WHITE PAPER
Intel® Enterprise Edition for Lustre* Software
High Performance Data Division

Big Data Meets High Performance Computing

Intel® Enterprise Edition for Lustre* software and Hadoop combine to bring big data analytics to high performance computing configurations.

Contents

Executive Summary ... 1
Exploiting High Performance Computing to Drive Economic Value ... 2
Challenges Implementing Big Data Analytics ... 2
Robust Growth in HPC and HPDA ... 3
Intel and Open Source Software for Today's HPC ... 5
Software for Building HPDA Solutions ... 6
Hadoop Overview ... 6
Enhancing Hadoop for High Performance Data Analytics ... 8
The Lustre Parallel File System* ... 8
Integrating Hadoop with Intel® EE for Lustre Software* ... 9
Why Intel® For High Performance Data Analytics ... 11

Executive Summary

Big Data has been synonymous with high performance computing (HPC) for many years, and has become the primary driver fueling new and expanded HPC installations. Today, and for the immediate term, the majority of HPC Big Data workloads will be based on traditional simulation and modeling techniques. However, the technical and business forces shaping Big Data will lead users to consider and deploy new forms of HPC configurations to unlock insights housed in unimaginably large stores of data.

It's expected that these new approaches for extracting insights from unstructured data will use the same HPC infrastructures: clustered compute and storage servers interconnected using very fast and efficient networking – often InfiniBand.

As the boundaries between High Performance Computing (HPC) and Big Data analytics continues to blur – it's clear that technical computing capabilities are being leveraged for commercial technical computing with algorithmic complexity as the common denominator. In fact, International Data Corp. (IDC) has coined the term, High Performance Data Analytics (HPDA), to represent this confluence of HPC and Big Data analytics deployed onto HPC-style configurations.

HPDA could transform your company and the way you work by giving you the ability to solve a new class of computational and data-intensive problems. It can enable your organization to become more agile, improve productivity, spur growth and build sustained competitive advantage. With HPDA, you can integrate information and analysis end-to-end across the enterprise and with key partners, suppliers, and customers. This enables your business to innovate with flexibility and speed in response to customer demand, market opportunity, or a competitor's move.

Exploiting High Performance Computing to Drive Economic Value

Across many industries, an array of promising HPDA use cases are emerging. These use cases range from fraud detection and risk analysis to weather forecasting and climate modeling. Here are a few key examples:

Using High Performance Data Analytics at PayPal.[1]

While internet commerce has become a vital part of the economy, detecting fraud in 'real time' as millions of transactions are captured and processed by an assortment of systems – many having proprietary software tools – has created a need to develop and deploy new fraud detection models. PayPal, an eBay company, has used Hadoop and other software tools to detect fraud, but the colossal volumes of data were so large their systems were unable to perform the analysis quickly enough. To meet the challenge of finding fraud in near real time, PayPal decided to use HPC class systems – including the Lustre file system on their Hadoop cluster. The result? In their first year of production PayPal saved over $700 million in fraudulent transactions that thy would not have detected previously.

Weather Forecasting and Climate Modeling

Many economic sectors routinely use weather[2] and climate predictions[3] to make critical business decisions. The agricultural sector uses these forecasts to determine when to plant, irrigate, and mitigate frost damages. The transportation industry makes routing decisions to avoid severe weather events. And the energy sector estimates peak energy demands by geography to balance load. Likewise, anticipated climatic changes – from storms and extreme temperature events – will have real impact on the natural environment as well as on human-made infrastructure and their ability to contribute to economic activity and quality of life. Governments, the private sector, and citizens face the full spectrum of direct and indirect costs accrued from increasing environmental damage and disruption.

In summary, Big Data has long been an important part of high performance computing – but recent technology advances, coupled with massive volumes of data and innovative new use cases – have resulted in data intensive computing becoming even more valuable for solving scientific and commercial technical computing problems.

Challenges Implementing Big Data Analytics

HPC offers immense potential for data-intensive business computing. But as data explodes in velocity, variety, and volume, it is getting increasingly difficult to scale compute performance using enterprise class servers and storage in step with the increase. One estimate is that 80% of all data today is unstructured; unstructured data is growing 15 times faster than structured data[4] and the total volume of data is expected to grow to 40 zettabytes (10^{21} bytes) by 2020[5]. To put this in context, the entire information on the World Wide Web in 2013 was estimated to be 4 zettabytes[6].

