

Accelerating Innovation, Productivity and Time to Value with HPC using the IBM Elastic Storage Server (ESS)

Sponsored by IBM

Srini Chari, Ph.D., MBA

August, 2016

<mailto:chari@cabotpartners.com>

Executive Summary

Big Data has become characteristic of every computing workload. From its origins in research computing to use in modern commercial applications spanning multiple industries, data is the new basis of competitive value. The convergence of High Performance Computing (HPC), Big Data Analytics, and High Performance Data Analytics (HPDA) is the next game-changing business opportunity. It is the engine driving a Cognitive and Learning organization with Data as its fuel.

Businesses are investing in HPC to improve customer experience and loyalty, discover new revenue opportunities, detect fraud and security breaches, optimize research and development, mitigate financial risks, and more. HPC also helps governments respond faster to emergencies, improve security threat analysis, and more accurately predict the weather – all of which are vital for national security, public safety and the environment. The economic and social value of HPC is immense.

But the volume, velocity and variety of data are creating barriers to performance and scaling in almost every industry. To meet this challenge, organizations must deploy a cost-effective, high-performance, reliable and agile infrastructure to deliver the best possible business and research outcomes. This is the goal of IBM's Elastic Storage Server (ESS).

IBM ESS is a modern implementation of software defined storage, combining IBM Spectrum Scale (formerly GPFS) software with IBM POWER8 processor-based servers and storage enclosures. By consolidating storage needs across the organization, IBM ESS improves performance, reliability, resiliency, efficiency and time to value for the entire HPC workflow – from data acquisition to results – across many industries.

Real world industry examples spanning HPC workloads in life sciences/healthcare, financial services, manufacturing and oil and gas are discussed in detail. These examples and recent industry standard benchmarks (IBM ESS is 6x to 100x faster than other published results for sample workloads relevant for HPC) demonstrate the unique advantages of IBM ESS.

Clients who invest in IBM ESS can lower their total cost of ownership (TCO) with fewer, more reliable, higher-performing storage systems compared to alternatives. More importantly, these customers can accelerate innovation, productivity and time to value in their journey to become a Cognitive business.

Copyright© 2016. Cabot Partners Group, Inc. All rights reserved. Other companies' product names, trademarks, or service marks are used herein for identification only and belong to their respective owner. All images and supporting data were obtained from IBM, NVIDIA, Mellanox or from public sources. The information and product recommendations made by the Cabot Partners Group are based upon public information and sources and may also include personal opinions both of the Cabot Partners Group and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. The Cabot Partners Group, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your or your client's use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors which may appear in this document. This paper was developed with IBM funding. Although the paper may utilize publicly available material from various vendors, including IBM, it does not necessarily reflect the positions of such vendors on the issues addressed in this document.

Cabot Partners Group, Inc. 100 Woodcrest Lane, Danbury CT 06810. www.cabotpartners.com

High Performance Storage Key to Extract Value from Data Deluge

The relentless rate and pace of technology-enabled business transformation and innovation are astounding. Several intertwined technology trends in Social/Mobile, Instrumentation and the Internet of Things (IoT) are making data volumes grow exponentially. In 2018, about 4.3 exabytes (10^{18} bytes) of data is expected to be created daily with over 90% unstructured.¹

To extract timely insights from this growing data, it must be stored, processed and analyzed. But storing all this raw and derived data quickly, reliably, and economically for the long term is very challenging. To generate valuable time-critical insights, it is also imperative to quickly prepare, analyze, interpret and keep pace with these rapidly growing data volumes.

This requires faster, large-scale, cost-effective, highly-productive and reliable High Performance Computing (HPC) servers and storage. Today, large-scale HPC systems with clustered servers and storage are very affordable to acquire. This is spurring investments in HPC and High Performance Data Analytics (HPDA). Fueled by HPDA², by 2019, the total HPC market is expected to reach \$31 billion.³ With 9.4% annual growth, the HPC storage⁴ market is the fastest growing component, and is expected to reach \$6.8 billion by 2019.

Clients across many industries – Healthcare/Life Sciences, Financial Services, Oil and Gas, Manufacturing, Media and Entertainment, Public Sector and others – are increasingly using HPDA. These use cases integrate Systems of Records (structured data) with Systems of Engagement (unstructured data – images, videos, text, emails, social, sensors, etc.) across multiple organizational silos to produce new *High Value Systems of Insights* (Figure 1).



Figure 1: High Value Insights from Integration and Analysis of Structured and Unstructured Data

As clients add more storage capacity (including Network Attached Storage – NAS), they are realizing that the operating costs (including downtime and productivity loss) of integrating, managing, securing and analyzing exploding data volumes are escalating. To reduce these costs, many clients are using high performance scalable storage with parallel file systems.

The IBM Elastic Storage Server (ESS) is a leading high performance scalable storage system that combines IBM Spectrum Scale (an enterprise-grade, high performance parallel file system) with IBM POWER8 processor-based servers and storage enclosures.

¹ <https://storageservers.wordpress.com/2016/02/06/how-much-data-is-created-daily/>

² Earl Joseph, et. al., "IDC's Top Ten HPC Market Predictions for 2015, January, 2015

³ <http://www.idc.com/getdoc.jsp?containerId=259211>

⁴ <https://hpcuserforum.com/presentations/tucson2016/EarlApril2016Meetingslides4.11.2016.pdf>

By 2018, 4.3 exabytes of data generated daily

At \$6.8B in 2019, Storage is the fastest growing HPC market component

IBM ESS lowers TCO and improves performance and reliability

Many traditional HPC applications (seismic analysis, life sciences, financial services, climate modeling, design optimization, etc.) are becoming more data-intensive with higher fidelity models and more time-critical interdisciplinary analyses.

