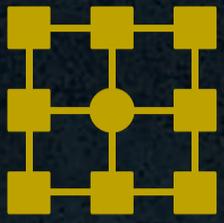


**NEWSFEED**

INTERVIEW **RATMIR TIMASHEV,**  
**CEO OF VEEAM**

**HOW TO**

**USING MICROSOFT AZURE**  
**FOR DATA BACKUP**



**HPC**

**REVIEW**

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

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

FLASH STORAGE  
**THE FINAL  
FRONTIER?**

**BREAKING THE \$1/GB BARRIER**





# 4 Blades

## HighServer XLR4 Blade

Virtualization Server with GPUs / Co-Processors

- High density
- Low power consumption
- Easy management

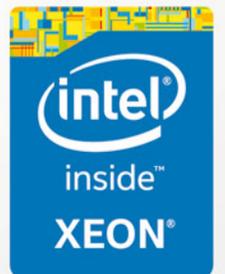
# GIGABYTE™

# 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.

# FLASH STORAGE, THE ULTIMATE FRONTIER

**WELCOME TO THIS 6TH ISSUE OF HPC REVIEW!** This is an outstanding issue in many aspects. You will certainly have noticed that we took advantage of the summer break to improve the readability and reading comfort of your magazine, while refreshing its graphic layout to add modernity and dynamism. It is also an opportunity to thank you, our readers, for your remarks and observations which have served us well to make your magazine evolve in the right direction.

**OUR COVER STORY EXAMINES THE IMPOSSIBLE EQUATION THAT STORAGE COMPANIES STRUGGLE TO RESOLVE, NAMELY** reconciling the performance improvements of flash technology over traditional mechanical hard drives and gaining in capacity in order to break the dollar per terabyte barrier before this year's end. Which will be synonymous of widespread adoption of full flash appliances, helping to give a welcome boost to numerous HPC and Big Data related areas. This story also gives you some hints of what is undoubtedly the future of storage once the technological obstacles have been overcome. While this ambitious objective is attained, the industry has quite convincing solutions up its sleeve. Happy reading!

**HAPPY READING!**

# FLASH STORAGE THE FINAL FRONTIER?



**NEWSFEED**

**LENOVO** UNVEILS FIRST SKYLAKE MOBILE XEON WORKSTATIONS

**RATMIR TIMASHEV,** CEO OF VEEAM

**INTEL AND MICRON** PRODUCE BREAKTHROUGH MEMORY TECHNOLOGY

**FUJITSU** DEVELOPS TECHNOLOGY TO VISUALIZE THE ENERGY REQUIRED TO EXECUTE SOFTWARE

**INTEL'S** COLLABORATIVE CANCER CLOUD USES BIG DATA TO FIGHT DISEASE

**8.4 MINUTES** TO DECODE THE HUMAN GENOME

**IBM & GENCI TEAM** TO DRIVE SUPERCOMPUTING CLOSER TO EXASCALE

**LAB REVIEW**

**DELL PRECISION WORKSTATION M3800**

**DELL PRECISION WORKSTATION M3800**

**ENGENIUS NEUTRON EWS7928P**

**VEEAM ENDPOINT BACKUP**

**HOW TO**

**USING MICROSOFT AZURE FOR DATA BACKUPS**

**VIEWPOINT**

**CYBER SECURITY IN HIGH-PERFORMANCE COMPUTING ENVIRONMENT**

**TECH ZONE**

**WHAT IS CODE MODERNIZATION?**

EDITOR'S NOTE

OPINION

SOURCES

THE HPC OBSERVATORY

**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.

# POWER AND COOLING: THE SWORD OF DAMOCLES?

**CONSISTENTLY RANKED AS THE NUMBER TWO CONCERN FOR HPC DATA CENTERS, POWER AND COOLING FACE BIG UNKNOWNNS**



**STEVE CONWAY**

IDC RESEARCH VICE PRESIDENT, HPC

**F**ifteen years ago, power and cooling didn't make the top 10 list of issues HPC data centers were facing. That changed quickly with the rise to dominance of clusters and other highly parallel computer architectures, starting in the period 2000 to 2001 and escalating from there. In IDC's worldwide surveys since 2006, power and cooling have consistently ranked as the number two concern for HPC data centers, right behind the perennial quest for bigger budgets.

## **SPENDING IS NOT RACING AHEAD YET**

Despite this elevated concern, during the period 2006 to 2013, the portion of HPC budgets devoted to power and cooling held steady at between eight and nine percent on average. True, the average budget increased substantially during this period, so the eight to nine percent figure was larger in 2013 than in 2006, and larger in 2006 than in 2000. But the absolute increases in spending on power and cooling pale in comparison with the explosive growth in the rated performance of HPC systems since 2000, especially supercomputer-

class systems. Between April 2000 and September 2013, the average peak performance of systems on the Top500 list skyrocketed from 154 GF to 652 TF, a factor of 4,217, while average TOP500 Linpack performance jumped 4,373-fold, from 102 GF to 446 TF.

## **MOST SITES ARE MANAGING WELL**

Rampant growth in HPC system sizes and processing power has not caused power and cooling budgets to spiral out of control yet, for a number of reasons.

HPC vendors have made their technologies and products substantially more power-efficient. The list of technologies contributing to improved power efficiency is long. It includes x86 and other base processors, along with GPGPUs and Intel Phi, both of which are notably energy efficient but have limited applicability today. Another important contributor is the trend toward liquid cooling, which is several times more energy-efficient than air cooling. Liquid cooling appears in many guises, from water-chilled doors to exotic immersive implementations.

# AS WE APPROACH THE EXASCALE ERA, IS THE POWER AND COOLING ISSUE A DISRUPTIVE SWORD OF DAMOCLES HANGING OVER THE HPC COMMUNITY?

HPC sites have been on a tear updating their power and cooling infrastructures or building new ones. In a recent IDC worldwide study, two-thirds of the HPC sites had budgets in place to upgrade their power and cooling capabilities, to the average tune of about \$7 million. Not surprisingly, government sites typically have the largest systems and are under the greatest energy and spatial pressure. Industrial sites are least constrained — half of them never see their energy bills, because someone else in the company receives and pays them. By and large, HPC data centers, with the help of vendors, have been coping so far with rising requirements for power and cooling, even at an average cost of about \$1 million per megawatt. The real question is what the future holds.

## THE SWORD OF DAMOCLES?

As we approach the exascale era, is the power and cooling issue a disruptive sword of Damocles hanging over the HPC community, or will it continue to be manageable?

The study I cited earlier asked HPC users and vendors whether they expected any revolutionary advances in power and cooling technology in the next five years. The users said no, the vendors said yes, and they were referring in most cases to the same advances (because the vendors had briefed the users about their plans).

The users are less optimistic about breakthroughs, but most do not seem heavily concerned yet, except on one important point: potential tradeoffs between productivity and power efficiency. Tradeoffs frequently mentioned by users include pressure to overbuy energy-efficient, harder-to-program coprocessors and accelerators; and accepting more service disruptions because of shorter upgrade cycles to deploy more energy-efficient systems.

But these tradeoffs seem manageable in the sense that they are based on decisions users will make. Even the vaunted goal of fitting an exascale computer into a 20 MW power envelope is a matter of time, of when rather than if. It might happen in 2020, 2024 or a different year, but it will happen.

## THE BIGGER UNKNOWN

Much less certain is when deeper, more integral energy-efficiency capabilities will become available. These are considerably more challenging than squeezing a peak exascale into a 20 MW package, however difficult that may be. These deeper capabilities will go a long way toward making exascale and lesser extreme-scale computing a reasonable proposition for funders and users alike.

In particular, sophisticated power management (“power steering”) will be needed throu-

# HPC SITES WILL LIKELY HAVE FELLOW TRAVELERS IN THE FORM OF MAJOR INTERNET PLAYERS.

ghout the system to dynamically shift power to where it's needed at every moment. Both hardware and software will need to be able to “learn” about power needs on the fly. Achieving this goal will require large investments to develop software that can power profile and power-steer many elements of the system, including:

- Cores and processors
- The interconnect and network interface
- The storage system
- The operating system, programming model and entire software stack
- Application codes (power-aware applications)

The HPC community is capable of developing these and other needed capabilities, given enough time, money and personnel. So, the real question, as with so many major HPC undertakings, is when will these necessary elements come together in sufficient quantity? If the past is any guide, a large chunk of the funding will need to come from government sources. And for that to happen, it has become increasingly clear that the HPC community will need to make a strong case for the returns government funders can anticipate from major HPC investments like this.

This discussion so far assumes that HPC data centers will have access to enough reliable energy, even in the exascale era. That is not a given. Today's largest HPC systems already consume as much electricity as a small city,

and their exascale successors promise to devour more, even with expected advances in energy-efficiency. Some of the biggest HPC data centers worry that their local power companies may balk at fully supplying their future needs. A few sites have “plan B” scenarios in place, in which they go off the grid and build small nuclear reactors. And in some parts of the world, reliable access to adequate power is already a major challenge for HPC data centers — an important reminder that power and cooling are concerns not only for sites marching toward exascale capacity, but for most HPC sites.

## FELLOW TRAVELERS

In their pursuit of energy efficiency at extreme scale, HPC sites will likely have fellow travelers in the form of major Internet players. A pattern is already forming in which these companies locate new data centers in geographical areas where power is comparatively cheap and plentiful. Google set the tone more than five years ago, by building a vast new data center along the Columbia River near Oregon's Dalles Dam, with its 1.8-gigawatt power station and relatively inexpensive hydroelectric power. A prominent HPC example is Oak Ridge National Laboratory, whose power appetite is fed by the Tennessee Valley Authority. The lab's data center hosts multiple petascale systems from DOE, NSF and NOAA. The biggest HPC and Internet data centers will likely have much to learn from each other in the coming years.

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)



LENOVO UNVEILS  
**FIRST SKYLAKE  
MOBILE XEON  
WORKSTATIONS**

**“WE ARE EXCITED TO BRING INTEL XEON’S LEGACY OF PERFORMANCE, MANAGEABILITY AND RELIABILITY TO MOBILE WORKSTATIONS,”** said Navin Shenoy, vice president and general manager, Mobility Client Platforms Group, Intel. **“NOW THE POWER OF INTEL XEON IS IN A NOTEBOOK THAT CAN HANDLE THE UNIQUE COMPUTE AND GRAPHICS INTENSIVE NEEDS OF ENGINEERS AND DESIGNERS.”**

Lenovo unveiled a new family of mobile workstation, the ThinkPad P50 and P70, at the SIGGRAPH 2015 Conference and Exhibition. These new ThinkPad P Series systems are the best-performing and most reliable mobile workstations available and are designed to run the most demanding Independent Software Vendor (ISV) applications.

“We’ve built features into these machines that were previously unachievable in a notebook, making them the most versatile and highest-performing mobile workstations ever,” said Victor Rios, vice president and general manager, Workstations, Lenovo. “We’re focused on making sure users have the tools necessary to drive innovation. That is why we are expanding our portfolio and raising the standard of mobile workstation performance. Introducing the ThinkPad P Series unites our mobile portfolio with our existing award-winning line of tower workstations, the ThinkStation P Series.”



### **MEET THE NEW THINKPAD P SERIES**

**ThinkPad P70:** Designed for high-end professionals in industries ranging from media and entertainment to oil and gas, the Lenovo ThinkPad P70 introduces a 17-inch mobile workstation to the ThinkPad family. The new offering comes packed with the most memory and storage ever found in a mobile workstation. The ThinkPad P70 is loaded with up to 64GB of DDR4 ECC memory and the ability to handle up to four storage devices and up to a terabyte of SSD storage, utilizing the latest PCIe technology for speeds up to five times faster than current SATA technology. Additionally, the ThinkPad P70 comes with two Intel® Thunderbolt™ 3 ports for ultra-fast connectivity and a 4K UHD display or optional FHD touch.

“Media and entertainment professionals invest in complex and incredibly resource-intensive applications to bring their creations to life. They are increasingly looking for the performance to run these applications in a mobile workstation,” said visual effects industry pioneer Scott Ross. “The Lenovo ThinkPad P70 gives these users all the performance they need and more to continue to raise the bar for graphics innovation and move their industry forward. It’s the best mobile workstation I’ve seen.”

**ThinkPad P50:** Lenovo’s thinnest and lightest full-function mobile workstation yet, the ThinkPad P50 features a beautiful 15.6-inch UHD 4K display and is certified to run users’ most requested ISV applications. The ThinkPad P50 is a feature-rich, highly powerful mobile workstation for performance-see-

**“WORKSTATION USERS NEED THE BEST GRAPHICS PERFORMANCE POSSIBLE TO HARNESS THE FULL POTENTIAL OF THEIR PROFESSIONAL APPLICATIONS,”** said Greg Estes, vice president of Enterprise Marketing, NVIDIA. **“BY USING NVIDIA QUADRO MOBILE GPUS, THE LENOVO THINKPAD P SERIES WILL PROVIDE CUSTOMERS WITH THE MOST ADVANCED GRAPHICS TECHNOLOGY AVAILABLE TO ACHIEVE THE HIGHEST QUALITY RESULTS, REGARDLESS OF THE TASK.”**

king users. As the follow-on to the ThinkPad W541, the P50 builds on the innovation and purposeful design of Lenovo’s industry-leading ThinkPad mobile workstations.

### **THE ULTIMATE IN MOBILE WORKSTATION PERFORMANCE**

The new ThinkPad P50 and P70 are mil-spec tested for maximum durability, and also come loaded with the most advanced components and features available in a mobile workstation. These machines are the first equipped with the new Intel® Xeon® Processor E3-1500M v5 product family, for lightning-fast performance and enhanced reliability for critical workstation applications. Both systems feature NVIDIA Quadro GPUs for unmatched graphics capability. The ThinkPad P Series also come with X-Rite Pantone color calibration, keeping colors accurate.

The ThinkPad P50 and P70 feature a brand new FLEX Performance Cooling system. The dual-fan design allows for cooler, quieter and faster performance. This revolutionary design allows for both optimal uptime, as dictated by



the needs of the application, giving the user the ability to push the system harder and for longer periods of time.

### **PRICING AND AVAILABILITY**

The new Lenovo ThinkPad P Series mobile workstations will be available starting in Q4 2015 through Lenovo business partners and on [www.lenovo.com](http://www.lenovo.com). Pricing for the P50 begins at \$1599, and the P70 begins at \$1999.

**“OUR CUSTOMERS ARE USING SOLIDWORKS TO MAKE GREAT DESIGNS HAPPEN. THE LENOVO THINKPAD P SERIES IS DESIGNED TO OFFER THE PERFORMANCE THEY NEED WITHOUT HAVING TO SACRIFICE FUNCTIONALITY OR RELIABILITY,”** said Nicholas Iwaskow, SolidWorks, Director Alliances and Partnerships, Dassault Systèmes. **“WE ARE EXCITED TO SEE HOW THE THINKPAD P SERIES CONTINUES TO EVOLVE IN THE YEARS TO COME.”**



# RATMIR TIMASHEV

CEO of Veeam



«In 2018, we plan to reach a billion dollars in revenue.»



eeam is the leading company specializing in backup software for virtualized environments. On the occasion of his visit to Paris, HPC Review interviewed its CEO Rattimir Timashev.

Can you describe your background?

Before founding Veeam, I founded the company Aelita, specializing in Windows Server management solutions, Active Directory and Exchange. Aelita was sold to Quest in 2004 for \$115 million. I left Quest Software in 2005 to start a new business, Veeam which I co-founded in 2006 with my associate Baronov with a simple concept, but a precursor for the time: provide systems management solutions for virtualized server infrastructure and so take advantage of the still emerging trend at the time, consisting in virtualized work environments based on VMware GSX and ESX.

How did you come to position yourself as early on the only emerging area of virtualization?

My partner and I have studied the market for a year, and we realized that virtualization began to take the form of a trend. In 2006, after founding Veeam, we listened companies to understand their needs and started showing a prototype of our solution that we made evolve accordingly. It took us a year and a half of development before launching our first product in 2008. In six years, our turnover has reached 275 million, of which 90% came from our flagship product, Veeam Backup & Replication.

Besides the corporate feedback, did you have other indicators of the emergence of virtualization?

Yes. By 2006, 69 percent of US companies had implemented virtualized environments. Another indicator was the growing success of VMware. So it appeared logical and relevant to focus on this sector that had concrete needs around the management and backup of virtualized infrastructure and no simple alternative to expensive and proprietary solutions.

How did Veeam's software offer evolve over time?

Veeam has developed a range of solutions for modern data protection specifically for virtualization and we have extended this approach to the cloud with our Veeam Cloud Connector, designed for Data Center and Cloud Service Providers. This from a simple finding related to the evolution of the market: Virtualization is now the new standard for enterprise data center - it is the norm rather than the niche.

To what factors do you attribute your business success?

We have a remarkable annual growth of our turnover from 62% in 2012 and 58% in 2013 and just crossed half a billion dollars in yearly revenue. We are rapidly gaining market share over the major established players in the enterprise backup such as CommVault and Symantec. We gain more than 3,000 new customers a month. At this rate, I plan to cross the billion dollars in sales in less than three years, in 2018.

Do you define the strategic direction of the company's development by yourself?

Almost. We are two, with my partner Andrei, who is a scientist and brilliant technologist cou-



pled with an outstanding businessman. It is at the heart of the definition of our products, and resolving development problems encountered along the way. We met when we were 18 years old at high school. We grew up together. Science had become a prestigious profession in Russia when we received our undergraduate degrees in the 80s.

What characterizes your products and sets them apart from the competition?

Besides the intrinsic qualities of our software in which we firmly believe, because it is the intersection of a continuous observation of competition and new ideas that we develop continuously, we distinguish two approaches. The first concerns a decision we made early on, namely an extremely short iterative development cycle of about three to six months. The second element relates to the method of dissemination, inspired from the video game: freemium. We are convinced of the quality of our software, and offer them for free and unlimited. The most advanced features are of course available in paid versions, but we have seen the benefits of this distribution channel, which does as much to convince our customers than our standard business practices.

You recently released a software that backups physical machines. For what reason?

