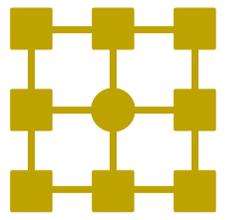


**EXCLUSIVE!** Win the Green 500 champion server from ASUS



# HPC REVIEW

The reference media in high-performance IT solutions

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

HPC | BIG DATA | CLOUD | STORAGE | VISUALIZATION | VIRTUALIZATION | NETWORKS | WORKSTATIONS

# MACHINE LEARNING

THE FUTURE OF HPC?

Newsfeed  
**Intel Xeon E7 v3**

the Big Data  
crunching  
processor

Exclusive  
Interview  
**George  
Teixeira,**  
DataCore  
CEO

Tech Zone  
**Green Mountain,**  
the greenest  
datacenter  
in the world

#3 GLOBAL EDITION

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)



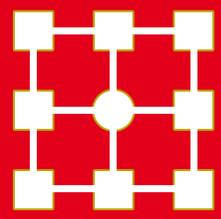
# Welcome !

**This issue focuses on Machine Learning**, another hot topic for HPC system engineers, researchers and developers nowadays. Achieving a scalable Machine Learning platform involves rethinking all the technology layers, starting from the algorithms and linking them to new database technologies able to stand the test of zettabytes of data. It is also a challenge on the hardware side : in order to reduce latency and maximize algorithm and application throughput for mission-critical sorts, filters and full text searches over distributed indexes requires a great deal of last generation hardware, but also a rethinking of network and storage infrastructure in a performance-oriented architecture. We have studied eBay's case of such an extreme environment.

**Finetuning can be profitable.** Achieving better performance is not only about shelling big bucks, but also eliminating the bottlenecks. Our Viewpoint contributor helps us understand some of the intricacies of dealing with last generation storage interfaces and first or second generation disks to achieve the best performances without needing to replace the existing storage resources. Fine tuning can be profitable ! As for many HPC usages, it essentially boils down to a single rule of thumb: adjust, refine the existing infrastructure by making it as coherent and balanced as possible can achieve positive results even in a tight and limited budget environment. All it takes is some understanding of the underlying technologies and how they combine within the global infrastructure. This applies to clusters of systems, as well as within a system. We welcome your suggestions and comments at [editorial@hpcreview.com](mailto:editorial@hpcreview.com)

Happy reading!

# BIG HPC REVIEW CONTEST



# HPC REVIEW

**A unique chance to win  
an ASUS ESC4000 G2S!**

The groundbreaking server that powers  
the World's Number One Green500 cluster



**Ready to run with Intel Xeon CPUs,  
AMD GPUs, memory and storage!**

Asus ESC4000 G2S (2U) features :

- Processor : Dual Intel® Xeon processor E5-2600 v2 product family
- Core Logic : Intel C602 PCH
- Memory : 16 x DDR3 DIMM Slots (Up to 512GB R-DIMM)
- HDD bays : 6 x hot-swappable 2.5in bays
- Networking : 2 x Intel GbE LAN (82574L)  
+ 1 x Management LAN
- Graphic Cards: 2 x AMD FirePro S9150 - 16Gb

**Click below to enter the contest!**

**I want to win!**

**The winner will be announced on September 1st, 2015**

Free contest with no purchase necessary. Participation open worldwide.





## COVER STORY

# MACHINE LEARNING

THE FUTURE OF HPC?



## NEWSFEED

**Intel Xeon E7 v3**

**Openstack and hyperconvergence with DataCore**

**Atlantis Computing**

**NASA triples the power of the NCCS to model the earth**

**Meet the Memristors**

**228TB on prototype tape**

**CubeSat concept**

**New SanDisk PCIe accelerators**

**Books | Moocs | The HPC observatory**

## LAB REVIEW



**HP Converged-System 240  
HC StoreVirtual**

**NVIDIA  
GTX Titan X**

**Netgear  
ReadyNAS 314**

**HPC Labs :  
How do we test**

## VIEWPOINT

**How to eliminate storage bottlenecks in a data growing world**

## TECH ZONE

**Green Mountain, the greenest data center in the world**



# HPC Cloud Computing is coming of age

In IDC's worldwide studies of high performance computing (HPC) end-user sites, the proportion of sites employing cloud computing—public or private—has steadily grown from 13.8% in 2011, to 23.5% in 2013, to 34.1% in 2015. Also represented in this mix is the growing contingent of hybrid clouds that blur the public-private distinction by combining on-premise and external resources. Our research shows that there are persistent concerns about data security in public clouds, even though data security and confidentiality have generally improved. For example, Amazon Web Services, the most popular public cloud among HPC users, is now compliant with HIPAA, the federal Health Insurance Portability and Accountability Act designed to safeguard the privacy and security of health information.

## SENDING WORK OFF PREMISE

Another important brake on sending work off premise has been the widespread perception that this means using public clouds that are suitable only for embarrassingly parallel workloads—ones that can readily be subdivided into smaller jobs, each of which can be run independently. San Francisco-based Rescale is one of the newer companies (2011) that is altering the landscape of what can be done effectively beyond an organization's firewall. Nearly all of the privately held company's principals have strong backgrounds in fields where structural analysis and fluid-structures interactions are important. Founder and CEO Joris Poort began his career as a structural and software engineer working on the Boeing 787 "Dreamliner" airplane, as did co-founding



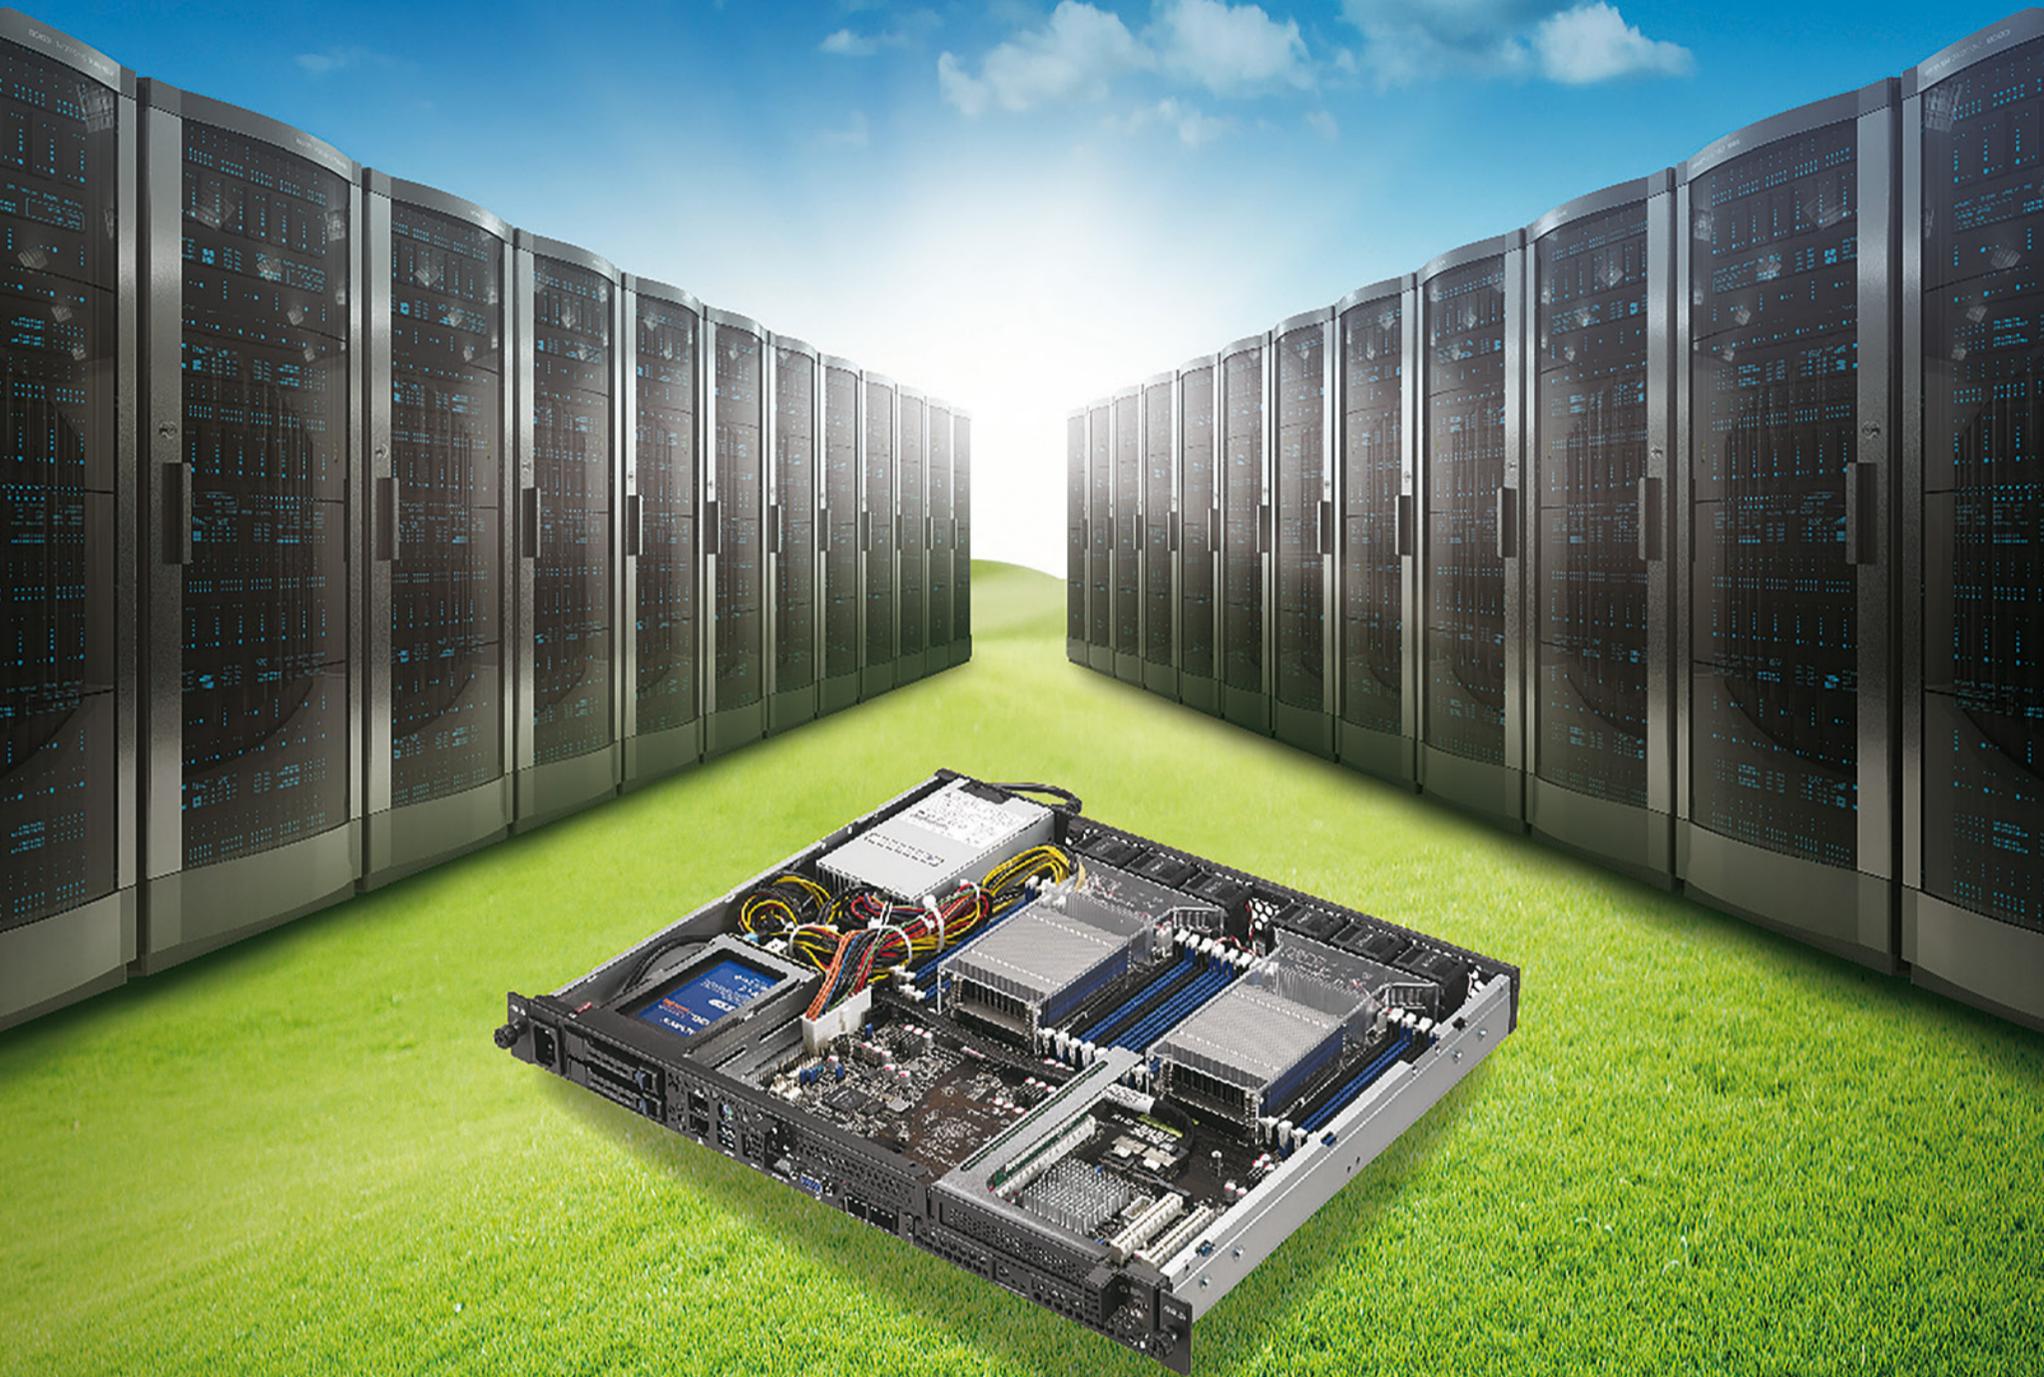
**STEVE CONWAY**  
is IDC Research  
Vice President,  
HPC

CTO Adam McKenzie. Sales VP Tony Spagnuolo headed aerospace sales at MSC.Software following stints at Rockwell International and General Motors. Marketing VP Shing Pan led solver product marketing at Altair Engineering. The company has deep bench strength in low-latency, IO-intensive HPC.

## A NEW BUSINESS MODEL

The result is a business model and products (ScaleX Enterprise, ScaleX Developer, and ScaleX Pro) that exploit contracts not only with public cloud service providers, but also with private clouds and large HPC centers. Rescale matches customer requirements to the appropriate resource or resources, leveraging ISV partnerships with Ansys, CD-adapco, Dassault Systemes, ESI, MSC.Software, and Siemens PLM, along with popular ISVs in fields including energy, life sciences and weather forecasting, plus open source software such as OpenFoam. Customers decide where their data is stored. Aerospace and defense organizations are among the company's initial customers. As expected, pricing options are risk-based, with advance bookings priced lower than last-minute requests.

