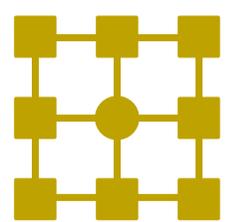


**FIRST REVIEW** NVIDIA'S FLAGSHIP QUADRO M6000



# 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

EXCLUSIVE  
INTERVIEW

**Ralph Schmitt**

OCZ CEO

**#2**  
GLOBAL  
EDITION



## STORAGE VIRTUALIZATION

A WHOLE  
NEW WORLD  
OF GROWTH  
OPPORTUNITIES

LAB REVIEW  
**Dell PowerEdge FX2**  
InforTrend  
EonStor 3000 DS  
**Dell Precision T5810**  
Toshiba W50

HOW TO  
**Create your  
Raspberry Pi  
supercomputing  
cluster**

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 !

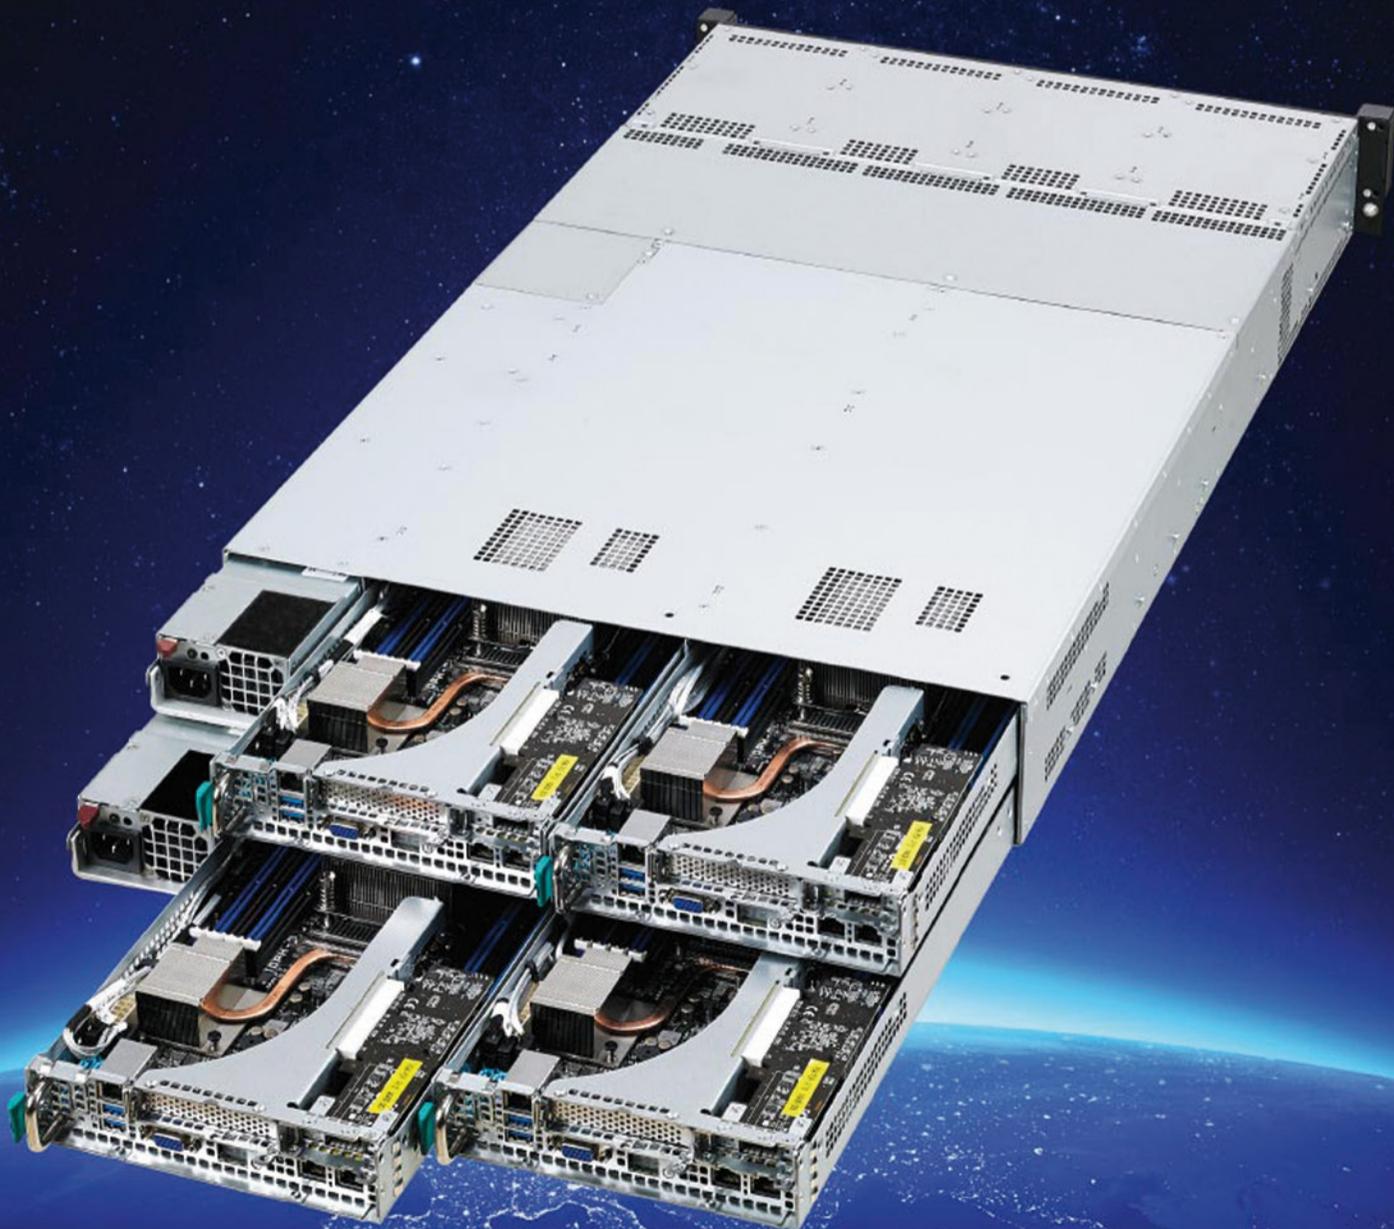
**This issue focuses on a hot topic nowadays,** storage virtualization. Being able to get the most of an existing storage infrastructure and to extend it while winning capacity and features has become key within tight budget constraints. And virtualizing the existing storage pools and islands within the enterprise is only possible through virtualization. If you have a storage problem, **there are good chances you will find the solution In this issue!**

**This issue is also packed with product reviews,** of which the outstanding – or not, depending on the use case – new flagship of NVIDIA, the Quadro M6000 – Our benchmarks show it certainly shines on the performance side, but at a price. **Our product reviews also focus on infrastructure elements** like Dell's modular and scalable PowerEdge FX2's powerful enough to make one rethink an enterprise equipment strategy. Not to mention InforTrend's innovative and performance-oriented storage arrays

**As always,** we welcome your suggestions and comments at [editorial@hpcreview.com](mailto:editorial@hpcreview.com)

Happy reading!

**Maximize Cabinet Density,  
Maximum Power Saving  
RS720Q-E8-RS8-P (145W)  
RS720Q-E8-RS12**





COVER STORY



# STORAGE VIRTUALIZATION

A whole new world of growth opportunities

NEWSFEED

Intel celebrates 50 years of Moore's Law

The first 3D printed car

Ralph Schmitt, CEO OCZ

HGST Active Archive

Books | Moocs | The HPC observatory



LAB REVIEW

NVIDIA Quadro M6000

Toshiba W50

Dell FX2

Infotrend EON Stor 3000 DS

Dell Precision Tower 5810

Qlik Sense Desktop

HPC Labs : How do we test

HOW TO

A Raspberry Pi-computing based cluster

VIEWPOINT

The best benchmark is you

TECH ZONE

CD-adapco, 35 years of innovation



# When Data Needs More Firepower: The HPC, Analytics Convergence

**B**ig data arguably originated in the global high-performance computing (HPC) community in the 1950s for government applications such as cryptography, weather forecasting, and space exploration. High-performance data analysis (HPDA)—big data using HPC technology—moved into the private sector in the 1980s, especially for data-intensive modeling and simulation to develop physical products such as cars and airplanes. In the late 1980s, the financial services industry (FSI) became the first commercial market to use HPC technology for advanced data analytics (as opposed to modeling and simulation). Investment banks began to use HPC systems for daunting analytics tasks such as optimizing portfolios of mortgage-backed securities, pricing exotic financial instruments, and managing firm-wide risk. More recently, high-frequency trading joined the list of HPC-enabled FSI applications.

The invention of the cluster by two NASA HPC experts in 1994 made HPC technology far more affordable and helped propel HPC market growth from about \$2 billion in 1990 to more than \$20 billion in 2013. More than 100,000 HPC systems are now sold each year at starting prices below \$50,000, and many of them head into the private sector.

## WHAT'S NEW?

It's widely known that industrial firms of all sizes have adopted HPC to speed the develop-



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

ment of products ranging from cars and planes to golf clubs and potato chips. But lately, something new is happening. Leading commercial companies in a variety of market segments are turning to HPC-born parallel and distributed computing technologies — clusters, grids, and clouds — for challenging big data analytics workloads that enterprise IT technology alone cannot handle effectively. IDC estimates that the move to HPC has already saved PayPal more than \$700 million and is saving tens of millions of dollars per year for some others.

The commercial trend isn't totally surprising when you realize that some of the key technologies underpinning business analytics (BA) originated in the world of HPC. The evolution of these HPC-born technologies for business analytics has taken two major leaps and is in the midst of a third. The advances have followed this sequence:

- **Phase 1** was the advance from the mainframe mentality of running single applications on traditional SMP servers to modern



clusters (i.e., systems that lash together homogeneous Linux or Windows blades to exploit the attractive economics of commodity hardware). The cluster was invented by two NASA HPC experts in 1994.

- **Phase 2** was the move to grids with the goal of supporting multiple applications across business units coherently. This enables enterprisewide management of the applications and workloads.
- **Phase 3** is the emerging move to cloud computing, which focuses on delivering generic computing resources to the applications and business units on an on-demand, pay-as-you-go basis. Clouds can be hosted within a company, by an external provider, or as a hybrid combination of both.

#### WHY BUSINESSES TURN TO HPC FOR ADVANCED DATA ANALYTICS

High-performance data analysis is the term IDC coined to describe the formative market for big data workloads that exploit HPC resources. The HPDA market represents the convergence of long-standing, data-intensive modeling and simulation (M&S) methods in the HPC industry/application segments that IDC has tracked for more than 25 years and newer high-performance analytics methods that are increasingly employed in these segments as well as by commercial organizations that are adopting HPC for the first time. HPDA may employ either long-standing numerical M&S methods, newer methods such as large-scale graph analytics, semantic technologies, and knowledge discovery algorithms, or some combination of long-standing and newer methods.

The factors driving businesses to adopt HPC for big data analytics (i.e., HPDA) fall into a few main categories:

- **High complexity.** HPC technology allows companies to aim more complex, intelligent questions at their data infrastructures. This ability can provide important advantages in today's increasingly competitive markets. HPC technology is especially useful when

there is a need to go beyond query-driven searches in order to discover unknown patterns and relationships in data — such as for fraud detection, to reveal hidden commonalities within millions of archived medical records, or to track buying behaviors through wide networks of relatives and acquaintances. IDC believes that HPC technology will play a crucial role in the transition from today's static searches to the emerging era of higher-value, dynamic pattern discovery.

- **High time criticality.** Information that is not available quickly enough may be of little value. The weather report for tomorrow is useless if it's unavailable until the day after tomorrow. At PayPal, enterprise technology was unable to detect fraudulent transactions until after the charges had hit consumers' credit cards. The move to high-performance data analysis using HPC technology corrected this problem. For financial services companies engaged in high frequency trading, HPC technology enables proprietary algorithms to exploit market movements in minute fractions of a second, before the opportunities disappear.
- **High variability.** People generally assume that big data is “deep,” meaning that it involves large amounts of data. They recognize less often that it may also be “wide,” meaning that it can include many variables. Think of “deep” as corresponding to lots of spreadsheet rows and “wide” as referring to lots of columns (although a growing number of high-performance data analysis problems don't fit neatly into traditional row-and-column spreadsheets). A “deep” query might request a prioritized listing of last quarter's 500 top customers in Europe. A “wide” query might go on to analyze their buying preferences and behaviors in relation to dozens of criteria. An even “wider” analysis might employ graph analytics to identify any fraudulent behavior within the customer base.

**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 :



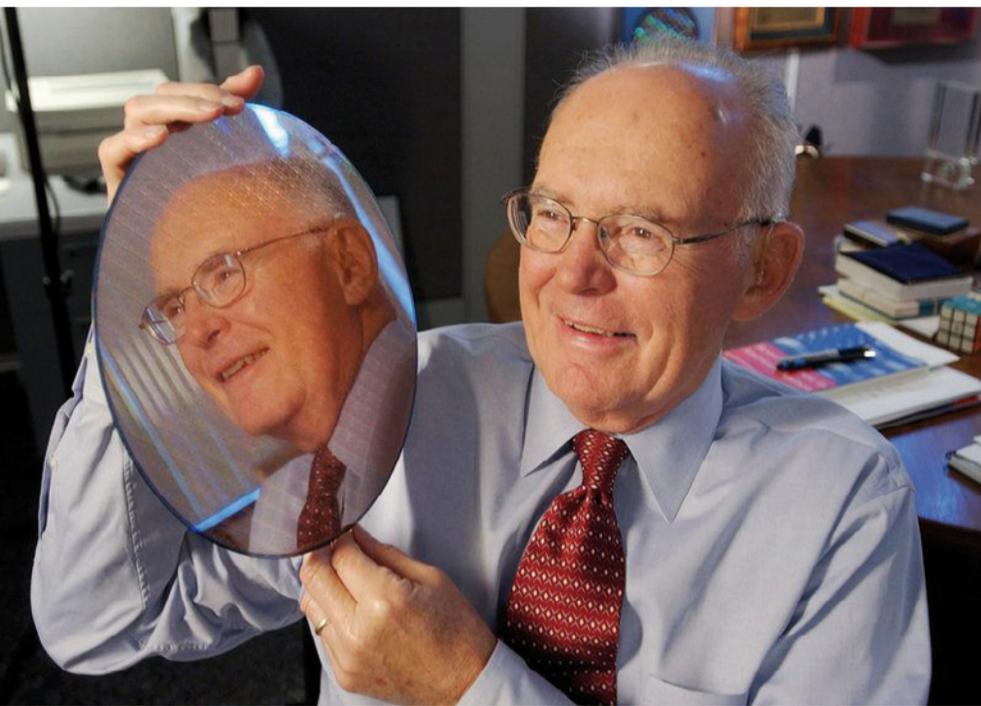
Silver Sponsors :





# 50 years of Moore's Law

**50 years ago, Gordon Moore, Intel co-founder decreed that the power of microprocessors would double every 24 months. Five decades later, this prediction continues to be true.**



PAUL SAKUMA (AP)

**« The number of transistors in a processor will double approximately every 24 months. »**

— Gordon Moore, Intel co-founder

**O**ver the investments in semiconductor manufacturing technology, Intel has made Moore's Law a reality. Relying on his own observations and his calculations Gordon Moore had no idea that his law would last as long. In the mid 60s, the manufacturing process used silicon wafers with a diameter of 2 inches. Today, the wafers have of a diameter of 12 inches, and we the industry is on the verge of switching to 18-inch wafers. With a slump in production costs, a single transistor

costs less than a billion times than 50 years ago. And a comical resulting enacted by Moore himself, who says that the number of transistors produced since finally exceeded the number of ants on Earth!

## **1969: AN INNOVATIVE RESPONSE TO AN UNUSUAL DEMAND**

In 1969, the Japanese company Nippon Calculating Machine Corporation Intel approach to design a set of 12 chips for its new Busicom 141-PF calculator. Engineers from Intel respond by suggesting a family of only four chips, of which one would be programmable. This set of four chips known as the MCS-4 included a central processing unit (CPU) 4004, a read only memory (ROM) for custom applications programs, a random access memory chip (RAM) for data processing and an input / output (I / O) circuit.

## **1971: THE BEGINNING OF THE ERA OF INTEGRATED ELECTRONICS**

With its 2,300 transistors Intel 4004 processor can be programmed for use in a variety of products. In this it differed from its predecessors and peers. After buying the rights of Nippon Calculating Machine Corporation, Intel launched the marketing of Intel 4004 processor and chipset with an advertisement in the issue of November 15, 1971 issue of Electronic News "heralding a new era in integrated electronics."