The reputation of our software in virtualized environments is at its cusp. We decided to extend the simplicity and efficiency of our solutions on client computers. It also has the virtue of allowing us to gain a foothold in other areas already occupied by our competitors, but with a simple, elegant and free solution. We want to convince the market, as we did with our virtualized backup solutions.

This first version does not seem functionally equipped to counter your competitors.

It's normal for us. We released the first version voluntary incomplete in order to get it out the fastest possible. This software will improve over time, incorporating our ideas and the feedback of our users and customers.

How is your revenue distributed in the world?

Approximately 35% of our revenue comes from the United States, 50% from Europe and 15% from the rest of the world. 15% of our customers use Hyper-V, and 85% use VMware. We protect 8.4 million virtual machines worldwide which amount currently to a total of 40 to 50 million.

What are your current and future projects?

We launched Veeam FastSCP for Azure, which is the equivalent of the legendary Norton Commander, but applied to the Cloud. Administrators worldwide have welcomed this initiative! We are also working on the next version of Veeam Backup and Replication 9 scheduled in November, and whose characteristic will be based on the notion of replicas and containers, to allow the establishment of disaster recovery plans by simply setting up a cloud connector.

How do you see the future of Veeam?

I am fortunate to have learned the rules of profitable and efficient sales and marketing and my partner is a world-class software developer. We operate in the hottest segment of the market of the moment. Together we are armed to make Veeam the next champion in virtualization and cloud technologies.

INTEL AND MICRON  
PRODUCE

# BREAKTHROUGH MEMORY TECHNOLOGY

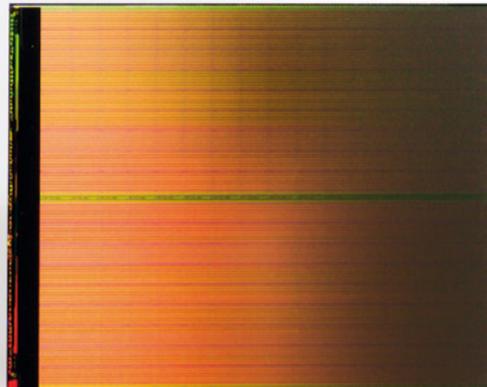
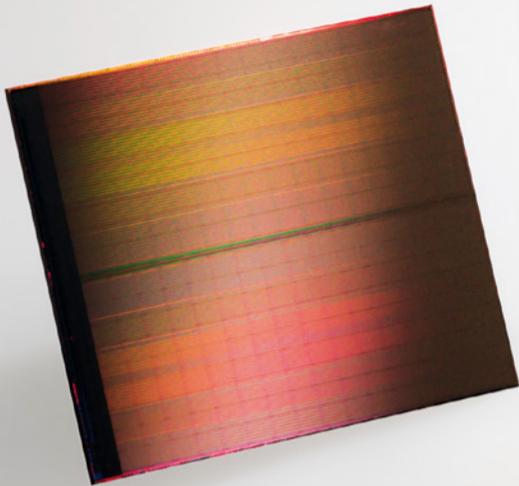
# A

nnounced earlier this year, Intel and Micron will begin production on 3D Xpoint, a new class of non-volatile memory, creating the first new memory category in more than 25 years.

3D XPoint technology brings non-volatile memory speeds up to 1,000 times faster than NAND, the most popular non-vo-

latile memory in the marketplace today. The companies invented unique material compounds and a cross point architecture for a memory technology that is 10 times denser than conventional memory. 3D XPoint technology is a major breakthrough in memory process technology and the first new memory category since the introduction of NAND flash in 1989.

The explosion of connected devices and digital services is generating massive amounts of new data. To make this data useful, it must be stored and analyzed very quickly, creating



As the digital world quickly grows – from 4.4 zettabytes of digital data created in 2013 to an expected 44 zettabytes by 2020, 3D XPoint technology can turn this immense amount of data into valuable information in nanoseconds. Businesses could use this new capacity to more quickly identify fraud detection patterns in financial transactions; healthcare researchers could process and analyze larger data sets in real time, accelerating

challenges for service providers and system builders who must balance cost, power and performance tradeoffs when they design memory and storage solutions. 3D XPoint technology combines the performance, density, power, non-volatility and cost advantages of all available memory technologies on the market today. The technology is up to 1,000 times faster and has up to 1,000 times greater endurance than NAND, and is 10 times denser than conventional memory.

“For decades, the industry has searched for ways to reduce the lag time between the processor and data to allow much faster analysis,” said Rob Crooke, senior vice president and general manager of Intel’s Non-Volatile Memory Solutions Group. “This new class of non-volatile memory achieves this goal and brings game-changing performance to memory and storage solutions.”

“One of the most significant hurdles in modern computing is the time it takes the processor to reach data on long-term storage,” said Mark Adams, president of Micron. “This new class of non-volatile memory is a revolutionary technology that allows for quick access to enormous data sets and enables entirely new applications.”

complex tasks such as genetic analysis and disease tracking.

### **NEW RECIPE, ARCHITECTURE FOR BREAKTHROUGH MEMORY TECHNOLOGY**

Following more than a decade of research and development, 3D XPoint technology was built from the ground up to address the need for non-volatile, high-performance, high-endurance and high-capacity storage and memory at an affordable cost. It ushers in a new class of non-volatile memory that significantly reduces latencies, allowing much more data to be stored close to the processor and accessed at speeds previously impossible for non-volatile storage.

The innovative, transistor-less cross point architecture creates a three-dimensional checkerboard where memory cells sit at the intersection of word lines and bit lines, allowing the cells to be addressed individually. As a result, data can be written and read in small sizes, leading to faster and more efficient read/write processes.

3D XPoint technology will sample later this year with select customers, and Intel and Micron are developing individual products based on the technology.

# FUJITSU

DEVELOPS TECHNOLOGY  
**TO VISUALIZE THE ENERGY  
REQUIRED TO EXECUTE  
SOFTWARE**



**F**ujitsu Laboratories Ltd. Has announced that it has developed technology that precisely calculates the energy required to execute various software programs running on server CPUs, for energy-efficient programming. Servers equipped with Intel-made CPUs include a power-control mechanism that can measure power consumption for the CPU as a whole. Until now, however, it was not possible to calculate the energy required to execute software on a core-by-core basis, so it has been difficult to take a



high levels, there is a need to reduce overall energy consumption.

## ISSUES

One way to reduce energy consumption is through the use of more energy-efficient hardware. Another is to reduce the energy required to run programs on servers. A precondition for energy-efficient programming is to have an understanding of the energy being consumed by existing software. Servers equipped with Intel-made CPUs include a power-control mechanism called RAPL (2) that can be used to measure power consumption for the CPU as a whole. But this has not extended to being able to analyze

software-based approach to reducing power consumption. Now Fujitsu Laboratories has developed technology that uses information that can be tracked at the individual core level, such as clock cycles and cache-hit percentages, to estimate energy consumption in detail, down to the program module level. This makes energy-efficient programming a more efficient process, contributing to both lower overall server energy usage and, by using surplus power, higher software performance. Details of this technology are being presented at the Summer United Workshops on Parallel, Distributed and Cooperative Processing 2015 (SWoPP 2015), opening in Beppu, Oita Prefecture, on August 4.

## ENERGY CONSUMPTION SURVEILLANCE

As the scale and processing volumes of systems such as datacenters and supercomputers expand, so too does their energy consumption. For example, in the case of a top-level, high-performance supercomputer, power consumption is thought to reach roughly 18 megawatts. According to a report by the Ministry of Internal Affairs and Communications (1), Japan's datacenters consume an average of 7.72 billion kWh per year. Given these

energy consumption at an individual core level, which is where software runs. This has made it difficult to get a detailed picture of the energy requirements of software.

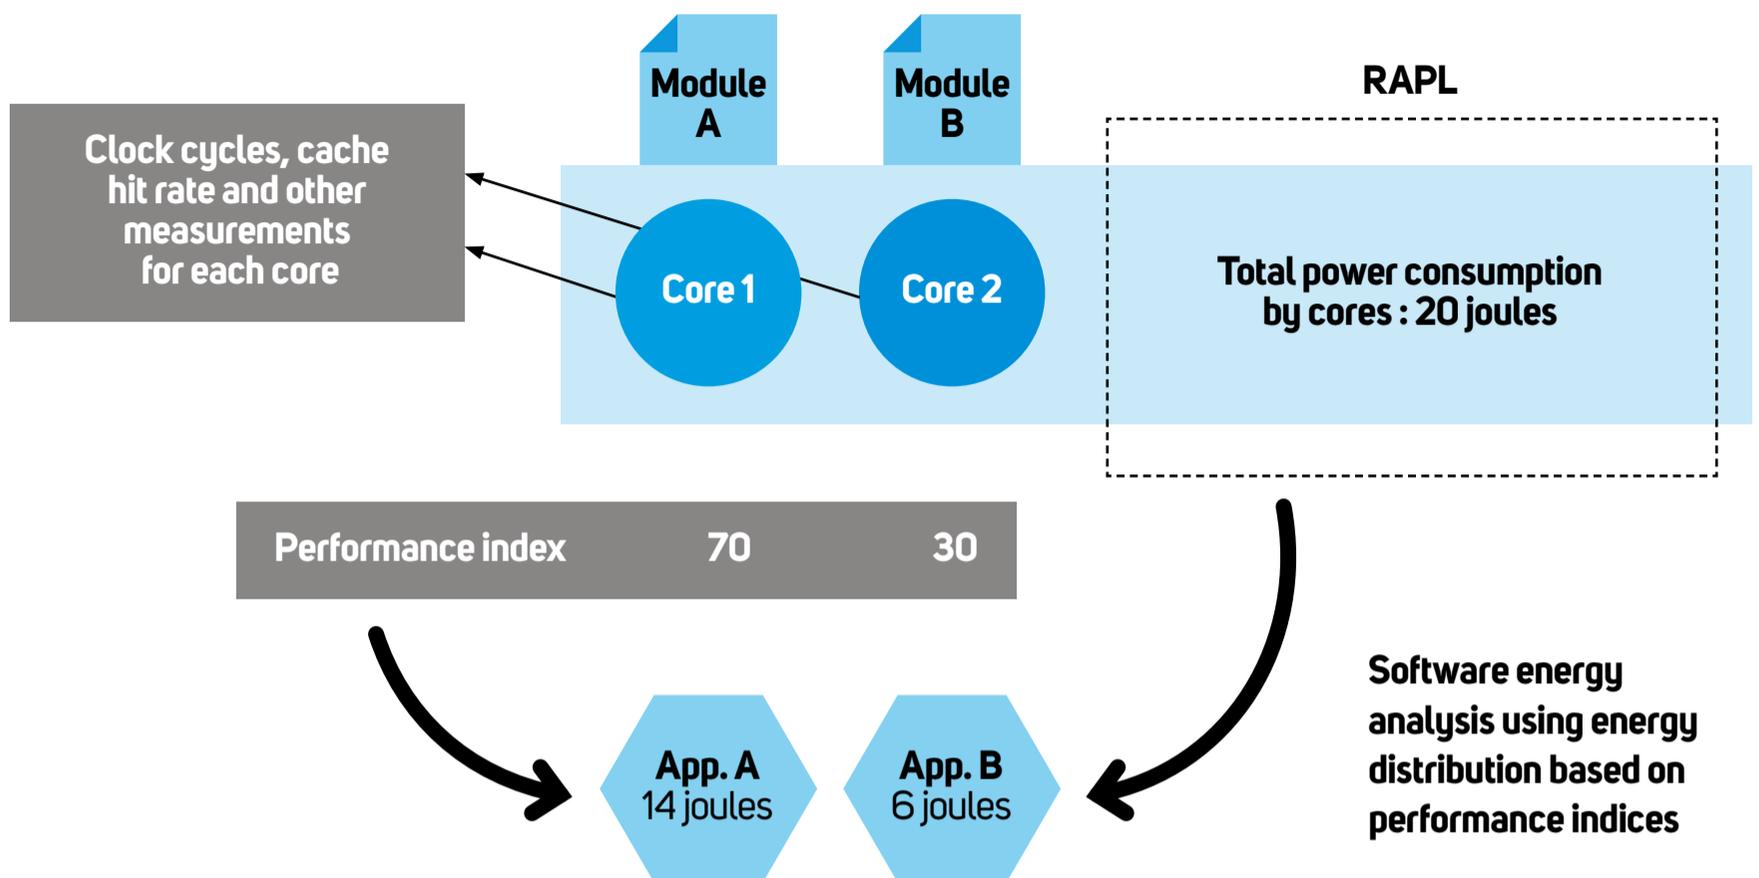
About the Technology

## AN ANALYTIC METHOD

Fujitsu Laboratories has developed an analytic method that can make precise estimates of software energy requirements (Figure 1) for servers equipped with Intel-made CPUs. The technology has the following characteristics:

### 1. PERFORM DETAILED ANALYSIS OF SOFTWARE ENERGY REQUIREMENTS BY USING POWER DISTRIBUTION BASED ON PERFORMANCE INDEX

Combining measurements that can be captured at the level of the individual CPU core, such as clock cycles and cache-hit percentages, Fujitsu Laboratories newly devised performance indices with a high degree of correlation to energy consumption. By distributing total CPU energy consumption across each core according to the index values calculated per CPU core, it becomes possible to get a detailed picture of energy consumption on a program module basis.



## 2. CALCULATING ENERGY INFORMATION WITH LOW OVERHEAD

The types of measurements used as performance indices for each CPU core are limited, which enables information capture down to the millisecond. The calculation of energy information imposes an overhead of roughly only 1% of the total, meaning that the measurement does not significantly impact performance. Additionally, fine-grained sampling makes it possible to get a highly accurate picture of the energy required to execute a given software program.

This technology promises to help software developers tune their software for lower energy requirements (energy-efficient programming) in order to reduce overall server energy consumption and, by using spare power and increasing parallelism, to boost software performance.

## FUTURE PLANS

Fujitsu Laboratories is proceeding with tests of reducing the energy required to execute software with the goal of a practical implementation of this technology in fiscal 2016. The company is also looking into applying the technology to Fujitsu's own datacenters, with the goal of analyzing datacenter energy consumption in more detail to improve datacenter energy efficiency.

[1] Report by the Ministry of Internal Affairs and Communications See «Report of the Study Group on ICT Policy to Address Global Warming,» published April 2008 by the Ministry of Internal Affairs and Communications.

[2] RAPL «Running Average Power Limit.» A mechanism for controlling and measuring electric power in Intel-made CPUs.

**THE COMPANY IS ALSO LOOKING INTO APPLYING THE TECHNOLOGY TO FUJITSU'S OWN DATACENTERS.**

# INTEL'S COLLABORATIVE CANCER CLOUD **USES BIG DATA TO FIGHT DISEASE**



Big Data tends to be uttered in the same breath as business – a lot of air is spent by experts talking about the commercial benefits of leveraging customer data. But at IDF in San Francisco, Intel has given a glimpse of how large-scale data analysis could make a real change to people's health, potentially revolutionizing the treatment of diseases such as cancer and Alzheimer's.

# AT THE HEART OF THE NEW SYSTEM IS GENOMIC DATA, WHICH INVOLVES THE SEQUENCE OF AN INDIVIDUAL'S DNA.

**D**uring a session on the Internet of Things (IoT) and Big Data, Intel revealed a new collaboration with Oregon Health & Science University (OHSU) on a way to bring greater precision to the treatment of cancer. The platform, called the Collaborative Cancer Cloud, will allow hospitals and research institutions to share important patient data to bring about precise individual healthcare.

At the heart of the new system is genomic data, which involves the sequence of an individual's DNA. Genomic data is key for creating precise treatments for cancer, but the length of time it takes to order and communicate this data is insufficient. Universities do not currently have the computing resources or infrastructure to efficiently share genomic and clinical data. The Collaborative Cancer Cloud aims to solve this problem by allowing institutions to analyze data in a distributed way.

## GENOMIC ANALYSIS IN JUST ONE DAY BY 2020

During the presentation Eric Dishman, Intel's general manager of Health & Life Sciences, was brought on stage to speak about the platform. Dishman is a cancer survivor, and spoke about how genomic analysis led to him being cured of cancer after 23 years of suffering from the disease.

"It took 23 years to develop a personalized treatment for Eric," said Dr Brian Druker, a leading researcher at OHSU. "We should be able to do this in one day."

Intel's aim is to make the treatment Dishman received possible in as little as one day by 2020. A key step in this direction is creating a way for researchers to communicate the vast quantities of genomic, imaging and clinical data, initially between OHSU and universities in Boston and Texas. Another important benefit to the Collaborative Cancer Cloud is the system's alleged security. Diane Bryant, senior vice president and general manager of Intel's Data Centre Group, said during the presentation that information is wiped from the platform after being shared by researchers, preventing confidential information from being misused.

## OPEN SOURCE PLATFORM

Bryant went on to announce that the Collaborative Cancer Cloud will be open source, which is intended to allow analytics across a broader set of data and lead to greater insights for personalized healthcare.

"Each year millions of people all over the world, including more than one million patients in the United States, learn that they have a cancer diagnosis," said Eric Dishman in a statement. "Instead of going through painful chemotherapy that can kill healthy cells along with cancerous cells, what would happen if those patients were able to be treated as individuals based on their specific genome sequencing, and a precision treatment plan could be tailored specifically for their disease? And what if it could happen within 24 hours?"

The implications for this platform extend beyond cancer treatment, towards any disease that encompasses genomics. This means it could be also be used to help treat Alzheimer's, autism and diabetes. This is enormous, and is a hugely encouraging indication of the positive societal applications of Big Data.

# 8.4 MINUTES TO DECODE THE HUMAN GENOME

A team from the Joint Genome Institute at Lawrence Berkeley National Lab and researchers from UC Berkeley have used 15,000 cores on the Cray XC30 Edison supercomputer to boost the complete assembly of the human genome, bringing the time down to 8.4 minutes.