The challenge in business computing is that much of the actionable patterns and insights reside in unstructured data. This data comes in multiple forms and from a wide variety of sources – information-gathering sensors, logs, emails, social media, pictures and videos, medical images, transaction records and GPS signals. This deluge of data is what is commonly referred to as Big Data. Utilizing Big Data solutions can expose key insights for improving customer experiences, enhancing marketing effectiveness, increasing operational efficiencies and mitigating financial risks.

It's very hard to linearly scale compute performance and storage to process Big Data. In a traditional non-distributed architecture, data is stored in a central server and applications access this central server to retrieve data. In this model, scaling is linear – as data grows, you simply add more compute power and storage. The problem is, as data volumes grow, querying against a huge, centrally located data set becomes slow and inefficient, and performance suffers.

Robust Growth in HPC and HPDA

Though legacy HPC has been focused on solving the important computing problems in science and technology – from astronomy and scientific research to energy exploration, weather forecasting and climate modeling – many businesses across multiple industries want to exploit HPC levels of compute and storage to solve important problems, such as predicting financial risk with greater accuracy, or bringing drugs to the market earlier.

High performance computers of interest to your business are often clusters of affordable computers. Each computer in a commonly configured cluster has between one and four processors, and today's processors typically have from 4 to 12 cores. In HPC terminology, the individual computers in a cluster are called nodes. A common cluster size in many businesses is between 16 and 64 nodes, or from 64 to 768 cores.

Designed for maximum performance and scaling, high-performance computing configurations have separate compute and storage clusters connected via a high speed interconnect fabric, typically InfiniBand. Figure 1 shows a typical complete HPC cluster.

