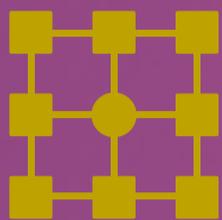


**LAB REVIEW**  
**WD ARKEIA RA4300**  
**BACKUP APPLIANCE**

**INTERVIEW**  
**SANDEEP GUPTA**  
NVIDIA SENIOR DIRECTOR  
PROFESSIONAL SOLUTIONS GROUP

**VIEWPOINT**  
**WHAT IT TAKES**  
**TO REINVENT**  
**SUPERCOMPUTING**



# HPC REVIEW

THE REFERENCE MEDIA IN HIGH-PERFORMANCE IT SOLUTIONS by HPC TODAY

[www.hpcreview.com](http://www.hpcreview.com)

# BIG DATA

THE TECHNOLOGIES TO ACHIEVE **SMART DATA**

-BIG DATA-



# GIGABYTE™

## A LEADING GPU SERVER, AT YOUR FINGERTIPS

### G190-H44 GPU Server

- Supports 4 x double slot cards
- Supports the Intel® Xeon Phi™ cards
- 4 x Storage bays and 2 x expansion card slots
- Redundant 2000W power supply, 80 PLUS Platinum certified



Available at:

**MICROTRONICA**  
A DIVISION OF ARROW

[b2b.gigabyte.com](http://b2b.gigabyte.com)

Supports the  
**Intel® Xeon® Processor E5-2600 V3**  
Intel Inside®. Powerful Solutions Outside.



Intel, the Intel logo, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel Corporation in the U.S. and/or other countries.

# BIG DATA: ACHIEVING A SUCCESSFUL TRANSITION TO SMART DATA

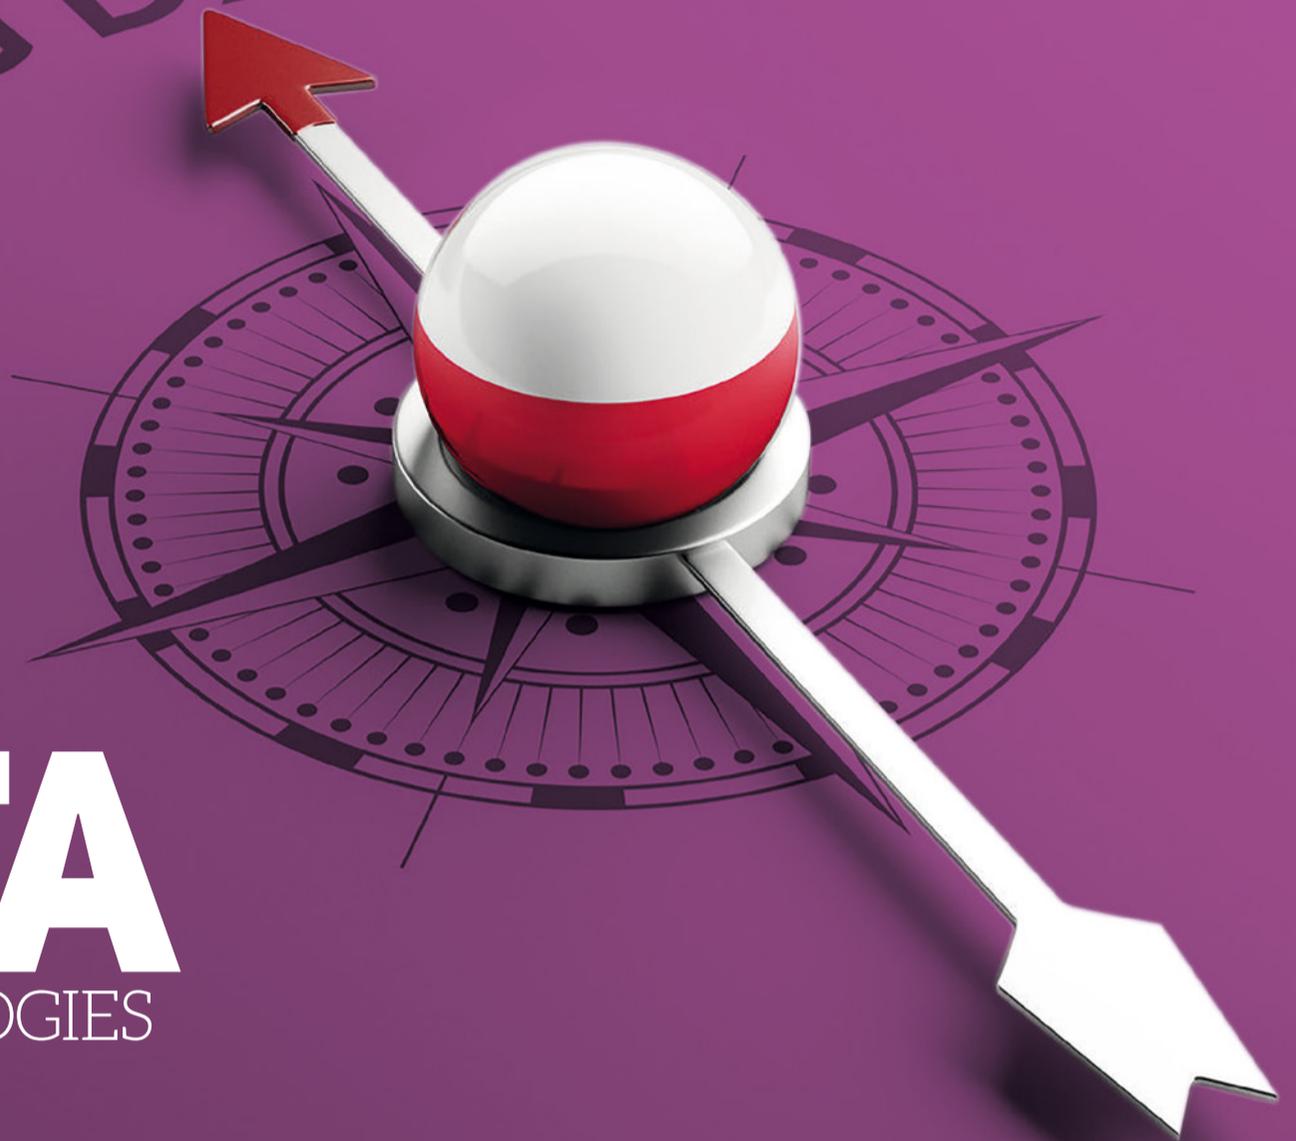
The fertile imagination of many science fiction writers has enabled them to devise pioneering concepts that surround us today: computers, the Internet, artificial intelligence and even the machines capable of learning, the beginnings of Machine Learning as we know today. However, a phenomenon that had not been foreseen by these visionaries, is the deluge of data generated by humanity and to which we have to face in all its forms. Whether to capture, store them, or analyze them to extract the most significant part, once again the machine, combined with the human genius, comes to our rescue.

## **THE CHALLENGE: TO GO FROM BAD DATA TO SMART DATA**

For big data is at risk if the software mechanisms are not sufficiently relevant to extract the good Data. This is crucial to extract what data scientists have already named the Smart Data, which represents only an infinitesimal part of the original data corpus. Because in this matter, in the end it is the quality, not the quantity of data that makes the difference. The aim of these treatments is tangible in analytical business intelligence matters and affects all aspects of modern process. Anticipation, production, logistics, improving the efficiency ratio of the resources involved and thereby the competitiveness. Big data future promises to be bright, thanks to the many technologies that underlie convergent efforts of researchers, developers and Data Scientists. The step after that? Solving this complex input / output equation in real time so that Smart Data can benefit everyone!

**HAPPY READING!**

-BIG DATA-



# BIG DATA

THE TECHNOLOGIES  
TO ACHIEVE  
SMART DATA

**NEWSFEED**

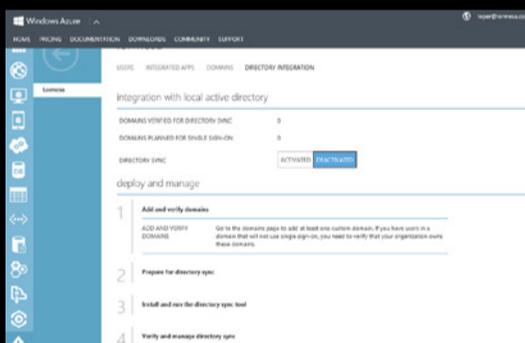
## INTEL OPTANE



### INTERVIEW: SANDEEP GUPTA

IBM INVENTS THE  
SUCCESSOR TO SILICON

**LAB REVIEW**



MICROSOFT AZURE  
ACTIVE DIRECTORY  
WD ARKEIA RA4300

**HOW TO**

OCZ'S NEW NVME-  
COMPLIANT Z-DRIVE  
6000 SSD SERIES

**VIEWPOINT**

WHAT IT TAKES  
TO REINVENT  
SUPERCOMPUTING  
OVER AND OVER AGAIN

**TECH ZONE**

INTEL SEES SIGNIFICANT  
LUSTRE MOMENTUM  
IN ENTERPRISE  
AND RESEARCH HPC

EDITOR'S NOTE

OPINION

SOURCES

THE HPC OBSERVATORY

# BIG HPC REVIEW CONTEST

## HPC REVIEW

### WIN AN OCZ SABER 1000 SATA ENTERPRISE SSD! (960 GB capacity model)

#### Saber 1000 Features

- Leverages OCZ's Barefoot 3 controller, in-house firmware, and next generation A19nm NAND from Toshiba
- Delivers sustained performance and consistent I/O responses for high-volume deployment enterprise SATA SSDs
- Supports Power Failure Management Plus (PFM+) to protect against unexpected power loss events
- Central management capability via OCZ StoragePeak 1000 software suite
- Energy savings resulting from low 3.7W operating power
- 2.5-inch, 7mm enclosure footprint
- Backed by a 5-year warranty



**CLICK BELOW TO ENTER THE CONTEST**

**I want to win!**

**THE WINNER WILL BE ANNOUNCED ON NOVEMBER 30TH, 2015**

No purchase necessary. Participation open worldwide. Drawing on November 30th, 2015.

# WHY MACHINE LEARNING NEEDS HPC

**STEVE CONWAY**RESEARCH VP, IDC HIGH PERFORMANCE  
COMPUTING GROUP

It seems to me that a reasonable definition of technology is artifice that augments our native human abilities. With it, we can lift better (lever), move better (vehicles), see better (microscope, telescope), hear better (radio, hearing aids), feel more sensitively (haptics), and more. I'm not aware of much progress in improving our interdependent senses of taste and smell.

**EXTENDING THE PREMISE**

Calculators and computers extended this premise from our physical beings to our brains, helping us to think better, at first by relieving us of mathematical grunt work. But the dream of true thinking machines, devices that can be taught to reason and act on their own, has long loomed in the human psyche. The dream dates back at least to ancient Greece and phenomena such as the tripods that the god Hephaestus purportedly crafted to walk on their own to Mount Olympus and back: early driverless vehicles.

Not surprisingly, in each historic era, the proposed mechanism for the machine «brain» closely mirrored the presumed mechanism for the human brain — from divinely created, to mechanical in the Machine Age, to electro-chemical-digital in our time (cf., the sub-headline of a 2013 published paper: «physicists have developed a technique that can tell which parts of the brain rely on analog signals and which rely on digital signals»). Once created, none of these powerful metaphors entirely disappears from our beliefs. Our imaginations today live with an amalgam of prior and contemporary brain metaphors.

**ADVANCED COMPUTER CAPABILITIES**

As computer capabilities have advanced, these machines — notably supercomputers — have been able to beat the best humans at increasingly complex games, from checkers to chess to television's Jeopardy! quiz show. (In the context of this discussion, we loosely define a game as a non-lethal activity associated with a

## HUMAN THINKING IS HEAVILY GOAL-ORIENTED, TELEOLOGICAL.

# MAYBE THAT'S IN OUR DNA.

goal and rules for reaching it.) One of the next major «games» will be driverless vehicles. In the short run, they may not outperform the best human drivers, but they will certainly outperform the worst, in part because driverless vehicles don't get drunk or distracted by tweets. Another important game will be personalized (precision) medicine, where computers are already starting to be used for decision support for difficult, near-real time diagnoses and treatment plans.

If it seems odd to use the term «game» for such serious pursuits, consider the source. Human thinking is heavily goal-oriented, teleological. Maybe that's in our DNA. If we hadn't learned the game and the rules, Darwin says, we would not have survived as a species. It's not surprising that we have created our proto-thinking machines in our own image, to follow the same fundamental behavior pattern.

### FROM GAMING TO MACHINE LEARNING

Today, the field of machine learning can be defined within the framework of game-playing.

At one end of the spectrum is mainstream machine learning, where the computer is given the goal and then painstakingly trained to play the game reasonably well. This falls under the general headings of supervised learning and its less explicit relative, reinforcement learning.

At the other end today is unsupervised learning, in which the computer is presented with data and must ferret out the goal and the rules on its own.

The best of contemporary unsupervised learning is often called cognitive computing, because it's thought to mimic the cognitive processes of the human brain. Mind you, we humans don't exactly know yet how human cognition works. Figuring out how we acquire and use knowledge is a question that continues to tantalize psychologists, philosophers and brain researchers. Presumably, we understand more about machine cognition, since we are the ones who've been teaching the machines how to think.

### A NARROWING GAP

The gap between human and machine brains will continue to narrow. Farther, probably much farther, in the future there will be artificial intelligence (AI) that passes the Turing test: an evaluator of a text exchange can't tell which participant in the conversation is the human and which is the machine. A more cynical perspective sees AI as a capability that will always lie in the future, because threatened humans will always move the goal post when a machine reaches the one-yard line.

Sci-fi visions to the contrary, we humans needn't worry about machine domination any time soon. It's trite but true that human thinking is far more intricate and nuanced than any realistic roadmap for machine cognition. No one seriously expects a thinking machine to replicate the fullness of Shakespeare, Einstein or Louis Armstrong. At most, we will get Rain Man, a machine (or collection of machines) with exceptional niche intelligence.

# BUT IF HISTORY IS ANY GUIDE, THE MAINSTREAM MARKET FOR MACHINE LEARNING WILL SOON ENOUGH EVOLVE TO THE POINT WHERE GOOD ENOUGH IS NO LONGER GOOD ENOUGH.

## **DISTRIBUTED INTELLIGENCE**

This intelligence likely will be geographically distributed, via the evolving Internet of Things (IoT), with concentrations of networked brain-power varying by location as needed. It may be akin to the view from an airplane traversing a nighttime landscape: bright fusions of light in large cities, more moderate brilliance in smaller towns, and isolated points of light in the countryside, with illuminated strands of roads and highways forming the interconnect fabric.

China, for one, is counting on an HPC architecture called REST 2.0 to manage the country's future IoT. As Dr. Zhiwei Wu, Chinese Academy of Sciences (CAS), outlined at the ISC'13 Big Data conference in Heidelberg, CAS is developing a new server they hope will scale to 1 billion threads, with a dramatically simplified architecture and hardware/software stack — along with storage and an «elastic processor» with a functional instruction set architecture (FISC) design. Closer to home, international e-commerce giant PayPal has been investigating the use of HPC to manage a global graph infrastructure that will embody increasing intelligence about the behaviors of retail shoppers.

Whether the objects of interest are human shoppers, driverless vehicles in traffic or smart electrical grids, it's clear that the IoT will

be far more than just an «Internet of stupid things.» HPC will be needed to move massive data, to process it as quickly and as locally as feasible, and to help maintain the wellness of this network with its distributed intelligence.

## **HPC'S MASSIVE PARALLELISM**

One of HPC's main contributions to the future of machine learning, including its use in the nascent IoT, will be the ability to exploit parallelism. As recent IDC studies have confirmed, few algorithms used by the mainstream machine learning industry have been parallelized. That community got its start in about the year 2000, in an era when generational jumps in the single-threaded performance of x86 processors often improved performance without the need to parallelize. Failure to parallelize algorithms can have important consequences, such as limiting applications to «good enough» speed and resolution, or forcing users to run problems at sub-optimal sizes or, in certain sectors, increasing reliance on costlier physical experimentation.

But if history is any guide, the mainstream market for machine learning will soon enough evolve to the point where good enough is no longer good enough to compete and survive. This is where the HPC community will come to the rescue.

**ASUS**<sup>®</sup>  
IN SEARCH OF INCREDIBLE



# ESC4000 G3

## Green500 Champion Successor

<http://www.asus.com/Server-Workstation/>

Intel, the Intel logo, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel Corporation in the U.S. and/or other countries.

# **INTEL OPTANE**

An ultra fast memory  
technology for personalized  
medicine and analytical  
Big Data

# ANNOUNCING INTEL® OPTANE™ TECHNOLOGY

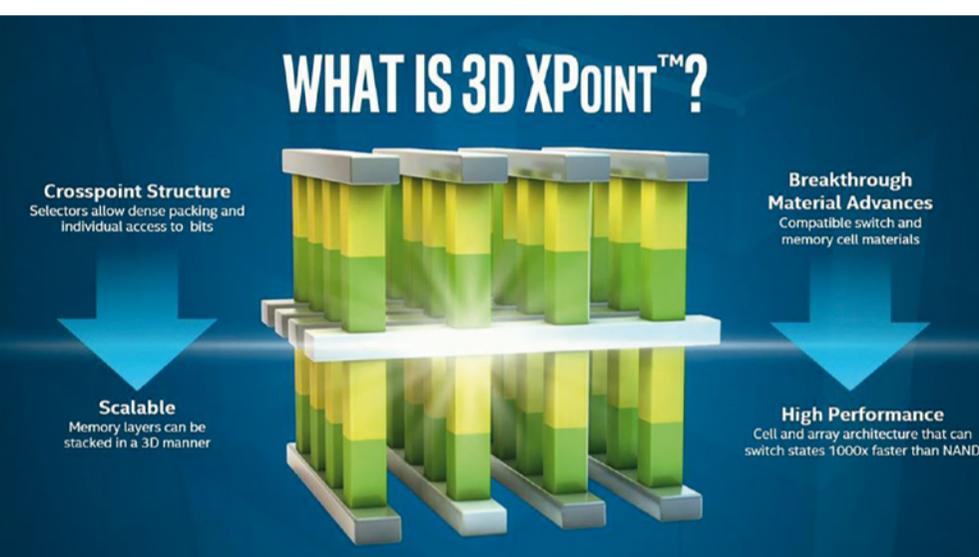
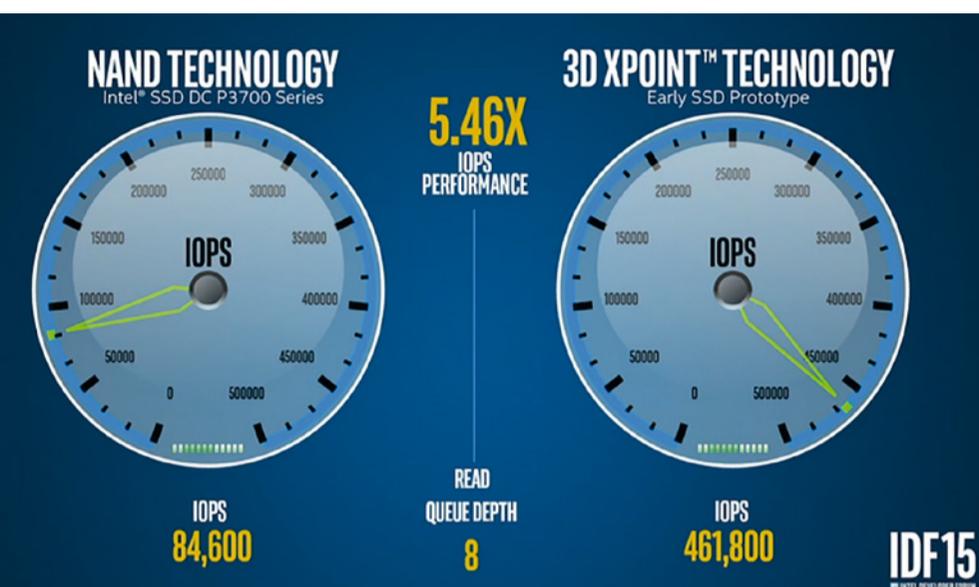
UNLEASHING BREAKTHROUGH PERFORMANCE FOR A NEW GENERATION OF COMPUTING

Based on 3D XPoint™  
STORAGE. MEMORY. CONVERGED

## THE LATEST MEMORY TECHNOLOGY WAS CREATED IN 1989, THE SAME YEAR AS THE WEB.

**D**espite the radical advances in terms of IT architecture over the last 50 years, memory technology has not evolved significantly. The NAND architecture is the latest technology, but it was created in 1989, the same year as the Web. NAND technology is economical and retentive, but its relative slowness requires a combination with DRAM, another old technology, launched a year after Moore's Law in 1966. The introduction of a new architecture is not only expected, but occurs at a critical moment in the expansion of the digital world. Indeed, the current storage is about 10,000 times slower than the fastest

processor memory. Taking the example of supercomputing, the current memory technology creates bottlenecks, which means that existing technology is not, and will never live up to the challenge of exponential growth ahead of volume data. Indeed, by 2020 the volume of data created should be multiplied by 10 to reach 44 Zb, or 44 billion TB. In association with Micron, Intel has developed a new class of memory to solve this problem. Called 3D XPoint, this technology is based on a cross-point architecture without transistors, like a chessboard in 3 dimensions. The memory cells are located at the intersection of word and bit lines, which allows to assign them individual addresses. The data can therefore be read and written in small units, which induces a much more efficient reading / writing process.



