trader and computationally intensive non-trader tasks. In this case, the workload uses technical computing because of the numerically intensive requirements of most applications and their association with economic modeling and simulation-based research.
- **Electronic design and analysis/IT (EDA/IT/ISV).** This workload area covers all electrical/electronic tasks, including schematic capture, logic synthesis, circuit simulation, PCB routing, and system modeling. It also includes the use of technical servers within IT manufactures for R&D, system development and testing, application development, software development, and other product design and testing. Systems used by independent software vendors for software development are also included.
- **Geosciences and geoengineering.** This workload includes earth resources–related applications such as seismic analysis, oil services, and reservoir modeling. These applications are used in both institutional research and commercial enterprises. Geosciences can also include areas such as mining, natural resource management, geographic information systems (GIS), and mapping.
- **Government laboratories and research centers.** This workload centers on government-funded research and development institutions. These organizations are generally funded at a state, national, or multinational level and may combine purely scientific research with research in areas of national priority (e.g., cancer research) and/or research for defense-related programs. These users are less bound by strict economic constraints than those in commercial development environments. These centers normally do not offer degree programs for students. Universities that receive major government funding are included under academic/university.

- **National defense.** This workload centers around applications such as surveillance and signal processing; encryption; command, control, communications, and intelligence (C3I); geospatial image management and analysis; defense research; weapons design; and other national security applications. In addition, national security organizations may field applications that work to identify and track potential security threats through database-oriented pattern-matching applications. Although these applications may not always be numerically intensive, they will be developed and used by organizations that are firmly rooted in technical computing markets. In many cases, these applications will be run in conjunction with traditional security applications such as cryptography and image analysis.
- **Academic/University.** This workload centers on scientific research and engineering R&D efforts conducted at public or private institutes of higher education and includes systems sold for both research and educational activities. Privately funded and/or nonprofit research institutes that have a strong academic mission (i.e., work to extend the bounds of public knowledge) are also included in this segment. Applications are typically compute or data intensive and often require high-performance graphics. These users are less bound by strict economic constraints than those in commercial development environments. This segment includes NSF sites that are located at universities.
- **Weather forecasting and climate modeling.** This workload centers on applications such as atmospheric modeling, meteorology, weather forecasting, and climate modeling. This segment includes systems dedicated to these tasks primarily in the government and defense segments.
- **Other.** This segment includes any technical computing workloads not otherwise specified by the previous definitions.

HPC Application/Industry Workload Sub-Segment Categories

Bio-Sciences

This workload centers on applications such as genomics, proteomics, pharmaceutical research, bioinformatics, drug discovery, bio-analytic portals, and agricultural research. Computational techniques include database searching and management, molecular modeling, and computational chemistry. These workloads appear in commercial, academic, and institutional research environments. Systems that are specifically targeted for these workloads should be included; systems purchased for more general scientific and R&D environments should be counted in the university and academic, national laboratory and research center, or national defense segments.

- **Genomics**: a branch of biotechnology concerned with applying the techniques of genetics and molecular biology to the mapping and DNA sequencing of sets of genes or the complete genomes of selected organisms, with organizing the results in databases, and with applications of the data (as in medicine or biology). HPC clusters are often used to process the results of purpose-built gene sequencing machines.
- **Proteomics**: a branch of biotechnology concerned with applying the techniques of molecular biology, biochemistry, and genetics to analyzing the structure, function, and interactions of the proteins produced by the genes of a particular cell, tissue, or organism, with organizing the information in databases, and with applications of the data. HPC systems are often used to process the associated large data sets.
- **Drug Discovery**: In the fields of medicine, biotechnology and pharmacology, drug discovery is the process by which new candidate medications are discovered. Historically, drugs were discovered through identifying the active ingredient from traditional remedies or by serendipitous discovery. HPC systems are used to quickly test large numbers (often millions) of drug candidates (small molecules) for lock-and-key fit with disease agents.

- **Bioinformatics**: the application of computer technology to the management of biological information. Computers are used to gather, store, analyze and integrate biological and genetic information, which can then be applied to gene-based drug discovery and development.
- **Agricultural Research**: the application of computer technology to the science of optimizing agricultural production and quality.
- **Epidemiology/Public Health:** the study of the etiology and spread of disease outbreaks that affect, or threaten to affect, large segments of a population. Public health is a general term for the study of phenomena that affect the health of a human population.
- **Precision Medicine:** a medical model based on customizing healthcare, with medical decisions, practices, and products tailored to the individual patient rather than on a set of standardized procedures applicable to all 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.

CAE: Product Design

This workload centers on applications such as finite element modeling and analysis, mechanical computer-aided engineering (CAE), civil engineering, structural analysis, computation fluid dynamics (CFD), crash testing, NVH, and solid modeling. Like CAD applications, these CAE tasks are used to design automobiles, aircraft, running shoes, ski equipment, sail boards, beer bottles, and other everyday items. Workloads include those tasks generally accomplished by engineers, not drafters.