## EDISON BY THE NUMBERS

332 **terabytes**  
**memory**

2.39 **petaflop/second**  
**peak performance**

124,608 **processing**  
**cores**

462 **terabytes/**  
**second global memory**  
**bandwidth**

11 **terabytes/second**  
**network bisection**  
**bandwidth**

7.56 **petabytes disk**  
**storage**

163 **gigabytes/second**  
**I/O bandwidth**

**T**he work represents how the coupling of high-powered computational capacity, matched with novel approaches to complex code parallelization, can significantly speed large-scale scientific research. For genomics, this speedup marks a dramatic improvement in what was possible before for using de novo assemblers to rebuild a genome from a selection of short reads. For instance, on the same machine, the unmodified Meraculous code took 23.8 hours. For the far more complex wheat genome, full assembly has been difficult for most standard de novo assemblers. The team used their HipMer approach to scale wheat genome assembly across 15,000 cores in just under 40 minutes.

The team says that while the new approach has “inherent advantages of discovering variations that may remain undetected when aligning sequence data to a reference genome,

unfortunately, new assembly computational runtimes cannot up with the data generation of modern sequencers.”

### **LARGE-SCALE PARALLELIZATION**

The large-scale parallelization effort, which the team calls HipMer (short for high performance Meraculous—the name of the de novo assembler) picks apart the algorithm for end-to-end parallelization, an effort that involved extensive code work on multiple steps of the genome assembly process. While the full details of the parallelization effort can be found here, the takeaway is that there might be a new approach for genome research centers to get around one of the most pressing computational problems—keeping up with the massive amount of data against extremely complex algorithms.

# FOR THE FIRST TIME, ASSEMBLY THROUGHPUT CAN EXCEED THE CAPABILITY OF ALL THE WORLD'S SEQUENCERS.

“For the first time, assembly throughput can exceed the capability of all the world’s sequencers, thus ushering in a new era of genome analysis...The combination of high performance sequencing and efficient de novo assembly is the basis for numerous bioinformatics transformations, including advancement of global food security and a new generation of personalized medicine.”

## PERFORMANCE BY DESIGN

The Edison supercomputer was designed to tackle this exact type of work—complex scientific codes that are oftentimes just as data-intensive as they are computationally demanding. The Edison supercomputer, which was put into production just over a year ago, was designed to be balanced for both data-heavy jobs and those that need a lot of powerful cores. When the team at NERSC worked with Cray to push out their XC30, the emphasis was on the all-important Cray Aries interconnect, but also on having high memory bandwidth and memory on each node. The end result is a machine with some impressive I/O speeds that match the computational boost from the 124,608 cores. Edison provides 163 gigabytes per second of I/O bandwidth and as a whole, is capable of 2.4 petaflops at peak.

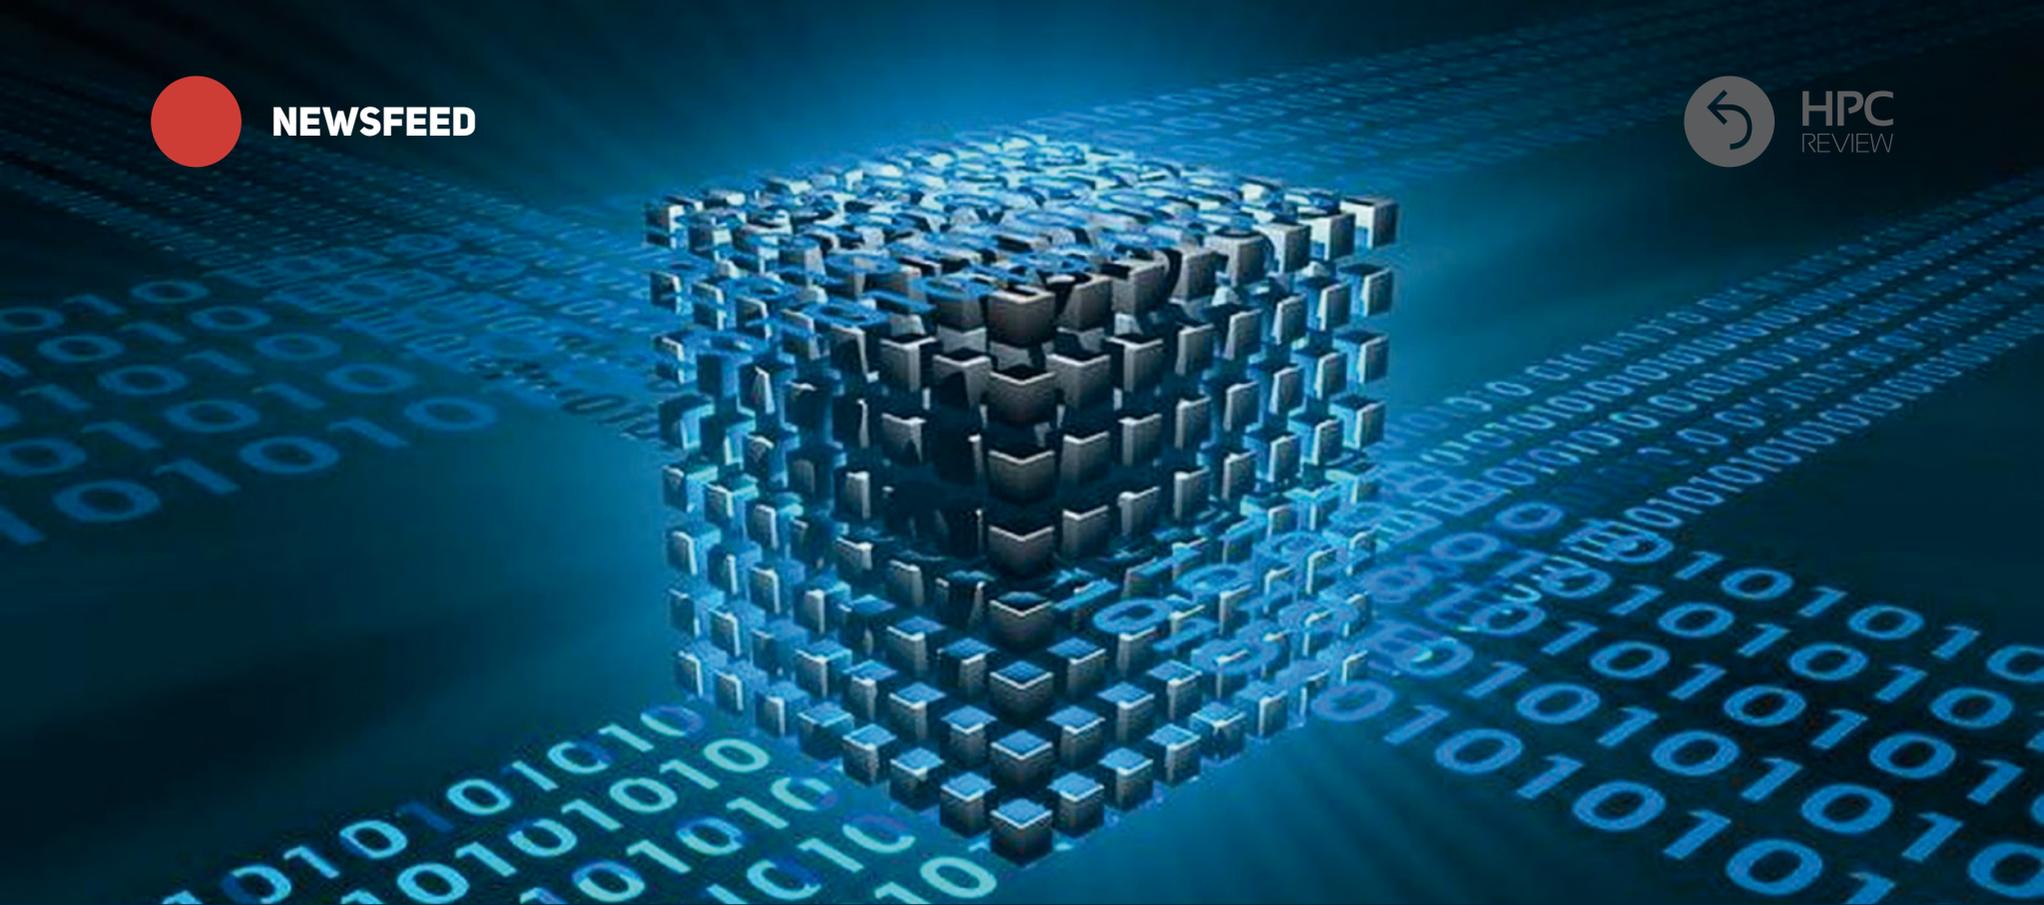
The end-to-end scaling of the team’s HipMer genome assembler approach showing the hu-

man genome scalability on Edison on the left and the more complex wheat genome on the right (both axes are in log scale).

## EFFICIENT CODE FOR SCALING

Naturally, the supercomputer cores and memory bandwidth are important, but so too is the work put into carving the code so it can efficiently be split over the 15,000 12-core Ivy Bridge cores. The work was built off the ubiquitous Meraculous assembler, which is generally thought to be the most advanced de novo genome assembler. The resulting HipMer approach is based on “several novel algorithmic advancements by leveraging the efficiency and programmability of UPC, including optimized high-frequency k-mer analysis, communication avoiding de Bruijn graph traversal, advanced I/O optimization, and extensive parallelization across the numerous and complex application phases.” Other improvements include altering the communication model and using a global address space for shared access for reads and writes from any of the many processors.

The team also developed a distributed memory approach for HipMer, which was able to tackle the wheat genome in less than 11 minutes across 20,000 cores. The next phase of the research with HipMer will involve applying it to broader metagenomics studies and extended assembler research. **NICOLE HEMSOOTH**

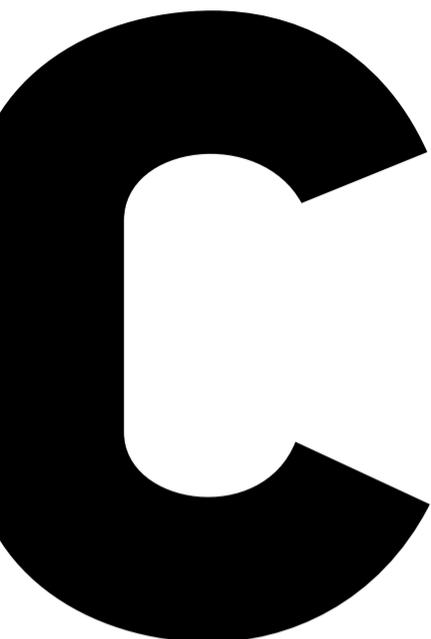


IBM & GENCI TEAM TO DRIVE SUPERCOMPUTING

# CLOSER TO EXASCALE



IBM and GENCI, the high performance computing agency in France, today announced a collaboration aimed at speeding up the path to exascale computing – the ability of a computing system to perform at least one exaflop, or a billion billion calculations, in one second.



Currently the fastest systems in the world perform between ten and 33 petaflops, or ten to 33 million billion calculations per second – roughly one to three percent the speed of exascale. Put into context, if exascale computing is the equivalent of an automobile reaching 1000 miles per hour, today's fastest systems are running within a range between ten and 33 miles per hour.

The collaboration, planned to run for at least 18 months, focuses on readying complex scientific applications for systems under development expected to achieve more than 100 petaflops, a solid step forward on the path to exascale. Working closely with supercomputing experts from IBM, GENCI will have access to some of the most advanced high performance computing technologies stemming from the rapidly expanding OpenPOWER ecosystem. Supported by more than 140 OpenPOWER Foundation members and thousands of developers worldwide, the OpenPOWER ecosystem includes a wide variety of computing solutions that use IBM's licensable and open POWER processor technology.

As part of the collaboration, GENCI will closely examine the impact and requirements of POWER's open architecture on scientific applications, intending to foster a deeper understanding of application requirements as the computing industry advances towards exascale computing with an increased interest in accelerator technologies.

### **OPENPOWER PROCESSORS, NVIDIA NVLINK AND MELLANOX INFINIBAND**

The collaboration will attempt to take full advantage of the impact of OpenPOWER-based innovations such as the connection of NVIDIA GPUs accelerators to POWER processors through the high-speed NVIDIA NVLink interconnect, as well as how Mellanox EDR 100Gb/s InfiniBand switches can exploit IBM's

Coherent Application Processor Interface (CAPI) to dramatically improve solution performance. Additionally, experts from GENCI and French research organizations together with IBM plan to work on understanding the evolution of programming models, considering MPI and OpenMP as a first step for shared memory multiprocessing programming. Alternative application program interfaces will also be considered, given potential changes may be required as systems move closer toward exascale.

"If we want to continue to address the challenges of the French scientists and engineers, we need to anticipate the rise of new high performance computing architectures that bring us closer to exascale and prepare our communities," stated Catherine Riviere, CEO of GENCI.

IBM will provide dedicated technical experts to support application porting and optimization efforts as well as organizing along with GENCI education and porting sessions. This collaboration will be supported by the newly created POWER Acceleration and Design Center in Montpellier as part of the partnership established with both NVIDIA and Mellanox. The Center will provide technical expertise around scientific applications, programming models and systems, as well as early access to forthcoming 2016 platforms and the latest innovative technologies (NVIDIA NVLink, IBM CAPI), the IBM high performance computing (HPC) software stack and the NVIDIA Tesla Accelerated Computing Platform.

"The work we are doing with GENCI -- bringing together some of the best minds in science and information technology -- is a collaborative effort on a grand scale involving not just GENCI and IBM, but thousands of developers contributing to the rapidly expanding OpenPOWER ecosystem worldwide," said Michel Teyssedre, CTO of IBM France. "We fully expect our collaborative efforts will produce innovations capable of moving the supercomputing industry that much closer to exascale."

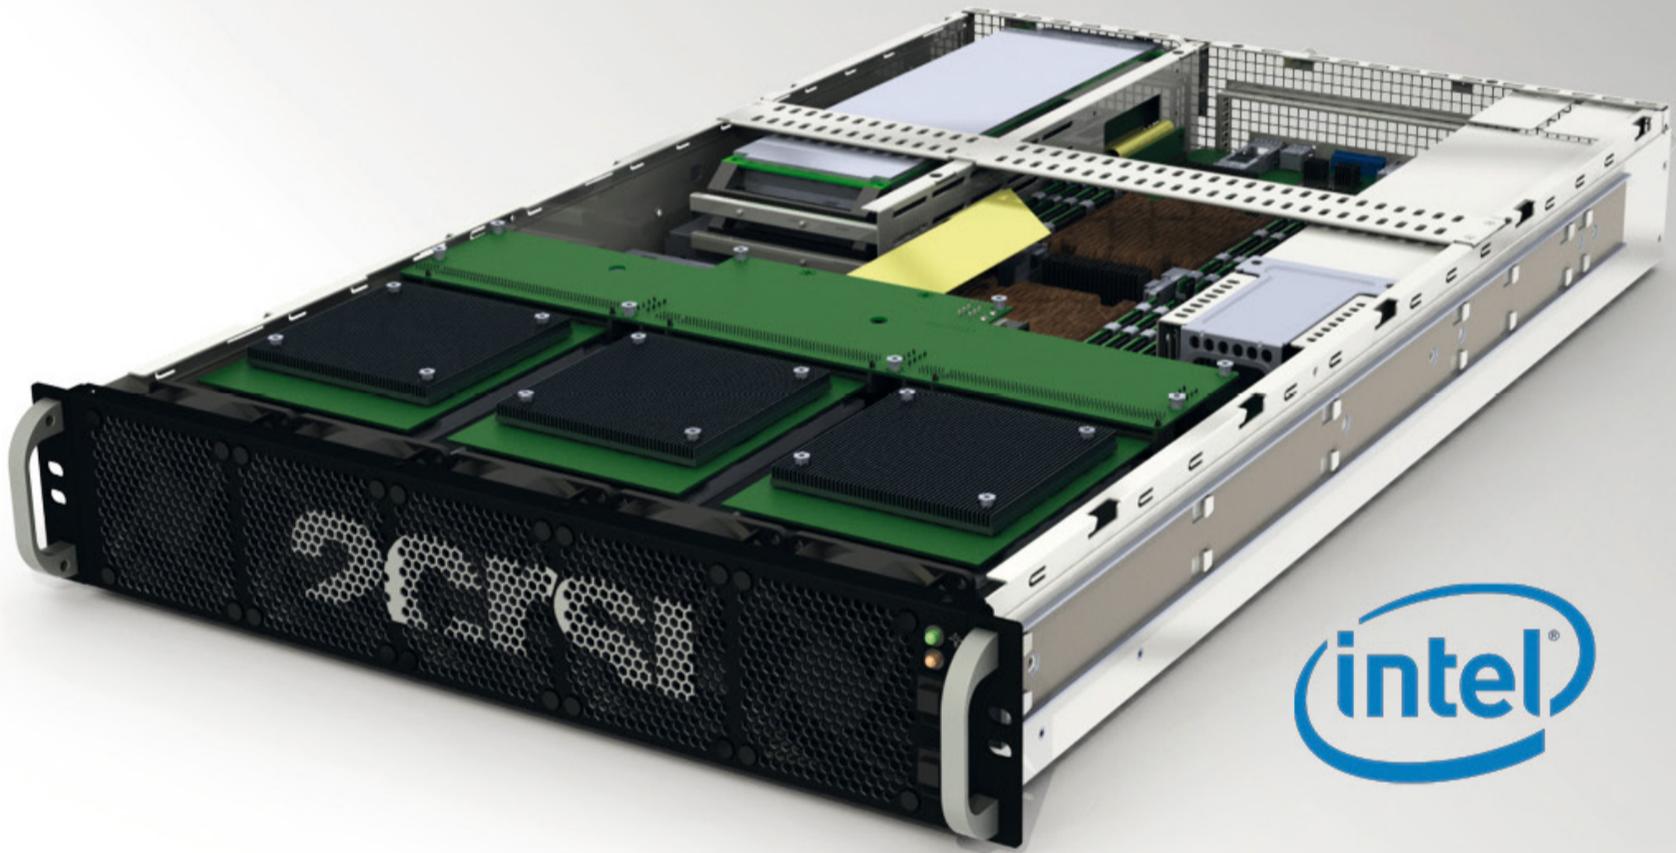
## HEXAPHI

“ *Innovating in computing...* ”



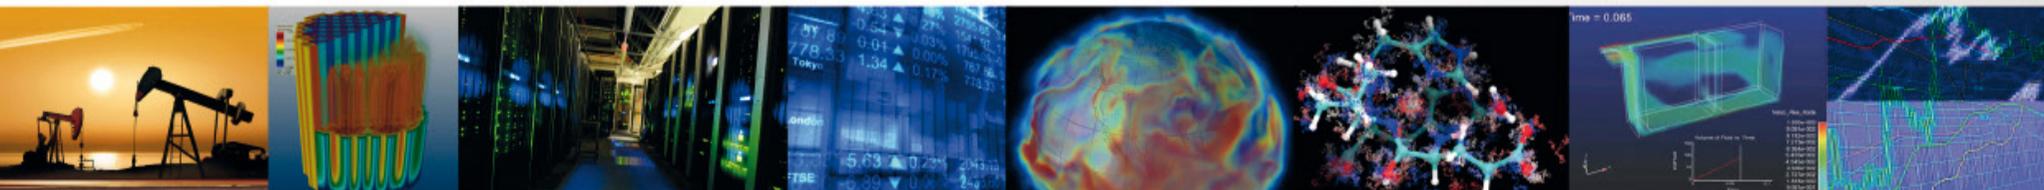
Découvrez la puissance HEXAPHI et sa performance de calcul pouvant atteindre 8 Tflops au sein d'un chassis 2U de 2000W.