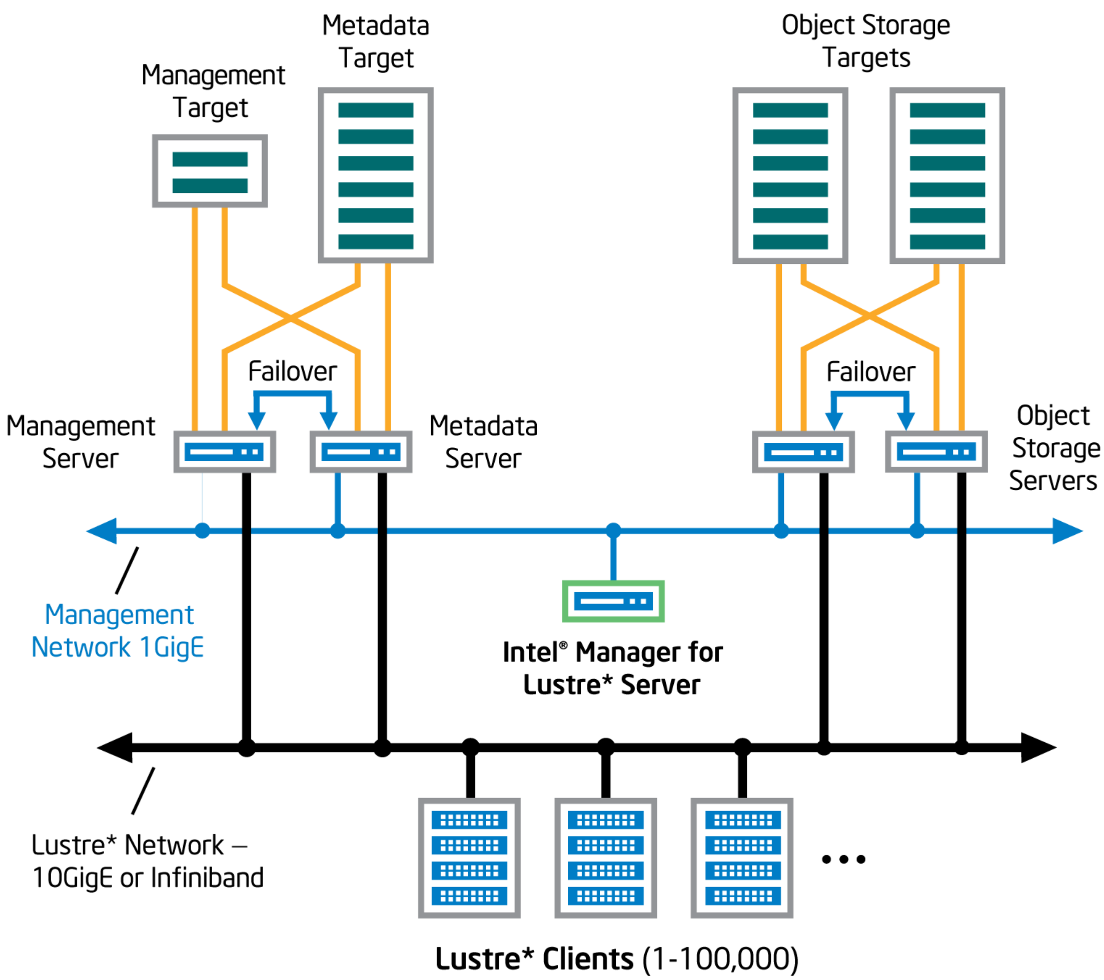

Figure 1. Typical High Performance Storage Cluster Using Lustre

According to IDC, the worldwide technical server market will continue to grow at a healthy rate of 7.3% CAGR from 2011 – 2016, and is projected to generate about $14.6 billion in revenues by 2016. Likewise, IDC estimates that storage will continue to be the fastest-growing segment within HPC, growing nearly 9% through 2016, and are projected to become a $5.6 billion market

A large part of this solid future growth is attributed to the increased use of HPC systems for high performance data analytics (HPDA) workloads within traditional HPC data centers. IDC forecasts that the market for HPDA servers will grow robustly at 13.3% CAGR from 2012 to 2016, and will approach $1.3 billion in 2016 revenue. HPDA storage revenue will near $800 million by 2016, with projected growth of 18.1% CAGR from 2011 – 2016.

From Wall Street to the Great Wall, enterprises and institutions of all sizes are faced with the benefits and challenges promised by big data analytics. But before users can take advantage of the nearly limitless potential locked within their data, they must have affordable, scalable, and powerful software tools to analyze and manage their data. More than ever, storage solutions that deliver sustained throughput are vital for powering analytical workloads. The explosive growth in data is matched by investments being made by legacy

HPC sites and commercial enterprises seeking to use HPC technologies to unlock the insights hidden within data.

Intel and Open Source Software for Today's HPC

As a global leader in computing, Intel is well known for being a leading innovator of microprocessor technologies; solutions based on Intel® Architecture (IA) power most of the world's data centers and high performance computing sites. What's often less well known is Intel's role in the open source software community. Intel is supporter of and major contributor to the open source ecosystem with a longstanding commitment to ensuring its growth and success. To ensure the highest levels of performance and functionality, Intel® helps foster industry standards through contributions to a number of standards bodies and consortia, including OpenSFS, the Apache Software Foundation, and Linux Standards Base.

Intel® has helped advance the development and maturity of open source software on many levels:

- As technology innovator, Intel® helps advance technologies that benefit the industry and contributes to a breadth of standards bodies and consortia to enhance interoperability.
- As ecosystem builder, we deliver a solid platform for open-source innovation, and help to grow a vibrant, thriving ecosystem around open source. Upstream, we contribute to a wealth of open-source projects. Downstream, we collaborate with leading partners to help build and deliver exceptional solutions to market.
- As project contributor, Intel® is the leading contributor to the Lustre file system, the second-leading contributor to the Linux kernel, and a significant contributor to many other open-source projects, including Hadoop.

Figure 2. A sample of the many open source software projects and partners supported by Intel

Software for Building HPDA Solutions

One of the key drivers fueling investments in HPDA is the ability to ingest data at high rates, then use analytics software to create sustained competitive or innovation advantages. The convergence of HPDA, including the use of Hadoop with HPC infrastructure is underway. To begin implementing a high-performance, scalable, and agile information foundation to support near real-time analytics and HPDA capabilities, you can use emerging open source application frameworks such as Hadoop to reduce the processing time for the growing volumes of data common in today's distributed computing environments.