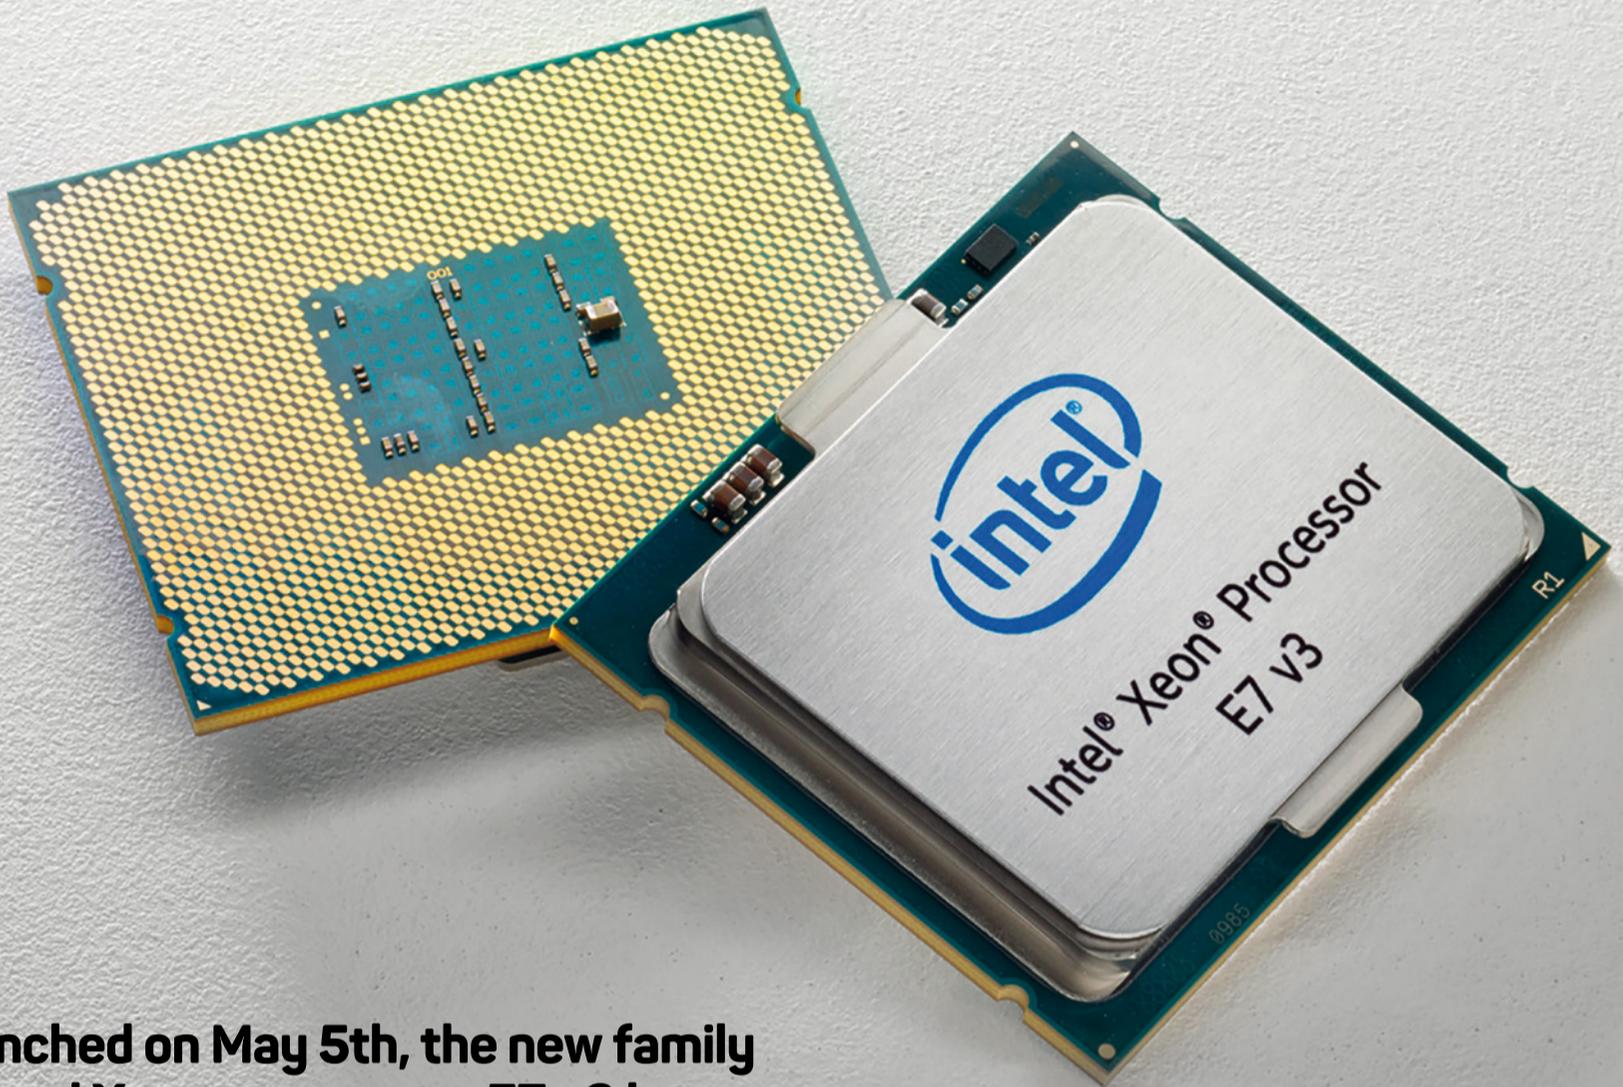
**ASUS**<sup>®</sup>  
IN SEARCH OF INCREDIBLE



# RS400-E8-PS2-F

Most Compact 1U Dual Xeon Flexible I/O Server

[http://www.asus.com/Commercial Servers Workstations/](http://www.asus.com/Commercial_Servers_Workstations/)

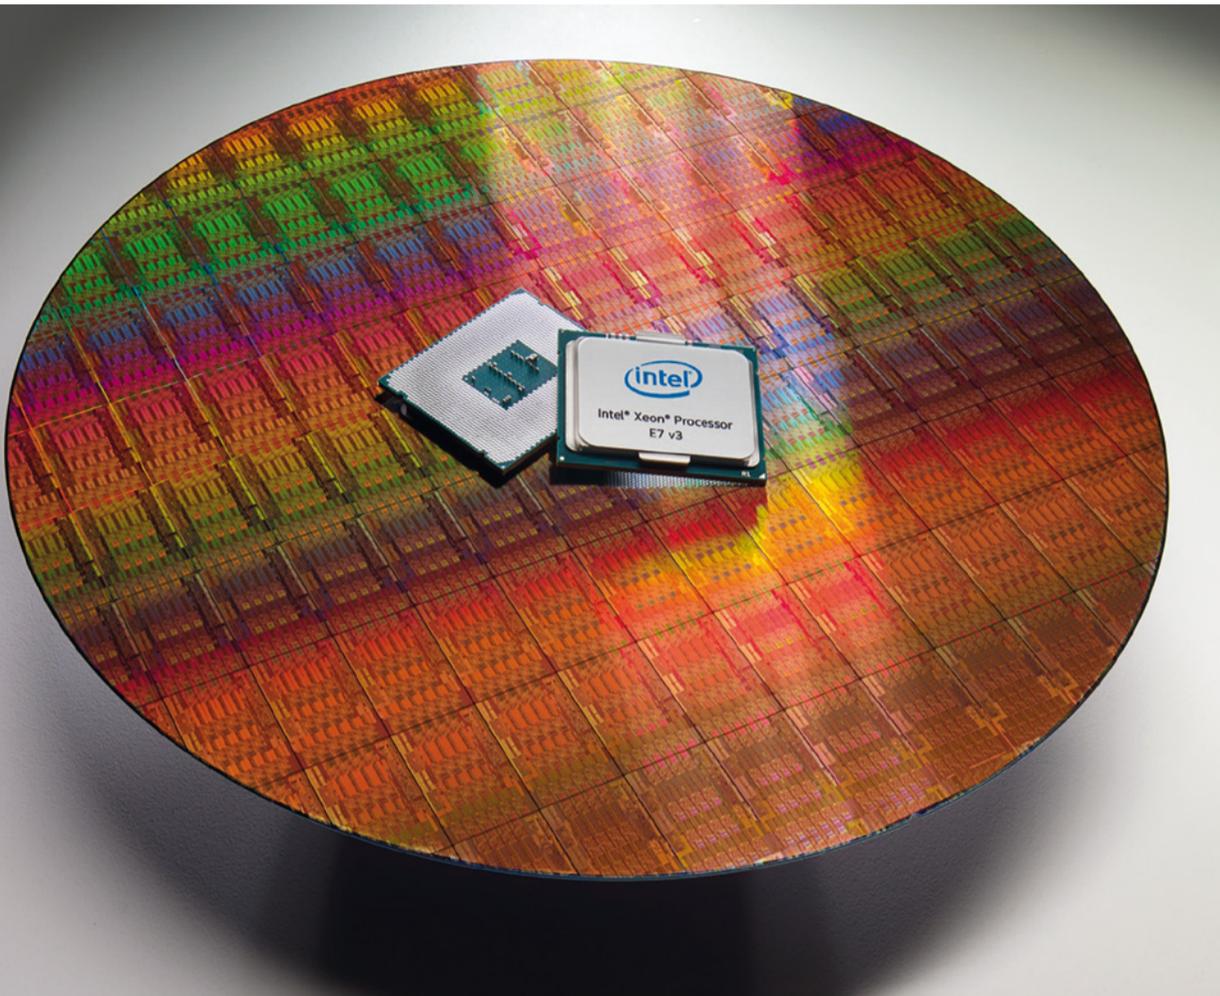


Launched on May 5th, the new family of Intel Xeon processors E7 v3 has built-in characteristics targeted towards real-time analysis in a big data context

# Intel Xeon E7 v3, the Big Data oriented processor

**E**ngraved in 22nm technology and counting 5.7 billion transistors, the Intel Xeon E7-8800 / 4800 v3 processor family is designed to improve access times to information through their analytical abilities in real time, transforming data into strategic advantages. The economic analysis in real time is a priority for many businesses sectors ranging from health to trade through te-

lecommunications. The need for immediate access to actionable information from large volumes of data is at the heart of the growth of new big data analysis technologies and in-memory computing. Given the current and future demand, revenues from this market are expected to exceed \$ 95 billion per year by 2018. 50% of large organizations will then have adopted in-memory processing to design their digital strategies.



**“In the digital services economy, success or failure depends on the speed at which a company can act on the information it has masses of data”**

**Diane Bryant, vice president and general manager of Data Center Group, Intel**

### **A PROCESSOR DESIGNED TO ACCELERATE THE ANALYSIS AND DECISION MAKING**

The family of Intel Xeon processors E7 v3 allows customers to process and analyze large amounts of data securely in memory systems for a faster decision making and enhanced operational efficiency to ensure the same business competitiveness. *“In the digital services economy, success or failure depends on the speed at which a company can act on the information it has masses of data,”* said Diane Bryant, vice president and general manager of Data Center Group, Intel. *“The Intel Xeon E7 family v3 is a driving force in accelerating this phenomenon. It provides real-time analysis that helps companies enhance customer satisfaction by offering more customized products and services, and create new revenue streams while improving operational efficiency.”* According to the manufacturer the highlights of this new generation are :

- **The new family of processors** provides an average performance improvement of 40% compared to the previous generation.
- **Performance is multiplied by six** in appli-

cations processing in-memory transactional workloads optimized with new Intel Transactional Synchronization Extensions instructions (Intel TSX).

- **With up to 18 cores**, 20% better compared to the previous generation, and up to 45 MB of last level cache, this processor family offers up to 70% more analysis sessions per hour compared to the previous generation.
- **The performance per dollar is up to ten times better**, with a total cost of ownership announced 85% lower compared to alternative RISC architectures, thanks to reduced hardware acquisition costs, power supply, and cooling expenditures.
- **The maximum configuration accepts up to 32 sockets** with the greater memory capacity per socket of the market (up to 12 terabytes of DDR3 or DDR4).
- **New security and reliability capabilities are present**, including increased cryptographic performance of the new Intel Advanced Encryption Standard New Instructions (AES-NI), as well as other improvements at the micro-architectural level. The Intel Xeon E7 family



**“Data that previously needed days or weeks to be analyzed are now available almost in real time.”**

**Justin Chen, IT director of Nippon Paint**

v3 includes twelve different models of processors, including many models optimized for various market segments, such as the two high-frequency processors designed for database applications requiring the fastest possible cores. 17 manufacturers worldwide have begun to announce platforms based on the Intel Xeon processor E7 v3: Bull, Cisco, Dell, Fujitsu, Hitachi, HP, Huawei, Inspur, Lenovo, NEC, Oracle, PowerLeader, Quanta, SGI, Sugon, Supermicro and ZTE. And Intel and Cloudera have at the same time announced new details about the impact of their partnership on the adoption of wide analysis process. They relate in particular to the implementation of Intel AES-NI instructions that multiply by 2.5 off-load encryption performance, allowing Hadoop data to be encrypted using only 1% of the processor.

### **TRANSFORMING THE ENTERPRISE WITH REAL-TIME ANALYSIS**

NipponPaint, one of the biggest paint producers in Asia, currently uses Intel Xeon E7 v2

processors with in-memory analysis software SAP HANA to capture consumer behavior more accurately and accelerate its analytical capacity and action to optimize their marketing campaigns. *“With the real-time analysis, NipponPaint can track and analyze consumer preferences in various fields such as color, style or designer to create products and services specifically to their needs.”* said Justin Chen, IT director of Nippon Paint. *“Data that previously needed days or weeks to be analyzed are now available almost in real time. We are currently testing systems based on the new Intel Xeon processor E7 v3, to enjoy its new features performance and reliability, and speed up our analysis.”*

The Xeon E7 processors v3 embed new instructions such AVX 2.0, Cache Monitoring and VMCS shadowing, with performances up to upper 35% in OLTP. DDR4 memory bus provides 16% more bandwidth, and up to 50% improved energy efficiency compared to the previous generation. **RAMON LAFLEUR**



# Software Defined Storage Openstack and hyperconvergence with DataCore

The CEO of DataCore, George Teixeira, delivers his vision of the software defined market storage through recent announcements.

**T**he new edition of SANsymphony PSP is a milestone for DataCore's flagship products, SANsymphony V10 and Virtual San. DataCore has a track record for many years now in unifying multiple available storage environments on the storage market. The software publisher is straightforward : whether we are talking about hyperconvergence, virtual or physical SAN or cloud storage, both Sansymphony V10 and Virtual San support them all. DataCore's software services are listed by the publisher in four categories : availability, performance, efficiency and of course storage infrastructures management.

## STORAGE AND CLOUD, AN ESSENTIAL TANDEM

Among the recent proposed improvements to the management part, DataCore has committed to integrate its software with cloud environments. After Microsoft Azure which was



the premier platform supported and Amazon Web Services, DataCore has added support of the de facto standard Openstack including Cinder. It has been impossible nowadays to ignore Openstack because the majority of equipment manufacturers and vendors operating in the storage sector support it. It therefore has become an obvious way to open up to many private and hybrid platforms. Openstack goes along with another big trend of hyperconverged servers where computing, storage and network are combined on standard x86 server hardware. On the datacenter side, DataCore plays on both fronts: distributed and multi-tenancy.

## SDS SERVICES EVERYWHERE

On the Software Defined Storage side, new uses like VDI have appeared that justify investing in SDS infrastructure. According to George Teixeira, the target in this market are small and medium-sized enterprises or departmental organizations with smaller installations up



**The idea of hyperconvergence is based on standard physical hardware servers which consolidate computing, storage and network to facilitate the expansion of data centers and company infrastructures which always need more power. The storage market is no exception to this trend.**

to 500 seats. One of the main obstacles to the adoption of VDI was that performance was inevitably linked to expensive storage infrastructure such as Flash. DataCore has specialized for over two years on the VDI business by combining the functions of its virtual storage infrastructure and the latest technology, enabling its customers to reduce the costs. For example, inputs / outputs are optimized by performing certain tasks directly by using memory as intelligent caching instead of relying on flash memory. Equally noteworthy improvements relate to the automation and management side of VDI deployments. Additional features are deduplication and compression performed in the background for some clustered virtual disks in pools. Support for Veeam environment (through scripts) helps improve the automation and centralization of these tasks on VMware environments. SANsymphony PSP also features a centralized dashboard for different geographically distributed SANsymphony environments while providing finer granularity in terms of administration on remote sites.

### **IMPROVED INPUT/OUTPUT PERFORMANCE**

DataCore has worked on I/O at different levels, starting with the arrival of new metrics for virtual disks and pushing alerts on physical environments. Technologically, the access to flash memory resources has also been reworked. Along the same lines, the Random Write Accelerator feature uses memory as a superfast cache

in order to write data sequentially on the drives and not in a dispersed manner, optimizing the latency induced by the access to the data.