## 50 in 10 milestones

1968

1971

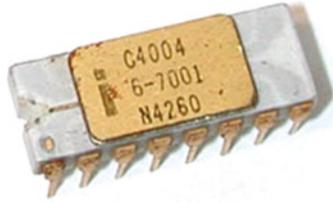
1974

1978

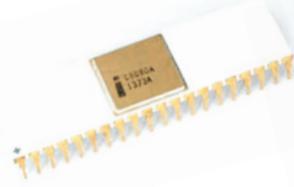
1982



Bob Noyce and Gordon Moore establish Intel



Intel introduces the 4004, the first integrated processor



Intel 8080, the first mass processor. 4500 transistors and ten times more powerful than its predecessors



Intel 8088, built in the heart of the IBM PC



Intel 80286, with 134,000 transistors

1985

1985

1989

1993

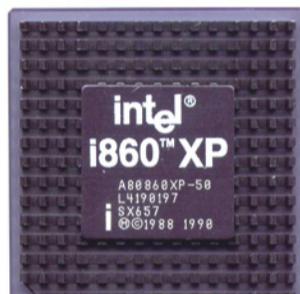
2008



iPSC / 1, the first supercomputer based on Intel 286 CPU



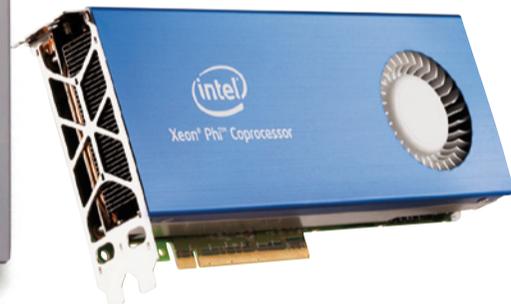
80386, 275000 transistors and multitasking



Intel i860 CPU over one million scientific use transistors



The first Pentium was launched, with more than 3 million transistors



Intel launches the first Xeon Phi dedicated to HPC

### THE PROCESSOR THAT GAVE BIRTH TO AN INDUSTRY

The 4004 has quickly become the first general purpose programmable microprocessor on the market and became the basic element around which the engineers could design of electronic devices of their own. About the size of a fingernail, the Intel 4004 processor provides the same computing power as the ENIAC, the first electronic computer built in 1946, which filled a whole room. To get an idea of the progress made since the launch of the Intel 4004, Intel Core processor contains 560 million transistors, is a factor of integration almost 250 000

times. The width of the circuits of the 4004 processor was 10 microns, or 10 000 nanometers. Today, the circuits of Intel microprocessors are between 45 and 32 nanometers wide. By comparison, the average human hair is 100,000 nanometers wide. **RAMON LAFLEUR**

**The 4004 has quickly become the first general purpose programmable microprocessor.**



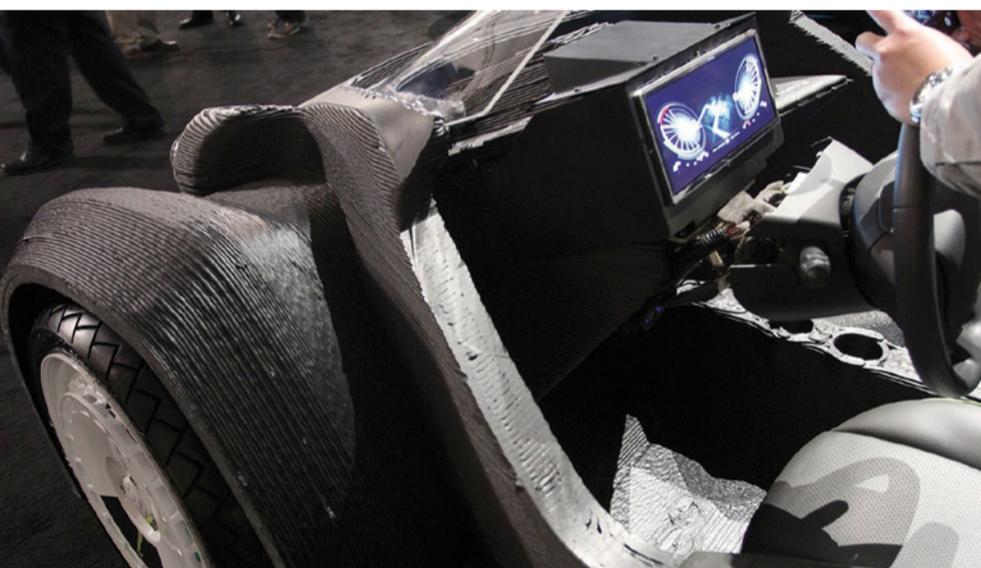
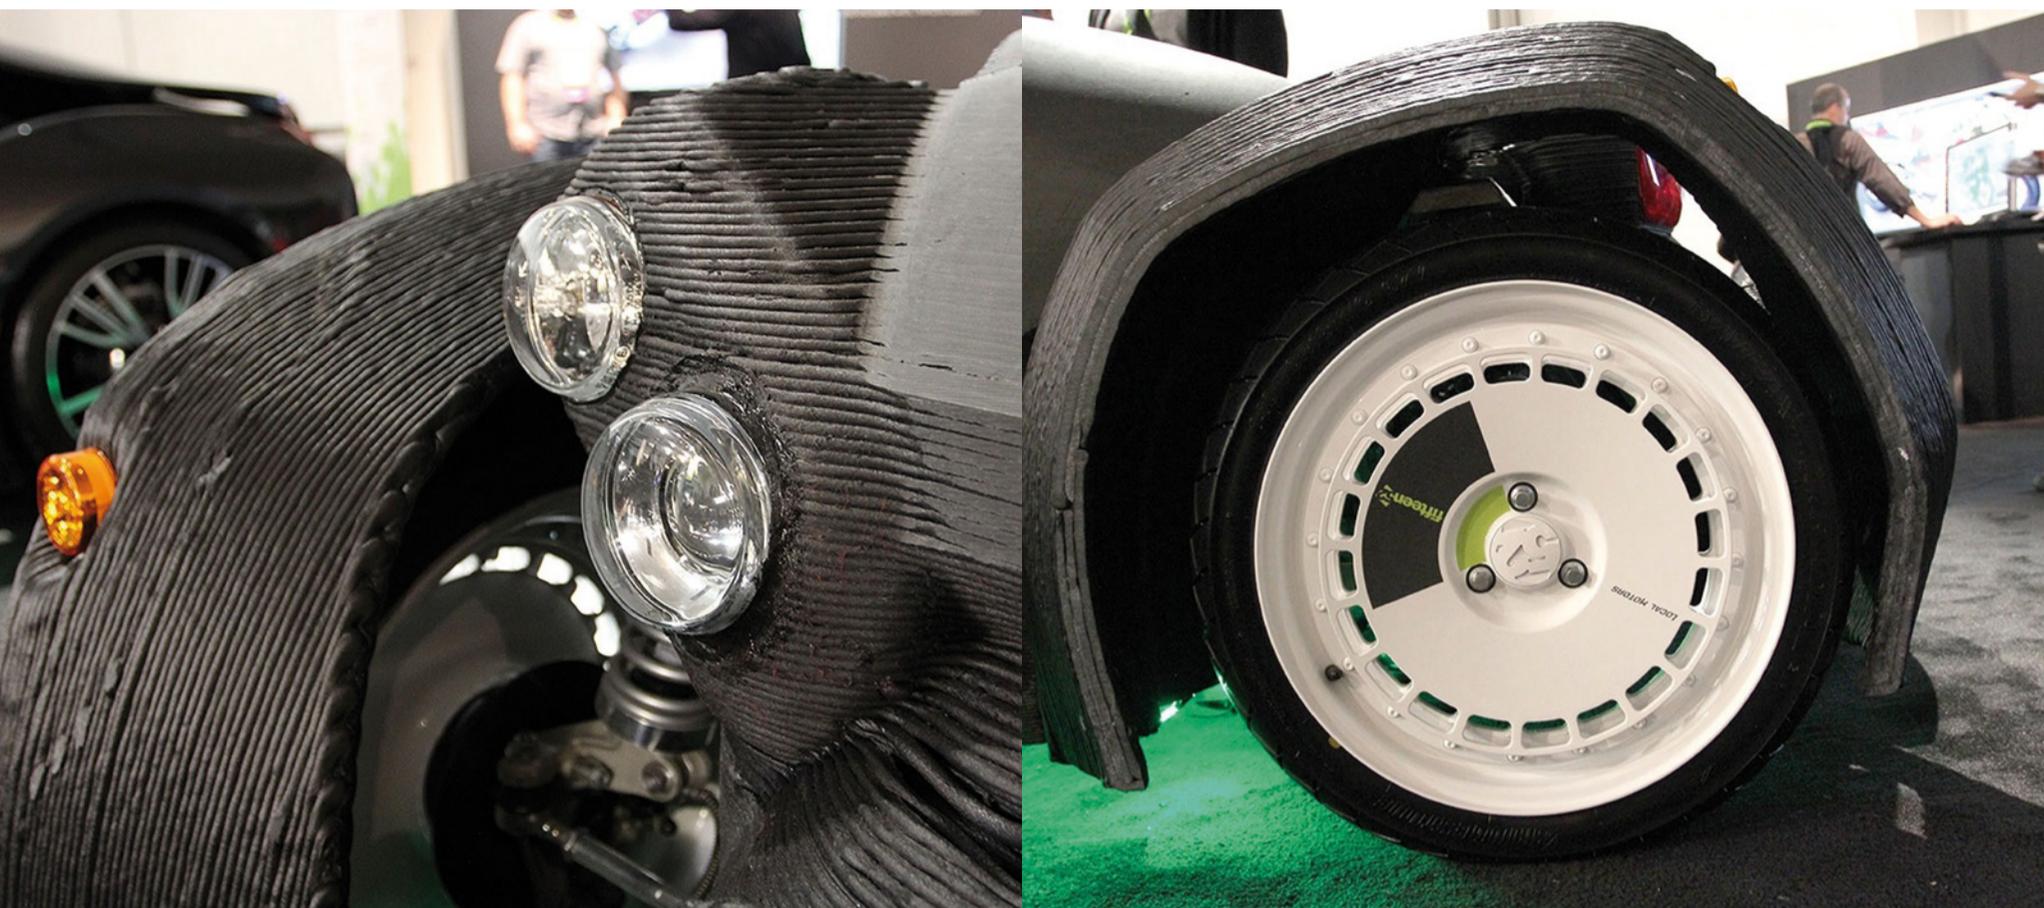
## Local Motors Strati

# The first 3D printed car

**What can transport you to the exhilarating speed of 65 km/h that you can (almost) print in 3D at home? A revolutionary concept vehicle named Strati.**

**N**owadays, it is possible to print almost anything in 3D. The small US company Local Motors located in Arizona has taken the concept even further, by designing the first functional vehicle entirely made of 49 3D printed elements. The chassis and the vehicle body are made of a single solid piece. The vehicle has a top speed of 65km / h. Powered by an electric motor and able to carry two passengers, the Strati is made from black plastic layer reinforced with carbon fibers. The

printer used to build the vehicle was built by a company called Cincinnati Inc. and was able to produce parts up to 90 x 150 x 300 cm. Local Motors hopes to be able to make cars like this for about \$ 17 000 in the near future. The seats, bodywork, chassis, dashboard, center console and cover are printed in 3D. The experimental process used for their printing was designed with accuracy in mind. Every one piece part was done in order to facilitate assembly, but also to ensure the best possible vehicle structural rigidity.

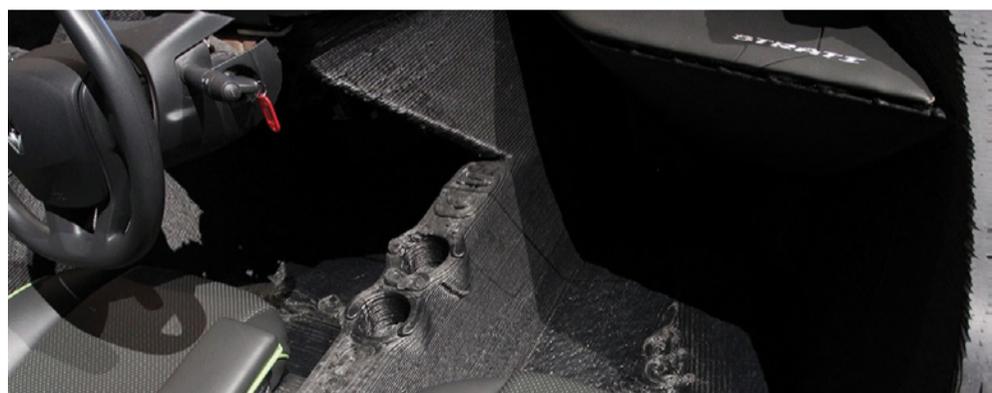


### 100 FACTORIES IN 10 YEARS

The Strati has a driving range anywhere between 190 and 240 km. The prototype version was demonstrated at the International Manufacturing Technology show in Chicago last month. Some mechanical components such as the battery, motor, wiring and suspension were borrowed from Renault's Twizy. Local Motors adopts a decentralized approach. Instead of having a huge factory, Local Motors plans to build 100 «micro factories around the world in the next 10 years.» Local Motors officials further indicate that each of these plants, located near urban centers may create more than 100 jobs each.

### HIGH-TECH DASHBOARD

Local Motors has teamed with NVIDIA to integrate NVIDIA Drive CX, digital cockpit computer just announced during CES 2015. Prior to integration of the Drive CX, the dashboard of the Strati was much more traditional. Now, the dashboard is made of a flat screen with gra-



phics worthy of a science fiction movie! NVIDIA DRIVE CX is a digital cockpit computer, capable of running a 3D navigation software, digital instruments, but also the processing of natural speech (for voice commands), and image processing. NVIDIA indicates that it is «a platform with a double benefit: market the Strati faster and with the lowest total cost. «The vehicle is not yet available commercially, but you can already sign up to be notified when it becomes available on the Local Motors website.

A photograph of a man in a white shirt leaning on a metal railing, with a large OCZ Storage Solutions logo in the background. The logo consists of the letters 'OCZ' in large blue font, with 'STORAGE SOLUTIONS' in a smaller blue font below it. Below that, the text 'Hibata Group Company' is partially visible in blue. A red horizontal bar is at the top of the page.

OCZ  
STORAGE SOLUTIONS

Hibata Group Company

Interview

# RALPH SCHMITT

**CEO of OCZ answers our questions  
about his vision for the future  
of flash-based storage**

**Mr Schmitt, can you tell us more about OCZ's enterprise product and market strategy?**

OCZ has gone through a lot in the past couple of years which ended on a high note when Toshiba announced that it will be acquiring OCZ's assets for \$35 million. Since then OCZ has been operating as an independent subsidiary under Toshiba. Unlike what usually happens after acquisitions, OCZ has not been integrated to Toshiba's existing units and all of OCZ's different units (such as marketing, product development etc.) have remained intact. The OCZ-Toshiba relationship resembles more like a strategic partnership to be honest as that indicates how separate the two are, but only in a good way.

The reason why Toshiba wanted to keep OCZ independent was to ensure that the company's talent doesn't get buried into Toshiba's massive organization. The common issue in acquisitions is that the acquiring company may not be able to use all the talent and resources of the acquired company because the culture of each organization can differ. In the end, individual talent isn't just something concrete and absolute like money that can easily be transferred from one place to another -- there are various environmental and cultural factors that go into it. For example, a change from liberal and relaxed culture to more tightly managed one can suppress some individuals because the way they work has to change, and the new way may not suit them. It is always a major challenge to integrate two companies with different organization cultures, so keeping the two separate is not always a bad idea. Fortunately, Toshiba saw eye-to-eye with us and we just haven't had enough time and resources to fulfill the potential. As a result, Toshiba decided to let OCZ continue its business and R&D as a separate unit by providing financial support for the company. The new OCZ had, and still has, a sustainable business plan, so there was no reason to break up the company into pieces.

The companies have also gone through the process of unifying product roadmaps. Ini-

tially, there was quite a bit of overlap because both had their own lineups and tried to cover as much of the market as possible, but by mid-2015 there should no longer be any major product overlaps. For instance, OCZ has stopped all SAS development as Toshiba has more expertise in that field and the same goes for certain OEM drives where Toshiba has always had a strong presence. At the same time, Toshiba will let OCZ handle the OCZ branded retail sales and on the enterprise side OCZ's focus will be in SATA and PCIe drives. The idea and goal is to supplement each other's lineup, not to compete.

**How has OCZ's time to market improved since Toshiba's acquisition?**

As part of the Toshiba Group, OCZ is now able to leverage a number of Toshiba's labs and run much more extensive testing throughout the design cycle, as well as do more specialized testing and validation to meet specific customer requirements. During the baby steps of the acquisition, the companies had a few different philosophies when it came to validation, which makes sense given the different backgrounds. For instance, Toshiba might focus more on the NAND during testing/validation (Toshiba is the inventor and major supplier of NAND, after all), whereas OCZ would focus on other aspects we could impact like the controller and firmware first.

The companies have then aligned their philosophies to use the same thinking during product development and testing. Toshiba holds OCZ drives to the same quality standards as its own and OCZ has adopted many of Toshiba's processes for its quality control. The testing/validation aspect is just one example of how the two companies are supplementing each other. The key here is that both companies have and are still learning from each other to ensure that both companies are putting out the best possible products.

The one major advantage OCZ has gained is access to Toshiba's NAND knowhow. At the high level NAND is a fairly simple semiconductor (you apply a high voltage that creates an



electric field that forces electrons to tunnel from the channel through the silicon oxide to the floating gate), but once you get down to the atomic level it becomes a highly complex piece of technology. Some of OCZ's engineers have been living in Japan since the acquisition and just learning things about NAND and its characteristics to help build better products. In this case, especially endurance tweaks that require a very deep understanding of NAND, so having access to that knowledge and applying it to product design will provide OCZ with a very important advantage in the market.

Another crucial element that Toshiba provides is the access to its NAND supply. In the past OCZ had to source NAND from the open market, which resulted in varying quality NAND being used. Like many components (and goods/stocks in general), NAND is sold on spot market where anyone can sell or buy NAND. In case a NAND manufacturer (or any company) has excess stock, they can dump a portion of that into the spot market and sell it at the day's price. It is possible to buy NAND straight from the manufacturer, but the required volume tends to be fairly high and in turn you don't get the benefits of the dynamic spot market (predicting supply/demand and the associated price changes can be very profitable, if successful).

**Along the same lines, how has your R&D evolved?** I would regard OCZ as the one-stop-solution for the demanding storage needs. We control all the aspects of the technology. Our flexibility to customize the solutions according to the client's needs in multiple places on the value chain gives us a leading edge in this competitive market. Moreover like answered in question two, its parent company Toshiba's leadership in NAND flash memory segment gives OCZ the access to next-gen solutions.

Aside from high-level roadmap comparison, there is cooperation in individual products too. For a TLC NAND based SSD, the companies decided that it's better to let Toshiba handle the development of that drive because the com-

pany has been developing one for quite some time now and building a competitive TLC SSD requires a hefty amount of NAND knowledge, which is Toshiba's area of strength. It simply makes sense to focus all resources on one platform instead of having OCZ and Toshiba both spend money on a design that ultimately has the same end goal. However, for the first time, the Toshiba-built TLC drive will also retail under OCZ brand, which makes sense as OCZ already has a widespread distribution network and the required support for retail sales. On the other hand, Toshiba is very interested in OCZ's upcoming JetExpress controller, so I wouldn't be surprised to see Toshiba branded SSDs with the JetExpress controller for OEMs. There is no substitute for quality

While OCZ has transitioned manufacturing operations to a Toshiba certified partner (PTI), a robust engineering team leverages innovative IP to design and develop differentiated flash storage solutions. Quality is embedded in the entire design cycle with supplier qualification and monitoring, meet ISO 9001:2008 certification.

Today, OCZ is focused on delivering the right balance of performance and quality ensuring superior reliability for customers. Quality is a priority in everything we do and this is highlighted throughout the organization to remind us what our #1 job is as a company. We've made significant improvements and investments to ensure quality throughout the organization. As result, the return/defective rates have and continue to drop significantly. Current products have an overall return rate of just 0.57%\* whereof 0.31%\* turned out to be really defective, for example only 0.05% confirmed failures of Vector 150.

Today, OCZ is very different from yesterday's OCZ. The management has changed the course to a sustainable path where focus is on long-term development and proper validation. The Toshiba partnership just further emphasizes the quality aspect and also gives OCZ the much needed NAND supply and knowledge access.



**How do you see the Flash storage market evolve in the next 5 years?** There are several good interfaces available for enterprise SSDs but PCIe has become one of the best. Its direct connection to the host CPU provides ‘near-zero’ latency. The current PCIe 3.0 version provides 1GB/s of bandwidth per lane enabling PCIe SSDs to use 4 or 8 lanes concurrently, equating to 4GB/s or 8GB/s of achievable bandwidth from each device.

Today’s enterprise servers have several PCIe slots available for SSD deployment and can accommodate up to 40 lanes of PCIe per CPU. A two-socket server can support 80 lanes of PCIe. With an ongoing performance gap between server processors and HDD subsystems, PCIe-SSDs are highly desirable and represent a fast-growing multi-billion dollar global opportunity over the next five years. At present, PCIe SSD deployments are at introductory levels.

To achieve these expected deployment levels, PCIe required a host-controlled interface standard. Without this, each SSD vendor has to develop their own proprietary drivers on how the SSD will interface with the host. By standardizing the interface, only one driver development is required. As such, the Non-Volatile Memory Express (NVMe) specification was created to deliver the full potential of non-volatile memory for enterprise and client platforms in support of PCIe-based SSDs. NVMe is a very hot topic and we should dive deeper into that in a specialized topic.

