



## IDC FutureScape

# IDC FutureScape: Worldwide High-Performance Data Analysis 2016 Predictions – Science and Industry Turn to HPC for Mission-Critical Big Data Challenges

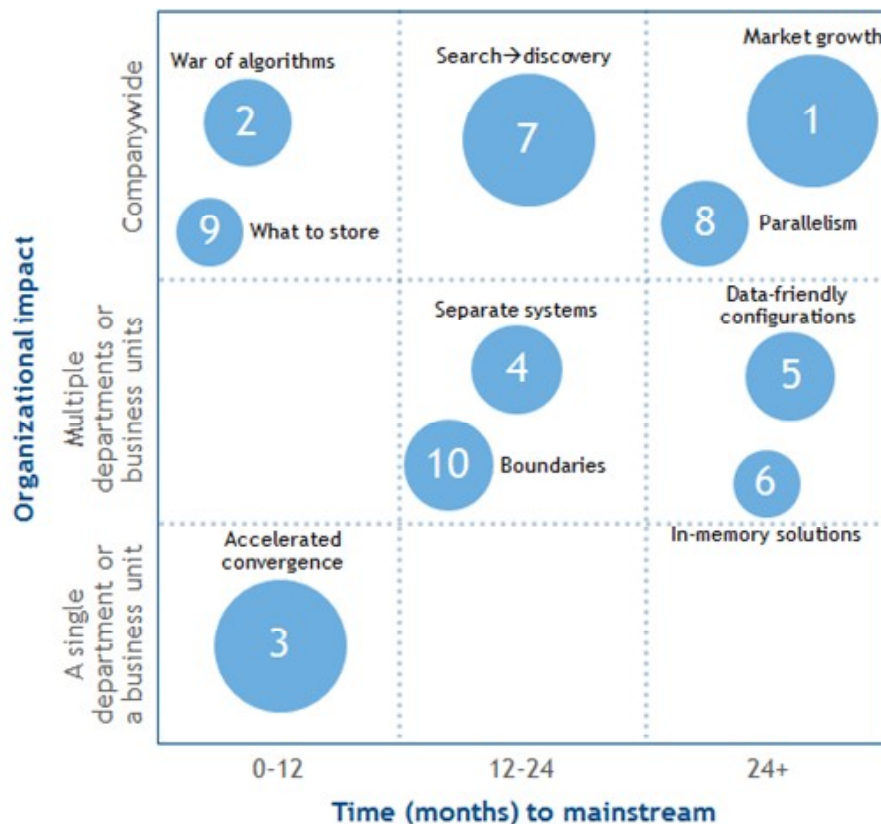
Steve Conway  
Robert Sorensen

Earl C. Joseph, Ph.D.

### IDC FUTUREScape FIGURE

FIGURE 1

## IDC FutureScape: Worldwide High-Performance Data Analysis 2016 Top 10 Predictions



Note: The size of the bubble indicates complexity/cost to address.

Source: IDC, 2015

Figure 1 presents IDC's high-performance data analysis (HPDA) top 10 predictions in terms of their likely impact across the enterprise and the time it will take for the predictions to reach mainstream. By mainstream, IDC means the broad middle of the bell curve of adoption (i.e., the 40-60% of enterprises that are neither the first movers and early adopters nor the last to act). Each bubble's size provides a rough indicator of the complexity and/or cost an enterprise will incur in acting on the prediction.

## IDC OPINION

---

The new edition of the *Strategy for American Innovation*, made public by President Obama on October 21, 2015, named high-performance computing (HPC) as one of the top investment priorities for growing the U.S. economy. A key role for HPC cited in the *Strategy for American Innovation* document is to address "the rise of extremely large data sets and attendant computational challenges."

High-performance data analysis is IDC's term, now widely adopted, for the use of high-performance computing technology to address very complex or time-critical big data problems – including jobs that enterprise server technology alone can't handle well. The HPC community was the birthplace of big data, but until recently, most data-intensive HPC (HPDA) was used for modeling and simulation (M&S) in government, academia, and companies ranging from automakers, petroleum firms, and big pharmaceuticals to the likes of Procter & Gamble (P&G), Walmart, Whirlpool, and Ping. The analytics side of HPDA has been less prominent. It was pioneered by the intelligence community in the post-World War II era, followed in the early 1990s by the back offices of large investment banks.

HPDA today is a convergence market in two respects. First, simulation and analytics are converging: Existing HPC sites are confronting much larger volumes of simulation data produced by bigger supercomputers, as well as massive analytics data from powerful new scientific instruments and sensor systems. Second, more commercial firms are turning to HPC for mission-critical, daunting analytics challenges – especially those that require very large shared memories and ultrafast interconnect fabrics for handling necessary data movement. Shining examples include PayPal, which added HPC for real-time fraud detection/prediction, and the U.S. Postal Service (USPS), which adopted HPC when its enterprise servers could not keep up with daily mail volumes.

Although big data analysis is sometimes assumed to depend only on software, this U.S. Postal Service evaluation shows how greatly an HPC parallel hardware-software architecture can speed solution times (the USPS is on its fourth generation of HPC systems):

The Postal Service reports that the ability to process very large data problems entirely in memory can boost performance by three to six orders of magnitude [for their application]. Conversely, having to process data outside of a shared memory can severely cut performance – by about 50 percent when moving data from one blade to a neighboring blade; by two-thirds when moving from the top blade in a rack to the middle blade; by a factor of 25 when moving data from the top blade to the bottom blade in the same rack, and by a far greater magnitude when transferring data between racks. ("HPC for Advanced Analytics at the USPS," *EnterpriseTech*, May 5, 2015)

IDC believes that the global HPC community's long-term experience with data-intensive computing will keep HPC users at the forefront of big data developments, increasingly benefiting the mainstream enterprise computing market in the process. It is no accident that the White House *Strategy for American Innovation* calls for coordinated investments in two priority areas: HPC and precision

medicine (a domain HPC is already becoming crucial for). The HPC community's primary contribution to enterprise big data may well be greater exploitation of parallelism, especially for cognitive computing and artificial intelligence (AI) initiatives. IDC's 2015 studies confirmed that many mainstream machine learning (ML) algorithms have never been parallelized to enable faster, higher-resolution solutions. As the big data market advances from today's predominantly static searches to the emerging era of higher-value, dynamic discovery of unknown patterns, the HPC hardware-software ecosystem will help lead the way.