### **SDS AND HYPERCONVERGENCE THROUGH PARTNERSHIPS**

The idea of hyperconvergence is based on standard physical hardware servers which consolidate computing, storage and network to facilitate the expansion of data centers and company infrastructures which always need more power. The storage market is no exception to this trend. Being agnostic to the hardware, DataCore's SDS solution has been discussed with different physical platform providers. Huawei, Fujitsu, Dell, Lenovo or Cisco are today partnering with DataCore. For the company's CEO and co-founder, hyperconvergence is a natural evolution of IT. Regarding SDS, the market's software evolution about software defined storage to which he was a pioneer in 1998, eventually proved him right. So much so that according to George Teixeira, the result of this convergence is inevitable, which he summed up in one sentence: *"the future of storage is the server."* George Teixeira also evoked very positive growth perspectives through partnerships like the one signed last March with Huawei and the ongoing talks with the second largest Chinese manufacturer Lenovo, which if successful, will make two new giants emerge with a very strong hunger to conquer the hyperconverged server market... **SOLANGE BELKHAYAT-FUCHS**



# Atlantis Computing slashes flash storage prices

**Atlantis Computing is an American company who arrives in Europe, specializing in flash storage appliances and which has just launched a fully converged flash appliance announced at a cost up to 90% lower than the rates currently available. And with four renowned partners to accompany it.**

**D**esigned for medium-sized companies, Atlantis' HyperScale is an hyperconverged appliance, meaning it pre-integrates a storage management layer defined by the company's patented software on any server at one of the four HP launch partners, Cisco, Lenovo or SuperMicro, and equipped with Intel Xeon or Intel SSD processors, and all VMware or Citrix hypervisor. "Medium-sized businesses need an extremely rapid economic storage, proposed by a provider of computer equipment, they know and trust. HyperScale Atlantis is the first ready-to-run appliance to integrate flash storage at less than half the current cost", says Chetan Venkatesh, founder and CEO of Atlantis.

Built on Atlantis USX, the SDS platform of the company, the CX-12 SuperMicro-based appliance offers 12 Terabytes and is sold at a record price of 78000 dollars, way lower than the competing alternatives from Nutanix or Simplivity. An appliance has 4 blade servers, each with two Intel Xeon E5-2680 v3 processors, 256 or 512GB SSD memory and 16 400 GB modules (16 x 800 GB for the CX-24 version of 24 TB). All versions include two standard 10GbE ports.

## THE ADVANTAGES PUT FORWARD BY THE MANUFACTURER

**All-flash performance:** Atlantis HyperScale performs deduplication memory before the data is written to storage, which minimizes



latency, delivers high performance in terms of IOs and reduces traffic on the storage network. **A revolutionary cost:** The hyper-converged appliances Atlantis HyperScale reduce costs using its IO acceleration technology and data reduction to reduce the amount of physical memory and flash memory needed by local TB of storage capacity.

**Freedom of choice:** Atlantis HyperScale gives companies the freedom to use their preferred server platforms and hypervisors, eliminating the complexity and additional investments in training, time and tools.

**Modular Scalability:** Atlantis HyperScale allows companies to scale their storage as and extent of the growth of their IT resources, rather than buy large expensive storage arrays and with limited performance.

Although little known in Europe yet, Atlantis Computing has more than 800 customers in its portfolio and the equivalent of 40 Petabytes of storage deployed. The support (three-year onsite intervention in less than 4 hours) is performed by Atlantis, even on HP, Supermicro, Cisco and Lenovo servers.



As with previous supercomputer upgrades, NCCS staff has worked closely with SGI for installing the new servers.

# NASA triples the power of the NCCS to model the earth

BRUCE PFAFF / MICHAEL CHYATTE





## Over 3.3 petaflops, or 3.361 trillion floating point operations per second.

**T**he NASA Center for Climate Simulation (NCCS) has almost tripled the peak performance of its supercomputer Discovery, reaching over 3.3 petaflops, or 3.361 trillion floating point operations per second. This unprecedented upgrade is necessary to meet the explosion of modeling applications of Earth science at NASA. First step was to choose the best equipment able to run the GEOS-5 (Goddard Earth Observing System Model) and NUCWRF (NASA Unified Weather Forecasting and Research) applications.

### MORE CORES FOR BETTER PERFORMANCE

The adopted SGI Rackable cluster-based solution is intended to replace parts of the Discoverer supercomputer built in 2011 to establish three Scalable Compute Units (SCU) exploiting the current generation of Intel Xeon Hawsell E5-2697 v3 processor with 14 cores each, for a cumulative total of 64,512 processor cores. Teams from the NCCS constantly keep in mind the needs of its user community during the technical definition of the new components of Discovery. *“Gradually, as the number of cores increases, memory requirements also grow forcing us to review the source code of our applications to take advantage”* said Daniel Duffy, the director of the unit HPC NCCS. Within each of the three SCUs, each processor core is flanked by 4 gigabytes of memory. The total being 138 Terabytes.

### AN EARTH MODEL THAT IS GAINING IN PRECISION

The results calculated by the overall calculation model GEOS-5 at 12 km resolution are carried on to the new clusters to calculate regional models in better definition of 24, 12 and

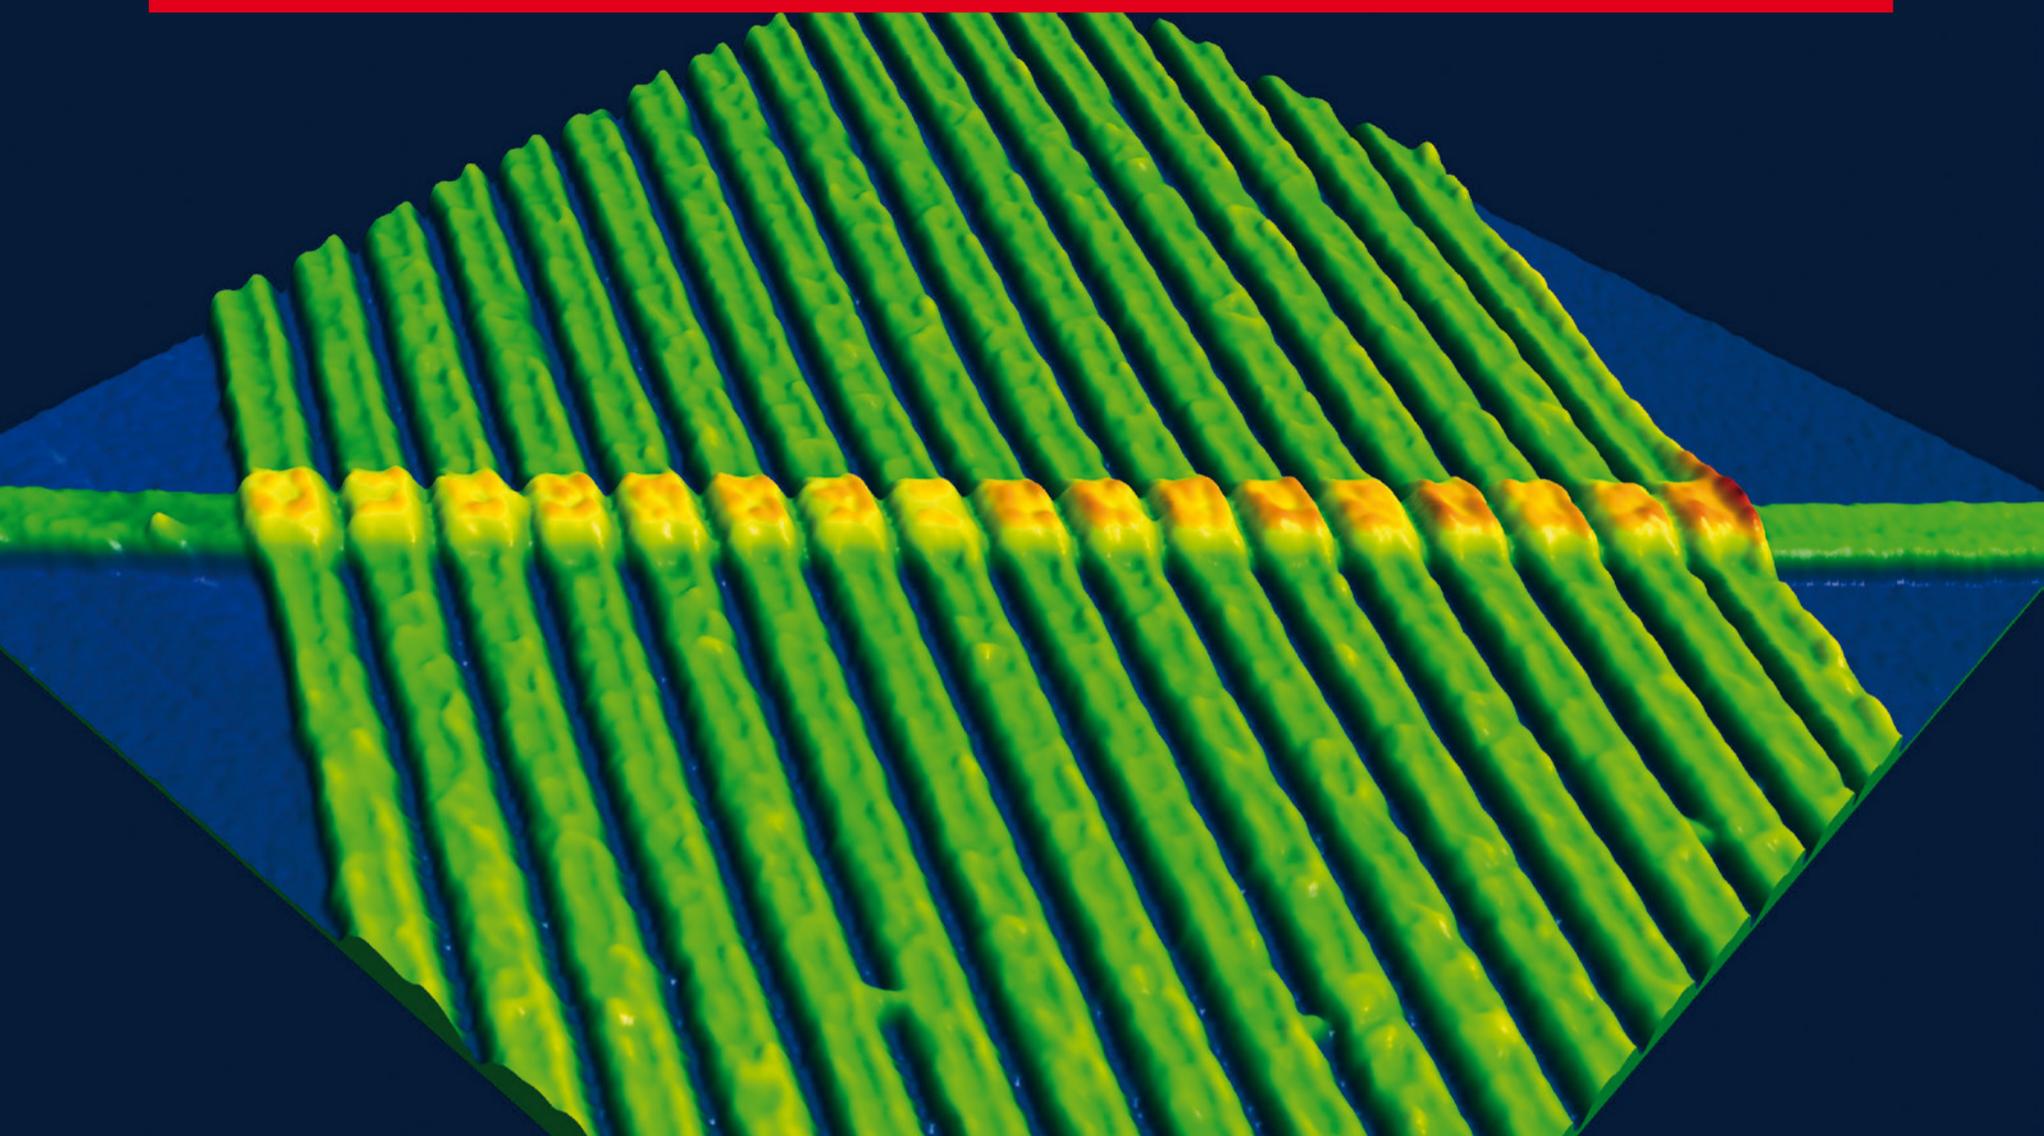


BRUCE PFAFF

**Nasa's Center for Climate Simulation (NCCS) Discovery supercomputer is based on SGI rack servers to house a total of 64,512 processor cores.**

4 km. The researchers then compare the accuracy gain on weather phenomena such as northwestern US snow storms, summer storms and the western states rivers in winter. Another architectural improvement is the presence of a full speed interconnect fabric. On newer servers, each node consists of 28 cores which communicate with other nodes via an FDR (Fourteen Data Rate) Infiniband link, permitting speeds of 56 gigabits per second.

*“GEOS-5 calculation models become very large and require extremely fast processors and communication buses. If they are insufficiently dimensioned, this may become a handicap to the smooth running of the simulations to be calculated”* insists Daniel Duffy. This data explosion due to the accuracy improvement of the calculated models also impacts Discovery's storage requirements, which were increased to 33 Peta-bytes. The normal time of installation of an UCS averages a year. *“When we're done, we will have installed three SCU in 7 months”*, said Mike Donovan's Discovery system administrator. **RAMON LAFLEUR**



# Meet the Memristors, circuits inspired by the human brain

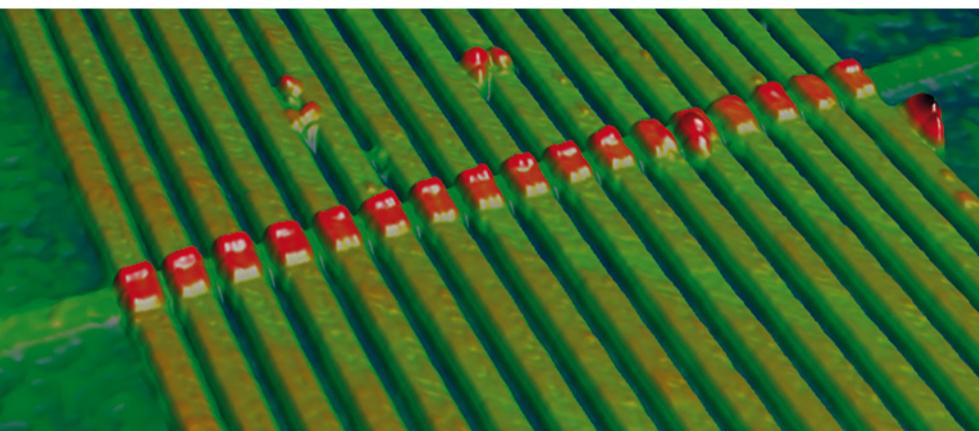
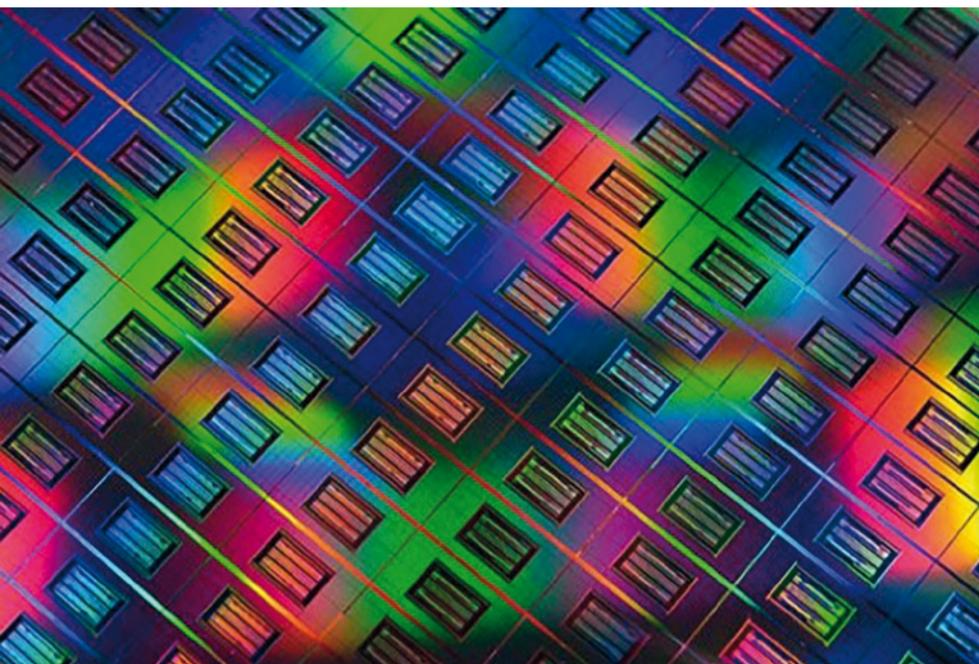
**T**he memristors are circuits of a new kind that go back to 2008. Their peculiarity is to be inspired by the human brain. Unlike traditional processors, the prototype of this new kind of processor only handle single tasks at the moment, like differentiating between black and white colors of presented objects. In the future, memristors should be able to handle way more complex data like speech, images and the surrounding environment.