**THE BEST OF BOTH WORLDS  
NAND AND DRAM**

3D XPoint combines the best features of existing technologies. Retentive and economic, it is 1000 times faster than NAND technology and multiplies by 10 the capacity of the DRAM. To get an idea of its speed, solid-state drive might take you to the moon in four days (384,000 km), while the 3D technology XPoint would you do the round trip to Mars in as many days (225 million km). Endurance is another plus point of the architecture of 3D XPoint. 1000 times more enduring than NAND, the new technolo-

gy means that businesses and advanced users will no longer have to worry about the read / write cycles. If 3D technology XPoint was a car engine oil, you could do 4.83 million km before having to drain. The first public demonstration of SSD units XPoint 3D technology took place during the 2015 edition of IDF and was conducted by the CEO of Intel, Brian Krzanich. This demonstration helped highlight the performance gains multiplied by 5 or 7 compared to SSD units based on the fastest NAND memories. Mr. Krzanich commented on the performance: «When you offer to developers of products with performance gains multiplied by 5 or 7, you open the way for inventions and innovations outside the norm.»

**A TECHNOLOGY AVAILABLE STARTING  
NEXT YEAR**

Intel confirmed that 3D technology will be commercialized by 2016. XPoint SSD units proposed under the Intel Optane product portfolio should be integrated into all systems, low-power Ultrabooks with supercomputing systems. A new line of Intel DIMMs will also be designed for data center platforms.

These technologies have the potential to revolutionize many industries. The Optane storage devices are ideal for storing and processing large volumes of multimedia content, as well as immersive games, grounded virtualization in the cloud, to resident databases in memory and high-performance computing. Intel plans significant breakthroughs in medical applications and analysis of Big Data. This memory technology also allows companies and cities to process and analyze large volumes of data to develop smarter systems.

**3D XPOINT COMBINES THE BEST FEATURES OF EXISTING TECHNOLOGIES. RETENTIVE AND ECONOMIC, IT IS 1000 TIMES FASTER THAN NAND TECHNOLOGY AND MULTIPLIES BY 10 THE CAPACITY OF THE DRAM.**



“A multi-product roadmap to bring unprecedented creation capabilities to millions of designers.”



**SANDEEP GUPTE**  
**SENIOR DIRECTOR OF**  
**NVIDIA PROFESSIONAL**  
**SOLUTIONS GROUP**



# ALL THESE PRODUCTS WILL WORK TOGETHER TO GIVE DESIGNERS IN A VAST ARRAY OF INDUSTRIES POWER THAT WAS — UNTIL NOW — AVAILABLE TO JUST A HANDFUL.

What products are you announcing?

We are announcing a multi-product roadmap to bring unprecedented creation capabilities to millions of designers. It has three main pieces:

**IRAY 2015** — the latest version of our GPU-accelerated rendering software, with new features to support exchanging materials across design applications, scalability outside of a workstation and much faster rendering speed.

**QUADRO M6000** — our most powerful professional GPU, featuring our Maxwell architecture and 12GB of graphics memory to support complex designs.

**QUADRO VISUAL COMPUTING APPLIANCE** — upgraded with eight M6000-class GPUs, this scalable appliance achieves unprecedented speed and visual fidelity, and is specifically tuned to accelerate our Iray software.

All these products will work together to give designers in a vast array of industries power that was — until now — available to just a handful.

What is NVIDIA Iray?

NVIDIA Iray is a highly interactive and intuitive physically based rendering technology that generates photorealistic imagery by simulating the physical behavior of light and materials. Unlike traditional production renderers, Iray delivers

results reflecting real-world behaviors. Designers don't need expert knowledge of computer graphics techniques to quickly achieve photorealistic results. Iray progressively refines its image until completion, providing consistent results from interactive editing to final frame. It's a highly predictive approach that marries with the scalable, world-class performance across NVIDIA GPUs to give constant feedback and rapid results. This slashes the time needed to perfect scenes and deliver images rivaling photographs.

Who is Iray aimed to?

Iray is primarily for designers and visual effects artists working with 3D content who require predictable photorealistic imagery. It delivers immediate visual feedback that results in stunning imagery for everything from architecture, engineering, and design to marketing and advertising visual effects.

Can you describe Iray's main features?

NVIDIA Iray rendering simulates real world lighting and practical material definitions so that anyone can interactively design and create the most complex of scenes. Predictable and physically based immediate visual feedback results in stunning imagery for architecture, engineering and design to marketing and advertising visual effects.



Iray consists of three rendering modes, Photoreal, Interactive and Realtime to produce varying degrees of physically based and photorealistic imagery with different performance characteristics. The rendering modes can be customized with unique parameters to achieve the desired interactivity and photorealism required for your application. They offer a seamless user experience when switching between rendering modes as each mode uses the same scene and material descriptions stored in memory within the Iray database.

### **IRAY PHOTOREAL**

Iray Photoreal is a high-performance, global illumination rendering mode that generates photorealistic imagery by simulating the physical behavior of light interaction with surfaces and volumes. Iray balances ease of use with the highest achievable quality for photorealistic final frame output. Images are progressively refined to provide the full global illumination including caustics, sun studies and luminance distributions.

Compositing elements that isolate lighting components using a technique called Light Path Expressions (LPE) and saved to output

buffers are now available in Iray. Rendering the compositing elements are done in parallel with the final image output. The buffers can be edited in post-processing using standard imaging or compositing software and the results can be integrated into the final output image without time consuming and expensive re-rendering. Iray generates the standard scene specific render elements such as alpha channel, z-depth, normals, texture coordinates, material id, object id and irradiance. Common LPEs elements are: indirect illumination, direct diffuse, direct specular, indirect diffuse, sub surface scattering, caustics and many more. LPEs can be defined on a per object and per light basis allowing for complete creative control in post.

### **IRAY INTERACTIVE**

Iray Interactive rendering mode uses advanced approximation algorithms to mimic realism while minimizing unwanted noise - maximizing performance during interaction. Path space filtering, a technique to intelligently remove noise in the final frame during navigation and interaction is available for Iray Interactive. Other computationally expensive features such



as indirect light can be kept on or off to support noise free interactive frame rates with single GPU systems. The ability to adapt quality to the type of GPU being used ensures that the Interactive render mode can deliver suitable image quality during interactive tasks. Iray Interactive can render screen space ambient occlusion effects generating more realistic and physically plausible results interactively.

### **IRAY REALTIME**

Iray Realtime is an OpenGL render mode which uses conventional GPU raster techniques to achieve realtime rendering speed. Iray Realtime targets a look for the materials that is closely matching the look in the other render modes. However, it has limitations when it comes to accurate reflections and refractions as well as lighting simulations. Iray Realtime leverages OpenGL 3.3 and NVIDIA extensions. It is ideal where rendering speed or large display support has priority and limited photorealism is acceptable.

### **A SCALABLE RENDERING SOLUTION**

With Iray, we bring Interactive, Scalable, Physically Based Rendering to Millions. Throughout 2015, NVIDIA will bring Iray to several more 3D creation applications, including Autodesk's 3ds Max, Maya, Revit and McNeel Rhinoceros. DAZ 3D has also made Iray available to its customers. This means millions of designers will now have access to Iray's capabilities, including Iray material definition language (MDL), which allows physically based materials to be interchangeable across apps, so designers can switch from one tool to another and get consistent results. Iray 2015 is supporting the latest measurement format from X-Rite, while MDL is being supported by a growing number of companies that allow designers to create physically based materials, including Allegorithmic and Old Castle.

### **MULTIPLE GPUS IN ONE MACHINE**

Iray supports any number or combination of GPUs and CPUs running within the same machine, and gives you control over which



are being used for rendering. Load balancing between GPUs of different sizes, types, and generations is automatic—often doubling offline rendering performance when doubling the number of equivalent processors. Even older systems quickly become powerful rendering nodes by adding NVIDIA GPUs. This makes multi-GPU systems the preferred choice among Iray artists. The superior performance of Iray Photoreal mode begins by using each processor in your machine to its fullest. Thorough multi-threading and SIMD optimization squeezes the most from each CPU core, while scaling across the thousands of cores within NVIDIA GPUs delivers maximum speed.

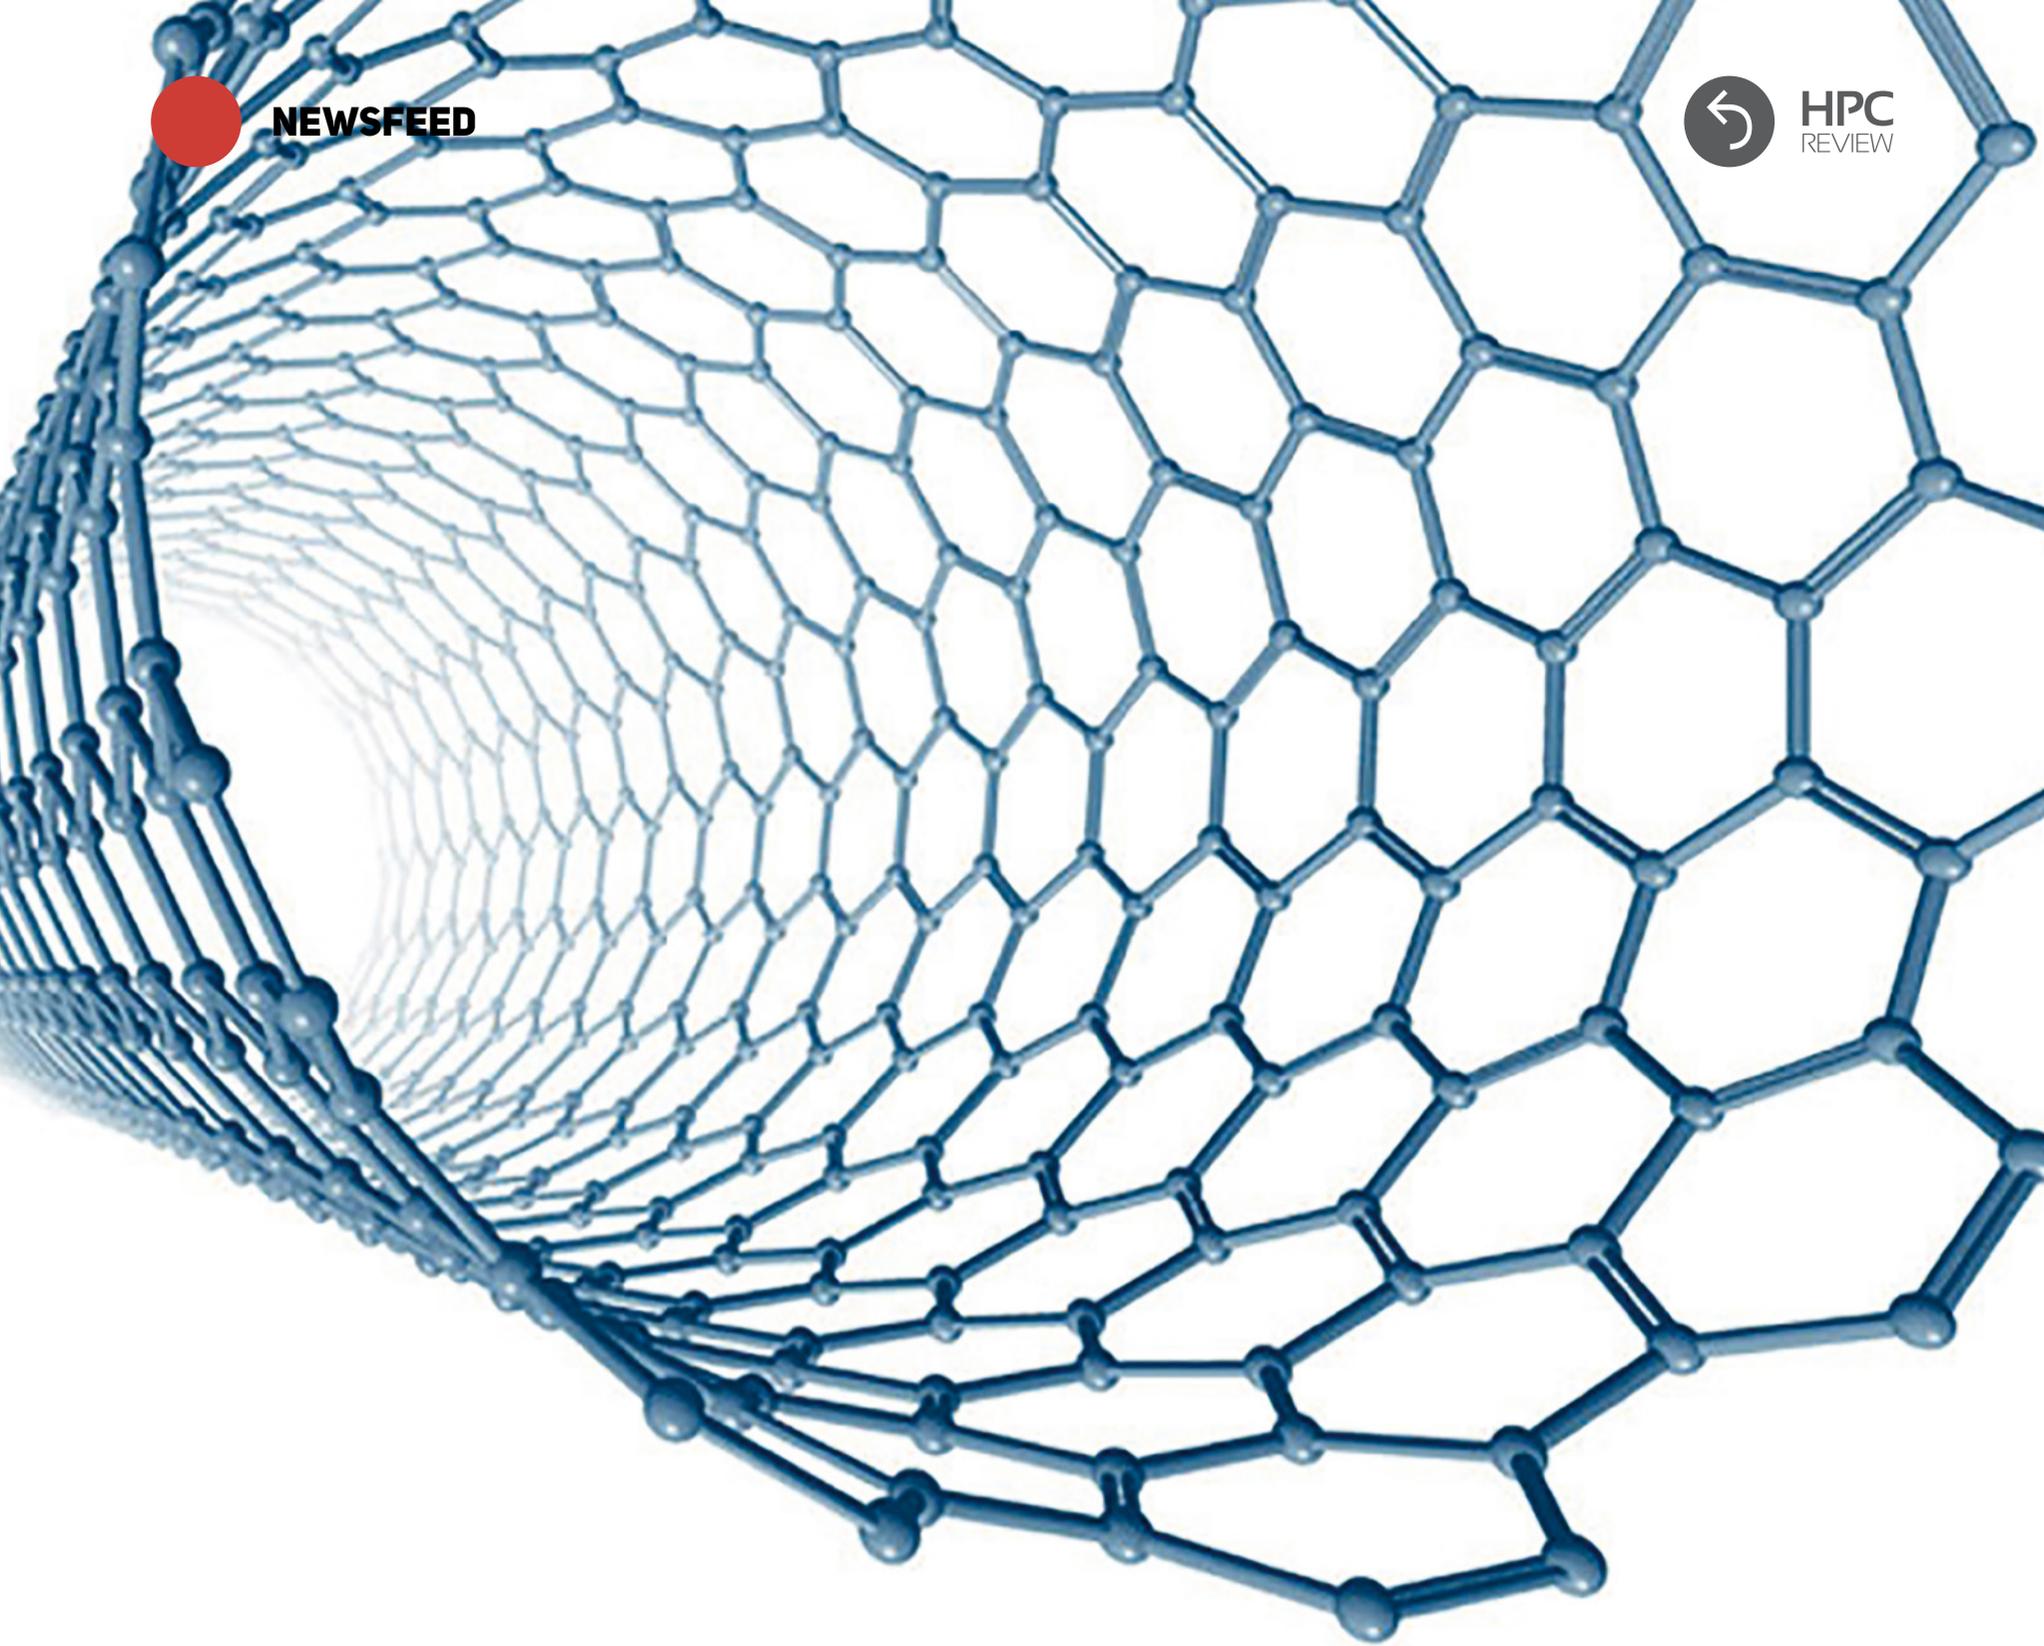
## What are the new Quadro Professional Graphics Cards main features ?