## IN THIS STUDY

---

This study presents IDC's 10 top predictions for the worldwide HPDA market, along with their motivating megatrends (key drivers) and projected impacts on IT organizations. It is important to keep in mind that HPC organizations differ from the enterprise organizations that are often associated with the term *IT*. HPC sites may be in private sector businesses, but they may also be in government or university environments. Established HPC datacenters most often are physically separate from IT datacenters that handle operations, in cases where an organization has both types of datacenters.

HPC datacenters are also philosophically different from general IT deployments. IT is generally about *provisioning* – equipping each of the company's knowledge workers with the basic computing tools they need to perform their jobs productively and providing as little beyond that as possible in order to stay within the budget. HPC, on the other hand, is about *enablement* – providing a small subset of specialized knowledge workers with the most powerful computational tools the company can afford. A typical IT worker's desktop or laptop system, with some help from a server system, is capable of fully supporting the worker's computing requirements, while there is often no limit to the amount of computing power an HPC user could exploit on the company's behalf.

## SITUATION OVERVIEW

---

*High-performance data analysis* is the term IDC coined to describe the convergence of the established data-intensive HPC market and the high-end commercial analytics market that is starting to move up to HPC resources. In addition:

- A good commercial example is PayPal, a multibillion-dollar company that integrated HPC servers and storage into its datacenter workflow to perform sophisticated fraud detection on credit card transactions in real time. Real-time detection can catch (and predict) fraud before it hits credit cards. IDC estimates that using HPC has allowed PayPal to save well over \$700 million to date. PayPal is now extending HPC use to affinity marketing and datacenter management.
- The U.S. Postal Service is on its fourth generation of supercomputers. Enterprise IT technology alone was unable to keep pace with the Postal Service's rapidly growing daily volumes of data. HPC technology has enabled the USPS not only to achieve near-real-time response rates on this expanding data but also to exploit mission-critical competitive opportunities.
- The Center for Pediatric Genomic Medicine at Children's Mercy Hospital, Kansas City, Missouri, has been using supercomputer power to help save the lives of critically ill children. In 2010, the center's work was named one of *Time* magazine's top 10 medical breakthroughs. For 1 infant suffering from liver failure, the center used 25 hours of supercomputer time to analyze 120 billion nucleotide sequences and narrowed the problem down to 2 genetic

variants. This allowed the doctors to begin treatment with corticosteroids and immunoglobulin. Thanks to this highly accurate diagnosis of the problem and pinpointed treatment, the baby is alive and well today.

Simulation-driven HPC is the longest-standing part of the HPDA market. Since the start of the supercomputer era in the 1960s, important HPC workloads, such as cryptography and weather and climate research, have been data intensive. The newer kid on the HPDA block is analytics, which comes in many flavors. The financial industry has been running analytics on HPC systems at least since the late 1980s. But newer methods, from MapReduce/Hadoop/Spark to graph analytics, have greatly expanded the opportunities for HPC-based analytics.

The common denominator underlying simulation- and analytics-based HPDA workloads is a degree of algorithmic complexity that is atypical for transaction processing-based business computing. With the help of sophisticated algorithms, HPC resources are already enabling established HPC users, as well as commercial adopters, to move beyond "needle in a haystack" searches in order to discover high-value, dynamic patterns.

## Summary of Key Drivers

Many dynamics affect the worldwide market for HPDA and its parent market, high-performance computing – including the evolution of the scientific, engineering, and advanced analytics domains that contribute to the use cases for these markets. At the highest level, IDC has identified the following six megatrends that drive the HPDA market:

- **21st century battleground:** The race for cyber supremacy
- **East-West:** Shifting global economic power, balance, and influence
- **Risky business:** International impact of economic volatility
- **Cy-Q:** An interconnected, informed, interactive, intrusive, intelligent, and cognitive ecosystem
- **DX:** Accelerating business disruption from digital transformation
- **Talent quest:** High demand for next-generation IT skills, but scarce supply

## IDC FutureScape Predictions

IDC's top 10 predictions represent IDC's view of the most important trends in the formative HPDA market. Because the HPDA market is largely (although not entirely) a subset of the global HPC market, these predictions are also important for an understanding of important developments affecting HPC.

### *Prediction 1: The HPDA Market Will Grow Robustly to \$4.9 Billion by 2019*

IDC forecasts that the worldwide HPC ecosystem market will expand from \$21.0 billion in 2014 to about \$31.2 billion in 2019 (at an 8.3% CAGR), with about half of the 2019 total, or \$15.2 billion, coming from sales of server systems and the rest from storage, software, and technical support service. The high-performance data analysis subset of the HPC market will grow at about three times the rate of the overall market to reach \$4.9 billion in 2019, with \$3.1 billion of that amount contributed by server sales (at a 22.1% CAGR). Storage will remain the fastest-growing component of the overall HPC market (at a 9.9% CAGR) and the HPDA subset (at a 26.5% CAGR). Key HPC growth drivers include scientific challenges, heightened private sector competition, national security, and the rapid rise of data-intensive simulation and analytics.