Newer HPDA applications are also being used for cyber-security, fraud detection, social analytics, emergency response, national security, and more. Deep Learning (Unsupervised Machine Learning leveraging HPDA) and Cognitive Computing are rapidly growing applications that can significantly benefit from HPC infrastructure.

Across many industries/applications, IBM ESS is helping clients enhance collaboration, innovation and productivity by optimizing HPC workflows across the entire data lifecycle.

Considerations to Optimize HPC Workflows across Data Lifecycle

Data volumes and access patterns intensify and vary widely as HPC applications, a crucial part of time/mission-critical workflows across many industries, become more data-intensive. What took days to analyze in a pure research context must now be done reliably in hours or less, even as larger number of projects and files must be tracked.

Storage architecture decisions are crucial to optimize the Total Cost of Ownership (TCO) for the entire collaborative workflow from data acquisition to data preparation to analysis and visualization/interpretation. Key considerations to enhance productivity and innovation are:

- **Data Location and Movement:** As data volumes grow exponentially, the costs of moving the data in and out of a compute processor becomes prohibitive. To move 1 byte from storage to the central processor, it could cost 3-10 times the cost of one floating point operation (flop).⁵ So, it is imperative to keep the frequently used active data on a high performance storage system close to the processor. This minimizes data motion and reduces access overheads especially when reusing the same data.
- **Applications/Workflow Performance:** After the raw data is acquired, it is typically consolidated, prepared and analyzed by multiple automated applications and analysis workflows working in tandem with end users; typically increasing the active data size by several-fold. These larger datasets must be processed on the hundreds of compute cores in a cluster. This means that the storage systems must perform and be able to feed these cores to keep the workflows operating at full-throttle.
- **Active and Archive Data:** Once data passes into the archive tier as part of a repository, it is important to quickly access data and metadata when needed, regardless of where it is and which operating system is requesting the file.
- **Data Security, Privacy and Protection:** How secure and private is the data? Is data stored in a redundant manner to ensure rapid recoverability? How much control does the user have over remote storage? These are especially critical in many commercial settings that have stringent regulatory compliance requirements.

Also, as disk drives become increasingly denser, traditional RAID is no longer an effective mechanism for data protection since it can take from several hours or even days to rebuild a failed drive, which can increase the chance of multi-disk failures.

The IBM Elastic Storage Server (ESS) can address many of the above challenges and more.

⁵ <https://www.nerac.gov/assets/NERSC-Staff-Publications/2010/ShalfVecpar2010.pdf>

IBM Elastic Storage Server (ESS) Overview

IBM Elastic Storage Server (ESS) is a modern implementation of software defined storage, combining IBM Spectrum Scale (formerly GPFS) software with IBM POWER8 processor-based servers and storage enclosures. Spectrum Scale is a widely used high-performance clustered/parallel file system that eliminates silos and simplifies storage management. It can be deployed in shared-disk or shared-nothing distributed parallel modes; providing concurrent high-speed file access to applications executing on multiple nodes of clusters.

Key Spectrum Scale features include:

- *POSIX file system* makes it easier to build workflows that include diverse workloads and data. Easier and quicker sharing and ingestion of data
- *Faster file system performance* that translates to workload acceleration through native exploitation of high performance networks
- *Run Analytics in-place*: The built-in Hadoop connector allows running Hadoop analytics in-place i.e. no need to copy data to HDFS to run Hadoop applications
- Encryption and secure deletion functions adds more security
- *Distributed metadata server* prevents a single point of failure and provides better performance than a single name node
- *Automated data management* features like, Active File Management (AFM), Information Lifecycle Management (ILM), and Multi-cluster promotes collaboration and improves operational efficiency
- *Access controls (ACLs)* allow better security control of data between multiple tenants in a shared infrastructure environment.

By consolidating storage requirements across the organization, IBM ESS reduces inefficiency and acquisition costs while simplifying management and improving data protection. Key features include:

- *Software RAID*: Runs IBM disks in a dual-ported storage enclosure that does not require external RAID storage controllers or other custom hardware RAID acceleration
- *Declustering*: Distributes client data, redundancy information and spare space uniformly across all disks of a JBOD – Just a Bunch of Disks. This reduces the rebuild or disk failure recovery process time compared to conventional RAID. Critical rebuilds of failed multi-terabyte drives can be accomplished in minutes—rather than hours or even days with traditional RAID
- *Data redundancy*: Supports highly reliable 2-fault-tolerant and 3-fault-tolerant Reed-Solomon-based parity codes (erasure coding) as well as 3-way and 4-way replication
- *Large cache*: Using a combination of internal and external flash devices along with the IBM Power server’s large memory cache, ESS is better able to mask the long latencies and inefficiencies of nearline SAS drives, while still leveraging these high density drives
- *Intuitive Graphical user interface (GUI)*: Allows management and monitoring of the system, both locally and remotely
- *Superior streaming performance*: Delivers over 25 GB/second of sustained performance
- *Scalability*: As server configurations are added to an installed configuration, the capacity, bandwidth, performance and single name space all grow. This means installations can start small, and grow as data needs expand.

A more detailed, industry-specific discussion follows. It highlights: the key industry trends, storage/data management challenges and how IBM ESS addresses these issues.

IBM ESS is a modern, high performance software defined storage system

Combines IBM Spectrum Scale parallel file system with Power Systems for high-performance and reliability

IBM ESS reduces inefficiency and acquisition costs while improving systems management and data protection

Overcoming Data Challenges in Life Sciences/Healthcare

All stakeholders in the healthcare/life sciences ecosystem – providers, payers, governments, biopharmaceutical companies, clinical research organizations (CROs), medical device and diagnostic firms, employers, and other public health organizations – are collaborating in innovative ways to drive better outcomes for the individual patient. Figure 2 describes this ecosystem and the substantial impact of a range of HPC disciplines. We detail key industry trends, storage/data management challenges and how IBM ESS overcomes these obstacles.