Hadoop functionality is rapidly improving to support its use in mainstream enterprises. Many IT organizations are using Hadoop as a cost-effective data factory for collecting, managing and filtering increasing volumes of data for use by enterprise applications. It is also being used for data archiving. Likewise, business users are deploying Hadoop for analytics applications that process and analyze large volumes of data, especially unstructured data such as logs, social media data, email, network, image, video, and sensor data.

But you also need robust Reliability-Availability-Serviceability (RAS), security, governance processes, and end-to-end support, normally found in enterprise-grade IT solutions provided by companies such as Intel®. By sourcing these key technologies from Intel®, you will get a trusted partner and support throughout your Big Data analytics implementation journey.

Hadoop Overview

Hadoop is an open-source software framework designed to process data-intensive workloads – typically having large volumes of data – across distributed compute nodes. A typical Hadoop configuration is based on three parts: Hadoop Distributed File System (HDFS), the Hadoop MapReduce application model, and Hadoop Common. The initial design goal behind Hadoop was to use commodity technologies to form large clusters capable of providing the cost effective high I/O performance MapReduce applications require.

HDFS is the distributed file system included within Hadoop. HDFS was designed for MapReduce workloads that have these key characteristics:

- Use large files
- Perform large block sequential reads for analytics processing
- Access large files using sequential I/O in append-only mode.

Having originated from the Google File System, HDFS was built to store massive amounts of data reliably and support efficient distributed processing of Hadoop MapReduce jobs. Since HDFS resident data is accessed through a Java-based application-programming interface (API), non-Hadoop programs cannot read and write data inside the Hadoop file system. To manipulate data, HDFS users must access data through the Hadoop file system APIs, or use a set of command-line utilities. Neither of these approaches is as convenient as simply interacting with a native file system that fully supports the POSIX standard. (POSIX, or Portable Operating System Interface, is an IEEE family of standards to ensure application portability between operating systems).

In addition, many current HPC clients are deploying Hadoop applications that require local, direct-attached storage within compute nodes which is not common in HPC. Therefore, dedicated hardware must be purchased and systems administrators must be trained on the intricacies of managing HDFS.

By default, HDFS splits large files into blocks and distributes them across servers called data nodes. For protection against hardware failures, HDFS replicates each data block on more than one server. In place updates to data stored within HDFS is not possible. To modify the data within any file, it must be extracted from HDFS, modified, before being returned to HDFS.

Hadoop uses a distributed architecture where both data and processing are distributed across multiple servers. This is done through a programming model called MapReduce, through which an application is divided into multiple parts, and each part can be mapped, or routed to and executed on any node in a cluster. Periodically, the data from the mapped processing is gathered or reduced to create preliminary or full results. Through data-aware scheduling, Hadoop is able to direct computation to where the data resides to limit redundant data movement.

The basic block model of the Hadoop components is shown below in Figure 2. Notice that most of the major Hadoop applications sit on top of HDFS. Also, in Hadoop 2, that these applications sit on top of YARN (Yet Another Resource Negotiator). YARN is the Hadoop resource manager (job scheduler).