We believe that the future lies in PCIe/NVMe, and our innovation lab is currently focusing on that. With the emergence of cloud computing, big data, virtualization and mobility, the data storage sector is going through a metamorphosis of its own. As a result, SSDs based on NAND flash memory chips, have emerged as the cardinal of storage interfaces, thus diminishing the hard drive market. We lead this transformation from hard drives to SSDs by

offering application specific solutions. OCZ’s flash-based storage and software solutions offer increased performance and responsiveness for a wide range of usage models and workloads, from virtualization to large-scale data centers to cloud computing. For example, our ZDXL PCIe solid state solution, which is an integration of hardware and software, can enhance performance of a SQL Server database environment,” claims Schmitt.

Striving to eliminate storage inefficiencies, OCZ’s SSDs take storage technology to the next level. We believe in amalgamating two basic factors in our solutions—cost and performance. In other words, to meet the performance requirements in a cost effective way and enhance the total computing experience. The company’s broad portfolio of SSD solutions includes several interfaces and enterprise feature-sets that enable unprecedented flexibility and efficiency for high throughput systems. Each solution is designed and built with the highest level of quality and reliability that enterprises and OEMs can count on for business-critical applications.

OCZ’s ability to marry high performance with cost effectiveness comes from the company’s long stint with the high performance computing and gaming industries. We already have a strong foothold in the gaming market, where it is important to push performance at low cost. We are looking forward to spark the same strategy in the enterprise world.

Let’s give you also an interesting example that reflects our expertise: 888.com, one of the world’s most popular online gaming providers was facing problems with their traditional data center and network with several I/O bottlenecks. Since multiple players access the resources at the same time, the key challenge was to reduce database I/O pressure on the SAN and improve the SQL server database performance. We simplified their storage system using our PCIe-based ZD-XL SQL Accelerator. OCZ also reduced the I/O load on the SAN by 50 percent thus alleviating the workload pressure and enhancing the efficiency of ap-

---

\* Data reflects active products as of December 15, 2014



plication data. We limited the problems and improved the SQL Server data throughput by 150 percent.

**... And compared to the mechanical HDD market?** Opportunities continue to unfold in the enterprise for flash-based solid-state storage as they provide faster I/O performance than hard disk storage, support large capacities and a variety of form factors and interfaces, and consume less power. The amount of data stored within enterprises and by client users has reached astronomical proportions creating a need to deliver data even faster and more efficiently.

Saying Flash-Drives will replace HDD completely will not be the right approach. Of course there will be Flash-only environments, when highest performance is required. But more often we will see a combination of both, as HDDs are still the medium to archive data the most affordable way. This will not change anytime soon.

### **What will be OCZ's next big product launch?**

With the increase of enterprise applications, SSDs have become an essential technology to ensure predictable delivery of Read I/O to the application layer. Failing to match such capabilities leads to increased workload and reduced performance.

One of OCZ's recently launched flagship storage solutions, the Saber 1000 Series, is a new class of enterprise SATA SSDs designed for read-intensive applications targeting high-volume deployment hyperscale, web-hosting and distributed computing environments.

Customers asked, we delivered. Recently OCZ announced the addition of a new enterprise SSD line to the Intrepid 3000 portfolio, the 3700 Series. Intrepid 3700 drives were specifically designed to deliver maximum usable capacity with no sacrifices in performance or endurance. It sounds too good to be true, but Intrepid 3700 SSDs provide a storage architect with the best of both worlds: ample storage and ample DWPD (Drive Writes Per Day).

Lastly, OCZ's PCIe solution Z-Drive 4500 is a full-featured flash solution that integrates robust performance, reliability, data protection, and monitoring functions for critical enterprise applications and demanding big data environments. After one OCZ Enterprise customer found themselves contending with data traffic their current storage architecture could not handle, they turned to OCZ Z-Drive 4500 SSDs to increase performance for their users.

Our very exciting next-generation NVMe-compliant Z-Drive 6000 PCIe SSD Series is currently in the evaluation state and will be launched in the second half of 2015. It is a native PCIe 3.0 NVMe 1.1b solution designed in an easy-to-deploy 2.5-inch form factor that not only delivers extremely high performance and near-zero latency, but also features a hot-pluggable SFF-8639 connector for easy deployment and superior serviceability. Catering to both the enterprise and consumer markets, OCZ's recently unveiled JetExpress SSD controller technology will be on display. Designed and developed in-house, JetExpress will be the heart and soul of OCZ's future product line being native SATA and PCIe/NVMe, supporting multiple form factors including M.2 and 2.5-inch SATA.

All in all, OCZ feels to be on the right path. The product roadmap, and especially the new JetExpress PCIe controller looks very promising and we are convinced the potential ends up being well executed. We are working continuously in earning back the consumers' trust, including those that received an unreliable drive in the past. We know this will take time. But the new OCZ has been injected with a new attitude and more resources when it comes to testing and validation. And we will show our expertise and innovation with our next generation products – so stay tuned!

INTERVIEW BY **JOSCELYN FLORES**

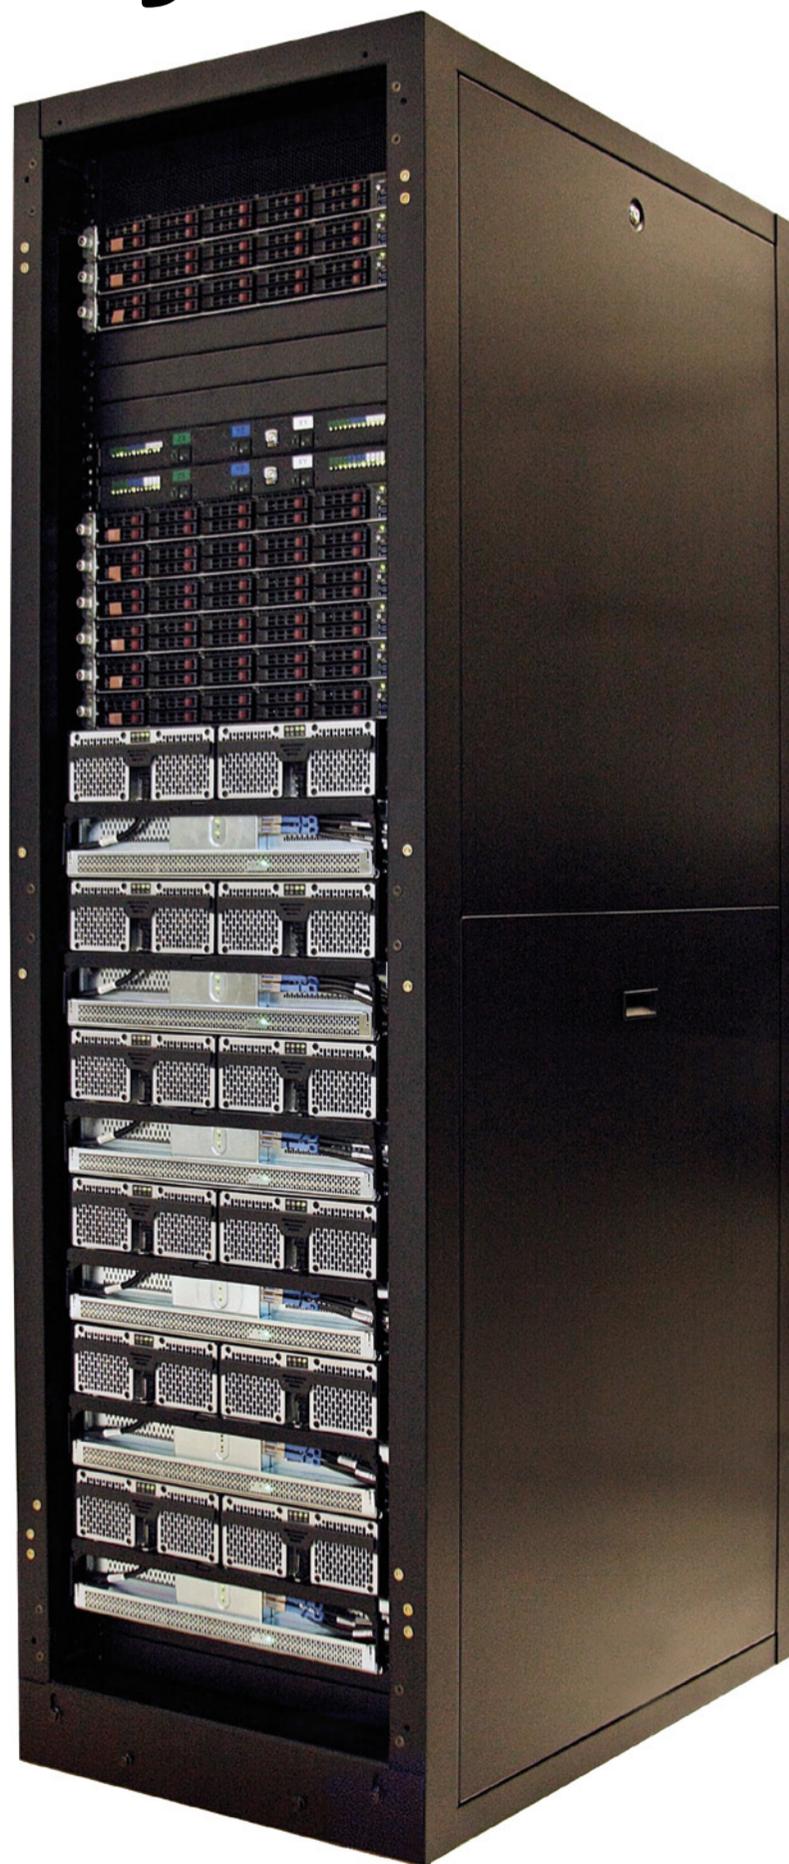


# HGST's Active Archive platform offers 4.7 Petabytes of object storage

**HGST taps vertical innovation in data storage and object storage software, delivering a system that rivals the Scale and TCO of traditional cloud infrastructure**

**F**or data centers battling massive data growth, flat budgets and limited IT resources, HGST announced its Active Archive System, an object storage system that enables businesses to easily store, retain and harness the power of data. Delivering 4.7 petabytes (PB) of raw data storage in a single rack, the HGST Active Archive System is a ready out-of-the-box object storage system that helps data centers easily evolve from siloed data storage to cloud-scale active archiving. For data that is past the create and modify phase of its life, and in need of long-term retention with fast access, the Active Archive System provides unprecedented levels of accessibility, scalability, simplicity and affordability. Its breakthrough TCO beats the white box economics of traditional cloud infrastructure to enable organizations to store and access more data, driving deeper insights and ultimately unlocking the value of their data.

“HGST’s Active Archive System supports our strategy to bring higher value to the market so customers can transform their data into busi-





**« Data has become the currency of the new economy. The ability to store, retain and interact with data enables businesses to gain new insights to help boost productivity, make better decisions and even predict the future. We recognized this tectonic shift. »**

ness value,” said Dave Tang, senior vice president and general manager, Elastic Storage Platforms Group, HGST. “Data has become the currency of the new economy. The ability to store, retain and interact with data enables businesses to gain new insights to help boost productivity, make better decisions and even predict the future. We recognized this tectonic shift and have delivered the only solution that can offer superior economics with industry leading rack density, resiliency, and access at petabyte-scale. Leveraging our investment in object storage software and our unique approach of vertical innovation, the Active Archive System is positioned to help users unlock the power of their data.”

### **VERTICAL INNOVATION POWERS THE THIRD PLATFORM**

The Active Archive System has been singularly designed to address the need for rapid access to massive data stores. HGST’s recent acquisition of Amplidata enables the delivery of advanced object storage software, providing the foundation for a resilient object storage system that delivers the highest level of data durability in the industry, surviving an entire data-center outage when deployed across multiple sites. Utilizing HGST’s most advanced, 8TB, second-generation HelioSeal hard drives, coupled with tightly integrated, tuned and optimized hardware, the scale-out object storage system delivers the performance, efficiency and scale required for public and private cloud data centers. The combination of industry leading power efficiency at one watt per TB, and an ac-

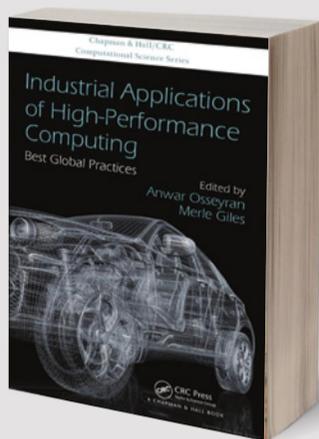
quisition cost that beats white box economics enables the system to deliver breakthrough TCO that is 50-70 percent lower than alternative open source and commercial object storage systems. Through this tight end-to-end system integration, HGST has delivered a 3-5X increase in density and power efficiency compared to other scale-out solutions.

The uniqueness of the HGST Active Archive System is its Simplicity at Scale foundation. It has never been as easy and cost-effective to scale-out in multiple petabyte increments so that businesses can focus resources on growing their revenue instead of managing their infrastructure. The HGST Active Archive System is a plug-and-play, S3 compliant scale-out object storage system; it only requires network and power connections to each rack to be put to work. The Active Archive System offers limitless scaling without the costly need for replication, and its open interface supports ease of integration and flexibility as needs evolve.

“Today, HGST has delivered on its strategy to expand into higher value storage platforms and systems,” said John Rydning, Research Vice President, IDC. “HGST has intersected a critical juncture in the market as companies look to accelerate the rollout of third platform solutions. The plug-and-play simplicity of an integrated hardware platform matched with the strategic acquisition of Amplidata’s object cloud storage software has resulted in an innovative solution that addresses the rapidly evolving needs of cloud service providers and data-intensive organizations.”



## BOOKS



## Industrial Applications of High-Performance Computing: Best Global Practices

This book offers a global overview of high-performance computing (HPC) for industrial applications, along with a discussion of software challenges, business models, access models (e.g., cloud computing), public-private partnerships, simulation and modeling, visualization, big data analysis, and governmental and industrial influence. Featuring the contributions of leading experts from 11 different countries, this book provides a brief history of the development of the supercomputer, describes the supercomputing environments of various government entities in terms of policy and service models. The book includes a case study section that addresses more subtle and technical aspects of industrial supercomputing, shows how access to supercomputing matters, and how supercomputing can be used to solve large-scale and complex science and engineering problems, emphasizes the need for collaboration between companies, political organizations, government agencies, and entire nations. **Chapman and Hall/CRC, Anwar Osseyran & Merle Giles, 410 pages (£63.99)**

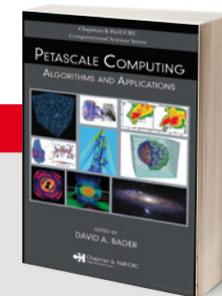
## High Performance Computing: Programming and Applications

This book presents techniques that address new performance issues in the programming of high performance computing (HPC) applications. Omitting tedious details, the book discusses hardware architecture concepts and programming techniques that are

the most pertinent to application developers for achieving high performance. Even though the text concentrates on C and Fortran, the techniques described can be applied to other languages, such as C++ and Java. After discussing architectural and software challenges, the book outlines a strategy for porting and optimizing an existing application to a large massively parallel processor (MPP) system. With a look toward the future,

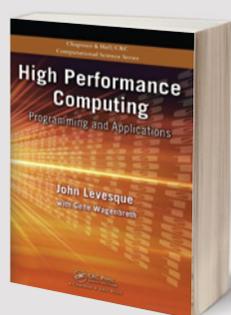
it also introduces the use of general purpose graphics processing units (GPGPUs) for carrying out HPC computations. A companion website at [www.hybridmulticoptimization.com](http://www.hybridmulticoptimization.com) contains all the examples from the book, along with updated timing results on the latest released processors.

**Chapman and Hall/CRC, John Levesque & Gene Wagenbreth, 244 pages (£63.99)**



## Petascale Computing: Algorithms and Applications

This book captures the state of the art in high-performance computing algorithms and applications and addresses the challenging problems of developing application codes that can take advantage of the architectural features of the new petascale systems in advance of their first deployment. The book illustrates how petascale computing can be applied to space and Earth science missions, biological systems, weather prediction, climate science, disasters, black holes, and gamma ray bursts. It details the simulation of multiphysics, cosmological evolution, molecular dynamics, and biomolecules. The book also discusses computational aspects that include the Uintah framework, Enzo code, multithreaded algorithms, petaflops, performance analysis tools, multilevel finite element solvers, finite element code development, Charm++, and the Cactus framework. **Chapman and Hall/CRC, David A. Bader, 616 pages (£63.99)**





## MOOCS

STARTING SOON

# Software Defined Networking

In this course, you will learn about software defined networking and how it is changing the way communications networks are managed, maintained, and secured. This course introduces software defined networking, an emerging paradigm in computer networking that allows a logically centralized software program to control the behavior of an entire network. Separating a network's control logic from the underlying physical routers and switches that forward traffic allows network operators to write high-level control programs that specify the behavior of an entire network, in contrast to conventional networks, whereby network operators must codify functionality in terms of low-level device configura-

tion. Logically centralized network control makes it possible for operators to specify more complex tasks that involve integrating many disjoint network functions (e.g., security, resource control, prioritization) into a single control framework, allowing network operators to create more sophisticated policies, and making network configurations easier to configure, manage, troubleshoot, and debug.

**Courses start :** May 25, 2015

**Length:** 8 weeks

**Effort:** 7 – 10 hours per week

**Institution:** University of Princeton

**Languages:** English

**Link :** <http://fr.coursera.org/course/sdn1>

# Scalable Machine Learning

Learn the underlying principles required to develop scalable machine learning pipelines and gain hands-on experience using Apache Spark. Machine learning aims to extract knowledge from data, relying on fundamental concepts in computer science, statistics, probability and optimization. Learning algorithms enable a wide range of applications, from everyday tasks such as product recommendations and spam filtering to bleeding edge applications like self-driving cars and personalized medicine. In the age of 'Big Data,' with datasets rapidly growing in size and complexity and cloud computing becoming more

pervasive, machine learning techniques are fast becoming a core component of large-scale data processing pipelines. This course introduces the underlying statistical and algorithmic principles required to develop scalable real-world machine learning pipelines. We present an integrated view of data processing by highlighting the various components of these pipelines, including exploratory data analysis, feature extraction, supervised learning, and model evaluation. You will gain hands-on experience applying these principles using Apache Spark, a cluster computing system well-suited for large-scale machine learning tasks. You will

implement scalable algorithms for fundamental statistical models (linear regression, logistic regression, matrix factorization, principal component analysis) while tackling key problems from various domains: online advertising, personalized recommendation, and cognitive neuroscience.

**Courses start :** June 29, 2015

**Length:** 5 weeks

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

**Institution:** University of Berkeley

**Languages:** English

**Link :** <http://www.edx.org/course/systematic-program-design-part-1-core-ubcx-spd1x>



## MOOCS

ALREADY STARTED

# Intro to Parallel Programming : Using CUDA to Harness the Power of GPUs

Learn the fundamentals of parallel computing with the GPU and the CUDA programming environment! In this class, you'll learn about parallel programming by coding a series of image processing algorithms, such as you might find in Photoshop or Instagram. You'll be able to program and run your assignments on high-end GPUs, even if you don't own one yourself. You'll master the fundamentals of massively parallel computing by using CUDA C/C++ to program modern GPUs. You'll learn the GPU programming model and architecture, key algorithms and parallel programming patterns, and optimization tech-

niques. Your assignments will illustrate these concepts through image processing applications, but this is a parallel computing course and what you learn will translate to any application domain. Most of all we hope you'll learn how to think in parallel.