Key Life Sciences Trends. Rapidly declining gene sequencing costs, advances in recording technology and affordable clustered compute solutions to process ever larger datasets is transforming life sciences research. Today, a human genome can be sequenced within a day⁶ and for about \$1000, a task that took 13 years and \$2.7 billion to accomplish during the Human Genome Project.⁷ Likewise, data from light-sheet fluorescence microscopy (LSFM) can be analyzed to relate neuronal responses to sensory input and behavior. These analyses can run in minutes on clusters⁸, turning brain activity mapping efforts into valuable insights.

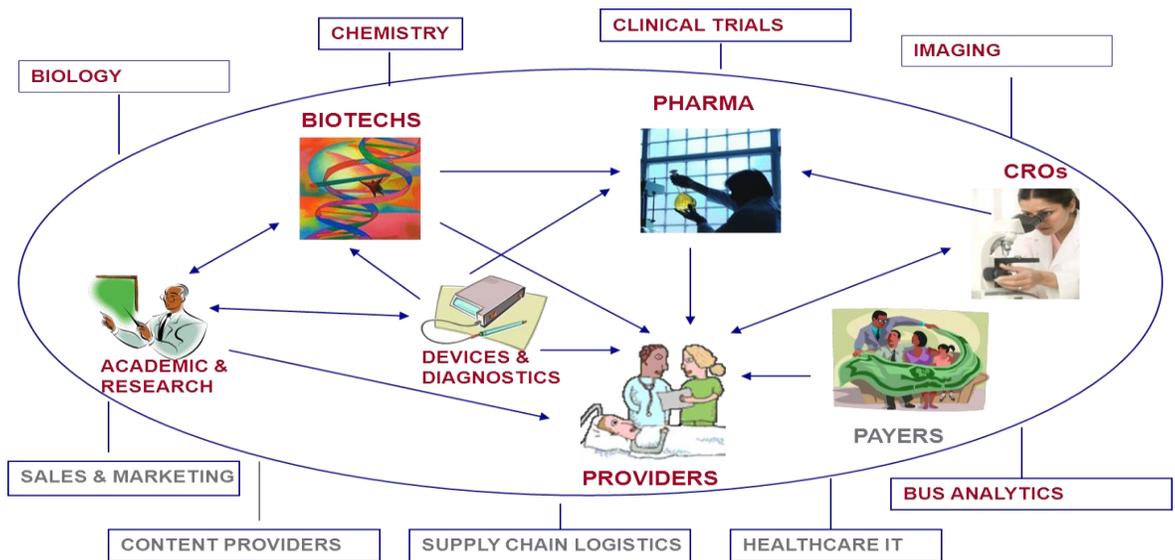


Figure 2: Healthcare/Life Sciences Disciplines/Industries (Red) Benefit from HPC

By 2025, the economic impact of next-generation sequencing (NGS) and related HPC technologies (Figure 2) could be between \$700 billion to \$1.6 trillion a year. Bulk of this value results from the delivery of better healthcare through personalized and translational medicine. NGS enables earlier disease detection, better diagnoses, discovery of new drugs and more personalized therapies. But it is crucial to overcome storage/data challenges.

Key Healthcare Industry Trends. Each person is expected to generate one million gigabytes of health-related data across his or her lifetime⁹, the equivalent of more than 300 million books. McKinsey estimates that if the US health care system were to use Big Data Analytics creatively and effectively to drive efficiency and quality, the potential value from

⁶ <http://www.prnewswire.com/news-releases/study-whole-genome-sequencing-technology-enables-26-hour-diagnosis-of-critically-ill-newborns-nearly-halving-previous-record-for-speed-300151018.html>

⁷ <http://www.veritasgenetics.com/documents/VG-PGP-Announcement-Final.pdf>

⁸ Jeremy Freeman, et. al., "Mapping brain activity at scale with cluster computing", Nature Methods, July 2014.

⁹ <http://www-03.ibm.com/press/us/en/pressrelease/46580.wss>

Healthcare /
Lifesciences
ecosystem
collaborating
to drive better
patient
outcomes

A human
genome can be
sequenced
within a day
for about
\$1000

By 2025, the
economic
impact of HPC
in Life Sciences
could be up to
\$1.6 trillion a
year

healthcare data could be worth more than \$300 billion every year, two-thirds of which would be in the form of reducing national health care expenditures by about 8 percent.

As Electronic Medical Record (EMR) systems become more affordable and widespread, data can be exchanged more easily. Recent advances in software are also making it simpler to cleanse data, preserve patient privacy and comply with Health Insurance Portability and Accountability Act (HPPAA). But there are still many obstacles with compiling, storing and sharing data reliably with high-performance and security.

Life Sciences/Healthcare Storage and Data Management Challenges. The rate of growth of genomics and imaging data continues to explode. For instance, the Illumina HiSeq X Ten System – designed for population-scale whole genome sequencing (WGS) – can process over 18,000 samples per year at full utilization. Each HiSeq X Ten System generates up to 1.8 terabytes (TB) per run. When the HiSeq X Ten System operates at scale, it can generate as much as 2 petabytes (PB) of persistent data in one year.