- **Crashworthiness:** the ability of a structure to protect its occupants during an impact. This is commonly tested when investigating the safety of aircraft and vehicles.
- **Structural Analysis:** the determination of the effects of loads on physical structures and their components. Structures subject to this type of analysis include all that must withstand loads, such as buildings, bridges, vehicles, machinery, furniture, attire, soil strata, prostheses, and biological tissue.
- **Fluid-Structure Analysis:** the determination of the combined effects of loads and the fluid flows (air, water, etc.) on physical structures and their components.
- **Noise/Vibration/Harshness:** the study and modification of the noise and vibration characteristics of vehicles, particularly cars and trucks. While noise and vibration can be readily measured, harshness is a subjective quality and is measured either via "jury" evaluations, or with analytical tools that can provide results reflecting human subjective impressions.
- **Environmental Friendliness:** a term that refers to goods and services, laws, guidelines and policies that inflict reduced, minimal, or no harm upon ecosystems or the environment. HPC systems are routinely used to improve the environmental friendliness of vehicle designs.
- **Materials Science:** the scientific study of the properties and applications of existing and new materials of construction or manufacture (such as ceramics, metals, polymers, and composites).

Mechanical Design

Focuses on computer-aided design (CAD) and drafting. This workload centers on applications such as mechanical computer-aided design; 2D, 2.5D, and 3D design and drafting; and 3D wire frame. Design and drafting applications require graphics capability but are less compute intensive than design engineering and analysis applications (see "CAE: Product Design" above). CAD tasks are typically performed by designers and drafters and are often used in conjunction with CNC machines. Users are

found primarily in manufacturing industries such as automotive, aerospace, heavy machinery, and consumer goods.

- **CAD and Drafting:** the creation and analysis of an edge or skeletal representation of a real-world object in 2D, 2.5D and 3D.
- **3D Wireframe:** the analysis of an edge or skeletal representation of a real-world object. 3D wireframe models consist of points, lines, arcs, circles, and other curves that define the edges or center lines of objects.
- **Storage and Access of Mechanical Designs:** the use of HPC servers to store and access mechanical designs.

Chemical Engineering

This workload centers on applications such as molecular modeling, computational chemistry, process design, and chemical analysis. It includes all chemistry applications that are not directly related to biosciences research and development. These workloads appear in commercial, academic, and institutional research environments.

- **Molecular Modeling:** a technique for deriving, representing, and manipulating the structures and reactions of molecules and those properties that are dependent on these three-dimensional structures.

DCC (Digital Content Creation) & Distribution

This workload category centers on applications such as 2D and 3D animation, film and video editing and production, and multimedia authoring for both CD and Web pages that utilize sophisticated graphics content. This category also includes servers used for image rendering, content management, and distribution of finished products for areas such as film, TV, commercial animation, advertising, product styling, and industrial design as well as servers used for large-scale games. These workloads are developed in large part in concert with scientific visualization research and technologies. In addition, the creation of special effects and animation for motion pictures requires significant amounts of computational capacity.

- **3D Animation/Special Effects:** the production of 3D animated video content or post-production special effects such as green screens or explosions. Production-quality effects often must be very detailed and require significant computing resources to achieve the desired resolution.
- **Film Editing/Production:** editing raw footage into a finished product such as combining videos from different camera angles or adjusting frame rate, color, lighting, or sound.
- **Advanced Rendering:** the generation of images from a model. This can be compute-intensive at high levels of resolution when taking into account details such as texturing, lighting, shading, and so on.
- **Film/Video Distribution:** the distribution of media such as streaming of movies, television shows, or short video clips. HPC may be used for the storage and real-time access of such media.
- **Advanced Gaming with Central Servers:** video games that rely on central servers for synchronization between players playing remotely, such as massively multiplayer online (MMO) games.

Defense

This workload centers around applications such as surveillance and signal processing; encryption; command, control, communications, and intelligence (C3I); geospatial image management and analysis; defense research; weapons design; and other national security applications. In addition, national security organizations may field applications that work to identify and track potential security threats through database-oriented pattern-matching applications. Although these applications may not always be numerically intensive, they will be developed and used by organizations that are firmly rooted in technical computing markets. In many cases, these applications will be run in conjunction with traditional security applications such as cryptography and image analysis.

- **Surveillance/Signal Processing:** surveillance is the close observation and tracking of the behavior of suspect individuals or groups. Signal processing is an enabling technology that encompasses the fundamental theory, applications, algorithms, and implementations of processing or transferring information contained in many different physical, symbolic, or abstract formats broadly designated as signals.
- **Encryption:** the process of encoding a message so that it can be read only by the sender and the intended recipient. Decryption is the process of reading the message by applying the secret code.
- **Communications Intelligence:** information gathered from the communications of individuals, including telephone conversations, text messages, and various types of online interactions.
- **Anti-Terrorism:** activity aimed at identifying and counteracting sources of manifest or potential terrorism.

Economics/Financial

This workload centers on applications such as econometric modeling, portfolio management, stock market and economic forecasting, and financial analysis. The segment includes both trader and computationally intensive non-trader tasks. In this case, the workload uses technical computing because of the numerically intensive requirements of most applications and their association with economic modeling and simulation-based research.