**Courses start :** Anytime, Self-Pace

**Length:** 12 weeks

**Effort:** 6 hours per week

**Institution:** NVIDIA

**Languages:** English

**Link :** <http://www.udacity.com/course/intro-to-parallel-programming--cs344>

## Cryptography I

Learn about the inner workings of cryptographic primitives and how to apply this knowledge in real-world applications. Cryptography is an indispensable tool for protecting information in computer systems. This course explains the inner workings of cryptographic primitives and how to correctly use them. Students will learn how to reason about the security of cryptographic constructions and how to apply this knowledge to real-world applications. The course begins with a detailed discussion of

how two parties who have a shared secret key can communicate securely when a powerful adversary eavesdrops and tampers with traffic. We will examine many deployed protocols and analyze mistakes in existing systems. The second half of the course discusses public-key techniques that let two or more parties generate a shared secret key. We will cover the relevant number theory and discuss public-key encryption and basic key-exchange. Throughout the course students will be exposed to many exciting open problems

in the field. The course will include written homeworks and programming labs. The course is self-contained, however it will be helpful to have a basic understanding of discrete probability theory.

**Courses start:** April 20, 2015

**Length:** 6 weeks

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

**Institution:** University of Stanford

**Languages:** English

**Link :** <http://fr.coursera.org/course/crypto>



# 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.

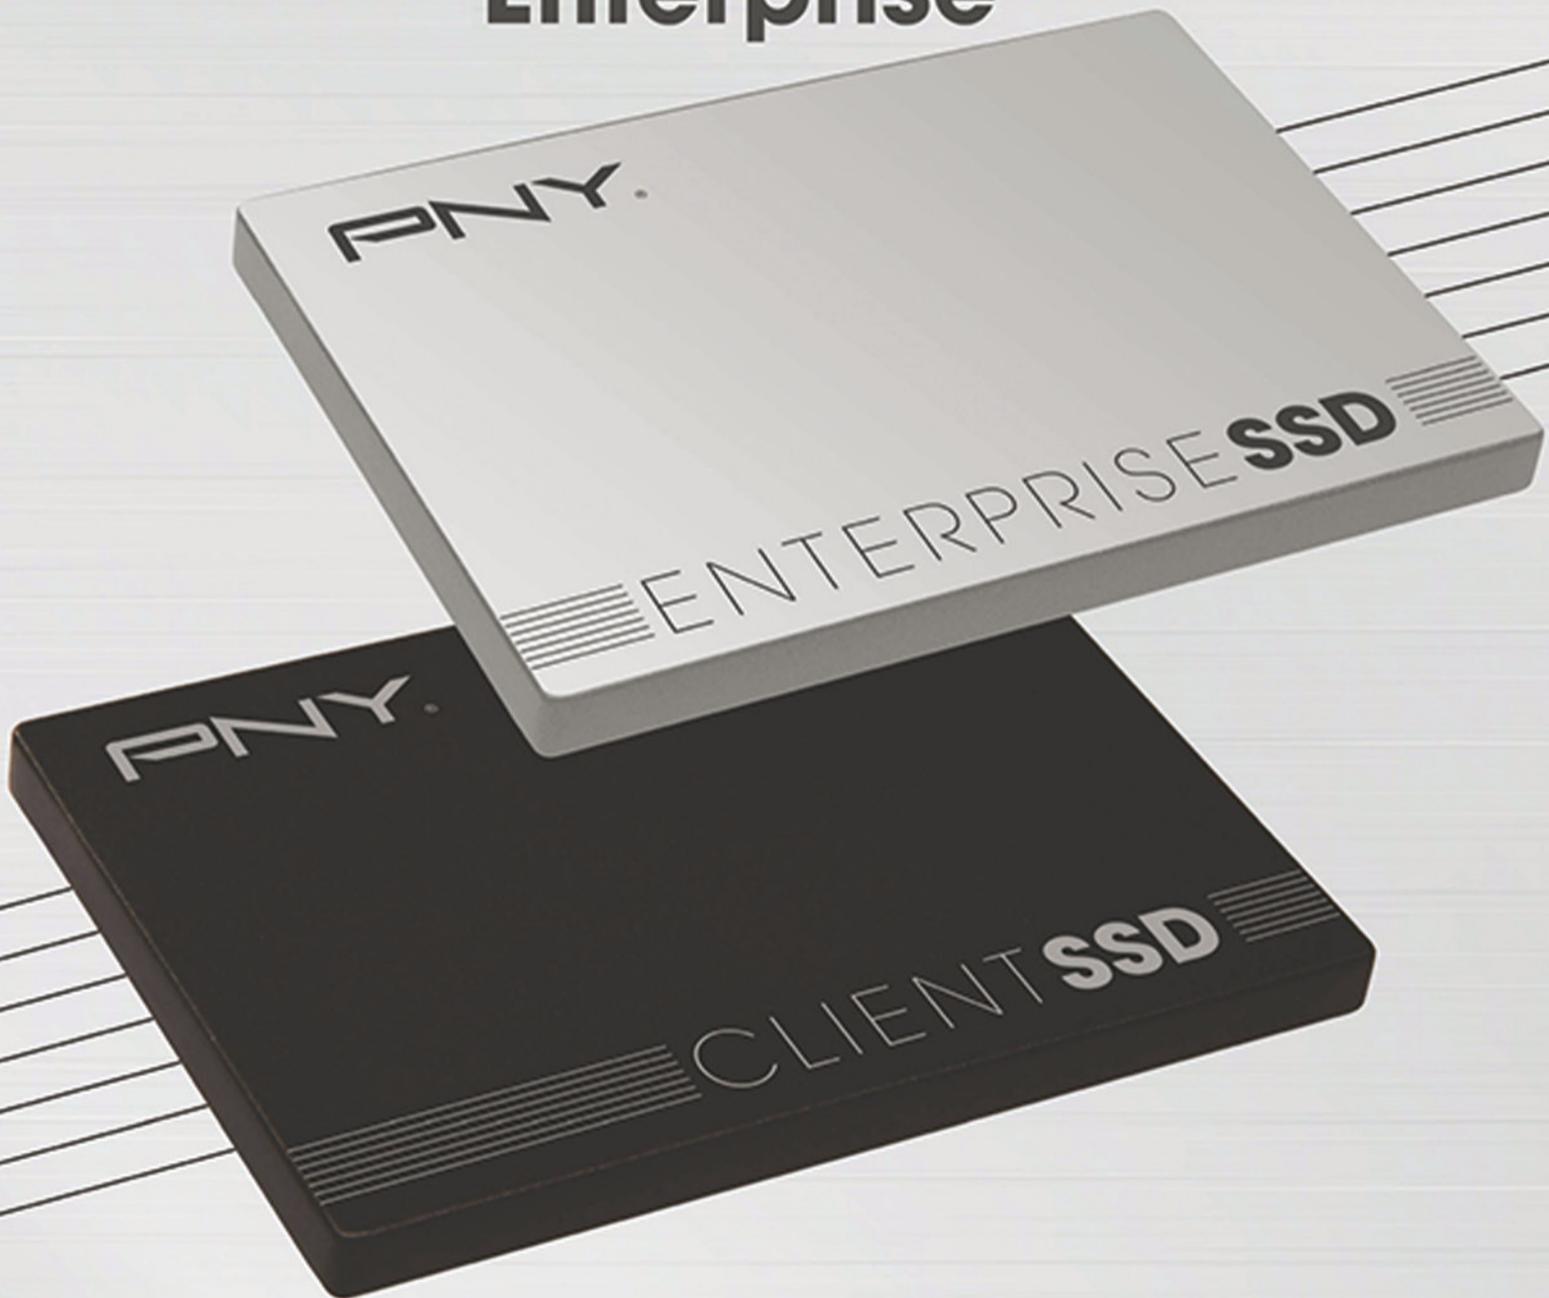
# PNY

## 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)



# STORAGE VIRTUALIZATION

## WHAT ARE THE BENEFITS ?



**A fast growing trend these last years has been virtualization. So much so that almost every IT product family can be virtualized these days. Networks, Storage, Servers and even Desktops and workstations through VDI-like mechanisms. We have decided to focus on storage, since it pretty much underlines the other market segments.**



**BENEFIT #1** **Flexibility** A single management console provides optimum visibility and knowledge of the storage infrastructure to help identify underperforming storage areas and fast enterprise storage devices, while adding new features of live migration. An impossible feat with traditional storage.

**T**oday, virtualization is a subtle blend of hardware and software. On the hardware side, the processors are evolving dramatically, with the number of physical and logical cores constantly increasing, which are immensely helpful for virtualization tasks. The third-generation Xeon processors (Haswell) have from 4 up to 18 physical cores for the Xeon E5-2699 v3 cpu. This multiplies all capabilities for a server to run dozens of virtual machines (VMs) under optimal conditions. For a two- or four-processor server, the number of active VMs that can be executed can be from dozens to hundreds. In addition to this horsepower, specific instructions called VT (for Virtualization Technology) are built into the processors that improve VM performance by granting them direct addressing capabilities to the hardware resources of the host machine via the PCI Express bus.

On the software side, the two leading virtualization platforms are Microsoft Hyper-V and VMware vSphere (also known as ESX Server). The former is a type 2 hypervisor, running over a functioning operating system (Windows Server 2012 or 2012 R2). The latter is a type 1 hypervisor, independent of any OS and running natively on the host server.

For this review, we turned to a precursor, DataCore, that invented the concept in the late 90s with their application SANsymphony, which just celebrated its 10th edition. Pascal Le Cunff, South Europe Manager of DataCore, helped us pinpoint the main benefits of storage virtualization

The primary benefit of this model is the separation of service and data layers. According to Mr. Le Cunff, storage virtualization is bene-



**Pascal Le Cunff, DG France Datacore**

ficial to any company. First on the application side within VMware environments, by enabling the establishment of an application cluster to improve performance, and secondly by adding fault tolerance. Storage virtualization also helps improve load distribution. On the hardware side, DataCore's approach improves memory performance with unprecedented acceleration by providing:

- independence from the hardware,
- a manageable quality of service (QOS),
- consolidation of input / output operations in memory.

## **MANAGING HETEROGENEOUS INTERFACES**

This virtualized approach allows freedom of choice with respect to hardware. No matter which brands of resources are dedicated to storage, they can all be managed from a single dashboard. This also applies to the varying interfaces that have often evolved over



**BENEFIT # 2** **Performance** With virtualized storage, performance does not only depend on the hardware. Better and more capable software such as SANsymphony v10 help increase the existing storage device's lifespan and access times with an in-memory cache logic that boosts performance for equal material!

the years. A virtualization solution allows the seamless coexistence - and intercommunication ! – of iScsi and Fibre Channel interfaces for instance. This is far from being the sole benefit : access to recent and innovative features such as auto-tiering (moving the most frequent-accessed data from traditional hard disks to fast flash-memory based SSDs), CDP (continuous backup) and thin Provisioning are also present. Added to this is a connection ability – along with snapshot features – with Cloud spaces like Amazon Web Services and Microsoft Azure Cloud platform.

### **PROVIDE STORAGE RULES TO ENSURE OPERATION**

Being able to manage QOS over the entire storage infrastructure allows the prioritization of input / output operations according to application needs. Financial or HR applications can be categorized as being critical and therefore benefit from guaranteed and priority access...despite the storage infrastructure load while still remaining within the permissible limits.

Another example is a company's business database, that through QOS manageability can keep a guaranteed flow and thus remain exploitable during favorable conditions or maintain a minimum percentage of bandwidth even under heavy load, thus ensuring high availability. As Mr. Le Cunff states, it is not a question of data volume, citing for example the case of a food operations company whose database is only 300GB, but hosts critical traceability data which must be accessible at all times with optimal flow rates.

### **HARDWARE INDEPENDENCE TO COMBINE EXISTING AND NEW TECHNOLOGIES**

An enterprise's infrastructure is almost always the accumulation of heterogeneous resources using separate connectors: iScsi SAN, FC SAN, disks, disk arrays... And as often happens, whenever a manufacturer decides to change its product line, connection and compatibility questions arise. Take for example a client who has purchased an IBM drive bay 3 years ago, a model since removed from the vendor's catalog. It very probably is not supported anymore, either on the hardware or the software compatibility side. Another typical example is an enterprise needing to evolve its existing infrastructure started years ago, but being forced to stay with the same technology, often at exorbitant costs (new equipment is expensive) and without any performance benefit. A virtualized storage architecture solves these two scenarios and provides access to the latest technologies without disturbing the existing hardware infrastructure. This is also the opportunity to equip dumb storage arrays (since the latter is offset at the storage hypervisor) and thus at lower cost.

### **MIGRATION WITHOUT THE HEADACHE**

Migration effects are also eased in a consolidation phase, during which many low capacity drives are replaced with more recent, higher capacity drives. Another benefit is energy consumption falls sharply! Usually data migration requires an IT department to first save data, and thus establish an equivalent storage space. Such a setup takes time, is cumbersome and risky, not to mention that it will need to keep track of the data transfers



**BENEFIT # 3** **Hardware Independence** Only large companies have an homogeneous storage infrastructure. Most of the time, enterprises pick the best of the flock, ending up with equipment from several vendors over time. Adding the appropriate software layer allows new supervision and management features to be extended and helps to avoid sacrificing older or less efficient resources.

between the backup and restore operations. When a company has a virtualized storage infrastructure, the abstraction layer that isolates the hardware allows, through a feature known as substitution, to copy the data to new storage spaces without interruption - and afterwards, leaving the freedom to keep the old, less effective devices for less critical needs or to withdraw them from service altogether.

### **AN OPENING TO THE SIMPLIFIED CLOUD AND A THREE-SPEED STORAGE**

Harnessing the Cloud as a third storage space for overflow purposes or BCP/BRP (Business Continuity Plan /Business Recovery Plan) is much simpler with a virtualized infrastructure. Depending on the speed of the network link, it is possible to maintain a synchronous (broadband) or asynchronous (low speed) backup mechanism. Even better, the Cloud is also integrated into the single management console in order to keep overview as clear and complete as possible. In addition to the seamless integration of cloud storage, a virtualized hypervisor enables visibility of storage space types. These are typically divided into three storage levels that are set up within the company:

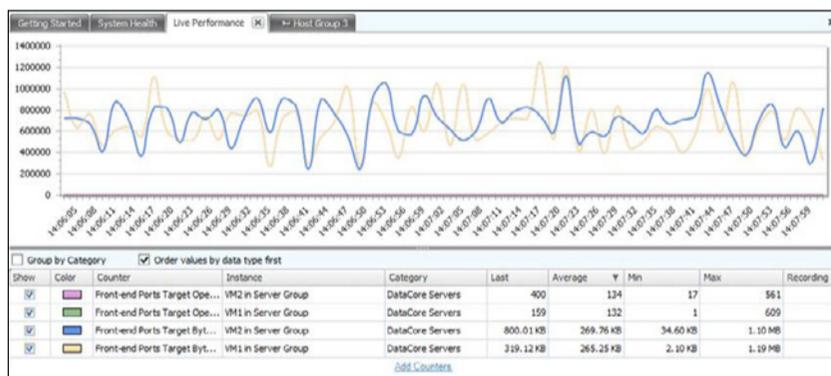
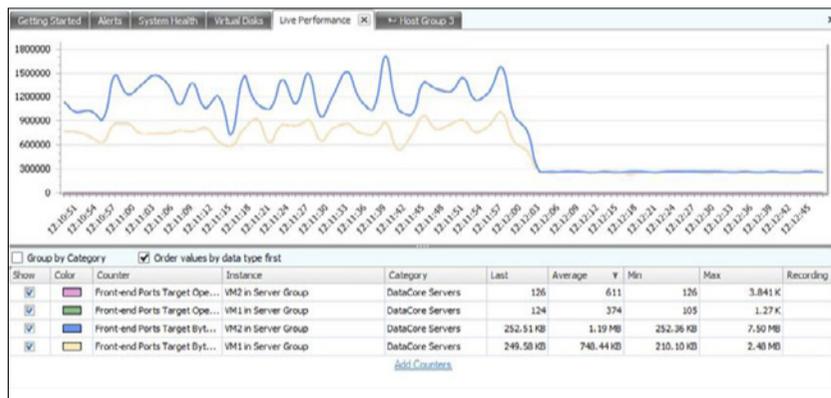
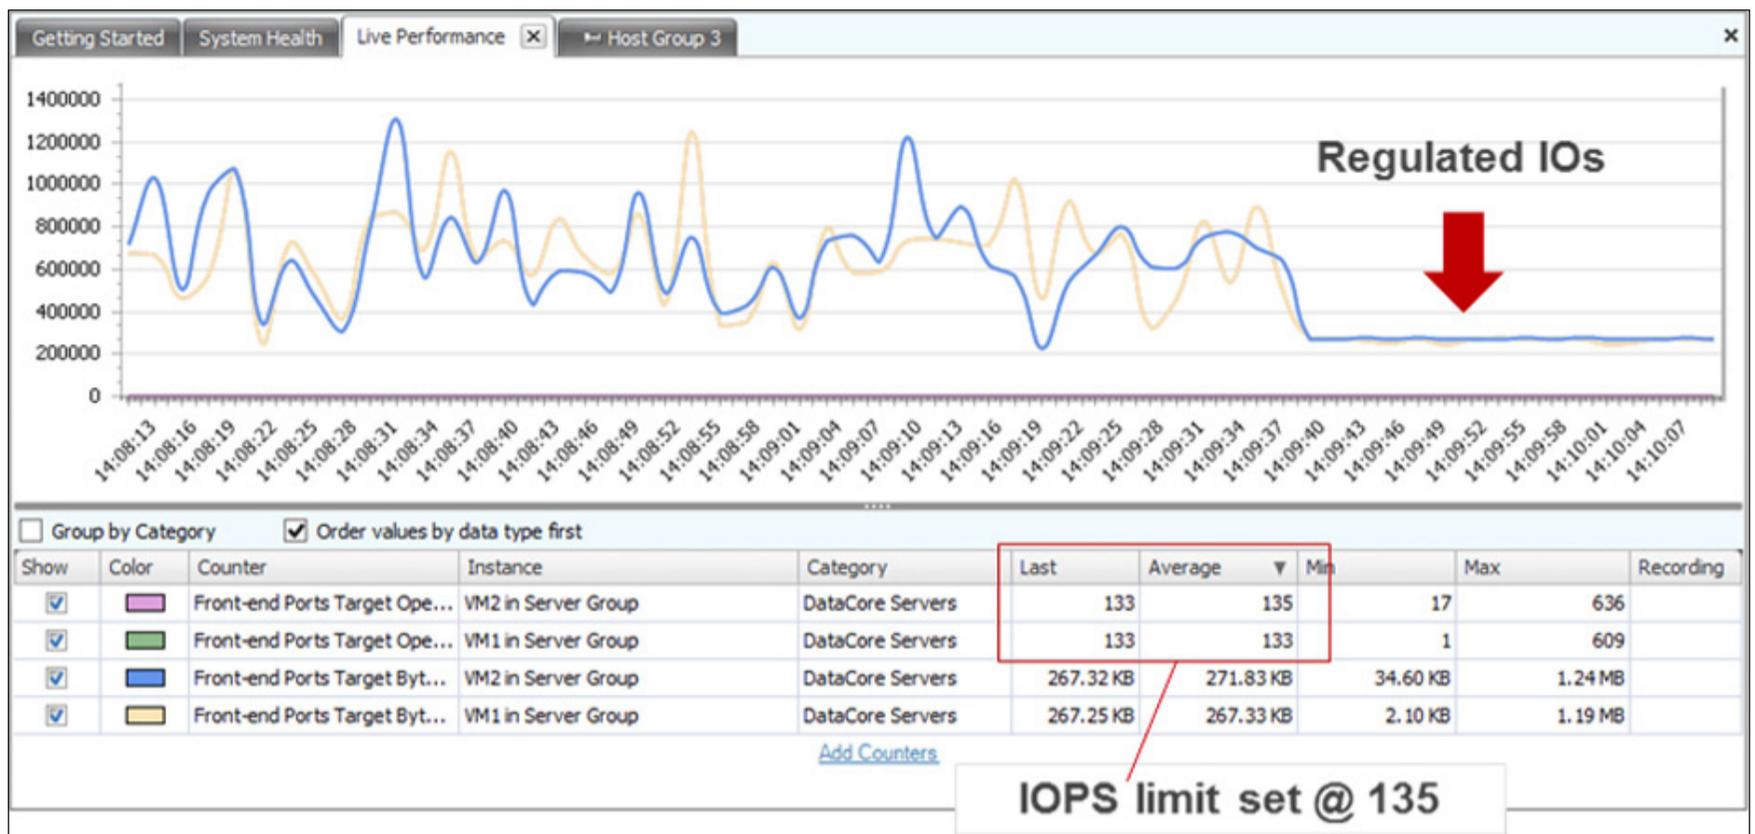
- Slow storage spaces are dedicated to local backups.
- Efficient storage spaces are reserved for office and production applications (accountancy, HR).
- High-performance storage-based SSD or Flash are reserved for critical business applications.