The memristor developed by researchers at the Universities of California and Stony Brook, processes data not with digital logic circuits, but with elements that mimic, in simplified form, the neurons and synapses of the biologi-

cal brain. When such a network is exposed to new data, it “learns” by dynamically rewiring the synapses that connect neurons together.

## **A NEW APPROACH TO AN EXISTING CONCEPT**

Software versions of neural networks have been around for decades and have recently been at the heart of innovations in facial recognition and speech, in the context of demonstrations by Facebook and Google. Another experiment conducted by Google has demonstrated a neural network learning to recognize cats from YouTube images. For this experiment, Google has used 16,000 processors during three days. If this experience can make us smile, we only



need to imagine the possibilities of a much smaller and more energy-efficient neural processor. Such processors could for example allow mobile robots to be able not only to learn from the existing environment, but also to get closer to a form of intelligence by correlating and contextualizing collected data.

### **SYNAPSES AND NEURONS IN HIGH-TECH VERSION**

IBM was the first to attempt to make “neuromorphic” circuits inspired by the human brain, but using the same materials than traditional processors, i.e. silicon-based transistors. But according to Dmitri Strukov, assistant professor at the University of California, Santa Barbara who led the work on the new Memristor chip, the usual components are not adapted to simulate such operation. One reason is the large number of transistors needed to represent a single synapse. Conversely, on the Memristor prototype, each of the 100 synapses is represented by a single Memristor. *“A biological synapse is an analog memory device, and*

*it really is not a satisfactory solution in terms of compactness and energy savings using conventional technologies,”* Strukov said. *“However, memristors are a device analog memory and are perfectly suited to achieve this goal.”*

### **THE POTENTIAL TO CHANGE COMPUTING AS WE KNOW IT**

Robert Legenstein, associate professor at the Graz University of Technology in Austria, believes that *“if this concept can be extended to neural networks of large size, it will affect the future of computing: laptops, mobile phones and robots could incorporate low-energy capable of neuromorphic powerful processors to process visual information, hearing and any type of sensory information in real time.”*

### **A 44-YEAR OLD MATHEMATICAL MODEL COMES ALIVE**

The memristor was predicted mathematically in 1971 by Leon Chua, a professor of electronics at the University of California, Berkeley. It was created in 2008, when a team of Hewlett-Packard researchers, of which was part Strukov, created simple circuits, with electrical resistance properties coded as current. These circuits were immediately considered as having the potential to be used for storage of denser data in a neural network.

### **MEMRISTOR-PROPELLED RRAM FOR TOMORROW'S COMPUTING**

HP and Hynix memory manufacturer began cooperating in 2010 to develop components based on Memristors as RRAM or ReRAM ultrafast memories. But until now, no one had created a neural chip built entirely with memristors says Strukov. His team has since developed ways to improve the memristors manufacturing process to produce more reliable devices than before. This is a bet on the future, memristors have the potential to radically redefine computer architecture, he says. Enough to hold according to hp, 100 TB of data on a smartphone in less than ten years !

**RAMON LAFLEUR**



# 228TB on prototype tape

**Fujifilm Recording Media U.S.A., Inc., a subsidiary of Fujifilm Corporation announced that in conjunction with IBM, a new record in areal data density of 123 billion bits per square inch on linear magnetic particulate tape has been achieved.**

**F**or the fourth time in less than 10 years, Fujifilm and IBM have accomplished record breaking storage capacities on tape, announcing the highest capacity storage media ever achieved, including HDD, BD or solid memory NAND flash technologies. This breakthrough in data density equates to a single tape cartridge capable of storing up to 220 terabytes of uncompressed data. 220 terabytes is more than 88 times the storage capacity of the current LTO Ultrium 6 tape. A tape of this size can provide enough storage to preserve the human genome of 220 people on a single cartridge.

*“With high performance computing and cloud storage services on the rise, this data density achievement is significant,”* said Peter Faulhaber, president, FUJIFILM Recording Media USA, Inc. *“Fujifilm and IBM are leading the technological development of advanced tape innovation that meets the market’s growing data requirements and delivers tape as the medium of choice for archival storage.”* This record breaking demonstration was achieved using an advanced prototype tape incorporating NANOCUBIC technology developed by Fujifilm, with advanced tape-drive technologies developed by IBM.

## **FUJIFILM TECHNOLOGY ENHANCEMENTS**

Fujifilm’s NANOCUBIC technology is enhanced to increase recording density by decreasing the magnetic particle size that is essential for high recording density. Fujifilm’s original BaFe synthesis method increases the uniformity of BaFe particle size and decreases 25% of the switching field distribution (SFD), which

is an important magnetic parameter for high density recording. The lower SFD leads to a high quality signal output due to the uniform magnetic property of each recorded bit. To ensure the stability of the ultra-fine BaFe particles, Fujifilm improved the magnetic coercivity, yielding an archival life of over 30 years.

A highly controlled dispersion process and newly developed chemical compound allows the BaFe particles to separate and disperse more uniformly and increases the perpendicular oriented ratio. Perpendicular orientation technology with BaFe produces a high signal to noise ratio and better frequency response. Enhanced NANO coating technology with a very smooth non-magnetic layer controls the tape surface roughness, providing a smooth magnetic layer for higher signal output. Fujifilm’s advanced servo writing technology decreases high frequency vibration of the servo tracks and enables a higher track density due to more precisely placed servo tracks.

## **IBM TECHNOLOGY ENHANCEMENTS**

A set of advanced servo control technologies enable more accurate head positioning and increased track density. An enhanced write field head technology enables the use of much finer barium ferrite particles. Innovative signal-processing algorithms for the data channel enable reliable operation with an ultra-narrow 90nm wide giant magnetoresistive (GMR) reader. Fujifilm will continue to lead the development of large capacity data storage media with BaFe technology to provide a cost-effective archival solution to preserve digital data.



# NASA creates CubeSat concept for planetary exploration

**A**lthough scientists are increasingly using pint-size satellites sometimes no larger than a loaf of bread to gather data from low-Earth orbit, they have yet to apply the less-expensive small-satellite technology to observe physical phenomena far from terra firma. Jaime Esper, a technologist at NASA's Goddard Space Flight Center in Greenbelt, Maryland, however, is advancing a CubeSat concept that would give scientists that capability. Dubbed the CubeSat Application for Planetary Entry Missions (CAPE), the concept

involves the development of two modules: a service module that would propel the spacecraft to its celestial target and a separate planetary entry probe that could survive a rapid dive through the atmosphere of an extraterrestrial planet, all while reliably transmitting scientific and engineering data.

Esper and his team are planning to test the stability of a prototype entry vehicle—the Micro-Reentry Capsule (MIRCA)—this summer during a high-altitude balloon mission from Fort Sumner, New Mexico. *“The CAPE/MIRCA concept is like no other CubeSat mission,”* Esper



## The CAPE/MIRCA spacecraft, including the service module and entry probe, would weigh less than 11 pounds and measure no more than 4 inches on a side.

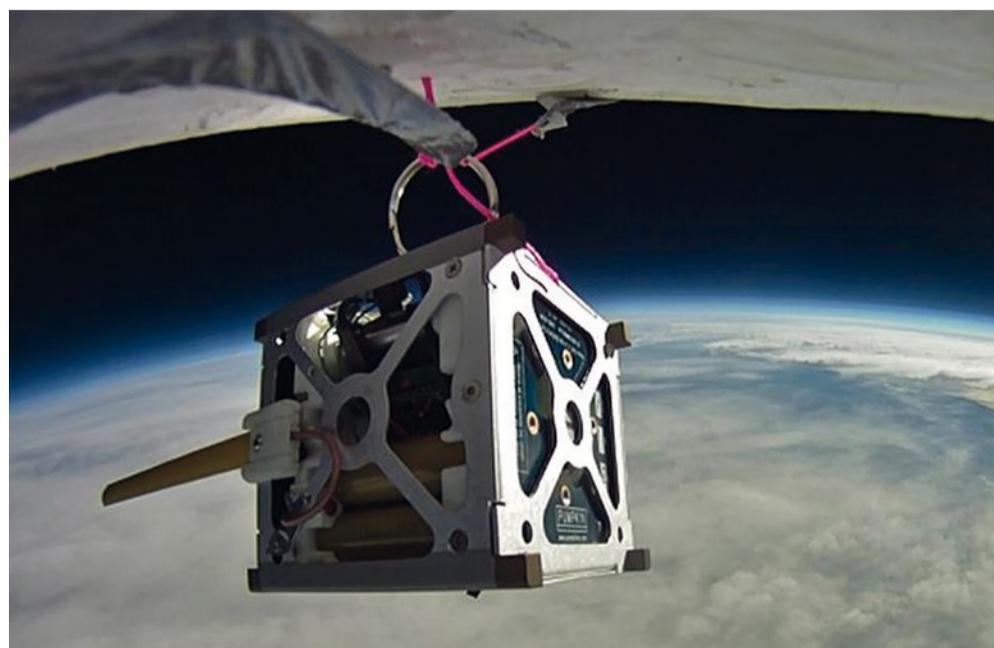
said. *“It goes the extra step in delivering a complete spacecraft for carrying out scientific investigations. We are the only researchers working on a concept like this.”*

Under his concept, the CAPE/MIRCA spacecraft, including the service module and entry probe, would weigh less than 11 pounds (4.9 kilograms) and measure no more than 4 inches (10.1 centimeters) on a side. After being ejected from a canister housed by its mother ship, the tiny spacecraft would unfurl its miniaturized solar panels or operate on internal battery power to begin its journey to another planetary body. Once it reached its destination, the sensor-loaded entry vehicle would separate from its service module and begin its descent through the target’s atmosphere. It would communicate atmospheric pressure, temperature, and composition data to the mother ship, which then would transmit the information back to Earth.

The beauty of CubeSats is their versatility. Because they are relatively inexpensive to build and deploy, scientists could conceivably launch multiple spacecraft for multi-point sampling — a capability currently not available with single planetary probes that are the NASA norm today. Esper would equip the MIRCA craft with accelerometers, gyros, thermal and pressure sensors, and radiometers, which measure specific gases; however, scientists could tailor the instrument package depending on the targets, Esper said.

### BALLOON FLIGHT TO TEST STABILITY

The first step in realizing the concept is demonstrating a prototype of the MIRCA design during a balloon mission this summer. According to the plan, the capsule, manufactured at NASA’s Wallops Flight Facility on Virginia’s



Eastern Shore, would be dropped from the balloon gondola at an altitude of about 18.6 miles (30 kilometers) to test the design’s aerodynamic stability and operational concept. During its free fall, MIRCA is expected to reach speeds of up to Mach 1, roughly the speed of sound. *“If I can demonstrate the entry vehicle, I then could attract potential partners to provide the rest of the vehicle,”* Esper said, referring to the service module, including propulsion and attitude-control subsystems. He added that the concept might be particularly attractive to universities and researchers with limited resources.

In addition to the balloon flight, Esper said he would like to drop the entry vehicle from the International Space Station perhaps as early as 2016 — a test that would expose the capsule to spaceflight and reentry heating conditions and further advance its technology-readiness level. *“The balloon drop of MIRCA will in itself mark the first time a CubeSat planetary entry capsule is flight tested, not only at Goddard, but anywhere else in the world,”* he said. *“That in turn enables new opportunities in planetary exploration not available to date, and represents a game-changing opportunity for Goddard.”*



# New SanDisk PCIe accelerators

Available in capacities from 1 to 6.4 TB, the new Fusion IO cards improve the price / performance by a factor of 4 and multiply the capacity by 2



**D**esigned for Big Data and Cloud Computing infrastructure, the new generation of PCIe Fusion IO Memory accelerators SanDisk integrate NAND flash memory and access acceleration software on virtual data storage layer (VSL). According to the manufacturer, the price / performance ratio is improved by a factor of 4 with a reduced price of 61% from the previous Fusion IO Drive 2 generation.

According to John Scaramuzzo, senior vice president and general manager, Enterprise Storage Solutions at SanDisk, *“At launch 8 years ago, the Fusion IO technology has fundamentally changed the expectations of data center performance and remains the reference for all other PCIe products by combining powerful technology Fusion IO Memory to SanDisk’s vertical integration strategy in the NAND flash memory.”* This offer supports the conclusions

of a recent report from Gartner that *“given the proliferation of virtual servers, virtual desktop infrastructure (VDI) and computer memory applications, to which is added the use essential analytical tools in real time, we have identified a growing interest in PCIe technology”*.

## ACCELERATE THE RECOVERY OF APPLICATIONS

By streamlining access to data, the new Fusion IO Memory cards increase performance of Microsoft SQL databases, Oracle and MySQL, but also new database technologies like Spark, Redis and MongoDB.

The application accelerators Fusion IO Memory portfolio include the following references:

- **SX350 and SX300 Series** - with a capacity of 1.25 to 6.4 TB, these cards are designed for mixed read / write workloads such as virtualization, databases, business intelligence (BI) and real-time financial data processing; intensive reading, such as web hosting, data mining, processing of seismic data, caching of content, 3D animation and CAD / CAM.
- **PX600 Series** - with a capacity of between 1 and 5.2 TB, the PX600 Series is designed for workloads in mixed use, such as virtualization, databases, business intelligence (BI) and treatment financial data in real time.
- **Mezzanine Series** - available on the latest HP Blade System, Gen9 and Gen8 servers, and new UCS B-Series blade servers from Cisco.
- **FlashSoft Offers** - New offerings combine the advantages of Fusion ioMemory flash devices and setting FlashSoft caching software regardless of the architecture or the storage operating system.

# BIG HPC REVIEW CONTEST

## HPC REVIEW

### Win an OCZ Saber 1000 SATA Enterprise SSD! (960 GB capacity model)

#### Saber 1000 Features

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



Click below to enter the contest

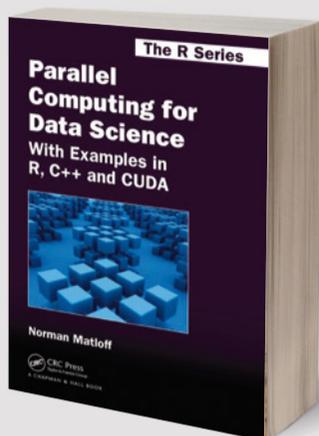
# I want to win!