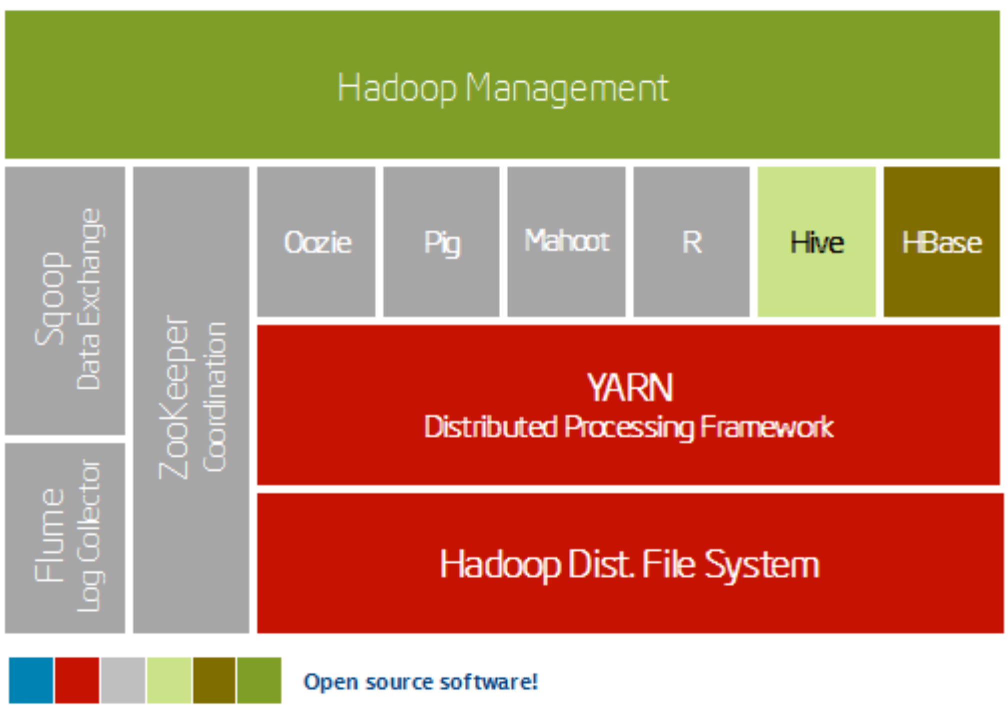

Figure 3. Hadoop 2.0 Software Components

Enhancing Hadoop for High Performance Data Analytics

High Performance Computing developed from the need for high computational performance and is normally characterized by high-speed networks, parallel file systems, and diskless nodes. Big Data grew out of the need to process large volumes of data typically with low speed networks, local file systems, and disk nodes. To bridge the gulf, there is a pressing need for a Hadoop platform that enables big data analytics applications to process data stored on HPC clusters.

Ideally, users should be able to run Hadoop workloads like any other HPC workload with the same performance and manageability. That can only happen with the tight integration of Hadoop with the file systems and schedulers that have long served HPC environments.

The Lustre* Parallel File System

Lustre is an open-source, massively parallel file system designed for high-performance and large-scale data. It is used by more than 60 percent of the world's fastest supercomputers. With Lustre, it's common for production class solutions to achieve 750 GB/s or more, with leading edge users reaching over 2 TB/s in bandwidth. Lustre scales to tens of thousands of clients and tens or even hundreds of petabytes of storage. The Intel® High Performance Data group is the primary developer and global technical support provider for Lustre-powered storage solutions.