*“C'est le meilleur rapport densité/vitesse de calcul jamais obtenu”*



6 Intel® Xeon PHI™ 7120D  
Jusqu'à 1To DDR3 1866Mhz  
Jusqu'à 2 Intel® Xeon® E5-2697V2

Châssis 2U  
8 Tflop /s double précision  
2000 W de consommation



HPC@2crsi.com

[www.2crsi.com](http://www.2crsi.com)





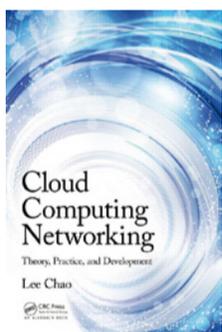
# books

## **CLOUD COMPUTING NETWORKING: THEORY, PRACTICE, AND DEVELOPMENT**

Lee Chao

**Chapman-Hall / CRC Press**

**507 pages, £54.39**



This book covers the key networking and system administration concepts as well as the vital hands-on skills you need to master cloud technology and is designed to help you quickly get started in deploying cloud services for a real-world business. It provides detailed step-by-step instructions for creating a fully functioning cloud-based IT infrastructure using the Microsoft Azure cloud platform. In this environment, you can develop cloud services collaboratively or individually. The book enhances your hands-on skills through numerous lab activities. In these lab activities, you will learn to implement the following services in a cloud environment: Active Directory, DHCP, DNS, and Certificate Services; Configure Windows Server so it can route IP traffic; Im-

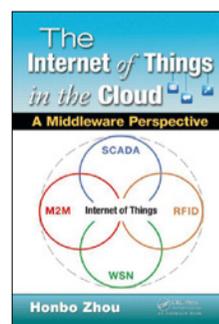
plement IP Security Policy and Windows Firewall with Advanced Security tools; Create a point-to-site connection between Microsoft Azure and a local computer; Create a site-to-site connection between Microsoft Azure and an on-premises network; Develop a hybrid cloud that integrates Microsoft Azure with a private cloud created on a local network

## **THE INTERNET OF THINGS IN THE CLOUD: A MIDDLEWARE PERSPECTIVE**

Honbo Zhou

**Chapman-Hall / CRC Press**

**391 pages; £59.49**



Although the Internet of Things (IoT) is a vast and dynamic territory that is evolving rapidly, there has been a need for a book that offers a holistic view of the technologies and applications of the entire IoT spectrum. Filling this void, *The Internet of Things in the Cloud: A Middleware Perspective* provides a com-



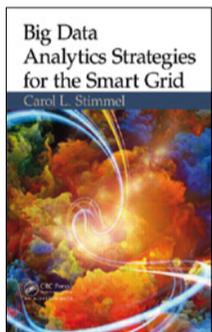
prehensive introduction to the IoT and its development worldwide. It gives you a panoramic view of the IoT landscape—focusing on the overall technological architecture and design of a tentatively unified IoT framework underpinned by Cloud computing from a middleware perspective. Organized into three sections, it describes the many facets of Internet of Things—including the four pillars of IoT and the three layer value chain of IoT; it focuses on middleware, the glue and building blocks of a holistic IoT system on every layer of the architecture, and it explores Cloud computing and IoT as well as their synergy based on the common background of distributed processing.

## **BIG DATA ANALYTICS STRATEGIES FOR THE SMART GRID**

Carol L. Stimmel

**Chapman-Hall / CRC Press**

**256 pages, £54.39**



Readable and accessible, *Big Data Analytics Strategies for the Smart Grid* addresses the needs of applying big data technologies and approaches, including Big Data cybersecurity, to the critical infrastructure that makes up the electrical utility grid. It supplies industry stakeholders with an in-depth understanding of the engineering, business, and customer domains within the power delivery market. This book addresses the need to apply big data technologies and approaches to the critical infrastructure that makes up the electrical utility grid, explains how to implement a data analytics program to meet the challenges of a modern grid that is responsive from an operational perspective and meets the demands of greenhouse gas legislation and describes how to solve collection and storage challenges and how to analyze and act on new forms of information to realize real returns on smart grid investments.

# MOOCS

## **ENGINEERING SOFTWARE AS A SERVICE (SAAS), PART 1**

UC BerkeleyX

Learn software engineering fundamentals using Agile techniques to develop Software as a Service (SaaS) using Ruby on Rails. This intermediate SaaS course uncovers how to code long-lasting software using highly-productive Agile techniques to develop Software as a Service (SaaS) using Ruby on Rails. Learners will understand the new challenges and opportunities of SaaS versus shrink-wrapped software. They will understand and apply fundamental programming techniques to the design, development, testing, and public cloud deployment of a simple SaaS application. Using best-of-breed tools that support modern development techniques including behavior-driven design, user stories, test-driven development, velocity, and pair programming, learners will see how modern programming language features like metaprogramming and reflection can improve productivity and code maintainability. Weekly coding projects and quizzes will be part of the learning experience in this SaaS course. Those who successfully complete the assignments and earn a passing grade can get an honor code certificate or verified certificate from BerkeleyX. The videos and homework assignments used in this offering of the course were revised in October 2013. The new class also includes embedded live chat with Teaching Assistants and other students and opportunities for remote pair programming with other students.

**Starts on** October 5, 2015 **Length** 9 weeks

**Effort** 12 hours/week **Price** FREE - Add a Verified Certificate for \$49 **Subject** Computer Science

**Level** Intermediate **Languages** English

**Video Transcripts** English **Link** [www.edx.org/course/engineering-software-service-saas-part-1-uc-berkeleyx-cs169-1x#!](http://www.edx.org/course/engineering-software-service-saas-part-1-uc-berkeleyx-cs169-1x#!)



## DATA SCIENCE AND MACHINE LEARNING ESSENTIALS

Microsoft

Learn key concepts of data science and machine learning with examples on how to build a cloud data science solution with R, Python and Azure Machine Learning from the Cortana Analytics Suite. Demand for Data science talent is exploding. Learn these essentials with experts from M.I.T and the industry, partnering with Microsoft to help develop your career as a data scientist. By the end of this course, you will know how to build and derive insights from data science and machine learning models.

You will learn key concepts in data acquisition, preparation, exploration and visualization along with examples on how to build a cloud data science solution using Azure Machine Learning, R & Python. Data Science is an essential skill for analyzing and deriving useful insights from data, big and small. McKinsey estimates that by 2018, a 500,000 strong workforce of data scientists will be needed in US alone. The resulting talent gap must be filled by a new generation of data scientists.

This course is organized into 5 weekly modules each concluding with a quiz. By achieving a passing grade in the final course assessment you will receive a certificate demonstrating that you have acquired data science skills and knowledge. Apart from answering your questions on the forum, faculty will host an office hour to address questions you may have while undertaking this course. Get an ID verified certificate to demonstrate your data science knowledge and share on Linked-in.

**Starts on** September 24, 2015 **Length** 5 weeks

**Effort** 3 - 4 hours/week **Price** FREE - Add a Verified Certificate for \$50 **Subject** Computer Science

**Level** Intermediate **Languages** English

**Video Transcripts** English **Link** [www.edx.org/course/data-science-machine-learning-essentials-microsoft-dat203x#!](http://www.edx.org/course/data-science-machine-learning-essentials-microsoft-dat203x#!)

## KNOWLEDGE MANAGEMENT AND BIG DATA IN BUSINESS

HKPolyUxyX

Learn why and how knowledge management and Big Data are vital to the new business era. The business landscape is changing so rapidly that traditional management, business and computing courses do not meet the needs for the next generation of workers in the business world. Most traditional methods are of a repetitive, rule-based nature and will be gradually replaced by Artificial Intelligence. In the knowledge era, the most value added job will be to manage knowledge, which includes how knowledge is created, mined, processed, shared and reused in different trades and industry. At the same time, the amount of data and information (prerequisites of knowledge) is exploding exponentially. By 2020, IDC projects that the size of the digital universe will reach 40 zetabytes from all sources including, websites, weblog, sensors, and social media. Big data will transform how we live, work and even think. These trends and more will have a profound effect on how we see the world and create policies. The course is offered by the Knowledge Management and Innovation Research Center (KMIRC) of The Hong Kong Polytechnic University. Most of our research are company and industrial based. Capabilities and competencies of the KMIRC are further strengthened by the international alliances it has formed with leading practitioners. Many of whom are regarded as the «Hall of Fame» of knowledge management, and renowned research centers worldwide. The course is suitable for participants with background in humanities, management, social science, physical science or engineering. No prior technical background is assumed.

**Started on** August 25, 2015 **Length** 6 weeks

**Effort** 4 - 6 hours/week **Price** FREE - Add a Verified Certificate for \$50 **Subject** Business & Management

**Level** Introductory **Languages** English

**Video Transcripts** English **Link** [www.edx.org/course/knowledge-management-big-data-business-hkpolyux-ise101x#!](http://www.edx.org/course/knowledge-management-big-data-business-hkpolyux-ise101x#!)

**CHIFFRES  
CLÉS**

**44** MILLIARDS  
DE  
DOLLARS

Projection du chiffre d'affaires mondial du HPC en 2020.

**8,3%**

Croissance annuelle du secteur du calcul intensif.

**220** MILLIARDS  
DE  
DOLLARS

Chiffre d'affaires cumulé sur la période 2015-2020.

Source :



**LE TOP 3  
DU TOP 500**

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

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

**3 SEQUOIA**  
Lawrence Livermore National Laboratory,  
Etats-Unis : **17173 / 20133 TFlops**  
Constructeur IBM Blue Gene/Q Architecture  
PowerPC A2

Le TOP500 classe tous les six mois les 500 superordinateurs les plus puissants au monde. Les deux valeurs retenues, RMAX et RPEAK, représentent la puissance de calcul Linpack maximale et théorique.

**LE TOP 3  
DU GREEN 500**

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

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

**3 6217 MFLOPS/W**  
Suien High Energy Accelerator Research  
Organization / KEK (Japon)

Le classement Green 500 liste les supercalculateurs les plus éconergétiques au monde. L'efficacité énergétique est appréciée par la mesure des performances par Watt. L'unité est ici le MFlops/Watt.



# AMD FirePro™ S9150 Server GPU

## The GPU of choice for high-performance computing

The most compute-intensive workloads in data analytics or scientific computing are no challenge for the AMD FirePro™ S9150 server GPU. With support for OpenCL™ 2.0, 16 GB GDDR5 memory, and up to 2.53 TFLOPS of peak double-precision and up to 10.8 GFLOPS-per-watt peak double-precision performance, the choice is clear.



### The AMD FirePro™ S9150 features:

- Up to 5.07 TFLOPS peak single-precision floating point performance
- Up to 2.53 TFLOPS peak double-precision floating point performance
- 16 GB of GDDR5 memory
- Up to 320 GB/s memory bandwidth
- AMD Graphics Core Next (GCN) Architecture
- Full Rate Double Precision
- AMD STREAM Technology
- AMD PowerTune Technology
- OpenCL™ 2.0 support

THE **GREEN**  
500 

#1 on the November, 2014 Green500 list<sup>1</sup>

Please visit [www.fireprographics.com/s-series](http://www.fireprographics.com/s-series) to find out where you can get the AMD FirePro S9150.

1. AMD FirePro™ S9150 server GPU powers the #1 supercomputer on the November, 2014 Green500 list. For more details, please visit <http://www.green500.org/news/green500-list-november-2014>

2. AMD FirePro™ S9150 max power is 235W and delivers up to 2.53 TFLOPS peak double and up to 5.07 peak single precision floating point performance. Nvidia's highest performing single-GPU server card in the market as of March 2015 is the Tesla K40, max power of 235W, with up to 1.43 TFLOPS peak double and up to 4.29 peak single-precision compute performance. FP-97

© Copyright 2015 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, FirePro, and combinations thereof are trademarks of Advanced Micro Devices, Inc. OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission by Khronos. Other names are for informational purposes only and may be trademarks of their respective owners.



# FLASH STORAGE THE FINAL FRONTIER?





# BREAKING THE **\$1/GB** BARRIER

Flash storage is the hottest buzzword today. Priceless not so long ago, manufacturers now claim that the union of capacitive and performance storage is now a matter of months. But are we talking about the same technologies? Decryption.



# ACCORDING TO SEVERAL STUDIES CONDUCTED BY IDC AND GARTNER, OVER THE NEXT FIVE YEARS, MORE THAN 95% OF IT INVESTMENTS WILL BE DEVOTED TO NETWORK AND STORAGE INFRASTRUCTURE.

It has become impossible nowadays to disregard Flash technology in all IT projects with a storage component. But let's not forget that this term covers several product families, each with their specific and privileged usage scenarios.

According to several studies conducted by IDC and Gartner, over the next five years, more than 95% of IT investments will be devoted to network and storage infrastructure. The latter is changing phase. So far mainly composed of mechanical drives, the data access needs

require much more performance, scalability and agility. Hybrid storage solutions first made their appearance, to balance the best of both worlds: capacitive spaces based on mechanical disks that still represent the best capacity / cost ratio and besides, SSD performance was out of

reach for years. In 2014, the combined sectors of racks full Flash and hybrid represented \$ 11.3 billion worldwide. For comparison, the turnover in Western Europe servers accounted for \$ 3.7 billion in 2014. With an estimated growth of this market in double digits over the next five years.

## IS ALL FLASH STORAGE AT HAND?

Since last year, several actors have generalized full flash storage arrays in their product portfolios. The aim is to improve performance in critical applications for the enterprise. Websites, data and big data processing bases and Internet of Things are the main reasons companies invest today. These solutions provide at least ten times higher performance compared to conventional disk drives. It is also the most innovative segment. Leveraging its specificities in terms of speeds and latency, manufac-

# LEGEND HAS IT THAT TO WIN TEN TIMES PERFORMANCE, THE FLASH STORAGE COSTS TEN TIMES AS MUCH.



turers add mechanisms to increase efficiency such as thin provisioning, compression, deduplication, snapshots, cloning and replication. With the stated goal to balance the difference in cost and capacity compared to traditional storage devices.

## A GAP SMALLER THAN IT APPEARS

Legend has it that to win ten times performance, the Flash storage costs ten times as much. This is likely if you only take cost into consideration. To effectively compare the two approaches, it should be taken into account all the costs during the life cycle of each technology through their TCO. And by doing so, the gap narrows significantly. The reason is simple: conventional disks have a considerably higher energy consumption, and must be constantly cooled further to retain their operation temperature. It is moreover not the only parameter to be considered. The space per square meter in a datacenter is expensive, very expensive. Over a period of three years, taking into account all the above parameters, this results in a lower TCO estimated by IDC 80% in favor of Flash storage!

## THE NEW ACTORS MAKE THE MOST AFFORDABLE FLASH STORAGE

But then, the entry ticket to acquire full Flash berries was until recently, prohibitive. The investment could reach several hundred thousand euros to have a latest generation array. This emerging market with longtime established actors like IBM, EMC or Nutanix and which had its own dynamic, is already being pushed by new entrants like Atlantis Computing and others. This company just arrived in Europe, and displays clear ambitions: to offer hyperconverged servers and appliances incorporating Flash storage at the lowest cost. Its new full flash HyperScale appliances 12TB boast reduced costs of 50-90% in comparison with appliances such as the hybrid NX-3461 (308 000 \$) and full flash NX-9240 (800 000 \$) of Nutanix. In comparison, Atlantis Computing has put a very aggressive price tag on its HyperFlash appliance: 78,000 dollars. This amounts to 6.5 dollars per gigabyte, against 25 and \$ 66 / gigabyte with Nutanix' models. A significant difference which, without tipping the market the next day, still promises to generate some questioning among users doing their market studies before acquiring Flash storage arrays.