The winner will be announced on April 30th, 2015

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



## BOOKS



## Parallel Computing for Data Science

This book one of the first parallel computing books to concentrate exclusively on parallel data structures, algorithms, software tools, and applications in data science. It includes examples not only from the classic « $n$  observations,  $p$  variables» matrix format but also from time

series, network graph models, and numerous other structures common in data science. The examples illustrate the range of issues encountered in parallel programming. With the main focus on computation, the book shows how to compute on three types of platforms: multicore systems, clusters, and graphics processing units (GPUs). It also discusses software packages that span more than one type of hardware and can be used from more than one type of programming language. Readers will find that the foundation established in this book will generalize well to other languages, such as Python and Julia. **Chapman and Hall/CRC, Norman Matloff, 328 pages, £38.99**

## GPGPU Programming for Maths and Science

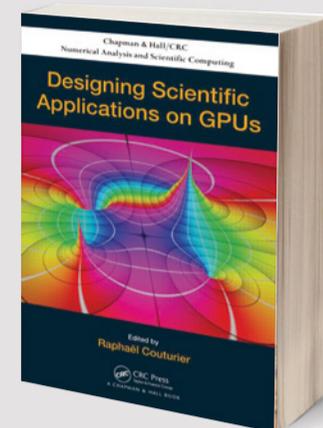
The book addresses programming on a graphics processing unit (GPU). It discusses many concepts of general purpose GPU (GPGPU) programming and presents practical examples in game programming and scientific programming. The core of the book focuses on the GPU from the perspective of Direct3D

11 and the High Level Shading Language (HLSL). This chapter covers drawing 3D objects; vertex, geometry, pixel, and compute shaders; input and output resources for shaders; copying data between CPU and GPU; configuring two or more GPUs to act as one; and IEEE floating-point support on a GPU. The book goes on to explore practical matters of programming a GPU, including code sharing among applications and performing basic tasks on the GPU. Focusing on mathematics, it next discusses vector and matrix algebra, rotations and quater-

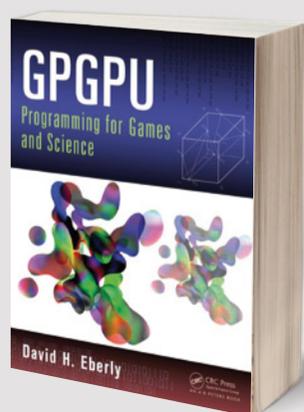
nions, and coordinate systems. The final chapter gives several sample GPGPU applications on advanced topics. **Chapman and Hall/CRC, David H. Eberly, 469 pages, £57.99**

## Designing Scientific Applications on GPUs

This books shows how to use GPUs for applications in scientific fields, from physics and mathematics to computer science and explains the methods for designing or porting your scientific



application on GPUs. It will improve your knowledge about image processing, numerical applications, methodology to design efficient applications, optimization methods, and much more. The first part of the book introduces the GPUs and Nvidia's CUDA programming model. The second part focuses on significant image processing applications on GPUs. The third part presents general methodologies for software development on GPUs and the fourth part describes the use of GPUs for addressing several optimization problems. The fifth part covers numerical applications, including obstacle problems, fluid simulation, and atomic physics models. The last part illustrates agent-based simulations, pseudorandom number generation, and the solution of large sparse linear systems for integer factorization. **Chapman and Hall/CRC, Raphael Couturier, 498 pages, £59.99**





## MOOCS

# Big Data in Education

Education is increasingly occurring online or in educational software, resulting in an explosion of data that can be used to improve educational effectiveness and support basic research on learning. In this course, you will learn how and when to use key methods for educational data mining and learning analytics on this data. You will learn about the methods being developed by researchers in the educational data mining, learning analytics, learning at scale, student modeling, and artificial intelligence in education communities, as well as standard data mining methods frequently applied to educational data. You will learn how to apply these methods, and when to apply them, as well as their strengths and weaknesses for different applications. The course will discuss how to use each method to answer education research questions and to drive intervention and improvement in educational software and systems. Methods will be covered both at a theoretical level, and in terms of how to apply and execute them using software

tools like RapidMiner. We will also discuss validity and generalizability; towards establishing how trustworthy and applicable the results of an analysis are. Some knowledge of either statistics, data mining, mathematical modeling, or algorithms is recommended. Experience with programming is not required. This course is comparable in difficulty to the first course in the Masters in Learning Analytics at Teachers College Columbia University, though it does not go into the same depth as that course. **Free.**

**Starts** July 1, 2015

**Level:** Advanced

**Length:** 8 weeks

**Effort:** 6-12 hours/week

**Subject:** Statistics & Data Analysis

**Institution:** ColumbiaX

**Languages:** English

**Video Transcripts:** English

<https://www.edx.org/course/big-data-education-columbiax-bde1x#!>

## Introduction to Big Data with Apache Spark

Organizations use their data for decision support and to build data-intensive products and services, such as recommendation, prediction, and diagnostic systems. The collection of skills required by organizations to support these functions has been grouped under the term Data Science. This course will attempt to articulate the expected output of Data Scientists and then teach students how to use PySpark (part of Apache Spark) to deliver against these expectations. The course

assignments include Log Mining, Textual Entity Recognition, Collaborative Filtering exercises that teach students how to manipulate data sets using parallel processing with PySpark. This course covers advanced undergraduate-level material. It requires a programming background and experience with Python (or the ability to learn it quickly). All exercises will use PySpark (part of Apache Spark), but previous experience with Spark or distributed computing is NOT required. Students should take

this Python mini-quiz before the course and take this Python mini-course if they need to learn Python or refresh their Python knowledge. **Free.**

**Starts** June 1, 2015

**Level:** Intermediate

**Length:** 5 weeks

**Effort:** 5 - 7 hours per week

**Subject:** Computer Science

**Institution:** UC BerkeleyX

**Languages:** English

**Video Transcripts:** English

<https://www.edx.org/course/introduction-big-data-apache-spark-uc-berkeleyx-cs100-1x>



## MOOCS

# Cyberwar, Surveillance and Security

Once heralded as the ultimate vehicle for open communication and self-expression, the internet is rapidly becoming a globally networked surveillance device. Serious threats to national security, combined with the seemingly endless capacity of digital processing and storage, have led to levels of data capture and 24/7 monitoring of individuals' activity that were unimaginable even a decade ago. With resistance to such practices rising, this course will equip you to take an active part in the debate. You will gain a broad understanding of the competing tensions of the laws related to national security and personal and commercial privacy

in the post-Snowden online environment, and grasp the looming consequences of this battle for peace, sovereignty, human rights and the internet itself. **Free.**

**Starts** May 21, 2015

**Level:** Introductory

**Length:** 6 weeks

**Effort:** 2-3 hours/week

**Subject:** Law

**Institution:** AdelaideX

**Languages:** English

**Video Transcripts:** English

<https://www.edx.org/course/cyberwar-surveillance-security-adelaidex-cyber101x#!>

## Big Data in Education

Education is increasingly occurring online or in educational software, resulting in an explosion of data that can be used to improve educational effectiveness and support basic research on learning. In this course, you will learn how and when to use key methods for educational data mining and learning analytics on this data. You will learn about the methods being developed by researchers in the educational data mining, learning analytics, learning at scale, student modeling, and artificial intelligence in education communities, as well as standard data mining methods frequently applied to educational data. You will learn how to apply these methods,

and when to apply them, as well as their strengths and weaknesses for different applications. The course will discuss how to use each method to answer education research questions and to drive intervention and improvement in educational software and systems. Methods will be covered both at a theoretical level, and in terms of how to apply and execute them using software tools like RapidMiner. We will also discuss validity and generalizability; towards establishing how trustworthy and applicable the results of an analysis are. Some knowledge of either statistics, data mining, mathematical modeling, or algorithms is recommended. Experience with programming

is not required. This course is comparable in difficulty to the first course in the Masters in Learning Analytics at Teachers College Columbia University, though it does not go into the same depth as that course.

**Free.**

**Starts** July 1, 2015

**Level:** Advanced

**Length:** 8 weeks

**Effort:** 6-12 hours/week

**Subject:** Statistics & Data Analysis

**Institution:** ColumbiaX

**Languages:** English

**Video Transcripts:** English

<https://www.edx.org/course/big-data-education-columbiabde1x#!>



# THE HPC OBSERVATORY



## \$44 Billion

This is the projection of the worldwide turnover of HPC in 2020. Market Research Media research firm expects the area of supercomputing will grow an average of 8.3% annually to reach \$ 44 billion in 2020 . This sector will also generate 220 billion dollars in cumulative sales over the period 2015-2020. Source : <http://www.marketresearchmedia.com/>

### THE TOP 3 OF THE TOP 500

- 1 Tianhe-2**  
 National Supercomputing Center in Canton:  
**33863/54902 TFlops** Manufacturer:  
 NUDT, architecture Xeon E5-2692 Xeon Phi  
 31S1P +, TH Express-2
- 2 Titan**  
 Oak Ridge National Laboratory, USA:  
**17590/27113 TFlops** Manufacturer: Cray  
 XK7, architecture Opteron 6274 + Nvidia  
 Tesla K20X Cray Gemini Interconnect
- 3 Sequoia**  
 Lawrence Livermore National Laboratory,  
 USA: **17173/20133 TFlops** Manufacturer:  
 IBM Blue Gene / Q architecture  
 PowerPC A2

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

### THE TOP 3 GREEN 500

- 1 5271,81 MFlops/W**  
**GSI Heimboltz Center (Germany)**  
 57.15 kilowatts consumption
- 2 4945,63 MFlops/W**  
**High Energy Accelerator KEK (Japan)**  
 37.83 kilowatts consumption
- 3 447,58 MFlops/W**  
**GSIC Center, Tokyo Institute of Technology (Japan)**  
 35.39 kilowatts consumption

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

**INNOVATION BY SIMULATION**  
SIMULER POUR INNOVER

# Forum Teratec

le rendez-vous international  
*The international meeting*

**HPCBIG DATASIMULATION**

23 & 24 juin/June 2015

ECOLE POLYTECHNIQUE  
PALAISEAU - FRANCE

[www.teratec.eu](http://www.teratec.eu)



Platinum Sponsors :



Gold Sponsors :



Silver Sponsors :





Take large amounts of data processed through Big Data platforms. Throw in a dash of Business Intelligence to add visualization and analysis capabilities. Make them apply by machines, you get the basics of a Machine Learning platform.

# MACHINE LEARNING



THE FUTURE  
OF HPC?



## Machine Learning is the latest hype in modern computing. One of the most commonly accepted definition refers to computer systems that improve in gaining experience from the analyzed data.

**M**achine Learning is the latest hype in modern computing. One of the most commonly accepted definition refers to computer systems that improve in gaining experience from the analyzed data. To reach this goal, researchers, statisticians and data scientists define software models capable of processing and analyzing large amounts of data through dedicated machines. The final objective: to establish predictive models based on the data correlation and the observed trends.

Because the conclusion is clear: this analytical work done by humans is already showing its limits, unable to cope with the flood of data that the Internet of Things promises to carry. For we must then be able to process and analyze massive amounts of data to identify emerging trends and apply the necessary measures whenever necessary. Or to identify discriminatory elements, or to detect weak signals that could influence a decision. In an increasingly digital-dependent world, the stakes are high!

### **TOMORROW, THE INTERNET OF THINGS**

Since the Internet of Things will generate 90% of data tomorrow, it is an appropriate response in the form of algorithms enabling machines to analyze these data of a new genre in unprecedented quantities. This is where the Machine Learning makes sense. But the Internet of Things is not the only reason for machine learning. Big Data is equally important. We already have a global amount of 8-Zetta Bytes of data worldwide. Some of this data is made up of static data, but the dynamic share will grow in coming years, to reach according to IDC, 40-Zetta Bytes in 2020.

### **A SOFTWARE + HARDWARE EQUATION FOR AN EFFICIENT MACHINE LEARNING SERVICE**

Developments in Machine Learning platforms are based on two pillars. On one hand by creating algorithms more efficient that can handle, filter, and analyze the data, and on the other hand to choose the best equipment capable of storing, reading, processing, analyzing and writing data. A set of algorithms, processors and storage devices thus provide the foundation of a system of performance Machine Learning.

### **THREE ASPECTS TO SOLVE A PROBLEM MACHINE LEARNING**

As Dennis DeCoste, Machine Learning Director at eBay stated, a Machine Learning problem should be treated in three ways: random algorithms capable of growing to scale (scalable), a choice of hardware capable of meeting the requirements to achieve an effective treatment, and a compilation model broken down into small and efficient modules tested and refined individually rather than a complex algorithm to optimize.

### **SOME CONCRETE EXAMPLES OF MACHINE LEARNING**

How useful can Machine Learning be in everyday life? Here are some answers.

**Social Networks:** this is a rich ground for extremely complex information to be processed for human teams given the amount of daily data. Facebook and Twitter are a network of human sensors from which it is possible to extrapolate trends, as demonstrated by IBM's research teams. So much so that they prefer



to call them «social sensor» rather than social networks. On several occasions, the influence of these social sensors was demonstrated in scholarship and public health (flu). To illustrate, Twitter alone represents 250 million tweets per day. The information considered by IBM's Cognos Consumer Insights engine on each Tweet IBM, are its virality (the number of retweets of a given tweet), seniority (based on date and time of publication), and correlated with the number of tweets' related content (based on a lexical analysis of its content).

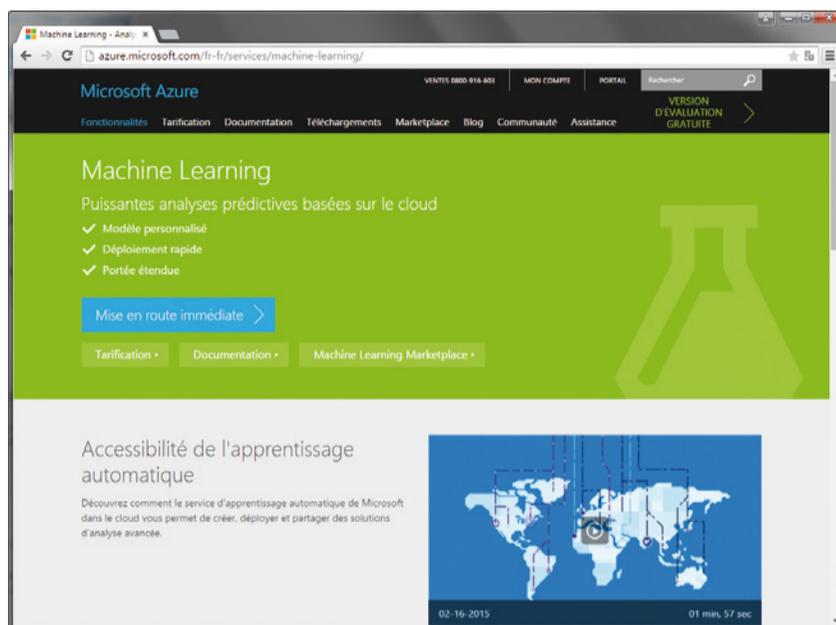
**Likes:** a network such as TripAdvisor used to rank restaurants and hotels worldwide generates an enormous amount of data. Machine Learning allows to treat such huge amounts of data by applying a semantic processing to separate the positive feedback from the negative ones. Then it is possible to extend or limit them on a given geographic area to apprehend the top ranked ones. Apply this treatment to another geographic area (city, region or country) and you have a multiple comparison basis for the list of most popular destinations, for example.

**Stock Exchange and Finance:** IBM demonstrated the causal link between potentially fraudulent financial transactions by examining the history of raw financial data on all global monetary exchanges, correlated with the financial information provided by the media (Bloomberg). The two types of data were the subject of an in-depth analysis by IBM's Watson Automated Integration engine, which helped pinpoint the truthfulness on the observed unusual financial transactions corresponding to a real event or not.

**Sale predictions :** this type of analysis is done to correlate the data pertaining to acts of purchase. Type of purchases (food, clothing, entertainment), products purchased, value of the purchased product, purchasing or seasonal recurrence are all elements that can be crossed to infinity to identify surging or drop-

ping trends, buying habits of consumers and therefore, increase inventory forecasts accuracy.

## MICROSOFT'S AZURE ML INITIATIVE



Microsoft Azure Machine Learning is an online service that developers can use to build predictive analytic models (using data sets from multiple data sources), and then easily deploy these models as Web services. Azure ML Studio provides features applicable to many flow scenarios for building predictive models, provides easy access to data sources, and includes exploration data functions, visualization tools and application of ready to use ML algorithms. The ML Azure platform includes a hundred of ready to use scenarios to solve a wide range of scientific problems. Moreover this scenario library can be extended through R language modules, in order to respond to situations not covered by the standard functions.

## AN EXTENSIBLE ML PLATFORM

Once the module R language finalized, it can be shared with colleagues via GitHub. These modules can be used on non-standard data processing processes, such as managing data formats specific to a given area, increase flexible data transformation, or dedicated to data extraction. Better yet, an R module may incorporate any of the hundreds of pre-installed scripts on Azure ML, and even incorporate them into your own module.



Here is an implementation example: creating an R module that aggregates data in JSON format and adapts to Azure ML dataset format. Such a module is composed of three parts:

- An R code file defines what the module must perform;
- Optional accompanying files (configuration files or scripts R)
- An XML file that determines the inputs, outputs and module parameters.