The Quadro M6000 is a ultra high-end card targeted towards extreme CAD/CAM users, photorealistic rendering, extreme DCC effects and seismic and 3D medical visual exploration as well as massive GPGPU compute tasks. It features 12GB of onboard GDDR5 memory and is about 61% faster than previous generation.



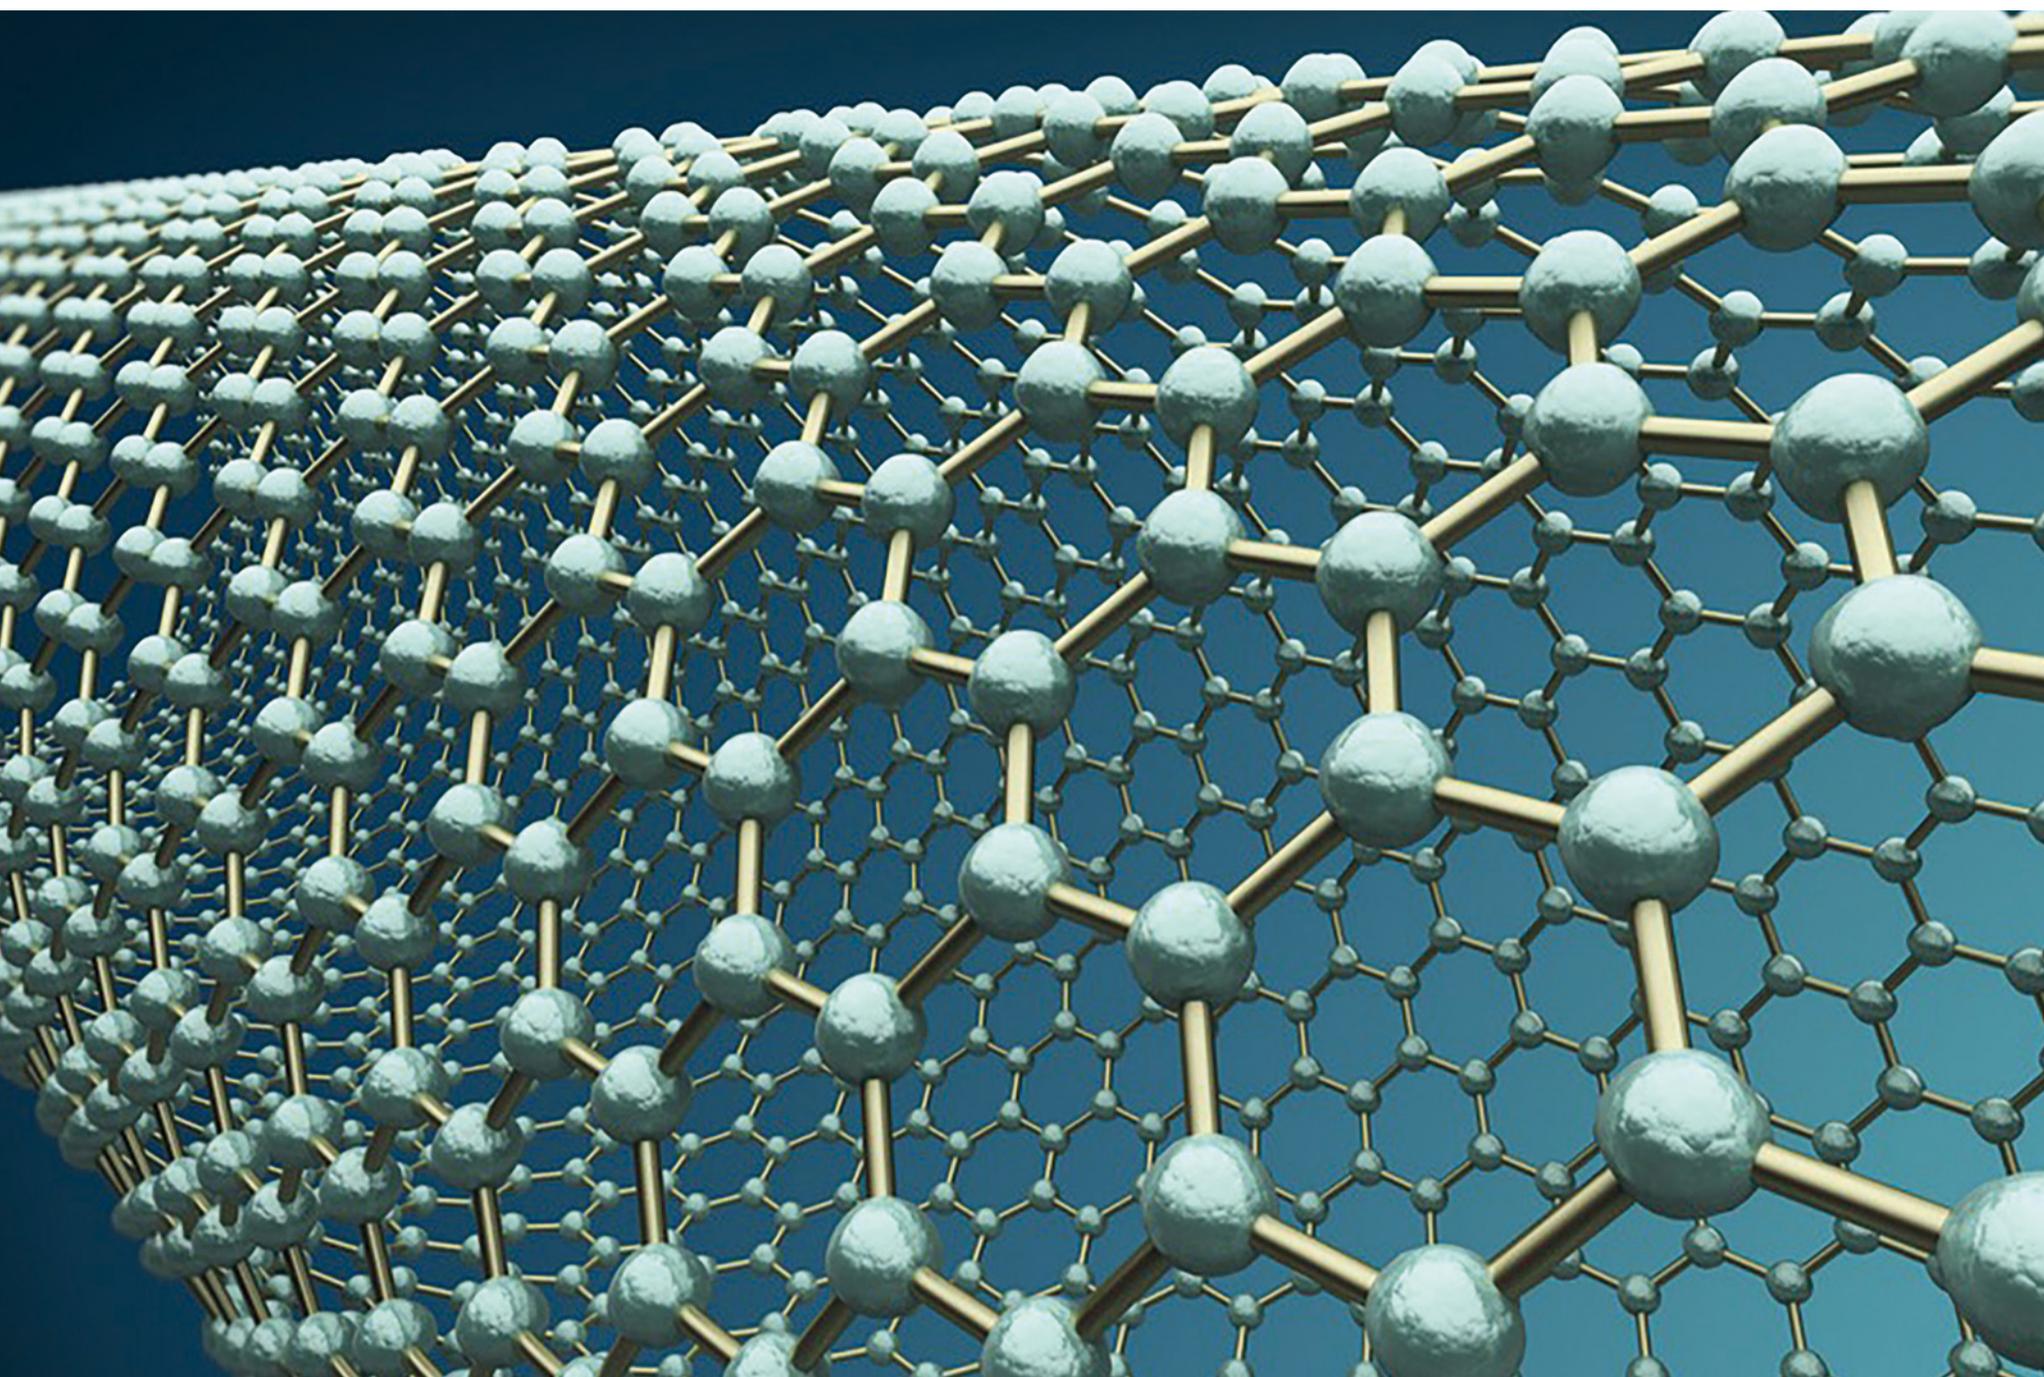
It features 3072 CUDA cores versus 2880 for the K5000. We also launched the Quadro M5000, built around the Maxwell Core and PCIe 3.0 compatible. Single precision performance is 4.3 Tflops (versus 3.6 for the K5200 built on the Kepler architecture). It has also 4 x Displayports and one DVI (versus 2 DP + 2 DVI), allowing up to four 4K displays. Power consumption is only 150W versus 225W.

The new M4000 replaces the K4200, and offers 1664 CUDA cores (versus 1344) and +50% performance in Iray with 2.6 Tflops versus 2.1. Memory has doubled to 8GB like the M5000, and up to four 4K displays can be connected simultaneously.



# IBM INVENTS THE SUCCESSOR TO SILICON

IBM announced a major engineering breakthrough that could accelerate carbon nanotubes replacing silicon transistors to power future computing technologies.



BM scientists demonstrated a new way to shrink transistor contacts without reducing performance of carbon nanotube devices, opening a pathway to dramatically faster, smaller and more powerful computer chips beyond the capabilities of traditional semiconductors. IBM's breakthrough overcomes a major hurdle that silicon and any semiconductor transistor technologies face when scaling down. In any transistor, two things scale: the channel and its two contacts. As devices become smaller, increased contact resistance for carbon nanotubes has hindered performance gains until now. These results could overcome contact resistance challenges all the way to the

1.8 nanometer node – four technology generations away.

Carbon nanotube chips could greatly improve the capabilities of high performance computers, enabling Big Data to be analyzed faster, increasing the power and battery life of mobile devices and the Internet of Things, and allowing cloud data centers to deliver services more efficiently and economically.

**SILICON: REACHING THE POINT OF PHYSICAL LIMITATION**

**SILICON: REACHING THE POINT OF PHYSICAL LIMITATION**

Silicon transistors, tiny switches that carry information on a chip, have been made smaller year after year, but they are approaching a point of physical limitation. With Moore's Law running out of steam, shrinking the size

of the transistor – including the channels and contacts – without compromising performance has been a vexing challenge troubling researchers for decades.

IBM has previously shown that carbon nanotube transistors can operate as excellent switches at channel dimensions of less than ten nanometers – the equivalent to 10,000 times thinner than a strand of human hair and less than half the size of today’s leading silicon technology. IBM’s new contact approach overcomes the other major hurdle in incorporating carbon nanotubes into semiconductor devices, which could result in smaller chips with greater performance and lower power consumption.

### **A \$3 BILLION R&D INVESTMENT**

Earlier this summer, IBM unveiled the first 7 nanometer node silicon test chip, pushing the limits of silicon technologies and ensuring further innovations for IBM Systems and the IT industry. By advancing research of carbon nanotubes to replace traditional silicon devices, IBM is paving the way for a post-silicon future and delivering on its \$3 billion chip R&D investment announced in July 2014. “These chip innovations are necessary to meet the emerging demands of cloud computing, Internet of Things and Big Data systems,” said Dario Gil, vice president of Science & Technology at IBM Research. “As silicon technology nears its physical limits, new materials, devices and circuit architectures must be ready to deliver the advanced technologies that will be required by the Cognitive Computing era. This breakthrough shows that computer chips made of carbon nanotubes will be able to power systems of the future sooner than the industry expected.”

### **A NEW CONTACT FOR CARBON NANOTUBES**

Carbon nanotubes represent a new class of semiconductor materials that consist of single atomic sheets of carbon rolled up into a tube. The carbon nanotubes form the core of a transistor device whose superior electrical proper-

## **CARBON NANOTUBE TRANSISTORS CAN OPERATE AS EXCELLENT SWITCHES AT CHANNEL DIMENSIONS OF LESS THAN TEN NANOMETERS, THE EQUIVALENT TO 10,000 TIMES THINNER THAN A STRAND OF HUMAN HAIR.**

ties promise several generations of technology scaling beyond the physical limits of silicon.

Electrons in carbon transistors can move more easily than in silicon-based devices, and the ultra-thin body of carbon nanotubes provide additional advantages at the atomic scale. Inside a chip, contacts are the valves that control the flow of electrons from metal into the channels of a semiconductor. As transistors shrink in size, electrical resistance increases within the contacts, which impedes performance. Until now, decreasing the size of the contacts on a device caused a commensurate drop in performance – a challenge facing both silicon and carbon nanotube transistor technologies.

IBM researchers had to forego traditional contact schemes and invented a metallurgical process akin to microscopic welding that chemically binds the metal atoms to the carbon atoms at the ends of nanotubes. This ‘end-bonded contact scheme’ allows the contacts to be shrunken down to below 10 nanometers without deteriorating performance of the carbon nanotube devices.

### **A TECHNOLOGY AVAILABLE WITHIN THE DECADE**

“For any advanced transistor technology, the increase in contact resistance due to the decrease in the size of transistors becomes a major performance bottleneck,” Gil added. “Our novel approach is to make the contact from the end of the carbon nanotube, which we show does not degrade device performance. This brings us a step closer to the goal of a carbon nanotube technology within the decade.”

When data grows & demand increases,  
but budgets don't.

# Meet the all-new **Saber 1000**

SATA 3.0 Enterprise SSD series.

## Features

- OCZ's Barefoot 3 controller, in-house firmware, and next generation A19nm NAND from Toshiba
- Delivers sustained performance and consistent I/O responses for high-volume deployment enterprise SATA SSDs
- Supports Power Failure Management Plus (PFM+) to protect against unexpected power loss events
- Central management capability via OCZ StoragePeak 1000 software suite
- Available up to 960GB



Delivers great performance for today's typical datacenter applications



Read Cache  
& Indexing



Virtual Desktop  
Infrastructure  
(VDI)



Front-End  
Web Server



Media  
Streaming



Video on Demand  
(VoD)



Decision Support  
System (DSS)



Cloud  
Infrastructure



Customer  
Relationship  
Management (CRM)



Online  
Archiving



Video Editing/  
Photo Sharing



Enterprise Content  
Management  
(ECM)



Virtual Tape  
Library (VTL)



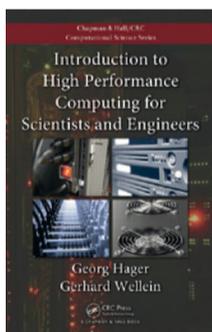
# books

## **INTRODUCTION TO HIGH PERFORMANCE COMPUTING FOR SCIENTISTS AND ENGINEERS**

Georg Hager, Gerhard Wellein

**Chapman & Hall**

**360 pages, £50.99 / eBook £35.69**



The text first introduces the architecture of modern cache-based microprocessors and discusses their inherent performance limitations, before describing general optimization strategies for serial code on cache-based architectures. It next covers shared- and distributed-memory parallel computer architectures and the most relevant network topologies. After discussing parallel computing on a theoretical level, the authors show how to avoid or ameliorate typical performance problems connected with OpenMP. They then present cache-coherent nonuniform memory access (ccNUMA) optimization techniques, examine distributed-memory parallel programming with message passing interface (MPI), and explain how

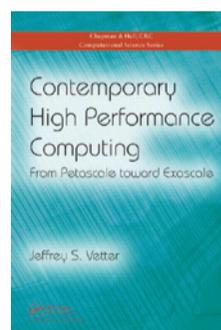
to write efficient MPI code. The final chapter focuses on hybrid programming with MPI and OpenMP.

## **CONTEMPORARY HIGH PERFORMANCE COMPUTING: FROM PETASCALE TOWARD EXASCALE**

Jeffrey S. Vetter

**Chapman & Hall**

**730 pages, £63.99 / eBook £44.79**



The first part of the book examines significant trends in HPC systems, including computer architectures, applications, performance, and software. It discusses the growth from terascale to petascale computing and the influence of the TOP500 and Green500 lists. The second part of the book provides a comprehensive overview of 18 HPC ecosystems from around the world. Each chapter in this section describes programmatic motivation for HPC and their important applications; a flagship HPC system overview covering computer ar-



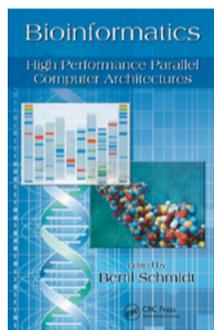
chitecture, system software, programming systems, storage, visualization, and analytics support; and an overview of their data center/facility. The last part of the book addresses the role of clouds and grids in HPC, including chapters on the Magellan, FutureGrid, and LL-Grid projects. With contributions from top researchers directly involved in designing, deploying, and using these supercomputing systems, this book captures a global picture of the state of the art in HPC.