## Associated Drivers

- **21st century battleground:** The race for cyber supremacy
- **East-West:** Shifting global economic power, balance, and Influence
- **Risky business:** International impact of economic volatility
- **Cy-Q:** An interconnected, informed, interactive, intrusive, intelligent, and cognitive ecosystem
- **DX:** Accelerating business disruption from digital transformation
- **Talent quest:** High demand for next-generation IT skills, but scarce supply

## HPDA IT Impact

- Servers will grow at 22.1% CAGR, while storage will grow faster than servers, at a 26.5% CAGR.
- Drivers: Organic growth of M&S + rapid rise of analytics.
- Commercial analytics adopters turn to HPC to address complexity and time criticality.

## Guidance

- HPC was the birthplace of big data.
- HPDA is fulfilling 1990s dreams of turning big data into actionable knowledge.
- Look to HPDA for emerging trends in big data (e.g., highly parallel machine learning).

## *Prediction 2: HPDA Competition Will Remain a War of Algorithms*

Algorithms embodying mathematical models are the main intellectual capital and competitive weaponry of the HPDA and overall big data markets. As a general rule, smart algorithms need less data/fewer iterations to home in on useful solutions, and less-clever algorithms require more data/iterations. IDC 2015 studies revealed a relatively sharp divide between machine learning algorithms in the enterprise big data community and those employed by HPDA (HPC) users. The enterprise ML algorithm scene began ramping up in earnest in about the year 2000 and has produced mature algorithms, but few of them have been parallelized to enable faster performance and higher resolution. This is because they grew up in the era when performance was advanced by Moore's Law-related generational jumps in the single-threaded speed of processors. By contrast, the HPDA ML algorithm scene has existed in the highly parallel world of HPC from the start. IDC believes that HPC parallelism will increasingly affect the enterprise big data market (see Prediction number 8). Skilled algorithm developers are in short supply globally.

## Associated Drivers

- **21st century battleground:** The race for cyber supremacy
- **Cy-Q:** An interconnected, informed, interactive, intrusive, intelligent, and cognitive ecosystem
- **DX:** Accelerating business disruption from digital transformation
- **Talent quest:** High demand for next-generation IT skills, but scarce supply

## HPDA IT Impact

- Algorithms will need continual revision to incorporate the latest science, engineering, analytics, and HPC technology advances.
- Algorithms developed for HPC will function well in commercial settings.
- Skilled algorithm developers will remain in short supply.
- Experienced HPC algorithm developers will be pursued by enterprise IT organizations.

## Guidance

- HPC user organizations will need to compete vigorously for skilled algorithm developers.
- Leading commercial firms will need to employ parallelized algorithms for competitiveness.
- HPC vendors can benefit from helping commercial firms transition to parallelized algorithms.

### ***Prediction 3: The Convergence of HPC-Hosted Simulation and Analytics Will Accelerate in the Next Five Years***

PayPal adopted HPC following a business (not IT) meeting aimed at reducing the large amount of revenue being lost to fraud. Someone in the room said, "Our problem's already been solved in the HPC market." This episode underscores both the potential benefits of HPC for advanced commercial analytics and the general lack of awareness of HPC-enabled analytics solutions. IDC nevertheless expects more commercial users to adopt HPC for advanced analytics problems that enterprise technology can't handle well alone – even as established HPC users increasingly add analytics to their simulation workloads. The HPC analytical engine will be part of the live workflow, unlike three-tier client/server architectures in which advanced analytics were performed on sample data off to the side. Since 2012, IDC has tracked revenue and produced forecasts for these repetitive commercial HPDA use cases: fraud and anomaly detection, marketing, and advanced business intelligence (BI). IDC expects precision medicine to be the next HPDA use case to emerge as a pursuable market segment.

## Associated Drivers

- **21st century battleground:** The race for cyber supremacy
- **Cy-Q:** An interconnected, informed, interactive, intrusive, intelligent, and cognitive ecosystem
- **DX:** Accelerating business disruption from digital transformation

## HPDA IT Impact

- Existing M&S users will add more analytics.
- More commercial users will turn to HPC for advanced analytics that enterprise technology can't handle well alone.
- HPDA moves beyond the three-tier client/server architecture to enter the live workflow.

## Guidance

- Most HPDA will remain on the simulation side.
- Repetitive use cases are emerging for commercial analytics using HPC technology.
- Many commercial firms are unaware of the benefits and affordability of HPDA.
- CIOs and IT heads (especially industry leaders) should evaluate potential benefits of HPC in the interest of remaining competitive.

### ***Prediction 4: HPC Sites Will Increasingly Buy Separate Systems for Simulation and Analytics***

For the foreseeable future, the bulk of the HPDA market will consist of HPC sites performing data-intensive modeling and simulation. The continual introduction of more powerful supercomputers, scientific instruments, and sensor networks will fuel the rapid growth of data-intensive M&S. This includes the nascent Internet of Things (IoT) (e.g., China's IoT plan will rely primarily on HPC for network management). HPC sites typically will run newer analytics workloads on their existing HPC systems, largely because existing budgets won't permit acquiring separate systems for analytics. IDC expects that as analytics workloads increase over time, however, larger numbers of established HPC sites will buy separate systems for simulation and analytics. Commercial firms adopting HPC for advanced analytics will buy data-friendly HPC systems from the start, to the extent that their needs can't be met by cloud computing (see Prediction number 10).

## Associated Drivers

- **DX:** Accelerating business disruption from digital transformation

## HPDA IT Impact

- Existing HPC users: Over time, more sites will arrange budgets to acquire separate HPC systems for data-intensive workloads (HPDA).
- Commercial analytics users will benefit from the drive to create more data-friendly HPC systems and attendant parallel software.