Similarly, new technologies based on imaging and multi-electrode arrays are making it possible to record simultaneously from hundreds or thousands of neurons and for some organisms, nearly the entire brain. For example, an hour of two-photon imaging in mouse can yield 50–100 gigabytes (GB) of spatiotemporal data, and recording from nearly the entire brain of a larval zebrafish using light-sheet microscopy can yield 1 TB or more.²

As the cost of sequencers and imaging instruments become more affordable, smaller institutions are increasingly deploying them. Even larger existing research organizations are purchasing more instruments. This only compounds the growth of distributed raw data. Raw data must be consolidated, aligned and packaged; making storage requirements even greater. Unfortunately, storage costs are not declining as fast as sequencing costs.¹⁰ Estimates are that in 2025, 2 to 40 Exabytes (EB) will be required just for the human genomes.¹¹

To provide higher value insights and holistic patient-centric healthcare, information across many data pools (claims, clinical, behavioral, genomic, imaging, etc.) must be rapidly integrated and analyzed. This requires a common high-performance, cost-effective analytical platform with low storage acquisition costs especially as data volumes continue to explode.

But many life sciences/healthcare organizations need are also realizing that the operating costs (including downtime and productivity loss) of managing, securing, tracking and cleansing these exploding volumes of data are growing even more. These organizations also need a storage solution that provides the bandwidth to scale to very large data volumes and allows users to collaborate across geographies within single name space.

The IBM ESS solution. It provides high performance (throughput) possible with native file access using POSIX client that traditional NFS or Scale-out NFS cannot match. In addition, Active File Management (AFM) enables global collaboration in a single name space.

Based on the POWER8 processor, ESS is a turnkey integrated solution that is quick to deploy, and is one hundred percent implemented in software and using standard servers and Just a Bunch of Disks (JBOD). This helps reduce the TCO compared to alternatives.

¹⁰ <http://www.genome.gov/sequencingcosts/>

¹¹ Zachary D. Stephens, et. al., "Big Data: Astronomical or Genomical?" PLOS Biology, 2015.

Declining costs of life sciences / healthcare data generation is driving storage requirements

But storage costs are not declining as fast as data generation costs

IBM ESS provides a cost-effective, scalable, high-performance storage for collaborative life sciences

IBM Systems can complete 65x coverage of the whole human genome

IBM Systems provide fast data ingestion rates from storage and superior performance to accelerate the entire workflow because of the unique architectural attributes of the POWER8: larger number of threads, greater memory size and bandwidth, higher clock rates and support for a Coherent Accelerator Processor Interface (CAPI).

For example, Burrows-Wheeler Aligner (BWA) is an efficient NGS program that aligns relatively short nucleotide sequences against a long reference sequence such as the human genome. With Power Systems and ESS, it is possible to complete 65x coverage of the whole human genome using the Broad Institute’s best practice pipeline consisting of BWA and other genomic tools (Samtools, PICARD, GATK) in less than 20 hours.

Addressing Data Management Challenges in Financial Services

Banks and Insurance companies are under intense pressure to cut costs yet improve the quality, accuracy and confidence of risk assessment. Integrated Financial Risk Analytics has become a core and pervasive part of these firms (Figure 3). Key industry trends, storage/data management challenges and how IBM ESS overcomes these obstacles are detailed here.

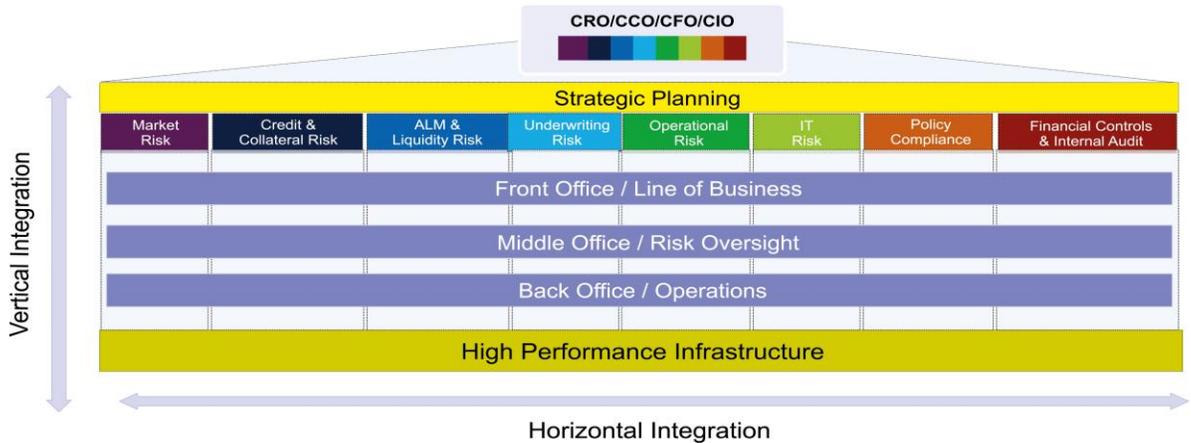


Figure 3: Better Outcomes with Vertical and Horizontal Integration of Risk

Key Trends. Increasingly, financial firms must adhere to an avalanche of stringent and complex regulatory requirements. Regulators now require tighter supervision of model risk management and are carefully dissecting failures from inadequately managing risk.

Besides traditional quantitative risks such as credit, market and liquidity risks; qualitative risks such as operational, reputation and strategic business risks are increasingly becoming important¹². Consequently, CEOs increasingly rely on their CFOs and Chief Risk Officers (CROs) for strategic advice and active risk management¹³ to gain a competitive edge.

In the past, many firms analyzed risk in silos or using ad-hoc approaches without structured governance processes. But now, with recent Basel III, Solvency II and Dodd Frank regulations aimed at stabilizing financial markets after the global financial crisis, firms have strong incentives to improve compliance so as to reduce capital requirements and reserves.