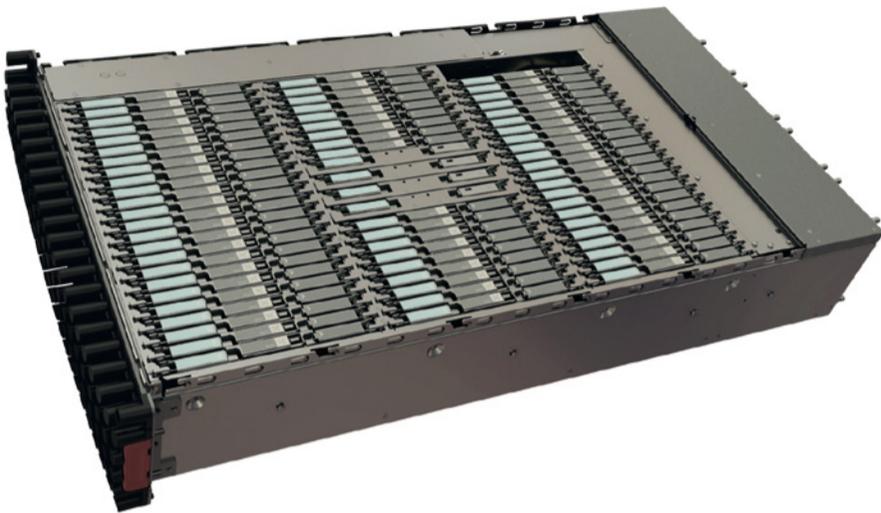
## THE LITTLE SECRET OF THE ACTORS OF FLASH STORAGE

It must be said that Atlantis Computing, like almost all the players in this sector, have a little





# HOW DO MANUFACTURERS TO SELL FEWER THAN FLASH STORAGE THAN THEIR CLAIMS? THEY ACTUALLY RELY ON A STRING OF MECHANISMS WHICH THE STORAGE WORLD HAS THE SECRET.



secret about which they communicate more or less widely: the «actual» capacity. Remember this word because it represents the difference between the displayed and the actual, useable capacity. By acquiring a 12 TB appliance if you open however often you count and recount, you will not find 12 TB of flash memory, but between 30 and 50% less, i.e. 6-8 actual gigabytes, for which the dedicated qualifier is «raw». An appliance marketed as having 12 TB will in fact only have a portion of what a company believes to acquire. This is not false advertising but merely different presentation. Not a problem per se, provided the customer is aware and informed.

## OPTIMIZATION MECHANISMS

How do manufacturers to sell fewer than flash storage than their claims? They actually rely on a string of mechanisms which the storage world has the secret: compression and thin provisioning are the keywords. In other words, they apply a compression ratio to sell 6-8 Gigabytes real, marketed as being roughly equivalent to 12GB of uncompressed data. It is a rule that it is better to know what you buy ... and it does relativize the sometimes outlandish claims of scientific calculations on

TCO, which are based on irrefutable technical properties, but whose calculation parameters may not be totally comparable across different manufacturer's models.

## STORTRENDS, ANOTHER TROUBLEMAKER - OBJECTIVE: 50 CENTS GIGABYTE

As we see, the race to go under the symbolic \$ 1 per gigabyte goes full swing. And now another actor in ambush, the US giant American Megatrends, bluntly announced an AFA (All Flash Appliance) to aim this time, 50 cents per gigabyte! Its model StorTrends 3500i offers 56 TB of storage expandable to 256 GB. This model has the distinction of being marketed in hybrid version at entry level, and can be upgraded to full flash later on. Aimed at small and medium enterprises, it uses conventional disk drives and flash SSD caching, and an automatic tiering mechanism takes care to offer the best performance for the most requested data are deported on the SSDs. And resulting with latencies of around 3 milliseconds against 10 ms for conventional discs. Switching to full-flash gets the latency down to 1 millisecond.

## INCUMBENTS CONS ATTACK

HP is no exception, and intends to be present in this new emerging market. Evidenced by the recent announcement of its StorVirtual 4335, a hybrid appliance that displays 12 times the performance with a 90% energy reduction to 12.4 Terabytes capacity for 59,000 dollars. At the





# THIS MUTATION IS EXPLAINED IN PART BY THE EMERGENCE OF CLOUD COMPUTING IN THE ENTERPRISE.

other end of the spectrum is the MSA 2040 SAN 2U storage array, capable of hosting 200, 400, 800 and 1.6 TB SSDs each for a total of 38.4 TB raw storage capacity, and sold 25% cheaper than its predecessors. The starting price of each SSD is \$ 1,599. The MSA 2040 SAN also has sophisticated properties tailored for integration with an enterprise infrastructure with features like snapshots, replication, copy volume and self-encryption of volumes and disks.

## HIGHER, STRONGER, FURTHER: FLASH FOR BIG DATA

A couple months ago, Sandisk, one of Flash storage historical players, has made an announcement that caused a stir by launching its InfiniFlash storage server, intended for Big Data markets. Admittedly, the dimensioning of the InfiniFlash is related to the ambitious target of the manufacturer. Available in three versions (IF500, IF700 and IF900) this 3U chassis is packed with Flash cards - up to 64 - each with 8 TB of Flash for a total of 512 TB.

## DATA CENTERS ALSO TAKE FLASH STORAGE INTO ACCOUNT

As emphasized by Pascal Cheyroux from Sandisk, flash technology is about to establish itself as an integral part of daily operations performed in large datacenters. Many companies seek to bridge the gap between faster processors and storage performance offered by traditional hard drives. In order to accomplish this, they adopt flash technology, which brings speed and performance to the infrastructure. This development is also highlighted by IDC, which estimates that last year, at least half of IT services of large companies have deployed servers and storage arrays with flash technology to manage their workloads. IDC also states that 80% of storage devices shipped

in 2015 will be compatible with flash technology, which is a sign we are entering the era of Flash-Transformed Data Center (FTDC).

This mutation is explained in part by the emergence of cloud computing in the enterprise, along with the increased pressure on cloud service providers that must offer solid guarantees on application performance predictability. These trends, plus Big Data analysis, social media and mobility are all challenges for CIOs who must keep pace with high demand and heavy workloads.

It is the combination of these challenges announcing the evolution of datacenters towards what we call the «flash-Transformed data center», where each floor is rapidly embracing flash technology. But what is the real significance of this transformation? What is the impact for companies whose business depends on the cloud?

## FLASH AGILITY ALLOWS FOR MORE COMPLEXITY

First, flash technology deployed in servers, storage and network infrastructure appliances provide faster processing, which results in the ability to process more complex analysis, a higher level of security and greater respect for the environment. Flash technology accelerates application performance in the cloud and can process large volumes of data without being penalized by the slowdown in I/O operations related to storage bottlenecks. It also supports the analysis of Big Data that identifies «models» within the data and produces executable data for the company.

This is why cloud service providers (CSPs) that offer a range of professional services in the form of «software as a service» (SaaS - Software as a Service) adopt flash technology to attain speed and performance. This approach translates directly into added value related to the provision of more efficient and faster business services. This



# FLASH TECHNOLOGY ALSO IMPROVES CORPORATE SECURITY BY INCREASING THE SPEED AT WHICH THE DATA IS ANALYZED.

trend will increase, since a growing number of data center hosts a growing volume of business workloads.

## SAFETY AND SPEED, THE TWO PRONGS OF TOMORROW'S DATA CENTER

Secondly, Flash technology also improves corporate security by increasing the speed at which the data is analyzed. In our information age, threats to the business assets are mainly digital, and various applications and monitoring systems designed to fight against these threats share a common challenge: a constantly increasing volume of data to be analyzed and an increasingly narrow reaction window. In addition, flash technology reduces the analysis times and speeds up automated responses.

## FLASH CHAMPIONS IN ENERGY CONSUMPTION

Finally, the energy consumption is another good point for Flash technology. According to Michael Bell, vice president of Gartner, «more than 50% of data centers consume more than 6 kilowatts per rack in two years and this number will rise to 70% and then to 80% in four years due to increasing the density of computer equipment.» Facing such expenses as Michael Bell calls them, is «fundamentally unsustainable», it is however possible to reverse the trend of ever increasing infrastructure costs by deploying flash technology in data centers. Flash technology reduces up to 90% energy consumption of data centers with equivalent application performance while supplying the data to a sustainably higher rate. It avoids the use of energy-hungry IT equipment, thus reducing the energy requirements of the data center while maintaining or even increasing, application performance, and lowering their huge operating costs.

For these reasons, Flash technology is a driving force for transforming the data center in order to respond to the next wave of computing requirements by speeding the performance of enterprise applications and databases, while making these data more quickly available to businesses.

It is for these reasons that Flash technology will play a decisive role in the computing landscape of tomorrow, as it continues to establish itself in servers and storage environments. As such, it will have a beneficial impact on all major categories of enterprise and cloud solutions, while following the pace of change in the world of business today.

## MARKET RESEARCH PERFORMANCE

We see, and no offense to his detractors, flash storage as having a very bright future. If the vast majority of storage spaces do not require millisecond low-latency offered by flash technology, some sectors will earn huge benefits. This is the case of content repositories, streaming media and services, large volumes of data for real-time analysis and public high-traffic web infrastructure. Another example cited by Pascal Cheyroux, Sales Director Southern Europe Sandisk regarding the video surveillance market. No conventional video surveillance, but rather airports or whole sea ports. Using dozens of 4K cameras, made necessary by the extent of the area and to be able to zoom in on an incident where applicable, means reviewing the rest of the infrastructure, he insists. The reason is the data access time. If the time needed for video consultation exceeds that of the automatic purge - and we talk about periods of six months - the result is counterproductive. Flash infrastructure allows bringing that period down to a mere few days.

---

**JOSCELYN FLORES**



# 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



# 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**

# DELL PRECISION WORKSTATION M3800



**T**he latest version of Dell Precision M3800 is a mobile workstation designed for the most demanding tasks, particularly in the fields of architecture, graphics, scientific calculations and tasks engineering. It revolves around a powerful Intel Core i7 processor and Nvidia Quadro professional GPU K1100M but its advantages do not stop there. A very high resolution 4K screen and a Thunderbolt port 2 are two included compared to last year's model. It is therefore

a decidedly upgraded mobile workstation, one of the most practical and luxurious solutions if you need a large work space in a small and lightweight case.

## DESIGN AND FEATURES

The Precision M3800 is an ultraslim mobile workstation. Measuring 18 x 374 x 254 mm, this machine weighs 1.8 kg, which is much lighter than the 17-inch Dell Precision M2800 or the 15.6-inch Lenovo ThinkPad W540. The chassis uses top quality materials, particularly a carbon fiber coating on the bottom. The 15.6 inch IGZO screen iridium zinc gallium oxide (IGZO) is protected by an edge to edge Corning Gorilla Glass coating. It is also touch capable on ten contact points. Ultra



HD resolution (3840x2160 pixels) offers a wider workspace than the Full HD screens at 1080p resolution. The display surface even supersedes the definition of 2880x1800 pixels of the Apple MacBook Pro 15 inch 2014 or even the 3K screen (2880x1620 pixels) of the Lenovo W540.

This large workspace is also a plus if you need to work on architectural drawings, scientific diagrams, huge spreadsheets and multiple HD videos in 1080p. The display quality is excellent, with fine gradation and beautiful colors. The machine, however, does not have an integrated calibration tool like the Lenovo W540 and the Toshiba Satellite P50T-BST2N01 - a sometimes necessary module if you need to calibrate the colors for print or editing videos.

Connectivity is excellent. Besides a Kensington safety lock, the M3800 features an SD card reader, a USB 2.0 port and a USB 3.0 port on the right machine. The USB ports are the same color, so you should watch closely to identify one that obeys the faster USB 3.0. On the left, are the headphone jack, an HDMI port, a Thunderbolt port 2 and a second USB 3.0 port. Good point, all three USB ports are able to charge

devices while on sleep mode. Dell includes a USB to Ethernet adapter, which is handy for enjoying a wired connection to access the network. The Thunderbolt port 2 also serves as a mini-DisplayPort connector, but chances are you will use it primarily to share data and devices with Mac or Windows workstations through network drives. As for wireless connections, the machine is compatible with 802.11ac signals and Bluetooth 4.0.

### **LIMITED ACCESS TO INTERNAL COMPONENTS**

Access to internal components is disappointing, however. Like the MacBook Pro, the Precision M3800 uses non-standard screws to lock the bottom, making it impossible to replace the RAM, storage or even the battery. Your IT department technicians should be able to access these items after purchasing the right tools, but you'd better leverage Dell's technical support and on-site support after remote diagnosis. The machine comes with a one year warranty, which is relatively short for a machine of this caliber.

## **THE BACKLIT KEYBOARD IS VERY COMFORTABLE, EVEN IF IT HAS NO NUMERIC KEYPAD.**

The backlit keyboard is very comfortable, even if it has no numeric keypad. The large touchpad is centered on the space bar, and is as responsive as the screen. The system integrates Windows 8.1 Pro with a USB key instead of the traditional recovery partition on the 256GB SSD. You can return to Windows 7 Professional on the same Microsoft licensing if your company has not standardized its procedures under Windows 8.1. The 16 GB RAM are sufficient to process large volumes of data and RAW images of several gigabytes.



### **PERFORMANCE**

The Precision M3800 features a quad-core Intel Core i7-4712HQ processor and Nvidia Quadro K1100M graphic processor. It gets 2664 points at PC Mark 8's Work Conventional and thus shows more than sufficient for all daily tasks. It took a minute and 19 seconds to complete the Handbrake video encoding test and 3min30 for the test with Photoshop CS6. Both results are very similar to other mobile workstations upscale, like the Apple MacBook Pro (1min17 on Handbrake, on 3min25 CS6) and the Lenovo W540 (1min20 on Handbrake, on 3min18 CS6). The machine gets 599 points on CineBench, which puts it behind the Toshiba P50T-BST2N01 (611 points) and the Lenovo W540 (637).

If you do not play probably video games on the M3800, its scores in 3D games match those of a professional workstation. Thus displays 22 frames per second (fps) to 29 fps in Heaven and Valley, both in medium level of detail. This is slightly better than the integrated graphics of the Apple MacBook Pro 15-inch Retina (22 fps in Heaven, 23 fps on Valley) although the Lenovo W540 (30 fps in Heaven, 42 fps on Valley) is faster at this level. In all cases, the M3800 has ample power to perform any type of work on the go.

### **BATTERY PERFORMANCE COULD BE BETTER**

Autonomy is more mixed. In our testing, the machine lasted 4 hours and 53 minutes. It is an hour less than the Lenovo W540 and the previous version of the M3800. The winner in this field is Apple's MacBook Pro, which lasts nearly nine hours on a full charge. The autonomy of the M3800 remains higher than the Dell Precision M2800 (4:26) and Toshiba P50T-BST2N01 (2:45).

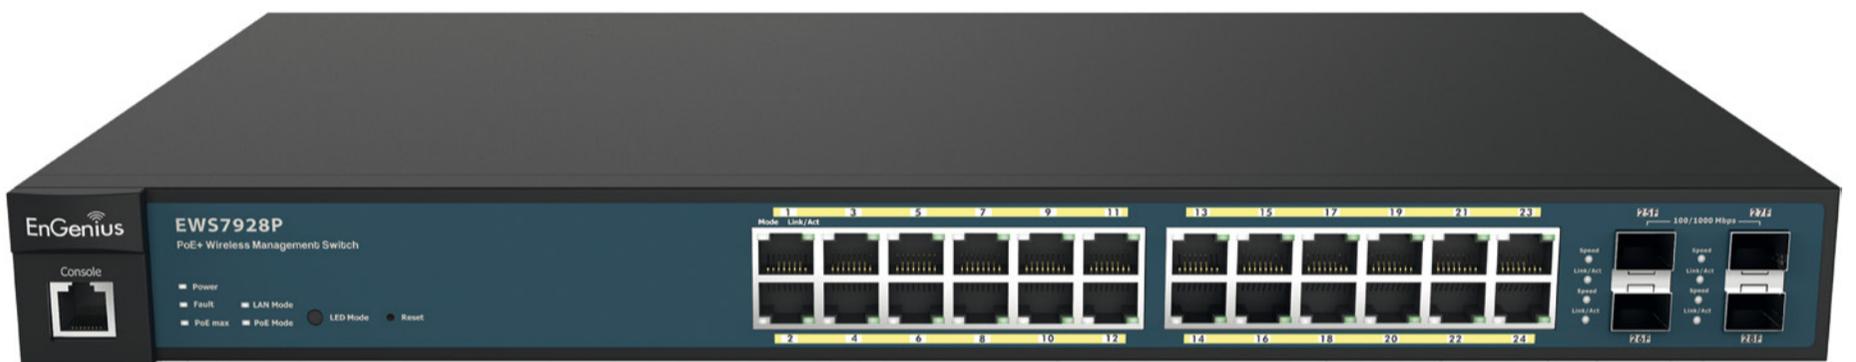
### **CONCLUSION**

Overall, the latest edition of Dell Precision M3800 is a great success in terms of connectivity, the quality / price ratio and the screen resolution. Its battery duration is less than that of last year's model, but this is hardly surprising given the highest display resolution. If you absolutely need the workspace of a 4K screen, and if you are willing to keep a power supply handy, the new M3800 is an excellent choice. The Lenovo ThinkPad W540 retains our preference in the field of mobile workstations, particularly through its integrated calibration tool, its greater autonomy (and its removable battery), and better access to its internal components. Ramon Lafleur [TEST]



# ENGENIUS NEUTRON EWS7928P

The EnGenius Neutron Series EWS7928P is a Managed Wireless solution 24-Port Gigabit PoE+ Layer 2 Switch with 4 Dual-Speed SFP ports ideal for Wireless Access Points, VoIP (Voice-over-IP) phones, and IP surveillance applications.



**T**he EWS7928P through its browser-based GUI (Graphical User Interface) can discover, configure, manage and monitor up to 50 wireless access points. The controller mode of the Switch provides a suite of wireless management features for IT managers. Features include clustering access points in clusters with identical policies, access point IP auto-discovery, remote access point rebooting, and visual monitoring features including topology and floor plan view modes for the mapping out and tracking of networks.

The controller also features a statistics section for record traffic tracking and comparison between units or over-time comparisons.

The interface also lets IT managers apply configuration changes to multiple Access Points or the entire network simultaneously. A dashboard provides instant access to a variety of clients, network information, and events, while a real time map displays Access Point locations and network topology for planning and deployment. Managing the wired portion of the network is also easy using the same browser-based interface. The Switch supports IEEE 802.1p QoS which automatically classifies and prioritizes compliant device traffic to ensure that time sensitive data such as VoIP is forwarded without delay.



# THIS SWITCH OFFERS CENTRALIZED MANAGEMENT, AND QUICK SETUP THROUGH A WEB-BASED USER INTERFACE WHICH SUPPORTS UP TO 50 NEUTRON SERIES ACCESS POINTS.