### **THE MOST EFFICIENT STORAGE FOR THE MOST CRITICAL MISSIONS.**

These three types of storage spaces can be combined with QOS rules to ensure their availability and thus respond to the changing needs of a company. For example, a business database which increases from 50 GB to 700 GB can be migrated to more efficient and capacitive storage resources. Not forgetting a golden rule for storage: better to spread data across multiple disks (or drive arrays) of lower capacity than putting them into a single drive bay or disk, for obvious reasons of redundancy and security. The other benefit is to provide better load balancing across multiple arrays and disks (it is better to use three 300GB disks than a single 1TB drive). Some application contexts require the fastest available resources. The greatest constraints are those required by virtualization workspaces (VDI and DaaS). Software such as SANsymphony can identify the fast and slow zones on multiple storage spaces. And also identify potential bottlenecks and move applications accordingly onto higher performance storage pools. Another added bonus of such visibility is to benefit from statistics to use predictive failure algorithms in order to avoid the risks of an operation halt.

### **A RULE-ORIENTED STORAGE STRATEGY**

The VMware approach is to define a rule-based storage: take the physical storage and add an abstraction layer on top of it to pool all the existing storage resources. Thereafter, the VMware administrator uses these resources to ensure quality of service. To achieve this, the rules are defined first (extremely fast storage



and storage highly available or redundant storage, etc.). In this case, the VMware administrator in charge of deploying virtual machines will not need to know which storage it will be actually physically attached to, but will indicate the needs of the of the deployed application (highly available and/or high-performance). Therefore, the application is deployed with the defined quality of service rules. Thus ensuring that the storage strategy is consistent with the real needs of applications or VMs.

## TWO SOLUTIONS TO THE SAME OBJECTIVE

Specifically, there are two solutions at VMware to implement this strategy. The first product was released last year and is called VSAN. It is a technology embedded directly into the VMware hypervisor that lets you define QoS rules by using local server disks. Thus the administrator will create a pool of servers and physical disks will be concatenated to form a logical unit. There, the QoS rules will be defined and the administrator of the VM will use and apply the available resources based on this approach. Relying on wizards to be deployed, no dedicated storage administrator is needed. A concrete, practical implementation of this concept is EVO: RAIL, a converged infrastructure consisting of servers, networking and storage using the VSAN technology. Another new technology called Virtual Volume has appeared (an extension of the rule-based VSAN), but is applied to external systems belonging to partners offering SAN or NAS solutions. Is is also based on rules (redundancy, performance, availability, etc.) but the administrator can use external and distant SANs or NAS as it can do within a VSAN (without referring to LUNs or volume). **JOSCELYN FLORES**



# 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



# NVIDIA Quadro M6000

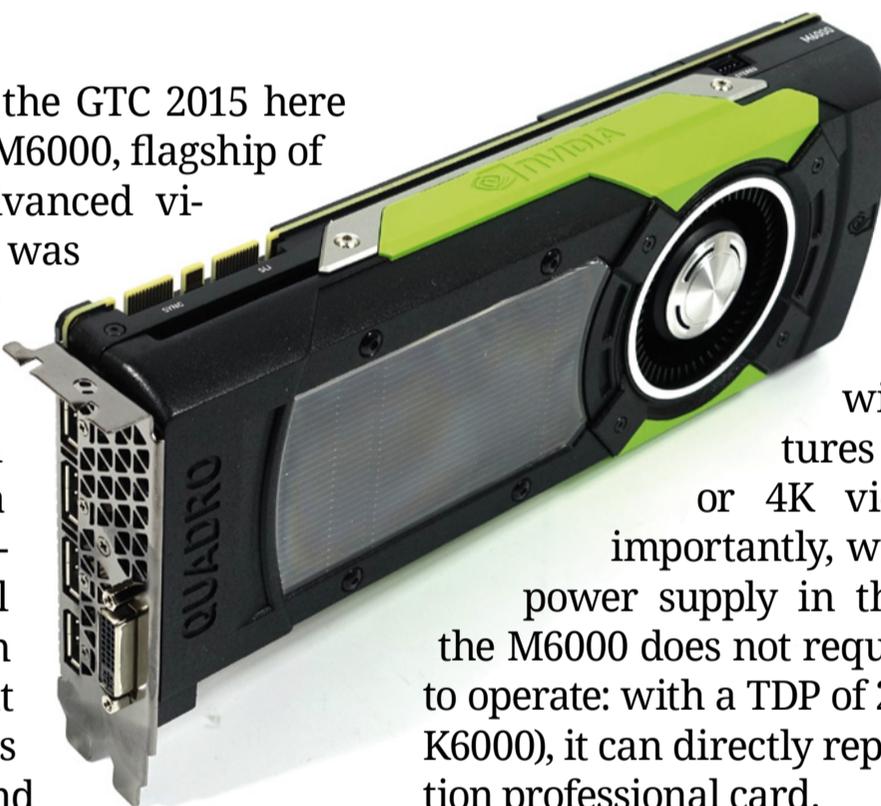
## A milestone. Or not.

**First card to use the Maxwell architecture dedicated to visualization, NVIDIA's Quadro M6000 shines in application environments optimized for NVIDIA technologies. The speedup / price ratios is more questionable...**

**A**s promised at the GTC 2015 here is the Quadro M6000, flagship of NVIDIA in advanced visualization. It was about time some would say, given the recent AMD's advances in this area, with cards such as FirePro W91xx with a very interesting cost / performance ratio. This small delay on the competition is mainly due the fact that the M6000 is designed, as its name suggests, around the brand new Maxwell architecture. So much for the Kepler architecture, which found its culmination on workstation cards with the K6000.

Developed primarily for HPC in mind, Maxwell brings to the M6000 a real added value in terms of acceleration of rendering and encoding tasks. As all our test reports show, will be discussed in more detail below, using an M6000 is like having a double Titan X. But first let's review the technical details of this first visualization oriented implementation.

Powered by a GM200 GPU, the Quadro M6000 embeds 3,072 CUDA cores and 12 GB of RAM Graphics to 317 GB / s of bandwidth (2 800/288 on the K6000), with a theoretical capacity of 7 TFLOPS in single accuracy (5.2 with K6000). For comparison, the FirePro W9100 is located



roughly in the same zone, with 2816 cores to 5.24 TFLOPS SP, but with 16 GB of memory and 320 GB / s, with 20% more of textures capacity treatment or 4K video streams. More importantly, with a properly sized power supply in the host workstation, the M6000 does not require anything special to operate: with a TDP of 250 W (225 W for the K6000), it can directly replace an older generation professional card.

On the interface side, in addition to a DVI port, the Quadro M6000 features four DisplayPort outputs, which can drive four 4K monitors or two 5K monitors. The Scalable Visual Solutions (SVS) function also allows it to control video walls or multi-node visualization clusters. Additionally, the Quadro Link function allows to chain four cards (Quadro only, no GeForce), which can manage up to 16 devices on a single display logic.

In our benchmarks, the Quadro M6000 does not disappoint. On basic tests like Cinebench, this card reached the highest framerate (110.7 fps) we have recorded so far on an NVIDIA card. It is also not far from the absolute record held by the AMD FirePro W9100 (113.8 fps).

Our further testings, whether using Maya, SolidWorks or SPECviewperf (see a complete



## Benchmark results

### SPECVIEWPERF 12.0.02

The benchmark results show the remarkable homogeneous conception of the Quadro M6000, with outstanding benchmark results making it perfectly suitable for CAD/CAM application environments.

Viewset	Composite	Window
catia-04	<b>117.55</b>	1900 x1060
creo-01	<b>93.32</b>	1900 x1060
energy-01	<b>13.18</b>	1900 x1060
maya-04	<b>107.63</b>	1900 x1060
medical-01	<b>59.83</b>	1900 x1060
showcase-01	<b>86.85</b>	1900 x1060
snx-02	<b>168.83</b>	1900 x1060
sw-03	<b>138.08</b>	1900 x1060



### CINEBENCH 15

The Quadro M6000 reaches a CineBench 15 score of 110 fps. About the double of a Quadro K4000 (67 fps).

Cinebench OpenGL

**110fps**

**Test configuration** ASUS P9X79E-WS motherboard (Bios 1602) Intel Xeon 1620v2 **Processor** (4 Cores/8 Threads - 3.7GHz) **Memory** 4 Crucial DIMMs x 4Gb DDR3 1600MHz (unregistered, unbuffered) **Storage** 120Gb SSD (Crucial M500 SATA) + 2Tb SATA Disk (Seagate Constellation ES3) **Power Supply** BeQuiet! DarkPower Pro 1200W, **Chassis** Corsair Obsidian 900D, PNY Quadro M6000 (Bios 84.00.1b.00.01) **OS** Windows 7\_64 SP 1 Pro (OEM) **Driver** nVidia Ver 347.88

list of benches below) show however strong speedups on applications able to use the on-board power when optimized for NVIDIA drivers. This is particularly noticeable in environments such as Catia, where the M6000 shows up to 15% more efficiency than any other card. If the intended use is mainly based on applications developed on CUDA, you can even expect performance boosts in the range of 30-40%

tion» as the ninjas say. Card types like Quadro K4x or FirePro W7x / 8x show a more attractive price / performance ratio.

The reason of our conclusion above is that the power of the M6000 is hampered by its retail price. It certainly is powerful enough to handle both rendering calculations and rendering itself, which ultimately means saving on an additional HPC accelerator, but with a price

## The power of the M6000 is hampered by its retail price.

compared to a K6000 (figures recorded via FluidMark and ratGPU, which are due to both the new silicon present on the board, and by optimizing the wired logic in the controllers and Drivers), which may be enough to justify the return on such an investment. However, in configurations where the application does not scale enough to fully use all the available horsepower, the usage profits plead for the «non-ac-

tag of more than € 4,500 (public prices have not yet been defined at the time of writing), the investment is only justified by a real application need. That is to say, by a software environment optimized for NVIDIA technologies. Because the AMD FirePro W9100 competitor remains positioned well below € 4,000... **MARC PARIS**



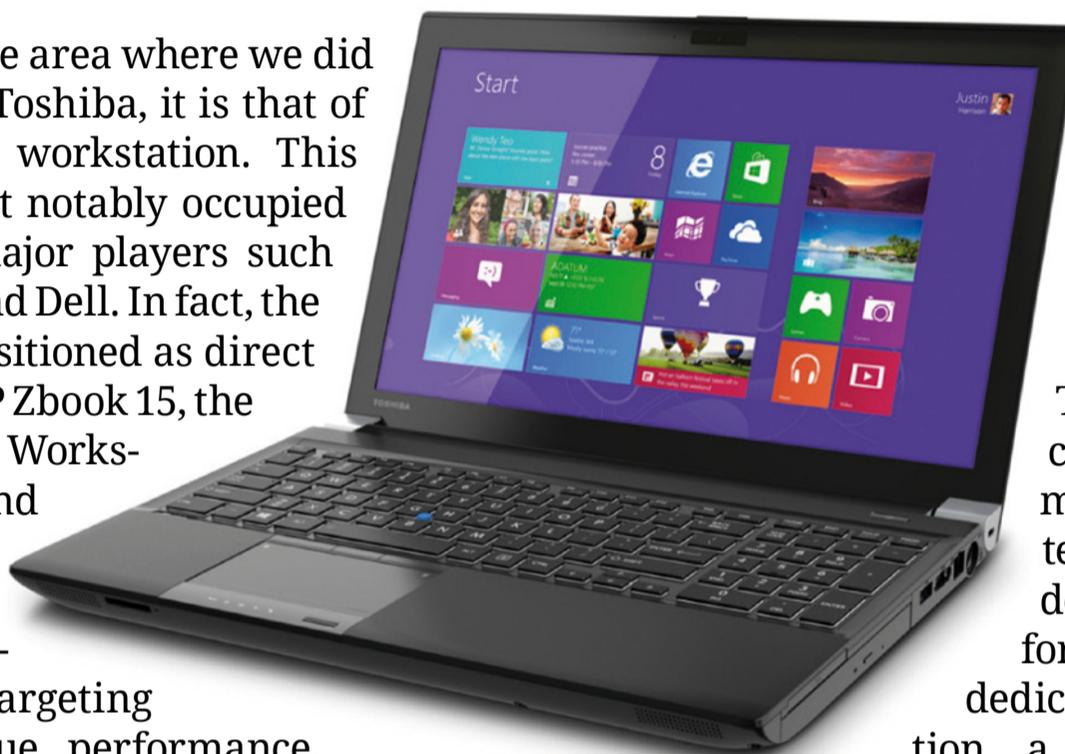
# Toshiba Tecra W50

## Toshiba innovates with a robust mobile workstation

If there is one area where we did not expect Toshiba, it is that of the mobile workstation. This area is most notably occupied by three major players such as HP, Lenovo and Dell. In fact, the Tecra W50 is positioned as direct competitor of HP Zbook 15, the Dell Precision Workstation M4800 and the Thinkpad W540. The Tecra W50 is a portable 15-inch targeting users who value performance, build quality, scalability and extended warranty options. This machine sets the tone when opening it : its sturdy frame reassures, with a choice of thick textured plastics and thick exposed steel hinges.

### A RICH AND CONSISTENT EQUIPMENT

The equipment is consistent with its positioning. It is powered by a powerful Intel Core Quad i7-4800MQ heart, 32 GB of GDDR5 memory, a 1TB drive and an NVIDIA Quadro K2100M graphics card. The full HD IPS pre-calibrated screen is excellent. The equipment includes a DVD burner, an HD webcam with microphone, a card reader and Wi-Fi connectivity, Bluetooth 4.0, DisplayPort, VGA, audio, Gigabit Ethernet and USB 3.0 ports 2Two and three USB 2.0 ports. And if it were not enough, the High-Speed Port Replicator III adds DVI-D and HDMI video outputs.



### SAFETY FIRST

The Toshiba Tecra W50 has a very complete safety equipment. Besides the usual password, Toshiba provides security applications to manage the integrated fingerprint reader, a Trusted Platform Module 1.2 chip dedicated to data encryption, a SmartCard reader and Computrace monitoring for theft. A fall sensor stops the hard drive in case of a fall detection to avoid data loss. The W50 also features a splash-resistant keyboard and a shockproof chassis, providing sturdiness enough to withstand drops from a height of 76 cm without damage. A Kensington Lock is of course included, anchored on one of its thick steel hinges.

The standard warranty is one year. It can however be extended to three years for a small fee (one hundred euros) provided it is subscribed within thirty days of purchase. The price of Tecra W50 is 2900 €.

RAMON LAFLEUR

HPC Bench  
Global index



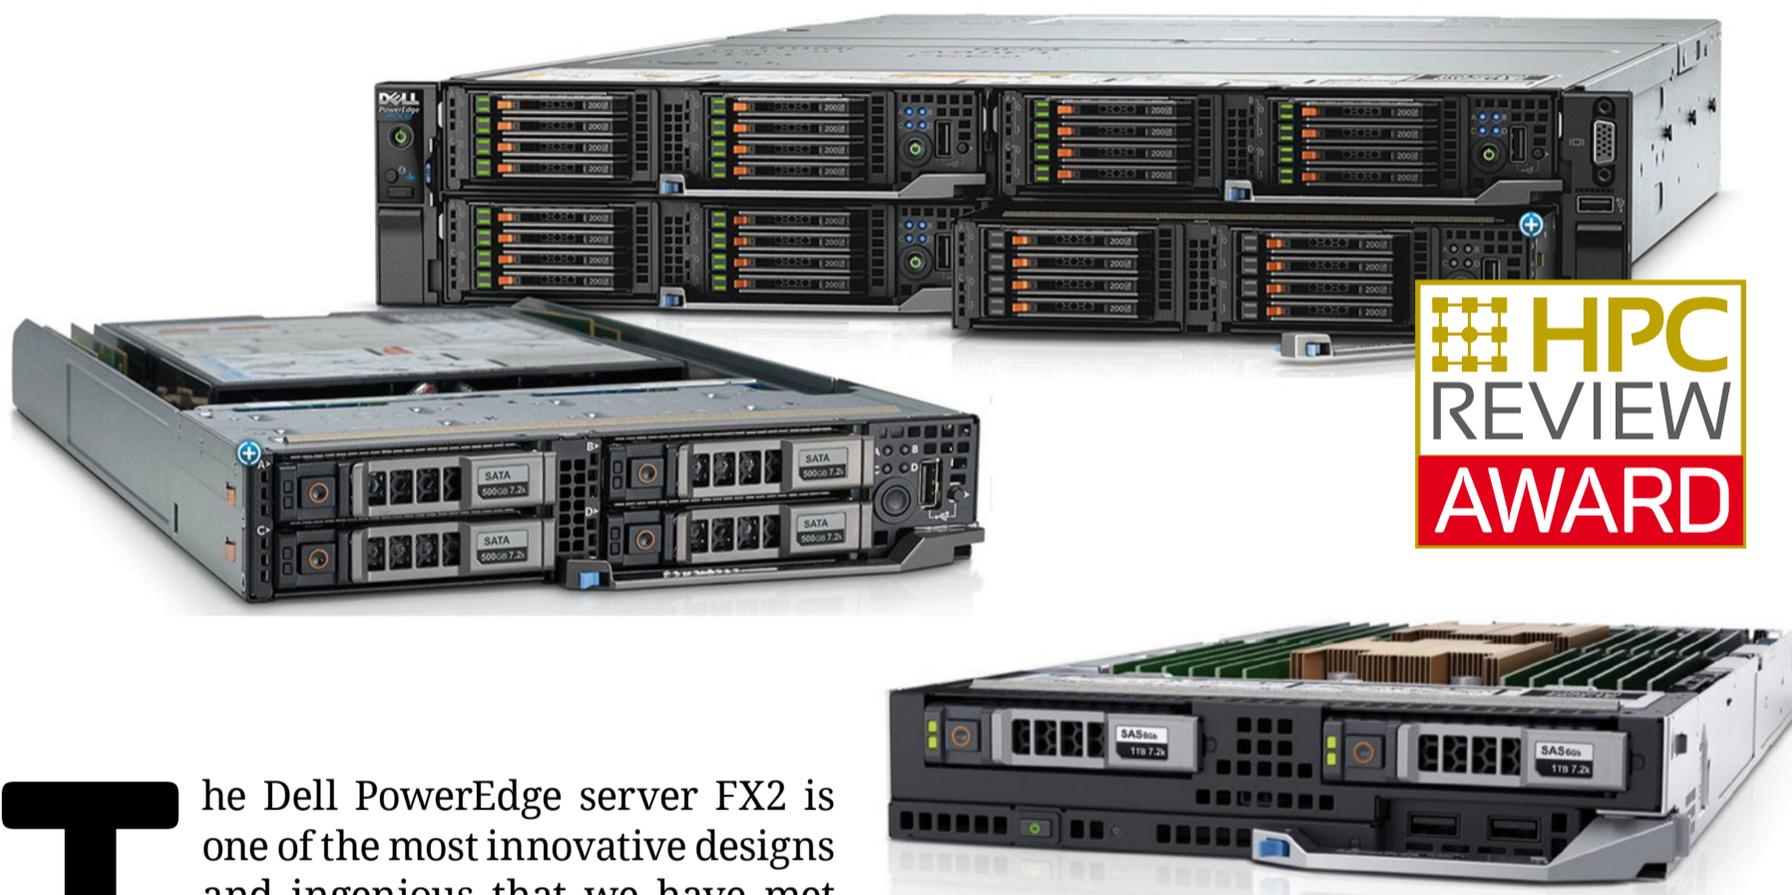
6 820

**Features:** Intel Core processor i7-4800MQ. 32 GB of memory. 1TB hard drive. NVIDIA Quadro K2100M Graphics Card. Full HD IPS screen with 1920x1080 pixels resolution



# Dell PowerEdge FX2

Manufacturers have pushed the concept of converged server quite far, with many benefits. Dell is no exception, with its PowerEdge server FX2 which represents the latest step in terms of integration. And what integration!