### **BIOINFORMATICS: HIGH PERFORMANCE PARALLEL COMPUTER ARCHITECTURES**

Bertil Schimdt

Chapman & Hall

370 pages, £102.00 / eBook £71.40



New sequencing technologies have broken many experimental barriers to genome scale sequencing, leading to the extraction of huge quantities of sequence data. This expansion of biological databases established the need for new ways to harness and apply the astounding amount of available genomic information and convert it into substantive biological understanding. This book discusses how to take advantage of bioinformatics applications and algorithms on a variety of modern parallel architectures. Two factors continue to drive the increasing use of modern parallel computer architectures to address problems in computational biology and bioinformatics: high-throughput techniques for DNA sequencing and gene expression analysis—which have led to an exponential growth in the amount of digital biological data—and the multi- and many-core revolution within computer architecture.

## mooCS



### **IMPLEMENTING REAL-TIME ANALYTICS WITH HADOOP IN AZURE HDINSIGHT**

Starts on November 2, 2015

Learn how to use Hadoop technologies like HBase, Storm, and Apache Spark in Microsoft Azure HDInsight to create real-time analytical solutions. Learn how to implement low-latency and streaming Big Data solutions using Hadoop technologies like HBase, Storm, and Spark on Microsoft Azure HDInsight. You'll learn to create and query HBase and Apache Spark tables to address problems and find solutions using Big Data. This course is the second in a series that explores big data and advanced analytics techniques with HDInsight; and builds on the batch processing techniques learned in DAT202.1x: Processing Big Data with Hadoop in Azure HDInsight. A third course in this series will follow this one to cover predictive analytics in HDInsight. Note: To complete the hands-on elements in this course, you will require an Azure subscription and a Windows client computer. You can sign up for a free Azure trial subscription (a valid credit card is required for verification, but you will not be charged for Azure services). Note that the free trial is not available in all regions.

However, it is possible to complete the course and earn a certificate without completing the hands-on practices.

**Length:** 3 weeks **Effort:** 3 - 4 hours/week  
**Price:** FREE  
**Institution:** Microsoft  
**Subject:** Computer Science  
**Level:** Intermediate  
**Languages:** English  
**Video Transcripts:** English  
**Link :** <https://www.edx.org/course/implementing-real-time-analytics-hadoop-microsoft-dat202-2x#!>

## **COMPUTATION STRUCTURES – PART 2: COMPUTER ARCHITECTURE**

Starts on November 17, 2015

In this interactive computer science course from MIT you'll learn the basic principles of computer architecture. Digital systems are at the heart of the information age in which we live, allowing us to store, communicate and manipulate information quickly and reliably. This computer science course is a bottom-up exploration of the abstractions, principles, and techniques used in the design of digital and computer systems. If you have a rudimentary knowledge of electricity and some exposure to programming, roll up your sleeves, join in and design a computer system! Using your browser for design entry and simulation, you'll implement a 32-bit computer using our gate library and write assembly language programs to explore the hardware/software interface. This is Part 2 of a 3-part program on digital systems.

**Length:** 12 weeks **Effort:** 6 hours/week  
**Price:** FREE / Add a Verified Certificate for \$49  
**Institution:** MITx  
**Subject:** Computer Science  
**Level:** Introductory  
**Languages:** English  
**Video Transcripts:** English  
**Link :** <https://www.edx.org/course/computation-structures-part-2-computer-mitx-6-004-2x#!>

## **PREDICTIVE ANALYTICS**

Starts on November 25, 2015

Master the tools of predictive analytics in this statistics based analytics course. Decision makers often struggle with questions such as: What should be the right price for a product? Which customer is likely to default in his/her loan repayment? Which products should be recommended to an existing customer? Finding right answers to these questions can be challenging yet rewarding. Predictive analytics is emerging as a competitive strategy across many business sectors and can set apart high performing companies. It aims to predict the probability of the occurrence of a future event such as customer churn, loan defaults, and stock market fluctuations – leading to effective business management. Models such as multiple linear regression, logistic regression, auto-regressive integrated moving average (ARIMA), decision trees, and neural networks are frequently used in solving predictive analytics problems. Regression models help us understand the relationships among these variables and how their relationships can be exploited to make decisions. This course is suitable for students/practitioners interested in improving their knowledge in the field of predictive analytics. The course will also prepare the learner for a career in the field of data analytics. If you are in the quest for the right competitive strategy to make companies successful, then join us to master the tools of predictive analytics.

**Length:** 7 weeks **Effort:** 4 - 5 hours/week  
**Price:** FREE / Add a Verified Certificate for \$25  
**Institution:** IIMBx  
**Subject:** Data Analysis & Statistics  
**Level:** Advanced  
**Languages:** English  
**Video Transcripts:** English  
**Link :** <https://www.edx.org/course/predictive-analytics-iimbx-qm901x#!>

**CHIFFRES CLÉS**

**44 BILLION DOLLARS**

Worldwide projected HPC market value by 2020

**8,3%**

Yearly growth of HPC market

**220 BILLION DOLLARS**

Compound market value over the 2015-2020 period

Source : Market Research Media



**TOP 500**  
TOP 3

**1 TIANHE-2**  
National Supercomputing Center, Canton :  
**33863 / 54902 TFlops** Manufacturer NUDT  
Architecture Xeon E5-2692 + Xeon Phi 31S1P, TH Express-2

**2 TITAN**  
Oak Ridge National Laboratory,  
USA : **17590 / 27113 TFlops** Manufacturer  
Cray XK7 Architecture Opteron 6274 + Nvidia Tesla K20X, Cray Gemini Interconnect

**3 SEQUOIA**  
Lawrence Livermore National Laboratory,  
USA : **17173 / 20133 TFlops** Manufacturer  
IBM Blue Gene/Q Architecture PowerPC A2

The TOP500 classes every six months the 500 most powerful supercomputers in the world. The retained values, RMAX and RPEAK represent the maximum and theoretical Linpack computing power.

**GREEN 500**  
TOP 3

**1 7031,6 MFLOPS/W**  
RIKEN Shoubu (Japan)

**2 6952,2 MFLOPS/W**  
Suiren Blue High Energy Accelerator  
Research Organization /KEK (Japan)

**3 6217 MFLOPS/W**  
Suiren High Energy Accelerator Research  
Organization /KEK (Japan)

Green 500 list ranks the most energy efficient supercomputers in the world. Energy efficiency is assessed by measuring performance per Watt. The unit here is the MFLOPS / Watt.



# 4 Blades

## HighServer XLR4 Blade

Virtualization Server with GPUs / Co-Processors

- High density
- Low power consumption
- Easy management

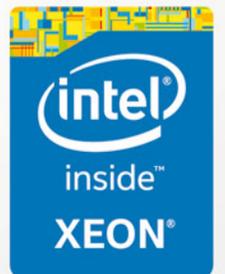
# GIGABYTE™

# 4 GPGPU

Supports the  
Intel® Xeon® Processor E3-1200 V4  
Intel Inside®. Powerful Solutions Outside.

Available at:

**CARRI**  
Systems



Intel, the Intel logo, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel Corporation in the U.S. and/or other countries.



# BIG DATA

## THE TECHNOLOGIES TO ACHIEVE SMART DATA

-BIG DATA-

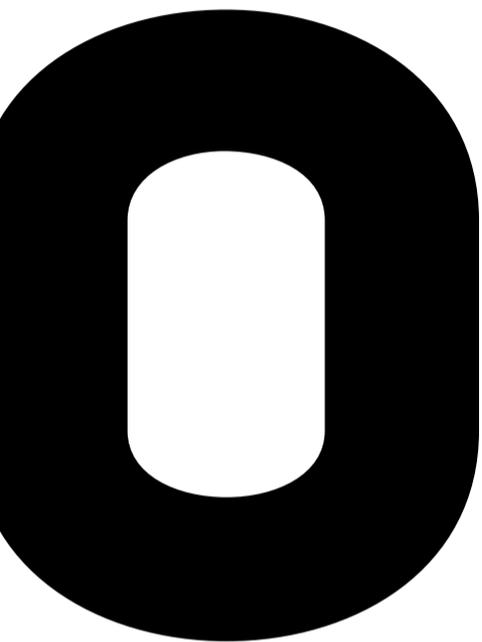




**EVERY MINUTE  
ARE EXCHANGED  
639.800  
GIGABYTES  
OF DATA ON  
THE INTERNET,  
INCLUDING 204  
MILLION E-MAILS  
AND TWO MILLION  
SEARCH QUERIES.**



# The market is growing rapidly and Gartner expects Big Data should create 4.4 million jobs worldwide by mid-2016.



Observing the mass of information that we create every day is enough to make your head spin: With over 2.5 trillion bytes of new data, 2.5 billion billion bytes, they gather each day as many information that the human mind was able to produce since its origin.

Sensors integrated in our immediate environment to trade on social networks, through our purchasing habits and online transactions, we participate every second enrich this voluminous common encyclopedia. Big Data is the result of this decision abiding word, an expression denoting the infinite mass of data collected every day by businesses worldwide. But the collection itself is only the sunny side of the mountain to climb. The shady side is the proper analysis of these data, through powerful and modern tools to root out the true meaning and reorient decision to more pragmatic choices in better agreement with customer demands. All business sectors are affected by the phenomenon, science or medicine, fields traditionally behind phenomenal databases, to smaller SMEs operating in a niche market and for whom Big Data is the fuel needed to their business. For example, remember that Facebook is raking every day 500 terabytes of data and Twitter 80MB... per second ! Unstructured, voluble, using multiple formats (text,

image, video, audio ...), stamped, geolocated and inherently «noisy», this information must be analyzed to extract the true meaning. Epidemiologists already use Twitter to map the evolution of flu or gastroenteritis, by correlating the keywords and geolocation of tweets. Analyzing this information allows predictive uses : like the famous «trends» of Twitter, companies can identify any type of movements and orientations from the collected data. But the Big Data revolution is not confined to purely external and public sources such as social networks. Companies are also invited to peruse their data sources, through channels of their own. Discover how this set of technologies can provide a substantial competitive advantage.

The Internet and IP networks have now mutated into huge pathways of infinite channels, carrying at any time of day or night unimaginable amount of data. It's simple: every minute are exchanged 639,800 gigabytes of data on the Internet, including 204 million e-mails, two million search queries and more than 600,000 electronic transactions, according to an IBM study ([www.ibmbigdatahub.com/video/big-data-speed-business](http://www.ibmbigdatahub.com/video/big-data-speed-business)). A prodigious source of information on the habits and expectations of consumers, but also on the fundamental changes in human activities! By entering fully into the digital age, businesses and society as a whole produce and brew constantly growing volumes of important data. In plain meaning the information is in the digital economy what coal was to the industrial economy: its main fuel and a phenomenal growth driver for all



companies. A recent EMC study indicates that 74% of French companies believe that Big Data facilitates decision-making. They are 47% to find that this set of technologies allows the ascension of the market leaders and 23% believe it creates competitive advantages. But they must still be able to collect quality data and give them meaning by analyzing and fine interpretation in light of business objectives! To some extent, companies are accustomed for decades to the feedback and the performance of their traditional distribution channels to drive their business, but the corresponding data was confined to internal databases, embedded in frozen applications on their own information system. With the rise of the Web, the sources are gaining in scope and diversity. The challenge of Big Data is precisely to give them meaning and to optimize decision-making, react almost in real time and enjoying a much more attractive return on investment.

### DECISION-MAKING DATA

Big Data demonstrates greater responsiveness and greater flexibility. In a network of ready-to-wear shops, for example, we can realize in light of sales that a majority of customers buying a certain dress goes with a specific shoe model. By bringing the two items on the shelves, it maximizes the effectiveness of cross-selling and increase profits accordingly. In the traditional structure of the sales force, such a decision could be taken only after the tedious manual analysis of receipts or intuition of a local store manager. This means that the dress season is likely to be almost over before doing the right inventory decisions! Using Big Data, it is conceivable to inject all instant sales data in a heterogeneous database, which may include customer tweets («how does one speak of my brand on social networks?») their Instagram pictures («how does one wear or use my products every day?») and the reactions of their relatives on Facebook («what is the popularity of my products to the community?»). By crossing requests and effectively questioning this mass of information, through the questio-

### TYPICAL BIG DATA PLATFORMS

Analytical Databases	42,10%
Operational Data warehouses	39,40%
Cloud-based data solutions	39,00%
On premise Data hosting solutions	33,60%
Datamart	30,10%
NoSQL platforms	21,60%
Hadoop & subprojects	16,20%
Other	0,40%

ning of matured intuitions and experiences, we obtain in an instant the major trends in order to take much more useful decisions for the growth of the company. But what seems clear in the field of marketing and trade applies with equal efficiency in all sectors. Human resources, health, automotive industry or aviation, utilities, leisure industry, connected objects, media and even internal processes: all these sectors increase their efficiency and productivity through Big Data. Or more specifically, Big Data enlighten their operation and contributes to optimize their performance.

### THE THREE DIMENSIONS OF BIG DATA

Specifically, the Big Data covers three dimensions that the company must master in order to effectively drive its decisions: the concept of volume, velocity and variety - we speak of the «rule of 3 V». The volume is firstly the huge amount of data we started talking and are now within the reach of companies. Like the publications exchanged tweets or on social networks, some of them are basically new. But many of them come from sensors and conventional recovery tools such as annual surveys of electricity meters for example. In this specific area, it is estimated that the analysis of such a volume of data helps identify faster, and even anticipate, incidents on the distribution network and to orchestrate a more rational energy consumption. In all cases, it is no longer uncommon to treat a

## WHY LAUNCH A BIG DATA PROJECT / BY SECTOR

Industry sector	CAS D'UTILISATION					
	Case treatment speed	Combine Hybrid Data	Anticipate Data Treatment	Use Diffused Data	Structure Data	Online Archiving
Finance	28,0%	15,9%	18,3%	11,0%	14,6%	12,2%
Sales	28,3%	21,7%	15,2%	10,9%	15,2%	8,7%
Indsutry	22,4%	20,4%	16,3%	16,3%	12,2%	12,2%
Public Services	21,6%	17,5%	17,5%	13,4%	12,4%	17,5%
DevOps	22,0%	13,6%	13,6%	13,6%	20,3%	16,9%
Health	20,8%	22,9%	12,5%	6,3%	16,7%	20,8%
Construction	20,3%	21,5%	15,2%	20,3%	10,1%	12,7%

By querying the DBMS with a classical web language like JavaScript you can enjoy a greater data treatment flexibility.

volume exceeding tens of terabytes or even petabytes (1000 terabytes). In comparison, it is estimated that one creates and exchanges more than two zettabytes (two million petabytes) on the Internet annually. The velocity refers to the speed of analysis and decision making. For chrono-sensitive processes, such as anomaly detection and fraud, but also for decisions that have an almost immediate impact on the level of sales (like our example of ready-to-wear shop), companies must be able to analyze data in real time. Finally, the range corresponds to the plurality of information collected. Unlike the conventional process of information processing, which aims to define clearly their field and will be treated together with similar equipment, the data from the Web are unstructured in nature and can include both text as sensor data, audio, video, location-based information and activity logs. By combining these three dimensions, big data allows not only new treatments, through the

analysis of data that were beyond the reach of the enterprises, but ensures greater responsiveness decision. Where it sometimes took several days to weeks of analysis to make sense of the information collected, Big Data offers a treatment response in the order of a minute.

### THE END OF A MYTH

However, Big Data technologies do not fall in any of the miraculous solution and it is not enough to store a greater volume of data than the competition to spontaneously enjoy a substantial advantage. The concern in the first place for companies to take the plunge, it is of course the ROI. For storing a large volume of data for processing, not to mention the operations to perform to interpret them, present a serious cost that must be profitable. For example, the offer of OVH for Big Data dedicated servers include a cluster of 48 TB of storage for 1000 euros per month ([https://www.ovh.com/fr/serveurs\\_dedies/big-data](https://www.ovh.com/fr/serveurs_dedies/big-data)). On the



## OBSTACLES, BY SECTOR

Industry sector	OBSTACLES						
	Share-holders	Strategy	Bad data management	Absence of Hadoop MongoDB Specialists	Complexity of deployment	No apps management	Other
Health	32,1%	22,6%	15,1%	11,3%	5,7%	11,3%	1,9%
DevOps	29,5%	20,5%	15,4%	16,7%	12,8%	5,1%	0,0%
Finance	28,3%	22,8%	16,3%	15,2%	12,0%	5,4%	0,0%
Industry	27,5%	20,0%	22,5%	12,5%	10,0%	7,5%	0,0%
Construction	25,6%	25,6%	18,2%	9,8%	11,0%	9,8%	0,0%
Sales	25,5%	23,6%	20,0%	14,5%	9,1%	7,3%	0,0%
Public Services	18,9%	24,2%	21,1%	12,6%	14,7%	7,4%	1,1%

side of Amazon Web Services, the supply is more fragmented but its integrated Amazon Redshift service, has a cost of around 1000 dollars per terabyte per year. According to the EMC study, 60% of companies confirm that the budget is the first factor in decision making, to engage in Big Data. They are almost 41% to delay their accession to this new wave of tools, arguing the lack of visibility into the return on investment. Another relatively difficulty to quantify : the ethical questions about the use of the collected information, and the specificity of local regulations on data protection.

### ISSUES AND TECHNOLOGIES

We are now immersed in a digital ocean, which includes both the data produced by traditional computers, but also increasingly so-called «digital noise» that is to say all data generated by our devices, such as a smartphone and any type of geolocation traces, but also the data that correspond to the actions you are doing on the web, our presence on social networks, connected objects, etc. With the phenomenon of Big Data, we have the means to record, capture, store and analyze everything, «says Bernard Ourghanlian, technical and security director at Microsoft France, during Microsoft's Techdays 2015. If companies are beginning to grasp the enhancement of information and trying to get it to help steer their decisions,

they must however retain the rule that 3V we mentioned previously, particularly the huge variety of data they are likely to peruse. Faced with these so heterogeneous information, which either include numbers of clicks on a web campaign, but also movies or business newspapers, traditional relational databases are pivotal in order to categorize information.