¹² Chartis, “The Risk Enabled Enterprise – Global Survey Results and Two Year Agenda”, 2013, <http://public.dhe.ibm.com/common/ssi/ecm/en/ytl03273usen/YTL03273USEN.PDF>

¹³ Pushing the frontiers: CFO insights from the IBM Global C-suite Study, 2014, <http://public.dhe.ibm.com/common/ssi/ecm/en/gbe03590usen/GBE03590USEN.PDF>

Integrated Risk Analytics is core and pervasive at Financial Services firms

Active risk management crucial for competitive edge

Pioneering firms investing in front-office real-time Risk Analytics that challenge traditional siloed storage

IBM ESS delivers very reliable, high-performance shared storage

Eliminate data and compute bottlenecks while ensuring business continuity

More than 2/3rd of losses sustained by financial firms between 2008 and 2011 were due to Credit Value Adjustment (CVA) mismatches¹⁴ rather than actual defaults. So, many leading firms are empowering their traders with investments in real-time risk analytics for better trading outcomes; extending their risk management operations from traditional end-of-day Value-at-Risk (VaR) reporting in the middle office to decision support in the front office.

The level of sophistication of risk models varies widely from relatively simple spreadsheet tools to complex mathematical models that can scale to thousands of economic scenarios and instruments. Complex real-time risk analytics pose many storage and data challenges.

Data/Storage Management Challenges. Many legacy risk systems are often ad-hoc and in silos, and cannot scale to handle the increased volume and frequency of analyses now demanded by regulators. The need to consistently apply accurate risk insights in making timely decisions throughout the enterprise is driving firms to improve standardize risk frameworks, consolidate risk systems, combine insights and share IT infrastructure.

The infrastructure must support a combination of large-scale compute and data-intensive analytics with real-time batch workloads. It must be reliable, flexible, agile and high-performance without overloading networks or letting costs go out of control.

The IBM ESS solution. It accelerates business results and delivers fully simulated near-real-time risk assessments. This innovative enterprise-grade, high-performance parallel storage solution allows users to right-size compute and right-place storage resources based on the importance and time-criticality of each analytic job.

It is fully POSIX compliant (so can run Analytics in-place instead of copying data to a different platform), delivering very high performance with no single point of failure and maintaining business continuity. Each compute node gets fast parallel read/write access to a common file system to accelerate the job.

IBM ESS dramatically improves simulation performance by avoiding filer “hot spots” common in network file sharing (NFS) or Server Message Block (SMB) file-sharing implementations. Fast parallel file system access is critical to speed up aggregation steps in risk analytics and also helps improve query performance.

It also provides efficient block-level data replication between multiple clusters in the same data center or in a remote center. Current data sets that are replicated between centers not only ensure business continuity if one center is unavailable, but also provide additional capacity to help meet periods of peak demand.

For example, major banks and insurance companies have seen dramatic reductions in aggregation time by replacing the traditional filers with ESS and increasing the data transfer speeds from compute hosts to shared storage. ESS eliminates compute and data bottlenecks by providing an independent path between compute nodes and storage to speed up data management – up to 10x improvement in raw file system I/O and 2x increase in scenario modeling capacity.

¹⁴ <http://www.shearman.com/~media/Files/NewsInsights/Publications/2013/11/BaselIIIFrameworkTheCreditValuationAdjustmentCVA-ChargeforOTCDerivativeTradesFIAFR111113.pdf>

Financial firms get a full report on all risk exposures on time every day

Engineering Simulation key to enhance manufacturer's product quality and reliability, productivity, innovation and profitability

Must quickly extract insights from large petabytes of simulation results data

This provides a very agile and scalable risk system, enabling analysts with on-demand capabilities for rapidly developing, testing, and deploying risk models while significantly improving overall system efficiency by effective sharing of critical resources. Firms can now reliably get a full report on their risk exposures on time every day. Mirrored data volumes to a second site can provide the business continuity critical to banks and insurance companies.

Solving Engineering Simulation Data Management Challenges

Many stakeholders in the manufacturing ecosystem – Automotive, Aerospace, Electronics and Heavy Industry, Suppliers, Governments and Academia – are collaborating to design and develop safer and better products. Figure 4 portrays how HPC applications in disciplines like structures, fluids, crash and design-optimization benefit manufacturing. We detail key industry trends, storage/data management challenges and the IBM ESS value proposition.