One could say that the XML file is the module skeleton, and the R code is its muscle. The module accepts input data in JSON format, and transforms it to generate a flat dataset. An accepted parameter is a string that specifies a zero replacement value. The corresponding R script is:

```
parse_json.R:
parse_json <- function(data_in, nullvalue=>»NA») {
  library(RJSONIO)
  library(plyr)
  data_out <- ldply(fromJSON(as.character(data_in[1,1]),nullValue=nullvalue,simplify=TRUE))
  return(data_out)
}
```

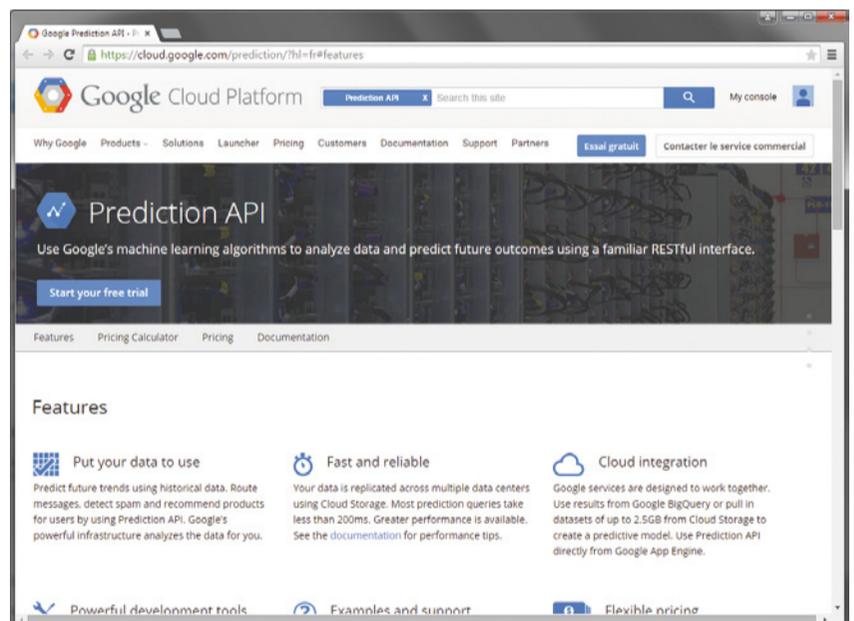
The XML description file defines the name of the following modules, the R function call starts the execution of the module, the input and output datasets, and any associated parameters.

```
parse_json.xml:
<Module name=>»Parse JSON Strings»>
<Owner>AzureML User</Owner>
  <Description>This is my module description. </
Description>
  <Language name=>»R» sourceFile=>»parse_
json.R» entryPoint=>»parse_json»/>
  <Ports>
    <Output id=>»data_out» name=>»Parsed
dataset» type=>»DataTable»>
    <Description>Combined Data</Description>
  </Output>
  <Input id=>»data_in» name=>»JSON formatted
dataset» type=>»DataTable»>
```

```
<Description>Input dataset</
Description>
  </Input>
</Ports>
<Arguments>
  <Arg id=>»nullvalue» name=>»Null replacement
value» type=>»string» isOptional = <true»>
  <Description>Value used to replace JSON
null value</Description>
  </Arg>
</Arguments>
</Module>
```

To add this module to your Azure ML account, simply create a zip file containing all the files and upload it by clicking on the menu + NEW> in your Azure Module ML Studio work environment. Once uploaded, your module will appear in the «custom» category of the range of modules, along with pre-installed modules.

## THE GOOGLE PREDICTION API



Google provides developers and researchers its Google Cloud platform within which lies the Prediction API. According to Google, this API allows already to achieve objectives such as:

- Predict whether a user will enjoy a movie based on movies he likes
- Categorize emails as spam or not
- Analyze feedback to determine if your products are popular or not
- Predict the expenses of users based on their consumption habits



## A prediction model is only as good as the data quality with which it is fed, and which will calibrate the model in order to return relevant information in any application.

Google makes available to users of this API two sites that exploit it, one for spam detection and the second for predicting film preferences. Prediction API is accessible via a RESTful interface. Google provides several libraries and scripts to access the API. The latter provides functions of pattern-matching and Machine Learning capabilities from data examples that you have to provide. The model supplied with the data evolves to become a trained model capable of responding to requests.

The number of direct call to the API functions Prediction version 1.6 is rather small, and is limited to:

- `prediction.hostedmodels.predict`  
Provides input data and requests data output from a hosted model.
- `prediction.trainedmodels.analyze`  
Gets the model analysis and data on which the model was trained
- `prediction.trainedmodels.delete`  
Deletes a trained model
- `prediction.trainedmodels.get`  
Checks the status of training your model
- `prediction.trainedmodels.insert`  
Trains a model of the Prediction API
- `List.prediction.trainedmodels.list`  
List of available models
- `prediction.trainedmodels.predict`  
Provides the id of a model and request a prediction
- `prediction.trainedmodels.update`  
Adds new data to a trained model

### THE MAIN STEPS IN THE USE OF THE PREDICTION API

A prediction model is only as good as the data quality with which it is fed, and which will calibrate the model in order to return relevant information in any application.

1 / **Create your training data.** You must create data to power your model adapted to the questions you need answers for. This is the most critical and complex step. You can also use a pre-trained model from the hosted template gallery and go to Step 4: Send a prediction query.

2 / **Send your workout data** to Google Cloud Storage using standard Google Cloud Storage tools.

3 / **Train the model with your data.** The Prediction API will train your model with your data of which you indicate the location. You must query the Prediction API to check it when the training is finished.

4 / **Send a prediction query.** Once the training phase is over, you can send a request to your model via the Prediction API. You will get back a response as a numerical value or as text according to data you have fed your model.

5 / *(optional)* **Send additional data to your model.** If you have a steady stream of new data consistent with the data used to initially feed your model, you will improve the relevance of results. This helps improve your model as you add new data to it.

The two most important aspects of the use of Prediction API is to create and structure the data which is supplied with the model, and to formulate a relevant question that the API is able to respond.

### RELEVANT PREDICTIONS?

The term «prediction» may sound misleading because the Prediction API is only able to accomplish two specific tasks:

- As a new element, to predict a numerical value based on similar values in the data used to train the model.
- As a new item, select a category that best describes it, given a set of similar items for inclusion in the data used in the lead.



**The data volume of a global platform such as eBay is enough to make your head spin: it amounts to 100 million users (buyers and sellers) and 10 million items per day that are processed and analyzed to determine their usage habits: research, clicks, wishes, and auction purchases.**

Both tasks can seem limited. However, if you formulate your request carefully and choose your data accordingly, the Google Prediction API will let inferred preferences of a user and keep the projection of future correlation values consistent with the data used to train your model.

#### **EBAY: 100 MILLION QUERIES PER DAY**

The data volume of a global platform such as eBay is enough to make your head spin: it amounts to 100 million users (buyers and sellers) and 10 million items per day that are processed and analyzed to determine their usage habits: research, clicks, wishes, and auction purchases. Regarding the object data, it is the price, titles, descriptions and images that are processed, with a history of several billion objects.

#### **THE LEARNING MACHINE TO ESTABLISH BASIC PROPERTIES**

These analysis for a given family of objects, the intrinsic properties. An aquarium is so defined by characteristics such as length, height, width, but also its volume, weight, accessories, etc. This allows to establish a complete taxonomy. And also to determine the distribution thus, 87% of the categories contain 17% of the items on eBay, and conversely, 1% of the categories contain 52% level.

#### **THE SEARCH ENGINE, EBAY'S QUEEN APP**

The application on eBay that attracts most of the Machine Learning developments is the search feature. Full-text seeking requires a

semantic decomposition, in addition to correlated criteria as the cheapest or closest item, etc. Behind the scenes this result is achieved by applying a temporal data mining algorithm based on Hadoop called Mobius. This massively parallelized mechanism can meet the 100 million daily queries with a very good response time of around a second.

#### **SCALABILITY? EBAY IS WORKING ON IT ...**

eBay's Learning Machine team does not intend to stop there, and is already thinking of implementing kernel, hashing and random projection algorithms in order to be able to cover more users and objects. This is not the only element needing to be adjusted to satisfy future needs. The ML teams have determined that the use of GPU acceleration solution allows improving performance by a factor of 20, by replacing 12 CPU cores at 3.5 GHz (ie 168 gigaflops) by an NVIDIA solution 690GTX capable of delivering 5.5 TeraFlops.

As we see, Machine Learning is a discipline that depends as much on algorithms than the quality of the analyzed data. And scalability in data volumes of a greater magnitude can only be achieved through hardware platforms making use of the most powerful CPUs and GPUs available to meet the increasing requirements of this new discipline. **RAMON LAFLEUR**



# 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.



# Lab Review

## How do 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



# HP ConvergedSystem 240 HC StoreVirtual



**T**hroughout its adoption, virtualization brings profound changes in the infrastructure of companies facing problems of development, consolidation and replacement of their IT. Hyperconvergent systems are one answer to these questions and can help companies simplify their IT operations. This is the target of HP's ConvergedSystem server 200-HC. It integrates in a single 2U rack a storage solution defined by software, monitoring tools and integration with VMware to create a fully virtualized environment and ready to use. The guiding principle of this solution is to be simple and open. It integrates HP OneView InstantOn technology, HP OneView for VMware vCenter and HP VSA StoreVirtual to ensure simple installation, full control and high availability storage servers.

## AN HYPERCONVERGED OFFER FOR SIMPLICITY

This offer is intended for midsize businesses, remote offices and, more generally, IT organizations wishing to adopt a modular approach to their IT infrastructure. Besides a unified management through vCenter hp OneView, the HC 240 ConvergedSystem StoreVirtual solution facilitates the implementation of a disaster recovery plan by providing network and storage redundancy. The integration within a single appliance ensures optimal response times in VDI environments to provide a fast start time and rapid response for virtualized workloads, adapted to the strong demand of the activity, as well as a simplified management and high availability of applications.





### **(ALMOST) NO EXPERIENCE REQUIRED**

Operational efficiency of the platform is the result of the convergence of StoreVirtual HP management tools, HP ProLiant servers and HP OneView and VMware vSphere. The skills and time required to manage separate servers, hypervisors, network and storage can discourage companies not yet equipped or unaware of the benefits of virtualization. These constraints are very much reduced in the case of the HC-240, a system designed to simplify the control of a virtual infrastructure and also to set up applications and virtual desktops. With OneView HP InstantOn, virtual machines can be deployed in minutes. Even better, the feature set is usable by IT staff even without advanced experience in storage. As Nariman Teymourian highlights, vice president and general manager of Converged Systems, HP, *“Small and medium enterprises are looking for ways to streamline its IT operations and deliver applications and services faster. HP ConvergedSystem 200-HC StoreVirtual enables businesses to accelerate response time and realize business opportunities.”*

### **A PRECONFIGURED AND PREINSTALLED SERVER IN UNDER 15 MINUTES !**

HP ConvergedSystem StoreVirtual HC-200 is a fully preconfigured and preinstalled cluster. To implement just type an IP address, the identification data and VMware vSphere license numbers. Once this is done, the system is ready. The promise of the manufacturer of a period of start-up less than 15 minutes is respected.

### **SEAMLESS VMWARE INTEGRATION**

To simplify administration, users can use HP VMware vCenter OneView for offering centralized server management, storage and virtual machines from VMware vCenter Server. This guarantees a single management interface without special knowledge in terms of ser-

vers, storage, networking and virtualization. Once the server is running, the administrator can provision virtual machines directly from the VMware vCenter Server, manage hyper-converged and traditional systems of network from a single console to OneView HP VMware vCenter, replicating a system hyper-converge or combine it with a server to get a recovery system flexible disaster, or optionally use VMware vRealize Operations to monitor the status of all items.

### **A CLUSTER OF CLUSTER FOR SCALABILITY PURPOSES**

The HC-200 server is itself a 4-node cluster, it may itself be grouped into a super 32 node cluster to expand the capacity and performance of the whole. These multiple clusters are managed from a single console. In addition, each additional system adds power and storage volume of four independent servers consolidated into one 2U chassis.

### **CONCLUSION**

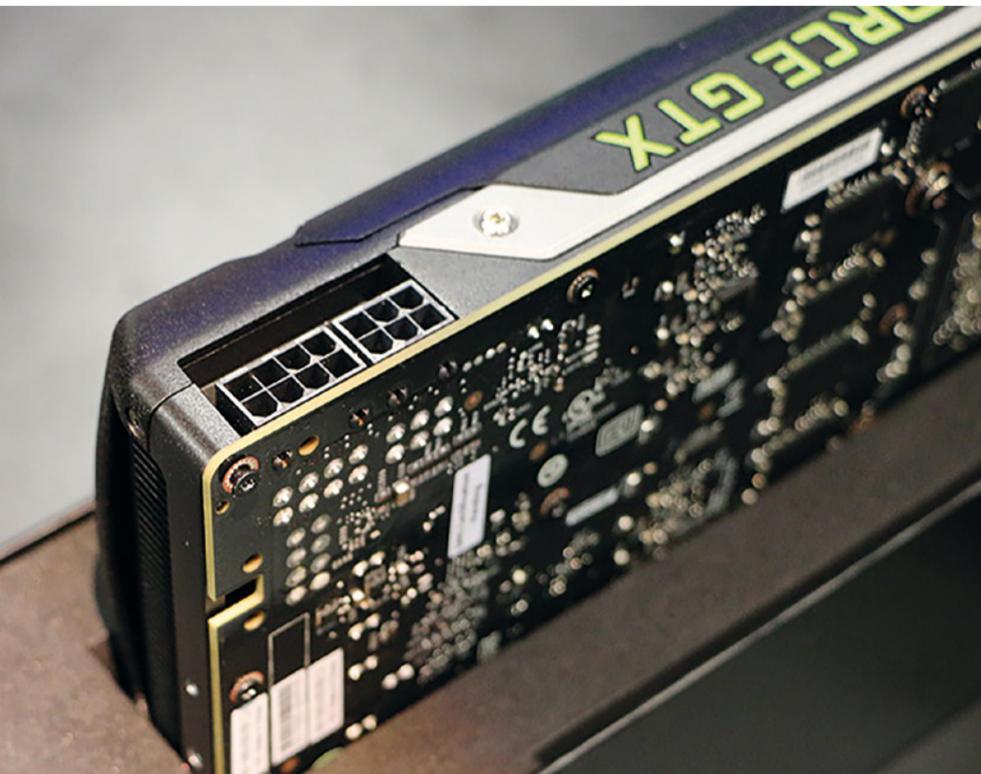
The HC-200 ConvergedSystem StoreVirtual keeps its promises and is characterized by a centralized administration regardless of the number of nodes and virtual machines and the ability to keep applications online without interruption. Add to this the possibility of migration or replication of data volumes, again without interruption, including x86 servers to even non converged long as they run any application and hypervisor along with hp's Storevirtual VSA. HP's ConvergedSystem HC-240 is available in two versions: with 512 GB of memory, 24 SAS 2.5-inch or 1.2 TB with 1 TB of hybrid memory, 16 disk 2; 5-inch SAS 1.2TB 8 and 400 GB SSD. **RAMON LAFLEUR**

HPC Bench  
Indice total



**10 800**

**Four blade servers in & 2U rack running VMware vSphere and StoreVirtual. 64 CPU cores at 2GHz. 512 GB Memory, 8x 1GbE and 10GbE network interfaces. 24 x 1.2 TB SAS 2.5 inches disks.**



SOURCE VIDEOCARDZ.COM

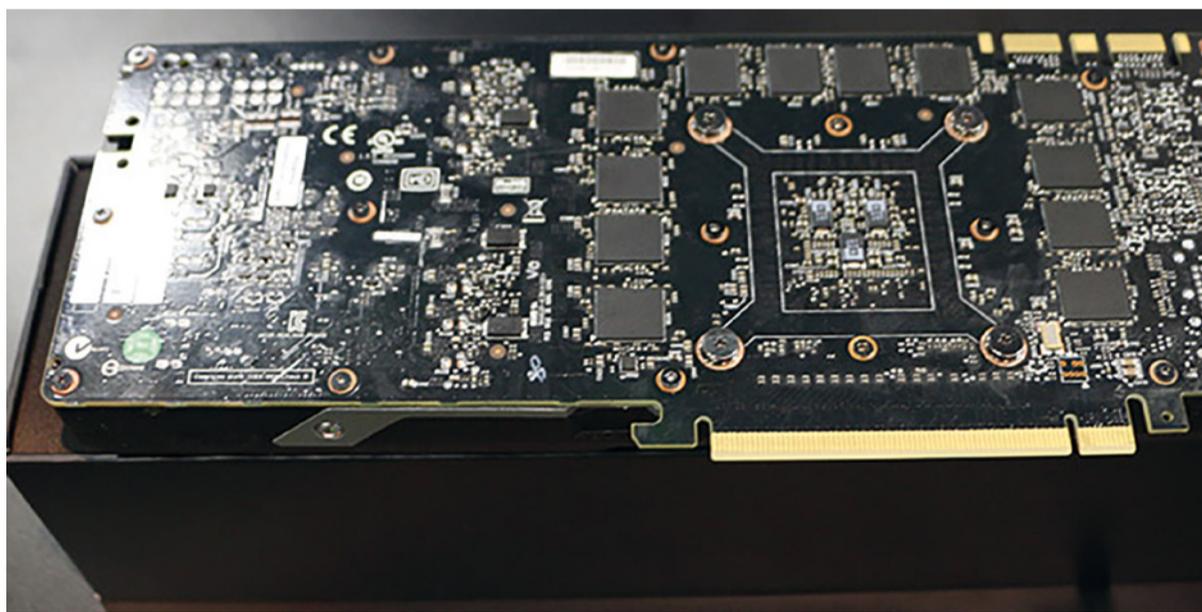
# NVIDIA GTX Titan X: + 37%

**Compared to the current flagship of the NVIDIA GeForce range, the GTX 980, the performance gain of the GTX Titan X is significant.**

**T**he graphics card GTX Titan X, presented to the public at the GDC (Game Developers Conference) 2015, should embody NVIDIA's know-how regarding video acceleration and thereby become the manufacturer's technology showcase at a time when direct competition, specifically AMD, is not far behind.