### HADOOP, THE TECHNOLOGICAL ANSWER FOR LARGE DATA VOLUMES

The technological response is born largely of Web giants like Google. While the search engine was still in its infancy and still faced competition from Altavista, Yahoo, Lycos, or Hotbot, the Mountain View company has developed a series of technologies to store, process and index close to five billion web pages. In 2001, it developed MapReduce, Google Big Table (compressed DBMS) and Google File System (Distributed File System), the three cornerstones of its algorithmic system to display search results. Doug Cutting, the developer of the free Lucene search engine in Java and distributed by the Apache Foundation, focuses on these projects and created the first prototype of Hadoop. Developed in Java, it is an open source framework designed to handle massive volumes of data, in the order of several petabytes. In the manner of the Google projects, it relies on a distributed file management to



## GOALS OF BIG DATA PROJECTS

Industry sectors	OBSTACLES			
	Operational analysis	Operational treatments	Social branding / perception analysis	Relational and comportemental analysis
Industry	58,2 %	21,8 %	9,1%	10,9 %
Public Sector	51,1%	12,5 %	21,6 %	14,8 %
DevOps	50,7 %	20,0 %	12,0 %	17,3 %
Finance	47,8 %	16,3 %	21,7 %	14,2 %
Health	47,3 %	21,8 %	7,3 %	23,6 %
Sales	47,1%	25,5 %	15,7 %	11,7 %
Construction	32,4 %	35,3 %	13,2 %	19,1 %

quickly process a permanent flow of information. Not to lose the battle of the Web, Yahoo! is very interested in such a solution and became the main technical and financial contributor, hiring Doug Cutting and turning its own search engine on this technological brick. The principle of operation of Hadoop is relatively simple. It revolves around the concept of «grid computing», distributing the execution of a treatment on several clusters of servers. In the manner of Google File System, it introduces its own file system, HDFS (Hadoop Distributed File System), which distributes data storage in the form of «blocks» on different nodes, while replicating in order to preserve non-altered copies. Distribution and management calculations are performed through MapReduce. As its name suggests, this technology combines two functions: «Map» that breaks an application into smaller subsets that lead to as many parts of the final result, and «Reduce» which consolidates the final result from the subsets obtained. Parallel processing helps saving considerable time, whereas traditional databases often use single batches. Through its modular architecture, Hadoop has four essential features to Big Data. First, it solves the problem of the cost of storage. To store more information, simply add additional nodes (as virtual machines, for example) and not to renew the company's storage arrays - a very

expensive proposition and difficult to anticipate. Hadoop is also scalable, that is to say it is easier to spread the processing solution based on the increase in its activity and scalability. Furthermore, by its distributed file system, Hadoop allows bulk storage of heterogeneous data. They do not have to be structured, unlike traditional relational databases, and there is no need to predict their use. Hadoop also ensures a higher safety, by its system redundancy and data replication, and high performance through parallel processing on a cluster nodes.

### THE BOOM OF NOSQL SOLUTIONS

As Google BigTable, Hadoop embeds a distributed database management system, HBase, which serves also as a base for Facebook since 2010. It is part of the «NoSQL» movement (for «Not Only SQL»), a category of DBMS that differ from traditional relational databases in the fact that the basic logic is no longer the table and the matrix representation of information, but the concept of binary document, and that their interrogation do s 'necessarily carries with SQL. Developed since 2007 by 10gen, MongoDB is one of the most famous DBMS and adheres to the same principle as HBase. It takes its name from the English «humongous», which means «huge» and he is able to be distributed over any number of nodes,



## THE IN-MEMORY TECHNOLOGY FOR REAL-TIME STATISTICS

The analytical research within Big Data projects need to be fast, continuous, and allow multiple instances. As we have seen through the exploration of different underlying technologies, NoSQL DBMS and distributed file systems like the Hadoop a framework exhibit such characteristics in part, by cutting information into subsets and treating them in parallel: one speaks of «cubes» of data, which

correspond to compilations of information necessary for the performance of a specific query. To further speed up treatments, in-memory technology store the «cubes» in memory, thus avoiding heavy physical process of writing and reading disk. As such, SAS announced the upcoming release of its environment «In-Memory Statistics for Hadoop,» which allows multiple users to manage, explore and analyze

data simultaneously. «This approach limits the memory data mixing induced by the disk I/O, which become very strict when analyzing large volumes of data,» says Oliver Schabenberger, responsible of the R & D division of SAS. The in-memory database engine thus accelerates the treatments and may be a real difference between Big Data project integration solutions.

which are added or removed at will. Objects are stored in BSON (binary JSON) without pre-determined pattern: it is thus possible to add new keys anytime, without reconfiguring the base. Specifically, the data correspond to «documents» stored in «collections»: the latter are similar to tables and relational databases and documents to different records. Within the same collection, the documents do not necessarily have to obey the same structure or present the same fields. As with the classic JSON, the documents consist of a series of key / value pairs and one can query them with well proven techniques like JavaScript and associative arrays. Here is an example of a typical MongoDB and NoSQL DBMS collection:

```
{
  «_id»: ObjectId(«2fa8c5db87c9»),
  «Name»:»Johnson»,
  «Firstname»:»Laura»,
  «Purchase»:»Blue Dress»
},
{
  «_id»:»ObjectId(«2fa8c6dv87c8»),
  «Name»:»Williams»,
  «Firstname»:»David»,
  «Adress»:
```

```
{
  «Street»:»12, Park Avenue»,
  «City»:»New York»,
  «Zip Code»:»12345»
}
}
```

Here, the key (field names) and values (that systematically follow the colon) are not preserved from one document to another. It is even possible to nest within a key document as the «Address» field in the above example. By querying the DBMS with a classical web language like JavaScript, you can enjoy a greater flexibility of treatment. Furthermore, the skills for this type of language are widespread and it becomes possible to «talk to» the data without a heavy expert recruitment process, still too rare on the market. The current offer of OVH and Amazon for storage and processing Big Data revolves around the couple also Hadoop / MongoDB. These technologies are a great ground for start-ups and start-ups, which offer complete and integrated distributions. Among the major players on the international level are Hortonworks, a subsidiary of Yahoo! receiving the support of Microsoft which integrates directly with Windows Server and Windows Azure.



## SMART CITY, THE CONNECTED CITY



**Rio de Janeiro, with its six million inhabitants, is about to be visited by tourists from around the world during the 2016 Olympic Games to improve its public services, particularly around the security of persons and problem prevention. The city has turned to IBM InfoSphere the solution to consolidate data collected by video surveillance cameras and**

**all the sensors of the city. In a single-center operations, officials access real-time video feeds from the cameras but also very diverse nature of information like tweets from visitors to visualize all the hot spots on a map. Predictive analytics algorithms anticipate problem areas; Weather Stations sensors, for example, complement statistical data to prevent**

**landslides and automatically send alerts to the population when necessary. The average response time is thus improved by 30% and public transport are more effective, avoiding bottlenecks and congestion. The concept of «Smart City», dear to IBM, should popularize in other international megacities.**

Cloudera has received \$ 740 million in funding from Intel and MapR, which is based on a native Unix file system instead of HDFS and reintroduced SQL-like queries on Hadoop data. The Hadoop framework has entered the galaxy of Apache projects and is distributed according to the principle of free software, at <http://hadoop.apache.org>. The main social networks like Twitter, Facebook and LinkedIn, but also Web giants like Amazon and PayPal, are based on this framework.

### **HOWEVER, HADOOP IS NOT ENOUGH**

As with Hadoop and its multiple components, Big Data is not confined to a single technology or technique. This is a trend that drastically transforms all companies and their relationship to information. It is therefore not confined solely to the IT department who must establish the technical infrastructure; there is a transverse upheaval in society, which must be carefully prepared by involving all services and staff, both in the way they consider the in-

## ZERO-DEFECT CARS



**By choosing a predictive analytics solution from IBM (SPSS), the**

**BMW automaker seeks to reduce and if possible, nullify the recalls of its vehicles. Its whole chain of production and after-sales service has been consolidated in a colossal project of Big Data. This huge pool of data includes both returns, service development and manufacturer tests, that helps dealers and repair services. More than 15,000 defects subtypes**

**have been identified, which are associated with driving styles or types of roads and paths. With predictive analytics, BMW is able to incorporate these results in real time into the design and production chain in order to detect anomalies and correct defects before the marketing of its vehicles.**

formation on the types of queries that can be run in order «to talk» to the raw data. In this sense, Big Data has nothing miraculous and therefore does not automatically arise indicators to assist in decision making. We must be able to question the ever increasing volume of data in the light of its own intuitions and the objectives we seek to achieve. Two types of attitudes now dominate the Big Data: operational research data, to grasp the immediate sense and in real time, and analytical research, where we retrospectively view the information as a whole through much more complex queries. Both trends complement each other and are to some extent contradictory: operational systems, such as NoSQL DBMS are capable of conducting concurrent requests and strive to reduce the response time for very specific searches, while systems Analytical facing very large volumes of data, uses treatments likely to take longer. The nature of information queried by companies varies widely between structured forms or not. Therefore, the tendency is to combine technologies and tools to query in parallel a wide variety of content. According to a recent study by IDC, nearly 32% of companies have already realized a Hadoop deployment and 31% of them intend to do so in the next twelve months. But Ken Rudin, the executive head of Facebook, recently said that «for companies seeking to exploit

large amounts of data, Hadoop is not enough». According to the same IDC study, nearly 36% of companies have deployed Hadoop and NoSQL DBMS to complement another type of databases, particularly MPP DBMS (Massively Parallel Processing, conform to SQL) as Vertica HP or Greenplum. To be most effective, the data are well correlated with traditional structured sets and unstructured information from new sources. The offer in the integration and deployment is now very broad and all the traditional players such as HP, IBM, Microsoft and Oracle offer solutions which fall around specific distributions of Hadoop to separate DBMS and proprietary data visualization tools. HP, for example, offers its own Vertica DBMS that integrates bi-directionally with all major distributions of Hadoop including Cloudera, Hortonworks and MapR. Microsoft highlights the HDInsight Services Hortonworks distribution, by integrating it with Windows Server and Windows Azure. The Redmond giant develops in parallel its Microsoft SQL Server DBMS, already widely used in traditional Business Intelligence. Oracle has developed an in-house solution with a Cloudera Hadoop distribution around NoSQL DBMS. Overlays SQL, but also alternative GlusterFS file systems such as Red Hat or Global File System (GPFS) from IBM, are sometimes preferred in HDFS Hadoop for deploying hybrid solutions.



## BIG DATA SERVICE LOGISTICS



**Big Data has a very significant impact in the field of retail and supply chain. The Dutch**

**decoration giant Hema has notably set up a platform revolving around Hadoop to thin the processing chain, logistics, and its supply chain. From any store located anywhere in Europe, customer purchase logs are sent back to the company and predictive analytics avoid out of stock situations by automatically renewing orders. Detailed statistics on cross-selling identify the best pairings of products that impact**

**directly sales. Some outlets are also equipped with interactive terminals that supplement these data and participate in better informing vendors about customer expectations and the information to issue their priority. Hema managed to streamline its customers data, the shopping experience and maximize its return on investment.**

### CHOOSING A BIG DATA SOLUTION

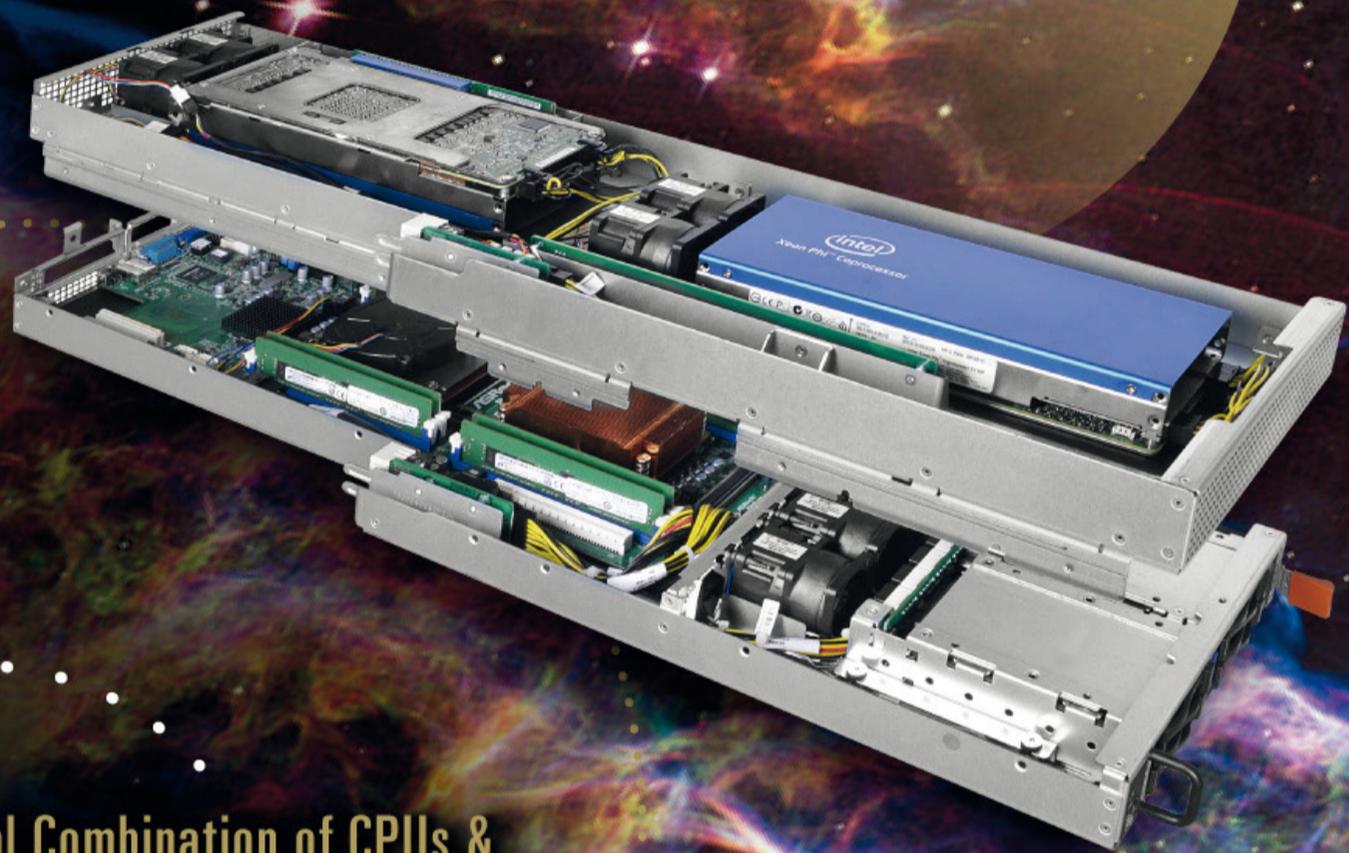
Although still young, Big Data technologies are evolving at high speed and many actors are confronted on the market. It must be said that it is particularly bright future! According to a study by Transparency Market Research firm, its turnover is expected to reach \$ 8.9 billion worldwide by end 2015 and known annual growth of the order of 40% in the next two years, reaching \$24,6 billion by 2016. IDC completes this view, stating that the area of Big Data services alone should grow by 21.1% per year. The most optimistic studies, particularly those conducted by the firm ABI Research, provide a turnover of around 114 billion dollars by the 2018-2020 period. Therefore, companies can turn to a huge variety of integrated solutions to deploy their Big Data project. Above a Hadoop distribution, they usually include a series of «packages» designed to automate and accelerate the treatments, and propose a set of APIs to develop internal applications in a familiar development environment to plan the execution of queries and also to ensure better data visualization. To choose wisely, companies must think about the nature of the data they already collect or that they intend to store, but above all adjust their expectations regarding their interpretation. Data visualization tools, for example, often differ from

one provider to another and can thus assist a public of non-statisticians. As such, market participants agree to add three new «V» to the rule already established: visibility, accuracy and value. The first component relates specifically to the data visualization tools; address high speed of colossal volumes of information is not enough, it is necessary that the decision unit is able to interpret them just as quickly. Truthfulness is a new trend which introduces verification algorithms of relevance and quality of information. On Twitter in particular, it has become essential to separate the useful data from the noise and so distinguish the best reviews of users' messages from robots. Finally, the value relates to the ultimate fulfillment of Big Data: being able to generate a really interesting return on investment and not be confined to mere technical performance. It is precisely at this stage that future Big Data experts will act. Data scientists, design engineers, statisticians, NoSQL and Hadoop development experts ... The market is growing fast and Gartner expects the Big Data should create 4.4 million jobs worldwide by the end of 2015. Young technologies, sustainable roots in academic training and real support government efforts to develop in France: more than a buzzword, Big Data could eventually become one the main business growth vectors.

**ASRock**  
— Rack —



# 2U2N-F 4GPU Server



Proportional Combination of CPUs &  
GPUs for Optimized Performance  
Virtualizing Graphic Workstations and  
Game Cloud Servers





# LAB REVIEW

## HOW WE TEST

### HPC LABS

HPC Labs is the technical unit of the HPC Media group and totally independent of the manufacturers. HPC Labs' mission is to develop methodologies and materials testing and software metrics in the high performance IT world. Capitalizing on best practices in the field, these tools are based on several decades of joint experience of the laboratorys' management.

### HPCBENCH SOLUTIONS

Specifically designed for HPC Review, the HPCBench Solutions assess not only performance but also other equally important aspects in use, such as energy efficiency, sound volume, etc. To differentiate synthetic protocols like Linpack, these protocols

HPC BENCH  
GLOBAL INDEX



9 108

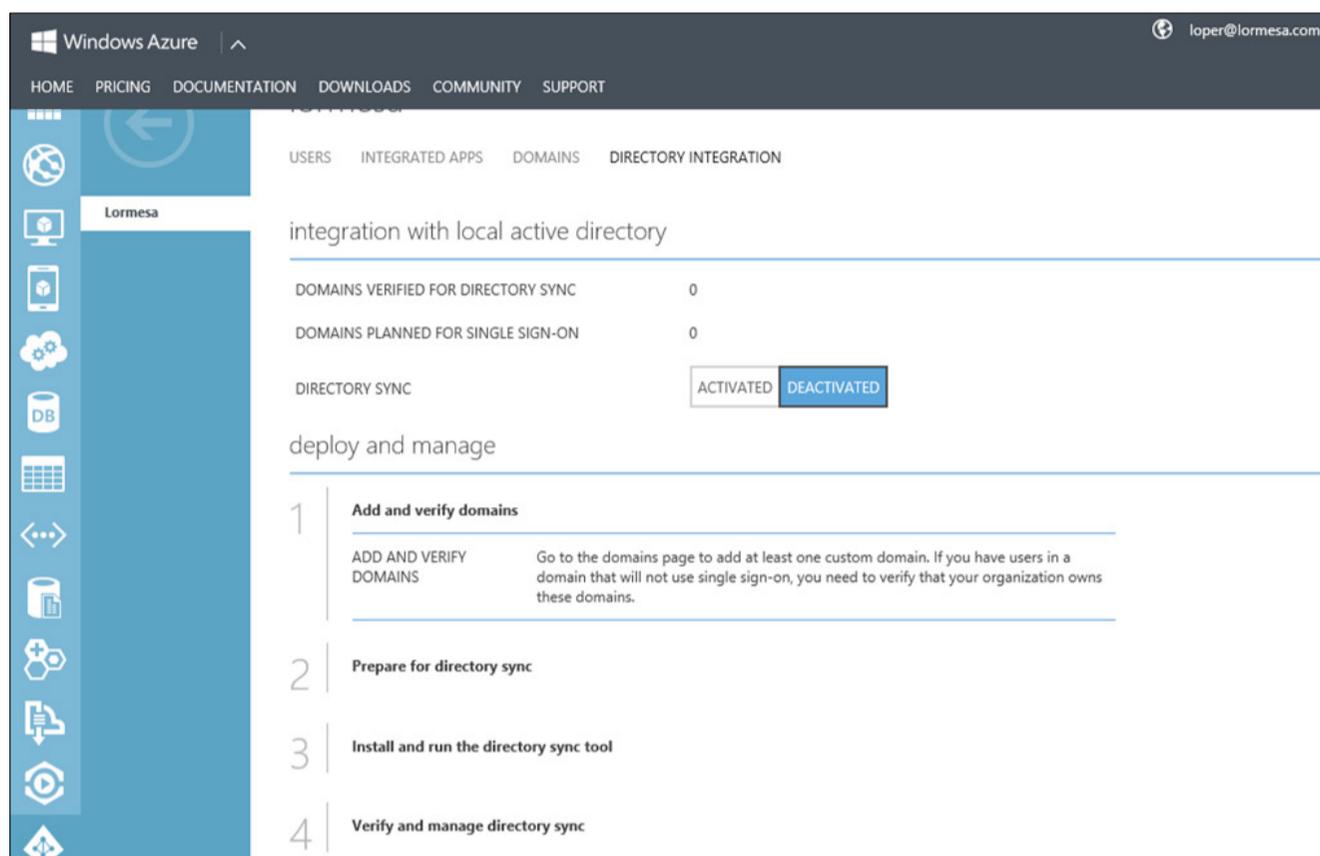
**A SINGLE SYNTHETIC INDEX TO HELP YOU  
COMPARE OUR TEST RESULTS**

allow direct comparison of solutions pertaining to the same segment, resulting in a single index taking into account the specific hardware or software tested. For example, an SSD will be tested with the HPCBench Solutions> Storage, while a GPU accelerator will be tested with the HPCBench Solutions> accels. Rigorous and exhaustive, these protocols allow you to choose what will be for you, objectively, the best solution.



