INDUSTRY DEVELOPMENTS AND MODELS

IDC's Worldwide HPC (Technical Server) Taxonomy, 2015

Earl C. Joseph, Ph.D. Steve Conway
Robert Sorensen

IDC OPINION

This taxonomy provides an overview of key technology and market definitions of the technical computing market. It is used by IDC analysts to generate IDC market sizing, forecasts, and company models in the technical computing space. Note that IDC uses a separate taxonomy for the overall server market and has a separate workloads taxonomy for the overall server market with a focus on the non-HPC segments. This document only addresses the high-performance computing (HPC) or technical server portion of the overall server market, highlighting the following:

- IDC's definitions and terminology used to track the technical computing market from the supply side (This market data is published in IDC's HPC QView.)
- IDC's measurement methodologies
- IDC's definitions for HPC technical industry/application workload categories

March 2015, IDC #255159

IN THIS STUDY

This IDC study 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 IDC defines and measures the markets they serve or are considering to enter. This taxonomy provides technology definitions and industry/application/workload categories as used in the technical computing market, as well as IDC's tracking methodology.

SITUATION OVERVIEW

Definition of Technical Computing

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

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

Tracking Methodology

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

IDC records all of the previously 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.

IDC 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. IDC refers to this data structure as the "HPC QView." In addition to the HPC QView, IDC maintains other HPC technical computing data structures such as:

- 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

IDC Technical Computing Server Revenue Accounting Rules

Initial System Shipment

Initial system shipment (ISS) characterizes the first sale of a system (previously referred to as a "new footprint"), and it includes major upgrades to existing systems. An ISS unit is made up of processors, memory, embedded disk storage, cluster interconnect hardware/software, any bundled operating system, a compiler, a math/statistical library, parallel computing, a 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). An ISS is recognized as a shipment only when the complete system or cluster is installed and accepted. In addition, major upgrades that include processors are treated as an ISS in the quarter that it is accepted. External user storage and all paid services are excluded from the ISS revenue value. If a system is paid for over a number of quarters, for example, via service or R&D contracts, IDC determines a value for the whole system when it is finally accepted by the buyer.

Average Selling Price

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

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

Neither ISS nor ASP should include the following:

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

IDC Technical Computing Market Segmentation

Based on input from HPC vendors and end users, IDC 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

Over the past two years, the worldwide HPC market has gravitated toward the high end of the segmentation, with supercomputers representing more than 50% of the market in 2012. In response to these developments, IDC held a meeting with key stakeholders to work toward redefining the segmentations. IDC is currently testing different segmentations, especially at the higher end. Any updates to the HPC segmentations will be effectively communicated across the HPC community.

Technical Computing Taxonomy for HPC QView

Data Elements in the HPC QView

- **Manufacturer.** Manufacturer identifies the technical computing server OEM.
- **Model.** Model refers to the name of a technical computer or the standard model name used by the OEM for a product line. In many cases, IDC combines multiple product models into one category to simplify the data collection process and provide a common category that lasts over time (e.g., models 510, 520, and 540 are combined into one set).
- **Competitive segment.** The competitive segment category is based on the average selling price of a technical server system. (For detailed definitions, refer back to the IDC Technical Computing Market Segmentation section.)
- **Price band.** A price band is a price range categorizing the average selling price of a technical server system. IDC 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 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.** The average processor package count per system refers to the 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 IDC Server Tracker. Also, note that a HPC "system" refers to the full cluster complex and not just a node.
- **CPU type.** IDC categorizes the technical server market based on the design of the CPU and the instruction set. The main CPU types are EPIC, RISC, x86, vector, proprietary (prop), and cell. IDC will add new types of CPUs as new designs come to the market. Note that CPU type refers to the processor package and not the core.
- **CPU brand.** CPU brand refers to the brand of CPU within each CPU type. IDC uses the processor model name as the "brand" in this data field (e.g., for x86 type CPUs, there are the Xeon brand and the Opteron brand). Note that CPU brand refers to the processor package and not the core.
- **Cores per processor.** Cores per processor refers to the number of cores in each processor package for a technical server system.
- **Node.** In technical computing terms, a node is a subpart of a server system that often contains a full copy of the OS (e.g., 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 IDC Server Tracker terms.
- **Average system node count.** Average system node count refers to the 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.** Unit refers to the total number of technical server models sold within a given quarter. (For a detailed definition, refer back to the Average Selling Price subsection in the IDC Technical Computing Server Revenue Accounting Rules section.) Note that IDC defines a system "unit" as the full server (e.g., a cluster with 2,000 nodes is considered one unit). In the overall IDC 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.** Revenue refers to the total factory revenue of a technical server model sold within a given quarter.
- **Average selling price.** Average selling price refers to the ASP of a technical server model for a given quarter. (For a detailed definition, please refer back to the Average Selling Price subsection in the IDC Technical Computing Server Revenue Accounting Rules section.)

Technical Computing Geographic Segmentation

IDC segments the technical server market by the following geographic regions:

- North America (including the United States, Canada, and Mexico)
- EMEA (including Western Europe, Central and Eastern Europe, the Middle East, and Africa)
- Asia/Pacific
- Japan
- The rest of the world (mainly Latin America)

Technical Server Operating System Segmentation

IDC assumes that all servers are shipped with an operating system. A portion of the operating system license fee that is 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.