- **Portfolio Optimization**: the process of choosing the proportions of various assets to be held in a portfolio in such a way as to make the portfolio better than any other according to some criterion. The criterion will combine, directly or indirectly, considerations of the expected value of the portfolio's rate of return as well as of the return's dispersion and possibly other measures of financial risk.
- **Pricing Exotic Instruments:** a derivative that is more complex than commonly traded "vanilla" products. This complexity usually relates to determination of payoff. The category may also include derivatives with a non-standard subject matter (i.e., underlying) developed for a particular client or a particular market.
- **Global Risk Management:** management of the worldwide financial risk incurred at any point in time by a company that trades stock, bonds, or other financial instruments, based on calculating the aggregate risk incurred by all of the firm's traders.

EDA (Electronic Design and Analysis)/IT/ISV

This workload area covers all electrical/electronic tasks, including schematic capture, logic synthesis, circuit simulation, PCB routing, and system modeling. It also includes the use of technical servers within IT manufacturers for R&D, system development and testing, application development, software

development, and other product design and testing. Systems used by independent software vendors for software development are also included.

- **Schematic Capture:** the creation of an electronic design model, called a schematic capture, that represents an electronic circuit.
- **Logic Synthesis:** the stage of electronic circuit design where a model called a register transfer level (RTL) is converted into an optimized logic gate representation.
- **System Modeling/Testing:** the production and testing of system models, such as functional or architectural models.
- **ISV and Other Software:** the development, testing, and commercial sale of application software.
- **System Vendors/OEMs:** the development and testing of new computer systems.
- **CSP (Cloud Service Providers):** only computers used for their internal operations, such as R&D.
- **IOT Devices and Systems:** a network of "things" (e.g., devices, sensors) capable of collecting, transmitting, and/or processing data to accomplish a task, usually needed in near-real time.
- **Other Types of IT & Telecom:** other workloads not listed above, such as the processing and transmission of data over the internet.

Geosciences

This workload includes earth resources-related applications such as seismic analysis, oil services, and reservoir modeling. These applications are used in both institutional research and commercial enterprises. Geosciences can also include areas such as mining, natural resource management, geographic information systems (GIS), and mapping.

- **Seismic Processing:** processing seismic data to reduce noise and enhance signal. Aims to improve the quality of the data to be used for interpretation.
- **Reservoir Modeling:** 3D modeling of a reservoir, such as petroleum, that takes into account the physical and chemical properties of both the material in the reservoir and the earth surrounding it. Such models are used to predict future production of the reservoir and inform development decisions.
- **Other Types of Earth Modeling:** includes all types of earth modeling except reservoir, weather, and climate modeling. Environmental agencies like the EPA also use HPC to develop a variety of models for assessing and predicting earth-related public health risks such as air quality or pathogens on recreational beaches.

Government Lab

This workload centers on government-funded research and development institutions. These organizations are generally funded at a state, national, or multinational level and may combine purely scientific research with research in areas of national priority (e.g., cancer research) and/or research for defense-related programs. These users are less bound by strict economic constraints than those in commercial development environments. These centers normally do not offer degree programs for students. Universities that receive major government funding are included under academic/university.

- **Scientific Research:** pursuing advances in fields of science – this is the main activity of most government labs.

- **Industrial Partnering:** refers to collaborations between government labs and private sector HPC vendors or users.

Academic/University

This workload centers on scientific research and engineering R&D efforts conducted at public or private institutes of higher education and includes systems sold for both research and educational activities. Privately funded and/or nonprofit research institutes that have a strong academic mission (i.e., work to extend the bounds of public knowledge) are also included in this segment. Applications are typically compute or data intensive and often require high-performance graphics. These users are less bound by strict economic constraints than those in commercial development environments. This segment includes NSF sites that are located at universities.

- **Scientific Research:** research into any scientific discipline using HPC. A primary purpose of most academic institutions. Oftentimes a single university conducts research into many different scientific disciplines simultaneously through different departments. This includes professors' and students' use of HPC.
- **Industrial Partnering:** a joint partnership between university and industry site that utilizes the university's HPC resources. Such partnerships combine R&D efforts of both parties for mutual benefit.

Weather/Climate

This workload centers on applications such as atmospheric modeling, meteorology, weather forecasting, and climate modeling. This segment includes systems dedicated to these tasks primarily in the government and defense segments.

- **Climate Research:** the use of HPC to study the dynamics and effects of past and ongoing climate change, and to predict future change.

Other

This segment includes any technical computing workloads not otherwise specified by the previous definitions. Retail and marketing is one of the main application workloads included in this category, as is transportation and logistics.

- **Transportation and Logistics:** transportation and logistics tasks such as package routing, traffic management, pattern recognition, and linear programming.
- **Retail and Marketing:** a variety of retail-related applications such as affinity marketing, targeted advertising, and prediction of future sales.
- **New Non-Traditional Applications:** new applications in traditionally non-HPC enterprise space, such as business intelligence.

FUTURE OUTLOOK

The Hyperion Research technical computing taxonomy will be updated regularly to reflect the changing nature of the marketplace. Hyperion Research encourages and welcomes any suggestions and recommendations from vendors and users in technical computing to help us better define and size the 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 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 and www.hpcuserforum.com