The GTX Titan X is based on the GM200 (Maxwell architecture), with 3072 CUDA cores, 192 TMU and a maximum of 12 GB of 384-bit GDDR5 memory. This card features two power connectors, like the GK-110 PGPU accelerators, which would place it in an maximum TDP of 250 Watts. It could even be less, given the improved energy efficiency of Maxwell over Kepler.

All these conjectures aside, the card tested by VideoCardz ([www.videocardz.com](http://www Videocardz.com)) and ChipHell ([www.chiphell.com](http://www.chiphell.com)) achieves outstanding results. Our colleagues note a 3DMark 11 Performance at 22 903 points, 35% better than a GTX 980, which is the current top of the line GeForce. Other benchmarks (Extreme,

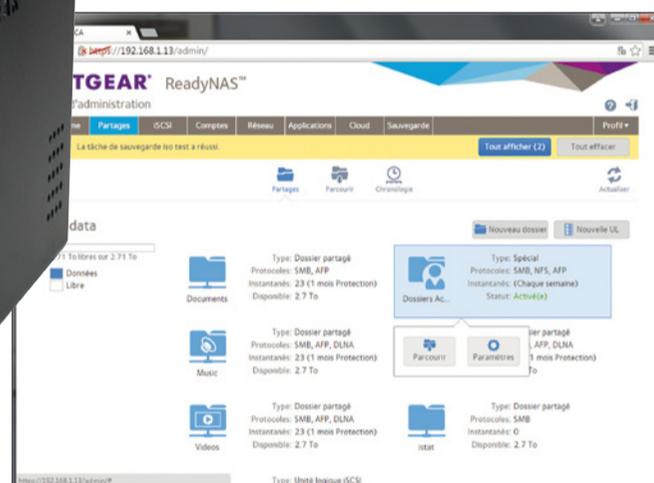
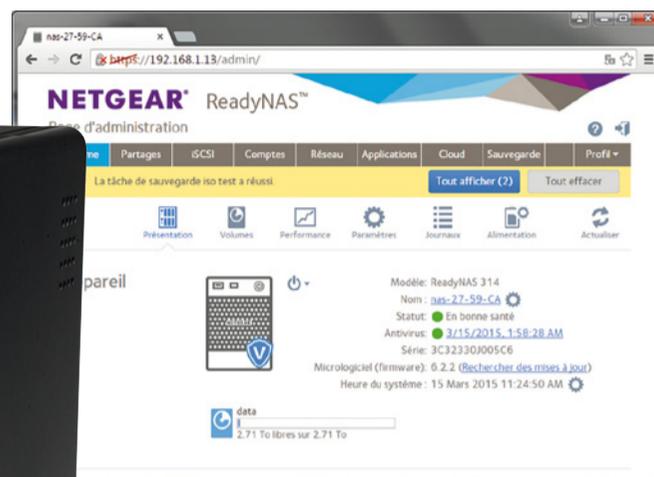


Extreme Firestrike and Firestrike) give similar speedups, always above 30% - of what, again, is probably a pre-production sample with drivers still developing.

As soon as AMD announced the upcoming release of the liquid-cooled Radeon X R9 390, which is believed to be the first to use HBM (High Bandwidth Memory) technology, NVIDIA has set the performance bar pretty high - which remains to be seen once production boards of both models become readily available. This generational leap is important, if only because this architecture is shared with the Quadro M6000. **FRANÇOIS REY**



# Netgear ReadyNAS 314



**T**he NAS market segment has been very prolific lately. So much so that sometimes becomes difficult to differentiate them, since they all share a very similar functional scope. Given the maturity of these products, they all have in common complete network management functions and comprehensive Raid options. On a closer look, however, some manufacturers are ahead of their own market. This is what we have found out while evaluating Netgear's ReadyNAS 314 NAS, an innovative model in many aspects. It is sold alone or preconfigured in 4-, 8 and 12TB capacities.

## EASY TO START

Designed for workgroups and businesses up to 50 users, the ReadyNAS 314 has a compact housing with 4 bays and an LCD screen. Hosted by

the manufacturer's own ReadyNAS OS 6.2.2 operating system, this model has a built-in assistant that guides you at start-up time on operating parameters, making it ready to use in minutes. What is striking is its full HTML interface, particularly fluid and intuitive. Ergonomics are thoroughly thought out in terms of navigation, mouseover help balloons and online documentation, so that one is never lost. The main menu bar features the major sections: System, Sharing, iScsi, accounts, network; Applications, Cloud and Backup. There is even a real-time antivirus available, that needs to be activated manually however.

## DATA PROTECTION ABOVE ALL

Another original feature of this model, related to data security, is the BTRFS file management system, much more advanced than the



ext3 or ext4 used by the vast majority of competing models. BTRFS (for B-Tree File System) has multiple integrated checksum of data and metadata mechanisms to ensure their integrity. Added to this is the possibility of activating unlimited snapshots. A replication mechanism locally or on a remote NAS (enough to establish a mini disaster recovery system) allows programming of regular backups. By default the initial backup is complete, the following being incremental, but you can opt for a full backup every time or every 1, 2, 3 or 4 weeks.

### SCALABILITY GUARANTEE

The scalability of the NAS is guaranteed by X-RAID2, a system that automatically takes into account the addition of disks that are configured transparently without hassle or loss of data. All you need is to swap out a disk with a larger one and the NAS takes care of the auto repartitioning and synchronization of the disks. If you need more storage space an eSATA port allows you to connect a 5-bay expansion enclosure (EDA500), which extends the total capacity up to 54 TB when using 6 TB disks. Scalability is also applicable on the software side. The ReadyNAS 314 may be functionally extended through a catalog of about 75 applications (php environments, bittorrent clients, Drupal and Joomla content management systems, video surveillance, MySQL, etc.), directly installable from the interface.

### MANY CLOUD FUNCTIONS

Netgear provides its users several cloud services, which can be activated from the NAS and accessed from the manufacturer's web portal. ReadyNAS Remote allows accessing the NAS content from any location and device via a secure VPN link. Mobile applications for Android and iOS also allow you to access

the contents of the NAS. ReadyCLOUD adds automatic synchronization of files from a PC or Mac and the possibility of sharing a file or a folder with a third party with a simple email link. Finally, ReadyNAS Replicate is an advanced backup and data synchronization mechanism, with extended settings for speed, capacity and backup mode (file or block).

### ISCSI AND PERFORMANCE

The iScsi volume creation function is very complete, and offers several interesting options: instant activation (every hour, day or week), an anti-bitrot copy-on-write mechanism, and even data compression (if the volume is declared dynamic). The result is immediately available for iSCSI LUN virtualization platforms like VMware vSphere or Microsoft Hyper-V. Performance is among the best we have measured with a 111 MB / second throughput (on a theoretical maximum of 120 over a Gigabit Ethernet link). The cooling of the NAS is done by a large fan, helping to bring down the operating noise to an acceptable level. The guarantee of the ReadyNAS 314 is five years on the discs and the unit with D+1 on site replacement and lifetime technical support.

### CONCLUSION

With its robust BTRFS file management system, fluid and ergonomic HTML 5 interface, many cloud features and built-in replication functions and scalability path, this model is in many key aspects ahead of its competitors. The many innovative features of the ReadyNAS 314 make this model a true innovator that shows what a next generation NAS should be... except that it is available today. This is a secure and scalable investment for companies seeking to build a NAS backup infrastructure.

**JOSCELYN FLORES**

---

**Specifications:** 4 Drive bays 4, expandable via eSata 5-bay expansion cabinet via optional EDA500 bays. **Processor :** Intel Atom Dual Core 2.1 GHz. **Drive bays :** SATA / SSD drives in 2.5 and 3.5 inches. 2 eSATA ports with a shared USB port. **Maximum capacity** 16 TB (56 TB with EDA500 box). 2 Gigabit Ethernet ports. **Price:** € 565 naked (€ 869 for 4TB, 8TB to € 917, € 1,353 in 12TB)



# 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



# How to eliminate storage bottlenecks in a data growing world

**Given the colossal avalanche of data, how to handle such volume profitably? And how to maintain or improve the performance of data storage? By eliminating bottlenecks. Here are some tips.**



## **PAVEL THOMAS**

*About the Author*  
is director of EMEA sales,  
storage products, Avago  
Technologies

## **ELIMINATION OF DATA STORAGE BOTTLENECKS**

The data deluge is not ready to end, especially with the surge of demanding applications such as the value and profitability analysis programs, social media, video distribution and grid computing to name a few examples. This data growth presents many challenges for system architects. How to handle such volume profitably data? And perhaps most important: how to maintain or improve the performance of data storage?

In fact, there are very few who are not affected by the data deluge. Even home users generate more video and high quality digital pictures which occupy considerable space on PCs and smartphones. This is one of the reasons why social networks such as Facebook are generating half a petabyte of data every single day (according to the notes of Facebook engineers in 2012).

And this increase in the quantity of data is accompanied by a longer access time. This article describes the ways in which the third-generation SAS technology eliminates the various data storage bottlenecks.

## **UNDERSTANDING WHERE BOTTLENECKS OCCUR**

When designing data storage systems, it is essential to understand where bottlenecks occur. Note also that these bottlenecks are changing with every new technology, along the data storage path. The three most important factors that affect performance storage are the server-side PCIe® bus (Peripheral Component Interconnect Express), the solution deployed by SAS HBAs (Host Bus Adapters) and expansion dispatchers whose interface can be either SAS or SATA (Serial Advanced Technology Attachment). SAS technology is now in its third generation, and performance doubled with each generation, the 3 Gb/s of the beginning has evolved to 6 Gb/s, and nowadays to 12 Gb/s. The SAS System, like PCIe, uses channels and high-performance storage systems typically combine multiple SAS ways to support the most important data rates. The HDDs (Hard Disk Drives) and SSD (Solid State Disks) are capable of supporting either SAS or SATA storage interfaces.



**The main problem when it comes to move to third-generation 12 Gb/s SAS technology, is well known: it is that of investment protection. Most organizations have invested heavily in SAS disks and want to preserve that investment as they pass the SAS 12 Gb/s barrier.**

### STORAGE SYSTEMS PERFORMANCE GUIDELINES

When designing a high-performance storage system, it is necessary to understand the performance limits of each. Typical yields of different types and generations of the three main elements discussed in this article are summarized in Table 1.

<b>HDD (6 Gbs/s SAS)</b>	230 MB/s
<b>SSD (6 Gb/s SAS)</b>	550 MB/s
<b>6 Gb/s SAS x4</b>	2 200 MB/s
<b>12 Gb/s SAS</b>	4 400 MB/s
<b>PCIe 2.x</b>	3 200 MB/s
<b>PCIe 3.0</b>	6 400 MB/s

#### Example of typical efficiency of different types of drives and SAS and PCIe generations

Table 2 summarizes some typical configurations and shows the location of the bottleneck when the system has a «suitable amount» of hard disks (slower elements of the system). As reported, to support more disks (for purposes of capacity) more PCIe, and / or more SAS channels should be used. If you look at it in a different way, in systems that have small numbers of records, the low cumulative return

Configuration	Bottleneck (MB/s)	# of HDDs	# of SSDs
<b>6 Gb/s SAS x4 / PCIe 2.x</b>	SAS (2 200)	9	4
<b>6 Gb/s SAS x8 / PCIe 2.x</b>	PCIe (3 200)	14	6
<b>12 Gb/s SAS x4 / PCIe 2.x</b>	PCIe (3 200)	14	6
<b>12 Gb/s SAS x4 / PCIe 3.0</b>	SAS (4 400)	19	8
<b>12 Gb/s SAS x8 / PCIe 3.0</b>	PCIe (6 400)	28	12

#### Storage configuration examples showing the respective bottlenecks and the number of drives supported in peak performance conditions

becomes the bottleneck, and it is not necessary to «overprovision» design by adding more recent technologies and / or more SAS lanes into the mix. The disks listed in Table 2 were all a 6 Gb/s interface and achieve 230 MB/s and 550 MB/s for 15 000 rpm HDDs and SSDs, respectively. Note that in Table 2, all disks are supposed to work to their peak efficiency simultaneously and that does not always happen. It is also important to note that the IOPS are often more critical than raw throughput in many current applications, depending on the circumstances. It is for these reasons that each configuration is normally able to support more disks.

### EVOLVING TO 12 GB/S SAS TECHNOLOGY

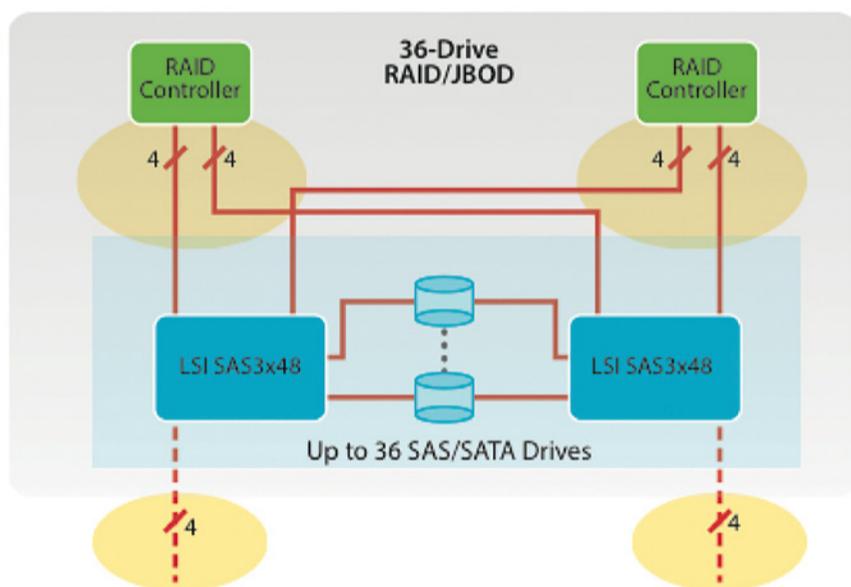
The main problem when it comes to move to third-generation 12 Gb/s SAS technology, is well known: it is that of investment protection. Most organizations have invested heavily in SAS disks and want to preserve that investment as they pass the SAS 12 Gb/s barrier. The difficulty being that third-generation SAS standard maintains backward compatibility by adopting the SAS slower system data rate.

In small point to point configurations, this is not always a problem, because migration would require upgrading the initiator and the target. However, in most organizations, this type of configurations are quite rare. Therefore, the performance limit at the slowest data rate at the system level does not allow these organizations without point to point configuration of obtaining the performance of 12 Gb/s since all disks are not SATA-III compliant.



## Data storage systems bottlenecks change depending on the generations of technologies connected end to end. With the advent of third-generation PCIe, for example, the second generation SAS technology has become the new bottleneck in data storage.

Fortunately, it is possible to overcome these limitations, and this requires a slightly better understanding of how SAS expanders work. A SAS expander enables one or more initiators to communicate simultaneously with multiple targets as shown in Figure 1. The expanders give the SAS its remarkable evolution faculties, and as each is able to support multiple disks they also allow to combine the performance of these drives.