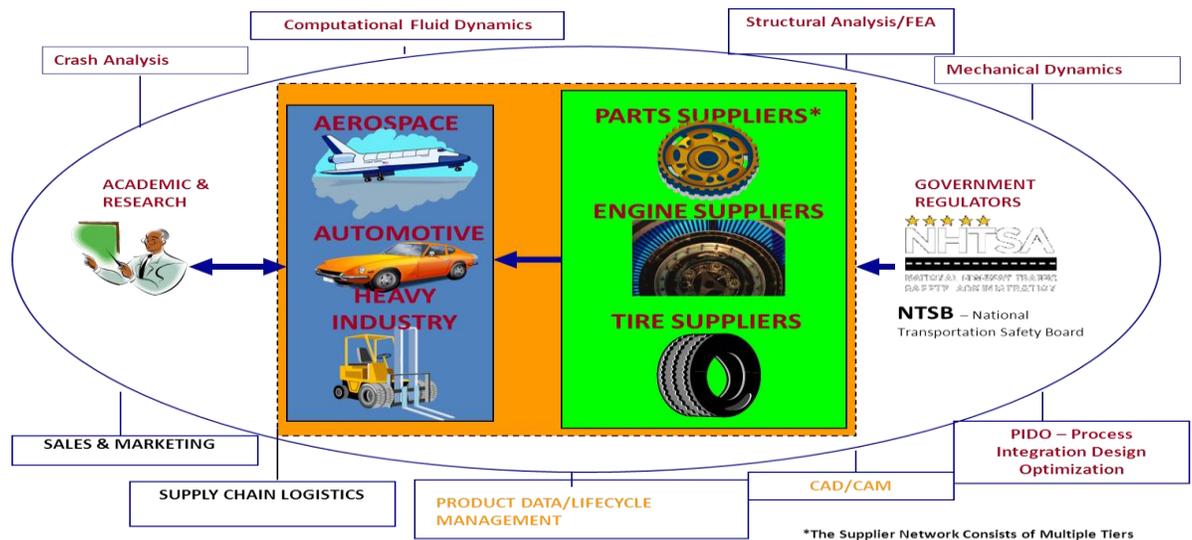


Figure 4: Manufacturing Industries/Disciplines (Red/Orange) Benefit from HPC

Mechanical Computer Aided Engineering (MCAE) Trends. Today's product development environment is global, complex and extremely competitive. Businesses race to improve product quality and reliability and to reduce cost and time-to-market to grow market share and profits. Complex cross-domain simulation processes must integrate with design throughout the product lifecycle. These realistic high-fidelity multidisciplinary simulations drive remarkable product innovation but cause a data deluge.

For instance, a single data set of Computational Fluid Dynamics (CFD) results, from one simulation, could run into 100s of gigabytes. During production analysis, when many such CFD simulations are necessary, these results can quickly aggregate to 100s of terabytes or even a few petabytes. Managing and drawing actionable business insights from this Big Data requires enterprises to deploy better data/storage management and simulation/analysis approaches to extract business value. This is critical to drive innovation and productivity.

Electronic Design Automation (EDA) Trends. A wide range of EDA solutions are used to collaboratively design, test, validate, and manufacture rapidly shrinking nanometer integrated chips leveraging advanced research, process technologies, and global R & D teams.

Predictive analytics using semiconductor physics' first principles drive up data sizes

Challenging to manage distributed data securely and efficiently, maintain data integrity and promote collaboration

IBM ESS improves availability, flexibility, performance and efficiency

Today, Static Timing Analysis (STA) for circuit simulation and Computational Lithography for process modeling are key HPC applications. The ultimate goal for many semiconductor R&D enterprises is to virtualize the full semiconductor development process. Doing so could reduce cost by requiring fewer silicon experiments and improving time to market for next generation semiconductor technologies. But these new predictive analytics applications based on near first principles of semiconductor physics could further drive up data volumes; placing even greater demands on HPC servers and storage.

Data/Storage Management Challenges. Engineering simulation data differs from other product design data in many crucial ways. It encompasses a large number of interconnected domain-specific tools, processes, and formats (e.g., EDA, crash, CFD). Simulation data is primarily unstructured and file sizes are often very large. Finally, simulation data must link with other closely related engineering datasets, and it has unique dependencies on HPC resources. These specific attributes make traditional data management approaches cumbersome and inadequate for the management of simulation data.

Unique storage/data challenges include: managing distributed data securely and efficiently, maintaining data integrity, promoting collaboration and communication, improving knowledge management, and reusing best practices. Engineering firms must also enhance the engineer's productivity with better user interfaces and timely, high-performance data access.

Simulations produce terabytes of data per day, and critical information and results are often buried in multiple files or documents. Engineering users require specific granularity for insights that are often hidden inside native data formats. This data variety requires advanced metadata management to enable searches and related sub-setting capabilities so users can extract the portions of datasets of interest to them.

The IBM ESS solution. It helps harness huge amounts of simulation data for greater engineering insight and productivity. These large shared data repositories can be co-located with the compute resource, and accessed remotely, to accelerate engineering workflows and enable increased collaboration and productivity throughout the manufacturing supply chain.

IBM ESS aggregates the power of multiple file servers and storage controllers, to provide:

- **Improved availability and disaster recovery** – IBM Spectrum Scale's advanced replication features allow data to be mirrored within a single location or across multiple sites. The file system can be configured to remain available automatically when disks or servers fail.
- **Flexibility** - Application data can be provided to different machines without moving data across storage devices or modifying existing file systems. Instead of having to copy the data to another machine for data access, that machine can be added to the ESS cluster. Data can be mounted and accessible from the new machine without requiring copies or movement. Active File Management (AFM) enables global collaboration in a single name space.
- **Continued high performance** - ESS allows storage expansion with automatic data rebalancing that is transparent to applications and also delivers performance on par with best of breed files systems.
- **Enhanced operational efficiency** - The simplified storage administration provided by ESS lowers total cost of ownership. Leading edge file system technologies such as integrated policy based storage management can help automate many storage management tasks; enabling high speed backups and restores and consolidation of multiple storage servers.

IBM ESS delivers dramatic reductions in cycle time and design / development costs

Companies are discovering large quantities of oil with better seismic data acquisition and processing

Requires massive petascale storage and HPC systems

With up to 10x better performance on global design and simulation tasks by eliminating storage related bottlenecks, IBM ESS helps provide dramatic reductions in the total cycle time (TCT). Many manufacturers have been able to bring new products to market ahead of competitors' offerings, and reduce design and development costs. This can significantly increase a manufacturer's revenues and profit margins. Better resiliency with ESS keeps simulation jobs running and also cuts systems administration time and cost.

Unlocking Value from Energy Exploration and Production Data

With diminishing conventional oil reserves, there is an urgent need to improve oil and gas exploration and production. HPC – particularly seismic processing – is being extensively used throughout the workflow (Figure 5) to accurately predict the location and nature of oil fields, eliminate delays and guesswork, and make better informed business decisions: where and how to drill, when to increase production, etc.

Seismic Processing Trends. Companies are discovering large quantities of oil by leveraging better seismic data acquisition and processing methods. This includes seismic data produced by sending sound waves from the earth's surface deep inside the earth and capturing the reflections from geological strata and formations, along with data from nearby oil wells. This processed and imaged data is then used with oil-well data to interpret and analyze the potential for oil and gas reserves. Figure 5 depicts a typical end-to-end workflow.