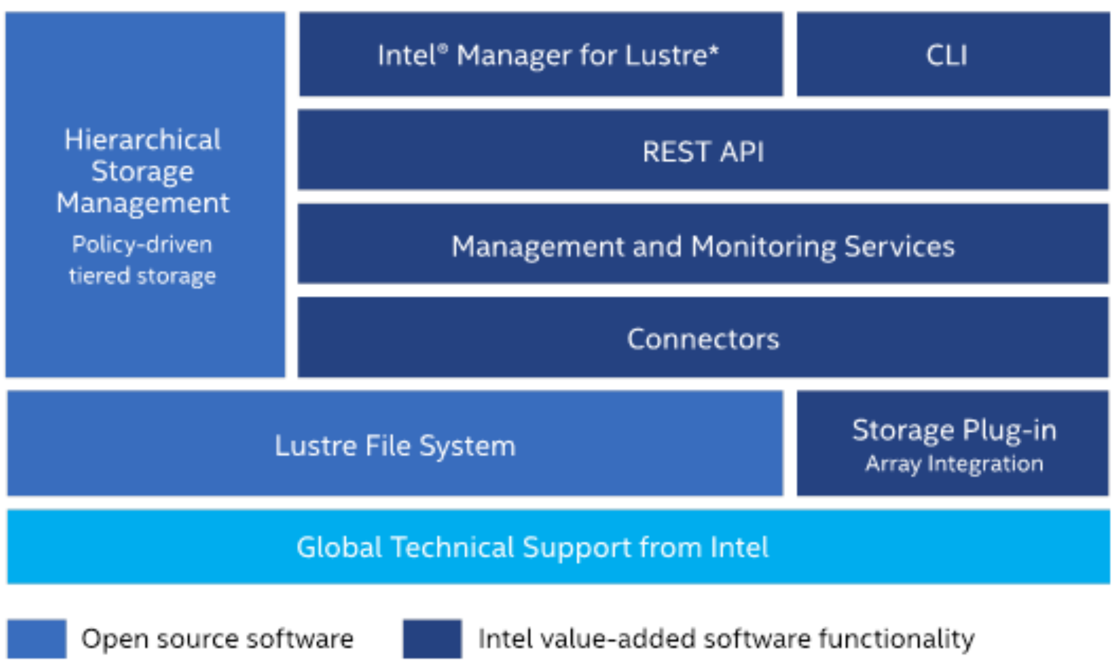

Figure 4. Intel® Enterprise Edition for Lustre software

To better meet the requirements of commercial users and institutions, Intel® has developed a value-added superset of open source Lustre known as Intel® Enterprise Edition for Lustre* software. Figure 4 illustrates the components of Enterprise Edition for Lustre* software. One of the key differences between open source and

Intel® EE for Lustre software are Intel® the unique software 'connectors' that allows Lustre to fully replace HDFS.

Integrating Hadoop with Intel® EE for Lustre* Software

Intel® EE for Lustre software is an enhanced implementation of open source Lustre software that can be directly integrated with storage hardware and applications. Functional enhancements include a simple but powerful management tools that make Lustre easy to install, configure, and manage with central data collection. Intel® EE for Lustre software also includes software that allows Hadoop applications to exploit the performance, scalability, and productivity of Lustre-based storage –without making any changes to Hadoop (MapReduce) applications.

In HPC environments, the integration of Intel® EE for Lustre software with Hadoop enables the co-location of computing and data without any code changes. This is a fundamental tenet of MapReduce. Users can thus run MapReduce applications that process data located on a POSIX-compliant file system on a shared storage infrastructure. This protects investment in HPC servers and storage without missing out on the potential of Big Data analytics with Hadoop.

In Figure 5 you'll notice that the diagrams illustrating Hadoop (Figure 3) and Intel Enterprise Edition for Lustre (Figure 4) have been combined. The software 'connector' in EE for Lustre replaces allows Lustre to replace HDFS in the Hadoop software stack, allowing applications to use Lustre directly and without translation. Unique to Enterprise Edition for Lustre, the software 'connector' is allows data analytics workflows, which perform multi-step 'pipeline' analyses, to flow from stage to stage. Users can run various algorithms against their data in series – step by step – to process the data. These pipelined workflows can be comprised of Hadoop MapReduce and POSIX compliant (non-MapReduce) applications. Without the 'connector', users would need to extract their data out of HDFS, make a copy for use by their POSIX compliant applications, capture the intermediate results from the POSIX application, before capturing those results and restoring them back into HDFS, then continue processing. These time consuming and resource draining steps are not required when using Enterprise Edition for Lustre.