**T**he Dell PowerEdge server FX2 is one of the most innovative designs and ingenious that we have met since the introduction of the HP Moonshot concept (already two years ago). With its FX2 platform, Dell houses in a 2U unit computing, storage and network, in addition to hot-swappable cooling and CMC and iDRAC administration interfaces. Its modular build it allows an enterprise to mix and match different modules (blocks) to suit specific needs. It is thus possible to insert up to 4 FM120x4 modules containing 4 servers each (with their own SSD) octo-core Atom-based (up to 128 cores and 32 DIMMs), or up to 4 FC630 modules based on Xeon E5-2600v3 processors for a total of 144 cores and 96 DIMMs or up to 8 FC480 blocks (224 cores, 64 DIMMs and 16 1.8" SSDs) or up to 3 FD332 storage blocks for a total of 48 SAS or SATA

**Each FC630 Module equips the PowerEdge two Intel Xeon E5-2600v3 18 cores each, 24 memory slots and two slots for 2.5-inch disks (or 8 slots for 1.8-inch disks) and a 10GbE connectivity .**

drives. The modules are combined as needed. A real construction set to build a «house» converged system fully shaped to your needs. And all is administered with a user-friendly Web interface.

The architecture of the Dell PowerEdge FX2 represents a significant evolution. It combines the flexibility of a rack module and the density and performance of a blade server. Which combined, gives this architecture unparalleled scalability to match a wide range of needs



## The extreme scalability of the PowerEdge FX2 architecture combines full-width modules, half width quarter or up to eight modules wide.



**At the rear, one of the configurations of the PowerEdge FX2 allows the addition of eight PCI Express cards. To the left are the connectivity modules.**

within the company. The PowerEdge FX2 is however, not intended to replace the range of traditional servers and blades of the manufacturer. It adds to the existing range of products.

The FX2 is available with two rear baskets: the first offer 8 low profile and half-length PCI Express slots, the second is limited to adding additional high performance network connectivity. Along the PowerEdge FX2, Dell has also launched several modules to configure the server.

### MODULAR BLOCKS TO CREATE THE SERVER ADAPTED TO YOUR NEEDS

Several blocks (or modules) are already available. The half-width unit PowerEdge FM120x4 (4760€) is an C2000 Atom-based MicroServer with 2, 4 or 8 cores and optimized for minimal energy consumption. Up to 16 servers housing 4 FM120x4 modules can fit in a FX2 rack, totaling 160 cores and 48 memory locations. A drive slot accommodates a 2.5-inch or two 1.8-inch SSDs.

The PowerEdge FC630 module (5560€) is designed to build a server intended for private cloud solutions and converged data center scenarios. A rack can accommodate up to 4 FX2 FC630 modules for a total of 144 cores and 96 memory locations.

### IN CONCLUSION

The new generation of converged server Dell PowerEdge FX2 accomplishes the feat of combining high-density blade servers and performance usually found in bulkier servers to concentrate computing power suitable for intensive tasks. The range of available modules will go on increasing, making the PowerEdge architecture FX2 a sustainable investment.

**RAMON LAFLEUR**



**In this configuration, the PowerEdge server FX2 contains 4 PowerEdge FM120x4 modules. It contains an Atom C2000 processor with 8 cores, a 2.5-inch drive slot (or two 1.8-inch) and a Gigabit Ethernet connection.**

HPC Bench  
Global index



**9 108**

**Chassis 2U Dimensions H: 8.68cm x W: 43.35cm x D: 85.16cm** **Optional blocks** up to four half-width blocks **FC630 Dell PowerEdge** or **PowerEdge FM120x4** **Power 1600W** or **1100W** hot-plug in redundant modes (1 + 1) or not (2 + 0) **Cooling 8** hot-removable fans **Network** up to two 1Gb or 10Gb modules **Slots** up to 8 PCI Express 3.0 slots (low profile, half-length) **Administration** Chassis Management Controller (CMC) to manage the shared infrastructure modules from a web console. iDRAC8 administration port



# Infortrend EonStor DS 3000



Infortrend is a specialized manufacturer in high-performance storage arrays. The EonStor DS 3000 range is divided into models with either one controller (suffix G) or two controllers (suffix R) providing additional redundancy.



## MODULAR DESIGN FOR MAXIMUM SCALABILITY

With its patented modular design, the EonStor DS arrays 3000 can see their controllers, power and backplanes exchanged in minutes – either in case of failure or in order to reach a higher level of performance. An automatic mechanism research and error self-correction ensures data integrity. The manufacturer has designed its bays with attention to storage capacity control, which includes an integrated thin provisioning mechanism.

## A WIDE RANGE OF MODELS

The standard storage array is has 2U height and has twelve 3.5-inch drives. Infortrend offers in each of its DS 3000 lineup 3U arrays / 16 bays and high density 4U arrays with 24, 48 and 60 disk bays. The latest versions of the firmware allows the secure remote replication for disaster recovery purposes and feature an automatic tiering mechanism to enable low response time demanding applications to take advantage of the installed SSD.

## BUILT FOR SPEED

Infortrend has taken particular care in the performance of its disk arrays, and boasts transfer rates of up to 5500 MB / sec transfer rate from the cache with 1.3 million IOPS. The performance is also visible in a data reconstruction scenario : a 2U 12-bay unit with 4TB disks is reconstructed in nine hours only. The SANWatch Management Software 3.0 provides visibility on the occupancy of storage areas, and displays continuous performance measures. **RAMON LAFLEUR**

HPC Bench  
Global index



**8 760**

**Features** Chassis 2U, 3U or 4U. 12, 16, 24, 48 or 60 drive bays. Connectors 8G or 16G Fibre Channel, iSCSI, 1G / 10G optional SAS controller 6G. Systems capable of aggregating 360 discs for 1440 To JBOD.



# Dell Precision Workstation Tower 5810

Beware, war machine!



**T**he Dell Precision T5810 is a workstation in a compact package combining an Intel Xeon processor E5-1650 v3 with 6 cores at 3.5 GHz and an Nvidia Quadro K6000 accelerator card. An explosive combination designed for the most demanding applications. This workstation will naturally find a place in a graphical design environment, whether it's 3D modeling or rendering calculations, for

example. Given its numerous certifications, it is also suitable for any architecture, energy exploration, and research.

## DESIGN AND FEATURES

The Precision 5810 is integrated into a relatively compact chassis measuring approximately 42 x 45 x 18 cm. The front and back panels are perforated, ensuring good air circulation in the chassis. With two carry handles (front



**If your business depends on powerful machines, this model is an excellent choice. The price of the tested configuration is ten times higher than the standard machine configuration, but demonstrates its scalability and adaptability in demanding production environments.**

and rear), the design of the chassis is made to simplify maintenance. Like the Lenovo ThinkStation P workstation range, all contact

areas within the chassis marked blue authorize tool-less replacement of elements. That goes for PCI Express cards, power supply and hard drives. A few seconds are enough to maintain it.

### **BUILT FOR PERFORMANCE**

The tested configuration included four SSD 400GB SAS configured in RAID 0 for a total of 1.6 TB of storage space, with transfer speeds suitable for intensive uses for which this machine

was designed. The standard model comes with 32 GB of memory. The only real downside of this model is its relatively limited memory and storage scalability, limited to two memory modules and two slots for 3.5-inch drives. The PCI Express side is a richer:

in addition to the x16 slot for graphics card the T5810 features three other x16 slots and even a short PCI slot. The interfaces are a plenty, with three USB 2.0 ports and an USB 3.0 port on the front panel, completed at the rear by three USB 2.0 ports and three USB 3.0 ports. The software environment is limited to the operating system and the Optimizer 2.0 Dell Precision utility. The latter is used to define the GPU usage thresholds, Hyper-Threading and even limit the number of active cores, in order to suit the intended application environment: Adobe Photoshop CS, Autodesk AutoCAD, Maya and others do not exploit the multi-core processors in the same way.

### **CONCLUSION**

If your business depends on powerful machines, this model is an excellent choice. The price of the tested configuration is ten times higher than the standard machine configuration, but demonstrates its scalability and adaptability in demanding production environments.

**RAMON LAFLEUR**



HPC Bench  
Global index



**7840**

**Intel® Xeon® Processor E5-1650 v3 (6 cores, 3.5 GHz, Turbo, HT, 15 MB cache, 140 W), Windows 7 Professional (supports 64-bit and Windows 8.1 license), 4 GB DDR memory 32 2,133 MHz, evaluated configuration: 4 SSD 400GB SAS drives, Nvidia Quadro K6000.**



# HPC Today

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...



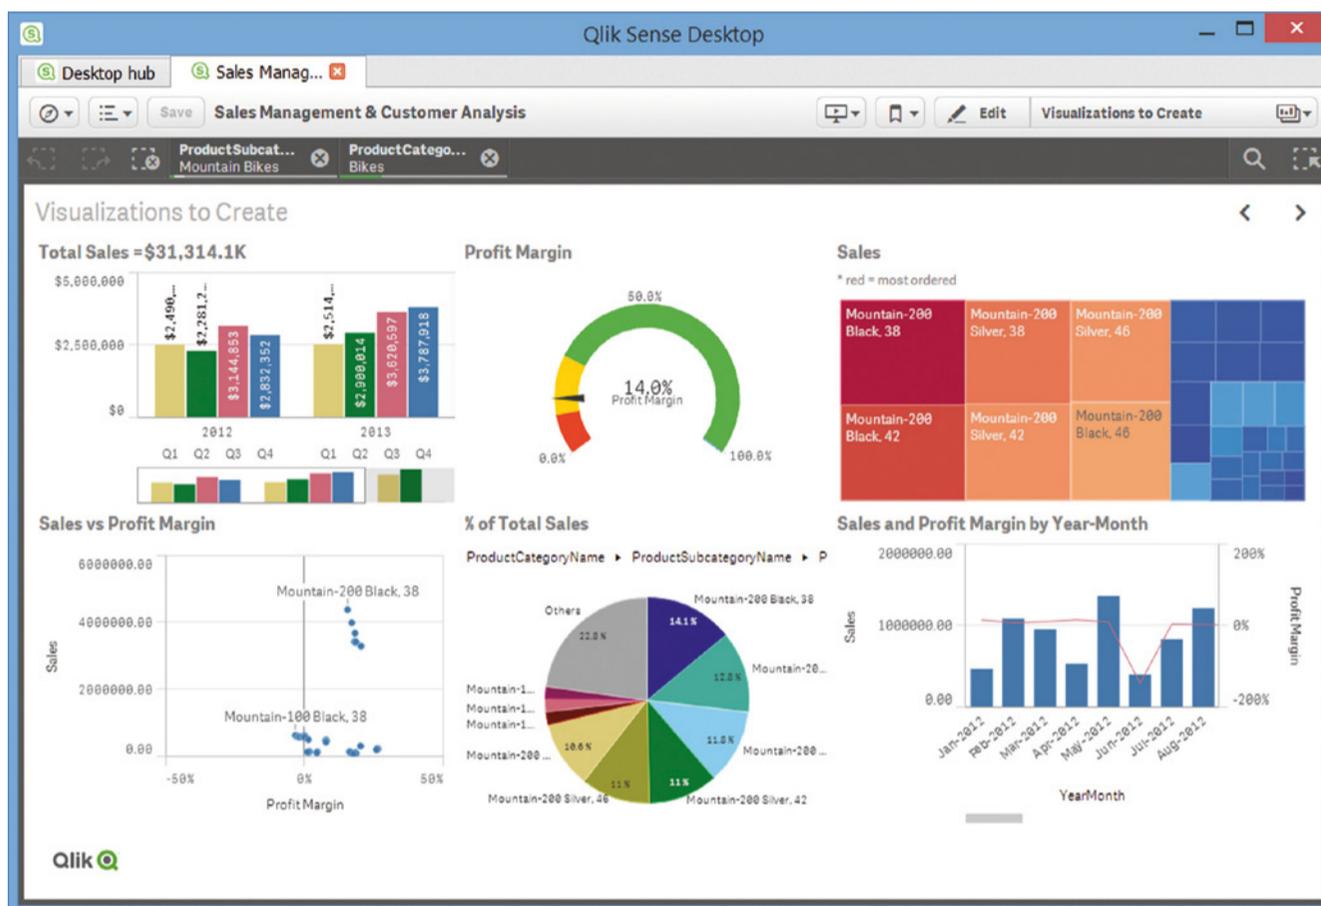
+ get yourself instant access to exclusive contents and services

# www.hpctoday.com



# Qlik Sense Desktop

**Qlik Sense Desktop is a new original tool for exploring and visualizing data rich, complete and... free.**



sources. Connectors are available for importing data of social networks, Oracle E-Business Suite, Microsoft Dynamics, Google Big-Query, SAP, etc. The only drawback we found is the lack of a search function of web data sources like Microsoft Power Query offers.

## IMPORTING DATA FROM VARIOUS SOURCES

The principle of busi-

**A**lready known for its QlikView software, Qlik launched QlikView Desktop, a self-service Business Intelligence tool dedicated to data visualization. The single-user version is free and complete and has no limitation except the inability to function in a collaborative mode. The visual richness and usability of the interface are a welcome surprise.

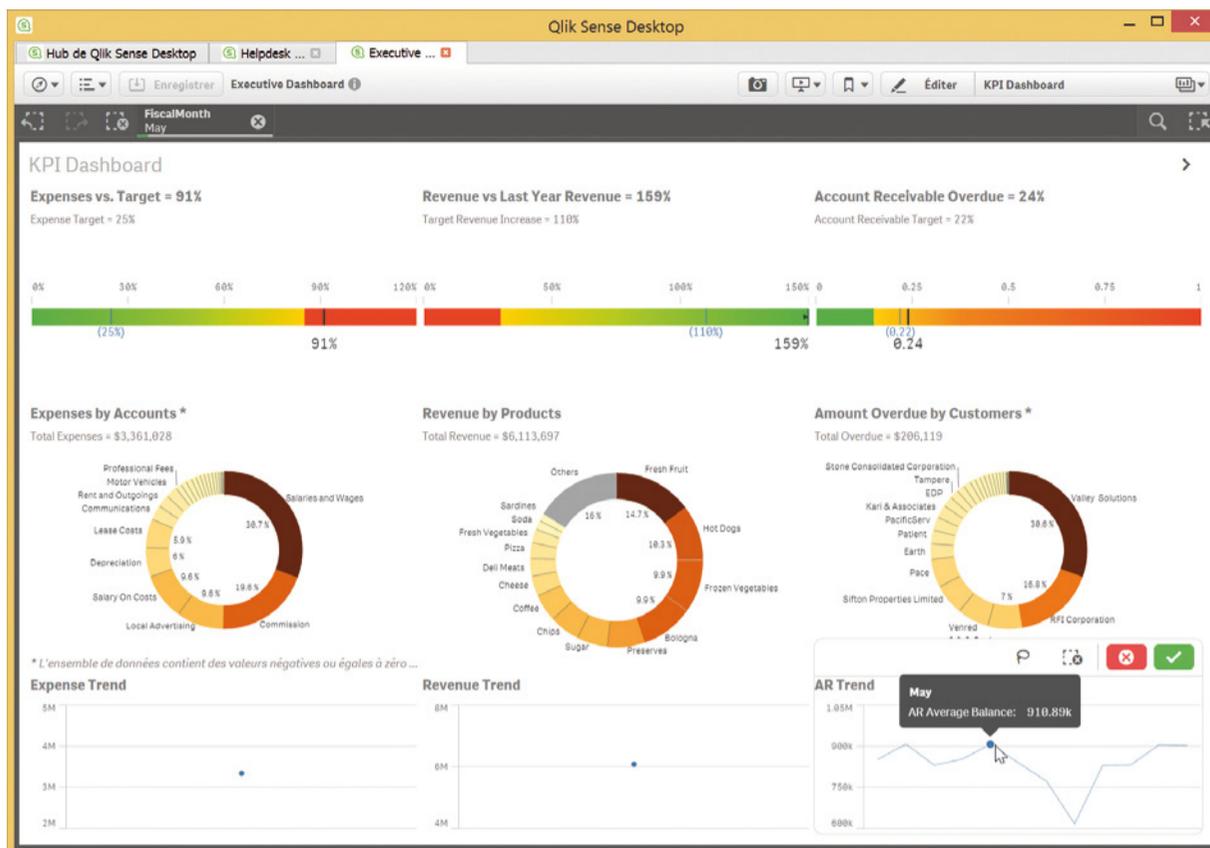
## USER-INTUITIVE AND EFFICIENT

This product allows you to explore, analyze, visualize and interact with data without any prior skill. This friendliness extends to the data import. This is achieved through a wizard allowing simple drag and drop of your Excel XML, CVS and QVD files. A more advanced «Loading Data» function can import data source ODBC and OLEDB or from cloud