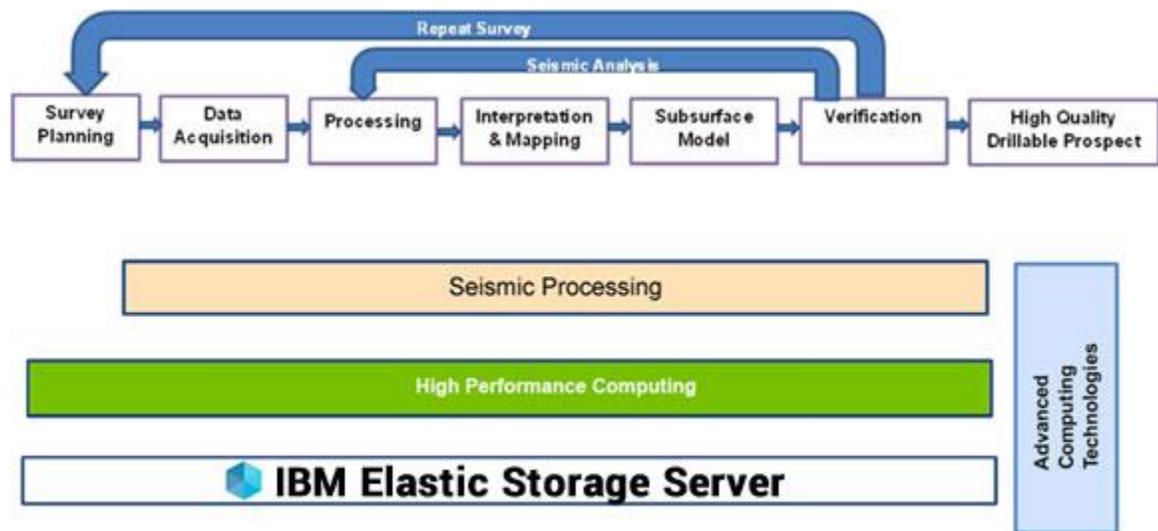


Figure 5: Iterative Seismic Survey Workflow with Embedded High Performance Data Analytics

Higher-fidelity imaging algorithms have vastly improved the ability to locate oil deposits and the probability of success of expensive deep-water drilling. But these advanced seismic processing techniques increase computational complexity and produce growing volumes of data; requiring the implementation of massive petascale storage and HPC systems.

Seismic Processing Storage/Data Management Challenges. Seismic survey and processing is highly data intensive and typically costs between \$20M - \$30M to collect survey data and about \$5M - \$10M to process it. The data capture itself occurs in odd locations such as ships at sea or vehicles in deserts over increasingly larger areas.

Storage critical for massive data ingest and for fast processing to produce multi-dimensional seismic maps

Must rapidly process petabytes of valuable data that accumulates over several decades

IBM ESS enables easy building, sharing, and reliable efficient acceleration of workflows

Storage is critical in two roles: first, in the ingestion rates in collecting data at the highest possible resolutions from the sensors - this can be at speeds upwards of 1 GB/sec, and second, in the fast processing of that data to produce 3D maps. Getting big data from shared storage to distributed computing processors results in high latency and reduced bandwidth.

Dataset Sizes: The use of higher-fidelity imaging algorithms for better interpretation requires a broader range of seismic data acquisition frequencies, larger data sampling rates and more acquisition locations. Frequencies range from low frequencies for deeper penetration to high frequencies for improved resolution. Multi-azimuth and wide-azimuth surveys make survey areas wider; requiring more boats, trawlers and sensors leading to even larger data volumes.

The raw survey data can be about 20-30 terabytes, but can quickly grow between 5 to 25 times as it is being processed. Portions of a single chunk of raw data are processed concurrently and then merged together once all processing has completed. The ‘live data’, meaning it’s on disk and is available for processing, is driven by the size and number of active seismic projects whereas ‘cold’ data is purged from disk storage and resides on tape.

Valuable data accumulates over years: it is not unusual for oil companies to retain and use data collected 20 years ago¹⁵ since the earth’s structure doesn’t change significantly in that time. This is why new storage requirements in the oil and gas industry easily touch tens of petabytes every year. As algorithms and processes continue to improve, oil companies will reprocess data to get better visibility into existing oil fields, leading to a growing dependence on very fast, highly scalable, easily managed storage platforms such as IBM ESS. Faster data processing enables quicker and better decisions about where to drill for oil.

The IBM ESS solution. It can reliably analyze an order of magnitude more seismic and reservoir data faster; producing better and more accurate decisions of where to drill and reducing the risk of “dry holes”. ESS also enables Oil and Gas companies to survey larger and more complex geological terrains faster. Large shared data repositories can be co-located with distributed compute resources to accelerate seismic survey workflows end-to-end.

IBM ESS aggregates the power of multiple file servers and storage controllers, to provide:

- **Easy building of complex workflows** – POSIX file system enables easier and faster sharing and ingestion of a diverse data types from multiple sources.
- **Run Analytics in-place** – The built-in Hadoop connector allows running Hadoop analytics in-place i.e. no need to copy data to HDFS to run Hadoop applications
- **More workflow acceleration** – Allows storage expansion with automatic data rebalancing that is transparent to applications and also delivers performance on par with best of breed files systems. Faster and more resilient storage that also provides faster recovery from failed disks with minimal impact to application performance.
- **Improved availability and disaster recovery** – Advanced replication features allow data to be mirrored within a single location or across multiple sites. The file system can be configured to remain available automatically when disks or servers fail.
- **Enhanced operational efficiency** - The simplified storage administration provided by ESS lowers total cost of ownership. Leading edge file system technologies such as integrated policy based storage management can help automate many storage management tasks; enabling high speed backups and restores and consolidation of multiple storage servers.