**The two 48-port SAS expanders presented in this example are each connected to 12 initiators and targets 36**

Figure 1 shows 48-port expanders connected to two redundant RAID controllers and a shared disk array. Arrays of additional drives can be daisy chained, and some SAS expanders can support up to 2,000 disks in such configurations. According to the third-generation SAS standard if a SAS or SATA drive operates at 6Gb/s with RAID controllers also working at 6Gb/s. But if the bandwidth of two 6 Gb/s disks must be combined into a single

12 Gb/s channel, how many channels can be combined with PCIe?

This combination of bandwidth and performance immediately boost IOPS while protecting the investment in first and second generation SAS storage systems. This allows to introduce 12 Gb/s SAS systems in an existing 6 Gb/s storage array in order to immediately double the overall system performance. In fact, with the combination of bandwidth, it is not even necessary to use last generation 12 Gb/s SAS drives to achieve 12 Gb/s SAS performance at the system level.

The combination of bandwidth works by using a buffer per 12 Gb/s port, to allow what appears to be a communication between 12 Gb/s SAS initiators and any speed (3 Gb/s or 6 Gb/s) SAS or SATA targets. Stacking technology is proprietary, but it operates entirely within the expander, allowing to standardize all interfaces. In this way, a SAS port 12 Gb/s SAS drive connected to a 6 Gb/s SAS port would operate as a second-generation backward compatible 6Gb/s, but the whole system would work at the third generation 12 Gb/s data rate.

## CONCLUSION

Data storage systems bottlenecks change depending on the generations of technologies connected end to end. With the advent of third-generation PCIe, for example, the second generation SAS technology has become the new bottleneck in data storage. The third-generation SAS technology now gets the best performance out of third-generation PCIe, making the new PCIe interface bottleneck on 12 Gb/s SAS systems.



# Green Mountain

The greenest data center in the world.

**Housed in a former NATO ammunition depot in the heart of the mountain of a Norwegian fjord, the Green Mountain data center has all the advantages of a highly secure and environmentally friendly place.**



**G**reen Mountain is a Tier III + data center, located in highly secure Rennesøy, near Stavanger, the «oil capital», on the southwest coast of Norway. A mirror site for replication and data security, was also built in Rjukan in Telemark, an historic industrial site located just over 150 km west of Oslo.

### **THE RENNESØY GREEN MOUNTAIN DATA CENTER**

What was originally a NATO ammunitions depot in the heart of a mountain provided a unique location for the data center, more than 20 000m<sup>2</sup> of IT rooms, somewhat unusual. It consists of six rooms accessible by long tunnels ensuring a high level of security. The standard Tier III and Tier IV were respected for its design.

The energy supply comes from several hydroelectric facilities upstream of the data

center. It is powered by 3 separate medium voltage networks. The entire power supply, medium voltage cells via the low voltage distribution to high-performance modular Symmetra MW UPS was provided by Schneider Electric. The latter was also involved with the Green Mountain AS Datacenter cooling to minimize the energy impact by using natural resources. Seawater is drawn from 100 meters deep and with a system of siphon it is brought to the fjord basins for feeding the heat exchangers (sea water / water datacenter) using effect of gravity, which reduces the consumption of pump systems. The water drawn from 100 meters deep at a constant temperature is at 8°C throughout the year.

This solution is a fully redundant cooling production (2 separate drilling systems, 2 pools) for distribution (double cold water loop, redundant air conditioning, separate networks) in an 2N type scheme.



**Schneider Electric, the global specialist in energy management, accompanied Green Mountain AS to design one of the most Green datacenters in the world to date, using water from the Baltic Sea for its cooling and powered by hydraulic energy, 100% renewable and low cost. Schneider Electric has addressed all of the energy infrastructure (cooling, electricity and urbanization of IT**

**rooms) on the project. «The volume of data doubles every two years and currently accounts for 2% of global CO2 emissions. We are proud to take part in this innovative project, decidedly environmentally friendly.» says Arild Bjørkedal, Vice President End User / ITB & Energy, Schneider Electric Norway. «Combining our flexible, scalable and reliable solutions and**

**the unique location of the data center ensures an extremely stable power supply to customers of Green Mountain. For Schneider Electric, this is a great achievement and we look forward to design new data centers in Norway. This order is indicative of our**

**expertise as a key data center provider. The combination of international and local expertise in our approach has ensured a timely delivery. It is the first to deliver an operational data center in such a short time, «says Arild Bjørkedal.**

The cold distribution is ensured by a precision cooling InRow cooling system to remove heat closer to IT sources, supplemented by Schneider Electric EcoAisle solution which ensures the thermal hot aisle containment, and increasing the performance of air conditioners. The environmental impact of cooling used in the data center is almost zero, in terms of cold production consumes very few resources and in terms of distribution. The data center is managed through Schneider Electric's StruxureWare for Data Center software suite monitoring tool that provides carriers and operators with the information they need to ensure the perfect balance between high availability and efficiency, and throughout data center life cycle.

### **THE RJUKAN GREEN MOUNTAIN MIRROR SITE**

Schneider Electric also designed the Rjukan data center. This mirror site whose construction began in December 2013 is already in service. The ability to deliver within a few months an operational data center was instrumental



in Green Mountain, as this required commitment has enabled it to win a major contract with one of its major customers.

This data center has been designed, built and delivered in less than five months: a technical feat made possible by the use of Powertrain prefabricated modular solutions and modular panels made by AST Modular, a Spanish company recently acquired by Schneider Electric. The speed of local authorities who validated the construction permits within 24 hours was also decisive for this important contract.

The Rjukan datacenter consists of prefabricated modules featuring medium voltage supply of cells with integrated SM6 power module including transformer, modular inverter and control cabinet and also batteries with a life of 30 minutes. The main building hosting the



## The Rennesøy Green Mountain Datacenter in figures



- 100% renewable redundant hydropower.
- Medium voltage cells, transformers, switchgear cabinet and high energy performance inverters

- reducing losses.
- Electrical distribution by prefabricated Canalis Sheath offering flexibility and reliability.
  - Modular Symmetra MW UPS scalable

- 1,2MW by 200 kW modules guaranteeing a yield of 97%.
- «Natural cooling» solution type using seawater and regulation by Schneider Electric PLCs, ensuring very low power consumption

- Integrated cooling units for highly efficient cooling.
- Units perimeter cooling Uniflair by Schneider Electric for cooling equipment rooms
- Thermal Containment of the hot aisle by the system EcoAisle to increase the performance of air conditioners by more than 30%
- Rehabilitation of a brownfield reducing the carbon footprint from construction
- StruXaware DCIM Suite for Datacenters software for continuous monitoring.



IT rooms was built with AST Modular «box in the box» modular room panels.

The cooling is entrusted to two BECF Free cooling chillers with 600 kilowatts of magnetic bearing «Turbocor» compressors and an array of onboard cooling exchangers. This solution is specially adapted to extreme temperatures (- 40 ° C) with coefficients of performance (COP) of over 30 in winter and above 15

on average over the year, traditional chillers typically offer between COP 3 and 5.

This installation is entirely powered by electricity generated by hydroelectric power stations located a few hundred meters below, whose proximity enhances reliability. The reliability of different sources of energy supply, the security inherent in their location and the respect of higher Uptime Tier III standards, make Green Mountain a perfect Datacenter for European companies concerned with protecting their data. The Rjukan data center is the first phase of an investment of several million of euros to build a large 'zero emissions' colocation center. Having delivered the project ahead of schedule was decisive.



## The Rjukan Green Mountain mirror datacenter in figures

- 100% renewable redundant hydropower
- Proximity single 6 hydroelectric plants
- Integration of last generation cooling solution Turbocor Uniflair free-cooling compressor by Schneider Electric to benefit from an exceptional energy performance
- Symmetra PX 500 modular UPS by 25kW modules.
- Prefabricated Schneider Electric

- Powertrain modules designed and pre-tested in the factory for rapid deployment on site. Reinforced insulation to sustain temperatures below -40 °C.
- Electrical distribution by prefabricated Canalis Sheath offering flexibility and reliability.
  - Fake Uniflair floor tiles by SE design adapted to a data center operating constraints.

- Thermal Containment of the hot aisle by the system EcoAisle to increase the performance of more than 30% of air conditioners
- Inrow cooling Units for highly efficient cooling.
- Free-cooling Uniflair air conditioning units by SE for cooling Powertrains
- «Modular Datacenter Room» Panels made by AST Modular SE

- By ensuring optimal partitioning of IT facilities at short notice.
- «Datacenter as a Service» and «Cloud Enabled» 100% modular and scalable solution that enabled the design and deployment in less than 5 months to support the customer's business model.
  - StruXaware DCIM Suite for Datacenters software for continuous monitoring.

### 98.5% HYDROPOWER ENERGY

98.5% of electricity in Norway comes from hydropower, which makes it the largest producer in Europe, an ecology feature consistent with the fundamental values of the country.

*“Our goal is to be a leading operator in the European market, recognized for the quality of its green data centers. Green Mountain AS displays a PUE (Power Usage Efficiency) of less than 1.2 in a type III configuration. In view of IT infrastructure solutions, cooling and management provided by Schneider Electric, our customers have the opportunity to reach this PUE, every hour, every day of the entire year”, says Knut Molaug.*



**“We have great ambitions for Green Mountain AS, which is why the choice of our partners was essential: we wanted the very best on the market in terms of energy solutions coupled with optimum use of the architecture and technology. Since we opened our first data center near Stavanger, we find that the security, the stability of electricity prices, and the ecology and scalability, help us meet strong demand. Thus we plan to invest £ 60 million (more than 75 million €) in the Rjukan data center over the next few years”, says Knut Molaug, CEO of Green Mountain AS.**



## LAB REVIEW

## Dell EqualLogic ps6210 Symantec Netbackup Appliance 5230

## VIEWPOINT

## The best practices to reduce infrastructure energy consumption

## EN COUVERTURE

# Flash Storage: the ultimate frontier

Breaking the \$1 per gigabyte barrier



**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

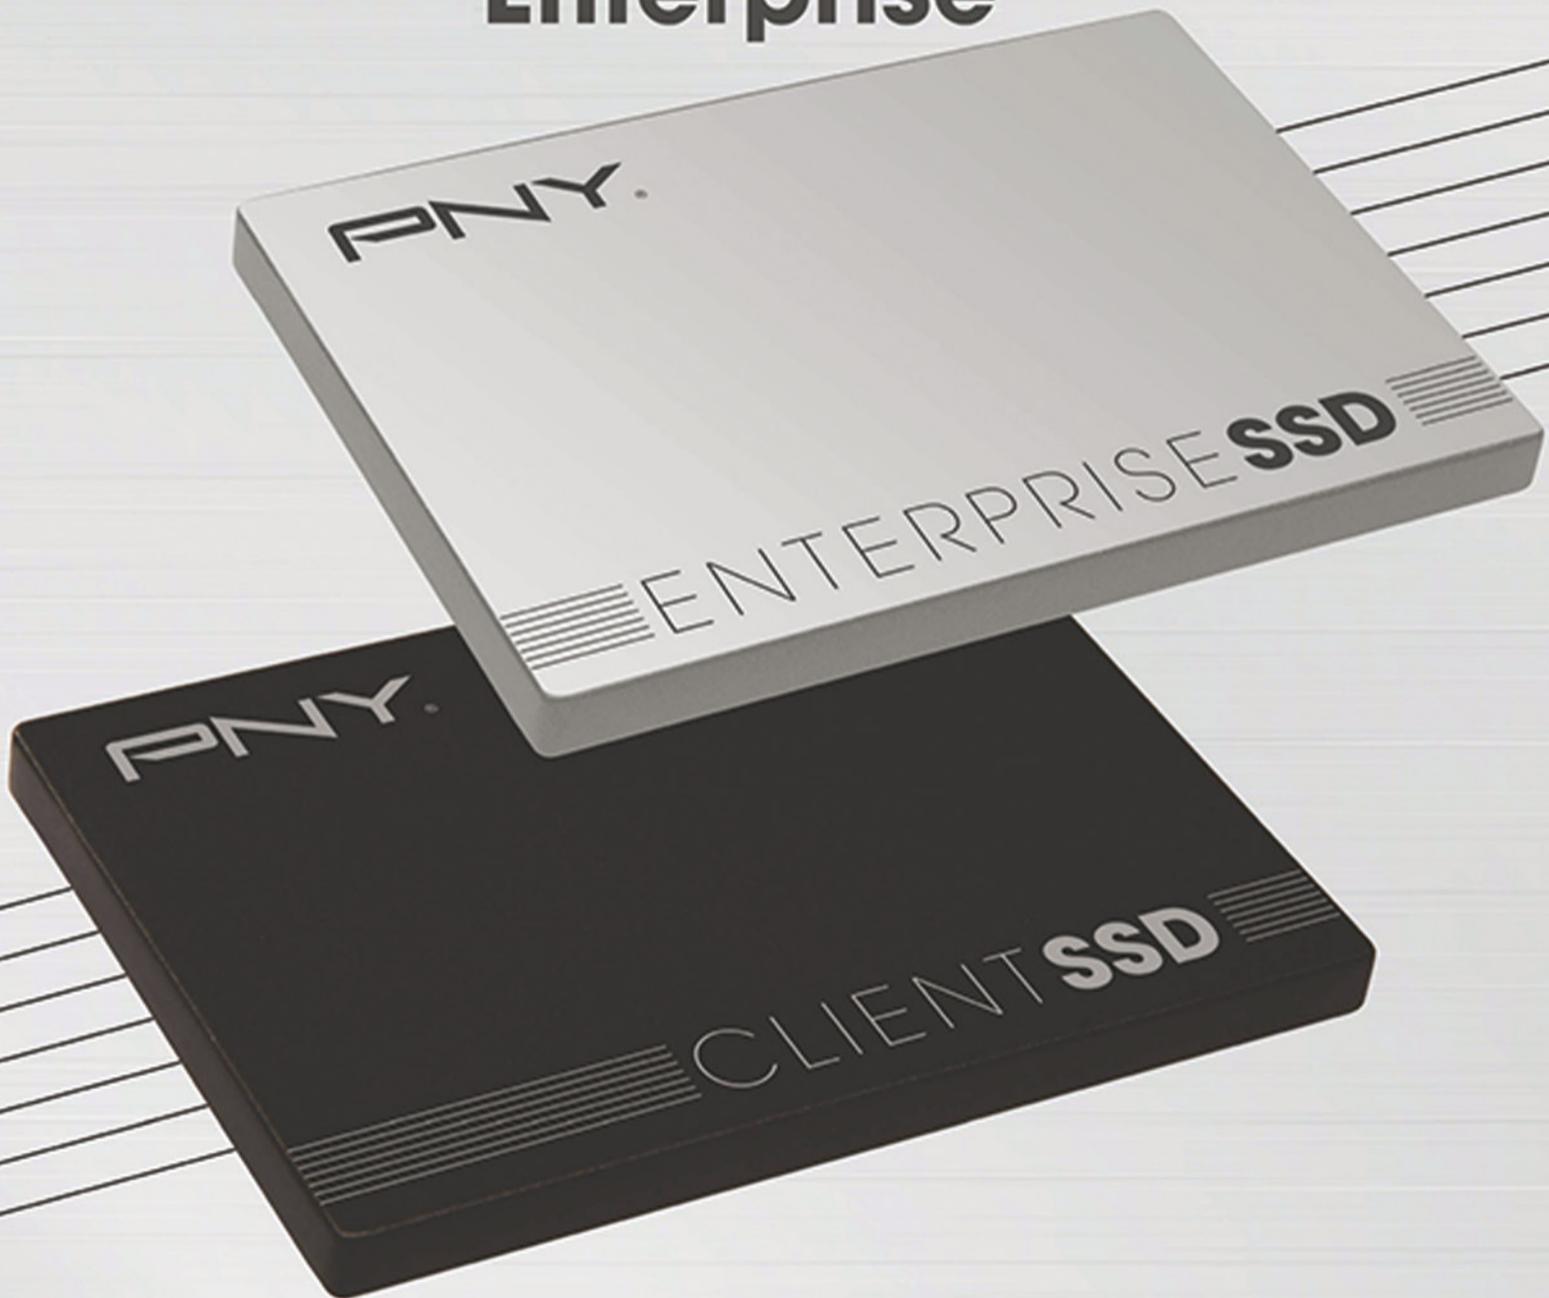
# PNY

## Professional Solutions

NVIDIA Quadro® / NVIDIA Tesla® / NVIDIA Grid® / SSDs

# Solid State Drives

## EP7000 Series Enterprise



## CL4000 Series Client

[WWW.PNY.EU](http://WWW.PNY.EU)

[PNYPROFR@PNY.EU](mailto:PNYPROFR@PNY.EU)