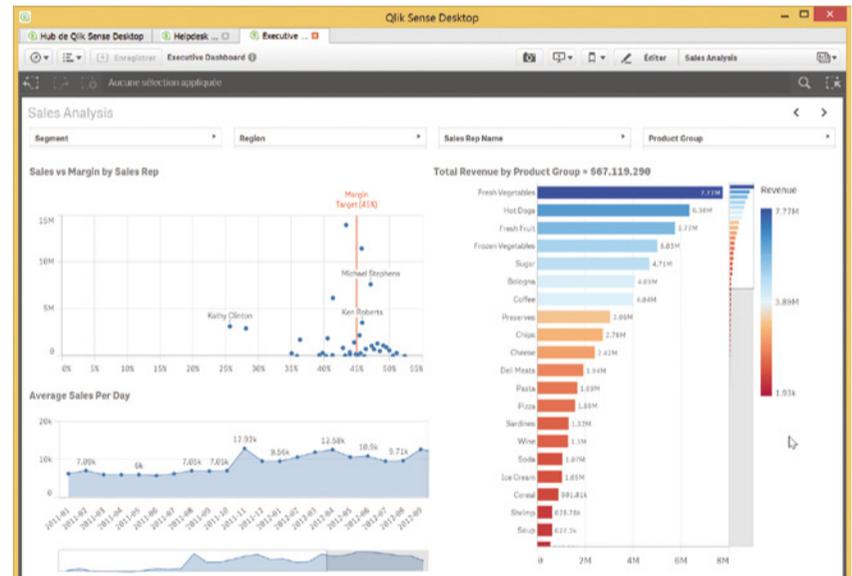
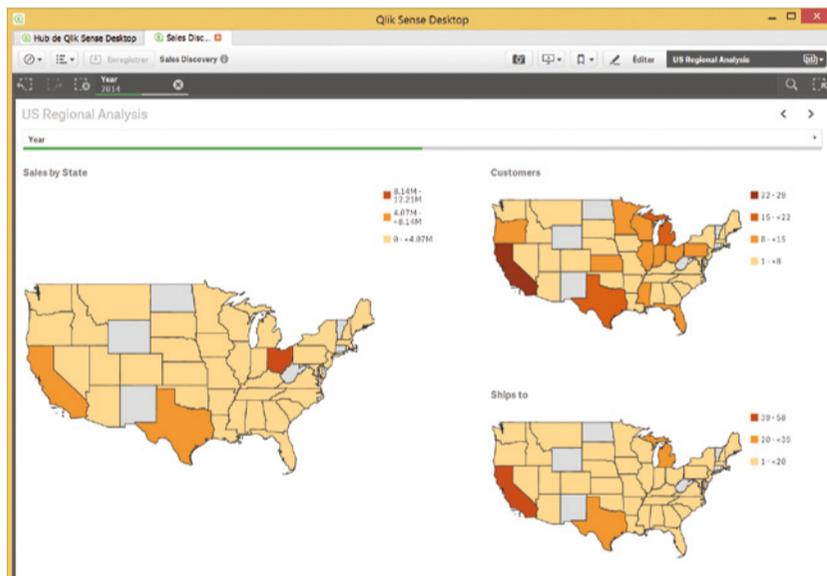
ness intelligence is to mix data sources. When importing data, Qlik automatically creates the appropriate joints from fields with same name. In the case of several fields common to both sources, Qlik Sense automatically creates synthetic keys to manage these relationships. By default, Qlik Sense imports the data into its engine in memory, which can be switched off if the amount of memory is insufficient for the volume of data imported by a live query method.

## AN INTUITIVE DATA EXPLORATION

Serious and amusing things start after the import data phase when one begins to create visualizations and playing with the data by filtering. The construction is done by simple drag and drop. With its responsive design interface, the screens can be adapted to different



**Intuitive and graphically efficient, Qlik Sense is an effective learning tool in a Self-service Business Intelligence scenario.**



screen sizes and orientations. The end result is always elegant and modern. The indexing engine allows to explore data without having to create a new visualization. Bookmarks safeguard the state of the filters you created, allowing you to come back later. Selecting items on a graph automatically adjusts the display of the corresponding graph. The result is a simple exploration of data with disconcerting simplicity.

## DATA STORYTELLING

Qlik Sense has an original storytelling feature that allows during sorts and filters, to so

screen captures allowing you to arrange them at will in a document for them to be annotated, commented and enriched with visual effects to focus on remarkable data variations. Compared to creating a powerpoint document the time savings are huge.

Intuitive and graphically efficient, Qlik Sense is an effective learning tool in a Self-service Business Intelligence scenario. Being moreover totally free, the Desktop version is surprisingly powerful and complete. **RAMON LAFLEUR**



# TinyTitan: a Raspberry Pi-computing based cluster

**We already knew the versatility and solid reputation of the Raspberry Pi platform. But what happens when one combines several Pi's to run them in parallel? The result is a massively parallel cluster! Here is how you can reproduce the Tiny Titan project.**

**T**he Oak Ridge facility is known to house the Titan supercomputer, capable of more than 20 quadrillion calculations per second. But this laboratory also houses lesser-known Tiny Titan, a cluster made of nine Raspberry Pi processing units, or «cores».

TinyTitan is a \$ 1,000 supercomputer designed to teach the basics of future scientific parallel computing. Unlike desktop and laptop

computers that hide the parallelism to their users, TinyTitan makes the concept of parallel computing explicit.

According to its creator Adam Simpson, creating a small-scale supercomputer is relatively simple, just wire the Pi's for them to work in parallel. The hardest part, explains this Titan support specialist, is to write the code for it. According to Adam Simpson, «there is no proper platform to achieve personal supercomputing projects.» So he and two of his colleagues decided to create a set of tutorials. Working together, the parallel cores can solve complex problems, and the tutorials are intended to explain the process. For example, a lesson shows how to add two numbers together lists, dividing the problem into individual nuclei components working simultaneously. The tutorials are available free online.

## The TinyTitan shopping list

**Here are the ingredients to build your own TinyTitan supercomputer**

QUANTITY	DESCRIPTION
9	Raspberry Pi Model B + with an 8GB MicroSD card
9	Animated Power Cables
9	LED on USB port
3	self-powered USB hubs
3	Raspberry Pi B + (Base)
6	Raspberry Pi B + (add-on)
5	1-meter Ethernet cable
1	Microsoft Xbox 360 Wireless Controller
1	HDMI to DVI Adapter
1	16-port Ethernet switch

**According to its creator Adam Simpson, creating a small-scale supercomputer is relatively simple, just wire the Pi's for them to work in parallel.**



Start your journey to stardom supercomputer by going to [github.com/TinyTitan](https://github.com/TinyTitan).

## CREATE YOUR OWN TINY TITAN

The Oak Ridge lab has created a set of scripts (shown below) to guide you in creating your own TinyTitan. Once your TinyTitan up and running, you will be ready to run the demo software.

### PREREQUISITES

- Two Raspberry Pi with the Raspbian distribution
- A router or network switch for connecting the Pi's
- An active Internet connection

## 01 FIRST STEP

The first step is to update each Raspberry Pi and configure the network settings. To start the each Pi that will be part of the cluster TinyTitan needs an active Internet connection. Download the initial configuration script and run it in a terminal as such:

```
$ Curl -kfsSLO https://raw.githubusercontent.com/TinyTitan/TinySetup/master/pi_setup.sh
$ Bash pi_setup.sh
```

When you are prompted to enter the number of nodes start at 1 and increment the value up to match the number of physical nodes present.

## 02 STEP TWO

Once the first stage achieved, connect all nodes in the network router or switch. On the Pi that which was assigned a connection node number 1, run the following script:

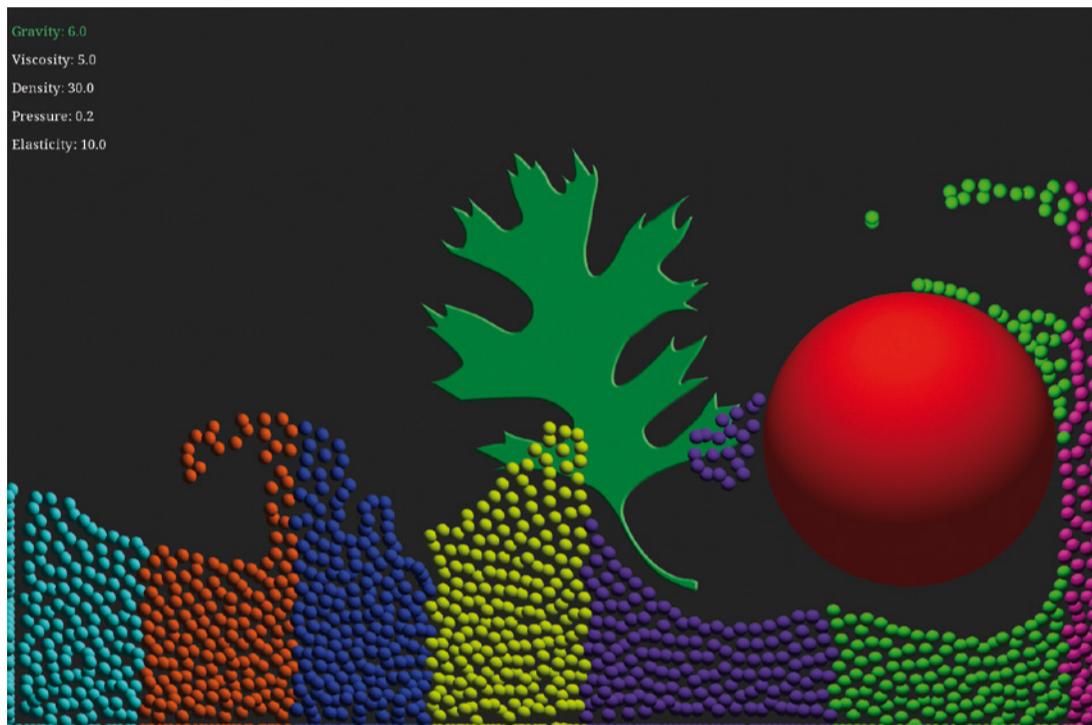
```
$ Git clone https://github.com/TinyTitan/TinySetup.git
$ Cd TinySetup
$ Bash pi_post_setup.sh
```

## 03 STEP THREE

Your Tiny Titan is now operational and ready to work. Extract the TinySPH and PiBrot examples available on the TinyTitan GitHub repository available there: <https://tinytitan.github.io/software/>

## SOME APPLICATION EXAMPLES

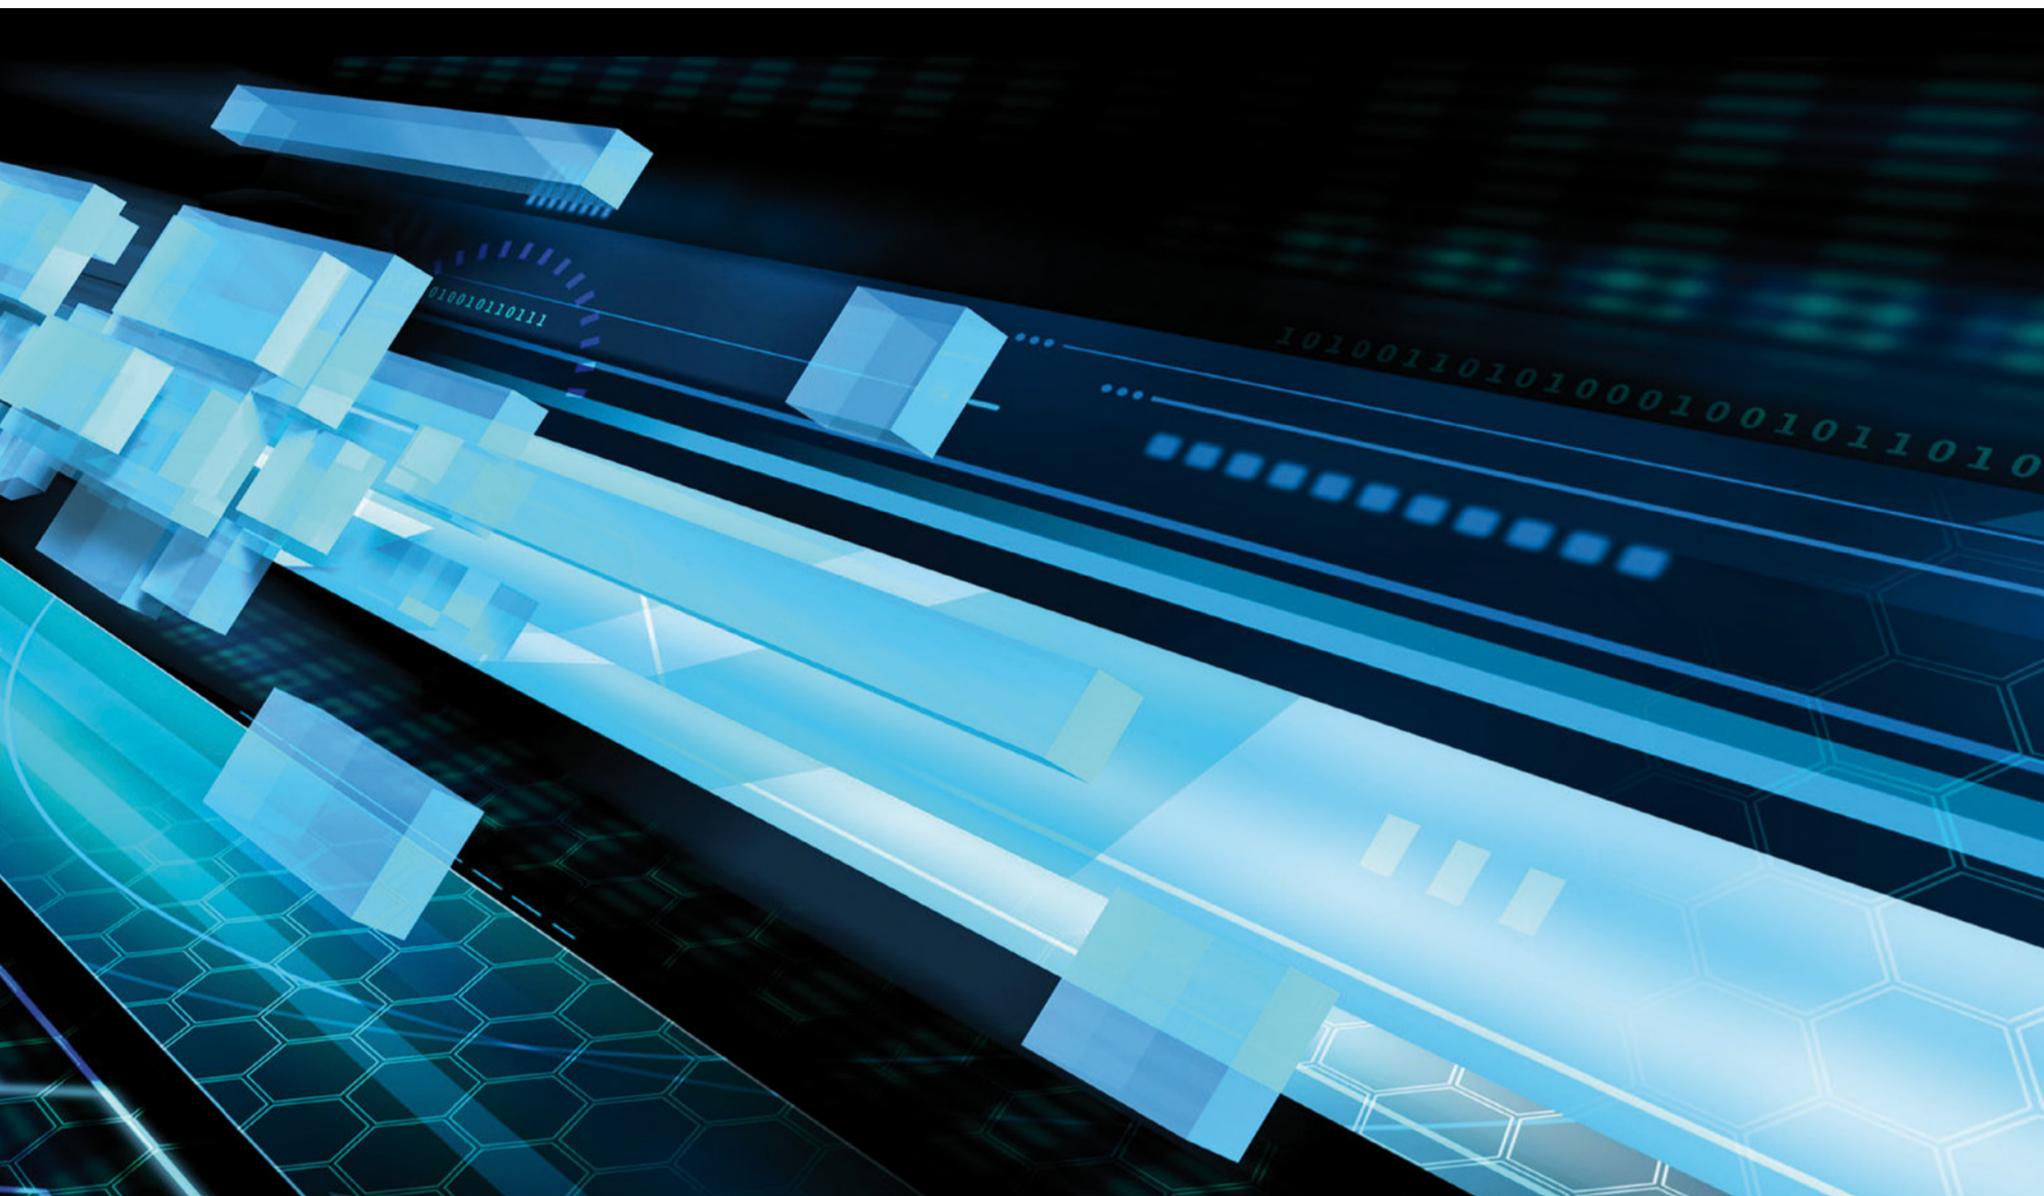
### TinySPH



TinySPH module is a 2D smoothed hydrodynamic particle calculation module designed to work in real time on the Oak Ridge Laboratory TinyTitan cluster. This application was designed to demonstrate the concept of distributed computing. Although it works better in a distributed environment, it is quite possible to run it on a shared memory system with a multicore processor. Although the code is designed to Tiny Titan it can, with some modifications, be compiled and run on Linux and OS X. The code makes use of MPI for distributed computing and requires at least two MPI processes (1 for rendering, one for the calculation). For more information: <http://github.com/TinyTitan/SPH>

### PiBrot

PiBrot is an example of exploiting the parallel computing code to calculate a Mandelbrot fractal. The screen is divided in half vertically with MPI task that displays the left image and the same image being calculated on the right side with N-2 MPI tasks, where N is the total number of MPI tasks. For more information: <http://github.com/TinyTitan/PiBrot>



# The Best Benchmark Is You

**Performance benchmarking has been an important part of evaluating computers for years**

**I**n today's hyper-connected world, technology buyers can find an answer to almost any question with a web browser and an internet connection. We can all do a little research before we sign the check or hand over the credit card. In the personal computer market, performance benchmarking has been an important part of evaluating computers for years, but what do benchmarks really tell us and which ones can we rely on?