## Guidance

- Existing HPC sites with HPDA workloads should evaluate the benefits and affordability of separate systems for simulation and analytics.
- Industry-leading commercial firms that rely on advanced analytics should begin tracking advances in data-friendly HPC systems.
- HPC vendors: Consider the value of separate SKUs for systems designed to support HPDA workloads.

## *Prediction 5: The Basic Architecture Will Remain, But Configurations Are Already Becoming More Data Friendly*

Recent IDC worldwide studies confirm that buyers expect clusters and other HPC systems to retain the basic von Neumann architecture, but configurations will shift from today's extreme compute centrism (the "memory wall") to more data-friendly variants. Look for growing tension between the urge to maximize peak/LINPACK flops ("machoflops") and the drive to satisfy user requirements that depend on large memories and fast data movement. This trend will especially benefit storage and interconnect vendors, and innovation in both of these areas has been heating up. Data-oriented benchmarks are cropping up to complement the still-useful high-performance LINPACK test, but unless buyers embed the fuller benchmarks into procurements, they will be of limited use for motivating vendors to address underserved HPDA needs. Meanwhile, important progress will continue on the processor and accelerator fronts. A natural market segmentation process will assign appropriate roles to CPUs, GPGPUs, Intel Phi coprocessors, and FPGAs. Accelerators/coprocessors will increasingly move from exploratory use into production settings.

## Associated Drivers

- **Cy-Q:** An interconnected, informed, interactive, intrusive, intelligent, and cognitive ecosystem
- **DX:** Accelerating business disruption from digital transformation

## HPDA IT Impact

- Simulation and analytics applications will benefit from improved system balance.
- Budgets will benefit from a more rational ability to match HPC system resources to user requirements, including not having to overprovision processors in order to corral their associated memories.
- Bravery will be needed in some cases to resist optimizing around TOP500/LINPACK goals.

## Guidance

- **Established HPC buyers/users:** Attempt to resist "machoflops" pressure and maximize user requirements in procurements.
- **Industry-leading commercial firms:** Increasing data friendliness should make HPC systems even more relevant for your competitiveness.

- **HPC vendors:** Offer more data-friendly configurations without disrupting the TAM-friendly cluster architecture.

### ***Prediction 6: In-Memory Solutions Will Become Dominant by 2019***

There are two primary strategies available for dealing with large data volumes: accelerate data movement and reduce data movement. Most system vendors will employ a mix of both. The thrust to accelerate data movement has resulted in increased interconnect innovation (see Prediction number 5), but even unavoidable data movement carries penalties in latency and energy use. Experts say that performing a calculation typically consumes about 1pJ of energy, while moving the results of the calculation consumes 40-100pJ. With average global energy prices at \$1 million per megawatt, the largest user organizations already face energy bills of \$5 million to \$20 million per year. Hence, there is considerable impetus to reduce data movement by solving problems in-memory. The ideal solution is a large (say, 16TB+) physically shared central memory, but that would be prohibitively expensive. The next-best solution is a large, logically shared memory, such as PayPal and the USPS have access to. Less desirable for HPDA users, but much more affordable, are the limited in-memory processing capabilities associated with enterprise server systems. These technologies often require more extensive data reduction, with attendant loss of problem resolution that may or may not be satisfactory.

#### **Associated Drivers**

- **21st century battleground:** The race for cyber supremacy
- **Cy-Q:** An interconnected, informed, interactive, intrusive, intelligent, and cognitive ecosystem
- **DX:** Accelerating business disruption from digital transformation

#### **HPDA IT Impact**

- HPDA vendor and user organizations will pursue a dual strategy: to accelerate data movement and, where practical, to reduce data movement via in-memory processing.
- Power-saving advances in other HPDA system technologies (e.g., advanced cooling technologies, energy-efficient processors, and energy-aware software) will help curb rising energy bills.

#### **Guidance**

- **HPC buyers:** Look for increased innovation in best-in-class interconnects and for memory hierarchies to deepen, with SSDs, burst buffers, and other near-line storage technologies reducing the need to go out to disk. Also look for the most daunting HPDA problems to be run on systems featuring modest-scale physically shared (SMP) memories in combination with large-scale, logically shared memories.
- **HPC vendors:** Expect the market for big-memory supercomputers to expand from its small base today, but at a gradual rate. Main demand in the HPDA market will be for HPC clusters that combine capable interconnects with memory hierarchies that include cache, near-line, and offline storage. (Note, however, that graph analytics cache hit rates may be so low that it makes more sense to turn off caches.)

### ***Prediction 7: Static Search Will Give Way to Higher-Value, Dynamic Discovery in the Next Five Years***

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 in order 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. IDC 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. Machine learning and its more profound relatives, deep learning and cognitive computing, will play increasingly important roles in the HPDA market – not only for advanced problem-solving analytics but also for helping relieve programmers by further automating software stack functions. HPC-supported machine learning (HPDA) differs from enterprise machine learning principally in the former's heavy use of parallelized algorithms. IDC expects HPC users to accelerate the evolution from today's mostly supervised machine learning to the emerging era of unsupervised machine learning, including cognitive computing and AI.

### Associated Drivers

- **21st century battleground:** The race for cyber supremacy
- **Cy-Q:** An interconnected, informed, interactive, intrusive, intelligent, and cognitive ecosystem
- **DX:** Accelerating business disruption from digital transformation

### HPDA IT Impact

- Even "*Jeopardy!* Watson" was a needle in a haystack search engine.
- Discovering unknown patterns in near real time has higher value, from cybersecurity to precision medicine and more.
- Graph analytics and related methods will become much more important.

### Guidance

- Organizations not taking this path risk losing competitiveness.
- Dynamic discovery needs clever algorithms, big memories, fast I/O.
- It is still early in this game, but not too early to investigate/evaluate.