The EWS7829P offers 24 Gigabit ports with IEEE 802.3at/af (PoE+) support and 4 SFP slots for fiber uplinks. The 24 Gigabit ports provide seamless connectivity for devices and applications that require high speed network access and reduce bottlenecks that can interrupt communication. The Switch delivers up to 30 watts per port over connected Ethernet cables to power essential devices like Wireless Access Points, IP Cameras, and VoIP (Voice-over-IP) Phone systems, making the EWS7928P a great choice for businesses that need to connect to a wide variety of devices for their company. Four SFP slots accept fiber transceivers for extending the wired network beyond the limitation of Ethernet cabling. This is especially useful for schools with satellite classrooms, or for hotels and resorts, or other businesses with multiple building properties.

## CENTRALIZED MANAGEMENT

This switch offers centralized management, and quick setup through a web-based user interface which supports up to 50 Neutron Series Access Points. IT managers can keep track of, map out, and plan the placement of Neutron Series Access Points and Switches in the network with a visual view of the topology as well as tools to search for and find Neutron Series Access Points via their IP or MAC address. The software also gives IT managers and network administrators the visibility they need to monitor, manage and quickly adjust the settings or performance of their network in real time. The feature rich inter-

face displays usage reports for real time and historical client connectivity to each Neutron Series Access Point as well as traffic flow and load over both wired and wireless portions of the network. PoE+ IEEE 802.3at/af Support enables PoE devices like Access Points, Client Bridges, IP Cameras, VoIP telephones and others to be powered directly from the Switch over an Ethernet cable. Up to 30 Watts per port with a total PoE budget of 185 Watts

## 4 SFP PORTS FOR EXTENSION

The EWS7928P includes 4 SFP ports to support transceivers for fiber cabling for connectivity to offices, buildings, or other parts of the network that exceed the distance capability of regular Ethernet. It has a full-Featured Layer 2 Manageability: VLAN, QoS, IGMP/MLD Snooping, STP/RSTP/MSTP, Link Aggregation Control Protocol (LACP), SNMP v1/v2c/v3, and more.

## MESH MODE

Under the AP Mesh mode which is slated to be available in the coming months, the Neutron Series Access Points can be used as the central connection hub for station or clients that support IEEE 802.11 a/b/g/n network. Under this mode, the Neutron Series APs can be configured with the same Mesh SSID and security password in order to associate with other Neutron Series APs. For example, you would use one band to connect Neutron Series Access Points in range with Mesh mode and the other band to broadcast traffic on the network.

# THE USER INTERFACE INCLUDES TWO EASY-TO-USE, DRAG AND DROP TOOLS TO VIEW A NEUTRON WIRELESS MANAGEMENT DEPLOYMENT.

Acting as a node within a web framework, each Neutron Series Access Point only needs to connect to the nearest node using the best path to transmit data, working collaboratively with other Access Points in the network infrastructure to function.

## FLOOR PLAN & MAP VIEWS

The user interface includes two easy-to-use, drag and drop tools to view a Neutron Wireless Management deployment. The Map View lets IT managers drag and drop a marker representing an Access Point that has been registered to a Neutron Switch onto a building within a campus topology to show the relative location of the specific Access Point. This Map View visual reference makes it easy to find the Access Point to monitor or reconfigure as necessary if the needs of the network change over time. Like the Map View, the Floor Plan View does much the same thing but now at the floor plan level. Scanned images of office or facility floor plans can be uploaded to the Switch interface so, that IT managers can drag and drop Access Point markers to their approximate locations. This tool also helps in the planning for additional Neutron Series Access Points and other related networked devices like IP Cameras.

## ROOM FOR GROWTH

The EWS7829P can support up to 50 Access Points. Each switch can connect to another via Gigabit Ethernet or fiber uplinks through their SFP ports. This enables network administrators to rapidly build out large wireless

device deployments over expansive properties with the added assurance that through the AP management User Interface, that they have total visibility of the deployed Neutron Series APs and Switches and can monitor and manage their performance, upgrade their firmware or make operation mode changes or even security changes like selective SSID-to-VLAN tagging as needed.

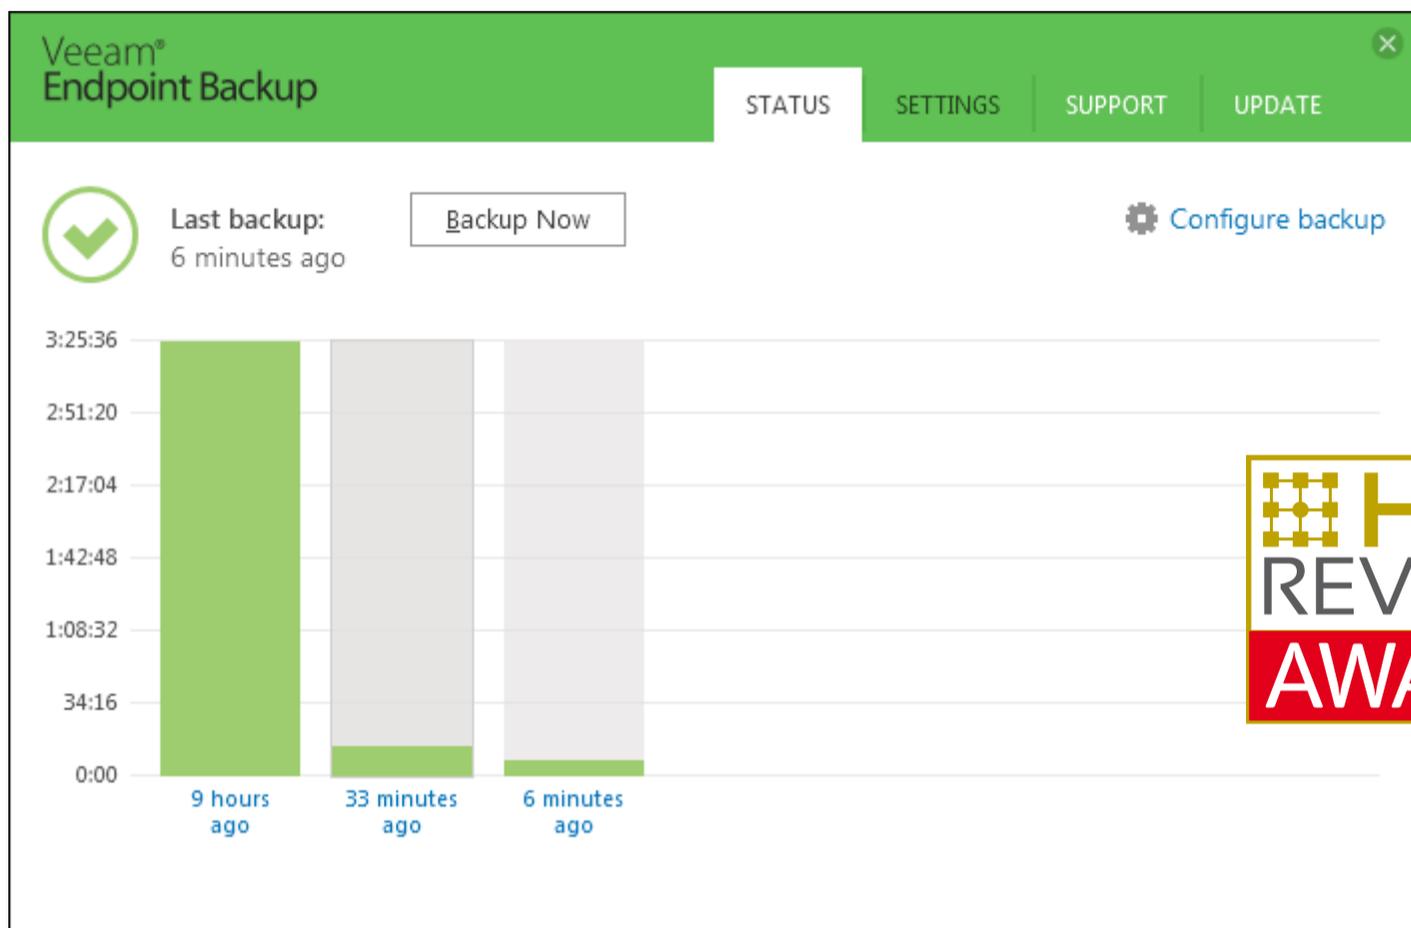
## COMPREHENSIVE SECURITY

The Wireless Management solution supports robust security features such as SSL Certificate. In addition, the Neutron series Switches also support a complete lineup of advanced Layer 2 features; including secure control connections between Switches and Access Points, Port mirroring, STP/RSTP/MSTP, Link Aggregation Control Protocol (LACP), SNMP v1/v2/v3, RMON, and ACL for extensive network security.

## WIRED AND WIRELESS NETWORK MANAGEMENT AND REPORTING

In addition to the network management features, the networking interface gives IT managers and network administrators the visibility they need to monitor, manage and quickly adjust the settings or performance of their network in real time. The interface displays usage reports for real time and historical client connectivity to each Access Point as well as traffic flow and load over both wired and wireless portions of the network.

# VEEAM ENDPOINT BACKUP



U

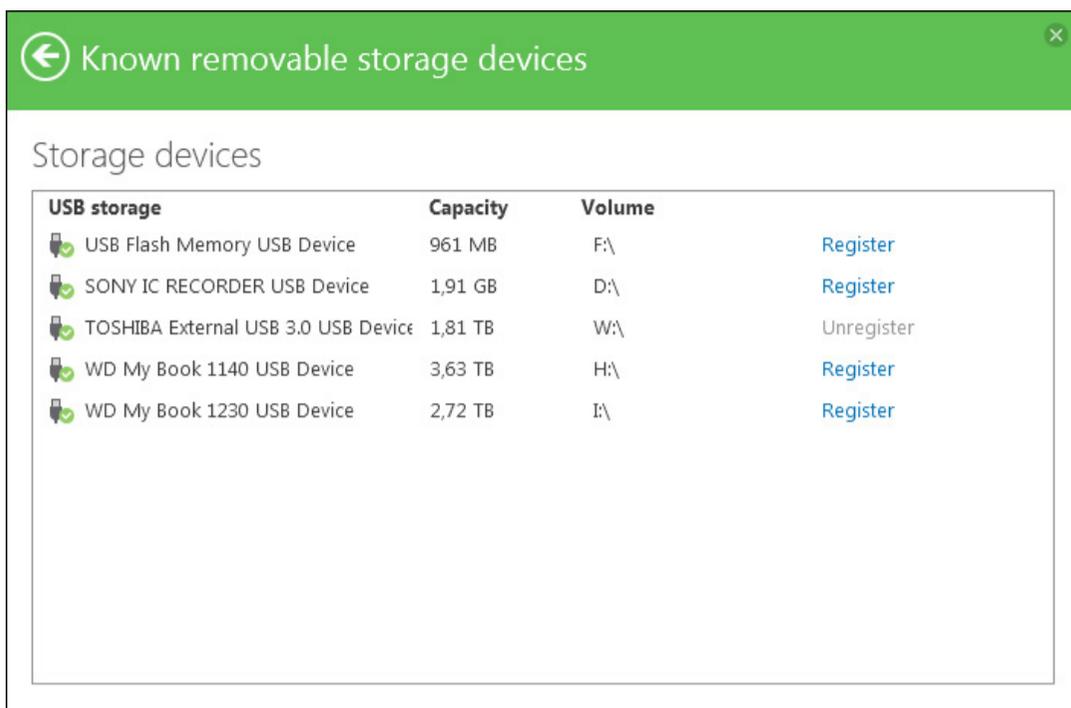
nless you have set up a real backup strategy, the preservation of fixed and mobile endpoints is too often a scarce practice in businesses of all sizes. Therefore, a disk crash or a laptop theft can lead to irreparable loss of crucial data that is impossible or extremely difficult to replace. To provide an answer to this thorny issue, Veeam has

developed Endpoint Backup, a totally free backup software for machines running Windows 7 SP1 or later, but also suited for Windows 2008 R2 SP1 or later servers (32 and 64 bits).

Historically, Veeam has always been confined since its creation in 2006, to virtualized environments (see the interview of CEO Rattimir Timashev in this issue). This incursion of the company's software portfolio addressing physical computers is a double novelty. First by the emergence of a solution out of its favorite area, and secondly because the software is totally free with no temporal or functional limits.

## A LIGHTWEIGHT AND INTUITIVE SOFTWARE FOR IMMEDIATE IMPLEMENTATION

Very light (about 150 MB), the software looks more like a mobile app than a traditional backup software, often synonymous with compli-



external USB drive connected to the PC, which helps to increase data security and also serenity, knowing that this is done automatically upon insertion of the drive!

### INVISIBLE BUT EFFICIENT BACKUPS

The dashboard software shows the status of backups in the form of vertical bars. Clicking on one of them, all details of the backup: origin, duration, space occupied and free disk space on the destination unit. It also displays the performance metrics in MB per second. The first complete backup takes the most time. This does not overload the CPU occupancy rate of the machine, the default setting being to run when the processor is idle. Our backup of a volume of 500 GB was done in just over three hours, with a compression ratio reducing the backup set down to 300 GB. The second backup performed the next day, incremental type was completed in less a quarter hour to 10 GB of files. A third, was made half a day later, and was done in a matter of minutes.

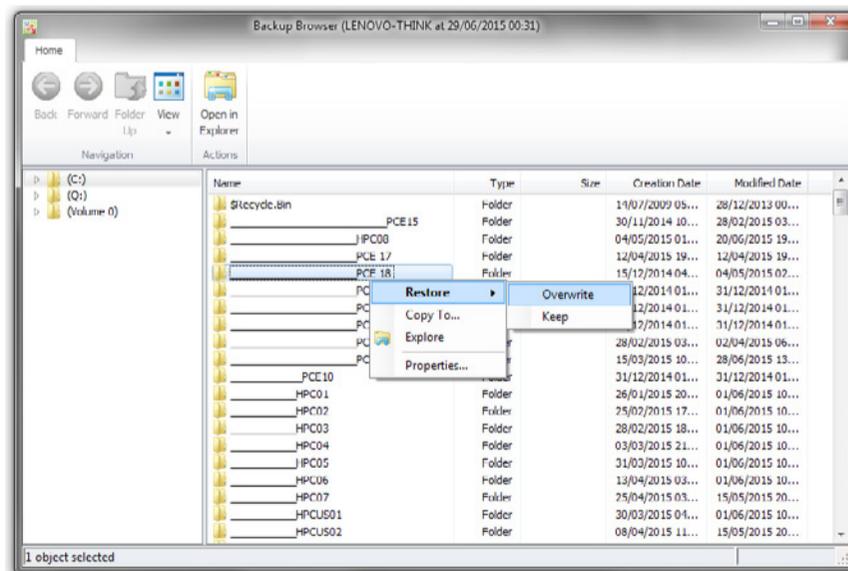
cated options and loaded dashboards. Veeam has opted for an uncluttered and intuitive interface to simplify the programming of the backup operations. At first launch, a wizard guides you through creating your first backup job. The simplicity comes from the automatic management of backup mode. The first backup is always a full backup, followed by incremental backups. Modern, the software performs compression and data deduplication in a transparent way without having to activate them.

### COMPLETE BACKUP OPTIONS

Three types of backup are offered: complete machine with all the disks, one or more volumes to choose from, or one or more files to choose from. The number of backup jobs is not limited, so you can schedule tasks with a very fine granularity. Once the selection is made, you only need to choose the backup destination. Here the choices offered feature local storage location such as a second drive or a USB drive, a network shared folder as a NAS, or a Veeam Backup & Replication repository. You can also set the backup retention period. By default, it is 14 days, but you can increase it if the drive capacity of your backup location permits. Finally, you set the date and time at which backups must be performed. Given the nature of incremental backups, daily default setting is entirely appropriate. An interesting option is to trigger the backup when a given

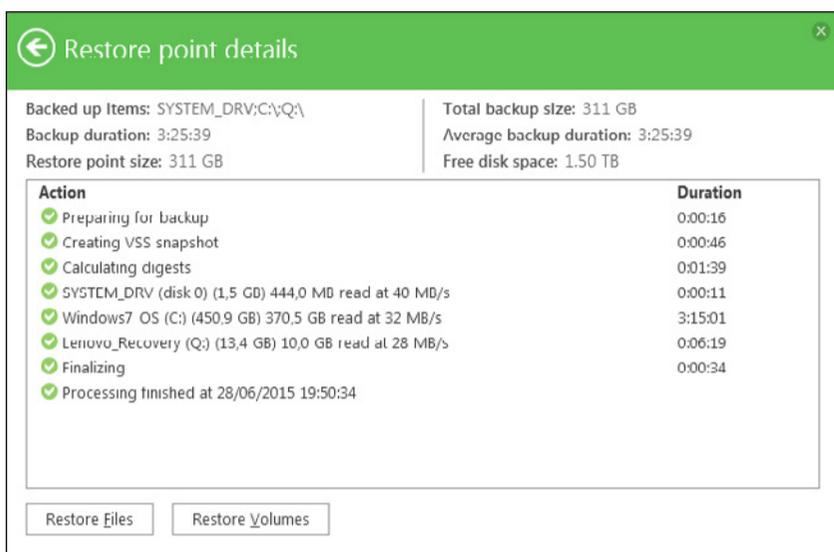
external USB drive connected to the PC, which helps to increase data security and also serenity, knowing that this is done automatically upon insertion of the drive!

### EVEN SIMPLER RESTORE OPERATIONS



Note that at any time it is possible to recover a file or an entire volume. The software displays a browser window to explore the content backups. We have not seen an easier implementation. You can recover a file from any backup

# VEEAM HAS MANAGED TO RESIST THE TEMPTATION TO PLAY THIS CARD AND OFFERS ENOUGH RICH FEATURES TO MEET ALL CASES.



set and supplement or replace the existing file on your system. Note an interesting timeout function that warns you after 30 minutes that without your intervention, the backup Explorer will close to prevent locks on files. The possibility of exploiting repositories of Veeam's big brother Backup & Replication is accompanied by a full range of new possibilities, such as partial restoration operations from Veeam Explorer for Microsoft Active Directory, Exchange, SharePoint and SQL Server. And also the ability to export Endpoint Backup backups to files VMware virtual disks (VMDK) and Microsoft Hyper-V (VHD and VHDX).