History shows the way computers have been evaluated continually changes. For decades computers were primarily sold on the basis of the clock frequency of its processor. AMD was, in fact, the first processor company to reach 1 GHz in 2000, but as the frequencies rose higher and more architectural differences were introduced, the link between clock frequency and performance experienced by the user became increasingly tenuous. Additionally, increases in power consumption and decreases in per-



**ROBERTO DOGNINI**  
*is Director of  
Commercial Sales  
EMEA of AMD*



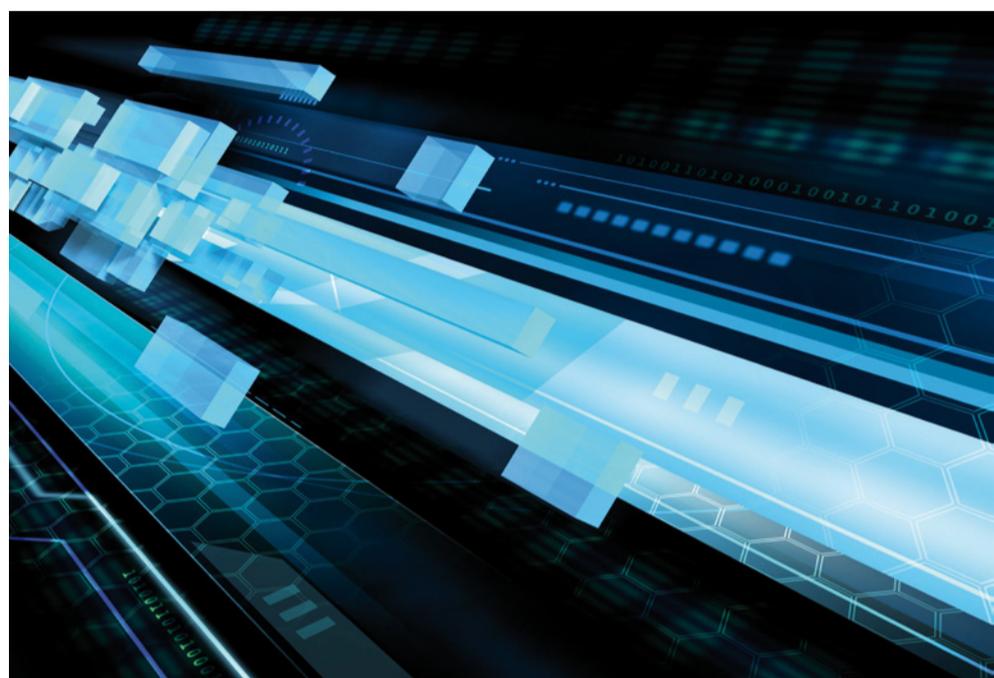
## Large enterprises are drawn to VDI because of its ability to reduce desktop support and management costs, as well as lower overall energy requirements of virtual desktops.

formance scaling with clock speed ultimately killed clock speed as a performance measurement. Microprocessor core counts became the next marketable way many mainstream users were sold computers, with AMD demonstrating the first x86 dual-core processor in 2004 and the first native quad-core x86 server processor in 2006.

### WHAT WERE BENCHMARKS CREATED FOR ?

Benchmarks were developed to help take the guess work out of how much frequency or core counts really deliver in terms of performance, and to provide objective guidance from parties outside of the hardware ecosystem itself. As these software companies maneuvered to become the gold standard of benchmarking against each other, cracks in the model began to appear with hardware companies fighting for optimizations to achieve the highest score resulting in diminished credibility of frequency as a measure of performance.

As processor architectures have evolved, some benchmark suites have not evolved with them; yet they remain a staple to decision makers when judging the performance of computers. To serve today's PC user, there is a need to tap into previously underutilized compute resources which are available in the form of graphics processing units (GPUs) and their massive parallel compute capabilities. Modern processors, like the Accelerated Processing Unit (APU), have both Central Processing Unit (CPU) and GPU processing engines as well as specialized video and audio hardware, all working together to contribute to the user's experience and process modern workloads efficiently for outstanding performance with minimal power consumption.



### ENGAGING WITH COMPUTERS LIKE NEVER BEFORE

Users today expect a rich visual experience and engage with computers like never before; consuming, creating, integrating and sharing high quality audio and video while interacting with their computers through touch, voice and gesture. There is seemingly less interest in what is happening behind the screen and more of an expectation that the system will just work to provide a great experience whether at home or at work. Systems now can include 64-bit multi-core APUs such as AMD's A-Series APUs . These advanced APUs utilize the latest CPU and GPU technologies in heterogeneous system architecture (HSA) optimized designs that allocate the software work to the processing engine best suited for the task.. While computing has reached levels with traditional microprocessor advancements where higher clock speeds require increased power and typically provide diminished rates of return on performance, HSA design is an improved approach to enabling the experiences that people expect today and tomorrow.



## **USERS' NEEDS HAVE CHANGED, BUT DID BENCHMARKS EVOLVE ?**

Given this seismic change in usage and user expectation, you would think that benchmarks would have changed to reflect the technology and user's needs and expectations. The sad reality is that many benchmarks haven't. Measuring just one task or one type of processing, such as single-core CPU performance, these benchmarks provide a limited view of system performance that does not easily translate to an evaluation of the experience of using the system the user is concerned with.

Is it valid to base buying decisions on benchmarks that only measure one aspect of the processor or are more heavily weighted toward a single, rarely used application? When buying a car is horsepower the only specification you consider on the window sticker at the dealership? Ultimately you are the best judge of what is good for you; in an ideal world, hands-on evaluation of a computer can determine if it will either satisfy your needs or it won't, simple as that.

### **THE NEED FOR A REALISTIC BENCHMARK**

In today's online world, hands-on evaluation of a PC is not always realistic, so benchmarks still have an important role to play. There are currently three benchmarks we believe provide a well-rounded picture of modern compute architectures based on the typical workload of today's consumer and commercial users. Two of them are produced by Futuremark, a European based organization that is open to the whole computing industry. Its recent PCMark 8 v2 benchmark suite was developed as the complete PC benchmark for home and business in collaboration with many of the industry's biggest names including Dell, HP, Lenovo, Microsoft and many semiconductor manufacturers. To gain a fuller view of system performance, a suite of modern benchmarks can be used adding the likes of Futuremark 3DMark® for graphics and GPU compute performance and Basemark CL from Rightware for total system compute.

Creating high-quality benchmarks that truly represent today's workloads and provide an accurate picture to users' needs should be inclusive of the whole industry. This is a view shared by industry experts, such as Martin Veitch of IDG Connect:

“We need experts who can measure the right things and be open and even-handed. Otherwise, we stand, as individuals and corporations, to waste vast sums: this applies to the private sector but also government and the public sector where tendering documents can often skew choices and lock out some vendors. That's a situation that is seeing more national and regional groups incorporate Futuremark in tendering documents. Even if they can be controversial and they never stand still for long, one thing is clear: benchmarks are important. And in any high-stakes game, you need somebody trustworthy to keep score.”

### **THE INDUSTRY NEED TO WORK TOGETHER**

When the industry does not work together it can result in benchmarks that are not representative of real-world tasks and can be skewed to favor one hardware vendor over another. This has happened before and while one hardware vendor may benefit, the real loser is the consumer who is presented with skewed performance figures and may pay for a perceived performance benefit.

Ultimately, it is the consumer that wins when the industry works together. The European Commission's recent decision to choose PCMark and make it a vital part of government procurement procedures is a win not just for Futuremark but those wanting to accurately measure system performance. Users can be assured that PCMark's reputation is that of an accurate, representative and unbiased benchmark that stands up to the highest scrutiny.

Benchmarks may be a valuable tool for your buying decision process and an important element in your final decision, but at the end of the day, the best and toughest benchmark is you.



# CD-adapCo, celebrating 35 years of innovation

**35** years ago, three young engineers started an engineering simulation business. Their intention was to apply the numerical simulation technology that was beginning to emerge from the nuclear industry to the wider set of engineering problems faced by industry as a whole. Those three engineers were Steve MacDonald, Bill Wheeler, and Marc Whittlesey and the company that they

formed was adapco (a company that would eventually become CD-adapco). I recently met up with them at the company headquarters in New York, with the aim of discovering how the company got started.

## **BEGINNINGS**

Adapco was born, almost fully formed, out of the Nuclear Industry, which in the late 60s had begun to employ numerical simulation as a tool for understanding some of the structu-



ral issues involved in designing nuclear power plants. Of the three founders, Steve MacDonald was the first to immerse himself in this brave new world of simulation. Graduating in 1967 with a degree in Engineering Physics, from the Colorado School of Mines, he was originally employed by Bettis Atomic Power Laboratory, which was operated by Westinghouse for the Department of Defense, to perform physical testing. “I was supposed to be working on the fatigue testing of materials, but the group that did that were also beginning to get their hands dirty with simulation, so pretty soon I ended up getting involved in that,” recalls Steve MacDonald.

Immediately seeing the opportunity, Steve MacDonald enrolled himself in Carnegie Mellon University at night and began using FEA tools in his job - initially a 2D FEA code called “FEAT”, and later 3D tools that were developed by his own lab. “I got involved in simulation early and did a lot of application development and I worked there for almost five years. I went to night school to get a masters degree. I did all my coursework for a PhD but didn’t finish my thesis’.

In the next few years, rather than pursuing an career in academia, Steve MacDonald began to use the tools at a succession of engineering companies around the Nuclear Industry, finally becoming an engineering supervisor for the FEA team at EDS Nuclear in Long Island, which he describes as “an extremely aggressive nuclear consulting company business that was entirely focused on generating profit”.

Into this team, Steve MacDonald recruited two engineers. Texan Bill Wheeler, a graduate of Texas A&M and Carnegie Mellon, who had previously worked for MacDonald Douglas and Westinghouse, and Marc Whittlesey, a graduate of Iowa State University, who had worked with MacDonald for three years at Foster Wheeler. Almost immediately, MacDonald, Wheeler and Whittlesey began to increase the scope of the engineering simulation employed by EDS.



“We were doing finite element analysis for the nuclear power industry,” says Bill Wheeler. “That was the primary business of the company. They had their own piping stress code, and they were doing lots and lots of piping stress analysis, although we didn’t get much involved with that. Steve had lots of connections from his previous jobs so we ended up working on steam generators, pressure vessels, and things like that. It was a fairly lucrative business because there was a federal law that said you had to do analysis of all the safety related equipment in a nuclear power plant. So, there was a lot of that kind of work being done, trying to show that components were going to hold up to all the postulated operating conditions and accidents without failing.”

### THREE MILE ISLAND

In 1979, the bottom fell out of the Nuclear Industry as a result of the “Three Mile Island” accident, in which one of the two Three Mile Island nuclear reactors in Dauphin County, Pennsylvania, suffered a partial meltdown. It was the worst accident in US commercial nuclear power plant history, and ultimately led to the cancellation of 51 US nuclear reactors that were planned to be built from 1980 to 1984.



“Steve and I were on a plane going to Pittsburgh to talk to some people at Westinghouse,” recalls Bill Wheeler. “And when we landed, we heard about the Three Mile Island accident. And we both just said, ‘That’s going to burst this bubble we’ve been working in.’”

“Steve convinced the company we were working for, EDS Nuclear, that maybe they should broaden their horizons with the gloomy outlook of the nuclear power industry,” continues Bill Wheeler. “EDS gave us a charter for about a year to try to work non-nuclear programs. Steve managed to somehow land us some contracts, with General Electric Aircraft Engines in particular, where we started working on fairly sophisticated three-dimensional finite element analysis of some bolted joints they were having problems with.”

“We did that for about a year, which really taught us how to run a company,” remembers Marc Whittlesey, “Steve, Bill and I had numerous discussions about doing exactly that, because we were doing budgets and planning and manpower and all this kind of stuff and so there was really very little about running a company that we didn’t know how to do.”

MacDonald, Wheeler and Whittlesey began formulating plans to start their own engineering services business sometime in 1981. However, they were forced to put those plans into action rather more quickly when the EDS management got wind of them. “We formed our company, went down and took out the name and everything, so we actually had a company before we left,” says Steve MacDonald. “And, of course, me being a person who tended to open his mouth too much, and still does, I’d talk about it to people. Word got back to the management of EDS Nuclear.”

“The three of us were in Napa Valley at a training session for this large 1,000- man company, and we were terminated,” continues Steve MacDonald. “Which was interesting because I had a wife and three kids and I had a mortgage, as did Marc and Bill, and so it was time for us to go to work. So we were pushed out



*Marc Whittlesey, Steve MacDonald & Bill Wheeler*

maybe a little early. We were probably going to go anyway three months later.”

“The thing that I remember is that when we left EDS Nuclear, they had about 900 people in the United States doing piping support stress analysis,” says Bill Wheeler. “By 1982, they had zero. So, it wasn’t that we were visionaries about forming a business. It was kind of a necessity. I think we all understood we had to do something because the nuclear bubble had burst. That was very obvious to all of us.”

### THE BIRTH OF ADAPCO

So, from the embers of the Nuclear Industry, adapco was born, initially as a four person outfit (the fourth person was Joe Sklerin, a minority shareholder in the original business) that operated out of Bill Wheeler’s attic in Long Island.

“We actually incorporated a company in the late summer of 1980,” remembers Bill Wheeler. “The company name, Analysis and Design Application Company, was about our fourth or fifth choice. I don’t remember what the first



few choices were. But the New York state government said they were too close to somebody else's name and wouldn't let us have that. So, we wound up using Analysis and Design Application Company Limited, and abbreviated to adapco. That's really where we got the name, it wasn't Steve's first choice."

"It was also obvious that there was only going to be one person who was electable as president. And that was Steve", says Bill Wheeler. "If anybody knew him, they would know instantly that he was the only choice. Steve is a natural leader, he had the unique ability to communicate and convince people to give him whatever he needed. That was his really strong point. He was also a very, very technically savvy engineer. But I'm not sure that he had the patience that Mark and I had to work a problem."

Steve MacDonald agrees: "Well, I always say that I wanted to be the president and I worked long enough to know that I would never make president in a highly structured company, so my solution to that was to go form my own company and appoint myself. That was the solution to my personal dislike of authority. Anyway, so I thought, well, I'll just start a company. That's how it happened. I did a lot of things that turned out to be good enough so that we didn't go out of business." "The three of us had things that we each brought to the mix," says Marc Whittlesey. "Steve was the ideas guy, the guy that liked to meet with the customers, go out, sell the work, that became very apparent right from the get-go. Bill Wheeler was the ultimate technocrat, he was always looking for the best way to skin the cat and so forth, and I was kinda the do-er, the guy that was really good at getting simulations running, and the results out of the door."

"So, that mix tended to work fairly well and it kept us out of technical trouble," remembers Marc Whittlesey. "Of course Steve's focus was on getting the work and thinking of the new methods and of course that caused me a lot of consternation at times, when he would sell something based on a promise and we had no

## «When performing numerical simulation you need boundary conditions. If you get your boundary conditions wrong, then you're in trouble.»

idea how to do this and it would drop on my desk. But that's what makes you smart, you know, it gives you a challenge. Something that you didn't know how to do at the time but you knew you had to get done in a limited amount of time and that 'Necessity really is the mother of invention'."

### PHILOSOPHY OF SUCCESS

This "pushing the boundaries" is a common theme. A large part of adapco's (and later CD-adapco's) success is based upon the principle that, right from the very beginning, Steve MacDonald realized the solution of real world engineering problems, which usually involve the interaction of many different physical phenomena, would ultimately require simulation tools that spanned a range of engineering disciplines.

"When performing numerical simulation you need boundary conditions," says Steve MacDonald. "If you get your boundary conditions wrong, then you're in trouble. So if you take an isolated component or area, let's say, the problem is that you are very dependent upon having exactly the right boundary conditions there or you're not going to get the right answers from your simulation. So my philosophy from the beginning was to model as far away from the object as I could get practically because then, even though I may not know the boundary conditions perfectly, I have a chance for the boundary conditions to adjust to a degree, some of them anyway, before the solution had propagated to the component that I was looking at."



“That was our philosophy,” agrees Wheeler. “We’re going to use as big a computing system as we could possibly find or afford. And we’re going to put as much detail and not make assumptions any more than we absolutely had to. That was the guiding principle from the day we started the company. I think that’s probably still reflective of the company overall today. We were not trying to use our experience to make the analysis tell us what our experience says. We were trying to get the analysis to tell us what was really happening and use our experience to try to debug our analysis work.”

“One of the first jobs was from Solar Turbines,” remembers Steve MacDonald. “Another was from General Electric, Lynn, which was the small aircraft engine division of GE. We got a remote job entry terminal, which was a card reader and a printer, and we prepared our problems. Then a little bit later, we bought a computer, which was about the size of a washing machine, it was a VAX 750, and we used that along with some discs. And, again, we were doing stress and thermal analysis, mainly, and a lot of it was on gas turbines, things that had thermal stress in them and were important.”

“And we moved into this building, we went from working in Bill Wheeler’s attic for a very short period of time to renting one-half of the bottom floor of this building at that time, and then we rented the whole bottom floor and then we rented the first floor. Eventually we ended up buying the whole thing. We just kept growing”, says Steve MacDonald.

### FROM STRESS TO CFD

Up until this point, adapco had been almost entirely involved with Finite Element Analysis of thermal and mechanical stress. However, Steve MacDonald’s vision of “pushing the boundary conditions as far away as possible from the region of interest”, began to necessitate the use of other simulation tools. “I remember Steve coming back from some of the meetings he was having with Ford where we were doing some heat transfer analysis,” says

**«We bought a computer, which was about the size of a washing machine, it was a VAX 750, and we used that along with some discs.»**

Bill Wheeler. «We didn’t know what the heat transfer coefficients were, but they could tell us how much the total heat rejection rate was. So, we would just say, ‘Okay, we’re going to put one heat transfer coefficient in the whole water jacket. And we’re going to tune that based upon the heat rejection.’, So that’s what we were doing. He was getting a lot of criticism saying, ‘Aw, you can’t, that’s not right. It’s not constant everywhere.’, so we put our heads together and kept saying, ‘What are we gonna do? What are we gonna do? What are we gonna do?’”

“I stuck my hand up and I said, ‘I think we could probably use CFD to get us some heat transfer coefficients that are not uniform.’, and so we started trying to do that,” says Bill Wheeler. “We did get some heat transfer coefficients out of those early CFD simulations. I don’t think they were very good but they were not uniform. We applied those to the stress model. Again, we tuned them up and down until we got the right heat rejection and kind of satisfied that. But in the interim, the people in the car companies, especially Ford, go, ‘Oh, wow. You can actually see where the flow is.’, which they were having a hard time visualizing back in the early ‘80s. This was probably 1983, we were doing this. And they started giving us contracts just to run the flow solutions without doing the heat transfer stuff in the blocks and stuff. And that’s really where adapco started doing CFD on an engineering services business, it was one of the first computer aided design tools that I think was universally adopted by an entire industry.”



# HPC REVIEW

Don't miss  
**HPC Review #3**  
Available on  
**May 30th, 2015**

## COVER STORY

**Is Machine Learning**  
mixing Big Data  
and Analytics  
**the future of high  
performance  
computing ?**

## LAB REVIEW

**Dell EqualLogic  
ps6210**  
**HP StoreOnce 4500**  
**NumScale analyzer**

## HOW TO

**Convert VMware  
ESX virtual  
Machines** to use  
them in a Microsoft  
Hyper-V environment

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

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

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

PUBLISHER  
**Frédéric Milliot**

EDITOR IN CHIEF  
**Joscelyn Flores**

EXECUTIVE EDITOR  
**Ramon Lafleur**

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

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

ART DIRECTOR  
**Bertrand Grousset**

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