### ***Prediction 8: HPC Parallelism Is Already Improving Solution Times and Accuracy in the Enterprise Big Data Market***

There are two main forms of parallel processing. One form, exemplified by MapReduce/Hadoop, permits a simple data-processing task to be run independently and concurrently on many processing elements, with the aggregate results assembled later on. The key here is that the task must be small enough to fit onto a single core or other processing element. (MapReduce/Hadoop is employed by about the same percentage of HPC users and enterprise IT users.) The other form of parallel processing, often found in HPC, harnesses the combined power of many cores to tackle a large, complex problem. In extreme cases, a single problem may exploit more than 100,000 cores and take hours, days, or weeks to run. Both forms of parallel processing are equally valid, but the second form is indispensable for tackling complex, nonuniform ("hairy") problems in reasonable time frames. IDC believes that this form of parallel processing will become increasingly important in the emerging era of dynamic pattern discovery and other complex data challenges, not only for traditional HPC users but also for leaders in the enterprise computing market. The user examples cited in the Situation Overview section illustrate this point.

### Associated Drivers

- **21st century battleground:** The race for cyber supremacy
- **Cy-Q:** An interconnected, informed, interactive, intrusive, intelligent, and cognitive ecosystem
- **DX:** Accelerating business disruption from digital transformation

## HPDA IT Impact

- **Existing HPC buyers/users:** Most already exploit multicore parallelism to tackle complex problems; they will push to higher levels of parallelism for competitive advantage.
- **Industry-leading commercial firms:** Competitive pressure will motivate more commercial leaders to master multicore parallelism in order to accelerate innovation.

## Guidance

- **Existing HPC buyers/users:** Moore's Law-related generational jumps in single-threaded processor performance have effectively halted. Progress now depends on exploiting parallelism in earnest.
- **Commercial firms:** It's time to evaluate the potential benefits of HPC technology for mission-critical big data challenges. Firms failing to do this may lose competitive advantage.
- **HPC vendors:** Pay attention to MapReduce/Hadoop/Spark, but also realize that your experience with multicore parallelism may provide even more powerful entry into the commercial big data market.

## *Prediction 9: The Dilemma of What to Store Will Grow as Data Volumes Double Every Two to Three Years*

Massive data doesn't always become big data – mainly because storing massive data can cost massive money. Tom Lange, director, Modeling and Simulation, corporate R&D of Procter & Gamble, which has used HPC to design consumer products including Pampers diapers and Pringles potato chips, said that although P&G manufactures billions of Pampers a year and they all have data sensors affixed to them, all this data is deleted once the diapers exit the manufacturing line. Businesses typically won't pay for storing data that doesn't make them money. Someone has to make the economic argument for storing data, especially when it comes in petabyte-size chunks. The same goes for science. The Large Hadron Collider at CERN generates 1PB of data per second when it's operating, and the upcoming Square Kilometre Array telescope will disgorge 1EB (or 1,000PB) of data per day. In both cases, only a small fraction of the data will be stored. Storing it all would break the budget. China's HPC-based Internet of Things initiative is targeting a 10,000-fold data reduction to avoid having to confront 100ZB of data from home sensors alone by 2030.

## Associated Drivers

- **21st century battleground:** The race for cyber supremacy
- **Cy-Q:** An interconnected, informed, interactive, intrusive, intelligent, and cognitive ecosystem
- **DX:** Accelerating business disruption from digital transformation

## HPDA IT Impact

- Datacenter officials will need to decide what data to store and what not to store, knowing that some data to be discarded might have value in the future.
- Industries that are highly regulated or highly exposed to litigation (e.g., aircraft/automotive manufacturers, healthcare firms) will need to "bite the bullet" and pay for nearly all data to be stored.

## Guidance

- Users should conduct organizationwide assessments before deciding what to store. One business unit's "garbage" data may be the lifeblood of another business unit.
- User organizations should enforce storage guidelines and police storage at regular intervals to free up expensive space.

## ***Prediction 10: The Boundaries Between On-Premise, Remote, and Cloud-Based HPDA Will Be Substantially Lower by 2019***

Increasingly, sites wanting to keep their computing on-premise are coming under pressure from senior managers to justify not moving it to a public or hybrid cloud environment. Management officials often assume that public clouds can handle all workloads as effectively as on-premise resources, with greater flexibility and at lower cost. But most public cloud architectures today were designed with "embarrassingly parallel" enterprise jobs in mind and can be much slower and costlier than on-premise resources when running complex, harder-to-partition HPC/HPDA workloads. IDC studies show that more than 25% of HPC sites exploit public clouds today, but, on average, for less than 10% of their total workloads. IDC expects that over time, public cloud architectures will become capable of handling a reasonable portion of communications-intensive jobs, at which time more of those jobs will be sent to cloud environments. For now, a new breed of cloud services providers is promising to take the worry out of this situation by running each customer workload in the optimal environment, using contracted resources at an on-premise facility, or custom virtual machines (VMs) in a cloud infrastructure, or a "vanilla" cloud.

### **Associated Drivers**

- **21st century battleground:** The race for cyber supremacy
- **Cy-Q:** An interconnected, informed, interactive, intrusive, intelligent, and cognitive ecosystem
- **DX:** Accelerating business disruption from digital transformation

### **HPDA IT Impact**

- HPC datacenter officials, especially in government organizations, will need to educate their senior managements and enterprise IT colleagues on workload requirements versus present cloud limitations.
- Commercial firms moving to HPDA for competitive reasons may need to augment public cloud use with access to higher-capability resources on-premise or at an outsourced facility.