## FREE BUT NOT LIMITED

Lightweight and easy is not synonymous with limited functions. Veeam has managed to resist the temptation to play this card and offers enough rich features to meet all cases. If your business requires it, no need to rummage in the sub-menus of options to temporarily disable backups, the second tab of the main interface takes you there directly. You can also disable notifications or force backups even when the computer is in use. This tab lets you access your backup devices identified by

the software, which assigns one of two states: registered or not. Endpoint Backup and keeps track of connected devices, and those that you have set as the backup destination are automatically recorded by the software, which saves you having to find the right drive for your backups. This is very useful if you commonly use multiple external drives, even more so if you use same brands or identical models.

## ADMINISTRATIVE FUNCTIONS DELEGATED BUT ABSENT

Simple to use, the software is however curiously devoid of administrative options with the exception of sending alerts and logs by email. In fact, this aspect is delegated to Veeam Backup & Replication, which provides the functions of supervision and administration for Endpoint Backup. Among these include the ability to move backups off-site on tape or on the cloud, to encrypt backup files located on Veeam repositories, or to assign different access rights and observing traffic network to activate Endpoint Backup backup tasks.

## BOOT AND RESTORE

The software is able to generate bootable media on any removable media (USB key or memory card), CD, DVD or Bluray, or as an ISO image. Once the PC booted from the boot media, the available restore options allow a system to be regenerated from any saved time.

In conclusion, Veeam Backup Endpoint represents a real progress for who wants to ensure their data through regular backups. Despite its free, it is far from being a low level software. Despite being free, it is likely to gain market share at the expense of traditional players in this sector. **JOSCELYN FLORES**

# USING **MICROSOFT** **AZURE** FOR DATA BACKUPS

Instead of relying on external hard disks prone to failure, why not use a powerful cloud infrastructure such as Microsoft Azure? Moreover, it has become reasonably easy to use, judge for yourself.

**T**oday, not a single IT administrator, or even an SMB should initiate an IT project or seek a solution to a computer problem without asking «What can the cloud do for me in this case?» And as long as you're interested so slightly to Microsoft technologies, the question should even be reformulated «What can Azure do for me? « Because Microsoft's Cloud is now so feature packed and accessible, that there are very few situations in which it does not provide a simple and economical solution to your projects and IT issues.

## **SIMPLIFY YOUR IT LIFE**

Among the more traditional challenges that any enterprise meets, backup remains a thorny topic that is suspended like a Damocles sword above your business. A backup infrastructure worthy of the name can quickly cost a fortune. Especially as the costs can skyrocket in terms of maintenance and management.

Instead, the cloud often provides super resilient storage capacity at low costs. Backing up data in the cloud may therefore seem a good idea, although it also raises questions around Internet bandwidth and time of restoration / recovery data from the Internet.

For small and medium sized businesses, a backup in the cloud is undoubtedly a very practical and economical solution to ensure that key data is well protected offsite, very advantageously replacing tape or even disk backups.

## **AZURE BACKUP CONTINUES TO EVOLVE**

Microsoft Azure provides a Backup service since 2013 and Windows Server 2012 R2 includes a cloud backup feature. This service was previously very data center oriented and offered a simple way to backup Windows servers, from versions 2008 to 2012 R2.

Since the end of 2014, access to Azure Backup has been extended to Windows computers and many innovations were introduced last February that justify a fresh look at this IaaS solution.



## A SIMPLE AND TRANSPARENT CLOUD SOLUTION

Azure Backup service can be used in a wide variety of scenarios. Typically the solution is handy for backing up PC Workgroups (thus not joined to a domain) as part of a SMB. It is also an elegant solution to protect servers and imagine disaster recovery scenarios based on a restart entirely in the cloud when an On-Premises Server has passed away. For more ambitious recovery scenarios, Azure also offers an «Azure Site Recovery» service.

The general operation of Azure Backup is simple, secure, reliable and efficient. Simply install on the required machines to backup an agent that adds itself to the Windows backup features. It then provides an incremental backup (of course starting with a full backup) in a predefined Azure Backup space. Data is both encrypted during transfer to the Azure DataCenter but also within the cloud infrastructure. Six copies of the data are maintained through two different DataCenters. Various deduplication and compression mechanisms are implemented in a completely transparent way to optimize backup and cloud occupation. Billing service has evolved since April 2015 and now has two components: firstly the use of the service (charged € 7.5 / month for every 500 GB saved), on the other hand the use of the storage space (roughly € 17 / TB / month). Typically, an SMB can backup all its jobs for less than 20 euros per month, with very strong data retention.

### Implement Azure Backup

#### 1 / CREATE THE CLOUD STORAGE SPACE

**BACKUPS** Start by logging into the Azure administration portal (manage.windowsazure.com) with your Microsoft account login ID. Then select New -> Data services -> Recovery Services -> Backup Archive -> Fast Creation. Enter a name, for example «Post-Windows-Backup» and the region «Northern Europe». Then click on «Create archive». Allow a minute before the space becomes operational.

#### 2 / RETRIEVE THE BACKUP CERTIFICATE

From the same Azure portal, click the previously created backup archive (click on «Post-Windows-Backup»). In the Dashboard tab, click the link «credentials vault» After 5 seconds, a «VaultCredentials» file will be downloaded. This identification file will be requested when configuring the backup agent.

#### 3 / DOWNLOAD THE CORRECT BACKUP

**AGENT** Return to the Azure portal and click your backup archive. Note that the dashboard offers two download links. The first is for Windows Server, System Center DPM and Windows clients (64-bit PC in Windows 7 and 8). The second is specific to Windows Server Essentials. Download the first executable and place it on a USB stick with the «VaultCredentials» backup certificate previously downloaded.

#### 4 / INSTALL THE AGENT ON THE WORKSTATION OR SERVER TO PROTECT

Now go to the machine to back up. Insert the USB drive and run the installer of the agent. Accept the default installation settings. Once the installation is complete, double-click the «Microsoft Azure Backup» icon on the desktop. Click «Add a server». Click Next. In «Safe Identification», click Browse and choose the file «VaultCredentials» recovered in step 2. This phase is crucial as it will directly link the agent to the backup space previously created. Click Next to generate the «passphrase» encryption. You can enter or generate it automatically. The software requires you to define a location to save this passphrase so that you never forget. If you forget, you would indeed be unable to restore the data (which is encrypted from end to end).

#### 5 / CONFIGURE BACKUP

Once registration is complete on the backup server, you will only need to configure and schedule backup options. Click Schedule backup and proceed as usual on Windows: Select the folders to back up, planning (daily or weekly), the retention policy (daily, weekly, monthly, yearly). Azure Backup now supports retention periods of up to 99 years.



## DPM: ANOTHER BACKUP MANAGER IN THE CLOUD

Azure Backup Agent backs up only files and directories. But one may have to back up other elements such as VM or databases. Microsoft System Center Data Protection Manager (DPM) is the tool of choice whenever you need to protect a wide variety of purely Microsoft VM Workloads

starting with Hyper-V, Exchange databases, SQL Server or Sharepoint. Obviously System Center DPM expands its backup mechanisms to encompass your Microsoft Azure cloud backup space via Azure Backup. Backups can be managed directly from the System Center user interface or via

PowerShell commands. Microsoft introduced a new management pack for DPM that provides advanced reporting capabilities with dashboards to effectively monitor your backup infrastructure in the cloud.

**BACKUP ON AZURE HARDLY DIFFERS FROM A TRADITIONAL BACKUP. IT FREES YOU FROM HARDWARE PROBLEMS AND IS AN «OFF-SITE» SOLUTION OF UNPARALLELED SIMPLICITY. IT ALSO ALLOWS YOU TO HAVE QUITE EXCEPTIONAL RETENTION PERIODS (UP TO 99 YEARS).**

6 / **FIRST BACKUP** The last scheduling option allows you to adjust the behavior of the initial backup. By default, it is made directly to the cloud which can take several days depending on your bandwidth and the amount of data to transfer. But Microsoft, for backups of several Terabytes of data has scheduled another option: the possibility of an «Off Line» backup on a SATA disk drive. You can then send this disc by mail (using a procedure described on the Azure website) directly to Microsoft to incorporate the backup without transferring its contents via Internet. Once the primary backup implanted, the backup agent will detect it and activate its differential backup principle to only send the changed items to your Azure cloud backup space.

7 / **CHECK BACKUPS** Remember that you are charged the amount of data backed up each month and the amount of data stored and kept in detention in the storage space. The Azure Backup Dashboard lets you keep an eye on the exact amount of your backups. It is particularly important given that the use of compression and deduplication makes the volume stored somewhat unpredictable (but much lower than a traditional backup).

8 / **RESTORE BACKUP** The backup is evidently only one of the two facets of the problem. The goal of any backup is also to allow the restoration of lost or accidentally modified data. The restoration includes having previously reinstalled the agent if it is to be to a new machine. The agent lets you restore data from backup space to which it was attached or another space (provided you have its certificate). The «Recover data» allows you to specify the amount and date of the backup to be used for recovery. Then select folders or files to restore. Then, specify whether the restore should be at the original location or to another location. The wizard will then provide an estimate of the recovery time.

As you can see, backup on Azure hardly differs from a traditional backup. It frees you from hardware problems and is an «off-site» solution of unparalleled simplicity. It also allows you to have quite exceptional retention periods (up to 99 years) and not worry about the quality of backups and their aging. Finally, it allows to refrain backup windows because the agent proves intelligent and works in the background without slowing down the machine. It is also possible to define bandwidth usage thresholds that raw backups do not impair too much your Internet connection...



**VIEWPOINT**



**PRAKASHAN KORAMBATH**  
INSTITUTE FOR DIGITAL  
RESEARCH AND EDUCATION, UCLA



**HPC**  
REVIEW



# **CYBER SECURITY IN HIGH-PERFORMANCE COMPUTING ENVIRONMENT**

Cyber attack is an unauthorized access to a computer information system or infrastructure with a malicious intent to steal sensitive document, compromise network, vandalize the resources or use the resources for further malicious actions by individuals or organizations or nations.

# THE EXPECTATION IN HPC ENVIRONMENT IS THAT THE RESEARCH DONE ARE MOSTLY OPEN AND THE RESOURCES SHOULD BE EASILY ACCESSIBLE AND THE POLICY SHOULD ACCOMMODATE THE NEED OF RESEARCHERS.

**F**or this discussion, we define a high-performance compute (HPC) environment as a compute resource running Linux operating system (OS) with around at least 500 to 1000 compute nodes having approximately 4000 to 12000 compute cores. They all invariably have high speed, low latency, high bandwidth interconnect network fabric such as Infiniband. They also have attached storage arrays of the order of hundreds of terabytes to petabytes. These kinds of resources are used simultaneously by an average of 200 to 300 users mostly from academic research environments in universities at any time although they may have over

1000 overall users. The complexity of securing the system increases with number of users, as there will be more cases of lost or compromised passwords. Usually in HPC clusters only few login nodes are open to the public network as the compute nodes are in a private network. Security breach involving computer viruses such as Trojan viruses or computer worms usually associated with Windows OS will not be discussed here. We will only discuss proactive mitigating steps to minimize interruptions in the operation of the resource.

The expectation in HPC environment is that the research done are mostly open and the resources should be easily accessible and the policy should accommodate the need of researchers who are collaborating around the globe. There is a need for balance between security and convenience. Because of the convenience

factor intrusion prevention is little bit harder on HPC systems and they are more vulnerable. However, there are a lot of positive benefits in operating an HPC environment in universities compared to the compute environments in financial institutions. Expectation is that typical researchers are not storing any personnel information such as social security number or private medical data on these systems. The biggest worry is that hackers may vandalize the system when they couldn't find any useful data or use this resource to stage criminal activities such as executing a distributed denial of service (DDoS). If the users are involved in any research with private medical data then they are required to do their research on HIPAA complaint compute environment. We will not address how to set up a HIPAA complaint system in this write up as it brings additional complexity of encrypting all the research data. HPC sites typically do not have to worry about attacks such as denial of service as these kinds of attacks are usually against high volume web portals such as news organizations or government web sites.

## **PROTECTING PASSWORDS AND DISABLING UNENCRYPTED NETWORK PROTOCOLS**

In the 1990s research compute environment used protocols such as telnet, ftp where the data between remote computers are communicated in clear text format. So, it was easy for anybody with reasonable expertise to intercept the communication and read the contents. It was easy to listen to an open port and record the keystrokes of users. None of the HPC sites that we know are running these kinds of protocols anymore. The traffic among HPC systems connected through public or private network now is exclusively through encrypted protocols using OpenSSL such as ssh, sftp, https etc. Since almost all HPC resources are running some version of Linux operating system they all invariably run Iptables based firewall at the host level, which is the primary tool to restrict access to service ports from

outside network. Many of them open only few ports such as port 22 for ssh. Iptables also help in operating the system when there are known zero-day vulnerabilities by isolating the resources from outside network.

Typical way the HPC systems compromised is through either users not protecting their password or using passwords that are easy to exploit such as 'test123'. By virtue of the design of the Linux OS, the exploit at the user level is often contained local to a particular user because regular users do not have elevated privileges and they do not have access to files of other users or users from a different group. Even though the security breach through compromised password is usually contained in a user environment, they become escalated in a situation where there is a flow in the Linux kernel itself, which will allow the hacker to trigger local root exploitation and elevate the privileges. In such situation the OS needs to be reinstalled with updated kernels. Linux kernels in the 2000s had frequent kernel flows and were susceptible to memory corruption (buffer overflow), which are becoming very rare these days. Another kind of problem is if the security package itself has flows such as Heartbleed bug (heartbleed.com) in OpenSSL, which was detected in 2014 even though the bug existed for many years prior to that.

In a scenario where users are hacked, often times owners of the accounts are unaware of the fact that they have been compromised. From our experience of running an HPC cluster for the past 12 years it is often the activity of the hackers that expose or alert the system administrators of the system about possible compromise meaning if somebody just login to the system and do nothing their actions are often overlooked. But as soon as the imposter or hacker start using the resources the monitoring tools that are often embedded in Linux OS can record the strange behavior of the system and an alert system administrator can execute remedial actions. Almost always the behaviors of hackers are completely different from that of the owner of the account. Activities such

as sudden burst of network activity, increased network latency, over loading the system with CPU usage, unauthorized jobs bypassing the job scheduler etc. are good indicators of possible compromise. Typically the hackers are exposed in 8 to 10 hours in such scenarios. Often times the affected systems are quarantined from outside network for forensic activities and all the logs are examined to trace the origin of attack such as time, frequency of attack, source host, source port, destination host, destination port and the protocol or application that is used in attacking the system. System will be put back to service after remedial actions are taken such as notifying the appropriate authorities if necessary, upgrading or removing the faulty application or kernel as well as any other upgrades.

Some of the HPC centers do not rely on users in protecting their password. So they implemented what is called One Time Password (OTP) where users are given a small calculator like devices to generate a random key to login to the system. However, that is inconvenient for users, as they have to carry this device all the time with them and also add to the operating cost of HPC systems. Those centers that do not use OTP often limit the failed access attempts to three or four times and restrict access to the account for a period of one or two hours to minimize the brute force attempt to crack user password on the system. Other safeguards HPC centers deploy include running only the necessary applications with administrative privileges that open TCP/IP ports to external network. Even if there is a need to open up Apache port or Database port the connection to these ports are restricted to certain user groups or hosts using Iptables.

The HPC environment also has few privileged account for system administrators and user support staff. These accounts have access to all the resources on the system in the sense that these users can read, write, or delete the contents in any of the accounts. So the account holders of these account needs to take extra caution on all the devices and network they operate. They should access their privileged

account from remote machines or untrusted network only after activating a virtual private network (VPN).

Other forms of attacks are little bit more sophisticated in the sense that hackers do a man-in-the-middle attack by making it appear the hacker controlled system to have similar credentials to the system that users are trying to access in that process tricking the users to give away their credentials. In other word hackers are trying to hijack the endpoints. Experienced hackers also try to cover their trails and continue their exploit by installing rootkits, modifying the RPM package repository, turning off or maneuvering the monitoring software and log files. The sophisticated hackers are usually capable of obfuscating their activities and disabling monitoring tools. In such situations system administrators have to rely on secondary effects such as unusual network bandwidth or unusual CPU load to detect the intrusion. It is possible that a sophisticated attack can remain undetected for months.

Phishing is another common way hackers get user credentials by enticing users to visit hacker controlled machines. Users who use popular social networking web sites are usually susceptible for this kind of attack and if they happen to use same password for their HPC accounts then they are inviting the hackers over from social networking sites to HPC systems. Social networking sites are usually good targets for hackers because they get access to the users contacts and access to their friends and continue their exploit. Intercepted e-mails are another source of compromise.

## Safeguarding the resource

1. **ALWAYS RUN IPTABLES** (firewalls) on the machines with public network connections.
2. **CHECK THE OPERATING SYSTEM LOGS PERIODICALLY** or have it checked by an application and report any anomalies in user behavior such as logins, source host, network bandwidth, resource usage, time and frequency of login etc.



3. **ENFORCE CHOICE OF PASSWORDS** that contain certain number of characters and combination of alphanumeric characters as well as special characters.
4. **FORCE PASSWORD CHANGE ONCE A YEAR** even though most users don't like to change passwords often.
5. **INACTIVATE EXPIRED ACCOUNTS**, unused accounts and accounts of employees when they leave the organization.
6. **TURN ON HOST BASED AUTHENTICATION** if possible. This is almost impossible for HPC systems as the users can be on any host at any time.
7. **ADVISE USERS NOT TO STORE PRIVATE KEYS** such as ssh private keys for passwordless logins on HPC resources because hackers who obtained elevated privilege can use it to access other machines.
8. **ALLOW ONLY AUTHORIZED MACHINES** with known MAC addresses in a network that issue DHCP IP addresses. This is also little bit impossible in today's world where people use all kinds of mobile devices such as smart phones to access the resources.
9. **USE SOFTWARE THAT USE PKI ENCRYPTION** by relying on public and private keys such as grid-ftp.
10. **ALLOW ONLY ENCRYPTED PROTOCOLS** such as ssh, sftp or https to access the system.
11. **ACTIVATE TOOLS SUCH AS SELINUX** that will control the access through predefined roles, but most HPC centers do not activate it as it breaks down the normal operation of the Linux clusters and it becomes harder to debug when applications don't work as expected.
12. **WHEN RUNNING A WEB-BASED APPLICATION**, add campus Shibboleth based authentication. This will serve at least as a spam filter and can curtail denial of service like attacks.
13. **ENFORCE RESOURCE QUOTA** such as total storage capacity or compute time when applicable.
14. **ISSUE ONE TIME PASSWORD** generating devices if possible.
15. **DO NOT ALLOW ANY HIPAA COMPLAINT** research on HPC clusters with lots of users.