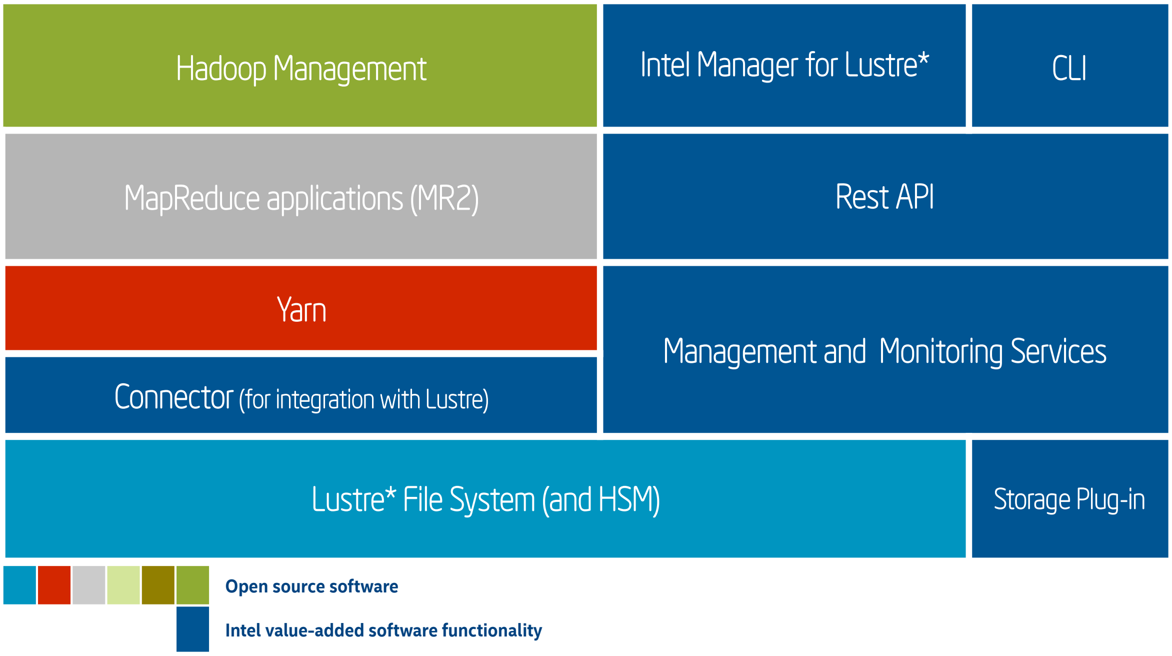

Figure 5. Software Connector allowing Intel® EE for Lustre software to replace HDFS

Further, the integration of Intel® EE for Lustre software with Hadoop results in a single seamless infrastructure that can support both Hadoop and HPC workloads thus avoiding islands of Hadoop nodes surrounded by a sea of separate HPC systems. The integrated solution schedules MapReduce jobs along with conventional HPC jobs within the overall system. All jobs potentially have access to the entire cluster and its resources, making more efficient use of the computing and storage resources. Hadoop MapReduce applications are handled just like any other HPC task.

The Intel® Enterprise Edition for Lustre* Software Advantage

Feature	Benefit
Shared, distributed storage	Global, shared design eliminates the 3-way replication HDFS uses by default; this is important as replication effectively lowers MapReduce storage capacity to 25%. By comparison, a 1 PB (raw) capacity Lustre solution provides nearly 75% of capacity for MapReduce data as the global namespace eliminates the need for 3-way replication. Allows pipelined workflows built from POSIX or MapReduce applications to use a shared, central pool of storage for excellent resource utilization and performance.
Sustained storage performance	Designed for sustained performance at large scale, Lustre delivers outstanding storage performance.
Maximized application throughput	Shared, global namespace eliminates the MapReduce *Shuffle phase* – applications do not

	need to be modified, and take less time to run when using Lustre storage - accelerating time to results.
Intel® Manager for Lustre*	Simple but powerful management tools and intuitive charts lowers management complexity and costs.
Add MapReduce to existing HPC environment	Innovative software connectors overcome storage and job scheduling challenges, allowing MapReduce applications to be easily added to HPC configurations.
Avoid CAPEX outlays	No need to add costly, low performance storage into compute nodes.
Scale-out storage	Shared storage allows you to add capacity and I/O performance for non-disruptive, predictable, well planned growth.
Management simplicity	Vast, shared pool of storage is easier and less expensive to monitor and manage, lowering OPEX.

Why Intel® For High Performance Data Analytics

From desk-side clusters to the world's largest supercomputers, Intel® products and technologies for HPC and Big Data provide powerful and flexible options for optimizing performance across the full range of workloads.