¹⁵ Adam Farris, “How big data is changing the oil & gas industry”, Analytics magazine, November-December 2012.

More precision in locating drillable prospects and increased utilization of existing wells

IBM ESS accelerates research, innovation and scientific breakthroughs

SPEC SFS is a realistic storage benchmark for real-life workloads

For example, with ESS, an end-to-end data analysis cycle now takes approximately a week at a major energy exploration company – a process that took 27 days to complete before. Oil and Gas companies can now be more precise in locating remote drillable prospects. Analytics, optimization and data virtualization techniques can render larger amounts of complex data in more intuitive ways, allowing engineers to improve their decision making and, ultimately, their production effectiveness. These new capabilities can help increase the utilization of the existing gas and oil fields. Given the enormous price tag of drilling a new well and the complexity of managing production facilities, even modest increases in oil exploration and production efficiency could tremendously improve energy affordability.

Examples Highlighting Unique Benefits of IBM ESS for HPC

Many clients benefit from the IBM Elastic Storage Server to boost performance, reliability and efficiency of their HPC workflows while lowering their Total Cost of Ownership (TCO) and accelerating the Time to Insights. Here are some client and real life performance examples.

Client: Research Center Deutsches Elektronen-Synchrotron (DESY)

Background	<ul style="list-style-type: none"> • DESY is a national research center in Germany that operates particle accelerators and photon science facilities used to investigate the structure of matter • Over 3,000 scientists from over 40 countries annually
Challenges	<ul style="list-style-type: none"> • New generations of sensing devices bring in ten times the amount of data in the same time as the previous devices did • Over 5 GB of data streams into computing center every second • Existing storage system was reaching the limits of capacity, lacked flexibility and exhibited slow performance • Very long delay between data acquisition and analysis by users.
Solution	<ul style="list-style-type: none"> • IBM Elastic Storage Server built on IBM Spectrum Scale combined with IBM POWER8 technology as new storage system.
Benefits	<ul style="list-style-type: none"> • Can easily scale to meet growing demand and remain an attractive research destination for top scientists worldwide • Rapid access in minutes to experiment data accelerates research, innovations and scientific breakthroughs • Increased operational efficiency with automated data management across multiple tiers throughout the lifecycle reduces burdens on systems administrators.

Performance. The 2014 [SPEC SFS benchmark](#) is a standard benchmark for measuring the maximum sustainable throughput that a storage solution can deliver. The prior 2008 benchmark measured throughput at the protocol layer – NFS or CIFS – and is now obsolete. The 2014 benchmark is protocol-independent and file system agnostic. It measures performance at the application system call level based on real-life data processing workloads.

IBM ESS first system to publish SPEC SFS benchmark

100x faster for VDA and 6x faster for SW Build

IBM ESS accelerates HPC workflows, enhances collaboration and efficiency, scales and protects IT investments

The Video Data Acquisition (VDA)¹⁶ workload reflects the data ingest phase typical in most data-intensive HPC applications with the number of Streams as the unit metric. The SW Build workload simulates large software compilation or build phase of a HPC workflow.

VDA	SPEC SFS reference	ESS GL6	SW Build	SPEC SFS reference	ESS GL6
# of Streams	15	1600	# of builds	26	160
Overall response time	48.79msec	33.98msec	Overall response time	0.96msec	1.21msec
# of load generators	32	8	# of load generators	32	8
# of storage nodes	4	2	# of storage nodes	4	2
# of drives	96	348	# of drives	96	348

Figure 6: IBM ESS is 100x faster for VDA and 6x faster for SW Build

IBM ESS is the first system to publish SPEC SFS 2014 results: **100x faster for VDA and 6x faster for SW Build** compared to the SPEC SFS reference system (Figure 6).¹⁷

Conclusions – Accelerating Time to Value

With the ever increasing volume of data, the boundaries between HPC and Analytics continue to blur. High Performance Data Analytics (HPDA) is growing rapidly. It is the engine of the next-generation of Cognitive and Deep Learning applications.

By 2019, global spending in Cognitive Computing is expected to reach \$31 billion with 17% in hardware.¹⁸ This hardware must perform and scale for HPC/Cognitive applications. The IBM Elastic Storage Server (IBM ESS) is designed for these very data-intensive workloads.

Across several industries, many HPC clients are already using IBM ESS to improve performance, reliability, resiliency, efficiency and time to value for the entire workflow.

Organizations should actively consider investing in IBM ESS to:

- *Accelerate HPC workflows* by over an order of magnitude by combining the unique attributes of the POWER8 architecture with IBM Spectrum Scale – the high-performance parallel file system that provides faster recovery from failed disks with minimal impact to application performance.
- *Enhance global collaboration* and improve organizational productivity and innovation by providing rapid access to information through a single name space.
- *Improve operational efficiency* with automated data management across multiple tiers throughout the lifecycle, and reduce burdens on systems administrators.
- *Scale and protect investments* in people, processes, platforms and applications throughout the Cognitive Computing journey from Data and HPC, to Analytics and Deep Learning.

Cabot Partners is a collaborative consultancy and an independent IT analyst firm. We specialize in advising technology companies and their clients on how to build and grow a customer base, how to achieve desired revenue and profitability results, and how to make effective use of emerging technologies including HPC, Cloud Computing, and Analytics. To find out more, please go to www.cabotpartners.com.

¹⁶ http://www.snia.org/sites/default/files/SpencerShepler_SPEC_Under-the-Hood_Review_Final.pdf

¹⁷ <https://www.spec.org/sfs2014/results/sfs2014.html>

¹⁸ <http://www.idc.com/promo/thirdplatform/RESOURCES/ATTACHMENTS/CognitiveSystemsInfographic.pdf>