Currently, IDC tracks the following operating systems by model:

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

Technical Computing Application Workload Data Structure

Note that this section refers only to the IDC HPC workloads segmentation. A separate IDC server workloads taxonomy provides a very different view of the total server market.

HPC Application/Industry Workload Categories

IDC 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, ASP-type service providers, 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 design (CAD) and drafting.** This workload centers on applications such as mechanical computer-aided design; 2D, 2.5D, and 3D design and drafting; 3D wire frame; and civil engineering design. Design and drafting applications require graphics capability but are less compute intensive than design engineering and analysis applications. CAD tasks are typically done by designers and drafters. Users are found primarily in discrete manufacturing industries such as automotive, aerospace, heavy machinery, and consumer goods.
- **Computer-aided engineering (CAE) and mechanical design and analysis.** 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, 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.
- **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 nontrader tasks. In this case, we placed this workload in technical computing because of the numerically intensive applications of most applications and their association with economic modeling and simulation-based research.
- **Electronic design and analysis/IT (EDA/IT).** 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.
- **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 national or multinational level and may combine both 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 performing applications in product development environments. These centers don't normally offer degree programs for students.
- **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, we believe that national security organizations are fielding 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 addition, we believe that these applications will be run in conjunction with traditional security applications such as cryptography and image analysis.
- **University and academic.** 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 performing applications in product 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.
- **Upcoming segments.** A set of high-performance data analysis (HPDA) workload segments will be coming soon.

FUTURE OUTLOOK

IDC technical computing taxonomy will be updated regularly to reflect the changing nature of the marketplace. IDC also encourages and welcomes any suggestions and recommendations from vendors and users in technical computing to help us better define and size the market.

IDC is planning to add additional segmentation based on the use of big data analysis. These may include areas like fraud detection, advance customer profiling, and large volume data generation.

ESSENTIAL GUIDANCE

This document presents IDC's current taxonomy for the technical computing market and should serve as a guide for the existing and future base of research in the target market.

LEARN MORE

Related Research

Additional research from IDC in the technical computing hardware program includes the following documents:

- *Experiences with Accelerators and Coprocessors in High-Performance Computing: HPC User Forum, September 15-17, 2014, Seattle, Washington* (IDC #251973, October 2014)
- *Major Global High-Performance Computing Initiatives: HPC User Forum, September 15-17, 2014, Seattle, Washington* (IDC #251971, October 2014)
- *Lenovo Completes Acquisition of IBM's x86 Server Business* (IDC #IcUS25176214, September 2014)
- *Worldwide Broader HPC 2014-2018 Forecast: Servers, Storage, Software, Middleware, and Services* (IDC #248835, June 2014)
- *Worldwide Technical Computing Server 2014-2018 Forecast* (IDC #248779, May 2014)
- *Global HPC Market Dynamics in 2013* (IDC #248137, April 2014)
- *Worldwide HPC Public Cloud Computing 2014-2017 Forecast* (IDC #247846, April 2014)
- *Summary of IDC's 2014 Research in the Use of HPC by Oil and Gas Organizations* (IDC #247704, March 2014)
- *IBM Sale to Lenovo Opens Opportunity for Other HPC Vendors* (IDC #IcUS24694314, February 2014)
- *IDC's Worldwide High-Performance Computing Predictions 2014* (IDC #WC20140211, February 2014)

- *Market Analysis Perspective: Worldwide HPC, 2013 – Directions, Trends, and Customer Requirements* (IDC #244742, December 2013)
- *HPDA Pulse Results: 2013 Hardware and Storage Market Analysis* (IDC #244493, November 2013)
- *HP FY13: Revenue Declines Abate on Stronger Core Business* (IDC #lcUS24466413, November 2013)
- *China Eyes 10,000-Fold Data Reduction for Internet of Things* (IDC #lcUS24392513, October 2013)
- *National and International Initiatives: HPC User Forum, September 2013, Boston, Massachusetts* (IDC #243776, October 2013)
- *High-Performance Data Analysis in the Life Sciences: HPC User Forum, September 2013, Boston, Massachusetts* (IDC #243774, October 2013)
- *Chinese Research in Processor Designs for High-Performance Computing and Other Uses* (IDC #243502, October 2013)
- *The Broader HPC Market 2012-2017 Forecast: Servers, Storage, Software, Middleware, and Services* (IDC #242742, August 2013)
- *IDC's Worldwide Technical Server Taxonomy, 2013* (IDC #242725, August 2013)
- *China Regains Top Supercomputer Title* (IDC #lcUS24190613, June 2013)
- *Worldwide High-Performance Data Analysis 2013-2017 Forecast* (IDC #241315, June 2013)
- *Top Issues for HPC Sites: HPC User Forum, April 29-May 1, 2013, Tucson, Arizona* (IDC #241463, June 2013)

Synopsis

This IDC study provides an overview of key technology and market definitions of the technical computing market. This taxonomy is used by IDC worldwide analysts to generate IDC market sizing, forecasts, and company models in the technical computing space. Vendors and users in the technical computing space will find this taxonomy useful in understanding how IDC defines and measures the markets they serve or are considering to enter.

"The technical computing market has emerged as one of the fastest-growing markets in the IT space in recent years," says Earl Joseph, program vice president for High-Performance Computing. "It is important for vendors and users as well as those who are looking for potential opportunities in the technical computing market to understand how the market is defined and measured by a set of commonly acknowledged terminologies and methodologies developed by the industry."

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-insights-community.com
www.idc.com