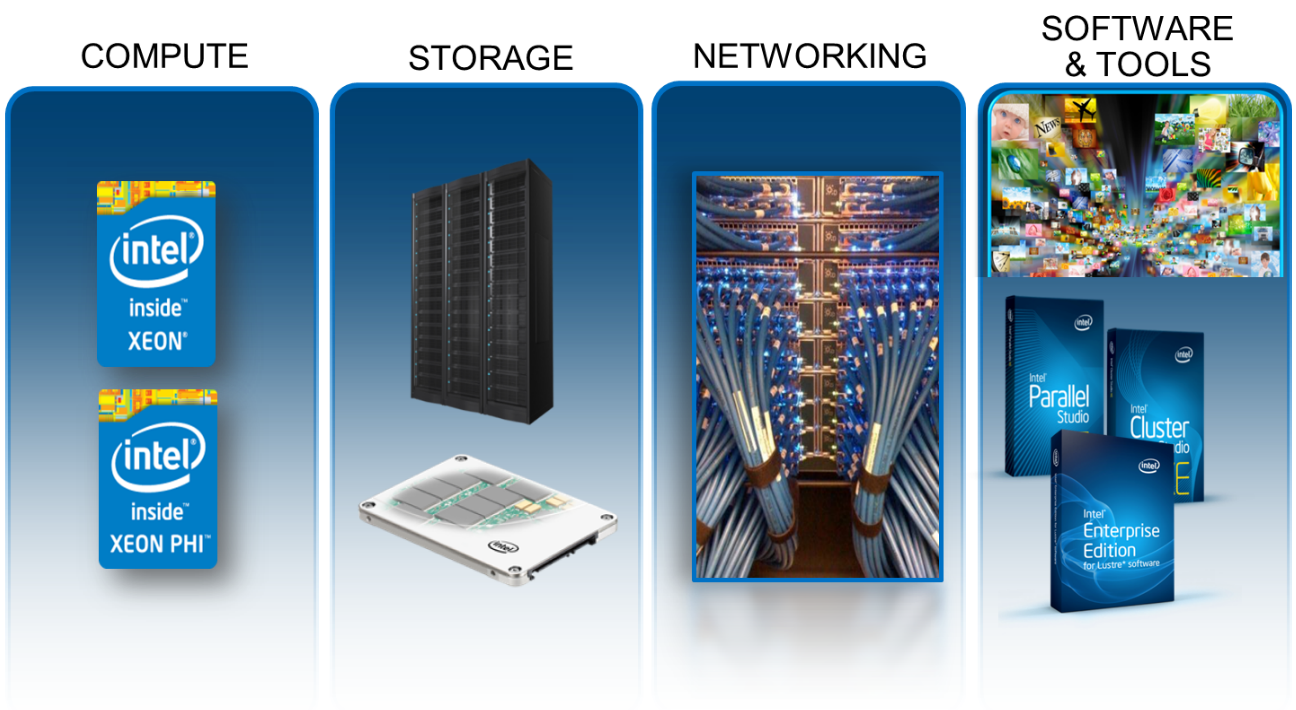

Figure 6. Intel® Portfolio for High Performance Computing

The Intel® portfolio for high-performance computing provides the following technologies:

- Compute - The Intel® Xeon processor E7 family provides a leap forward for every discipline that depends on HPC, with industry-leading performance and improved performance per watt. Add Intel® Xeon Phi coprocessors to your clusters and workstations to increase performance for highly parallel applications and code segments. Each coprocessor can add over a teraflops of performance and is compatible with software written for the Intel® Xeon processor E7 family. You don't need to rewrite code or master new development tools.
- Storage - High performance, highly scalable storage solutions with Intel® Lustre and Intel® Xeon Processor E7 based storage systems for centralized storage. Reliable and responsive local storage with Intel® Solid State Drives.
- Networking - Intel® True Scale Fabric and Networking technologies – Built for HPC to deliver fast message rates and low latency.
- Software and Tools: A broad range of software and tools to optimize and parallelize your software and clusters.

Notices and Disclaimers

[1] http://www.datanami.com/2013/06/19/idc_talks_convergence_in_high_performance_data_analysis/

[2] Laurie L. Houston, Richard M. Adams, and Rodney F. Weiher, "The Economic Benefits of Weather Forecasts: Implications for Investments in Rusohydromet Services", Report prepared for NOAA and World Bank, May, 2004.

[3] "Financial Risks of Climate Change", Summary Report, Association of British Insurers, June, 2005.

[4] Ramesh Nair, Andy Narayanan, "Benefitting from Big Data: Leveraging Unstructured Data Capabilities for Competitive Advantage". Ramesh Nair, Andy Narayanan. Booz & Company. 2012. http://www.booz.com/media/uploads/BoozCo_Benefitting-from-Big-Data.pdf

[5] IDC Digital Universe 2020.

[6] Richard Currier, "In 2013 the amount of data generated worldwide will reach four zettabytes"