**A TECHNICAL  
RECOGNITION  
AWARD**

# MICROSOFT AZURE ACTIVE DIRECTORY



**A**

zure Active Directory (AD) is tailored for companies needing to integrate an existing AD domain with cloud apps. Microsoft's IDaaS solution integrates tightly with AD. Attribute synchronization can be configured with Azure AD Connect and can later be mapped within individual SaaS

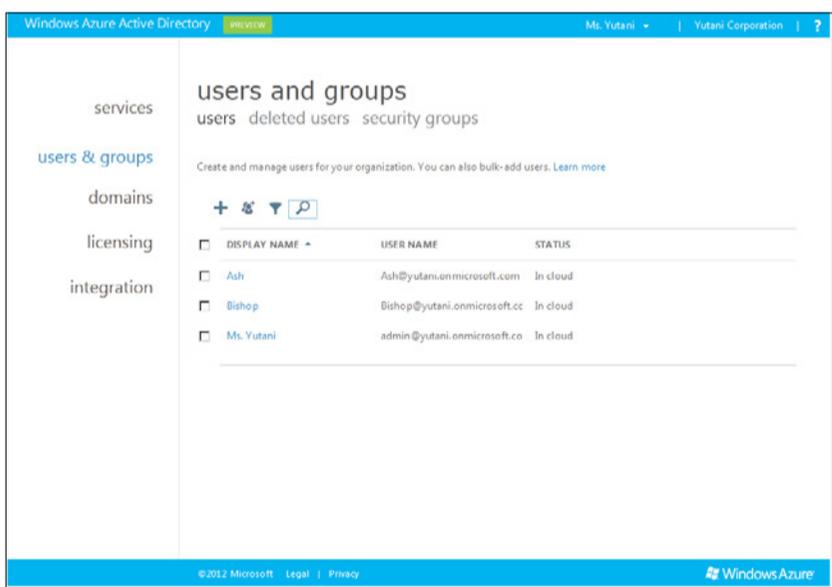
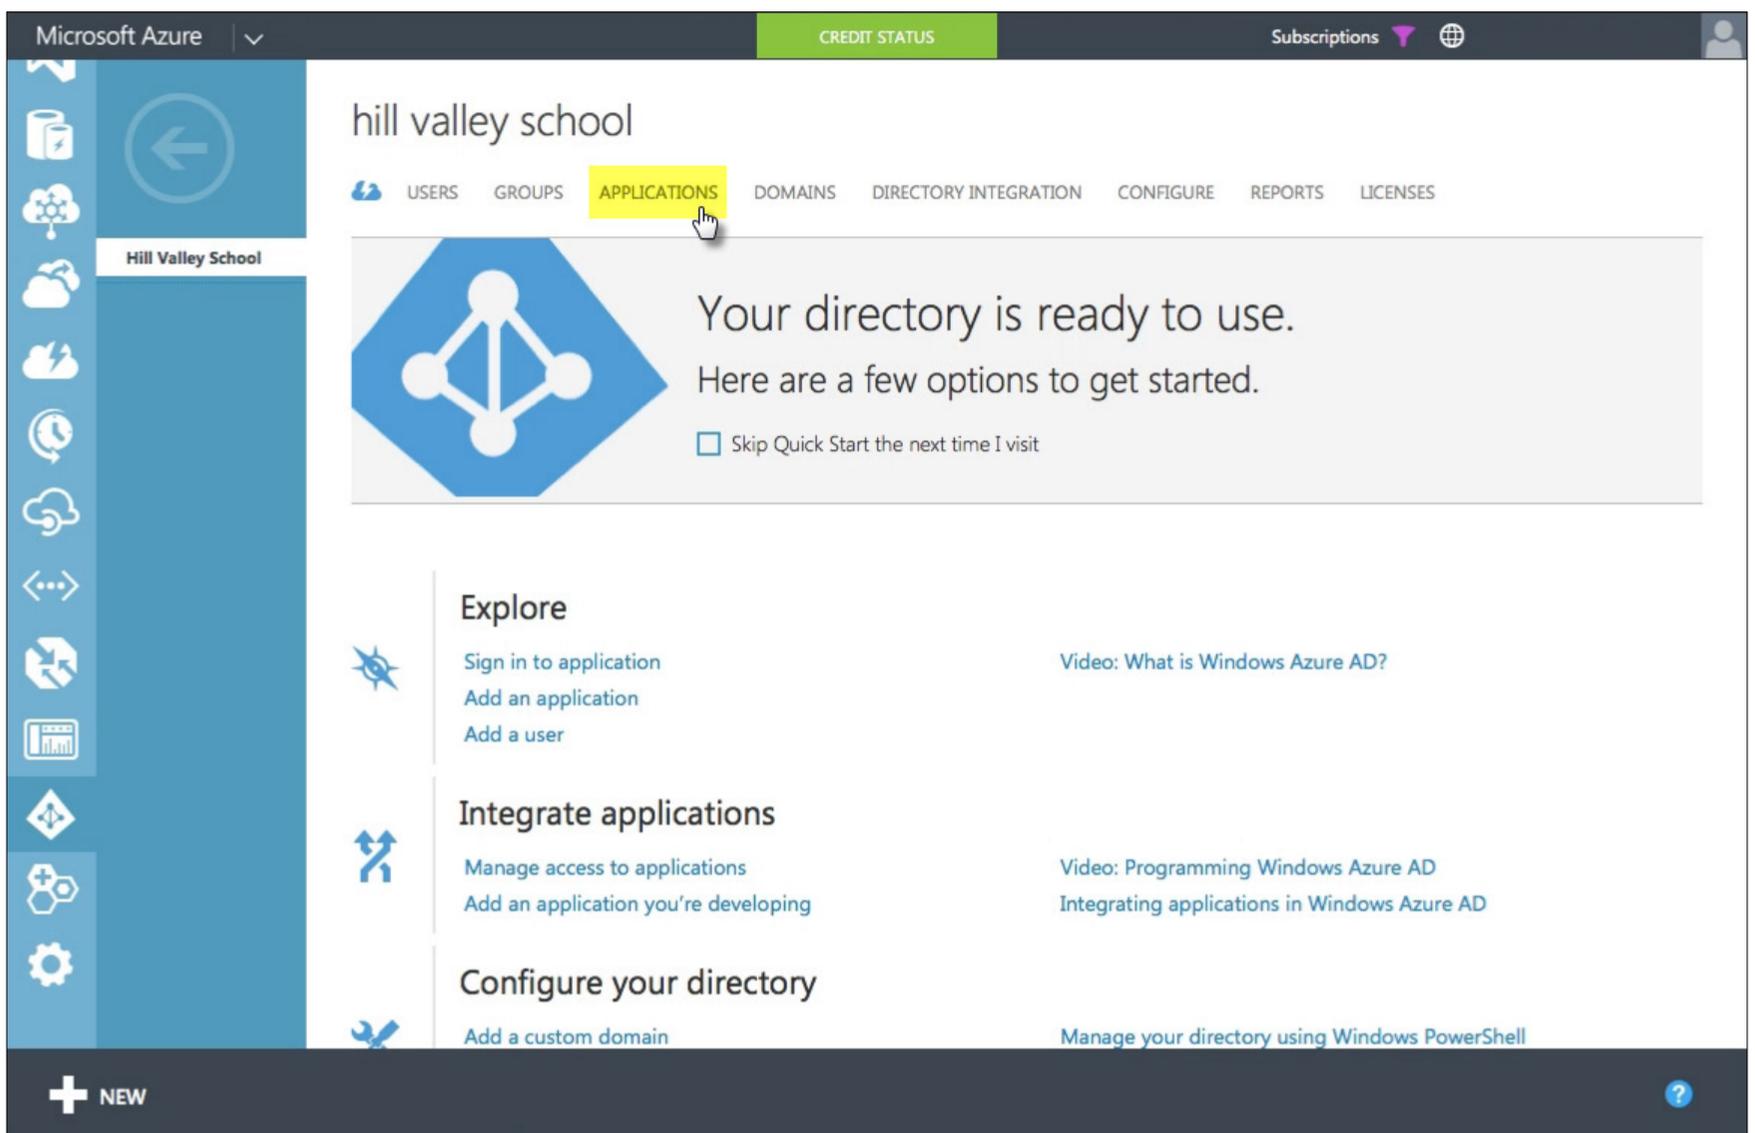
application configurations. Azure AD also integrates users of Exchange 2013 for their mail services (including Exchange Online) in conjunction with Office 365. On installation, Azure AD Connect will recognize an Exchange installation and will automatically synchronize the according attributes.

Windows 10 also brings new capabilities to integrate with Azure AD. Windows 10 supports joining devices to Azure AD as an alternative to the corporate AD. The big benefit for Azure AD users is that authentication to the user portal is seamless as the user is already authenticated to the device, and Windows 10 apps such as Mail and Calendar will recognize if an Office 365 account is available and be automatically configured. The log-in process is very similar to the default log-in style in Windows 8 where it asks for your Microsoft account details.

## USER MANAGEMENT

Both users and security groups can be synchronized using Azure AD Connect, or users and groups can be added manually within Azure AD. However, customers in large enterprises will need to frequently use the search





**PRICING**

Azure AD’s pricing begins with a free tier that supports up to 500,000 directory objects (in this case, that means users and groups) and up to 10 single sign-on (SSO) apps per user. The Free version of Azure AD is automatically included with Office 365 subscriptions, in which situation the object limit does not apply. With a retail price of \$0.50 per user per month, the Basic tier of Azure AD is extremely competi-

tive. The Basic service adds capabilities such as branding for the user portal and group-based SSO access and provisioning so, in order to automatically create user accounts in SaaS applications, you’ll need the Basic tier but it retains the 10 app per user limit. The Premium tier in Azure AD removes the limits from the amount of SSO apps users can have and adds self-service and multifactor capabilities for \$6 per user per month. Azure AD Premium also includes user Client Access Licenses for Microsoft Identity Manager which can be used to synchronize and manage identities in databases, applications, other directories, and more.

**CONCLUSION**

Azure AD covers the majority of the core features you should be looking for in an IDaaS provider, and also features some enterprise-level tools. The pricing is very competitive, and integration with Office 365 and other Microsoft products and services are solid.



# HPCToday

The global media in high-performance IT

## Subscribe now > for FREE

Subscribe now and receive, every month, an expert, actionable coverage of HPC and Big Data news, technologies, uses and research...



+ get yourself instant access to exclusive contents and services

# www.hpctoday.com



# WD ARKEIA RA4300

**A**

imed at small and medium-sized enterprises, the WD Arkeia RA4300 backup appliance is an integrated package including all hardware, software and storage. Delivered ready for use, it is compatible with more than 200 real and virtual platforms.

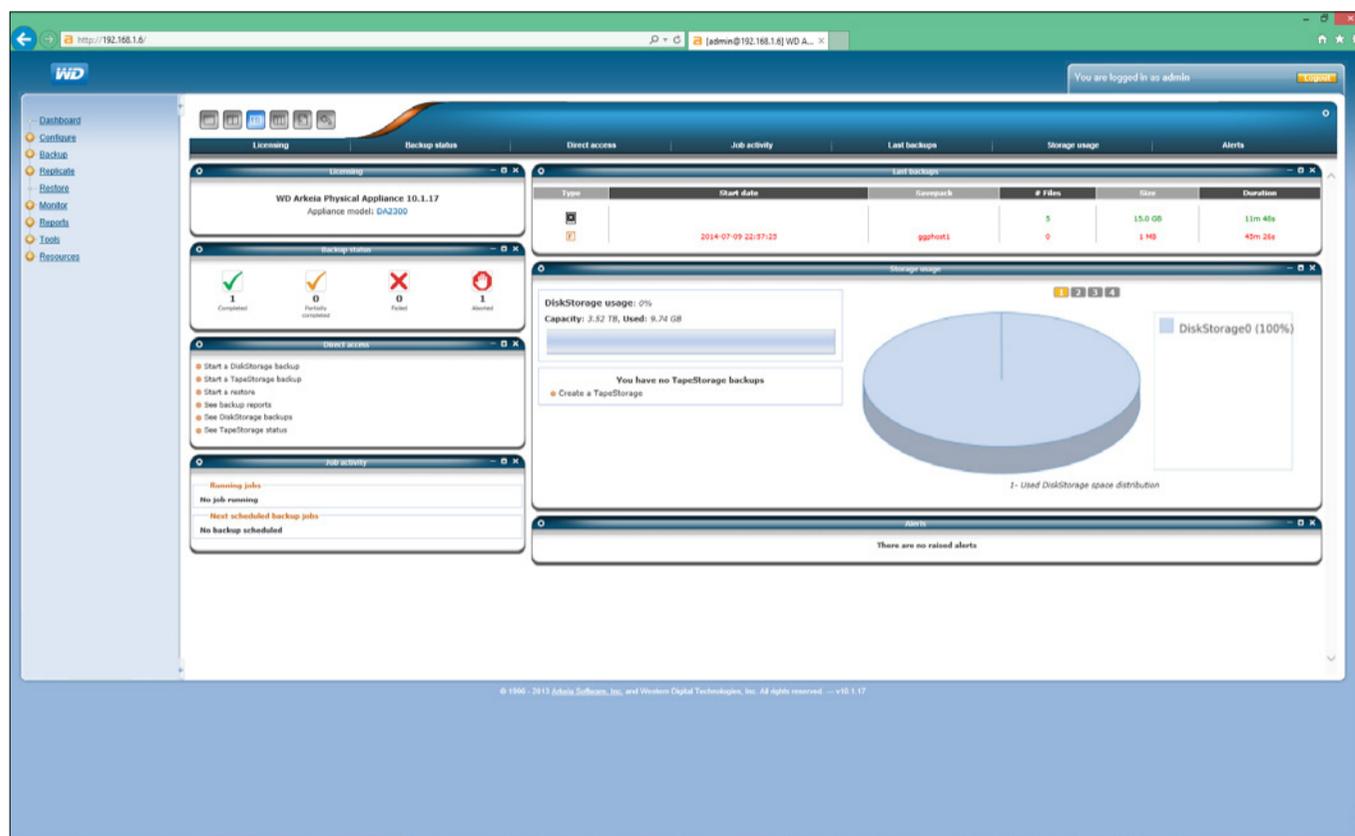
Arkeia, acquired by Western Digital in 2012, is a recognized specialist in corporate backup. The RA4300 appliance is a mid-range model «all in one» unit with integrated hardware, software and storage that can operate in standalone mode once configured. In the evaluated model, this appliance has 6 slots for disks,

all with 4 TB drives for a total storage of 24 TB. This space can be configured in RAID 6 to favor security. In this case, up to two drives may be inoperative without risk for the data.

## **FULL HYPERVISOR COMPATIBILITY**

The appliance is compatible with the four major virtualization platforms: VMware vSphere, Microsoft Hyper-V, Citrix Xen Server and Red Hat RHEV. On the client side, more than 200 environments are supported. This compatibility does not extend so far to all functions in a disaster recovery framework, since Windows 2000 platforms are no longer covered. Like many of his peers, this appliance contains a sophisticated deduplication mechanism to reduce the volume of data in transit and stored. This progressive deduplication acts on the source side and has a granularity down to the

# THE WD BACKUP ARKEIA BACKUP ENVIRONMENT IS SIMPLE TO USE. ALL OF THE ACTIONS OUTLINED IN THE TREE AND SETTING TAKES PLACE WITHOUT MAJOR DIFFICULTIES.



data block. With the benefit of saving time and reduced disk footprint. This helps to maximize the storage space that is, once the data is de-duplicated and compressed, to the storage space multiplied by a factor 2 to 5.

## A SCALABLE APPLIANCE

The WD Backup Arkeia backup environment is simple to use. All of the actions outlined in the tree (configure, back up, replicate, restore etc). and setting takes place without major difficulties. As with any use of this type, the administrator in charge of deployment needs to know on the fingertips of his company's infrastructure. The dashboard displays on a single screen disk space and its distribution, «health status» backup (successful or had problems) and the table of current tasks. Note that the backup application, WD Backup Server is also

available off appliance. RA4300 The appliance can be configured as a bare metal disaster recovery unit. The DR Arkeia software agent that relies on snapshots VSS (Volume Shadow Copy Service ) containing all the information stored on volumes (data, metadata, ACLs) and takes into account the partitions, volumes and native filesystems. An ISO image based on Windows PE 3.1 can be generated. We welcome the effort to maintain a WD Arkeia documentation space and a very complete and updated wiki, including recent patches against the OpenSSL heartbleed bug. The one-year warranty includes maintenance and covers hardware, software updates and support. Configurable and expandable at will, the WD Arkeia RA4300 provides a solid and well-integrated offering that provides a level of benefits without hidden costs nor surprises. **JOSCELYN FLORES**



# HPC ACADEMY

Developing High-Performance Computing Expertise in Europe



**NONPROFIT ORGANIZATION**

## Our missions:

- To promote HPC (High Performance Technologies) across Europe
- To develop HPC training courses for industrial and scientific audiences.
- To develop collaborative projects and codes with committed partners
- To initiate the academic and student communities to parallel programming in order to favor career evolutions
- To promote professional insertion among our members
- To encourage those who dare create and innovate.

## Our first partner members:



**Become a member**  
[www.hpc.academy](http://www.hpc.academy)

# OCZ'S NEW NVME-COMPLIANT Z-DRIVE 6000 SSD SERIES

INCLUDES NVME OVERVIEW, STRATEGIC POSITIONING, PERFORMANCE METRICS AND ENDURANCE

# A

According to IDC and its 'Worldwide Big Data Technology and Services' research, digital data will reach eight zettabytes in 2015 and forty zettabytes by 2020. With this amount of data generated and accessed, the performance requirement from storage devices has significantly grown. Archaic storage devices, such as

spinning disks or hybrid storage, suffer through severe I/O bottlenecks that greatly hamper applications and services. Hence, flash based solid state storage adoption in datacenters and cloud services is on the fast track.

With an ever-growing gap between application performance requirements and hard disk drive (HDD) capabilities, PCIe-based SSDs are becoming highly desirable and represent a fast-growing multi-billion dollar global opportunity over the next five years. At present, PCIe SSD deployments are at an adoption stage (see Figure 1), but forecast is to achieve an esti-



mated 2.75x growth in units shipped over the next 5 years according to TrendFocus market research. To achieve these forecasted deployment levels, new technologies must be developed that standardize and unify the storage stack, eliminate the existing proprietary driver model, and create a serviceable package.