### **Guidance**

- Datacenter managers should be prepared to defend their planned mixes of on-premise, remote, and cloud computing environments against undifferentiated cloud imperatives coming down from above.
- Cloud services providers should ensure that their offerings address communications-intensive workloads or provide access to noncloud facilities that can handle these jobs well.

## ESSENTIAL GUIDANCE

---

- For **existing HPC users**, the world is changing in a mostly evolutionary way. Many have lived with data-intensive computing, now dubbed "big data," for several decades. Their computing resources (HPC hardware and software) have been expressly designed to exploit parallelism in order to solve extremely large, complex big data problems in reasonable time frames – say, a fraction of a second to a week, instead of a few days to a decade (although there is a need to exploit parallelism to an even greater extent). The present trend is largely an intensification of what has been: much larger volumes of simulation data produced by powerful new scientific instruments and sensor networks, as well as by brawnier supercomputers. Newer to many HPC sites is data analytics and the challenge of finding meaning not only in computer models simulating physical phenomena but in the data itself. Yet here, too, elements within the global HPC community – especially intelligence agencies and large investment banks – long ago pioneered data manipulation methods that spread into the business world and are now becoming more pervasive in the HPC world as well. A major challenge for many existing HPC sites is to merge seemingly orthogonal findings from simulation and analytics methodologies. Another challenge is to determine which HPDA workloads to run on-premise and which to send to the cloud, even as cloud capabilities remain a moving target for these workloads.
- **Newer commercial HPC adopters** encounter a more revolutionary path. Moving into the HPC world brings with it innovation-enhancing opportunities to tackle more complex business problems with greater accuracy and speed. But this move also involves a transition to a more parallel universe that may seem alien at first. In truth, many of the commercial firms that are now benefiting from adding HPC to their resources had one or more HPC-knowledgeable personnel on site. The transition will seem more daunting for companies that don't. The good news, however, is that most HPC vendors have prepared themselves to make this transition easier.
- For **HPC vendors**, the forecast strong growth of the HPDA market presents a welcome opportunity to increase revenue. Because the term *HPC* can be foreign and off-putting in commercial circles, some HPC vendors have struggled with whether to hide their HPC backgrounds. Doing so, in IDC's opinion, is almost always a bad idea. HPC vendors approaching the commercial big data market should not position themselves as arrowing out into an unfamiliar market. They should note instead that requirements at the high end of the commercial data analytics market are moving into the vendor's HPC competency space. It may be prudent to use a term other than HPC – that's why IDC coined the term *high-performance data analysis* – but HPC firms should actively promote their HPC-enabled experience and capabilities.

## KEY DRIVERS

---

This section includes detailed descriptions of the key drivers, or megatrends, that are contributing to the evolution of the worldwide HPDA market.

### 21st Century Battleground: The Race for Cyber Supremacy

Accelerated military spending for IT-enabled drones, robots, laser-guided projectiles, cyberwarfare, and space-based weapons accelerates technology evolution. Nonmilitary applications in robotics and other emerging technologies proliferate and ripple through associated industries.

In last year's IDC FutureScape drivers, we noted the institutionalization of cyberwarfare by governments, businesses, and criminal enterprises. In the ensuing year, the challenge will become more pronounced and public. Technologically advanced countries – the United States, Japan, the nations of the eurozone and, to a lesser degree, China and India – have placed much of their military, financial, manufacturing, energy, and communications infrastructure under computer control. As a result, there is an unprecedented temptation and opportunity for adversaries to gain the means to manipulate those systems.

The emergence of new cyberweapons, like laser-guided projectiles, drones, satellite-based systems, and robots, provides unprecedented ability to manipulate warfare from a distance. At the same time, as in all previous races to invent new weapons technologies, IDC anticipates a trickle-down effect (i.e., businesses will benefit from IT advances that are reapplied for nonmilitary use).

Cautionary voices have raised concerns about the potential for a future in which wars are started and run without human input. More than 1,000 artificial intelligence researchers have cosigned an open letter calling for a ban on autonomous weapons, or "killer robots." The letter – published by the Future of Life Institute, and whose signatories include Elon Musk and Stephen Hawking – warns that "artificial intelligence technology has reached a point where the deployment of such systems is – practically if not legally – feasible within years, not decades, and the stakes are high: autonomous weapons have been described as the third revolution in warfare, after gunpowder and nuclear arms." In the past 20 years, the balance of world power has shifted several times and is in the process of finding its new stasis. IDC believes that the next 5-10 years will witness a race to gain cyber supremacy both as a protective strategy and as a preemptive capability should global uncertainty morph into direct conflict.

In this new battleground, targets are ubiquitous. The notion of military or civil domains does not exist: every entity – including businesses, individuals, and governments – is a target. Business leaders must understand what is at stake in this new context and assess the potential value that they and their businesses represent – as valuable, perhaps vulnerable, targets.

## **East-West: Shifting Global Economic Power, Balance, and Influence**

The emergence of Asia as a counterbalance to Western economic dominance is changing the flow of money, people, skills, industry, and influence while creating a dynamic tension and competitive imbalance regionally and internationally.

According to a report by the Economist Intelligence Unit (EIU) Ltd. in 2015, China is expected to overtake the United States in 2026 in nominal GDP (in U.S. dollar terms). India is expected to replace Japan in third place, and countries in the European Union will continue their slide.

The consequence of shifting economic growth and influence from West to East has many moving parts, including the flow of IT talent (the so-called reverse brain drain), business leverage, and currency dominance.

IDC's 2Q15 macroeconomic forecast states, "The swing to emerging geographies is evident. The number of scientists and engineers in the United States and Western Europe is falling compared with the number of scientists and engineers in China and India, while growth in the number of IT-related employees in those countries is three times the world average."