Isolate any HIPAA type research within restricted networks.

16. **ALLOW RESEARCH WITH SENSITIVE DATA** only within restricted networks.
17. **PREVENT ANY MOBILE DEVICES** from going back and forth between open networks and restricted networks.
18. **COMPARTMENTALIZE NETWORKS** so that it is easy to quarantine the compromised part of the network. This is also helpful in zero-day-vulnerability because the developers are still working on possible patches to fix those vulnerabilities.
19. **IT IS NOT USUALLY NECESSARY TO UPGRADE THE OS** to the latest because it takes a while for the community to test for various flaws. Upgrade the OS only if necessary.
20. **APPOINT WELL-TRAINED STAFF** with operating system and Internet knowledge with awareness of threat conditions and cyber forensic skills.

## **CLOUD AND VIRTUAL ENVIRONMENT**

Since cloud and virtualization technology became popular in an age when there is a universal awareness of cyber security, the developers of this technology have been under heavy scrutiny to make the technology harder to compromise. In the private cloud environment the resource owners usually have elevated privileges such as root password and the ability to open or close port to public network. The host providers usually don't have much control on the activities of the virtual host and virtual network, but they do have control on the hosted machines and hosted network. The cloud administrators have sufficient privileges even to examine the virtual instances running on a host machine. In a cloud environment through virtual networking and virtual instances different users are in their own isolated network and compromise on one virtual instance is isolated to only the resources owned by that group. In a way security is little bit better in cloud environment than a HPC cluster where all the nodes have identical set up. It is also easy to discard a compromised

# THE SECURITY IN PUBLIC CLOUD ENVIRONMENT IS LESS WELL DEFINED COMPARED TO A PRIVATE CLOUD ENVIRONMENT DEPENDING ON THE SERVICE LEVEL AGREEMENT.

image in a cloud environment and build and deploy a new one. Also, it is not necessary to provide public Internet connection to all virtual instances in a private cloud-computing environment.

The security in public cloud environment is less well defined compared to a private cloud environment depending on the service level agreement (SLA). The data is already moved to a public resource where employees of the service provider may have access to the data and if the data is unencrypted, there is a possibility of data getting intercepted during the transfer process. Also, access to data is not guaranteed all the time.

## **FIREWALLS AT THE CAMPUS BORDER**

A commonly used practice in university environment is network filtering (firewalls) at the campus border as the first line of defense starts there. The campus network administrators are continuously monitoring high volume and high frequency traffic for abuse and periodically block the network traffic from those IP addresses until the authenticity of the network behavior is investigated (blacklis-

ting/whitelisting). Network administrators can also block traffic specific to a certain communication protocol and ports if there are known vulnerabilities until mitigation steps have been taken.

Another commonly adopted approach is to block all inbound traffic to a group of compute resources or devices as a matter of policy and punch holes in the filter only to the resources, which needs both inbound and outbound traffic. Where possible private subnets are connected to Internet using network address translation (NAT) methodology, which involves rewriting the source and destination IP addresses when the packet traverses a firewall. This process is called IP masquerade used to hide an entire private subnet from Internet. In this way private subnet can access public network and not vice versa.

Implementing intrusion detection systems (IDS) also results in many false positive alerts because application developers do not follow any strict guidelines and anomalous behavior in one organization may be within the acceptable limits of another organization. Because of this reason many HPC centers do not rely on IDS.



# 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)



# HPC ACADEMY

Developing High-Performance Computing Expertise in Europe

The HPC Academy and its partner Graduate engineering school ESIEA intend to democratize parallel programming and promote best practices in the management of high performance IT infrastructures. To that effect, the organization has developed training courses for scientists, professionals and job seekers.

## NEXT HPC TRAINING

### HPC01-1: C++ and beyond: evolution of the language and best practices

#### About the training:

Using a sample application, we will resume the various elements emphasized in training to upgrade an application code. Participants can also attend the hands-on sessions their own code base. The training will end on a prospective point by examining the novelties already present in C++ 14 and the fundamental changes promised by C++ 17.

#### Target audience:

Engineers, technicians, scientists with intermediate knowledge in C++.

**Language:** French

**Duration:** 5 days

**Location:** Laval, France

**Price:** 4 000 euros

For more informations



Next session:

**November 2nd to November 6th 2015**

Our first partner members:



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



# HPC ACADEMY

Developing High-Performance Computing Expertise in Europe

The HPC Academy and its partner Graduate engineering school ESIEA intend to democratize parallel programming and promote best practices in the management of high performance IT infrastructures. To that effect, the organization has developed training courses for scientists, professionals and job seekers.

## NEXT HPC TRAINING

### HPC02: Developing for performance

**About the training:**

This training focuses on understanding the relationship between modern hardware systems and algorithms. It aims firstly to highlight the characteristics of the architecture of modern computing systems and the impact of these architectures on software development strategies. It will be completed with hands-on sessions implementing theoretical concepts.

**Target audience:**

Scientists, technicians and engineers in software development.

**Language:** French

**Duration:** 10 days

**Location:** Laval, France

**Price:** 8 000 euros

For more informations



Next session:

**October 19th to October 30th 2015**

Our first partner members:



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

# WHAT IS CODE MODERNIZATION?



# WHEN IT COMES TO PERFORMANCE, YOUR CODE MATTERS!



**MIKE PEARCE**  
INTEL HPC DEVELOPER EVANGELIST

**M**odern high performance computers are built with a combination of resources including: multi-core processors, many core processors, large caches, high speed memory, high bandwidth inter-processor communications fabric, and high speed I/O capabilities. High performance software needs to be designed to take full advantage of these wealth of resources. Whether re-architecting and/or tuning existing applications for maximum performance or architecting new applications for existing or future machines, it is critical to be aware of the interplay between programming models and the efficient use of these resources. Consider this a starting point for information regarding Code Modernization. When it comes to performance, your code matters!

Building parallel versions of software can enable applications to run a given data set in less time, run multiple data sets in a

fixed amount of time, or run large-scale data sets that are prohibitive with un-optimized software. The success of parallelization is typically quantified by measuring the speedup of the parallel version relative to the serial version. In addition to that comparison, however, it is also useful to compare that speedup relative to the upper limit of the potential speedup. That issue can be addressed using Amdahl's Law and Gustafson's Law.

## **GOOD CODE DESIGN TAKES INTO CONSIDERATION SEVERAL LEVELS OF PARALLELISM**

The first level of parallelism is Vector parallelism (within a core) where identical computational instructions are performed on large chunks of data. Both scalar and parallel portions of code will benefit from the efficient use of vector computing.

A second level of parallelism called Thread parallelism, is characterized by a number of cooperating threads of a single process, communicating via shared memory and collectively cooperating on a given task.

The third level is when many codes have been developed in the style of independent cooperating processes, communicating with

# **POOR DATA ALIGNMENT FOR VECTOR PARALLELISM WILL GENERATE A HUGE PERFORMANCE IMPACT. DATA SHOULD BE ORGANIZED IN A CACHE FRIENDLY WAY.**

each other via some message passage system. This is called distributed memory Rank parallelism, so named as each process is given a unique rank number. Developing code which uses all three levels of parallelism effectively, efficiently, and with high performance is optimal for modernizing code. Factoring into these considerations is the impact of the memory model of the machine: amount and speed of main memory, memory access times with respect to location of memory, cache sizes and numbers, and requirements for memory coherence.

Poor data alignment for vector parallelism will generate a huge performance impact. Data should be organized in a cache friendly way. If it is not, performance will suffer, when the application requests data that's not in the cache. The fastest memory access occurs when the needed data is already in cache. Data transfers to and from cache are in cache-lines, and as such if the next piece of data is not within the current cache-line or is scattered amongst multiple cache-lines, the application may have poor cache efficiency.

Divisional and transcendental math functions are expensive even when directly supported by the instruction set. If your application uses many division and square root operations within the run-time code, the resulting performance may be degraded because of the limited functional units within the hardware; the pipeline to these units may be dominated. Since these instructions are expensive, the developer may wish to cache frequently used values to improve performance.

There is no "one recipe, one solution" technique. A great deal depends on the problem

being solved and the long term requirements for the code, but a good developer will pay attention to all levels of optimization, both for today's requirements and for the future.

Intel has built a full suite of tools to aid in code modernization - compilers, libraries, debuggers, performance analyzers, parallel optimization tools and more. Intel even has webinars, documentation, training examples, and best known methods and case studies which are all based on over thirty years of experience as a leader in the development of parallel computers.

## 5 STAGE FRAMEWORK FOR MULTI-LEVEL PARALLELISM

The Code Modernization optimization framework takes a systematic approach to application performance improvement. This framework takes an application through five optimization stages, each stage iteratively improving the application performance. But before you start the optimization process, you should consider if the application needs to be re-architected (given the guidelines below) to achieve the highest performance, and then follow the Code Modernization optimization framework. By following this framework, an application can achieve the highest performance possible on Intel® Architecture. The stepwise approach helps the developer achieve the best application performance in the shortest possible time. In another words, it allows the program to maximize its use of all parallel hardware resources in the execution environment.

# AT THE BEGINNING OF YOUR OPTIMIZATION PROJECT, SELECT AN OPTIMIZING DEVELOPMENT ENVIRONMENT. THE DECISION YOU MAKE AT THIS STEP WILL HAVE A PROFOUND INFLUENCE IN THE LATER STEPS.

## THE STAGES

**LEVERAGE OPTIMIZATION TOOLS AND LIBRARIES:** Profile the workload using Intel VTune Amplifier to identify hotspots and Intel Advisor XE to identify vectorization & threading opportunities. Use Intel compilers to generate optimal code and apply optimized libraries such as Intel Math Kernel Library, Intel TBB, and OpenMP when appropriate.

**SCALAR, SERIAL OPTIMIZATION:** Maintain the proper precision, type constants, and use appropriate functions and precision flags.

**VECTORIZATION:** Utilize SIMD features in conjunction with data layout optimizations. Apply cache-aligned data structures, convert from arrays of structures to structure of arrays, and minimize conditional logic.

**THREAD PARALLELIZATION:** Profile thread scaling and affinitize threads to cores. Scaling issues typically are a result of thread synchronization or inefficient memory utilization.

**SCALE YOUR APPLICATION FROM MULTI-CORE TO MANY CORE** (distributed memory Rank parallelism): Scaling is especially important for highly parallel applications. Minimize the changes and maximize the performance as the execution target changes from one flavor of the Intel architecture (Intel® Xeon® processor) to another (Intel Xeon Phi Coprocessor).

## THE 5 STAGES IN PRACTICE

### STAGE 1

At the beginning of your optimization project, select an optimizing development environment. The decision you make at this step will have a profound influence in the later steps. Not only will it affect the results you get, it could substantially reduce the amount of work to do. The right optimizing development environment can provide you with good compiler tools, optimized, ready-to-use libraries, and debugging and profiling tools to pinpoint exactly what the code is doing at the runtime. Check out the webinars on the Intel® Advisor XE tool, that can be used to identify vectorization & threading opportunities.

### STAGE 2

Once you have exhausted the available optimization solutions, in order to extract greater performance from your application you will need to begin the optimization process on the application source code. Before you begin active parallel programming, you need to make sure your application delivers the right results before you vectorize and parallelize it. Equally important, you need to make sure it does the minimum number of operations to get that correct result. You should look at the data and algorithm related issues such as:

- Choosing the right floating point precision
- Choosing the right approximation method accuracy; polynomial vs. rational

# THE MEMORY ACCESS OF THE OUTER LOOP IS MORE LIKELY TO BE DIVERGENT THAN THAT OF AN INNER LOOP.

- Avoiding jump algorithms
- Reducing the loop operation strength by using iteration calculations
- Avoiding or minimizing conditional branches in your algorithms
- Avoiding repetitive calculations, using previously calculated results.

You may also have to deal with language-related performance issues. If you have chosen C/C++, the language related issues are:

- Use explicit typing for all constants to avoid auto-promotion
- Choose the right types of C runtime function, e.g. doubles vs. floats: `exp()` vs. `expf()`; `abs()` vs. `fabs()`
- Explicitly tell compiler about point aliases
- Explicitly Inline function calls to avoid overhead

## STAGE 3

Try vector level parallelism. First try to vectorize the inner most loop. For efficient vector loops, make sure that there is minimal control flow divergence and that memory accesses are coherent. Outer loop vectorization is a technique to enhance performance. By default, compilers attempt to vectorize innermost loops in nested loop structures. But, in some cases, the number of iterations in the innermost loop is small. In this case, inner-loop vectorization is not profitable. However, if an outer loop contains more work, a combination of elemental functions, strip-mining, and `pragma/directive SIMD` can force vectorization at this outer, profitable level. SIMD performs best on “packed” and aligned input data, and by its nature penalizes control divergences. In addi-

tion, good SIMD and thread performance on modern hardware can be obtained if the application implementation puts a focus on data proximity.

If the innermost loop does not have enough work (e.g., the trip count is very low; the performance benefit of vectorization can be measured) or there are data dependencies that prevent vectorising the innermost loop, try vectorising the outer loop. The outer loop is likely to have control flow divergence; especially if the trip count of the inner loop is different for each iteration of the outer loop. This will limit the gains from vectorization. The memory access of the outer loop is more likely to be divergent than that of an inner loop. This will result in gather / scatter instructions instead of vector loads and stores and will significantly limit scaling due to vectorization. Data transformations, such as transposing a two dimensional array, may alleviate these problems, or look at switching from Arrays of Structures to Structures of Arrays.

When the loop hierarchy is shallow, the above guideline may result in a loop that needs to be both parallelized and vectorized. In that case, that loop has to both provide enough parallel work to compensate for the overhead and also maintain control flow uniformity and memory access coherence.

## STAGE 4

Now we get to thread level parallelization. Identify the outermost level and try to parallelize it. Obviously, this requires taking care of potential data races and moving data declaration to inside the loop as necessary. It may also require that the data be maintained in a cache efficient manner, to reduce the overhead of

## IT HELPS TO HAVE AS LARGE A PARALLEL EFFORT IN EACH THREAD AS POSSIBLE.

maintaining the data across multiple parallel paths. The rationale for the outermost level is to try to provide as much work as possible to each individual thread. Amdahl's law states: The speedup of a program using multiple processors in parallel computing is limited by the time needed for the sequential fraction of the program. Since the amount of work needs to compensate for the overhead of parallelization, it helps to have as large a parallel effort in each thread as possible. If the outermost level cannot be parallelized due to unavoidable data dependencies, try to parallelize at the next-outermost level that can be parallelized correctly.

If the amount of parallel work achieved at the outermost level appears sufficient for the target hardware and likely to scale with a reasonable increase of parallel resources, you are done. Do not add more parallelism, as the overhead will be noticeable (thread control overhead will negate any performance improvement) and the gains are unlikely.

If the amount of parallel work is insufficient, e.g. as measured by core scaling that only scales up to a small core count and not to the actual core count, attempt to parallelize additional layer, as outmost as possible. Note that you don't necessarily need to scale the loop hierarchy to all the available cores, as there may be additional loop hierarchies executing in parallel.

If step 2 did not result in scalable code, there may not be enough parallel work in your algorithm. This may mean that partitioning a fixed amount of work among many threads gives each thread too little work, so the overhead of starting and terminating threads swamps the useful work. Perhaps the algorithms can be scaled to do more work, for example by trying on a bigger problem size.

Make sure your parallel algorithm is cache efficient. If it is not, rework it to be cache efficient, as cache inefficient algorithms do not scale with parallelism. Check out the Intel Guide for Developing Multithreaded Applications series for more details.

### STAGE 5

Lastly we get to multi-node (Rank) parallelism. To many developers message passing interface (MPI) is a black box the "just works" behind the scenes, to transfer data from one MPI task (process) to another. The beauty of MPI for the developer is that the algorithmic coding is hardware independent. The concern that developers have, is that with the many core architecture with 60+ cores, the communication between tasks may create a communication storm either within a node or across nodes. To mitigate these communication bottlenecks, applications should employ hybrid techniques, employing a few MPI tasks and many OpenMP threads.

A well-optimized application should address vector parallelization, multi-threading parallelization, and multi-node (Rank) parallelization. However to do this efficiently it is helpful to use a standard step-by-step methodology to ensure each stage level is considered. The stages described here can be (and often are) reordered depending upon the specific needs of each individual application; you can iterate in a stage more than once to achieve the desired performance.

Experience has shown that all stages must at least be considered to ensure an application delivers great performance on today's scalable hardware as well as being well positioned to scale effectively on upcoming generations of hardware.

NEXT ISSUE



#7  
**SEPTEMBER**  
**30<sup>th</sup>, 2015**

**TECH ZONE**

**SSDS FOR BIG  
DATA: FAST  
PROCESSING  
REQUIRES HIGH-  
PERFORMANCE  
STORAGE**

**COVER STORY**

**HYPERCONVERGED  
INFRASTRUCTURES**

Do they live up to the hype?

**VIEWPOINT**

**WHEN MASSIVE  
DATA NEVER  
BECOMES BIG  
DATA**

**HOW TO**

**BUILD A SUPERCOMPUTER WITH 1PFLOPS  
OF PEAK COMPUTING PERFORMANCE  
USING MANY-CORE PROCESSORS**

 **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