With this opportunity, OCZ Storage Solutions introduces the Z-Drive 6000 SSD Series

# THE Z-DRIVE 6000 SERIES IS BUILT WITH PERFORMANCE, ROBUSTNESS AND LONGEVITY IN MIND AND SUPPORTS PCIe 3.0 X4 CONNECTIVITY WITH FULL NVME 1.1B COMPLIANCE AND IN-BOX OPERATING SYSTEM DRIVERS.

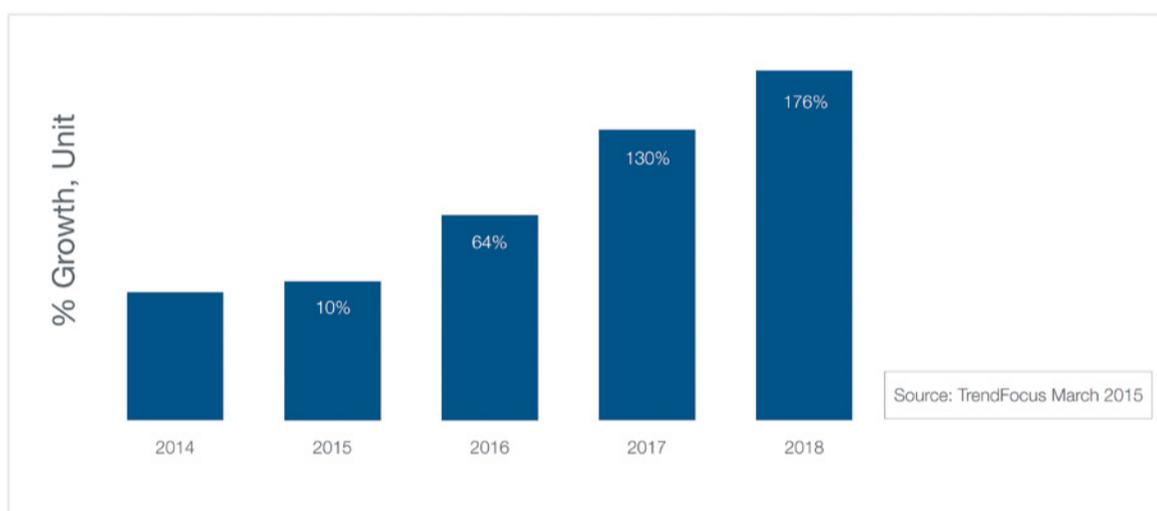


Figure 1: Percentage of PCIe SSD Volume Growth

## Percentage of PCIe SSD volume growth

that builds on its proven Z-Drive technology, pairing next generation PCIe and NVMe (Non-Volatile Memory express) technologies with robust reliability and an extensive feature set. The portfolio is available in multiple configurations - offering different form factors, capacity points, and endurance ratings to suit a wide range of applications and ecosystems as presented over the next few pages.

## NVME OVERVIEW

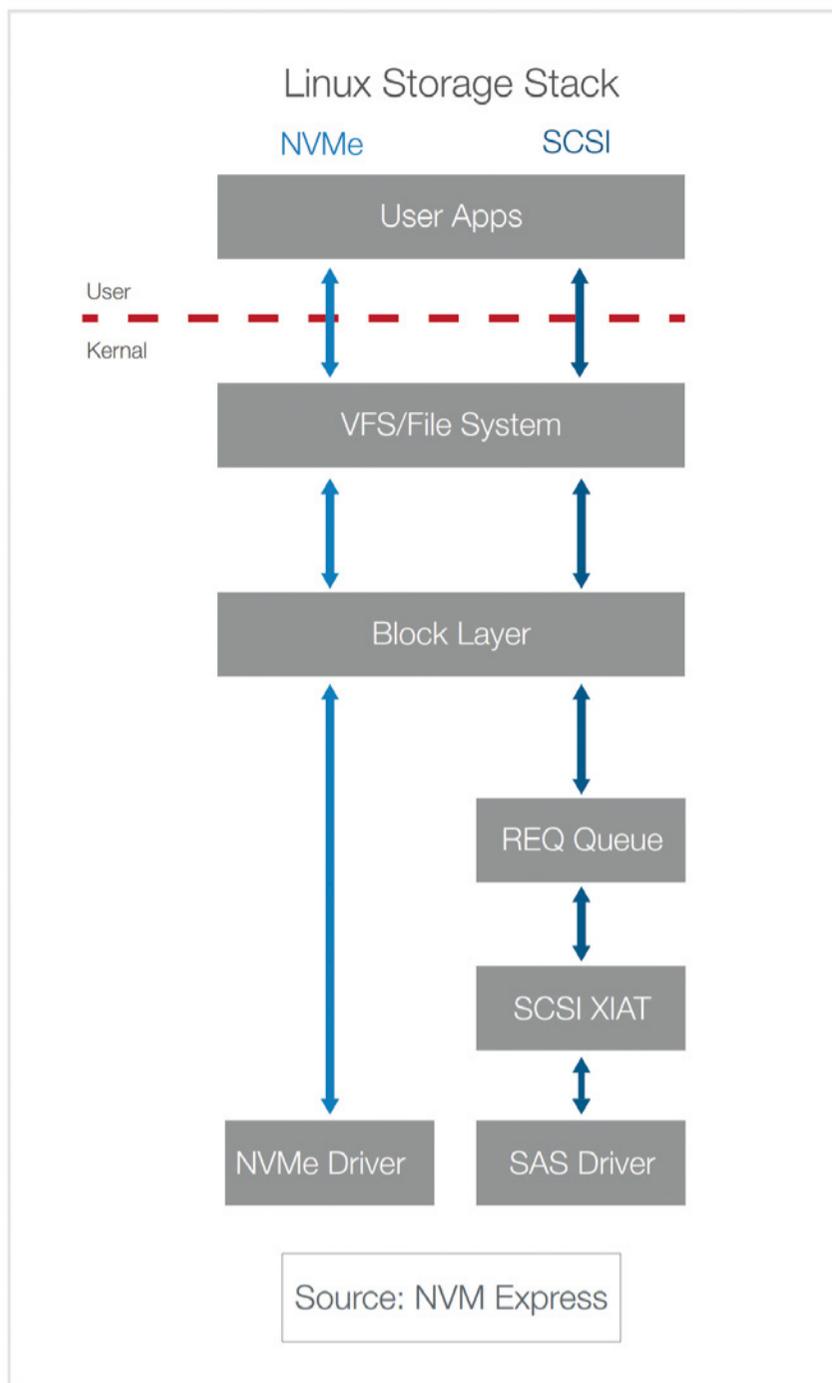
Today's enterprise compute platforms support up to forty lanes of the Generation 3.0 PCIe interface, where each lane has a theoretical bandwidth under 1GB/s. The reduced interface overhead and proximity of PCIe endpoints to the CPU allows PCIe-based SSDs to take advantage of extremely low latencies and high bandwidth not possible with SAS or SATA interfaces. Traditional PCIe flash storage devices offer high performance and high capa-

city but require proprietary software support. The NVMe specification extends traditional PCIe flash storage to new levels. Architected and optimized from the ground up specifically for non-volatile memory based solid-state storage, NVMe features a streamlined memory interface, command set and queue design that is well suited for

today's virtual operating systems. These features result in faster access to critical data and highly resilient storage capabilities. The NVMe host control interface allows for both system builders and storage vendors alike to develop the different parts of a storage ecosystem to the same specification enabling broad interoperability between storage devices, host platforms and supporting software. NVMe is an ongoing development effort coordinated through an open industry consortium of over 90 members under the direction of a 13-company promoter group.

## UNDERSTANDING NVME BENEFITS

I/O requests in an enterprise ecosystem spend much of their time in the hardware infrastructure (that includes NAND flash media, the flash controller, the host bus adapter, related hardware, etc.), and only a small portion of time working through the software I/O stack. The NVMe standard streamlines the software I/O stack by reducing unnecessary legacy



**This figure depicts the streamlined NVMe storage stack**

overhead and supporting multiple queues, and many more commands per queue than any other commonly used storage protocol as outlined in Figure 2. NVMe supports 64,000 commands per queue (as well as 64,000 queues) to enable extremely fast hardware responses especially when compared to the SAS protocol that can only support 254 commands per queue or the SATA protocol that supports 31 commands per queue.

Published tests results indicate that the Linux NVMe software stack reduced software overhead by more than 50% (from 6µs to 2.8µs) when compared to a Linux SCSI/SAS software stack.

	OCZ Z-Drive 6000	OCZ Z-Drive 6300		
Intended Application	Read Intensive	Mixed Use		
Usable Capacity (TB)	0.8, 1.6, 3.2	0.8, 1.6, 3.2, 6.4*		
NAND	Toshiba A19 MLC	Toshiba A19 eMLC		
Interface	PCIe 3.0 x 4 NVMe 1.1b	PCIe 3.0 x 4 NVMe 1.1b		
Form Factor	2.5" x 15mm	2.5" x 15mm HHHL		
4K Random Read (KIOPS)	600 / 700 / 700	600 / 700 / 700		
4K Random Writes (KIOPS)	115 / 160 / 160	75 / 120 / 120		
128K Sequential Read (MB/s)	2200 / 2900 / 2900	2200 / 2900 / 2900		
128K Sequential Write (MB/s)	1300 / 1900 / 1900	1000 / 1400 / 1400		
70%/30% Read/Write (KIOPS)	200 / 330 / 330	230 / 280 / 280		
Endurance (Drive Writes Per Day)	1	3		
4K Read/Write Latency (µs)	80	25	80	30
Warranty (Years)	5	5		

\*6.4TB to be released later in 2015

The number of instruction cycles was also reduced from 19,500 to 9,100. From an SSD vendor perspective, the NVMe interface has a number of benefits. First, standardization makes the development effort easier with a common protocol between client and enterprise, as well as a plethora of optional features that open many new opportunities for firmware differentiation. Second, protocol efficiency streamlines the I/O stack to reduce stack latency as extensive queue depth and queue quantity support the inherent parallelism of flash based storage. Third, the implementation of a Data Integrity Field (DIF) delivers redundancy checks for transmitted read/write errors initiated by the host, which in turn, reduces system downtime as well as total operating cost (TOC).

### Z-DRIVE 6000 SSD SERIES OVERVIEW

The Z-Drive 6000 Series is OCZ's first NVMe-based SSD portfolio that addresses those enterprise applications requiring high-performance and low latency I/O responses coupled with the data protection, endurance and reliability that IT managers expect in enterprise environments. The series / models that comprise the Z-Drive 6000 SSD portfolio include the following solutions: The Z-Drive 6000 Series is built with performance, robustness and longevity in mind and supports PCIe 3.0 x4 connectivity with full NVMe 1.1b compliance and in-box operating system (OS) drivers to ensure broad interoperability and performance.

## ENTERPRISE / DATACENTER OPTIMIZED DRIVE

The performance delivered by the Z-Drive 6000 SSD family is designed to fulfill the stringent requirements of OEM customers for enterprise-class centralized storage systems, server caching applications, VMware vSphere virtualized applications, and other applications that require high I/O performance and low latencies. The portfolio offers many benefits for enterprise-class deployments:

Z-Drive 6000 Series Features	Benefits
NVMe 1.1b Support	A streamlined I/O stack that reduces latency and supports massive parallelism, enabling users to meet current and future storage needs
3rd Party Driver Support	Full compliance with NVMe standard allows for 3rd party NVMe driver support across any platform, allowing OEMs and system integrators to only qualify a single driver of their choice
'Flow-Through' Case Design	OCZ's innovative 'Flow-Through' case design allows more airflow to critical components, keeping the device cool and reducing airflow requirements
End-to-End Data Protection	DIF* (host appended or drive appended), ECC, and CRC support delivers superior data protection and reliability
Power Loss Protection	Full power loss protection ensures that all in-flight data has reached the device at the moment of power loss and is not lost
Hot-Swap Support	Hot-Swap support enables unprecedented serviceability for PCIe storage devices by no longer requiring that a node be brought offline to service the storage device
Dual Port Support*	Dual Port support allows two data paths to the device, and up to four separate name spaces to ensure availability of critical data
User Selectable Power Modes	Variable power modes, in 15W, 20W, and 25W modes, reduce wasted power when maximum performance is not required
AIC MD2 Form Factor	HHHL AIC form factor enables OEMs and system integrators to populate sockets where an FHHL device would not fit, or where SFF-8639 bays are not available

\*Dual Port and DIF features available through future firmware update

### APPLICATIONS SUPPORTED

The Z-Drive 6000 NVMe-based PCIe SSD Series delivers extraordinary steady state performance to support applications that require large amounts of low-latency, high-bandwidth storage.

Application	Application Z-Drive 6000 SFF Read-intensive	Z-Drive 6300 SFF Mixed Workload	Z-Drive 6300 AIC Mixed Workload
Boot Operations (WS and file server)	Yes		
Decision Support System (DSS)	Yes		
Front-end Web Servers	Yes		
Media Streaming	Yes		
Read Cache and Indexing	Yes		
Video on Demand (VOD)	Yes		
Cloud Computing (SaaS, PaaS, IaaS)	Yes	Yes	Yes
Virtual Desktop Infrastructure (VDI)	Yes	Yes	Yes
Backup Operations		Yes	Yes
Big Data Processing & Analysis		Yes	Yes
MS Exchange		Yes	Yes
SQL Database OLTP		Yes	Yes
SQL Database OLAP		Yes	Yes
High Performance Computing (HPC)		Yes	Yes
Professional Media Editing / Post Production		Yes	Yes
Database & Error Logging		Yes	Yes
Enterprise Resource Planning (ERP)		Yes	Yes
Oil and Exploration		Yes	Yes
Medical Imaging		Yes	Yes
Computer-Aided Design		Yes	Yes

### SUMMARY

OCZ's premium enterprise-class Z-Drive 6000 SSD Series portfolio are cutting edge PCIe/NVMe drives with exceptional performance and lowest latency in its class that supports SFF and AIC form factors, read-intensive and mixed workload applications galore, and enterprise-class endurance and data reliability for any business looking to overhaul or simply improve their data storage systems.

Additional information on the Z-Drive 6000 SSD Series, as well as OCZ's complete line-up of enterprise SSDs and software solutions are available at [www.ocz.com/enterprise](http://www.ocz.com/enterprise).

1 - "First NVMe SSDs Announced," Dennis Martin, Storage Newsletter, September 20, 2013.

# OCZ'S PREMIUM ENTERPRISE-CLASS Z-DRIVE 6000 SSD SERIES PORTFOLIO ARE CUTTING EDGE PCIe/NVME DRIVES WITH EXCEPTIONAL PERFORMANCE AND LOWEST LATENCY IN ITS CLASS.



# WHAT IT TAKES TO REINVENT SUPERCOMPUTING OVER AND OVER AGAIN

I'm not usually a big fan of anniversaries (except my wedding day, of course), but I make an exception when it comes to IBM's collaboration with the US Government on supercomputing. Today is the 20th anniversary of the Accelerated Strategic Computing Initiative--a Department of Energy program that has safeguarded America's nuclear weapon arsenal and, at the same time, helped IBM assert ongoing leadership in this most demanding of computer domains.

With help from National Laboratories scientists, teams of IBMers have produced five generations of supercomputers--repeatedly ranking among the fastest machines in the world. The journey led us to where we are today: developing a sixth generation of computers, data-centric systems designed from the ground up for the era of big data and cognitive computing. The program was also instrumental in IBM's

rebound after the company's near-collapse in the early 1990s.

I remember the day the original ASCI contract was signed. IBM and DOE people had gathered in a conference room at the IBM headquarters north of New York City. Unexpectedly, Lou Gerstner, IBM's then-new CEO, popped in and gave off-the-cuff remarks. I remember him saying, "IBM is all about solving hard problems. This is the hardest problem there is. We're all in." I was sitting in a chair and he was standing behind me. He put his hands on my shoulders and said, "Here's the guy who will do it."

Gulp.

The task of creating computers that are capable of simulating nuclear explosions so countries don't have to test with actual bombs turned out to be difficult indeed.

The first years were the toughest. I had been with IBM for 20 years by then and had expe-

## **WE ALSO HAD TO DEVELOP A NEW PROCESS FOR DEVELOPING AND MANUFACTURING SUCH COMPLEX SYSTEMS – WITH THOUSANDS, AND, LATER, MILLIONS, OF PROCESSORS.**

rience in both hardware and software development. Most relevantly, I had been involved in an effort to transform IBM mainframes into supercomputers. That didn't pan out, but in the process we learned a lot about what it would take to build high-performance computers. We had relaunched our supercomputing effort with a new technology strategy just before we engaged with the Department of Energy. To ramp up the ASCI project development team quickly, I cherry-picked people from IBM's offices and labs all over the Hudson Valley. Some of them were green, in their 20s, but they had the nerve to rethink computing.

We made a series of radical choices. We adapted processors and systems technologies that IBM had developed for its scientific workstation business. UNIX would be the operating system. We had to invent new networking to hook all the processors together. And we were one of the first groups at IBM to use open source software. We had to move too quickly to code everything ourselves. We also had to develop a new process for developing and manufacturing such complex systems--with thousands, and, later, millions, of processors.

With each new generation, the requirements increased dramatically. The first machines produced 3 teraflops of computing performance, or 3 trillion floating point operations per second. The current generation produces 20 petaflops; 20 quadrillion operations per second. That meant we had to invent not just individual technologies but whole new approaches to computing. For instance, in the early 2000s, IBM Research and scientists at Lawrence Livermore National Laboratory teamed up to create a new supercomputing architecture which harnessed millions of simple, low-powered processors. The first systems based on this architecture, called Blue Gene/L,

were incredibly energy efficient and exceeded the performance of Japan's Earth Simulator by greater than a factor of 10, helping the US recapture leadership in supercomputing.

Today, we're developing yet another generation of supercomputers for the National Laboratories. They're based on the principle that the only way to efficiently handle today's enormous quantities of data is to rethink computing once again. We have to bring the processing to the data rather than follow the conventional approach of transmitting all of the data to central processing units.

When we first proposed this solution, we were practically laughed out of the room. But, today, data-centric computing is becoming accepted across the tech industry as the way to go forward.

Through the ASCI project, I learned lessons that I think are critical for any large-scale development project in the computer industry. First, you must assemble an integrated team of specialists in all of the hardware and software technologies. You can't negotiate to get the technologies and skills you need from a half dozen vice presidents who have their own priorities. Second, you must see the big picture. Don't think of a server computer in isolation. Plan so you can integrate servers and other components in large systems capable of taking on the most demanding computing tasks.

I guess there's one more critical lesson I learned from this tremendous experience: recruit bright and fearless people and ask them to do nearly impossible things. Chances are, they'll rise to the challenge.