## Risky Business: International Impact of Economic Volatility

Global economic activity is linked by IT and becomes increasingly sensitive to localized disruption, including Greece's default, the Middle East uncertainty, and other regional factors.

Stranded pockets of debt in the developed world, coupled with extreme variations in the distribution of personal and institutional wealth, are key factors in constraining the target and size of potential investments. Inter-enterprise IT connectivity accelerates the impact on business performance of market shifts and drives a greater need for continuous adaptation and change.

Macroeconomic forecasting and analysis is an exercise in evaluating the impact of trade-offs. In 2015, a nearly 30% plunge in China's stock market raised greater concerns about whether the fastest-growing major economic power's growth platform is built on a fragile foundation. Mitigating those concerns, however, is the reality that very few Chinese own stocks and the government's immediate willingness to steady the market with a large infusion of capital.

Of greater long-term concern is the widening gap in wealth between social strata, between young and old, between East and West. For example, in the United States, the richest 1% share of national income has risen from about 7% to nearly 20% since 1980 (source: Alvaredo, Atkinson, Piketty, and Saez, The World Top Incomes Database, 2013).

At the same time, the developed world is carrying a massive amount of unfunded institutional debt that is owed to aging workers in the form of pensions, healthcare benefits, and retirement plans. The U.S. federal government carries over \$18 trillion in debt, over \$150,000 per taxpayer (source: [usdebtclock.org](http://usdebtclock.org)).

Global economic volatility is one product of the trade-off between increasing liability in the West, coupled with slower growth; the rising tide of growth (and new financial responsibilities) in the East will drive a shift in money flows and investments for the foreseeable future.

CIOs and technology leaders play a major role in driving the enterprise through this volatility. For example, CIOs in the oil and gas industry must adapt to the consequences of oil barrel price variation. The agility and elasticity of the IT environment has become a critical differentiator. Leaders must build resilience and flexibility in their organizations and processes. New technologies offer new possibilities to shift from fixed to flexible costs and provide rapid adaptation to changing business needs.

## Cy-Q: An Interconnected, Informed, Interactive, Intrusive, Intelligent, and Cognitive Ecosystem

Cognitive advances, combined with robotics, IoT, artificial intelligence, augmented reality/virtual reality (AR/VR), and massive data sets, increase the ability of systems to mimic and surpass human intelligence while also creating an anonymous social platform that reinforces human-machine relationships. A lifestyle environment of omni-experiences becomes as much a function of interacting with intelligent interfaces as with real people.

What IDC calls Cy-Q is an intelligent, emergent cognitive network in which sensors and devices (the Internet of Things) create and operate a kind of meta-society based on the sharing and analysis of data about people and their lives that is updated continuously but largely invisible in its decision-making and cognitive processes.

For decades, popular culture has presented variations of a future in which smart machines develop the ability to intelligently interact with human society. IDC believes that we have already reached a point where a globally connected network ecosystem of massive volumes of data, combined with advanced and increasingly automated cognitive capabilities, is becoming capable of recognizing individuals through their digital identities on a daily basis.

From *Minority Report*, in which mall displays are able to recognize individuals and market to them in real time, to *Ex Machina*, in which a robot is so "human" that it can seduce a person, we have imagined advanced machine intelligence that complements or competes with our own. Ray Kurzweil (source: *The Singularity Is Near*) suggests that smart systems will surpass human capabilities within 20 years and advance far beyond from there. In another sense, they will become the next stage of human evolution.

A National Institute of Standards and Technology (NIST) report from 2012 states that the capacity of facial recognition technology has improved by two orders of magnitude over the past 10 years, while recognition errors have been dropping by 50% every 2 years.

U.S. programs are modest compared with those in China. The national Golden Shield program aims to use facial recognition to track every one of China's 1.3 billion people. The state of the art now in China is the ability to verify identity using facial recognition software from up to 500ft away (source: World Future Society, [www.wfs.org](http://www.wfs.org)). The simpler forms of connectedness (e.g., Amazon's suggestion of products based on a buyer's history and Google algorithms that create customized slide shows of pictures) might seem helpful or intrusive depending on the individual.

Robot helpers, or replacements for a growing number of skills and jobs, will similarly have a dual-edged impact. The question of whether a global "smart" ecosystem is perceived as intrusive or inclusive, hurtful or helpful, will play out over a long period of time; nonetheless, its ability to insert itself into the inner working of society, business, government, and personal life has already begun in earnest.

Business leaders must take stock of their products and services to see in what ways their businesses are vulnerable to competitors that may take advantage of advanced networks or robotic replacements. For every product or service, business leaders must assess the potential level of disruption from a cognitive and connected system.

## **DX: Accelerating Business Disruption from Digital Transformation**

Digital transformation (DX) experimentation for businesses goes mainstream and creates new business models with a seamless global reach. DX becomes a competitive requirement and the source of a massive wave of new investments in digitalizing business operations, communications, and services.

The term *digital transformation* has become so pervasive that it's easy to get confused about what it is and how it's impacting business strategy. IDC sees DX as an "outside in" process of recognizing the accelerating external adoption of social digitalization and adapting business initiatives to leverage IT's increasingly sophisticated ability to provide machine intelligence and intelligent analytic cognition to B2B, B2C, and C2C relationships. Examples abound:

- Automobile manufacturers, having already integrated digital communications and the Internet of Things and incorporated customer-led design, are exploring driverless vehicle technology.

- Department stores are deploying digital-enabled dressing rooms that enable shoppers to "try on" clothes by standing in front of a screen.
- Media companies routinely interact with customers to collect real-time feedback on changing tastes and habits.

The DX revolution is still in an early stage of its growth and not without its risks. High-profile examples of data hacking frequently make the news. Nevertheless, IDC believes that we are at a tipping point in DX adoption.