*First publication* : **THINK blog**  
[www.ibm.com/blogs/think/2015/10/21/what-it-takes-to-reinvent-supercomputing-over-and-over-again/](http://www.ibm.com/blogs/think/2015/10/21/what-it-takes-to-reinvent-supercomputing-over-and-over-again/)

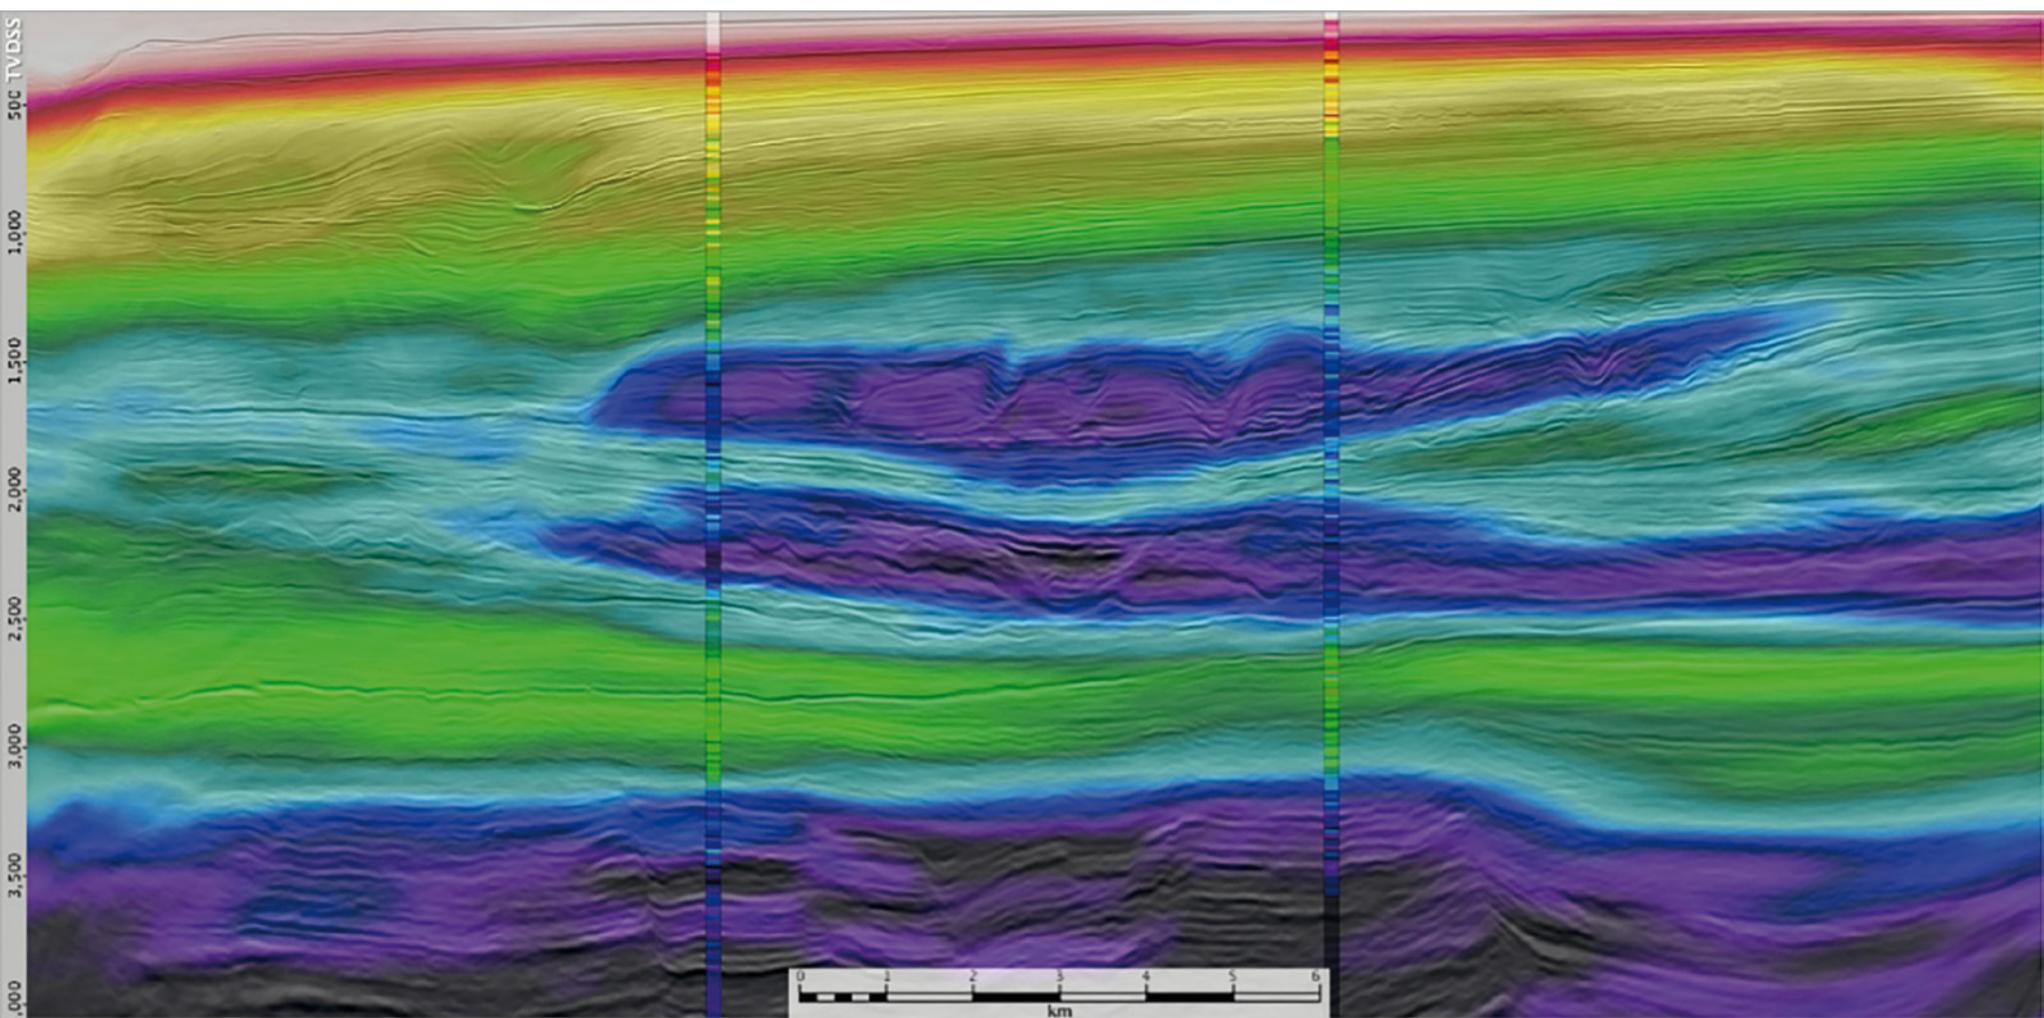


# INTEL SEES SIGNIFICANT LUSTRE MOMENTUM IN ENTERPRISE AND RESEARCH HPC

**T**he Lustre parallel file system software has always seen healthy adoption in the research labs around the world, in spite of a challenging start across multiple owners before being released to the open source community. A boost from OpenSFS, contributions from Whamcloud beginning in 2010, Intel's acquisition of Whamcloud in 2013, and introduction of commercial distributions based on the open

source version from a number of suppliers, including Intel, have helped accelerate the relevance and today are contributing to a growing adoption of Lustre in enterprise HPC.

According to Earl Joseph, a top HPC industry analyst with IDC, "Along with IBM's General Parallel File System (GPFS), Lustre is the most widely used file system. But Lustre is experiencing healthy growth in terms of market share while GPFS remains flat. Lustre is also supported by a large number of OEMs, providing the HPC community with a strong base for growth."



## **INTEL CONTINUES TO SEE SIGNIFICANT LUSTRE ADOPTION**

“Intel has experienced considerable growth in sales, measured by support contracts for the Intel editions of Lustre, year-over-year from 2013 to 2014,” said Bret Costelow, Director of Global Sales for Lustre Solutions at Intel. “This is representative of those Lustre adopters who have moved away from unsupported, roll-your-own versions, and the competition to Intel’s editions,” said Costelow. “And, we continue to see growth momentum in 2015.”

Since the formation of Intel’s High Performance Data Division (HPDD), the company has also grown its Lustre channel—covering OEMs and system Integrators—from a handful in 2013 to more than 170 today. “This momentum is evidence of our efforts at Intel to penetrate the enterprise market with a high-performance data solution that offers HPC-class performance with the enterprise-grade reliability, stability, manageability, and high availability features that IT departments require,” stated Costelow.

**DownUnder GeoSolutions uses Lustre to help handle the extraordinarily large datasets generated by seismic processing, such as this section revealing geological structure up to 4km below the surface (colors represent the varying speed of sound).**

Photo courtesy DownUnder GeoSolutions.

According to Costelow, Intel has seen multiple design wins with major OEMs, like Bull, Cray, DataDirect Networks, Dell, EMC, Huawei, HP, Inspur, NetApp, SGI, and SuperMicro. Their efforts emphasize the growing relevance for Lustre in the marketplace against other storage solutions—both in the academic and national labs and also in enterprise HPC—and thus the momentum Lustre is gaining as a leading parallel file system for high-performance applications.

## **EXPANDING ENTERPRISE DEPLOYMENTS**

“We’re seeing Lustre adoption in more enterprise and enterprise-influencing markets, including oil and gas, financial services, and genomics,” commented Costelow. For example,



**Iowa State University researchers are using HPC systems to understand how weather patterns affect crop plantings, such as these soybeans standing in water due to heavy Midwest rains.** Photo: Photo by Palle Peterson, Iowa State University, Published in the ICM newsletter, June 2008.

TATA Consulting Services (TCS) has been working with Intel to look at the convergence of HPC and Big Data in the Financial Services Industry (FSI) with Intel® Enterprise Edition for Lustre software as the underlying file system. Instead of testing with industry standard benchmarks, which don't always reflect the real impact of system performance for a particular usage, TCS is testing with a range of real-world financial applications. "Our objective was to come up with a platform for Hadoop data analysis using an HPC cluster that would give us good performance," said Rekha Singhal, Senior Scientist with TCS. "TATA was able to achieve a 3X performance gain using Hadoop on Intel's version of Lustre compared to Hadoop on the Hadoop Distributed File System (HDFS)."

At the Bank of Italy, the IT division of the economics and statistics department is modernizing its HPC infrastructure that supports user applications and provides some database storage. In particular, the newest system currently being readied for production is based

on Intel® Xeon® processors and the Intel® Enterprise Edition for Lustre software. It will support about 700 users in the department with 100 TB of storage. The new system is being designed to enable collaboration across the organization by storing and sharing from the Linux cluster's Lustre file system all of the users' documents created with their Windows applications. Additionally, this installation is representative of the smaller file sizes Lustre is having to move with the same high performance and efficiency as it has supported in traditional HPC installations. Instead of typical massively large file sizes in the hundreds of gigabytes to terabytes range, the new system will be supporting an average 4K file size, but many, many millions of them. Thus, Intel Enterprise Edition for Lustre version 2.5, with the ability to support many more meta data servers than possible in earlier versions, is being deployed.

Oil and Gas industries run some of the largest clusters around the world, according to Brent Gorda, General Manager of Intel's HPDD. With the massive data sets in seismic research, oil and gas represents commercially the types of workloads Lustre was designed for when it was first funded under U.S. Department of Energy grants.

"At DownUnder GeoSolutions (DUG), we operate Intel Lustre filesystems totaling 9 PB across our worldwide processing centers," stated Phil Schwan, Head of Software at DUG. "And we're adding multiple PB each year." Some of their codes are extremely I/O intensive. And the demands on the cluster's files system become more extreme because they run Intel® Xeon Phi™ coprocessors in many of their compute nodes. "Keeping them all busy at 90+ percent of their peak capability requires a lot of file-system bandwidth," commented Schwan.

## **LUSTRE ACHIEVEMENTS IN THE CLOUD**

"One of the most exciting opportunities we're seeing for Lustre expansion is with Intel® Cloud Edition for Lustre software," commented Costelow. Amazon has integrated the

## **THE RESEARCHERS HOPE TO INFORM AGRICULTURAL PRACTICE AND POLICYMAKING IN THE FACE OF CHANGING WEATHER PATTERNS. THESE KINDS OF SIMULATIONS DEMAND A LEVEL OF I/O THROUGHPUT FROM DATA SETS THAT ONLY A PARALLEL FILE SYSTEM LIKE LUSTRE CAN SERVE.**

Cloud Edition into their Amazon Web Services (AWS) to offer high-performance, scalable storage using Lustre in their Elastic Compute Cloud (EC2). SAS, the Business Analytics software company, delivers clustered analytics services through the Amazon Web Services Marketplace and recommends using Lustre on AWS for their analytics software. “With Amazon and SAS, we are seeing and expect to see even more significant and exciting traction taking place in the near future with Lustre in the Cloud,” added Gorda.

### **LUSTRE IN ACADEMIA INFLUENCES COMMERCE**

Academic research can have a significant influence on commercial enterprise. The HPC systems used in this arena often use Lustre to support their compute- and I/O-intensive workloads. For example, Iowa State University (ISU) in Ames, Iowa recently installed its Condo cluster, based on Intel Xeon processors, Intel® True Scale fabric, and the Intel Enterprise Edition for Lustre software with Intel support. Condo is being used across many scientific disciplines with highly potential commercial impacts. ISU researchers use Condo in genomics studies for the purpose of improving brood stocks, and for genomic analysis of ancestral varieties of corn to gain insight into what might improve the genotype of corn for food crop production. Corn is the most studied plant around the globe.

Another area of research is simulating and analyzing climatic changes across the U. S. Midwest. Scientists are interested in weather changes relative to impacts from known issues, such as global warming, but also loo-

king at how changes in land use and farming methods over 100 years might be a factor in how the Midwest has evolved to a wetter climate over the last few decades. The researchers hope to inform agricultural practice and policymaking in the face of changing weather patterns. These kinds of simulations demand a level of I/O throughput from data sets that only a parallel file system like Lustre can serve. These kinds of proofs of concept with Big Data and HPC cluster convergence, stretching Lustre beyond its historical usages, and cloud, commercial, and academic deployments further create momentum for Lustre across the evolving marketplace.

### **INTEL TO BUILD MOMENTUM BY ENHANCING OPPORTUNITIES FOR CHANNEL PARTNERS**

To improve its reach further and accelerate Lustre adoption, Intel is developing an entirely new world-class Intel® Lustre Solutions Reseller Channel Program. According to Gorda, the goals of this program are to develop an elite worldwide network of resellers with the best trained and technically competent staff; to deliver the most current, standardized, technically comprehensive, and efficient web-based progression of modular curriculum and testing in the parallel file systems reseller arena; and to deliver compelling incentives and unparalleled preferential access to Intel Solutions for Lustre Software staff, resources, support and information. “We are going to bring Lustre resellers the very best information and tools to further their work and continue the momentum we are experiencing with Lustre adoption,” he said.

## **A NEW PROGRAM INCLUDES SUPPORTING INDEPENDENT DEVELOPER INNOVATIONS FOR LUSTRE THROUGH GRANTS FUNDED BY INTEL.**

### **INTEL ADVANCING LUSTRE INNOVATION THROUGH INTEL® PARALLEL COMPUTING CENTERS GRANTS**

Intel® Parallel Computing Centers are universities, institutions, and labs that are leaders in their field. The primary focus of the Centers is to modernize applications to increase parallelism and scalability through optimizations that leverage cores, caches, threads, and vector capabilities of microprocessors and coprocessors. A new program includes supporting independent developer innovations for Lustre through grants funded by Intel. The goals of the program include the following:

- Fund universities, labs, and institutions to perform Lustre development.
- Provide funding via grant proposals that get new talent involved in developing features that the community cares about.
- Require that the resulting work must end up in the community tree for the benefit of the community.
- Seek a broader development community advancing Lustre.

The recipient of the first approved grant is André Brinkmann of Johannes Gutenberg University, Mainz Germany for his proposal “Lustre QoS: Network Request Scheduler and Monitoring Revisited.”

### **LUSTRE LEADING THE CHARGE IN EXASCALE COMPUTING**

Of course, Lustre continues to be a key component of the next-generation of exascale supercomputers. To do this, Lustre has taken position along with other Intel technologies in the Intel scalable system framework, a flexible blueprint for developing high performance, balanced, power-efficient and reliable systems capable of supporting both compute- and data-intensive workloads. “To create the next generation of hi-

ghly efficient supercomputers, we need to make sure the right ingredients are combined in precisely the right way,” stated Al Gara, Intel Fellow and Chief Exascale Architect for Intel’s Technical Computing Group. “To do that, we must look at things from a holistic view, from a total system perspective. And out of that comes the definition and the development of those ingredients.” Intel Enterprise Edition for Lustre is a key element of this scalable system framework, which also includes next-generation Intel Xeon processors and Intel Xeon Phi coprocessors, Intel® Omni-Path fabric, silicon photonics, and innovative memory technologies, along with the ability to efficiently integrate them into a broad spectrum of system solutions.

The new Aurora system, being designed and built for Argonne National Laboratory under the Collaboration of Oak Ridge, Argonne, and Lawrence Livermore (CORAL) initiative, is designed around the Intel scalable system framework, and will include the Intel Enterprise Edition for Lustre software. According to Gara, the Intel scalable system framework helps OEMs, SIs, and end users to break through the performance wall that has challenged exascale computing. “The framework will help drive even more momentum behind Lustre, offering both our channel partners and end users an efficient blueprint for building out their next-generation clusters for the important work they want to do,” added Costelow.

### **NO END IN SIGHT**

“Lustre momentum is growing,” said Intel’s Gorda. “And we will continue to strongly press the applications of Lustre throughout the HPC community, whether in the labs, in academia, or the enterprise,” he added. “We think Lustre yet has a lot of opportunity to enhance HPC.”

**KEN STRANDBERG**

NEXT ISSUE



#9  
**NOVEMBER**  
**30<sup>th</sup>, 2015**

**VIEWPOINT**

**WHY FULL  
FLASH STORAGE  
IS GAINING  
MOMENTUM**

**COVER STORY**

**SOFTWARE DEFINED  
DATACENTERS**  
POWERING TOMORROW'S  
INFRASTRUCTURE

**HOW TO**

**THE WORLD'S  
FIRST ARM BASED  
HPC CLUSTER**

**TECH ZONE**

**SIGNIFICANT LUSTRE  
MOMENTUM IN ENTERPRISE  
AND RESEARCH HPC**

 **HPC Media**  
The HPC and Big Data Global Media

**HPC Media**  
11, rue du Mont Valérien  
F-92210  
St-Cloud, France

**CONTACTS**  
[editorial@hpcreview.com](mailto:editorial@hpcreview.com)  
[subscriptions@hpcreview.com](mailto:subscriptions@hpcreview.com)  
[sales@hpcreview.com](mailto:sales@hpcreview.com)

**PUBLISHER**  
**Frédéric Milliot**

**EDITOR IN CHIEF**  
**Joscelyn Flores**

**EXECUTIVE EDITOR**  
**Ramon Lafleur**

**CONTRIBUTORS  
TO THIS ISSUE**  
**Amaury de Cizancourt**  
**Steve Conway**

**Roberto Donigni**  
**Ramon Lafleur**  
**Stacie Loving**  
**Renuke Mendis**  
**Marc Paris**

**ART DIRECTOR**  
**Bertrand Grousset**

**SUBMISSIONS**  
We welcome submissions. Articles must  
be original and are subject to editing for  
style and clarity