In IDC's recent *Digital Transformation MaturityScope Benchmark Survey*, we asked respondents to characterize their own DX maturity. While over 65% of respondents positioned themselves as being in the early stages of DX adoption (digital explorers/players), only 14% characterized themselves as digital "resisters" – the earliest stage and one characterized by an unwillingness to seriously consider DX initiatives. Clearly, businesses of every description have recognized the impact of DX on their sales and operations, and nearly all have begun to implement some degree of digital transformation.

Awareness of digital innovation is essential. Above all else, business and technology leaders must stay abreast of digital business initiatives that may have had or are having an impact on their business ecospheres. Every business of every size risks fundamental disruption because of new technologies, new players, and new ways of doing business.

## **Talent Quest: High Demand for Next-Generation Business/IT Skills, But Scarce Supply**

The ability to acquire 3rd Platform competency is constrained by a talent pool that is growing too slowly to meet business demand and too concentrated geographically to access for many enterprises. Scarce IT skills are concentrated generationally, requiring leadership to manage distinct demographic workstyles and expectations about career and the future. Jobs optimized for robotic machine intelligence and repetitive work impact a growing segment of the talent pool.

In 2013, IDC surveyed CIOs about the skills they'd need to support enterprise strategies until 2018. The talent that was identified as most urgently needed included expertise in business intelligence, analytics, mobile development, social developers, security, and business analysis. Not surprisingly, these are the skills necessary to drive digital transformation.

IDC has identified the most significant business areas in which increased IT spending is driving transformational change (see *Services Innovation: Future of IT Spending*, IDC #256556, June 2015).

CIOs and other technology leaders have found the previously mentioned skills to be in short supply, particularly when their business is remotely located from major cities or tech centers. As a result, there has been a dual-faceted trend to rely more on cloud-based service providers for scarce talent while also updating HR practices to make onboarding new hires more flexible and pertinent to the expectations of specific demographics (e.g., millennials).

IDC's 2Q15 macroeconomic forecast states, "IDC believes that the availability and skill level of talent have a direct impact on markets as diverse as network security and outsourcing. In the long run, the optimization of the slow-growth labor pool argues for cloud computing."

The search for talent is shaped and constrained by the way businesses and IT organizations accomplish work. However, CIOs and business leaders must keep abreast of technology trends in their business ecosystem – the talent needs driving your business may have already changed since you last formulated your staffing strategy. As talent cannot be procured or created easily, anticipation and strategic planning drive how leaders will prepare for their next talent needs.

## LEARN MORE

---

### Related Research

- *Emergent Use Cases in High-Performance Data Analysis: HPC User Forum, September 15-17, 2014, Seattle, Washington* (IDC #251976, October 2014)
- *USPS Touts Benefits of HPC for Big Data* (IDC #lcUS25123014, September 2014)
- *When Massive Data Never Becomes Big Data* (IDC #lcUS24922014, June 2014)
- *Perspectives on High-Performance Data Analysis: The Life Sciences* (IDC #248348, May 2014)
- *Micron Demonstrates Technologies to Address Emerging Challenges in Big Data Applications* (IDC #244843, December 2013)
- *HPDA Pulse Results: 2013 Hardware and Storage Market Analysis* (IDC #244493, November 2013)
- *Catalyst Supercomputer Heralds Shift to More Balanced Architectures* (IDC #lcUS24437513, November 2013)
- *China Eyes 10,000-Fold Data Reduction for Internet of Things* (IDC #lcUS24392513, October 2013)
- *Tools and Techniques for Technical Computing in Life Sciences: HPC User Forum, September 2013, Boston, Massachusetts* (IDC #243778, October 2013)
- *High-Performance Data Analysis in the Life Sciences: HPC User Forum, September 2013, Boston, Massachusetts* (IDC #243774, October 2013)
- *10 Things CIOs Should Know About High-Performance Computing* (IDC #241565, June 2013)
- *Worldwide High-Performance Data Analysis 2013-2017 Forecast* (IDC #241315, June 2013)

### Synopsis

This IDC study presents IDC's top 10 predictions for the formative high-performance data analysis (HPDA) market (i.e., the market for big data requiring HPC technology). This market covers the convergence of long-standing data-intensive simulations and newer methods for handling advanced data analytics. Market participants include established HPC sites, many of whom are adding analytics to their customary simulation workloads; commercial firms adopting HPC for the first time to tackle mission-critical data analytics challenges that enterprise server technology can't handle well alone; and TAM-expanding vendors moving in the HPC enterprise direction and vice versa.

"IDC forecasts that the global HPDA ecosystem market will continue growing robustly to reach about \$4.9 billion in 2019," said Steve Conway, IDC research vice president for HPC and HPDA. "Data-intensive simulation will expand rapidly, driven by powerful new scientific instruments, sensor networks, and brawnier HPC systems. Advanced analytics at existing HPC sites and a growing contingent of new commercial HPC adopters will help power this convergence market."

## 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](http://idc-insights-community.com)  
[www.idc.com](http://www.idc.com)

---

### Copyright Notice

This IDC research document was published as part of an IDC continuous intelligence service, providing written research, analyst interactions, telebriefings, and conferences. Visit [www.idc.com](http://www.idc.com) to learn more about IDC subscription and consulting services. To view a list of IDC offices worldwide, visit [www.idc.com/offices](http://www.idc.com/offices). Please contact the IDC Hotline at 800.343.4952, ext. 7988 (or +1.508.988.7988) or [sales@idc.com](mailto:sales@idc.com) for information on applying the price of this document toward the purchase of an IDC service or for information on additional copies or Web rights. [trademark]

Copyright 2015 IDC. Reproduction is forbidden unless authorized. All rights reserved.

